inSiDE

inSiDE • Vol. 8 No. 1 • Spring 2010

Innovatives Supercomputing in Deutschland

Editorial

Welcome to this new issue of InSiDE the journal on innovative Supercomputing in Germany published by the Gauss Center for Supercomputing.

The focus of the last months all over Europe was on the European HPC initiative PRACE (Partnership for Advanced Computing in Europe) and the Gauss Center for Supercomputing has played a very active and important role in this process. As a result of all these activities PRACE is now ready for the implementation phase. Here you find an overview of the results achieved in this issue. In our autumn edition we will be able to report on the next steps which will include the founding of a formal European organization.

PRACE is an important driving factor also for German supercomputing. Hence the Gauss Center for Supercomputing has again stepped up its effort to provide leading edge systems for Europe in the next years. HLRS should be able to announce its system soon after we go to press – so stay tuned for further information. LRZ has started its procurement process. InSiDE expects to report on the results in the autumn issue.

As we enter the Petaflops-era applications are the driving factors of all our activities. InSiDE presents the three winners of a Golden Spike Award of the HLRS. The applications that were honored by the HLRS steering committee are reporting on direct numerical simulation of active separation control, the modeling of convection over West Africa, and the maturing of galaxies by black hole activity. Further application papers deal with the numerical simulation and analysis of cavitating flows, the driving mechanisms of the geodynamo, the physics of galactic nuclei and the global parameter optimization of a physics-based united-residue force field for proteins.

Projects are an important driving factor for our community. They provide a chance to develop and test new methods and tools. They allow further improving existing codes and building the tools we need in order to harness the performance of our systems.
A comprehensive section of projects describes among other the recent advances in the UNICORE 6 middleware software. This is a good example how a nationally funded project that was supported by European projects finally turned into a community tool and is now taking the next step in the improvement process.

Let us highlight a typical hardware project, which is an excellent example of how the community can drive also hardware development. QPACE (Quantum Chromodynamics Parallel Computing on the Cell) is a project driven by several research institutions and IBM. The collaboration is a good example how we can deviate from main stream technology and still closely work with leading edge technology providers. The resulting merger of ideas and technologies provides an outstanding solution for a large group of application researchers.

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

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

Contents

PRACE: Ready for Implementation

In October 2009 PRACE [1], the Partnership for Advanced Computing in Europe, demonstrated to a panel of external experts and the European Commission that the project made "satisfactory progress in all areas" [2] and "that PRACE has the potential to have real impact on the future of European HPC, and the quality and outcome of European research that depends on HPC services".

This marked an important milestone towards the implementation of the permanent HPC Research Infrastructure in Europe, consisting of several world-class top-tier centres managed as a single European entity. It will be established temporarily as an association with its seat in Lisbon in spring 2010. The project has prepared signature-ready statutes and other documents required to create the association as a legal entity. Among them is the Agreement for the Initial Period that covers the commitments of the partners during the first five years where national procurements are foreseen by the hosting partners. Germany, France, Italy, and Spain have each committed themselves to provide supercomputing resources worth at least 100 Million Euro to PRACE during the initial period. The Netherlands and the UK are expected to follow suit. The access to the PRACE resources will be governed by a single European Peer Review system.

PRACE is one of 34 Preparatory Phase projects funded in part by the European Commission to prepare the important European Research Infrastructures identified in 2006 by ESFRI [3]. Having started in January 2008, PRACE is among the first projects to complete its work after only two years. This includes not only the legal and administrative tasks and the securing of adequate funds but due to the specific challenges of the rapidly evolving nature of High Performance Computing also very important technical work.

PRACE identified and classified a set of over twenty application that constitute a cross-section of the most important scientific applications that consume a large portion of the resources of Europe's leading HPC centres. When necessary, these applications were ported to the architectures that PRACE identified as likely candidates for near term Petaflop/s systems. The focus was, however, to evaluate how well the applications scale on the different architectures, how well the systems enable petascaling, and how difficult it is to exploit the parallelism provided by the hardware. Ten of these applications constitute the core of a benchmark suite that may be used in the procurements of the next Tier-0 systems.

In parallel PRACE conducted an in-depth survey of the HPC market involving all relevant vendors. The objective was to classify and select architectures promising to deliver systems capable of a peak performance of one Petaflop/s in 2009/10. As one important result a comprehensive set of architectures was defined and prototypes covering this set were deployed as prototypes at partner sites. In detail theses were: IBM Blue Gene/P (MPP [4]) at FZJ, CRAY XT5 (MPP) at CSC, IBM Power 6 (SMP-FN) at Sara, IBM Cell at BSC (Hybrid), NEC SX-9 coupled with Intel Nehalem (Hybrid) at HLRS, and Intel Nehalem clusters by Bull and Sun respectively at CEA and FZJ (SMP-TN). All prototypes were evaluated in close to production condition using synthetic benchmarks to measure system performance, I/O bandwidth, and communication characteristics. Most applications from the above set of were evaluated on at least three different prototypes. The results contain a wealth of information for the selection of future systems and where vendors and application developers should focus to improve performance. The Blue Gene/P installed at FZJ in 2009 is the first European Petaflop/s system that qualifies as a PRACE tier-0 system. In addition, PRACE made the prototypes available to users through open call for proposals. In the first and second round, over 8.9 million CPU hours were granted to nine research teams. The results will be published on the PRACE web site.

To stay at the forefront of HPC, PRACE has to play an active role: The project evaluated promising hardware and software components for future multi-Petaflop/s systems, including: computing nodes based on NVIDIA as well as AMD GPUs, Cell, ClearSpeed-Petapath, FPGAs, Intel Nehalem and AMD Barcelona and Istanbul processors; systems to assess memory bandwidth, interconnect, power consumption and I/O performance as well as hybrid and power-efficient solutions; programming models or languages such as Chapel, CUDA, Rapidmind, HMPP, OpenCL, StarSs (CellSs), UPC and CAF. PRACE co-funded work to extend the world's "greenest" supercomputer QPACE [5] for general purpose applications.

Finally, promising technologies, especially with respect to energy efficiency, will be evaluated with the ultimate goal to engage in collaborations with industrial partners to develop products exploiting STRATOS, the PRACE advisory group for Strategic Technologies created in the PRACE Preparatory Phase project.

PRACE Management Board welcomes Bulgaria and Serbia to the PRACE Initiative (Toulouse, September 8, 2009). Photo by Thomas Eickermann.

PRACE started an advanced training and education programme for scientists and students in advanced programming techniques, novel languages, and tools. Among them were a summer school (Stockholm, 2008) and a winter school (Athens, 2009), six code porting workshops and two seminars for industry. This will be continued and intensified in the permanent Research Infrastructure to ensure that users can effectively harness the 100thousand-fold parallelism and exploit the opportunities of novel or heterogeneous architectures in systems containing Cell processors or GPUs. Over 48 hours of training material are available on the PRACE web site as streaming videos.

The full impact of PRACE for European Science and industry can only be achieved if it serves informed users. PRACE partners gave over 90 presentations at scientific, technical, and policy setting events. PRACE showcased its achievements in the exhibitions of the Supercomputing Conferences in the US, the International Supercomputing Conferences (ISC) in Germany, and the EU initiated ICT conferences. The PRACE award for the best paper on petascaling by a student or young European scientist at ISC is already a recognized tradition.

The successful achievements of the PRACE project were important prerequisites to become eligible to apply for a grant of 20 Mio Euro supporting the implementation phase. PRACE submitted a related proposal for a PRACE First Implementation Project in November of 2009. The number of partners grew from 16 in 14 countries to 21 in 20 states [6] demonstrating the importance of HPC to Europe. Of course, the permanent PRACE Research Infrastructure will be open to all European states. The Gauss Centre for Supercomputing represents Germany both in the legal entity and in the project. The project partners asked the Research Centre Jülich to coordinate also the Implementation Phase project and to continue the very successful management provided by JSC's project management office.

The implementation phase project is expected to start in July 2010. To keep the momentum among the project partners and bridge the gap an extension of the project until June 2010 has been proposed by the project and granted by the European Commission. This will allow to the continuation of dissemination and training, porting and scaling of additional applications and evaluation of future technologies.

PRACE is well positioned and ready to shape the European HPC eco system.

References

[1] The PRACE project receives funding from the European Community's Seventh Framework Programme under grant agreement no RI-211528. Additional information can be found under www.prace-project.eu

[2] The reviewers may only tick yes or no aside to this predefined text when evaluating a project.

[3] ESFRI: European Strategy Forum on Research Infrastructures report at ftp://ftp.cordis.europa.eu/pub/esfri/docs/esfri-roadmap-report-26092006_en.pdf

[4] MPP: Massively Parallel Processor; SMP-FN/TH: Symmetric Multi Processor-Fat Node/Thin Node

[5] See: http://www.green500.org/lists/2009/11/top/list.php

[6] Germany is represented by GCS as member and FZJ as coordinator to comply with EC regulations.

New JUVIS Visualization Cluster

JSC offers a new visualization service which was made possible by additional BMBF-funding dedicated to the D-Grid project. The compute cluster named JUVIS is intended to be used as a parallel server for the visualization of scientific data. JUVIS allows the parallel processing of data as well as parallel rendering while the final image is displayed on the desktop of the user's local workstation. The cluster has been procured, installed and tested at JSC and is ready for operation since a couple of months. It consists of one login node, 16 compute nodes, and one fileserver with a 10.5 TB RAID system. In total the compute nodes are equipped with 128 CPU cores (32 Quad-Core Intel Xeon Processors E5450, 3.00 GHz), 256 GB main memory and Nvidia Quadro FX 4600 GPUs. In a first step, the visualization software ParaView was installed on JUVIS. ParaView is an open-source, multi-platform data analysis and visualization application based on the Visualization Toolkit VTK. It is able to run in parallel on the new visualization cluster, allowing the post-processing and remote visualization of very large data sets without the need to transfer the data to the local workstation.

In addition to ParaView, many other visualization software systems can be used on JUVIS with enabled remote rendering capability. For this purpose, a Virtual Network Computing (VNC) server and the VirtualGL package are installed. This infrastructure allows the usage of any OpenGL based visualization software with full 3D hardware acceleration on JUVIS, while the resulting image is displayed on the user's desktop in a VNC client window.

Golden Spike Award by the HLRS Steering Committee in 2009

Direct Numerical Simulation of Active Separation Control Devices

Aircraft design has always been a struggle for the most “efficient” form or method. The most efficient has been defined either by increased peak performance or in more recent times by reduced operational cost. In engineering terms reduced operational cost equals reduced or constant drag with constant or increased lift. Today's commercial airliners commonly encounter so called flow separation on the wing flaps during landing which led to the employment of rather complex and heavy flap mechanics which became necessary to compensate for the reduced lift. Therefore the suppression of such flow separation constitutes a vital step towards more efficient wing and flap design. The cause for flow separation are positive pressure gradients along the streamline which act as obstacle to the oncoming fluid. In order to enable the fluid to overcome larger adverse pressure gradients it has to be energized (see Figure 1). One method to do so is to move high energy fluid out of the undisturbed freestream towards the surface. The necessary vortical motion is induced by a large longitudinal eddy created with so called Vortex Generators (VG). Already in use are passive vortex generators, i.e. small sheets attached to the wing surface. Indeed these VGs are well able to generate aforementioned eddies but there are also disadvantages due to the fact that the passive VGs are optimized for a specific point of operation and induce parasitic drag at all other flight attitudes. In 1992 experimental work by Johnston et al. [1] has shown the general ability to suppress flow separation using so-called Jet Vortex Generators (JVG). The major advantage over solid VGs stems from the fact that JVGs are actively controllable and thus dispose of any additional drag outside of the point of operation. These promising findings led to intensive research in JVGs and cumulated in exhaustive experimental parameter studies covering many aspects such as velocity ratio, radius and blowing angle.

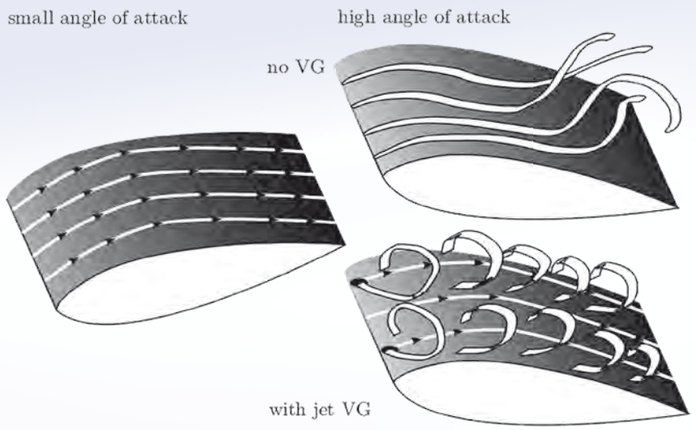

Figure 1: Scheme of Jet Vortex Actuator

Albeit the outcomes of these experiments yield a very good general idea of the mechanisms of active flow control devices there are still a number of open questions involved as no detailed picture of the formation of the vortex and its interaction with the boundary layer could be gained from experiments yet. Therefore, any design suggestions rely heavily on empirical data and are difficult to transpose to different configurations. During industrial design processes parameter studies are thus often undertaken numerically by means of Reynolds Averaged Navier-Stokes (RANS) methods. The crux of RANS lies in the fact that the involved model assumptions have to be adapted to every new configuration and in the fact that the underlying equations are inherently steady state and thus do not really allow for simulation of transient processes. Within this context the presented work covers simulations of a generic JVG configuration by means of highly accurate Direct Numerical Simulation (DNS) methods. The DNS approach is used for its lack of any model assumptions. Therefore, it is well suited to provide a reference solution for coarser or “more approximate” numerical schemes. Furthermore, DNS allows for a computation of the unsteady flow formation especially in the beginning of the vortex generation and detailed analysis of the fluid dynamics involved.

Navier-Stokes 3D

The numerical solver Navier-Stokes 3D (NS3D) has been developed at the Institut für Aero- und Gasdynamik (IAG) at Stuttgart University [2]. The program is based on the complete Navier-Stokes equation, i.e. no turbulence or small scale models are used. The computational domain consists of a rectangular box. Boundary conditions are applied to define an inflow, outflow, freestream and wall region. The fundamental differential equations are converted into approximate difference equations using finite differences on a structured grid. The finite difference stencils are chosen to be of compact form, i.e. information of both flow variable and its derivative is taken into account resulting in a numerical scheme of spectral-like wave resolution in space. The evolution in time is simulated by an explicit Runge-Kutta time integration scheme which maintains the accuracy of the spatial resolution. The resulting linear system of equations can be solved very efficiently on vector CPU supercomputers like the NEC SX-9 at HLRS Stuttgart because of the strongly structured form. Furthermore the domain can be split into boxes of equal dimensions which are then assigned to one computational

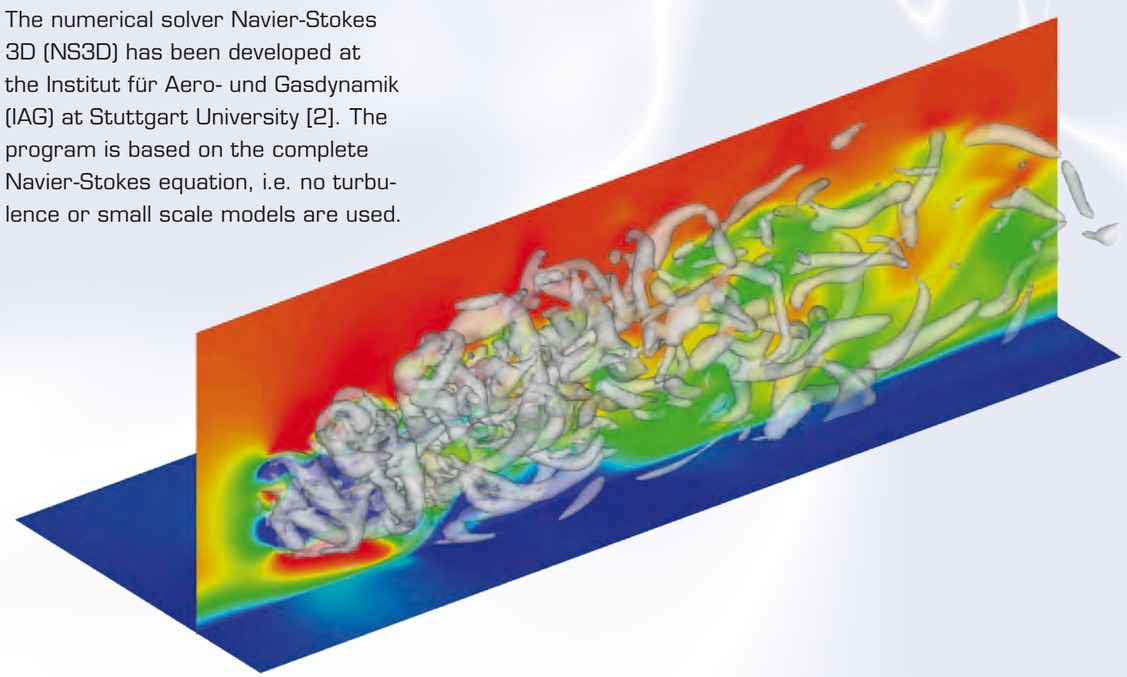

Figure 2: Instantaneous streamwise velocity contours and isosurfaces of vorticity visualized by λ2 criterion (β=150°)

process each. Interprocess communication is realized by use of Message Passing Interface (MPI) routines. Within each MPI process the domain is furthermore parallelized employing NEC Microtasking shared-memory parallelization. All computations described here have been run on one node of the NEC SX-9 within 12hrs computing time.

Jet Vortex Generator Simulations

Numerical simulations have been done for two Jet-in-Crossflow configurations which differ in the jet exit's orientation only. The configurations were chosen to closely match experiments described in [3] and are representative for separation control devices. The crossflow represents a laminar boundary layer on a flat plate with zero pressure gradient at freestream Mach number of Ma=0.25 which corresponds to landing speed for commercial airliners. The jet is described by a steady velocity distribution at the wall boundary of the computational domain and resembles the profile of a pipe flow. The jet-to-crossflow velocity ratio is set to R=3. The jets are inclined by an angle of attack α=30° and skewed to the downstream direction by angles β=30° and β=150° respectively. One can imagine an oblique jet to be blown either against the oncoming main flow or in line with it. Altogether 18.4m grid nodes are used and computations have been carried out for 31 flow through times. Albeit this does not suffice to provide data for statistical analysis, it yields a good picture of the evolution of the perturbed flow field. Even though both the jet and the laminar boundary layer are initially in a steady state the resulting flow regime becomes highly unsteady with the jet exhibiting unstable modes leading to the formation of large transient vortex structures reaching far out of the boundary layer (Figure 2). The jet affects the boundary layer by mainly two mechanisms namely the blockage of the oncoming fluid and the strong shear at the jet slopes. The blockage leads to flow structures similar to the wakes found behind solid obstacles, i.e. periodic eddy shedding close to the wall and a so called horseshoe vortex which wraps around the jet. Furthermore the oncoming flow deflects the jet into the main streamwise direction which leads to additional flow features unique to jets in crossflow such as the entrainment of near wall fluid layers upwards into the jet wake. Both the induced perturbations of the crossflow behind the jet exit as well as the fluid entrainment evidently depend on the jet parameters. The shear on the slope of the jet on the other hand induces a rotational motion which is prolonged and fortified behind the jet exit and leads to the formation of a large longitudinal eddy above the wall. In order to gain an insight in the fluid dynamics snapshots of the flow field are recorded and evaluated. Figures 3 and 4 depict the flow at the last recorded time step. In case of a skew angle of β=30° two longitudinal vortices establish in the flow and wrap around each other due to the induced velocities in the transversal plane. The development of a high-speed streak close to the wall takes place with an inclination to the centerline of the flow. The two longitudinal vortices seem not subject to instabilities at that point in time. They stretch rather well-defined downstream and upwards along the plate. The oncoming freestream deflects the jet in both spanwise and wall-normal direction. Thus a strong shear layer develops on the top side of the jet. The sweep of the jet over the boundary layer results in the near-wall fluid layers being en-

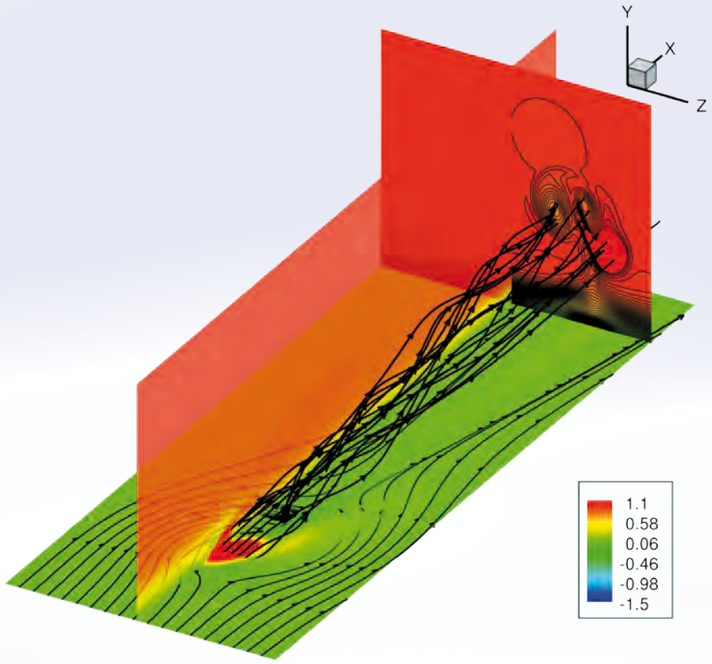

Figure 3: Instantaneous downstream velocity locity and streamlines (β=30°)

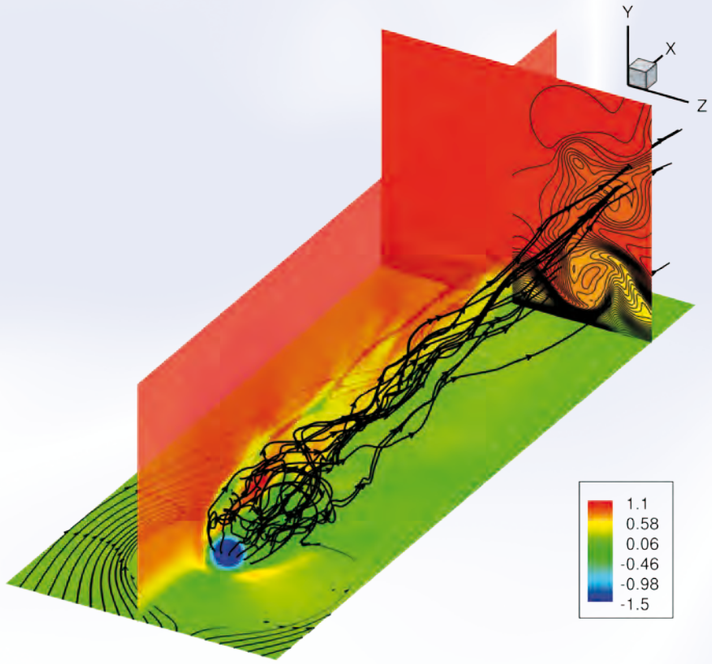

Figure 4: Instantaneous downstream velocity and streamlines (β=150°)

trained upwards into the trajectory of the jet. The entrainment region does not extent very far downstream though and the exchange of high- and low-speed layers inside the boundary layer is not very strong. A different flow can be observed when the jet is blown against the main flow. Firstly the blockage is a lot stronger than in the previous case. Secondly the wake exhibits instabilities which lead to a mixing of the fluid layers in all directions. Again the jet wake gets deflected but additionally the entrainment region extents further downstream and the wake rolls into a longitudinal vortex.

Boundary Layer Control

As already mentioned, Jet-in-Crossflow configurations are investigated as a means to suppress boundary layer separation. This is to be achieved by moving faster fluid layers closer to the wall and therewith increasing the wall friction. Figure depicts a comparison of the time-averaged wall-shear-stress distribution for the simulated cases. Both configurations lead to a net increase of wall shear stress in a confined stripe behind the jet. This area overlaps the trajectory of the jet only marginally though in case of a jet angle of β=30°. The additional energy is fed to the boundary layer in an overly straightforward fashion by adding momentum. In order to exploit such a mechanism with reasonable input-to-output ratio a tangential jet seems more advisable. On the other hand when the jet is directed against the main flow the region of increased shear extends promisingly in both downstream and spanwise direction. The reason being the different fluid dynamics at work. Firstly the jet acts as turbulator which leads to fuller velocity profiles compared to the unperturbed flow and furthermore the jet-crossflow interaction leads to the formation of a longitudinal vortex which transports fast fluid layers closer to the wall. Therefore this configuration seems feasible for active flow control devices even in already fully turbulent shear flows as they are found on aircraft airfoils and flaps.

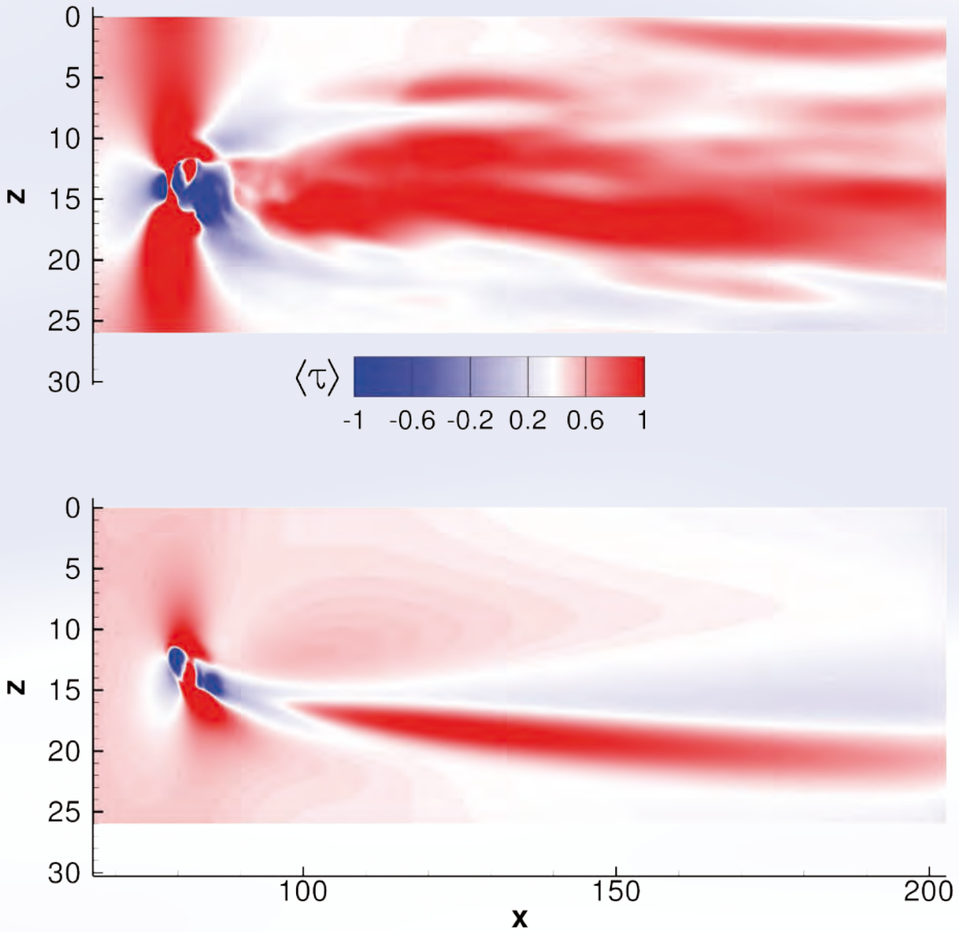

Figure 5: Comparison of mean wall shear distribution in actuated flow (top: β=150°, bottom: β=30)

Conclusion

Simulations of two different Jet-in-Crossflow configurations have been carried out in order to investigate the effect on a boundary layer at *Ma*=0.25. The jets have been skewed and inclined by values typically found in active-flow-control device setups. The simulations gave insight into the generation of a jet-vortex system and wake perturbations. In case of a jet aligned with the mean flow a stable flow established in which the boundary layer was only poorly energized by the jet. A skew of the jet against the main flow led to the generation of a longitudinal vortex which in turn led to significantly increased wall friction. Following simulations may include further variation of the pitch and skew angles of the jet as well as velocity ratio variation. Also of interest is a change of initial conditions towards a fully developed turbulent boundary layer.

Acknowledgements

We thank the High Perfomance Computing Center Stuttgart (HLRS) for provision of supercomputing time and technical support within the project “LAMTUR”.

References

[1] **Johnston, J. P., Nishi, M.**
“Vortex Generator Jets – a Means for Flow Separation Control”, AIAA Journal, 28(6), pp. 989–994, 1990

[2] **Babucke, A., Kloker, M., Rist, U.**
“Direct Numerical Simulation of a Serrated Nozzle End for Jet-noise Reduction”, in M. Resch (Eds.), High Performance Computing, in Science and Engineering 2007, Springer, 2007

[3] **Casper, M., Kähler, C. J., Radespiel, R.**
“Fundamentals of Boundary Layer a Control with Vortex Generator Jet Arrays”, in 4th Flow Control Conference, number AIAA, pp. 2008-3995, 2008

- Björn Selent
- Ulrich Rist

Institute of Aero- and Gasdynamics, University of Stuttgart

Applications

Numerical Simulation and Analysis of Cavitating Flows

Cavitation is a phenomenon which occurs in liquid flows if the pressure decreases locally below the vapor pressure. The liquid evaporates and a cavity of vapor is generated. With respect to fluid dynamics within hydraulic machines such as injection nozzles, pumps or ship propellers, local decompression of the liquid towards or even below vapor pressure commonly occurs [3]. Thereby, the liquid partially evaporates and completely alters its thermodynamic and fluid dynamic properties. Aside from harmful characteristics like cavitation erosion and noise production an increasing number of modern applications benefit from controlled cavitation. Modeling, simulation and analysis of cavitating flows is a challenging task due to the strongly increased complexity as compared to pure liquid or pure gas flows. The main challenge is to take compressibility effects of the liquid and of the vapor into account. This is necessary in order to predict the formation and propagation of violent shocks due to collapsing vapor bubbles. Furthermore, the physics of cavitating flows contain multiple spatial and temporal scales that need to be sufficiently resolved by the numerical method.

To investigate the fundamental dynamics of cavitation we compute the cavitating flow around a standardized shaped hydrofoil (NACA 0015) within a rectangular test-section [1]. The simulations were performed by applying the in-house CFD-Tool CATUM (CAvitation Technische Universität München) which is described in [4]. This tool relies on the Finite Volume Method and operates on block structured meshes. It is written in FORTRAN 95 and it is parallelized with MPI. A special feature of CATUM is its low Mach number consistency which is necessary to simulate wave propagation with 1500 m/s even for very low speed liquid flows. To reduce the computational effort a grid sequencing technique is applied as follows. The simulation is initiated on a coarse grid consisting of 2 x 10^5 cells.

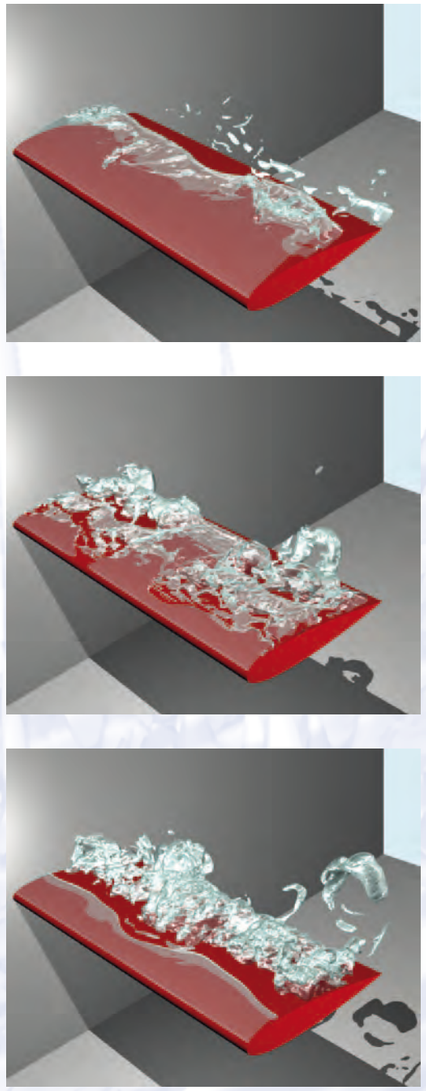

Figure 1: Three instances in time of the predicted shedding of sheet and cloud cavitation around a NACA 0015 hydrofoil (dt_{pic} = 3.4 · 10^{-3}s). The iso-surfaces represent the vapor volume fraction of α = 5%, that means 5% of the computational cell volume is covered by vapor. (The images were rendered with Blender using 8 CPU cores on the remote visualization server.)

The solution of this simulation is interpolated onto a refined grid and provides initial conditions for the next run on the refined grid. This grid sequencing procedure is repeated two times until the finest grid is reached. The finest grid consists of 2.4 x 10^7 cells, 250 in spanwise direction and 600 around the hydrofoil. The complete computation required about 10^5 CPU hours and was performed using 64 to 192 CPU cores on the Linux-Cluster of the Leibniz-Rechenzentrum (LRZ). In order to resolve wave dynamics we apply numerical time steps of 8.5 x 10^{-8}s. Every 1000th time step was stored on disk, giving a total of 480 data sets.

Only this high resolution in time allows us to investigate the propagation of shock waves in the flow field in the postprocessing phase. Due to the high resolution in time and space, 700 GB of data where produced. To analyze this amount of data the new remote visualization server of the LRZ is used. This machine is directly connected to the high performance file system of the compute servers and, therefore, no transfer of data to a local machine is necessary.

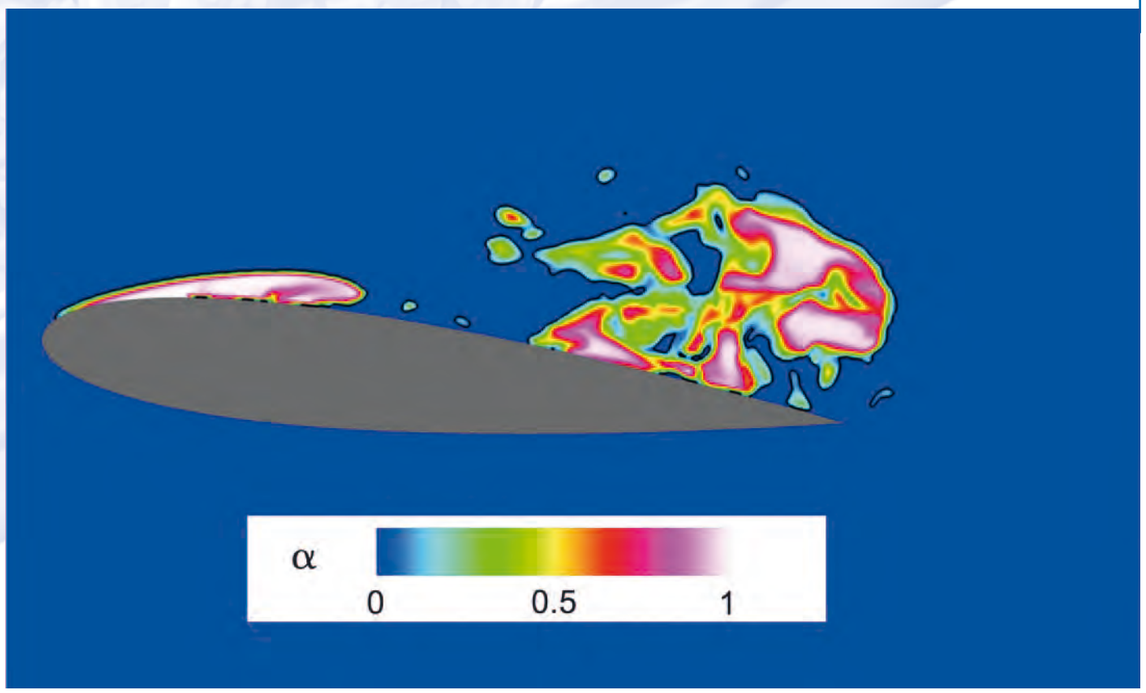

Figure 2: Vapor volume fraction α within a cavity at a cut plane at midspan position. The detached cavity shows a complex inner structure.

Additionally, the throughput rate is significantly higher as the one of a local disc which is crucial for time dependent flow calculations. For data analysis, Visit (using up to 24 CPU cores) and Tecplot (up to 8 CPU cores) on the remote visualization server were used to extract flow features like iso-surfaces of the vapor volume fraction α, slices of the flow field and flow vectors.

The grid sequencing technique enables an investigation of grid dependency of the solutions. We observe that large scale structures like the shedding of vapor clouds are grid independent and the cavity length and the shedding frequency remain constant for all grids. However, the flow computed on the finest grid exhibits more small scale structures like horseshoe vortices and vortices in spanwise direction as shown in Figure 1. A detailed study of the inner structure of the vapor regions where carried out on the finest grid. Figure 2 shows the vapor fraction at a cut plane at midspan position of the hydrofoil at one instance in time. It is observed that the attached cavity at the leading edge consists of almost pure vapor (white color), whereas the detached cloud exhibits a complex inner structure with strong variations of the vapor content. Each time when such a cloud collapses strong shocks are generated. Figure 3 depicts a collapse directly at the trailing edge of the hydrofoil which generates a maximum pressure of p = 225 bar. These pressure peaks at the surface are responsible for the damage of the material which is known as cavitation erosion. By recording the impact positions and the maximum pressures at the surface of the hydrofoil during the whole simulation we obtain detailed information of erosion sensitive areas [2] as shown in Figure 4. We observe instantaneous maximum loads up to 2500 bar. Our investigation confirms experimental observations that the regions close to the trailing edge are highly vulnerable to cavitation erosion.

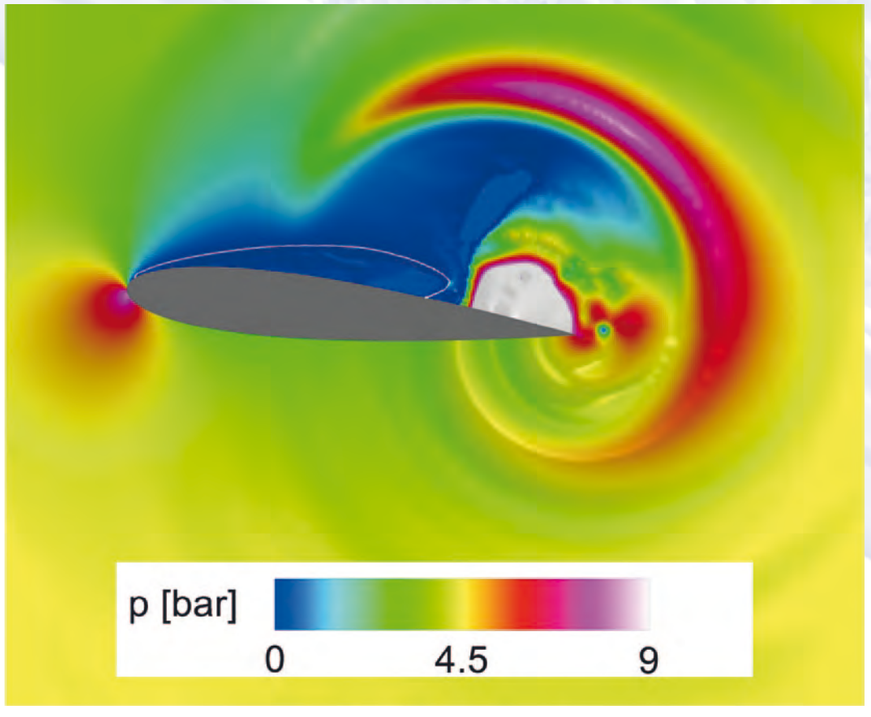

Figure 3: Shock formation and propagation due to collapsing vapor clouds. The collapse directly at the trailing edge generates a pressure of p = 225 bar.

Acknowledgments

Our investigations on numerical simulation of cavitating flows are supported by the KSB Stiftung Frankenthal, especially the prediction of erosion sensitive areas. We also would like to thank the Deutsche Forschungsgemeinschaft (DFG) for supporting further investigations on cavitation in injection nozzles of automotive engines. The support and provision of computing time by the LRZ is also gratefully acknowledged.

References

[1] **Schmidt, S. J., Thalhamer, M., Schnerr, G. H.**
"Inertia Controlled Instability and Small Scale Structures of Sheet and Cloud Cavitation", Proceedings of the 7th International Symposium on Cavitation, CAV2009 - Paper No. 17

[2] **Schmidt, S. J., Thalhamer, M., Schnerr, G. H.**
"Density Based CFD-Techniques for Simulation of Cavitation Induced Shock Emission", Proceedings ETC - 8th European Turbomachinery Conference, Graz, Austria, pp. 197-208, 2009

[3] **Sezal, I. H., Schmidt, S. J., Schnerr, G. H., Thalhamer, M., Förster, M.**
"Shock and Wave Dynamics in Cavitating Compressible Liquid Flows in Injection Nozzles", Journal of Shock Waves (Springer), DOI 10.1007/s00193-008-0185-3, 2009

[4] **Sezal, I. H.**
"Compressible Dynamics of Cavitating 3D Multi-Phase Flows", PhD Thesis, Technische Universität München, 2009

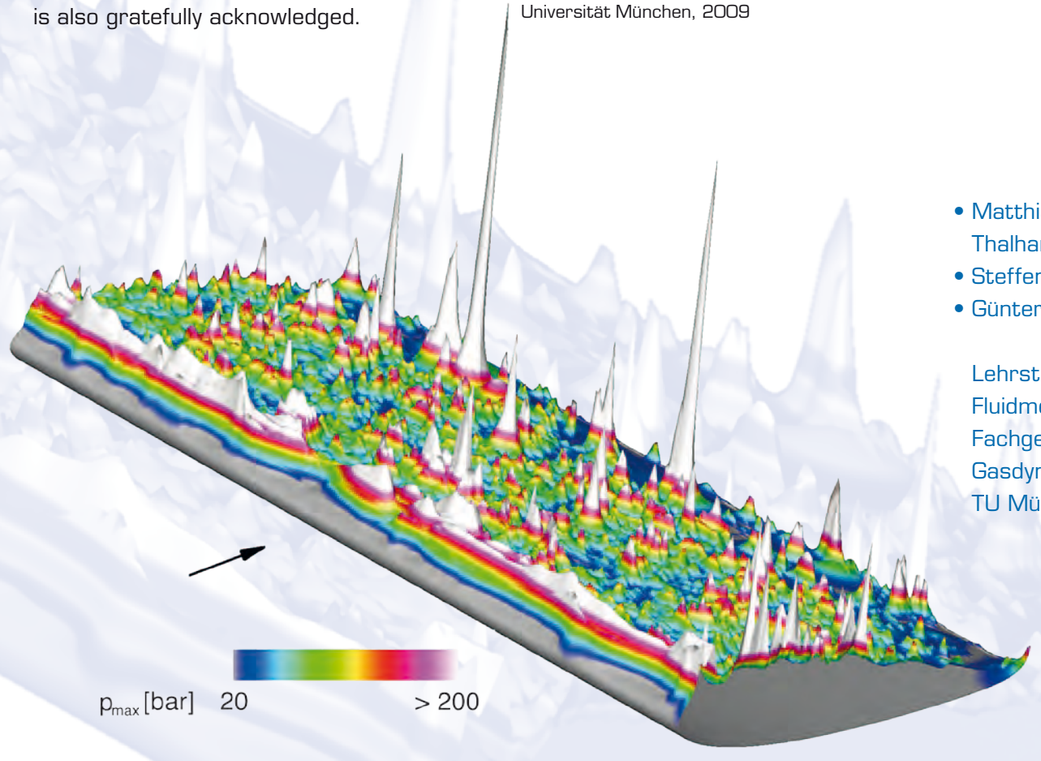

Figure 4: Instantaneous maximum loads of up to p = 2500 bar at the surface of the hydrofoil during the complete computation. High peaks indicate areas where cavitation erosion is likely to occur.

• Matthias Thalhamer
• Steffen J. Schmidt
• Günter H. Schnerr

Lehrstuhl für Fluidmechanik, Fachgebiet Gasdynamik, TU München

Applications

Golden Spike Award by the HLRS Steering Committee in 2009

Modelling Convection over West Africa

In the African Sudanian (10°N-15°N) and Sahelian climate zones (15°N-18°N) convective systems play a key role in the water cycle, because they contribute to about 80% to the annual rainfall. The convective systems are thunderstorm complexes with a horizontal extent of several hundreds of kilometers. They are part of the West African monsoon (WAM), which also impacts the downstream tropical Atlantic by providing the seedling disturbances for the majority of Atlantic tropical cyclones [1].

The WAM system is characterized by the interaction of the African easterly jet (AEJ), the African easterly waves (AEWs), the Saharan air layer (SAL), as well as by the low-level monsoon flow, the Harmattan, and the mesoscale convective systems (MCSs). The AEJ can be observed between 10-12°N and at a height of about 600 hPa and has a typical wind speed of about 12 m s^{-1}. The AEJ develops as a result of the reversed meridional temperature gradient due to the relatively cool and moist monsoon layer and the hot and dry SAL, which is located above the monsoon layer.

The AEWs are synoptic-scale disturbances that propagate westwards across tropical West Africa toward the eastern Atlantic and the eastern Pacific. The AEWs are characterized by propagation speeds of 7-8 m s^{-1}, a period of 2-5 days, and a wavelength of about 2,500-4,000 km. Important parameters for the initiation of convection are the spatial distribution and temporal development of water vapour in the convective boundary layer (CBL). Besides advective processes, water vapor is made available in the atmosphere locally through evapotranspiration from soil and vegetation. Many research findings show that the soil moisture exerts greater influence on the CBL than vegetation.

In the scope of the African Monsoon Multidisciplinary Analyses (AMMA) project we investigate the development of MCSs, the sensitivity of their life cycle to differing surface properties, the role of larger-scale weather systems (AEWs, the Saharan Heat Low) in their evolution, and the development of tropical cyclones out of such systems. In addition, we study the interaction of the SAL with African monsoon weather systems.

To simulate the weather systems over West Africa we use the COSMO (COnsortium for Small scale MOdelling, **www.cosmo-model.org**) model [2,3]. COSMO is an operational weather forecast model used by several European weather services, e.g. the German Weather Service (DWD). Additionally, we use COSMO coupled with the aerosol and reactive trace gases module (COSMO-ART) to investigate the interaction of the SAL with WAM systems. COSMO-ART [4,5] was developed in Karlsruhe and computes the emission and the transport of mineral dust. We used the computer facilities at the HP XC 4000 at the Steinbruch Centre for Computing (SCC) as the COMSO model requires substantial supercomputer resources. In the following, we focus on two different topics. On the one hand the life cycle of an MCS over West Africa is modelled with respect to the soil conditions. On the other hand the focus lies on comparison between the simulated MCSs over West Africa and over the Eastern Atlantic.

The first part of this study investigates the sensitivity of the life cycle of an MCS to surface conditions [6]. The analysis is based on simulations of a real MCS event on 11 June 2006 which occurred in the pre-onset phase of monsoon when vegetation cover is low and the impact of soil moisture is assumed to be dominant. Different conditions for soil moisture were applied for initialization of the soil model TERRA-ML. The run based on the COSMO soil type distribution and on original ECMWF fields was denoted with MOI.

However, comparison with AMSR-E satellite data showed that the MOI field contained too much soil moisture in the upper surface layer. Therefore, we reduced the volumetric soil moisture content in all layers by 35% compared to the initial conditions of MOI. This resulted in a similar soil moisture content in the uppermost level of TERRA-ML, compared to the soil moisture values of about 12% derived by the AMSR-E satellite and in-situ measurements of about 18% for the uppermost 5 cm taken at Dano (3°W and 11°N) [7] for the region around 11°N, where the MCS was observed. The corresponding simulation was designated as CTRL experiment. To eliminate the effect of spatial soil moisture variability on the initiation of convection, an additional simulation with a homogeneous (HOM) distribution of soil moisture and soil texture was performed. In this case the volumetric soil moisture was specified as a mean value of the volumetric soil moisture in the CTRL experiment along 11°N from 4.5°W to 4.5°E. To investigate the effect of dry regions on convective systems, where the soil moisture structure is less complex than the conditions present in the CTRL run, a dry band of 2 degree longitudinal extension was inserted into the homogeneous soil moisture field. In this band the volumetric soil moisture was reduced by 35% compared to the homogeneous environment. The corresponding experiment was denoted as BAND.

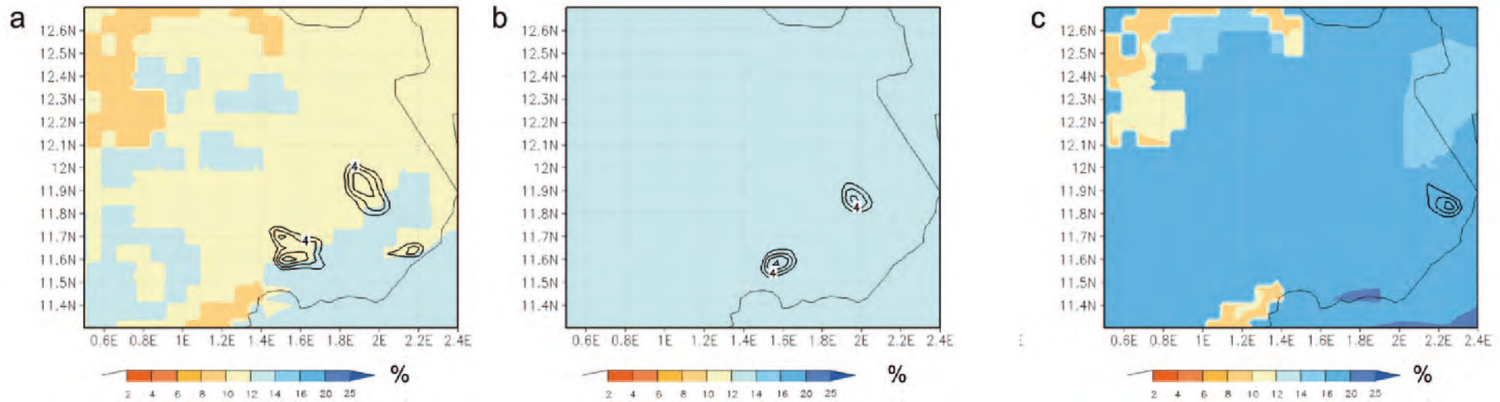

Figure 1: Volumetric soil moisture in % at 15 UTC in the uppermost layer (color shaded) and precipitation in mm h^{-1} (isolines, interval 2) on June 11, 2006 at 17 UTC for CTRL (a), HOM (b), and MOI (c) case. Taken from [6].

In the CTRL case three separate cells were initiated in the south-eastern part of Burkina Faso. Precipitation of up to 6 mm h^{-1} was simulated at 17 UTC (Figure 1a). The south-western most cell developed in the lee of an area with orographically induced upward motion. Triggering of convection often occurs in this way in West Africa. Two further cells developed in the east, at 1.9°E and 11.9°N and at 2.2°E and 11.6°N. In Figure 1a the precipitation pattern of the CTRL case at 17 UTC is overlaid on the soil moisture distribution in the uppermost layer at 15 UTC. This figure shows that all three cells developed in the transition zone from a wetter to dryer surface, while the centres of the precipitating cells were positioned over the dryer surface. In comparison to the CTRL case, only two precipitating cells had developed at 17 UTC in the HOM case (Figure 1b).

These two cells were observed at roughly the same locations as the most intensive cells in the CTRL case. However, the precipitation of both cells was less intense than that of the CTRL case at the same time. In the MOI case the favorable conditions with high convective available energy (CAPE) and low convective inhibition (CIN) values were more limited in space than in the other cases. In addition, the surface temperature in the region of interest in the MOI case was about 3°C lower than in the CTRL case. Under these conditions only one weak precipitating cell had developed at 17 UTC.

Once triggered, the convective cells developed quickly in the CTRL and HOM case and moved with the AEJ towards the west. About two hours after its initiation, the cells had already organized into an MCS in the CTRL run. In the HOM case three separate cells could still be distinguished at 19 UTC, which were less intense than in the CTRL case. The MOI run showed only weak convective activity. The impact of the dry band with a volumetric moisture content of 8.3%, surrounded by a homogeneous moisture content of 12.7% on the modification of a mature MCS is shown in Figures 2a and 2b. Precipitation was significantly reduced between 0 and 2°W, i.e. shifted by about one degree to the east of the dry band. Precipitation was interrupted, when the MCS approached the dry band, but re-generated when the system reached the western part of the band (Figure 2b). Convection was also initiated in the western part of the dry band (around 3°W) at about 19 UTC. The cloud cluster was accompanied by significant rainfall and moved to the west. At 18 UTC an area with lower CAPE values had developed within the dry band, while CIN increased to the east and west of it (Figure 2c). East of the band (0.83°W, 12.2°N) CIN further increased during the subsequent hours. Inside the band (1.8°W, 12.2°N) CIN was significantly lower. The lower CAPE inside the dry band resulted mainly through lower near-surface humidity. East of the dry band a lower near-surface temperature led to a higher CIN, which inhibited the convection of the approaching MCS. The increase of CIN inside the dry band later that night yielded from the nocturnal decrease of near-surface temperature and the passage of the MCS in the surrounding.

The simulations showed that convection was initiated in all model experiments, regardless of the initial soil moisture distribution. The area where

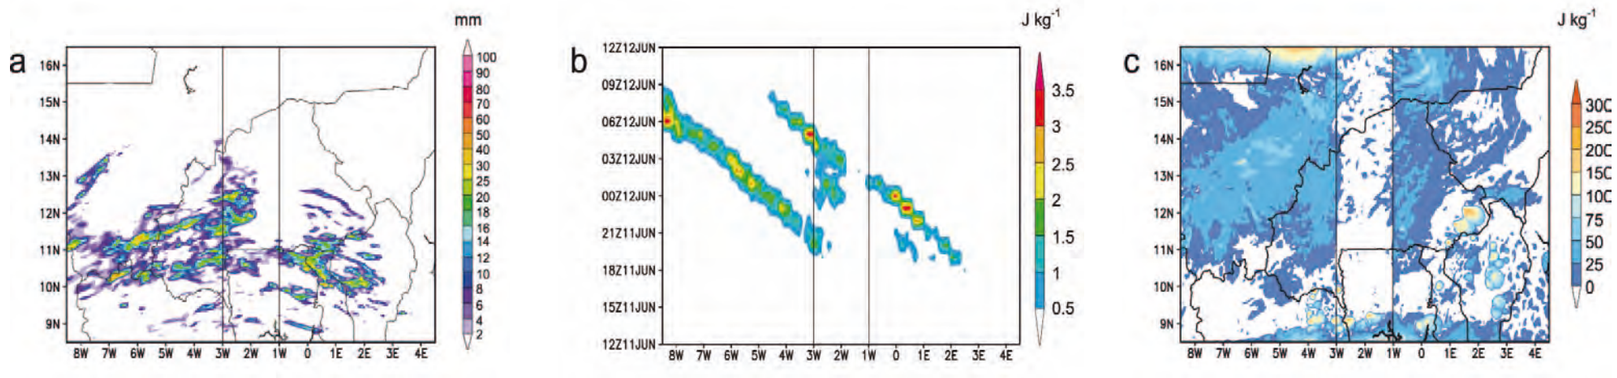

Figure 2: 24-h accumulated precipitation in mm (color shaded) starting from 06 UTC on 11 June 2006 (a), Hovmöller diagram of precipitation in mm h^{-1} (color shaded) averaged between 10.5°N and 13.5°N on June 11 and 12, 2006 (b), and CIN (c) in J kg^{-1} (color shaded) on June 11, 2006 at 18 UTC for BAND case. The solid lines enframe the area of the dry band. Taken from [6].

convection was initiated in the simulations corresponded roughly with the observations. In the CTRL case all three cells were initiated along soil moisture inhomogeneities and shifted towards the dry surface. Triggering of convection and optimal evolution was simulated in areas with low CIN, high CAPE and low soil moisture content (< 15%) or soil moisture inhomogeneities. In CTRL and HOM the cells developed quickly and merged into organized mesoscale systems while moving westwards. The MCS in the CTRL run experienced a significant modification. The precipitation disappeared when the MCS reached a region which was characterized by high CIN values (>150 J kg^{-1}), reduced total water content and soil moisture inhomogeneities.

In conclusion, triggering of convection on this day was favoured by drier surfaces and/or soil moisture inhomogeneities, while a mature system weakened in the vicinity of a drier surface. This means that a positive feedback between soil moisture and precipitation existed for a mature system whereas a negative feedback was found for triggering of convection. The two mechanisms are illustrative for the complexity of soil-precipitation feedbacks in an area where high sensitivity of precipitation on soil moisture was proven.

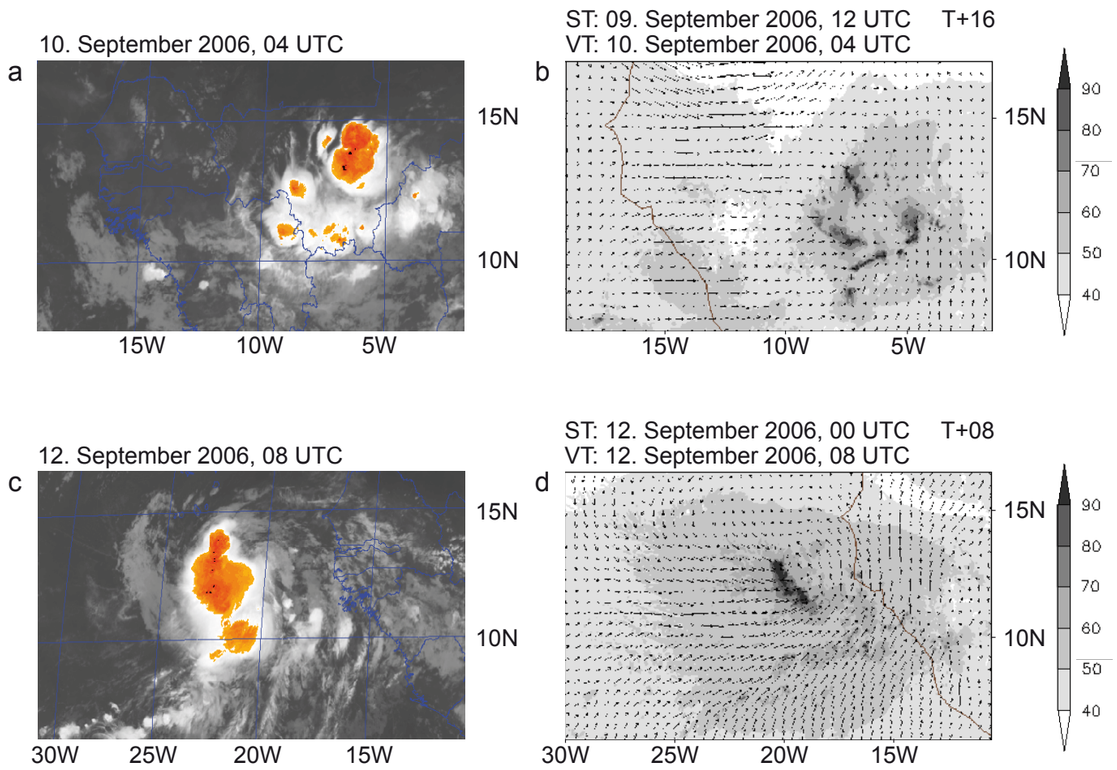

Figure 3: Comparison between the Rapid Developing Thunderstorm Product (RDT) (left) and COSMO model runs (right) on September 10, 2006 at 04 UTC (upper row) and on September 12, 2006 at 8 UTC (lower row). Left: Convective objects are superimposed over the Meteosat infrared images using shadings of grey up to -65°C, orange–red colors between -65° and -81°C, and black above -81°C in the RDT images. Courtesy of Météo France. Right: The vertical integral of cloud water, cloud ice and humidity (kg m^{-2}), indicating the convective updraught cores, from the 2.8-km COSMO run which was initialized on September 9, 2006 at 12 UTC. Taken from [10].

The second study investigates convective systems over West Africa and the eastern Atlantic. The convective systems are embedded in the AEW out of which a hurricane developed. In the afternoon hours on September 9, 2006 an MCS was initiated over land ahead of the trough of this AEW. The convective system increased quickly in intensity and size and developed three westward moving arc-shaped convective systems (Figure 3a,b). The near-surface winds depict the mostly westerly monsoon inflow and the weak easterly inflow into the system. As the system decayed, new convective bursts occurred in the remains of the old MCS. This lead to structural changes of the form of a squall line crossing the West African coastline at around midnight on September 11, 2006. During the next 24 hours, intense convective bursts occurred over the eastern Atlantic. These convective bursts were embedded in a cyclonic circulation which intensified and became a tropical depression on September 12, 2006, 12 UTC out of which Hurricane Helene developed. The largest and longest lived convective burst in this intensification period (Figure 3c,d) was analyzed here and compared to the convective system over land. The convective systems modify their environment and these changes can be related to the structure of the system itself.

A series of COSMO runs over large model domains (1000 x 500 gridpoints) with 2.8 km horizontal resolution were carried out such that the model area was centred around the MCS. As the system moved across West Africa the position of the model region was adjusted for each subsequent run. This enabled us to identify structural features of convective systems over West Africa and the Eastern Atlantic, and to analyze and discuss their differences. All the runs are 72 h in duration. The parameterization of convection is switched off. The model source code was adapted to provide information for moisture, temperature and momentum budgets [8,9,10].

The model was able to capture the above described convective systems as well as their structure and intensity changes very well. We identified three stages in the life cycle of the MCS over West Africa: the developing, the mature, and the decaying phase. To analyze the structure of these convective systems in more detail, an east-west cross section (Figure 4a) is drawn through the convective system shown in Figures 3a,b. During the mature phase, low-level convergence occurs between a strong westerly inflow and an easterly inflow from behind the system and is partly enhanced due to the descending air from the rear system. The associated ascent region extends up to about 200 hPa and is tilted eastwards with height (Figure 4a). The convergence continues up to 700 hPa. At this stage strong downdraughts have developed around 7.5°W just behind the low-level convergence zone that is located under the tilted updraught region. The area

of the westerly inflow reaches deep into the system up to about 700 hPa and the AEJ has its maximum around 600 hPa. Divergence associated with the upper-level outflow can be seen at around 200 hPa with very strong easterly winds ahead of the system and westerly winds in the rear. Thus, there is strong mid-level convergence east of the tilted updraught.

The cross section through a convective system embedded in the developing tropical depression over the eastern Atlantic is shown in Figure 4b. A major difference between the MCS over West Africa and the convective bursts that are embedded in the circulation over the Atlantic is their lifetime. The MCS over land lasted for about 3 days and the succession of MCSs over the ocean only for about 6 to 24 hours. The region of maximum heating and ascent is vertical and not tilted as in the convective system over land. The AEJ seems to occur at around 700 hPa west of the convective system. This is about 100 hPa lower than for the MCS over West Africa where the air is accelerated due to air that exited the updraught core at lower levels and slightly descends. It is also apparent from the cross section that the downdraughts are not as strong as for the MCS over land. Furthermore, the low-level convergence covers a much broader region than over the continent.

To quantify the difference between the convective systems over land and over water and to assess the influence of the convective systems on the environment, potential temperature, relative vorticity, and potential vorticity budgets for regions encompassing the convective region were calculated. Details and the results for the budget calculations can be found in [10]. In future studies the analysis could be applied to other cases in order to generalise these results.

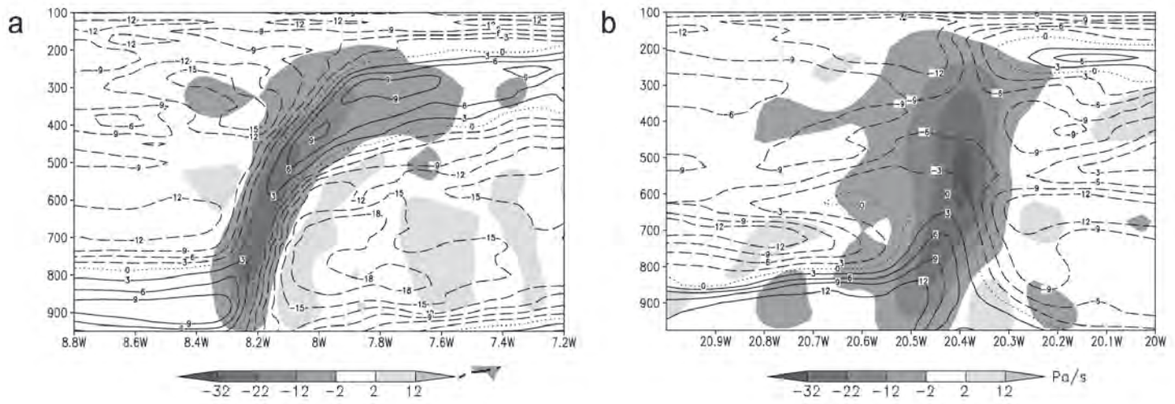

Figure 4: Cross sections through the convective system over land (a) and over the ocean (b). The vertical velocity (shaded) and zonal wind (contour interval is 3 m s^{-1}) are shown. The cross section (a) is along 13.1°N on September 10, 2006, 05 UTC, and (b) along 12.8°N on September 12, 2006, 09 UTC. Taken from [10].

Acknowledgment

This project received support from the AMMA-EU project. Based on a French initiative, AMMA was built by an international scientific group and funded by a large number of agencies, especially from France, UK, US and Africa. It has been the beneficiary of a major financial contribution from the European Community's Sixth Framework Research Program.

Detailed information on scientific coordination and funding is available on the AMMA International web site:

http://www.amma-international.org

References

[1] Avila, L. A., Pasch, R. J.
"Atlantic Tropical Systems of 1993", Monthly Weather Review, 123:887-896, 1995

[2] Steppeler, J., Doms, G., Schättler, U., Bitzer, H. W., Gassmann, A., Damrath, U., Gregoric G.
"Meso-gamma Scale Forecasts using the Nonhydrostatic Model LM", Meteorology and Atmospheric Physics, 82:75–97, 2003

[3] Schättler, U., Doms, G., Schraff, C.
"A Description of the Nonhydrostatic Regional Model LM", part VII: User's guide, Deutscher Wetterdienst, www.cosmo-model.org, 2008

[4] Vogel, B., Hoose, C., Vogel, H., Kottmeier, C.
"A Model of Dust Transport Apllied to the Dead Sea Area", Meteorologische Zeitschrift, 15:DOI: 10.1127/0941–2948/2006/0168, 2006

[5] Vogel, B., Vogel, H., Bäumer, D., Bangert, M., Lundgren, K., Rinke, R., Stanelle, T.
"The Comprehensive Model System COSMO-ART – Radiative Impact of Aerosol on the State of the Atmosphere on the Regional Scale", Atmospheric Chemistry and Physics, 9:14483–14528, 2009

[6] Gantner, L., Kalthoff, N.
"Sensitivity of a Modelled Life Cycle of a Mesoscale Convective System to Soil Conditions over West Africa", Quarterly Journal of the Royal Meteorological Society, 136(s1): 471-482, 2010

[7] Kohler, M., Kalthoff, N., Kottmeier, C.
"The Impact of Soil Moisture Modifications on CBL Characteristics in West Africa: A Case-Study from the AMMA Campaign", Quarterly Journal of the Royal Meteorological Society, 136(s1) : 442-455-, 2010

[8] Grams, C. M.
"The Atlantic Inflow: Atmosphere-land-ocean Interaction at the South Western Edge of the Saharan Heat Low", Master's thesis, Institut für Meteorologie und Klimaforschung, Universität Karlsruhe, Karlsruhe, Germany, March 2008

[9] Grams, C. M., Jones, S. C., Marsham, J. H., Parker, D. J., Haywood, J. M., Heuveline, V.
"The Atlantic Inflow to the Saharan Heat Low: Observations and Modelling", Quarterly Journal of the Royal Meteorological Society, 136(s1): 125-140, 2010

[10] Schwendike, J., Jones, S. C.
"Convection in an African Easterly Wave over West Africa and the Eastern Atlantic": a Model Case Study of Helene (2006). Quarterly Journal of the Royal Meteorological Society, 136(s1):364-396, 2010

- Juliane Schwendike
- Leonhard Gantner
- Norbert Kalthoff
- Sarah Jones

Institut für Meteorologie und Klimaforschung, Karlsruher Institut für Technologie

Driving Mechanisms of the Geodynamo

The Earth's outer core is a region of our planet extending from a depth of 2,900 km to about 5,100 km. This region is not accessible to direct measurements. Geophysicists have at least partially revealed the nature of the Earth's outer core. According to today's knowledge, fluid motions in the Earth's metallic liquid outer core drive a machinery which is called the Geodynamo. Potential energy, which stems from the accretion of the Earth is transformed into kinetic energy and further into magnetic energy by magnetic induction in an electrically conducting fluid.

The liquid outer core of the Earth is assumed to be a mixture composed predominantly of iron and nickel and a small fraction of lighter elements (e.g., sulphur, oxygen or silicon). Two potential sources are available for driving convection and thereby the dynamo in the liquid outer core: density variations due to temperature and due to compositional differences (cf. Jones, 2007 for an overview).

Temperature differences result from secular cooling of the planet. Once the planet has sufficiently cooled to grow a solid inner core, the latent heat being released in the solidification process, provides another source for thermal convection. Radioactive elements such as potassium (^{40}K) are an additional possible thermal source. From the material physics point of view it remains an open question whether indeed radioactive elements can be present in the core. The additional compositional differences arise from light elements being released from the solidifying material at the Inner Core Boundary (ICB). It is still unclear which driving source dominates core convection and thus the dynamo mechanism. Additional potential sources of dynamo action, but likely less effective, are precession induced flows due to inclined rotation axis as well as tidal forces.

Various numerical dynamo simulations have demonstrated that convective flows in rotating fluids can maintain a self-sustaining dynamo. However, due to limited computational resources, these models come even not close to realistic planetary conditions. On the other hand they are able to produce magnetic fields which show Earth-like features. Typically, numerical dynamo simulations do not take into account the effects that may arise from the presence of two buoyancy sources with significantly different diffusion time scales. Both driving sources differ in terms of their respective molecular diffusivities. The Lewis number $Le=\kappa_T/\kappa_C$ gives the ratio of thermal diffusivity κ_T to compositional diffusivity κ_C. For the Earth's outer core this ratio is estimated to be of $O(10^3)$. This implies that thermal instabilities diffuse much faster than compositional ones. It is well known that the spatial and temporal behaviour of convection depends strongly on the diffusivity of the driving component (e.g., Breuer et al., 2004).

For this reason we study thermo-chemically driven core convection and dynamos in rotating spherical shells by means of numerical models. Considering different thermal to chemical forcing ratios we hope to get a better general understanding of convection and magnetic field generation in the Earth's core, with regard to the observed spatial and temporal variations. These simulations are computationally very demanding, since, both, high spatial grid resolutions as well as long time-series are required. Therefore, high performance clusters are an essential working tool for such an approach and efficient parallel codes are required.
In the following sections we will give a short description of the applied model and numerical methods and present some representative results of our work.

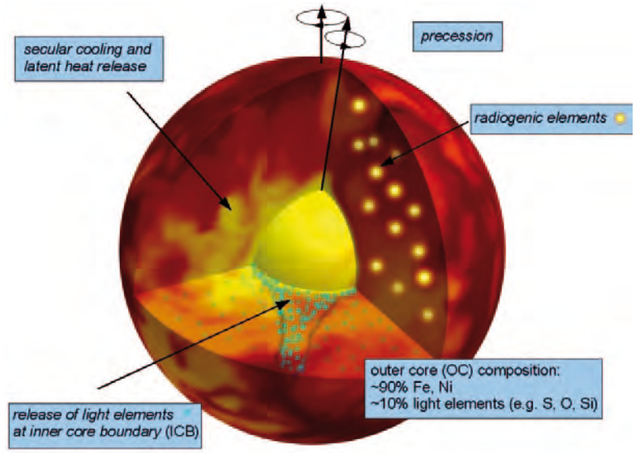

Figure 1: Illustration of potential driving sources for flows in the Earth's fluid outer core.

Mathematical Model and Numerical Methods

For modeling core convection and the Geodynamo we consider a thermally and chemically driven flow within a spherical shell, rotating with a constant angular frequency. We employ the Boussinesq approximation where density changes are only retained in the buoyancy terms. Furthermore, all material properties in the core are assumed constant. The governing magneto-hydrodynamic equations in non-dimensional form, using d as the characteristic length scale and the viscous diffusion time scale d^2/ν, where ν is the kinematic viscosity, are expressed as:

These are the *Navier-Stokes equation (1a)*, the *equation of continuity (1b)* which ensures incompressibility, the *heat transport equation (1c)*, the *compositional transport equation (1d)*, and the *magnetic induction equation (1e)* with the condition of divergence free magnetic field (1f). Π, T, C are perturbations in pressure, temperature and composition, with respect to an adiabatic well mixed reference state. $\underline{u}$ and $\underline{B}$ are the velocity and magnetic field vector, respectively. $\underline{e_r}$ is the unit vector of position and $\underline{e_z}$ is the unit vector indicating the direction of the axis of rotation.

$$\partial_t \underline{u} + \underline{u}\cdot\nabla\underline{u} = -\nabla\Pi + \nabla^2\underline{u} - 2\,Ek^{-1}(\underline{e_z}\times\underline{u}) + (Ra_\nu^T T + Ra_\nu^C C)\frac{r}{r_O}\underline{e_r} + 1/(Pm\,Ek)(\nabla\times\underline{B})\times\underline{B} \quad (1a)$$

$$\nabla\cdot\underline{u} = 0 \quad (1b)$$

$$\partial_t T + \underline{u}\cdot\nabla T = 1/Pr_T\,\nabla^2 T \quad (1c)$$

$$\partial_t C + \underline{u}\cdot\nabla C = 1/Pr_C\,\nabla^2 C \quad (1d)$$

$$\partial_t \underline{B} - \nabla\times(\underline{u}\times\underline{B}) = 1/Pm\,\nabla^2\underline{B} \quad (1e)$$

$$\nabla\cdot\underline{B} = 0 \quad (1f)$$

The non-dimensional control parameters appearing in Eq. 1(a)-(f) are:

- Ek: **Ekman number**, relative importance of viscous to Coriolis force
- Ra_ν^T: **thermal Rayleigh number**, relative strength of thermal buoyancy
- Ra_ν^C: **compositional Rayleigh number**, relative strength of compositional buoyancy
- Pr_T: **thermal Prandtl number**, ratio of viscosity to thermal diffusivity
- Pr_C: **compositional Prandtl number**, ratio of viscosity to chemical diffusivity

In the present analysis we assume that both temperature and composition provide de-stabilizing buoyancy. Temperature and composition are fixed on $T_i = 1$ and $C_i = 1$ at the Inner Core Boundary (ICB) and on $T_O = 0$ and $C_O = 0$ at the Core Mantle Boundary (CMB). The difference in temperature (composition) between the two boundaries defines the thermal (compositional) scale. For the compositional component the applied boundary condition at CMB is unrealistic. It allows light material to pass from the outer core into the mantle. However for a starting point, we wanted to investigate the influence of different diffusivities isolated from distinct boundary conditions. Regarding the model parameters the most critical one turns out to be the Ekman number *Ek* (a measure for the relative importance of viscous to Coriolis force in the momentum equation Eq. 1(a)). Motions in the Earth's fluid outer core are assumed to be strongly dominated by Coriolis force, expressed by a very low Ekman number of $Ek \approx 10^{-15}$. Theoretical analysis predict that the smallest scales in the convective flow pattern are strongly decreasing with decreasing *Ek*. For this reason global Geodynamo models are so far restricted to Ek > O(10^{-7}). In order to solve the system of partial differential equations Eq. 1 (a)-(f) numerically we apply two methodically distinct numerical codes in dependence of the required grid resolution and computer platform. The first one is based on a finite volume (FV) discretization with an implicit time-stepping scheme and efficient spatial domain decomposition for parallel computations using the Message Passing Interface (MPI) (Harder & Hansen, 2005). This sophisticated method has been developed over years at the Institute for Geophysics in Münster with the intention to reach more Earth-like parameter regimes, i.e., smaller Ekman numbers, by means of High Performance Computations on massively parallel computing architectures (>1000 nodes) like Blue Gene at Jülich Supercomputing Centre (JSC). The second code is based on a pseudo-spectral method with a mixed implicit-explicit time stepping scheme (Dormy et al., 1998). Pseudo-spectral means spectral discretization in longitude and latitude and finite differences in radial direction. Parallelization is implemented via decomposition in radial direction using MPI. These methods are commonly used for Geodynamo simulations since they are very efficient on shared memory architectures with fast processors. Therefore, we use this method for extensive parameter studies in more "moderate" parameter regimes.

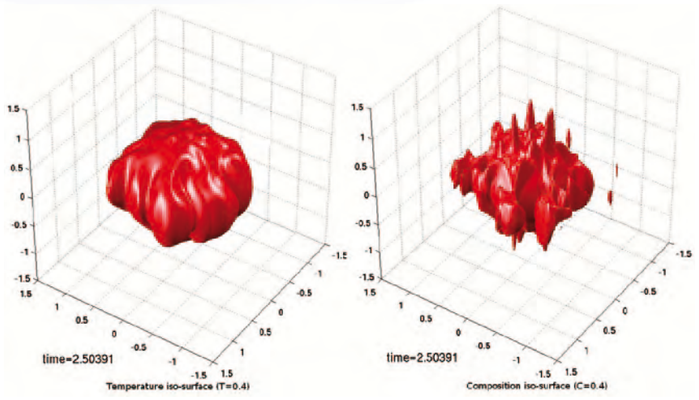

Figure 2: Snapshot of iso-surfaces of temperature field (left) and composition (right) for the predominantly chemically driven case δ = 20%

Parameter Study of Thermo-chemically Driven Convection and Dynamos in a Rotating Spherical Shell at Moderate Ekman Numbers

To investigate the influence of the driving mechanism on the convective flow pattern, we considered different scenarios including the two extreme cases of purely thermally and purely compositionally driven convection and the more likely situations of a joint action of both sources. Since we focus mainly on the influence of the driving mechanism we keep the Ekman number constant at a moderate value of $Ek = 10^{-4}$ for simplicity. The thermal and compositional Prandtl numbers are also fixed with $Pr_T = 0.3$ and $Pr_C = 3.0$, respectively. To parameterize different driving scenarios we define a total Rayleigh number as

$$Ra_\nu^{tot} = Ra_\nu^T + Ra_\nu^C = \delta\, Ra_\nu^{tot} + (1-\delta)\, Ra_\nu^{tot},$$

where $\delta = Ra_\nu^T / Ra_\nu^{tot}$: denotes the fraction of thermal contribution to the total Rayleigh number. We have considered different values of δ including the two end members of purely thermal and purely compositional convection, which merely differ in the Prandtl numbers. For the results presented in this section, the value of Ra_ν^{tot} is kept constant at a value of 4.8×10^6. One has to keep in mind that this fixed total Rayleigh number implies unequal super-criticality for the different scenarios defined by δ. This is due to the different diffusivities of the thermal and compositional component. For the purely thermally driven case this leads to a super-criticality of $\approx 3 \times Ra^T_{crit}$ and $\approx 16 \times Ra^T_{crit}$ for the purely compositional driven case.

Figure 2 illustrates the difference between spatial structure of the temperature and the compositional field for the predominantly compositionally driven case δ = 20%. The temperature field is much more diffusive compared to the compositional field where coherent small scale up-wellings develop, mainly in the polar and equatorial regions. This is due to the difference in their diffusivities expressed by different Prandtl numbers Pr_T and Pr_C. As a result the overall flow behvaviour depends strongly on the individual strength of thermal and chemical buoyancy in the momentum equation Eq. 1(a). This context is exemplarily illustrated in Figure 3 for the two cases δ = 20%, 80% by means of the spatial distribution of azimuthally and temporally averaged helicity $h = \langle \underline{u} \cdot \nabla \times \underline{u} \rangle_\Phi$ and differential rotation $U_\Phi = \langle u_\Phi \rangle_\Phi$, both being characteristic properties of convection in rotating fluids. These structures are mainly aligned parallel to the axis of rotation reflecting the relative importance of the Coriolis force for the flow dynamics. Helical flows and an azimuthal mean flow component varying with distance from the axis of rotation are thought to be essential for the dynamo mechanism

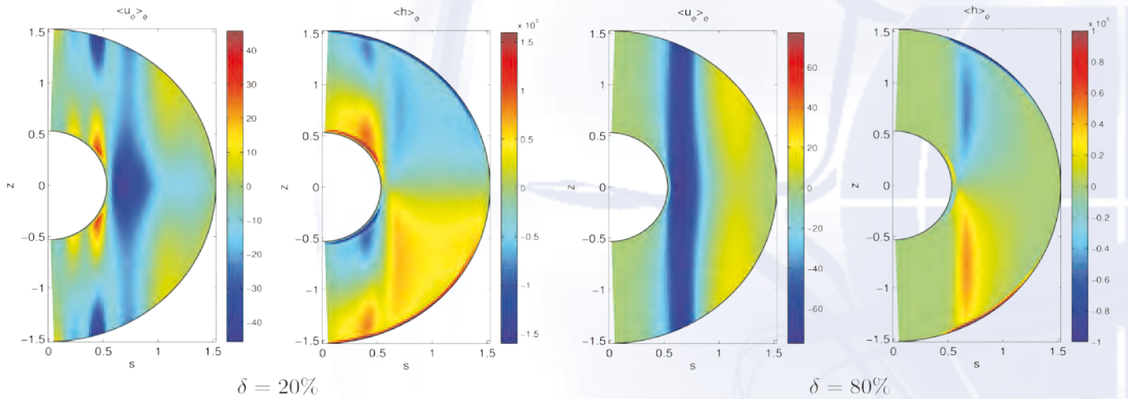

Figure 3: Spatial distribution of azimuthally and temporally averaged differential rotation and helicity for the predominantly chemically driven case δ = 20% (left) and the predominantly thermally driven case δ = 80% (right).

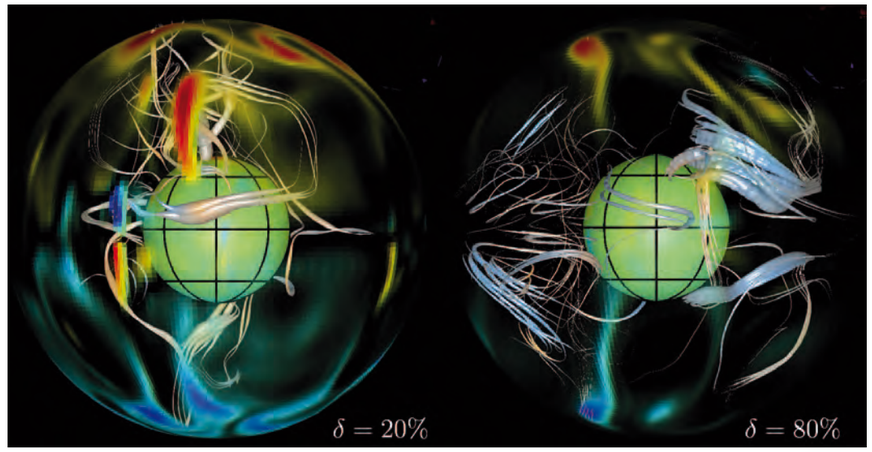

Figure 4: Snapshots of internal magnetic field lines for the predominantly compositionally driven case δ = 20% and for the predominantly thermally driven case δ = 80% (figures created with DMFI (Aubert et al., 2008))

in planetary cores. In this respect the most apparent difference between the two cases δ = 20%, 80% is the strong contribution of differential rotation and helicity in the polar regions in the predominantly compositionally driven case. In the case of dominant thermal forcing these flow structures evolve mainly in small bands (Busse-Taylor columns) outside the tangential cylinder (tangent to the inner core). The importance of these structures for the dynamo mechanism is illustrated in Figure 4 showing snapshots of dominant magnetic field lines for the two distinct driving scenarios at arbitrary times. Here, the width of the field lines reflects the relative local strength of the magnetic field. In the predominantly compositionally driven case magnetic field is mainly generated in the polar regions forced by strong compositional up-wellings in these regions (cf. Figure 2). By contrast, in the mainly thermally forced situation the dynamo process acts primarily in the bulk region outside the tangential cylinder. These two distinct driving scenarios demonstrate quite plainly the strong importance of the considered driving scenario on the flow structure and the dynamo process.

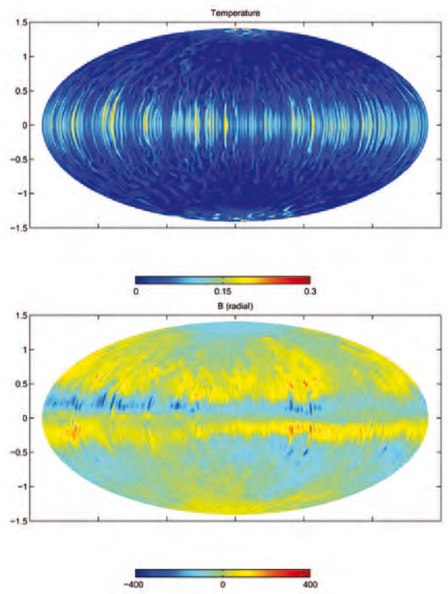

Figure 5: Example of a High Performance Computing dynamo simulation with ≈ $5x10^7$ volume cells and 10^5 timesteps. Shown are snapshots of temperature T (top) and radial magnetic field component B_r (bottom) near the core mantle boundary (CMB) at low Ekman number $Ek = 2x10^{-6}$.

High Performance Computing Application of Low Ekman Number Dynamo Simulation

Since the liquid outer core is characterized by fast rotation and low viscosity, its dynamical behaviour is dominated by an extremely low Ekman number. This puts severe demands on numerical simulations of core flows due to the evolving small spatial scales in the solution. This is demonstrated in Figure 5 which displays temperature and the radial magnetic field component below the core mantle boundary for a low Ekman number of $Ek = 2x10^{-6}$ in a convective flow driven by basal heating. This particular simulation was calculated within the framework of the DEISA (Distributed European Infrastructure for Supercomputing Applications) project 3D-Earth utilizing massively parallel computing (1,920 processes). With such computing power it was possible to calculate a time series of nearly 10^5 timesteps with a spatial resolution of ≈ $5x10^7$ volume cells. This high resolution is needed in the low Ekman number regime to resolve the dominant scales of the solution. As displayed in Figure 5 (top), the temperature field and thus buoyancy is extremely fine-scaled in the equatorial region and drives nearly hundred Busse-Taylor columns. However, the strong driving creates also strong upflows in the tangential cylinder. This vigorous convection inside the tangential cylinder has a strong impact on the generated magnetic field, which is no longer dominated by the dipole component but has strong higher multipole components (Figure 5, bottom). Dipole dominance is restricted to lower forcing, i.e., lower Rayleigh number, where buoyancy is insufficient to drive vigorous convection inside the tangential cylinder by breaking the Proudman-Taylor constraint. This particular result suggests that the observed different flow fields for thermal and compositional convection will have a strong influence on the generated magnetic field. A more detailed discussion of obtained results on thermally driven dynamos is given by Wicht et al. (2009).

Outlook

The main ambition for the next years will be to push the dynamo models continuously towards more realistic parameter regimes. In this respect, we have to answer the question whether the observed distinct flow behaviour dependent on the driving scenario will even hold in a regime with strongly increased weight of the Coriolis force. This is a challenging task and is basically constrained by the computational resources available and their future technical improvements. Therefore, modern High Performance Computer Centres with innovative and substantial technical support are essential for a further progress in understanding the secrets of the Geodynamo.

References

[1] **Aubert, J., Aurnou J., Wicht, J.**
"The Magnetic Structure of Convection-driven Numerical Dynamos", Geophysical Journal International, 172, pp. 945-956, 2008

[2] **Breuer, M., Wessling, S., Schmalzl, J., Hansen, U.**
"Effect of Inertia in Rayleigh-Bénard Convection", Physical Review E, 69, 026302/1-10, 2004

[3] **Dormy, E., Cardin, P., Jault, D.**
"MHD Flow in a Slightly Differentially Rotating Spherical Shell with Conducting Inner Core in a Dipolar Magnetic Field", Earth and Planetary Science Letters, Volume 160, pp. 15-30, 1998

[4] **Harder, H., Hansen, U.**
"A Finite Volume Solution Method for Thermal Convection and Dynamo Problems in Spherical Shells", Geophysical Journal International, 161, pp. 522-532, 2005

[5] **Jones, C. A.**
"Thermal and Compositional Convection in the Outer Core", Vol. 8 of Treatise on Geophysics, chap. 5, pp. 131-185, Elsevier Science, 2007

[6] **Wicht, J., Stellmach, S., Harder, H.**
"Geomagnetic Field Variations", (Eds: K.H. Glaßmeier, H. Soffel and J. Negendank), chap. Numerical models of the Geodynamo: from fundamental Cartesian models to 3D simulations of Field reversals, pp. 107-158, Advances in Geophysical and Environmental Mechanics and Mathematics

• Martin Breuer
• Tobias Trümper
• Helmut Harder
• Ulrich Hansen

Institute for Geophysics, University of Münster

Applications

The Maturing of Giant Galaxies by Black Hole Activity

Extragalactic jets are amongst the most spectacular phenomena in astrophysics: these dilute but highly energetic beams of plasma are formed in the environs of active black holes, moving with speeds very near the speed of light. They run into the hot gas surrounding the galaxy, where they are eventually decelerated and heat up the gas considerably, digging large cavities (the so-called jet cocoon) into the extragalactic gas.

Both the jets and the cavities are observable today with radio and X-ray telescopes on earth or in space. These observations have revealed that the power of these jets is even greater than was thought before: more than one trillion times the total power of our sun (10^{39} watts) for the most powerful ones – with activity durations of some ten million years. This power ultimately originates from the active supermassive black hole of the galaxy, since the jet taps the enormous rotational energy of the black hole. This huge amount of energy clearly has a considerable impact on the energy budget of the respective galaxy and its environment. Since extragalactic jets are most frequent in the most massive galaxies (observed in roughly every third among them), they have become a common explanation for cosmological simulations failing to reproduce giant galaxies correctly, despite their otherwise great success in modeling the evolution of our universe and its galaxies. While astronomical observations show that giant elliptical galaxies consist of old and red stars, showing almost no signs of current formation of new stars, these galaxies still grow in cosmological simulations and therefore are much bluer, actively forming stars and have considerably greater masses. This has caused an increased interest in jet physics and in research projects focusing on the interaction of jets with their host galaxy and its environment, now referred to as "jet feedback". Yet, the detailed processes involved and the exact importance of jets is still unknown. This becomes obvious in two opposing processes that are invoked in the current literature: positive and negative feedback.

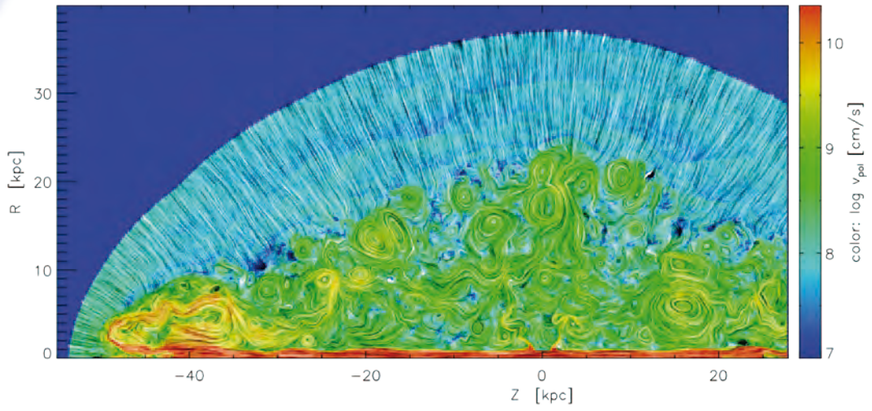

Figure 1: Velocity field of a jet (density contrast 0.001) after 30 million years. The jet beam (in red) reaches out to the jet head and then turns back, inflating a turbulent cocoon (mostly green).

Think Positive – or Negative?

Dense and very cool clouds of gas are the birthplaces of stars since only then can gravity surmount thermal pressure and cause the clouds to collapse, ultimately forming new stars as nuclear fusion sets in. If clouds are hit by a jet or its preceding bow shock, they are compressed by the impact and can become gravitationally unstable. This would result in an increased rate of star formation (positive feedback). On the other hand, the same impact could as well disrupt the cold clouds and thereby destroy the seeds necessary for the formation of new stars (negative feedback). Additionally, the jet would heat up the hot gas surrounding the galaxy to higher temperatures and prevent it from cooling down, getting compressed and forming new cold clouds. Both processes have been demonstrated to be possible, and only a detailed study with a realistic underlying model can help us reveal their real importance.

Observations of distant radio galaxies, showing them as they were 10 billion years ago, give additional hints on the interaction between the jet and the disk. There, extended emission-line nebulae aligned with the jet axis are observed and it is conjectured that these nebulae are actually created by the interaction of the jet with a galaxy's gaseous disk.

Computational Challenges

We have conducted a series of magnetohydrodynamic jet simulations to examine the interaction of jets with their environment at very high resolution assuming axisymmetry. These computations are extremely demanding if realistic parameters are used: although the jet plasma moves with almost the speed of light, its density is much lower than the ambient density of the extragalactic gas (in our simulations typically 1,000 times smaller), which results in a much slower propagation of the jet. Also, a considerable range of spatial scales has to be covered. While the jets extend over more than 100 kpc (1 kpc = 3 x 10^{16} km), the jet beams are 100 times narrower and have to be resolved in detail. All this results in simulations with more than 6 million cells and more than one million time steps necessary. The very large number of time steps makes simulations time-consuming but does not allow a massively parallel approach since the "problem per processor" would become too small then and communication costs between processors would become unwieldy. These simulations were

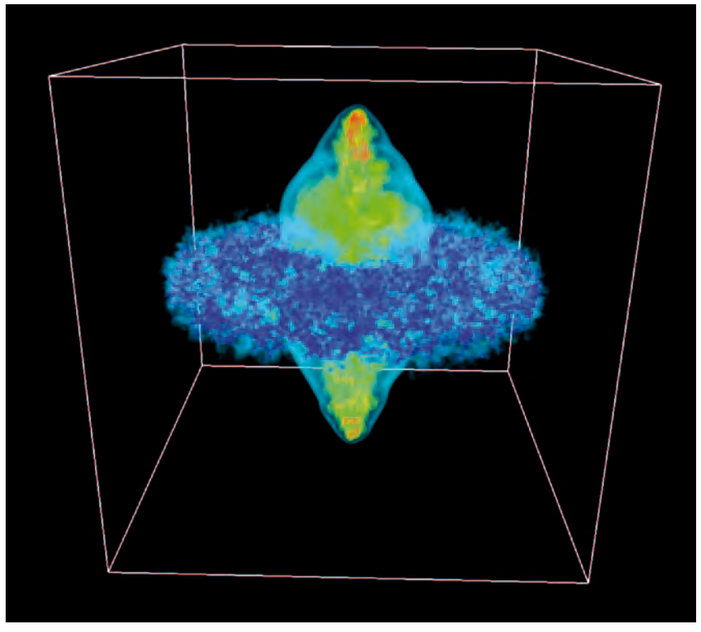

Figure 2: 3D simulation showing the interaction of a jet and a galactic gaseous disk (volume rendering of density)

therefore only possible on a powerful supercomputer such as the NEC SX-8 at the High Performance Computing Center Stuttgart (HLRS), to which we have adjusted and optimized our code. Due to its vector capabilities and shared memory architecture per node, the code ran very efficiently and allowed us to reach realistic sizes for the jets.

We found that the expansion of the jet cocoon follows self-similar models only at early times, but deviates from that significantly once the cocoon approaches pressure balance with its environment. We found this to be consistent with observations and also achieved bow shock shapes and strengths such as those observed typically. Magnetic fields turned out to be important to stabilize the contact surface between the jet and the ambient medium and were moreover found to be efficiently amplified by a shearing mechanism in the jet head.

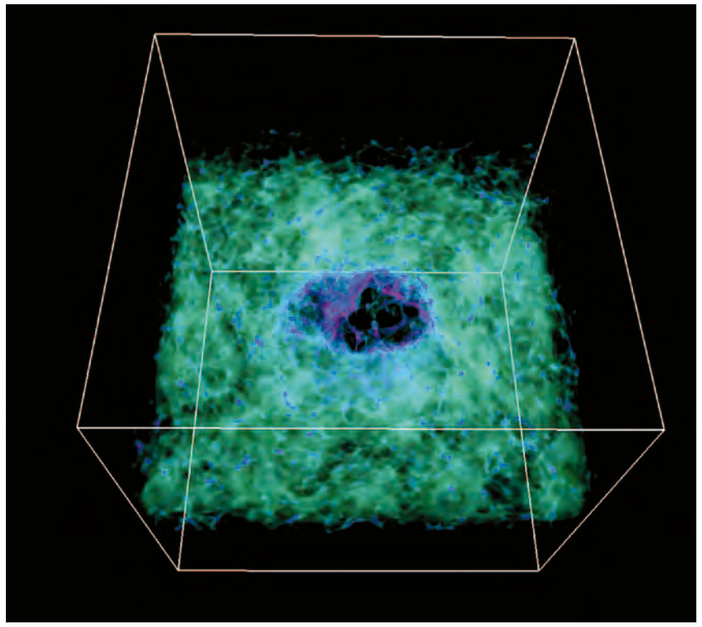

Figure 3: The gaseous clouds in the disk (green) are compressed (blue) by the action of the jet and the remaining gas is pushed outwards.

The Mist of Distant Galaxies

With respect to the riddle of the emission-line nebulae, we were able to test two models for the location and kinematics of the line-emitting clouds against observed properties and found that clouds embedded in the jet cocoon were able to much better reproduce observed velocities and morphologies than clouds embedded in the shocked ambient gas.

About 10 percent of the total jet power was measured as kinetic energy in the jet cocoon, which made it possible to link these findings to simulations of multi-phase turbulence that model the amounts and emission power of the cool gas phase embedded in these turbulent regions. The expected emission line luminosities, however, disagree considerably with the observed range of luminosities and we concluded that the models are still too simple to include all the necessary physics. This was, admittedly, not too surprising since the models relied on clouds passively advected with and spread all across the cocoon plasma, while the interaction of real clouds of cold gas would be considerably more complex.

Moving on – Simulations in 3D

To improve our model of the jet – cloud interaction, we had to extend our simulations to full three dimensions since a clumpy galactic or intergalactic gas cannot be modeled properly within axisymmetry. This also made it necessary to move to another code: RAMSES, an adaptive mesh refinement code, that includes cooling, gravity, star formation, cosmological evolution and magnetic fields. The mesh refinement is highly suitable for resolving small clouds in an otherwise large computational domain, in contrast to a uniformly resolved mesh. The current implementation allows a maximum dynamic range in length scale of 6 orders of magnitude. The code is parallelized by MPI and includes dynamic load-balancing between the different MPI processes. The 3D simulations are clearly computationally very demanding, since an additional dimension results in a much larger number of cells, even if the resolution is somewhat lower. On the other hand, a higher number of cells (in contrast to time steps) generally can be handled by massive parallelization if an efficient method is used.

Outlook

We have successfully tested RAMSES on both the Nehalem cluster at the HLRS and the HLRB-II at the Leibniz-Rechenzentrum (LRZ), and found it to be scaling almost linearly with up to 4,000 cores, if sufficient bandwidth is available. On the Nehalem cluster, the performance was ~50 percent smaller than expected from the single-core performance since memory bandwidth become a bottleneck for more than 4 cores per node; scaling then increased almost linearly for a larger number of nodes. HLRB-II, however, did not suffer from this and showed excellent scaling behaviour. We conjecture that the memory bandwidth bottleneck on the Nehalem cluster may be related to details of the MPI implementation, since MPI adjustments on a local Harpertown Beowulf cluster were able to get around this limitation.

For a small test simulation, we have successfully set up a clumpy galactic gaseous disk, similar to disks actually observed in distant galaxies. A jet is positioned in the center of the galaxy and interaction of the jet with the cold gas clouds embedded in the ambient medium is computed explicitly. Preliminary results indicate that actually both positive feedback by compression of clouds in the center as well as negative feedback by removal of gas along the jet axis may be acting at the same time. So far, the simulations have only covered a time much shorter than our previous simulations, but runs on a large number of processors will allow us to eventually examine the action of jet feedback in detail and over realistically long times.

References

[1] Gaibler, V., Krause, M., Camenzind, M.
"Very Light Magnetized Jets on Large Scales – I. Evolution and Magnetic Fields", Monthly Notices of the Royal Astronomical Society, 400, pp. 1782-1802, 2009

[2] Krause, M., Gaibler, V.
"Physics and Fate of Jet Related Emission Line Regions", Conference contribution, at arXiv:0906.2122

[3] Gaibler, V., Camenzind, M.
"Numerical Models for Emission Line Nebulae in High Redshift Radio Galaxies", Wolfgang E. Nagel, Dietmar B. Kröner, Michael M. Resch (Eds.), Springer 2010, High Performance Computing in Science and Engineering '09

[4] Gaibler, V.
"Very Light Jets with Magnetic Fields", PhD Thesis, Ruprecht-Karls-Universität Heidelberg, 2008

- Volker Gaibler

Max-Planck-Institut für extraterrestrische Physik, Garching

The Physics of Galactic Nuclei

Active Galactic Nuclei are powered by supermassive black holes. Whenever enough gas reaches the centre of otherwise normal galaxies, a fast rotating, thin and hot gaseous accretion disc is able to form and the centre lights up to become as bright as the stars of the whole galaxy. A surrounding ring-like, dusty, geometrically thick gas reservoir (the *torus*) provides mass supply and anisotropically blocks the light. This gives rise to two characteristic observational signatures, depending whether the line of sight is obscured by the torus (edge-on view) or not (face-on view). These nuclear regions of nearby active galaxies, as well as our own Galactic Centre have been observed in great detail with the most up-to-date instruments at the largest available telescopes and interferometers, yielding unprecedented resolution. As the occurring physical processes in these regions are still poorly understood, the aim of our project is to obtain a deep understanding of the origin and evolution of the complex stellar, gaseous and dusty structures and their relation to the observed phases of nuclear activity.

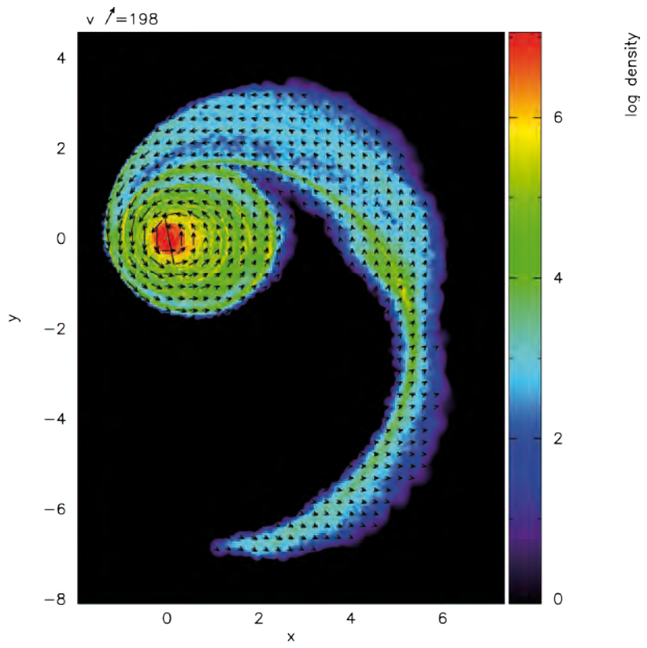

Figure 1: Density (colour coded in solar masses per pc^3) and velocity (overlayed arrows) distribution of the cloud interacting with the central supermassive black hole after 250,000 years

Nuclear Disc Formation in Galactic Nuclei

Recent observations of our Galactic Centre reveal a warped clockwise rotating stellar disc with a mean eccentricity of 0.36 ± 0.06 and a second inclined counter-clockwise rotating disc extending from 0.04 pc to 0.5 pc (e.g. [1] and references therein), made up of roughly 100 young and massive stars with a total mass of roughly $10^{4\text{-}5}$ solar masses [2,3]. Due to the hostility of the environment, they cannot be explained by normal means of star formation, as tidal forces would disrupt typical molecular clouds in the vicinity of the central black hole. We present the first numerical realization of a scenario proposed by [4], where an infalling cloud covers the black hole in parts during the passage and forms a self-gravitating, eccentric accretion disc, before it starts fragmenting. For our simulations we use the N-body Smoothed Particle Hydrodynamics (SPH) Code GADGET2 [5], which shows very good scaling on the Altix 4700 supercomputer at the LRZ. We gain roughly a factor of 1.8 in speed when doubling the number of CPUs. Most of our simulations were done using 128 CPUs, with typically 10,000 to 20,000 CPU hours per run.

Our simulations start with a spherical and homogeneous molecular cloud with parameters adapted from observations (radius 3.5 pc, H_2 gas density of 10^4 cm^{-3} gas temperature of 50 K). Figure 1 shows a cut through the midplane of the gas density distribution of our standard model (80 km/s cloud velocity, 3 pc offset from the x-axis) after an evolution time of 0.25 Myrs. The cloud has approached the black hole from the right hand side along the x-axis. After passing the black hole, material with opposite angular momentum collides, the angular momentum is efficiently redistributed and a rotating spiral structure builds up, whereas a part of the gas is transported through the inner boundary. The amount of accreted gas is mainly determined by the cloud's impact parameter relative to its radius. By that time, the inner gas disc has already relaxed to an approximate equilibrium state. The instability of such pressure stabilized discs can be assessed with the help of the so-called Toomre parameter. Figure 2 shows those regions of the disc, which are unstable to gravitational forces according to this Toomre criterion and will form stars. Defining the stellar disc in this way enables us to directly compare our results to the observed properties of the stellar system. Thereby we find that the initial cloud velocity determines the resulting stellar mass of the disc (the higher the velocity, the lower the disc mass). Our standard model shows a disc eccentricity of order 0.4, an outer radius of 1.35 pc and a disc mass of 0.7-6.95 $\times 10^4$ solar masses, in good agreement with the observations mentioned above. For details of this project we refer to [6].

Radiation Pressure Driven Dust Tori

One of the main questions of current AGN torus research is the stabilizing mechanism of their vertical scale height against gravity, as needed for the distinction between the two types of active galaxies. Several mechanisms have been proposed recently. We concentrate on one of the most promising solutions, namely the effects of infrared radiation pressure [7,8]. We are currently approaching the problem from two directions: (i) A parallel N-body like code is developed, which additionally includes radiation pressure forces and collisions between dusty gas clouds. This will enable us for the first time to assess the dynamical state of clumpy torus models, which are vitally discussed in literature nowadays, mostly treated with radiative transfer codes and show excellent agreement with observed properties of tori in nearby active galaxies. (ii) Tori are actually thought to be made up of a multi-phase medium including gas and dust at various temperatures. To be able to treat this mixture correctly, we are currently implementing a frequency resolved radiative transfer module

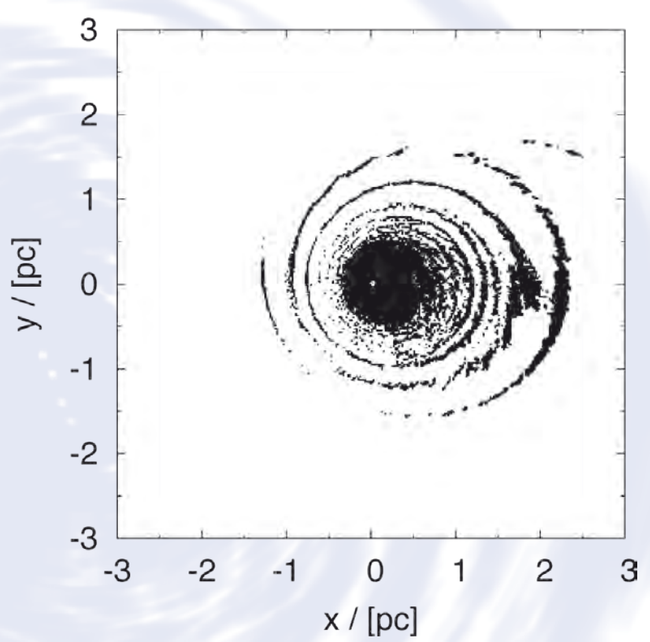

Figure 2: Gravitational unstable parts of the inner gas disc after 250,000 years, which will later form stars

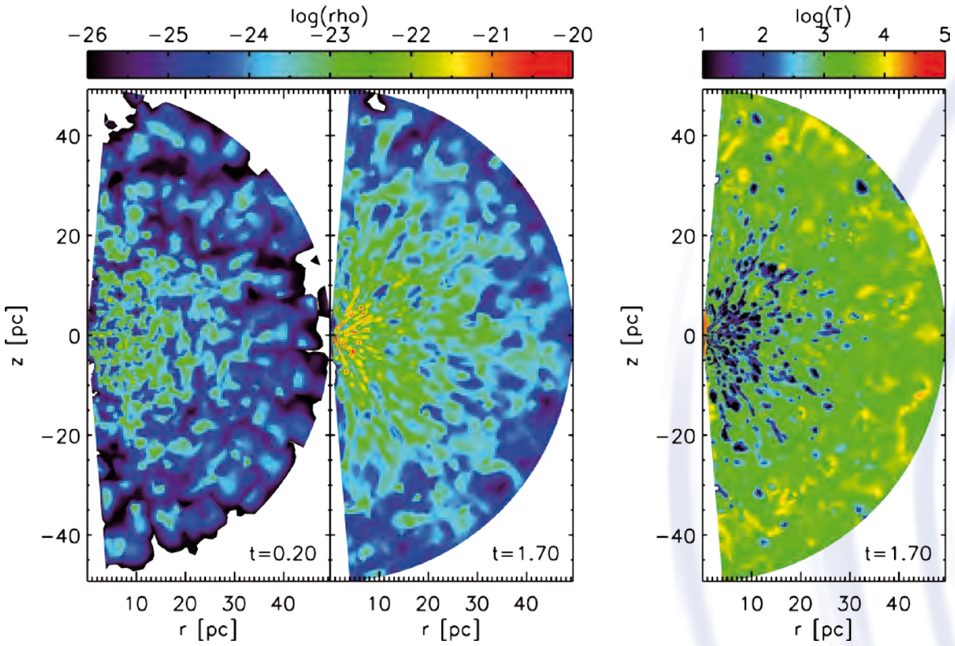

Figure 3: Gas density distribution (first two panels) and temperature distribution within a meridional slice through our 3D torus simulation of the Seyfert Galaxy NGC 1068.

into the Magneto-Hydrodynamics code NIRVANA. Radiation heats dust grains via absorption and accelerates them with the help of radiation pressure. We propagate forward and backward six radiation fields (one for each direction) and relate the thermodynamic quantities by equations of state for the gas and dust separately. As the light propagation timestep is often much smaller than the gas cooling and hydrodynamics timesteps, the code can subcycle the radiative transfer within the same hydrodynamics timestep. Test runs at LRZ showed that the code scales well down to a domain size of 16 cells and conserves the total energy (radiation, kinetic and thermal) very well. Using 16 processors at LRZ for two days, we are able to compute a periodic box of 60×64^2 cells for almost 15,000 timesteps. The code will then be applied to global gas and dust tori, irradiated by a central accretion disk.

Feeding Supermassive Black Holes with Nuclear Star Clusters

Recently, high resolution observations with the help of the near-infrared adaptive optics integral field spectrograph SINFONI at the VLT proved the existence of massive and young nuclear star clusters in the centres of nearby Seyfert galaxies. With the help of three-dimensional high resolution hydrodynamical simulations with the PLUTO code, we follow the evolution of the cluster of the Seyfert 2 galaxy NGC 1068. The gas ejection of its stars provide both, material for the obscuration within the so-called unified scheme of AGN and a reservoir to fuel the central, active region and it additionally drives turbulence in the interstellar medium on tens of parsec scales. This leads to a vertically wide distributed clumpy or filamentary inflow of gas on tens of parsec scale (see Figure 3 a,b,c), whereas a turbulent and very dense disc builds up on the parsec scale (see Figure 4). PLUTO [9] is a multiphysics, multialgorithm modular code which was especially designed for the treatment of supersonic, compressible astrophysical flows in the presence of strong discontinuities in a conservative form, relying on a modern Godunov-type High-Resolution Shock-Capturing (HRSC) scheme. PLUTO's parallel performance on the SGI Altix 4700 supercomputer at the LRZ shows almost ideal scaling.

In order to be able to treat viscosity effects and star formation within the disc correctly, we complement our 3D PLUTO simulations with an effective disc model, where we calculate the evolution of the surface density distribution. At the current age of NGC 1068's nuclear starburst of 250 Myr, our simulations yield disc sizes of the order of 0.8 to 0.9 pc, gas masses of 10^6 solar masses and mass transfer rates of 0.025 solar masses per year through the inner rim of the disc. This is in quite good agreement with the disc size as inferred from interferometric observations in the mid-infrared and compares well to the extent and mass of a rotating disc structure as inferred from water maser observations. Several other observational constraints are in good agreement as well. On basis of these comparisons, we conclude that the proposed scenario seems to be a reasonable model and shows that nuclear star formation activity and subsequent AGN activity are intimately related. For details we refer to [10].

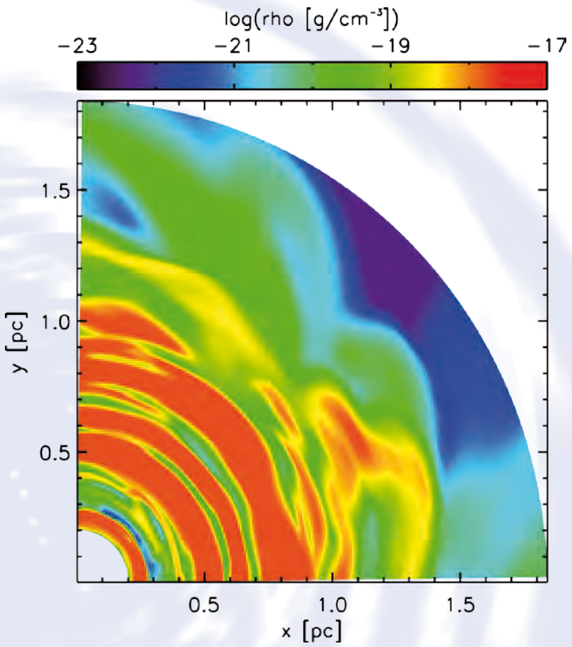

Figure 4: Central gas density distribution in an equatorial slice through our 3D torus simulation.

References

[1] **Bartko, H., et al.**
"Evidence for Warped Disks of Young Stars in the Galactic Center", ApJ 697, pp. 1741-1763, 2009

[2] **Nayakshin, S., Cuadra, J.**
"A Self-gravitating Accretion Disk in Sgr A* a Few Million Years Ago: Is Sgr A* a Failed Quasar?", A&A 437, pp. 437-445, 2005

[3] **Nayakshin, S., et al.**
"Weighing the Young Stellar Discs around Sgr A*", MNRAS 366, pp. 1410-1414, 2006

[4] **Wardle, M., Yusef-Zadeh, F.**
"On the Formation of Compact Stellar Disks around Sagittarius A*", ApJL, 683, pp. 37-40, 2008

[5] **Springel, V.**
"The Cosmological Simulation Code GADGET-2", MNRAS 364, pp. 1105-1134, 2005

[6] **Alig, Ch., et al.**
MNRAS, 2010, submitted (arXiv:0908.1100)

[7] **Pier, E. A., Krolik, J. H.**
"Radiation-pressure-supported Obscuring Tori around Active Galactic Nuclei", 399, pp. 23-26, 1992

[8] **Krolik, J. H.**
"AGN Obscuring Tori Supported by Infrared Radiation Pressure", 661, pp. 52-59, 2007

- Marc Schartmann
- Christian Alig
- Martin Krause
- Andreas Burkert

Max-Planck-Institute for Extraterrestrial Physics

Global Parameter Optimization of a Physics-based United Residue Force Field (UNRES) for Proteins

One of the major and still unsolved problems of computational biology is to understand how interatomic forces determine how proteins fold into three-dimensional structures. The practical aspect of the research on this problem is to design a reliable algorithm for the prediction of the three-dimensional structure of a protein from amino-acid sequence, which is of utmost importance because the experimental methods of the determination of protein structures cover only about 10 % of new protein sequences. Even more challenging is simulation of the dynamics of protein folding and unfolding and of the processes for which the dynamics of the proteins involved is critical (e.g., enzymatic reactions, signal transduction), because these processes occur within millisecond or second time scales, while the integration time step in MD simulations is of the order of femtoseconds [1]. In the case of protein structure prediction, methods that implement direct information from structural data bases (e.g., homology modeling and threading) are, to date, more successful compared to physics-based methods [2]; however only the latter will enable us to understand the folding and structure-formation process and protein dynamics.

The underlying principle of physics-based methods of protein-structure prediction and protein-folding simulations is the thermodynamic hypothesis formulated by Anfinsen [3] according to which the ensemble called the "native structure" of a protein constitutes the basin with the lowest free energy under given conditions. Thus, energy-based protein structure prediction is formulated in terms of a search for the basin with the lowest free energy; in a simpler approach the task was defined as searching the conformation with the lowest potential energy [4], and the prediction of the folding pathways can be formulated as a search for the family of minimum-action pathways leading to this basin from the unfolded (denaturated) state. In neither procedure do we want to make use of ancillary data from protein structural databases.

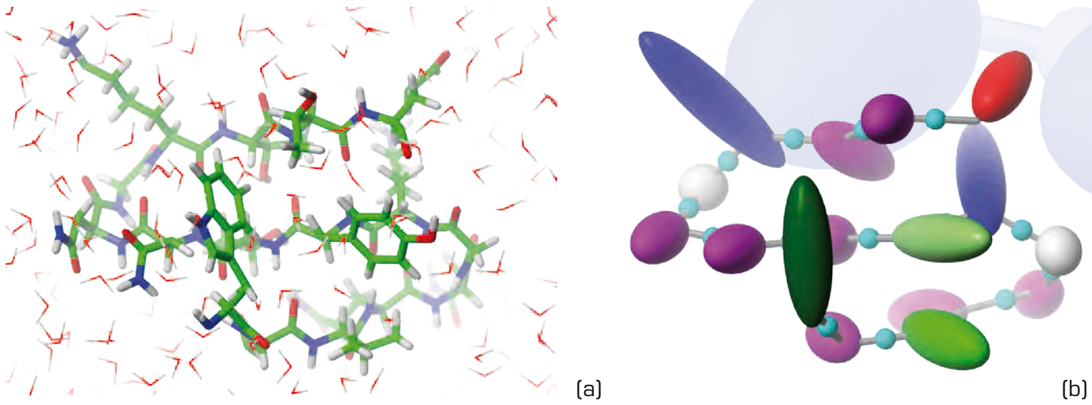

(a) (b)

Figure 1: Illustration of the correspondence between the all-atom polypeptide chain in water (a) and its UNRES representation (b). The side chains in part (b) are represented by ellipsoids of revolution and the peptide groups are represented by small spheres in the middle between consecutive α-carbon atoms. The solvent is implicit in the UNRES model.

The complexity of a system composed of a protein and the surrounding solvent generally prevents us from seeking the solution to the protein-structure and folding pathway prediction problem at atomistic detail [1]. Therefore, great attention has been paid in recent years to develop coarse-grained models of polypeptide chains and other biomolecular systems [5]. As part of this effort, we have been developing a physics-based united-residue force field, hereafter referred to as UNRES [6,7]. In the UNRES model [7] a polypeptide chain is represented by a sequence of α-carbon (C^α) atoms linked by virtual bonds with attached united side chains (SC) and united peptide groups (p). Each united peptide group is located in the middle between two consecutive α-carbons. Only these united peptide groups and the united side chains serve as interaction sites, the α-carbons serving only to define the chain geometry. The correspondence between the all-atom and UNRES representations is illustrated in Figure 1.

The UNRES force field has been derived as a Restricted Free Energy (RFE) function of an all-atom polypeptide chain plus the surrounding solvent, where the all-atom energy function is averaged over the degrees of freedom that are lost when passing from the all-atom to the simplified system (i.e., the degrees of freedom of the solvent, the dihedral angles χ for rotation about the bonds in the side chains, and the torsional angles λ, for rotation of the peptide groups about the $C^\alpha \cdots C^\alpha$ virtual bonds) [7]. The RFE is further decomposed into factors coming from interactions within and between a given number of united interaction sites. The energy of the virtual-bond chain is expressed by eq. (1).

$$
\begin{aligned}
U &= \sum_{i<j} U_{SC_iSC_j} + w_{SCp}\sum_{i\neq j} U_{SC_ip_j} + w_{pp}^{VDW}\sum_{i<j-1} U_{p_ip_j}^{VDW} + w_{el}f_2(T)\sum_{i<j-1} U_{p_ip_j}^{el} \\
&+ w_{tor}f_2(T)\sum_i U_{tor}(\gamma_i) + w_{tord}f_3(T)\sum_i U_{tord}(\gamma_i,\gamma_{i+1}) \\
&+ w_b\sum U_b(\theta_i) + w_{rot}\sum U_{rot}(\alpha_{SC_i},\beta_{SC_i}) + w_{corr}^{(3)}f_3(T)U_{corr}^{(3)} \\
&+ w_{corr}^{(4)}f_4(T)U_{corr}^{(4)} + w_{corr}^{(5)}f_5(T)U_{corr}^{(5)} + w_{corr}^{(6)}f_6(T)U_{corr}^{(6)} \\
&+ w_{turn}^{(3)}f_3(T)U_{turn}^{(3)} + w_{turn}^{(4)}f_4(T)U_{turn}^{(4)} + w_{turn}^{(6)}f_6(T)U_{turn}^{(6)} + w_{bond}\sum_{i=1}^{nbond} U_{bond}(d_i)
\end{aligned}
\qquad (1)
$$

The term $U_{SC_iSC_j}$ represents the mean free energy of the hydrophobic (hydrophilic) interactions between the side chains, which implicitly contains the contributions from the interactions of the side chain with the solvent. The term $U_{SC_ip_j}$ denotes the excluded-volume potential of the side-chain – peptide-group interactions. The peptide-group interaction potential is split into two parts: the Lennard-Jones interaction energy between peptide-group centers ($U_{p_ip_j}^{VDW}$) and the average electrostatic energy between peptide-group dipoles ($U_{p_ip_j}^{el}$); the second of these terms accounts for the tendency to form backbone hydrogen

bonds between peptide groups p_i and p_j. The terms $U_{corr}^{(m)}$ represent *correlation* or *multibody* contributions from the coupling between backbone-local and backbone-electrostatic interactions, and the terms $U_{turn}^{(m)}$ are correlation contributions involving m consecutive peptide groups; they are, therefore, termed turn contributions. The multibody terms are indispensable for reproduction of regular α-helical and β-sheet structures. The terms $U_{bond}(d_i)$, where d_i is the length of ith virtual bond and $nbond$ is the number of virtual bonds, are simple harmonic potentials of virtual-bond distortion; it has been introduced recently for molecular-dynamics implementation. The factors $f_n(T)$ account for the dependence of the respective free-energy terms on temperature [6].

The w's are the weights of the energy terms, and they can be determined (together with the parameters within each cumulant term) only by optimization of the potential-energy function, which is the subject of our present work. To search the conformational space with UNRES and to study the dynamics and thermodynamics of protein folding, we implemented Langevin dynamics in UNRES [7]. To sample the conformational space more efficiently than by canonical MD, we extended [8] the UNRES/MD approach with the replica exchange and multiplexing replica-exchange molecular dynamics method (MREMD). In the REMD method, M canonical MD simulations are carried out simultaneously, each one at a different temperature. After every $m < M$ steps, an exchange of temperatures between neighbouring trajectories $(j = i + 1)$ is attempted, the decision about the exchange being made based on the Metropolis criterion, which is expressed by eq 2.

$$\Delta = (\beta_j - \beta_i)\left[U(\mathbf{X}_j) - U(\mathbf{X}_i)\right] \qquad (2)$$

where $\beta_i = 1/RT_i$, T_i being the absolute temperature corresponding to the ith trajectory, and $\mathbf{X}_i$ denotes the variables of the UNRES conformation of the ith trajectory at the attempted exchange point. If $\Delta \leq 0$, T_i and T_j are exchanged, otherwise the exchange is performed with probability $\exp(-\Delta)$. The multiplexing variant of the REMD method (MREMD) differs from the REMD method in that several trajectories are run at a given temperature. Each set of trajectories run at a different temperature constitutes a *layer*. Exchanges are attempted not only within a single layer but also between layers. In this project we enhanced the MREMD by implementing multidimensional extension, in which pairs of replicas run at different temperatures and/or different sets of energy-term weights are exchanged in the multidimensional extension, based on the Metropolis criterion using the following Δ:

$$\Delta = \beta_j\left[U_k(\mathbf{X}_{j,k}) - U_k(\mathbf{X}_{i,l})\right] - \beta_i\left[U_l\left(\mathbf{X}_{j,k}\right) - U_l\left(\mathbf{X}_{i,l}\right)\right] \qquad (3)$$

where U_k and U_l denote different parameters of the potential energy, and $\mathbf{X}_{i,l}$ denotes the variables of the UNRES conformation of the replica simulated at temperature T_i using potential energy parameters U_l other symbols being the same as in eq. 2. This modification was necessary to carry out runs with multiple sets of energy parameters in global optimization of the parameters of the UNRES force field. In order to obtain a force field with folding properties, we developed a novel method of force-field optimization [6,7] which makes use of the hierarchical structure of the protein energy landscape. However, we realized that global and not local optimization (as performed so far [6]) of parameter space is needed for this purpose. Our recent preliminary results [9] of random scanning the parameter space and selecting the best parameter set by means of MREMD simulations of two peptides: 1L2Y and 1LE1, which have a well-defined α-helical and β-hairpin structure, respectively, produced a force field with only about 5 - 6 Å resolution, but very good transferability. Therefore, in the present project, we aimed at extending this approach to include a genetic algorithm, larger proteins, and experimental data of protein folding which are being carried out in our laboratory.

The goal of our current project is to optimize the UNRES force field by global search of the parameter space to reproduce structural and thermodynamic properties of proteins as functions of temperature. The force field will be applicable to MD simulations of large proteins at millisecond time scale in multiple-trajectory mode and to generating canonical ensembles of proteins. These features will enable us to study protein folding in both single-molecule and ensemble (kinetics) mode, chaperone interaction with proteins and such biologically important processes and systems as signal transduction, enzyme action, molecular motors, and amyloidogenesis, as well as to pursue physics-based prediction of protein structures, and the prediction of the effect of mutations on thermodynamic and kinetic stability.

The procedure of global optimization of the UNRES energy-function parameters developed within our project consists of the following 5 steps:

1. Run the multidimensional MREMD simulations on the training proteins for all sets of energy parameters in the initial population and evaluate the target function which makes use of the hierarchical structure of the protein energy landscape and selected experimental observables.

2. Select a subset of energy parameters to generate a new trial population by applying standard genetic operators: mutations and crossovers on vectors representing energy parameters.

3. Run the multidimensional MREMD simulations on training proteins after prescreening of trial population of energy parameters on two peptides with a minimal α-helical and a minimal β-hairpin fold.

4. Update the current population of energy parameters with best sets from trial population. If a new set P_{new} has a lower score function than a previous set P_i and the Cartesian distance between these sets is lower than a pre-set cut-off value $D_\circ$, P_{new} replaces P_i, if the distance is greater it is added and the "worst" previous set removed. $D_\circ$ is decreased as optimization progresses.

5. Iterate points 2. – 4.

The target function mentioned in points 1. and 4. is essentially the sum of the fourth powers of the deviations of the points of the calculated and experimental free-energy differences between the folded, unfolded, and partially folded forms of the training protein(s). The reader is referred to our earlier work for details [6]. To test the genetic algorithm we used the 1LE1 protein, which

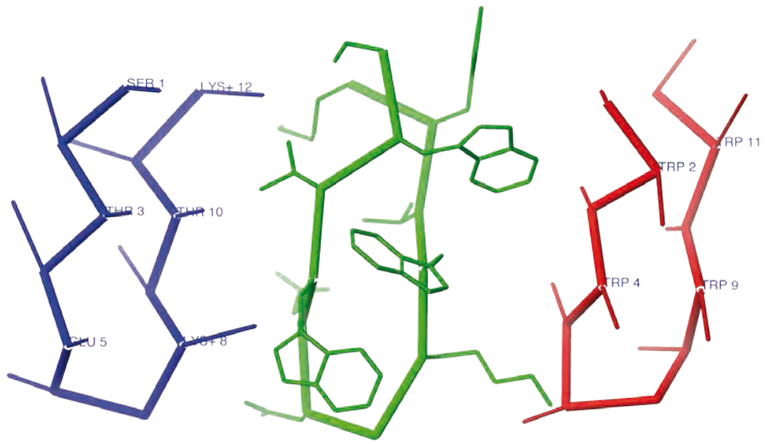

Figure 2: The most probable (left, blue; rmsd=2.6 Å) and the lowest-rmsd (right, red; rmsd=1.6 Å) conformation of 1LE1 corresponding to the best initial set of energy-term weights and the experimental (center, green) conformation of this peptide. In the UNRES structures the side chains are represented by virtual $C^{\alpha} \cdots$ SC bonds, while they are represented at atomistic detail in the experimental structure.

forms a single 12-residue β-hairpin [10]. The initial population consisted of 100 randomly-generated sets of energy-term weights. The most probable and the lowest-C^{α}-rmsd conformation, together with the experimental 1LE1 structure are shown in Figure 2. It can be seen that the most probable conformation has the β-hairpin shape; with the β-turn at Asn^6-Gly^7; however, the strands are flipped with respect to their native positions. The evolution of the fitness-score function with cycle number is shown in Figure 3. It can be seen that the score dropped by more than 2 orders of magnitude. The most probable and the lowest-C^{α}-rmsd conformations corresponding to the best set of energy-term weights are shown in Figure 4. It can be seen that not only the C^{α}-rmsd dropped but also the strands now have correct orientation with respect to the direction of the hairpin. Using the developed methodology, we are now optimizing the UNRES force field with bigger training proteins, which have more complicated folds. While the reported test required only modest computational resources (20 trajectories and 100 energy-term-weight sets, which makes a total of 2,000 cores), the actual optimization typically implies much larger jobs. Each energy energy and force calculation in a MREMD run is fine-grained to 8-16 cores and the number of MREMD trajectories is 128, multiplied by 100 energy-term-weight sets run simultaneously. Consequently, the optimization job requires 20,000+ cores. Consequently, these calculations can be accomplished only with the massively parallel resources provided by the Jülich Supercomputing Centre (Blue Gene/P). The optimized force field will be tested in 9th CommunityWide Experiment on the Critical Assessment of Techniques for Protein Structure Prediction (CASP9) which will take place between March and August 2010.

Figure 3: Evolution of fitness score with the number of the cycle for the optimization of the population of 100 sets of weights with the 1LE1 peptide.

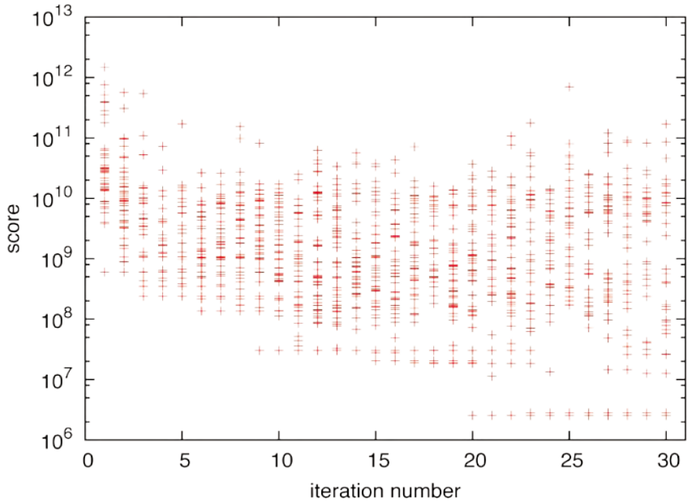

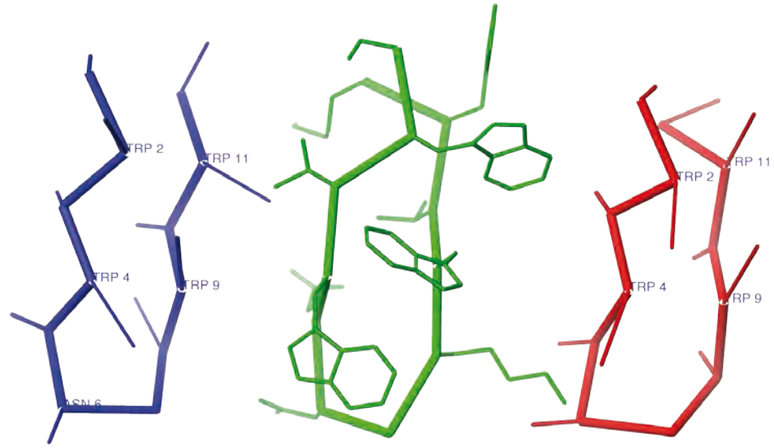

Figure 4: The most probable (left, blue; rmsd=2.0 Å) and the lowest-rmsd (right, red; rmsd=1.2 Å) conformation of 1LE1 corresponding to the best final (after 30 cycles) set of energy-term weights and the experimental (center, green) conformation of this peptide. In the UNRES structures the side chains are represented by virtual $C^{\alpha} \cdots$ SC bonds, while they are represented at atomistic detail in the experimental structure.

References

[1] **Scheraga, H. A., Khalili, M., Liwo, A.** "Protein-folding Dynamics: Overview of Molecular Simulation Techniques", Annual Review of Physical Chemistry, 58, pp. 57-83, 2007

[2] **Rohl, C. A., Strauss, C. E. M., Misura, K. M. S., Baker, D.** "Protein Structure Prediction Using ROSETTA", Methods in Enzymology, 383, p. 66, 2004

[3] **Anfinsen, C. B.** "Principles that Govern the Folding of Protein Chains", Science, 181, pp. 223-230, 1973

[4] **Scheraga, H. A., Liwo, A., Ołdziej, S., Czaplewski, C., Pillardy, J., Ripoll, D. R., Vila, J. A., Kazmierkiewicz, R., Saunders, J. A., Arnautova, Y. A., Jagielska, A., Chinchio, M., Nanias, M.** "The Protein Folding Problem: Global Optimization of Force Fields", Frontiers in Bioscience, 9, pp. 3296-3323, 2004

[5] **Kolinski, A., Skolnick, J.** "Reduced Models of Proteins and their Applications", Polymer, 45, pp. 511-524, 2004

[6] **Liwo, A., Khalili, M., Czaplewski, C., Kalinowski, S., Ołdziej, S., Wachucik, K., Scheraga, H.A.** "Modification and Optimization of the United-residue (UNRES) Potential Energy Function for Canonical Simulations. I. Temperature Dependence of the Effective Energy Function and Tests of the Optimization Method with Single Training Proteins", The Journal of Physical Chemistry B, 111, pp. 260-285, 2007

[7] **Liwo, A., Czaplewski, C., Scheraga, H. A. Rojas, A. V., Murarka, R. K., Ołdziej, S., Makowski, M., Kazmierkiewicz, R.,** "Simulation of Protein Structure and Dynamics with the Coarse-grained UNRES Force Field", In Coarse-Graining of Condensed Phase and Biomolecular Systems, 1st edition; Voth, G., Ed, CRC Press, Taylor & Francis; Boca Raton, FL, pp. 107–122, 2008

[8] **Czaplewski, C., Kalinowski, S., Liwo, A., Scheraga, H. A.** "Application of Multiplexing Replica Exchange Molecular Dynamics Method to the UNRES Force Field: Tests with α and α + β Proteins", Journal of Chemical Theory and Computation, 5, pp. 627–640, 2009

[9] **He, Y., Xiao, Y., Liwo, A., Scheraga, H. A.** "Exploring the Parameter Space of the Coarse-grained UNRES Force Field by Random Search: Selecting a Transferable Medium-resolution Force Field", Journal of Computational Chemistry, 30, pp. 2127–2135, 2009

[10] **Cochran, A. G., Skelton, N. J., Starovasnik, M. A.** "Tryptophan Zippers: Stable, Monomeric β-hairpins", Proceedings of the National Academy of Sciences U.S.A., 98, pp. 5578–5583, 2001

- Dawid Jagiela[1]
- Cezary Czaplewski[1,2]
- Adam Liwo[1,2]
- Stanislaw Ołdziej[2,3]
- Harold A. Scheraga[2]

[1] Faculty of Chemistry, University of Gdansk

[2] Baker Laboratory of Chemistry and Chemical Biology, Cornell University

[3] Intercollegiate Faculty of Biotechnology, University of Gdansk, Medical University of Gdansk

Recent Advances in the UNICORE 6 Middleware

The UNICORE Grid system provides a seamless, secure and intuitive access to distributed computational and data resources such as supercomputers, clusters, and large server farms. UNICORE serves as a solid basis in many European and international research projects that use existing UNICORE components to implement advanced features, higher-level services, and support for scientific and business applications from a growing range of domains. Since its initial release in August 2006, the current version UNICORE 6 has been significantly enhanced with new components, features and standards, offering a wide range of functionality to its users from science and industry. After giving a brief overview of UNICORE 6, this article introduces some of these new features and components.

Overview of UNICORE 6

UNICORE 6 is built on industry-standard technologies, using open, XML-based solutions for secure inter-component communication. UNICORE 6 is designed according to the principles of Service Oriented Architectures (SOAs). It uses SOAP-based stateful Web services, as standardized by the standards consortia W3C, OASIS and the Open Grid Forum (OGF). The security architecture is based on X.509 certificates, the Security Assertion Markup Language (SAML) and the eXtensible Access Control Markup Language (XACML). UNICORE 6 is implemented in the platform-neutral languages Java and Perl, and UNICORE 6 services successfully run on Linux/Unix, Windows and MacOS, requiring only a Java 5 (or later) runtime environment.

UNICORE 6 Architecture

A bird's eye view of the UNICORE 6 system is given in Figure 1. The system can be decomposed into three layers: the client, services and system layers. On the client layer, a variety of client tools are available to the users, ranging from the graphical UNICORE Rich client (URC), the UNICORE command-line client (UCC) to an easy-to-use high-level programming API (HiLA). Custom clients such as portal solutions can be easily implemented.

The middle layer comprises the UNICORE 6 services. Figure 1 shows three sets of services, the left and right one containing services at a single site while the middle shows the central services, such as the central registry, the workflow services and the information service, which serve all sites and users in a UNICORE Grid. The Gateway component acts as the entry point to a UNICORE site and performs the authentication of all incoming requests. It acts as a HTTP and Web services router, forwarding client requests to the target service, realizing a secure firewall tunnel. Therefore, in a typical UNICORE installation, only a single open firewall port is necessary.

The central UNICORE server component is based on UNICORE 6's Web services hosting environment. The XNJS component is the job management and execution engine of UNICORE 6. It performs the job incarnation, namely the mapping of the abstract job description to the concrete job description for a specific compute resource according to the rules stored in the so-called IDB (Incarnation Database). The functionality of the XNJS is accessible via two sets of Web service interfaces. The first set of interfaces is called UAS (UNICORE Atomic Services) and offers the full functionality to higher level services, clients and users. In addition to the UAS, a second set of interfaces based on open standards defined by the Open Grid Forum is available (depicted as "OGSA-*"). Some of the OGSA-* services are still under development or await official ratification, and the UNICORE community actively contributes to these discussions and provides reference implementations for emerging standards. The service layer includes a flexible and extensible security infrastructure, that allows interfacing to and integrating a variety of security components. Instead of the default UNICORE user database, Virtual Organization (VO) management systems can be accessed for getting user attributes such as role, project membership or local login. To enable interoperability scenarios, proxy certificates that are common in other Grid software are optionally supported.

On the lowest layer the TSI (Target System Interface) component is the interface between UNICORE and the resource management system and operating system of a Grid resource. In the TSI component the abstracted commands from the Grid layer are translated to system-specific commands, e.g. in the case of job submission, the specific commands like *llsubmit* or *qsub* of the resource management system are called. The TSI component is perform-

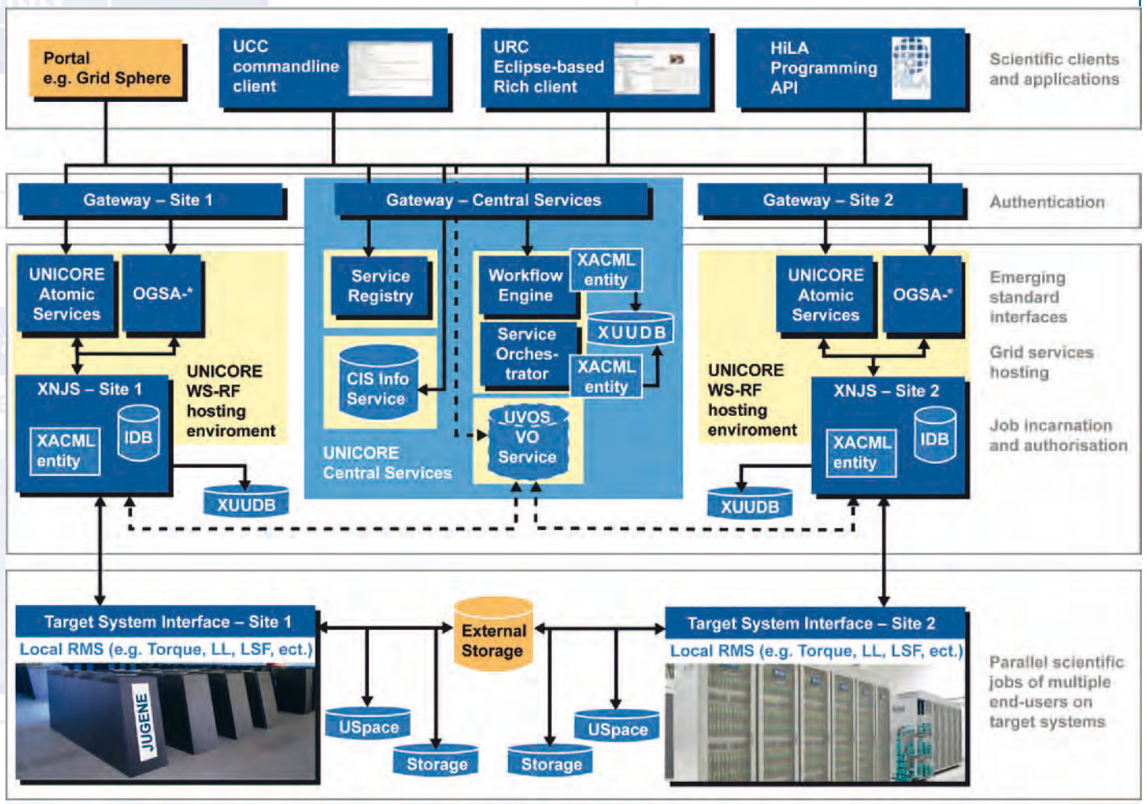

Figure 1: Overview of the UNICORE 6 system

ing tasks under the users UID and is the only UNICORE component that needs to be executed with root privileges. The TSI is available for a variety of commonly used resource management systems such as Torque, LoadLeveler, LSF, SLURM, and Sun GridEngine. In addition to job submission and job management services, UNICORE 6 provides unified access to storages via its Storage-Management and FileTransfer service interfaces. File systems are accessed directly via the TSI, while other storages can be integrated using custom adaptors. For example, an adaptor to the Apache Hadoop cluster storage system has been realized.

Recent Development Highlights

Workflow System

The UNICORE 6 workflow system offers powerful workflow features such as sequences, branches, loops and conditions, while integrating seamlessly into UNICORE Grids and into the UNICORE clients. The workflow system supports automated resource selection and brokering, and is scalable due to its two-layered architecture. Recently a for-each loop construct has been added that allows to process large file sets or perform parameter studies in a scalable and very user-friendly fashion. Figure 2 shows a screen shot of the URC, with the workflow editor shown on the right-hand side.

Improved Data and Storage Management

UNICORE's data and storage management capabilities have been integrated into the UNICORE Rich Client in an easy to use fashion. The URC allows common operations like editing, deleting or renaming, as well as drag and drop of files between the remote sites and the user's desktop. A new StorageFactory service allows to create and manage remote storage resources.

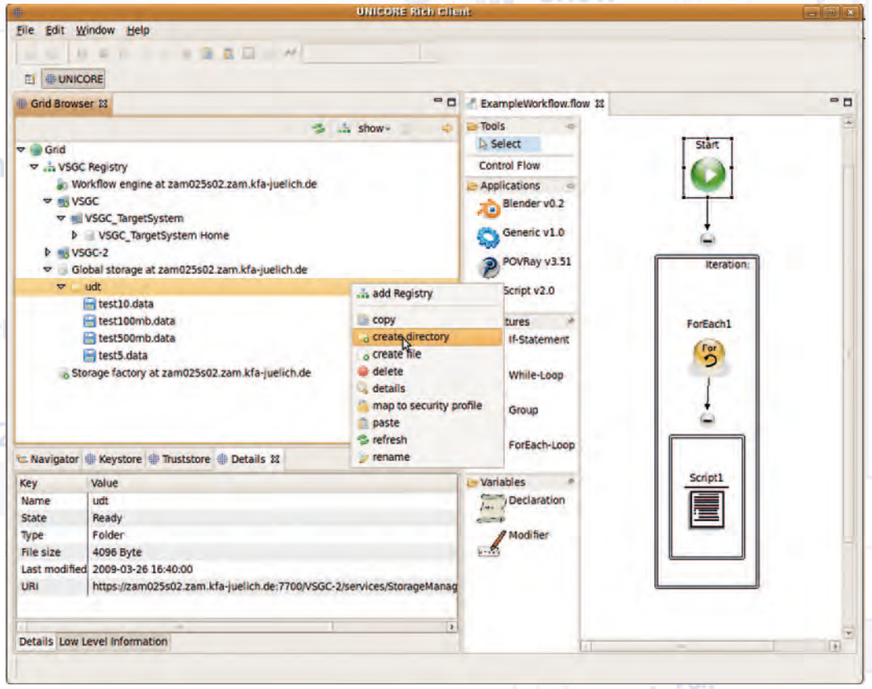

Figure 2: Screen shot of the UNICORE Rich Client

Execution Environments: Improved Support for Parallel Applications

Users running parallel applications are usually required to know about the invocation details of the parallel environment (such as OpenMP or MPI). A recently added UNICORE feature allows administrators to easily parametrize the available parallel environments in the UNICORE IDB, so that users can easily choose the required environment and set options and arguments, without being required to know the system specifics.

Remote Administration

It is often convenient to check site health, get up-to-date performance metrics and even deploy and un-deploy services remotely using the standard communication channels (i.e. Web services), without having to rely on other tools such as Nagios, or needing to actually log in to the UNICORE 6 server. For this reason, an Admin service has been developed that offers various remote administration and monitoring features. It can be used through a graphical URC plugin and through the commandline client. Using the UNICORE XACML based access control, access to this service is of course limited to privileged users.

Shibboleth Integration in the UNICORE Rich Client

Shibboleth is a standards-based identity management system aiming to provide single sign-on across organisational boundaries. It is considered a prime candidate for simplifying end user experience by hiding the specifics of the X.509 certificate based security. Therefore a plugin for the URC was developed that allows to access Shibboleth-enabled attribute authorities, which create short-lived certificates that enable access to the Grid resources.

Interactive Access to Grid Sites

Using the single sign-on feature of the UNICORE Rich Client, it is desirable to provide an interactive terminal access to the remote site. A URC plugin has been developed that can be used to open an interactive connection using a variety of means such as standard SSH using username and password, GSI – SSH and a custom X.509 based solution. The client-side plugin offers VT100 emulation and port forwarding.

Open Source Model

UNICORE 6 is available as open source under BSD license from SourceForge. At **http://www.unicore.eu** more information can be obtained and lightweight UNICORE 6 installation packages can be downloaded. A small UNICORE 6 testgrid is available for immediate testing. UNICORE 6 is continually evolved, with regular releases approximately every three months. The main contributors in the international UNICORE open-source community are: ICM Warsaw, CINECA, CEA, Technische Universität Dresden, and Forschungszentrum Jülich.

- Bernd Schuller
- Morris Riedel
- Achim Streit

Jülich Supercomputing Centre, Forschungszentrum Jülich

LIKWID Performance Tools

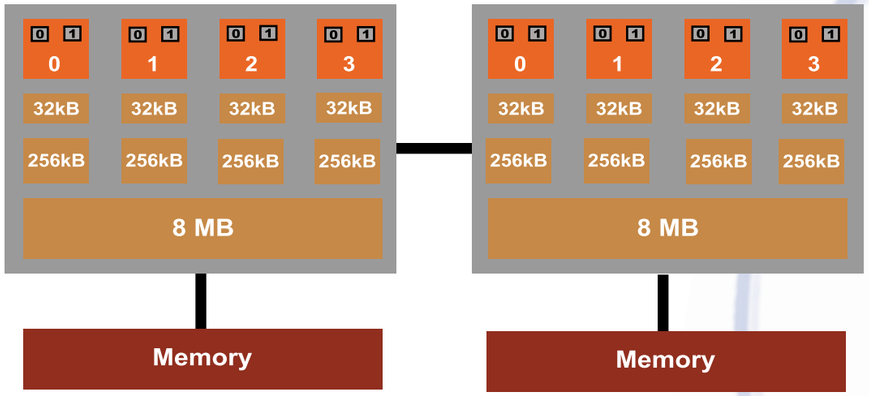

Figure 1: Topology of typical Intel Nehalem based Cluster Node

Today's multicore x86 processors bear multiple complexities when aiming for high performance. Conventional performance tuning tools like Intel VTune, OProfile, CodeAnalyst, OpenSpeedshop etc. require a lot of experience in order to get sensible results. For this reason they are usually unsuitable for the scientific user, who would often be satisfied with a rough overview of the performance properties of their application code. Moreover, advanced tools often require kernel patches and additional software components, which makes them unwieldy and bug-prone. Additional confusion arises with the complex multi-core, multi-cache, multi-socket structure of modern systems; affinity has become an important concept, but users are all too often at a loss about how to handle it.

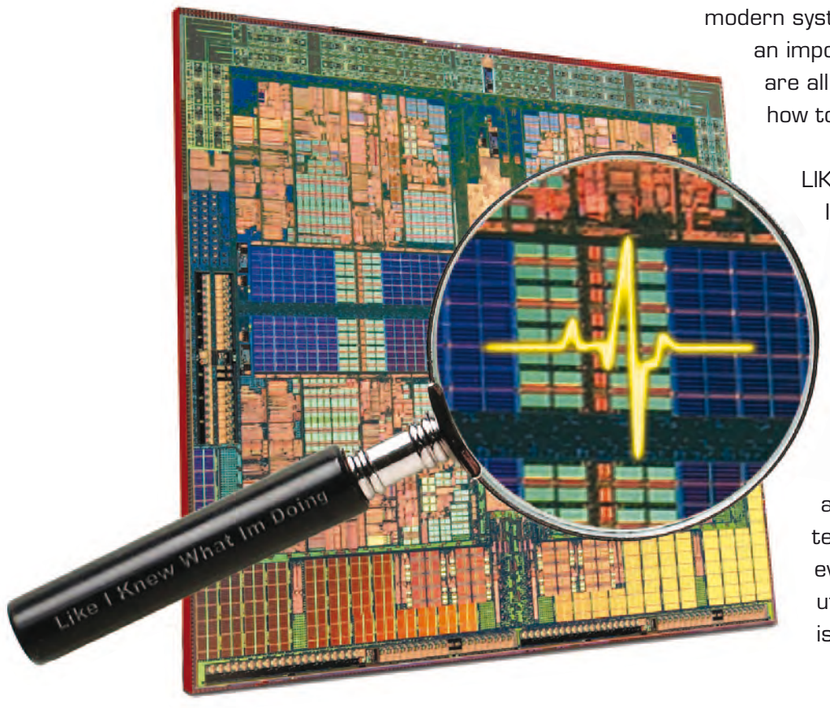

LIKWID ("Like I Knew What I'm Doing") is a set of easy to use command line tools to support optimization and is targeted towards performance-oriented programming in a Linux environment. It does not require any kernel patching and is suitable for Intel and AMD processor architectures. Multi-threaded and even hybrid shared/distributed-memory parallel code is also supported. Building it only requires a GNU C compiler and is a matter of seconds. In the following we describe all four tools in detail:

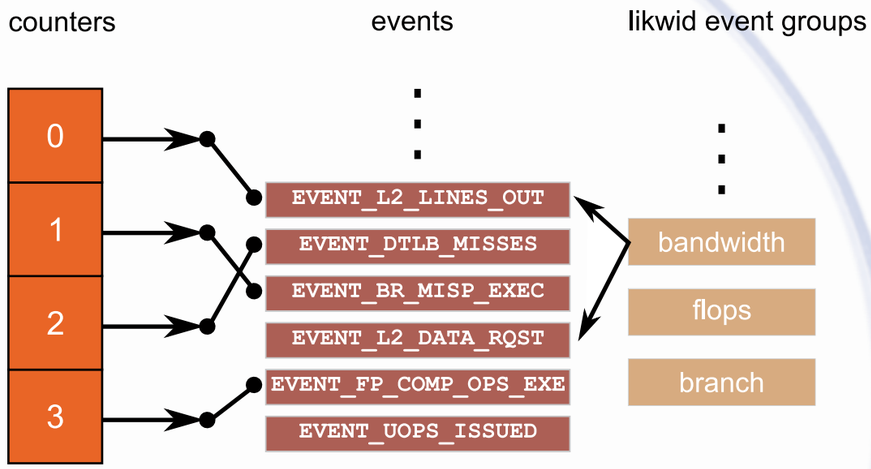

Figure 2: Relationship between counter register, hardware events and LIKWID event groups

likwid-perfCtr

Hardware performance counters are facilities to count hardware events during code execution on a processor. Because this mechanism is implemented directly in hardware there is no overhead involved. All modern processors provide hardware performance counters, but their primary purpose is to support computer architects during the implementation phase. Still they are also attractive for application programmers, because they allow an in-depth view on what happens on the processor while running applications. Tools for hardware performance counters have the reputation to be complex to install and even more complex to use. In almost all cases a special kernel module is necessary. Many tools are restricted to one processor vendor, others require code changes in the user's application. Hardware performance counters are controlled and accessed using processor-specific hardware registers (also called model specific registers (MSR)). likwid-perfCtr uses the Linux msr module to modify the MSRs from user space. It can be used as a wrapper or with code markers to probe only parts of the application. A major problem with hardware performance events is how to choose and interpret them correctly. The events differ between processors and it is difficult for a beginner to choose the right events to help in their optimization effort. likwid-perfCtr provides preconfigured groups with useful, ready to use event sets and derived metrics like bandwidth and event ratios. Still likwid-perfCtr is fully transparent, i.e. it is clear at any given time on which events the performance groups are based. This ease of use sets likwid-perfCtr apart from existing tools.

Supported architectures:

- Intel Pentium M (Banias, Dothan)
- Intel Core 2 (all variants)
- Intel Nehalem (Xeon, i7, i5)
- AMD K8 (all variants)
- AMD K10 (Barcelona, Shanghai, Istanbul)

likwid-topology

Multicore/Multisocket machines exhibit complex topologies, and this trend will continue with future architectures. Performance programming requires in-depth kowledge of cache and node topologies, i.e. which caches are shared between which cores and which cores reside on which sockets. The Linux kernel numbers the usable cores and makes this information accessible in /proc/cpuinfo. Still how this numbering maps on a processors topology depends on BIOS settings and may even differ for otherwise identical processors. The processor and cache topology can be queried with the **cpuid** machine instruction. likwid-pin is based directly on the data provided by **cpuid**. It extracts machine topology in an accessible way and can also report on cache characteristics. There is an option to get a quick overview of the node's cache and socket topology in ASCII art.

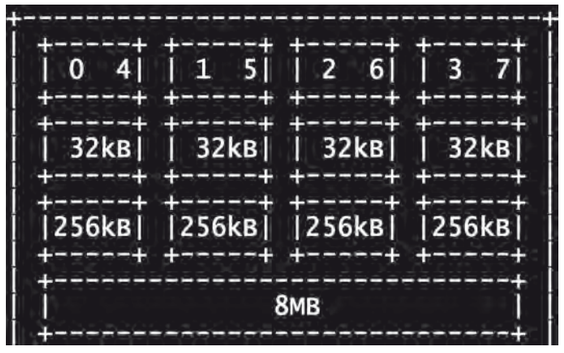

Figure 3: Intel Core i7

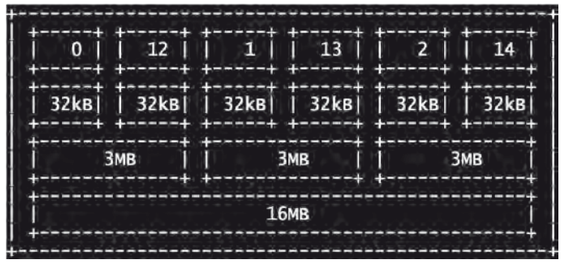

Figure 4: Intel Dunnington

Thread/Process pinning is vital for performance. If topology information is available, it is possible to pin threads according to the application's resource requirements like bandwidth, cache sizes etc. likwid-pin supports thread affinity for all threading models that are based on POSIX threads, which includes most OpenMP implementations. By preloading a library wrapper to the **pthread_create** API call, likwid-pin is able to bind each thread upon creation. It configures a dynamic library using environment variables and starts the real application with the library preloaded. No code changes are required, but the application must be dynamically linked.

```
$ export OMP_NUM_THREADS=4
$ ./likwid-pin -c 0-3 -t intel  ./a.out
```

Supported threading implementations:

- POSIX threads
- Intel OpenMP
- gcc OpenMP

Figure 5: Output of likwid-topology for Intel Xeon node with two sockets

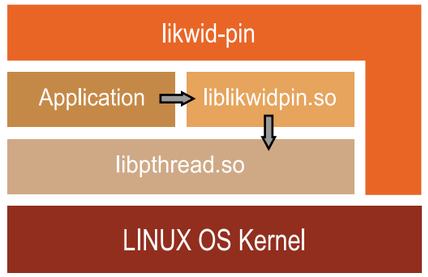

Figure 6: Pinning threads with likwid

Other threading implementations are supported via a "skip mask." This skip mask specifies which threads should not be pinned by the wrapper library. For instance, the Intel OpenMP implementation always runs **OMP_NUM_THREADS**+1 threads but uses the first created thread as management thread. The skip mask makes it possible to pin hybrid applications as well by skipping MPI shepherd threads. In addition, likwid-pin can be used for sequential programs as replacement for the (inferior) **taskset** tool.

Some compilers provide their own mechanisms for thread affinity. In order to avoid interference effects, those mechanisms should be disabled when using likwid-pin. In case of the Intel compiler, this can be achieved by setting the environment variable **KMP_AFFINITY**. The big advantage of likwid-pin is that the same tool can be used for all applications, compilers, MPI implementations and processor types.

likwid-features

An important hardware optimization on modern processors is to hide data access latencies by hardware prefetching. Intel processors not only have a prefetcher for main memory, several prefetchers are responsible for prefetching the data inside the cache hierarchy. Often it is beneficial to know the influence of the hardware prefetchers. In some situations turning of hardware prefetching might even increase performance. On the Intel Core 2 processor this can be achieved by setting bits in the IA32_MISC_ENABLE MSR register. Besides the ability to toggle the hardware prefetchers likwid-features reports on the state of switchable processor features as e.g. Intel speedstep. likwid-features only works for Intel Core 2 processors.

Future Plans

Some enhancements to the suite are imminent. Due to the limited number of performance counter registers on current CPUs, a code must often be run multiple times in order to get complete information. Event multiplexing, i.e. switching between sets of event sources on a regular basis, will enable the concurrent measurement of many performance metrics. This will introduce a statistical error to the event counts, which is however well controllable. The API for enabling and disabling performance counters in the current version of likwid-perfCtr is only rudimentary, and will be enhanced to become more flexible.

Further goals are the combination of LIKWID with the well-known IPM tool for MPI profiling to facilitate the collection of performance counter data in MPI programs, and the integration of a simple bandwidth map benchmark, which will allow a quick overview of the cache and memory bandwidth bottlenecks in a shared-memory node, including the ccNUMA behaviour.

How to get it

The LIKWID performance tools are open source. You can download them at: **http://code.google.com/p/likwid/**

- Jan Treibig
- Michael Meier
- Georg Hager
- Gerhard Wellein

Regional Computer Centre (RRZE), University of Erlangen-Nuremberg

Project ParaGauss

The interdisciplinary project ParaGauss in the framework of the Munich Center of Advanced Computing aims at developing and implementing new and improved strategies for calculating the properties of molecular systems at the level of quantum mechanics, i.e. from "first principles" and without empirical assumptions. The project relies on the density functional approach to calculate the electronic structure of atoms, molecules or solids and builds on the implementation in the parallel program ParaGauss [1]. To reach the goals of the project, alternatives to current parallel implementations have to be developed, to increase the number of efficiently usable processors and especially to improve the usability of modern multi-core architectures.

Compared to other areas of simulation, like fluid dynamics or continuum mechanics in engineering or the geosciences, quantum chemistry methods, such as those based on Density Functional Theory (DFT), lack a common data structure which would allow an efficient and homogeneous parallelization of the entire algorithm.

At the heart of the problem are the solutions $\varphi_i(\vec{r})$ of the Kohn-Sham equation $H\varphi_i(\vec{r}) = \varepsilon_i\varphi_i(\vec{r})$ where H is the Kohn-Sham operator which represents the interaction energies of one electron with all the others as well as with the atomic nuclei of the system. Summing over a sufficient number of one-electron orbitals $\varphi_i(\vec{r})$ from lower to higher energies ε_i yields the electronic density

$$\rho(\vec{r}) = \sum_i^{occ} |\varphi_i(\vec{r})|^2$$

The Kohn-Sham operator H itself depends on the electron density ρ so that the Kohn-Sham equation has to be solved iteratively until "self-consistency" has been reached. One starts from an initial guess for $\varphi_i(\vec{r})$, constructs the Kohn-Sham operator H and solves the Kohn-Sham equation. With the resulting density, one constructs a new operator, and so on, until the initial and resulting orbitals are sufficiently close to each other.

Various subtasks of mathematically different structure need to be tackled in this procedure. First, the Kohn-Sham equation is transformed to a generalized matrix eigenvalue problem by representing the orbitals as linear combinations of a finite set of localized, most often Gaussian-type function χ_k. The resulting matrix elements are then obtained as integrals, in fact scalar products $\langle \chi_k | V | \chi_l \rangle$, over the various terms V of the Kohn-Sham operator, representing the kinetic energy of the electron, the attractive potential due to the atomic nuclei, the repulsion of an electron by the classical Coulomb potential due to electron density ρ and the so-called exchange and correlation terms.

These latter terms comprise the quantum mechanical aspects of the many-body effects. The more demanding tasks in the Kohn-Sham procedure are the determination of the matrix elements according to analytical integrals and the evaluation of the electron-electron interaction during the self-consistency cycle from such integrals (Coulomb term) or by numerical integration (exchange and correlation terms). A smaller task, besides others, is the diagonalization of the generalized eigenvalue problem. To determine the electronic structure at a given molecular geometry, i.e. a given set of nuclei assumed to be fixed in space, one has to carry out these tasks repeatedly and consecutively during the iteration process. In a typical application, one determines the geometry of a molecular system (Figure 1), i.e. one has to minimize the energy by varying the molecular geometry, which requires first-order derivatives (forces) of the energy with regard to nuclear displacements, hence additional analytical and numerical integrations. Often one also has to determine second-order derivatives of the energy with respect to nuclear coordinates, e.g. when one needs an accurate description (frequencies) of the vibrations of the framework of nuclei.

In the program ParaGauss the procedure sketched above is parallelized in a coarse-grained fashion using message passing via MPI. Each individual subtask is divided in as few as possible consecutive blocks to minimize communication. Crucial for an efficient parallelization is that a large amount of "integrals" have to be handled when constructing the matrix elements of H. For larger molecules this data set reaches several GB. Also, the computational effort of the different tasks scatters widely, resulting in bottlenecks that limit the number

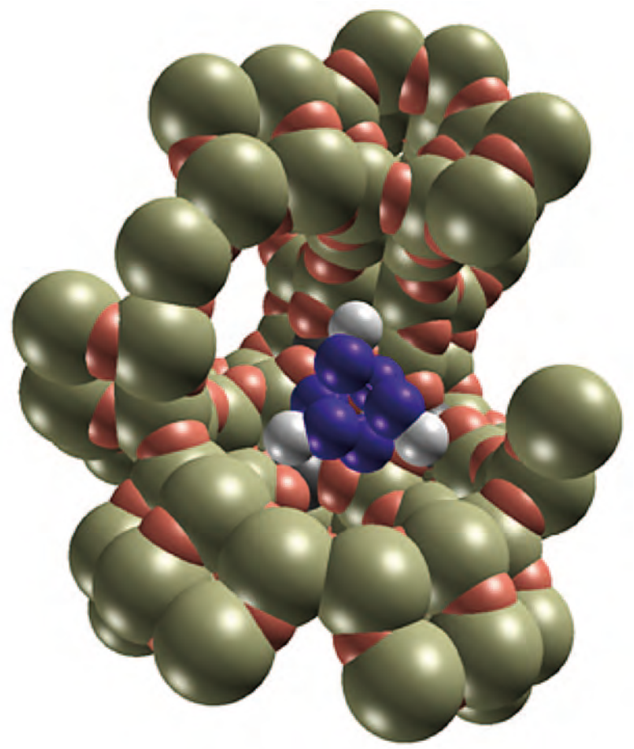

Figure 1: A model catalyst: transition metal hydride cluster Rh_6H_3 (rhodium purple, hydrogen white) adsorbed at the wall of a SiO_2 framework (zeolite of faujasite structure; silicon brown, oxygen grey)

of efficiently usable processors, especially for the smaller subtasks. To improve the usage of multi-core architectures as well as for a generally higher efficiency when engaging a larger number of processors, several strategies will be combined in a new implementation of ParaGauss. Multi-core architectures provide less memory per core and typically slower access than single-core hardware. This issue will become even more important as memory access is already a bottleneck at current computing platforms.

One way out is to limit the amount of data of parallel tasks at the expense of communication overhead. Currently we are exploring an alternative solution which uses message passing parallelization between multi-core shared memory nodes and, on each multi-core node, low-level parallelization of single MPI tasks by means of OpenMP. On a shared-memory multi-core system, we will compare the efficiency of an OpenMP parallelized variant of a module that evaluates analytical integrals with the existing version on the basis of message-passing parallelization. This module will serve as prototype for MPI parallelization between nodes.

As algorithmically simple tasks, e.g. numerical integration, are easily parallelized, computationally less demanding tasks, that are harder to parallelize, consume a growing part of the total run time with growing number of processors. Therefore it is necessary to optimize all tasks which take up more than ~1 % of real time in a sequential run. Otherwise these tasks will dampen the acceleration for a larger number of processors.

The solution of the generalized matrix eigenvalue problem is such a task. In general the matrix to be diagonalized is decomposed into diagonal blocks for symmetric molecular frameworks or calculations that account for spin polarization of open-shell systems. These densely filled submatrices are symmetric (or hermitian) with typical dimensions ranging from 100 to a few 1,000, reaching 10,000 for large problems. Currently these matrices are diagonalized by a serial algorithm, but submatrices are treated in parallel. As parallel eigensolvers are not very efficient for smaller matrices, we constructed a mixed procedure for diagonalizing a set of matrices. First one determines a hardware-dependent optimum number of cores for diagonalizing matrices of increasing size. Then one optimizes the order in which the matrices are diagonalized, employing varying numbers of processors to achieve a low overall execution time (Figure 2).

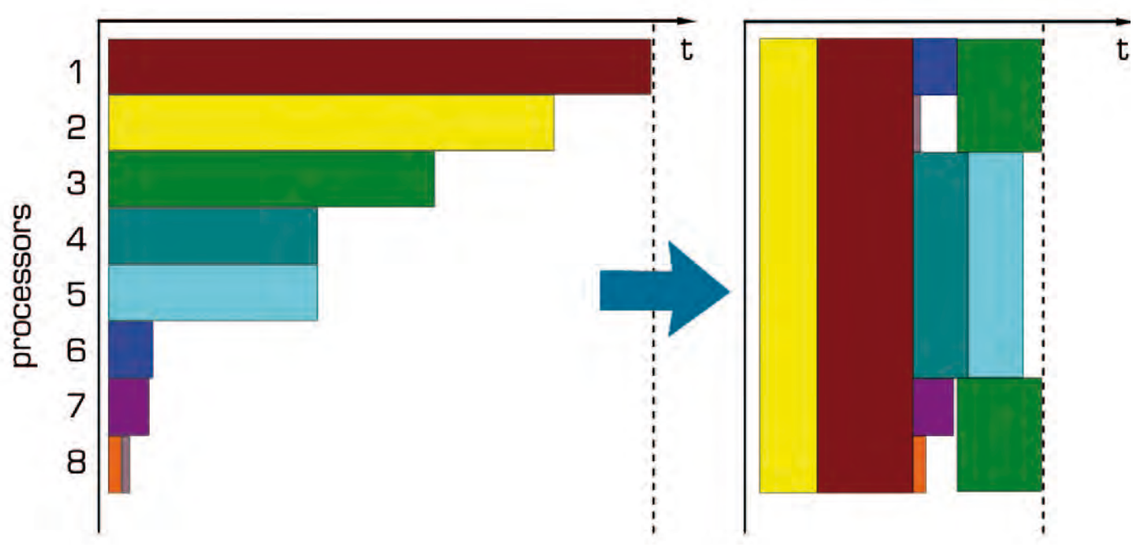

Figure 2: Scheme of parallel solution of several matrix eigenvalue problems.
Left: serial solution of individual problems in parallel, right: parallel solution of problems, partially parallelized.

In this spirit, other tasks also may be solved more efficiently by allowing variable numbers of cores. In this way, dynamically variable numbers of cores will reduce the overall runtimes while it also wastes hardware resources to a certain degree. In typical applications, one hardly ever carries out just a single huge simulation; rather, several if not many similar calculations are needed. This observation may help to resolve the conflict just described by tackling several such application tasks in parallel on a larger pool of cores. Typical problems benefitting from such a pooling strategy are the determination of energies of educts and products of a reaction, the determination of a reaction path that is represented by a finite set of intermediate structures, or the comparison of various isomers of a given compound. Combining this approach with an optimized dynamical usage of the numbers of cores for different subtasks of the algorithm will increase the overall efficiency of a simulation. Implementation of this combined strategy requires a strict modularization of the overall procedure into separate modules for each subtask. The number of processors used at a time will then be determined taking into account the optimum number of processors for specific tasks and the number of currently available cores.

Two-level parallelization, modularization, optimization of memory usage and pool-parallel execution of tasks taken together will yield a new software structure for quantum chemical calculations that is also ready for distributed computing. We are confident that the MAC project ParaGauss will deliver a simulation tool of enhanced usability and efficiency on shared-memory multi-core environments.

References

[1] **Belling, Th., Grauschopf, T., Krüger, S., Mayer, M., Nörtemann, F., Staufer, M., Zenger, C., Rösch, N.,**
"High Performance Scientific and Engineering Computing", in Bungartz, H.-J., Durst, F., Zenger, C. (Eds.), Proceedings of the First International FORTWIHR Conference 1998, Lecture Notes in Computational Science and Engineering, Springer, Heidelberg, Vol. 8, p. 439, 1999

• Notker Rösch
• Sven Krüger

Technische Universität München, Theoretische Chemie

Real World Application Acceleration with GPGPUs

Future microprocessor development efforts will continue to concentrate on adding cores rather than increasing single-thread performance. Similarly, the highly parallel Graphics Processing Unit (GPU) is rapidly gaining maturity as a powerful engine for computationally demanding applications [1]. GPUs, designed primarily for 3D graphics applications, are starting to become an attractive option to accelerate scientific computations. Modern GPUs do not only offer powerful graphics but also a highly parallel programmable processor featuring peak arithmetic performance and memory bandwidth to the on-board graphics card memory that substantially out-paces the CPUs counterpart. However, making efficient use of GPU clusters still presents a number of challenges in application development and system integration. GPUs are optimized for structured parallel execution, extensive ALU counts and memory bandwidth. GPUs are designed to provide high aggregate performance for tens of thousands of independent computations. This key design choice allows GPUs to spend the vast majority of chip die area (and thus transistors) on arithmetic units rather than on caches.

A particularly interesting application for GPU acceleration is the field of molecular dynamics simulation, which is dominated by N-body atomic force calculations. This article re-examines the performance of GPUs and compares the results to pure CPU performance. This comparison will be carried out using the NAMD molecular dynamics simulation package. NAMD (NAnoscale Molecular Dynamics) is a free-of-charge molecular dynamics simulation package written using the Charm++ parallel programming model, noted for its parallel efficiency and regularly used to simulate large systems (millions of atoms). It has been developed by the joint collaboration of the Theoretical and Computational Biophysics Group and the Parallel Programming Laboratory at the University of Illinois at Urbana-Champaign. At the LRZ, NAMD (version 2.7b1) has been compiled both with and without GPU support. CUDA acceleration mixes well with task-based parallelism, allowing NAMD to run on clusters with multiple GPUs per node.

GPUs at the LRZ

The LRZ operates two powerful visualization servers since early 2009. Both machines are SunFire-X4600-M2 servers with 32 CPU cores (8 quad-core Opterons with 2.3 GHz) and 256 GB RAM each. Each server has 4 Nvidia Quadro FX5800 graphics cards attached. The graphics cards feature 4GB on-board memory and 240 CUDA parallel processing cores (the Nvidia Quadro FX5800 graphics card is identical to the Tesla C1060 card). CUDA (Compute Unified Device Architecture) is a parallel computing architecture developed by NVIDIA [3]. CUDA has unlocked the computational power of GPUs and made GPUs much more accessible to computational scientists.

NAMD with CUDA

NAMD is a parallel molecular dynamics (MD) code for large biomolecular systems [2]. MD simulations compute

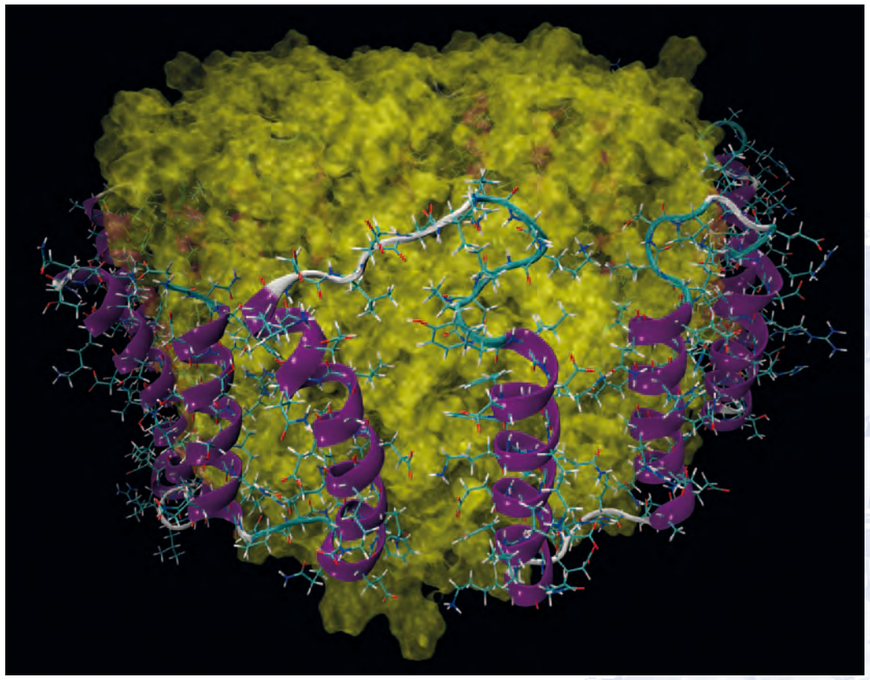

Figure 1: The Figure depicts a representation of the file apoa1 that was used in this article. This file consists of 92,224 atoms, including the water box surrounding the molecule (water is not shown in the Figure). The full name of the molecule is "ApoA-1 Milano". It is a naturally occurring mutated variant of the apolipoprotein A1 protein found in human HDL, the lipoprotein particle that carries cholesterol (depicted in yellow in the middle of the protein) from tissues to the liver and it is associated with protection against cardiovascular disease.

atomic trajectories by solving equations of motion numerically using empirical force fields [5,6], such as the CHARMM force field [7], that approximate the actual atomic force in molecular systems. NAMD is distributed free of charge to academics, pre-compiled for popular computing platforms or in source code form [4]. The velocity Verlet integration method [8] is used to advance the positions and velocities of the atoms in time. To further reduce the cost of the evaluation of long-range electrostatic forces, a multiple time step scheme is employed. The local interactions (bonded, van der Waals and electrostatic interactions within a specified distance) are calculated at each time step. The longer range interactions (electrostatic interactions beyond the specified cutoff distance) are computed less frequently. This reduces the cost of computing the electrostatic forces over several time steps. The application decomposes and schedules the computation among CPU computation and GPU kernel. The GPU computes the non-bonded interactions, bonded interactions such as angle, dihedral and torsion are still done on the CPU.

At every iteration step, NAMD has to calculate the short-range interaction forces between all pairs of atoms within a cutoff distance. By partitioning space into patches slightly larger than the cutoff distance, NAMD ensures that all interactions of one atom are with atoms in the same or neighboring cubes.

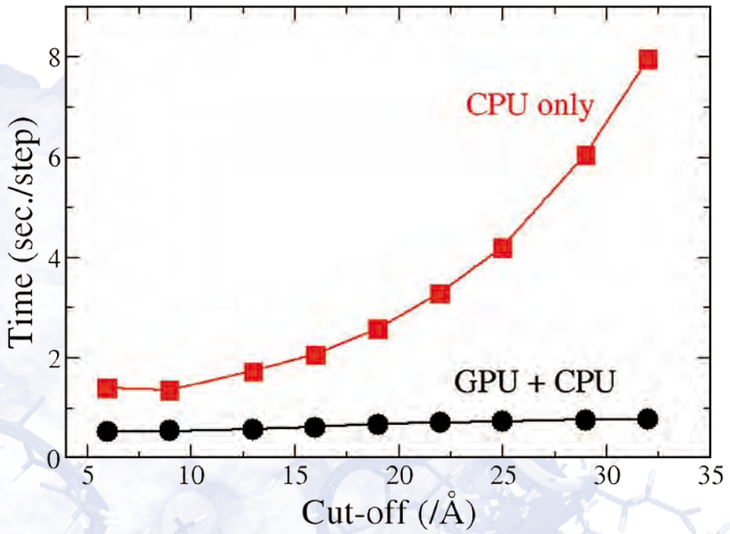

Figure 2: Run-time as a function of the cutoff value. Values for the CPU-only version of NAMD are represented by black circles, values for the CPU-GPU version of NAMD are depicted by red squares.

Results and Discussions

The performance results as a function of the cutoff is shown in Figure 2, where we compare the results of the CPU-only version of NAMD (using 4-Cores Opteron 2.3 GHz) with the CPU-GPU version of NAMD, using the Nvidia-FX58000 graphics cards. A larger value for cutoff means more computation is transferred to the GPU. To generate the results in Figure 2, one has to switch off the flags twoAwayX, twoAwayY and twoAwayZ. The configuration files of the apoa1 benchmark were used by varying the cutoff values from 6\Å to 32\Å. To increase parallelism on GPUs, every on or off combination of the twoAwayX, twoAwayY and twoAwayZ flags was tested. For example, twoAwayX=yes means that the simulation box size is halfed in the x direction. The results for different simulation box sizes are summarized in Figure 3. The majority of calculations are of the non-bonded type, and this is exactly what is off-loaded to the GPU in the CUDA-enabled version of NAMD. The non-bonded interactions (including electrostatic and van der Waals interactions) are calculated beyond the cutoff distance on the GPUs. The rest of the calculations such as the bonded interaction are done on the CPU. As the cutoff gets larger, the ratio of non-bonded to bonded calculations increases. This greatly favors the GPU as can be seen in the figures. However, most real-world simulations use cutoff values that lead to the CPU being used more heavily than the GPU. As such, the huge possible performance gain using GPUs cannot be fully passed on to the scientist. Using the CUDA streaming API for asynchronous memory transfers and kernel invocations to overlap GPU computation with communication and other work done by the CPU yields speedups of up to a factor of nine times faster than CPU-only runs.

Conclusion

Application acceleration using GPUs has shown promising results for the NAMD molecular dynamics simulation software at the LRZ Supercomputing Centre. The speed-up of the GPU-enhanced version of NAMD depends most strongly on the cutoff parameter and also on the twoAwayX, twoAwayY, and twoAwayZ options. These options reflect how the patches are arranged and the non-bonded interactions are passed on to the GPU (some arrangements are much better than others). To get the optimum performance out of GPUs, fine-grained parallelism and highly optimized code are needed.

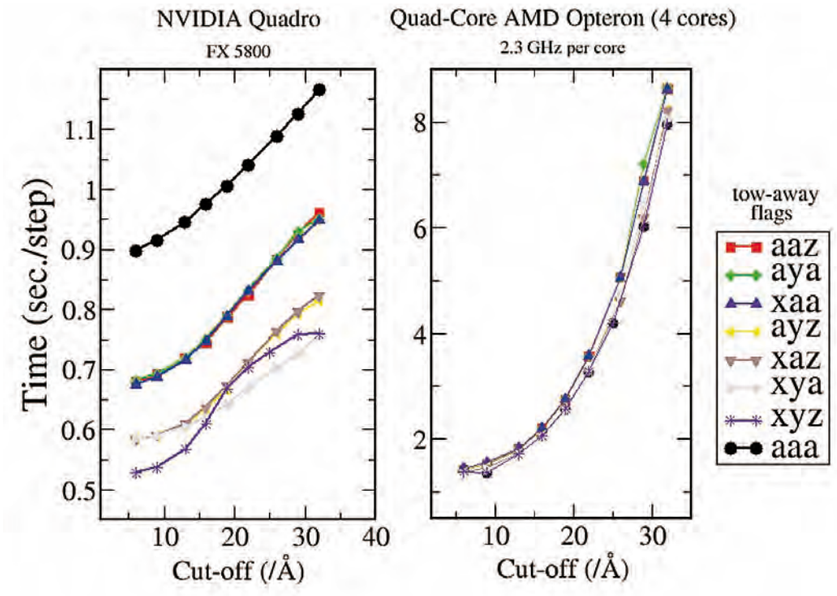

Figure 3: Variation of time performance as a function of Cut-off and two away flags.

References

[1] **Owens, J. D., Houston, M., Luebke, D., Green, S., Stone, J. E., Phillips, J. C.** "GPU computing" Proceedings of the IEEE, 96:879-899, May 2008

[2] **Phillips, J. C., Braun, R., Wang, W., Gumbart, J., Tajkhorshid, W., Schulten, K. Chipot, Ch., Skeel, R. D., Kale, L., Villa, E.,** "Scalable Molecular Dynamics with NAMD", Journal of Computational Chemistry, 26:1781-1802, 2005

[3] Nvidia CUDA (Compute Unified Device Architecture) "Programming Guide", Nvidia, Santa Clara, CA 2007

[4] http://www.ks.uiuc.edu/Research/namd/

[5] **Allen, M. P., Tildesley, D. J.** "Computer Simulation of Liquids", Oxford University Press, New York, 1987

[6] **McCammon, J. A., Harvey, S. C.** "Dynamics of Proteins and Nucleic Acids", Cambridge University Press, Cambridge, 1987

[7] **MacKerell, Jr. A. D., Banavali, N., Foloppe, N.** "Development and Current Status of the CHARMM Force Field for Nucleic Acids", Biopolymers 56: 257–265 (2001).

[8] **Ma, Q., Izaguirre, J. A., Skeel, R. D.** "Verlet-I/r-RESPA/Impulse is Limited by Nonlinear Instabilities Siam Journal on Scientific Computing, 24:1951-1973, 2003

- Momme Allalen
- Helmut Satzger
- Ferdinand Jamitzky

Leibniz Supercomputing Centre

Storage and Network Design for the JUGENE Petaflop System

In conjunction with the installation of the IBM Blue Gene/P Petaflop system JUGENE at Forschungszentrum Jülich (inSiDE Vol. 7 No. 1, p. 4), the supporting storage and networking infrastructure has been upgraded to a sustained bandwidth of 66 GByte/s, matching the enhanced computational capabilities.

Starting with the installation of the initial 16-rack JUGENE system in 2007, scalable parallel file services have been moved outside of the supercomputer systems and are now provided by a dedicated storage cluster JUST, which is based on IBM General Parallel File System (GPFS). There are three main operational considerations leading to this approach:

1. Data is more long-lived than the compute systems it is generated on.

2. There is a need to attach more than a single supercomputing resource to the parallel file system (locally and in different projects, e.g. DEISA).

3. Cutting edge supercomputers pick up many performance-related soft ware and firmware improvements that a pure storage cluster would not need.

Separating out the HPC storage greatly reduces the above interdependencies, and also makes it easier to offer site-wide parallel file services to a more heterogeneous landscape of computing and visualization resources. On the other hand, designing the networking infrastructure around it then becomes more challenging than solutions for a single supercomputer. The JUST cluster uses 10-Gigabit Ethernet as its networking technology, for several reasons:

- The I/O nodes of the Blue Gene/P (as the main client) have on-chip 10 GbE.

- Ethernet is by far the most stable, interoperable, and manageable networking solution.

- 10 GbE bandwidth is adequate, and the ultra-low latencies of other HPC interconnects are not needed for typical HPC storage traffic.

Expanding this infrastructure to support Petaflop systems requires a solution design which addresses some unique challenges which are not present (or not as relevant) at smaller scale.

- The 10 GbE network needs to provide very large port counts, and has to sustain Terabits of bandwidth between the supercomputer(s) and the storage cluster.

- The storage cluster needs to be perfectly balanced in every aspect of the hardware and software stack to be able to deliver the raw performance of thousands of disks to the end users.

- The solution needs to be resilient to faults and manageable. This becomes ever more critical with the exploding number of components both on the compute and the storage side.

- Memory is the most precious resource on petascale supercomputers, and the memory-efficient organization and scaling of application I/O patterns is an active area of our petascale research.

The Networking Layer

The networking solution needs to support a sustained GPFS bandwidth of 66 GByte/s and has to provide 10 GbE ports for 700 to 800 participants: the 600 I/O nodes of the 72-rack Blue Gene/P, at least 100 ports in the JUST storage cluster, plus connectivity to additional systems. To achieve this, JSC together with IBM started the design and planning phase well ahead of the actual deployment. Even today there is no switch solution on the market that offers 700 to 800 10 GbE ports on a single backplane, and the combination of multiple switches introduces several design and bandwidth limitations, e.g. due to port channel restrictions (IEEE802.3ad) and the spanning tree algorithm.

After discussing different network designs and even routing protocol based network solutions (e.g. equal-cost multi-path), a rather flat network topology was chosen, avoiding the risk of badly converging routing protocols and additional costs for big pipes between switch fabrics. Based on the knowledge of the behaviour of GPFS as the main application, a network layout was invented that uses four switch chassis but keeps all high-bandwidth storage access traffic local within each of the switch fabrics.

To achieve this, quad-homed GPFS servers were deployed that offer direct storage access local at each switch fabric. The high number of Blue Gene/P I/O nodes are equally distributed over those 4 switch fabrics. Rather than attaching whole BG/P racks to a single fabric (which would limit the rack to that one fabric), the 8 I/O nodes of a BG/P rack are distributed over all four fabrics, so smaller BG/P jobs can also benefit from the aggregate bandwidth of all four planes.

The following candidates have been evaluated for the switch fabric:

- Cisco Nexus 7000 in its 18-slot version

- Force10 E1200i as a 14-slot switch

- Myricom Myri-10G switches

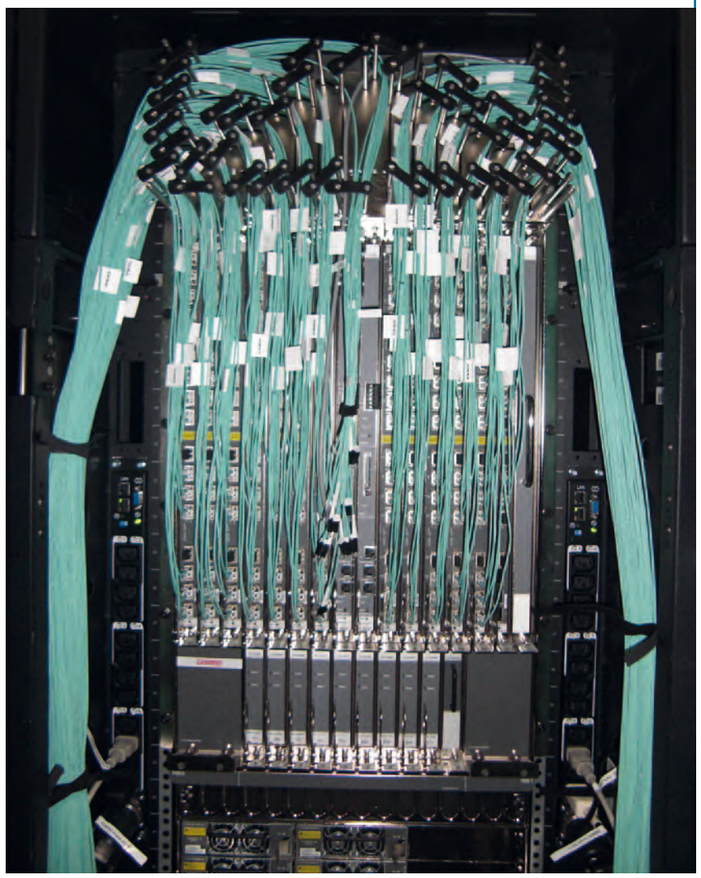

Figure 1: One of the four Force10 E1200i 10 GbE switches

While each alternative had some advantages and some disadvantages, especially in the area of new upcoming DataCenter Ethernet standards like PerPriority Flow Control (IEEE802.1Qb) and Multichassis Etherchannel, at the end the major decision criteria were the proven stable functionality of the existing single E1200 chassis for just the functions needed for JSC's network setup in combination with the assured availability of the desired hardware in sync with the overall Blue Gene/P and storage cluster deployment timeline.

Using the knowledge of Force10's module and ASIC design in combination with the known GPFS traffic patterns, JUST servers and BG/P I/O nodes are connected to the linecard ports in a layout that minimizes the congestion probability on the 4x over-subscribed linecards (which are required to achieve the large portcounts). At the end of the multi-stage deployment phase and some additional network optimizations (introduction of Jumbo frames, deployment of end-to-end flow control, tuning of network options and buffers), the chosen network design has proven to be the ideal solution for JSC's Blue Gene/P and POWER6 based storage cluster deployment. The traffic patterns observed in production mode via JSC's RMON and SFlow/Netflow based monitoring match the assumptions that were used as a base for this specific tailored network solution.

Using Ethernet as the base technology is a perfect fit for heterogeneous aspects like integrating the Nehalem-based JUROPA cluster, and interconnecting with different campus and project networks as well as visualization solutions.

GPFS Storage Design and Usage Model

The design of the hardware building blocks for the expanded JUST storage cluster is driven by three factors:

1. The target of a sustained aggregate bandwidth of 66 GByte/s.

2. The requirement to use quad-homed GPFS/NSD servers with sufficient bandwidth to drive the four 10 GbE network fabrics.

3. The requirement of a seamless migration from the initial JUST storage cluster installed in 2007 (and the goal to re-use parts of its hardware, in particular the SATA disks).

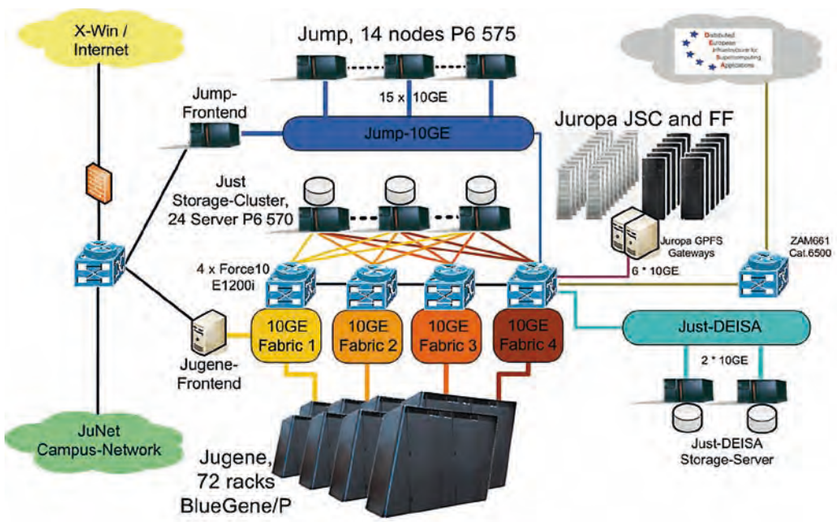

Figure 2: JSC Network Overview

As in the initial JUST cluster, the GPFS metadata is stored on dedicated metadata building blocks, with two IBM DS5300 storage subsystems and 15k FC disks. Metadata is protected both by Raid1 arrays and by GPFS replication.

Several combinations of servers and disk storage subsystems have been evaluated for the GPFS data building block. For best price/performance, nearline SATA disks are used and protected as 8+2P Raid6 arrays. The resulting GPFS data building block is shown in Figure 3:

- 3 x GPFS-DATA servers p6-570, each: 2 x CEC, 8 core, 8 x dual-port FC4 adapters, (8 ports used), 4 x 10 GigE link into BG/P network

- 2 x DS5300 storage subsystems, each: 16 x FC host ports (12 ports used), 24 x EXP5000, each @ 16 x SATA drives (5/6 are 1 TB drives, 1/6 are 500 GB drives reused from JUST-1)

To achieve the desired aggregate GPFS bandwidth, eight of these building blocks are deployed providing a total of 4.2 PByte usable capacity. This storage is partitioned into three classes of filesystems with different usage characteristics:

1. A large scratch filesystem $WORK with roughly half the capacity and bandwidth, not backed up or archived. Blue Gene/P production jobs should generally use this filesystem.

2. $HOME filesystems, with daily incremental backups.

3. $ARCH filesystems, with daily incremental backups and migration to tape storage through a combination of the IBM GPFS Information Lifecycle Management (ILM) features and IBM TSM/HSM.

Figure 4 depicts the user view of the GPFS filesystems on the expanded JUST cluster. Note that to cope with the performance impact of the increased number of files resulting from an 5x enlarged Blue Gene/P system, the number of $ARCH filesystems has been increased from one to three and users are distributed across those.

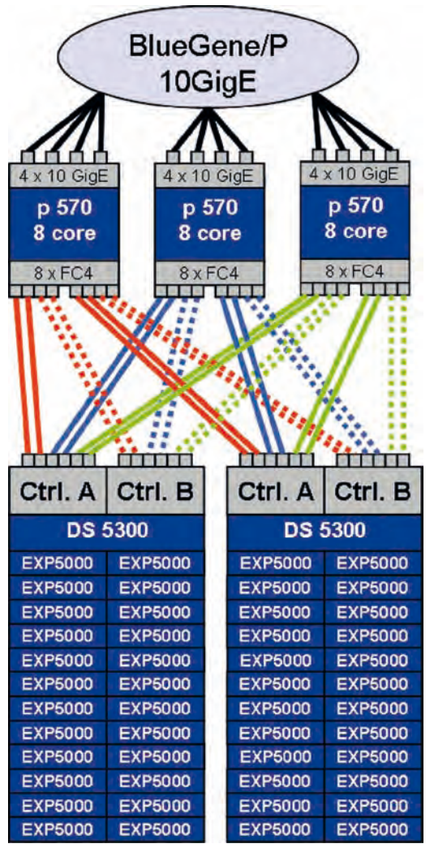

Figure 3: One JUST Building Block ($WORK uses four, $HOME and $ARCH share the other four building blocks)

GPFS Performance Aspects

To reach the maximum performance of a GPFS cluster, a number of "GPFS Golden Rules" need to be followed:

- Performance primarily depends on the number of disks within the file system (disks do not get faster anymore).
- The hardware configuration should be symmetric and all system components must be balanced.
- Settings of the GPFS parameters like the filesystem blocksize must be suitable and should match the physical layout.
- Changed characteristics from one generation of storage HW to the next (e.g. DS4000 to DS5000) should be reflected in the GPFS configuration.
- Changed usage patterns (e.g. an 5x increase in the number of files generated) will also have an impact on the filesystem performance that needs to be considered.

During the deployment phase, a number of these well-known best practices have been re-visited. Some of the design and configuration choices that were optimal with the 16-rack Blue Gene/P and generation-2007 storage hardware are no longer appropriate for a 72-rack Petaflop Blue Gene/P and generation-2009 storage hardware.

On the system level, these invaluable lessons learned (during installation, from early adopters and studies by the JSC application support team) have already been incorporated into the production setup. On the application level, efforts are ongoing to optimize the applications I/O patterns so they scale on Petaflop systems to the same degree that the MPI communication scaling has reached.

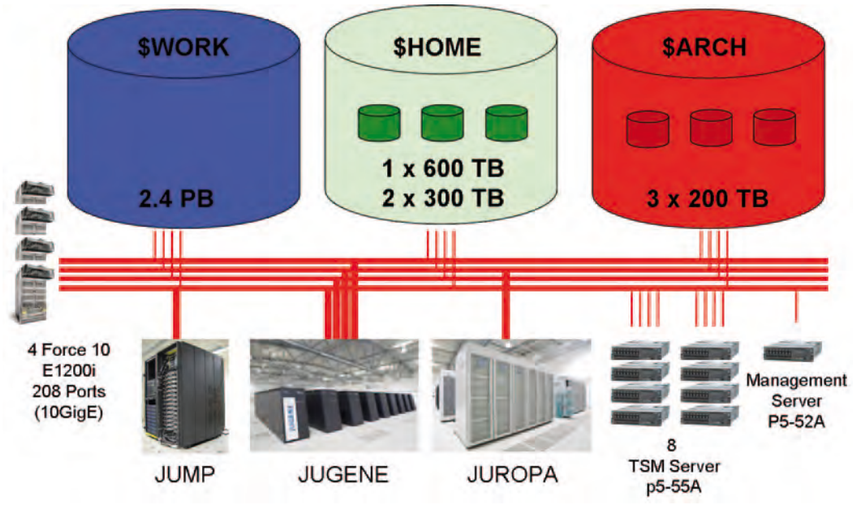

Figure 4: User view of the JUST GPFS filesystems and the connected Supercomputers at JSC

- Olaf Mextorf[1]
- Ulrike Schmidt[1]
- Lothar Wollschläger[1]
- Michael Hennecke[2]
- Karsten Kutzer[2]

[1] Jülich Supercomputing Centre, Forschungszentrum Jülich,

[2] IBM Deutschland GmbH

QPACE

QPACE (Quantum Chromodynamics Parallel Computing on the Cell) is a massively parallel and scalable computer architecture optimized for Lattice Quantum ChromoDynamics (LQCD). It has been developed in cooperation between several academic institutions (SFB TR 55 under the leadership of the University of Regensburg) and the IBM Deutschland Research and Development GmbH.

At JSC, a 4-rack QPACE system is installed. The building block is a node card comprising an IBM PowerXCell 8i processor and a custom FPGA-based network processor. 32 node cards are mounted on a single backplane and eight backplanes are arranged inside one rack, hosting a total of 256 node cards. The closed node card housing, which is connected to a liquid-cooled cold plate, acts as a heat conductor thus making QPACE a highly energy-efficient system.

To remove the generated heat a cost-efficient liquid cooling system has been developed, which enables high packaging densities. The maximum power consumption of one QPACE rack is about 32 kW. The 3D-torus network interconnects the node cards with nearest-neighbor communication links driven by a lean custom protocol optimized for low latencies. For the physical layer of the links 10 Gigabit Ethernet PHYs are used providing a bandwidth of 1 GB/s per link and direction.

In the context of PRACE work-package 8, promising technologies for future supercomputers have been evaluated and QPACE was identified as one of the advanced prototypes. In order to extend the range of applications for QPACE beyond LQCD, an implementation of the high-performance LINPACK (HPL) benchmark was a first step. Here, different from QCD applications, the transfer of large messages and collective operations must be supported efficiently. This required extensions of both network communication and the software stack.

A special FPGA bit-stream and a library of MPI functions for message passing between compute nodes has been developed at JSC in co-operation with IBM. The HPL benchmark produced excellent results (43.01 TFlop/s on 512 nodes) and the 4 rack system in Jülich corresponds to an aggregate peak performance of more than 100 TFlop/s.

Providing 723 MFlop/s/W QPACE was recognized as the most energy-efficient supercomputer worldwide and ranked top in the Green500 list at the Supercomputing Conference 2009 in Portland, Oregon.

- Willi Homberg

Jülich Supercomputing Centre, Forschungszentrum Jülich

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

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

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

Contact:
Leibniz Supercomputing Centre

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

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

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

Compute servers currently operated by LRZ are

System	Size	Peak Performance (TFlop/s)	Purpose	User Community
SGI Altix 4700 "HLRB II" Intel IA64 19 x 512-way	9,728 Cores 39 TByte	62.3	Capability Computing	German Universities and Research Institutes, DEISA
Linux-Cluster Intel Xeon EM64T/ AMD Opteron 2-, 4-, 8-, 16-, 32-way	4,438 Cores 9.9 TByte	36.3	Capacity Computing	Bavarian and Munich Universities, D-Grid, LCG Grid
SGI ICE Intel Nehalem 8-way	512 Cores 1.6 TByte	5.2	Capability Computing	Bavarian Universities, PRACE
SGI Altix 4700 & SGI Altix 3700 BX2 Intel IA64 128 & 256-way	384 Cores 1 TByte	2.5	Capability Computing	Bavarian Universities
Linux Cluster Intel IA64 2, 4- and 8-way	220 Cores 1.1 TByte	1.3	Capacity Computing	Bavarian and Munich Universities

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. Dr. hc. Michael M. Resch
Nobelstraße 19
70569 Stuttgart
Germany

Phone +49-711-685-8 72 69
resch@hlrs.de
www.hlrs.de

Compute servers currently operated by HLRS are

System	Size	Peak Performance (TFlop/s)	Purpose	User Community
NEC Hybrid Architecture	12 16-way nodes SX-9 with 8 TByte main memory + 5,600 Intel Nehalem cores 9 TB memory and 64 NVIDIA Tesla S1070	146	Capability Computing	German Universities, Research Institutes and Industry, D-Grid
IBM BW-Grid	3,984 Intel Harpertown cores 8 TByte memory	45.9	Grid Computing	D-Grid Community
Cray XT5m	896 AMD Shanghai cores 1.8 TByte memory	9	Technical Computing	BW Users and Industry
AMD Cluster	288 AMD cores 1.6 TByte memory	3.7	Technical Computing	Research Institutes and Industry

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

View of the HLRS BW-Grid IBM Cluster (Photo: HLRS)

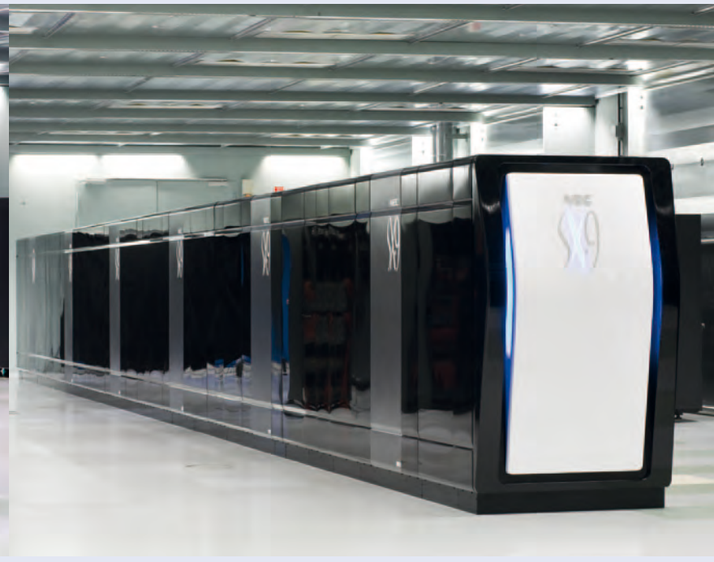

View of the HLRS hybrid NEC supercomputer (Intel Nehalem / NVIDIA S1070 / SX-9) (Photos: HLRS)

JÜLICH
FORSCHUNGSZENTRUM

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

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

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

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

Higher education for master and doctoral students in cooperation e.g. with the German Research School for Simulation Sciences, a joint venture of Forschungszentrum Jülich and RWTH Aachen University.

Contact:
Jülich Supercomputing Centre (JSC)
Forschungszentrum Jülich

Prof. Dr. Dr. Thomas Lippert
52425 Jülich
Germany

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

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

Compute servers currently operated by JSC are

System	Size	Peak Performance (TFlop/s)	Purpose	User Community
IBM Blue Gene/P "JUGENE"	72 racks 73,728 nodes 294,912 processors PowerPC 450 144 TByte memory	1,002.6	Capability computing	European Universities and Research Institutes, PRACE, DEISA
Intel Linux CLuster "JUROPA"	2,208 SMT nodes with 2 Intel Nehalem-EP quad-core 2.93 GHz processors each 17,664 cores 52 TByte memory	207	Capacity and Capability Computing	European Universities, Research Institutes and Industry, PRACE, DEISA
Intel Linux CLuster "HPC-FF"	1,080 SMT nodes with 2 Intel Nehalem-EP quad-core 2.93 GHz processors each 8,640 cores 25 TByte memory	101	Capacity and Capability Computing	EU Fusion Community
IBM Cell System "QPACE"	1,024 PowerXCell 8i processors 4 TByte memory	100	Capability Computing	QCD applications SFB TR55, PRACE
IBM Cell System "JUICEnext"	35 Blades 70 PowerXCell 8i processors 280 GByte memory	7	Capability computing	German Research School
AMD Linux Cluster "SoftComp"	125 compute nodes 500 AMD Opteron 2.0 GHz cores 708 GByte memory	2.5	Capability computing	EU SoftComp Community
AMD Linux Cluster "JUGGLE"	44 compute nodes 176 AMD Opteron 2.4 GHz cores 352 GByte memory	0.85	Capacity and capability computing	D-Grid Projects

NIC Symposium 2010 in Jülich

The John von Neumann Institute for Computing (NIC) was established in 1998 by Forschungszentrum Jülich and Deutsches Elektronen-Synchrotron DESY to support the supercomputer-oriented simulation in natural and engineering sciences. In 2006, GSI Helmholtzzentrum für Schwerionenforschung joined NIC as a contract partner. The core task of NIC is the peer-reviewed allocation of supercomputing resources to computational science projects from Germany and Europe. The NIC partners support supercomputer-aided research in science and engineering through a three-way strategy:

- Provision of supercomputing resources for projects in science research and industry
- Supercomputer-oriented research and development by NIC research groups in selected fields of physics and natural sciences.
- Education and training in all areas of supercomputing by symposia, workshops, summer schools, seminars, courses and guest programmes for scientists and students.

Every two years the NIC organizes a symposium to give an overview of the activities and results obtained by the NIC projects and research groups in the last two years. The fifth NIC Symposium took place at the Forschungszentrum Jülich from February 24 - 25, 2010. It was attended by more than one hundred scientists that had been using the supercomputers JUMP, JUROPA, and JUGENE in Jülich, and the APE topical computer at DESY in Zeuthen in their research. Fifteen talks and a large poster session covered selected topics. To accompany the conference, an extended proceedings volume was published that can also be viewed online at

www.fz-juelich.de/nic/symposium

Participants of the NIC Symposium 2010

17th EuroMPI Conference at HLRS

EuroMPI 2010 will be held at Stuttgart, Germany on September 12 - 15, 2010. HLRS has a long tradition in MPI and has contributed to the definition of the standard and to the proliferation of MPI over the years. It was therefore only natural to accept an invitation to host the next EuroMPI meeting at Stuttgart. The 17th European MPI Users' Group Meeting will be a forum for users and developers of MPI and other message passing programming environments. Through the presentation of contributed papers, poster presentations and invited talks, attendees will have the opportunity to share ideas and experiences to contribute to the improvement and furthering of message passing and related parallel programming paradigms. With HLRS focusing on industrial co-operations and applications EuroMPI will aim at bringing science and industry in the field closer together. The potential of MPI for industrial applications will be one of the issues discussed at the HLRS. Furthermore the key questions of scalability will be addressed. Handling one million cores is a task that is beyond currently available implementations and the workshop will be a good chance to discuss new ideas and strategies. The link to industry will also be given in the planned social event. The Mercedes Museum at Stuttgart is the perfect place to get an overview of more than 100 years of automobile history embedded in the overall world history. This visit will bring us close to the industrial heart of Stuttgart that beats for cars and everything automotive engineers can imagine.

Invited Speakers:
Jack Dongarra, UTK, USA
Bill Gropp, UIUC, USA
Rolf Hempel, DLR, Germany
Jesus Labarta, BSC, Spain
Jan Westerholm, AAU, Finnland

See also: www.eurompi2010.org

Euro MPI 2010

September 12-15, 2010 Join us!

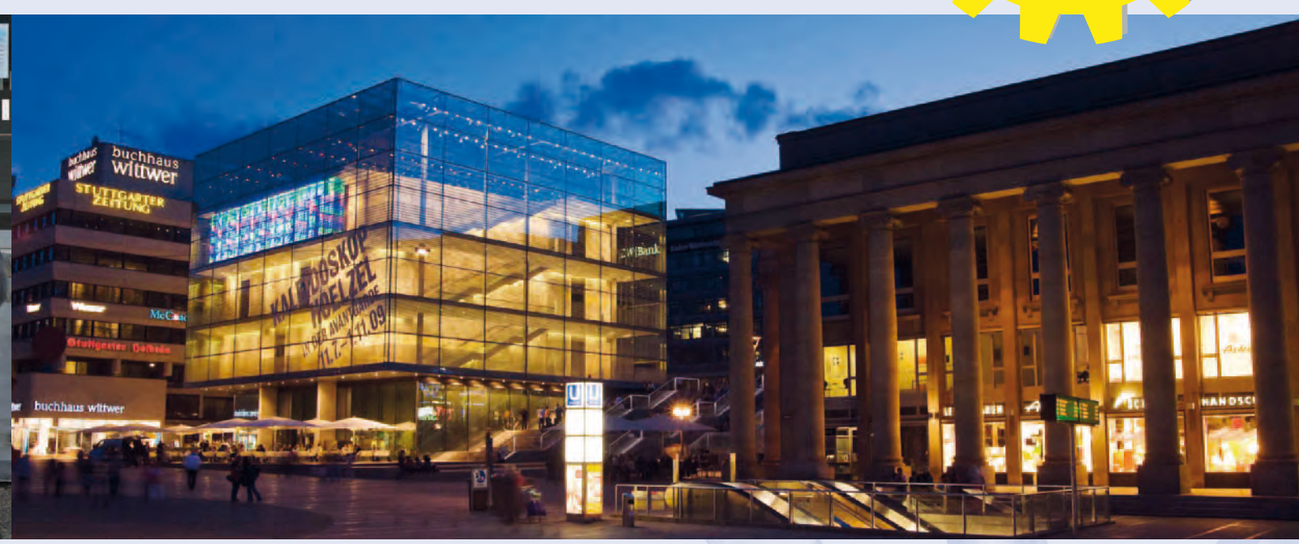

Schlossplatz Stuttgart at night

Storage Consortium at HLRS

On March 11, 2010, the 7th Conference for IT Professionals and Data Center Operators was held on the grounds of the University of Stuttgart in the large lecture hall of the HLRS. This conference was designed and organized by the Storage Consortium – information and communication platform of leading IT manufacturers operating in the German-speaking market. The event was held under the slogan "Energy Efficiency, Virtualization and Cloud Computing".

According to the Managing Director and initiator of the Storage Consortium, Mr. Norbert Deuschle, the one-day conference turned out to be a great success, especially due to the dedicated support by both the management and staff involved at HLRS. Mr. Deuschle quote: "Thanks to the willingness and proactive support by Prof. Dr. Michael Resch, Dipl.-Ing. Peter W. Haas as well as Dr. Thomas Bönisch of HLRS with his exciting keynote speech, we could now carry out our seventh successful user conference on the grounds of a large data center operator. After Leibnitz Rechenzentrum, DESY, Max-Planck Institute for Extraterrestrial Physics and other exciting conference locations, here at HLRS over 100 participants have given us a clear confirmation by ways of their positive feedback on the symposium that we do meet the needs of IT-/Storage representatives. Therefore, another conference in 2010 is still in the pipeline".

The event was supported by the following partners of the Storage Consortium: Riverbed, Cisco Systems, NetApp, Quantum, Bull, Symantec, F5 Networks, HP StorageWorks, 3PAR and Fujitsu Technology Solutions. For user presentations (practical relevance), the companies Sonepar Germany, Freudenberg IT, T-Systems International and, of course, the HLRS data center itself could be won.

Due to its IT infrastructure, HLRS is very well suited to host such an event, in particular, of course, in the areas of storage networks (10-/40-Gbit/s Ethernet) as well as systems (disk, tape). During the lunch break, visitors to the conference had the opportunity of going on a total of four tours in order to convince themselves. In addition to the large data center with NEC Super Computing and storage systems, especially the data management aspects and the Storage Management (File Management, HSM) found special attention. This has to be seen, of course, in the context of ever-increasing capacities of unstructured data to companies. In this context, quite naturally, visitors were highly interested in having a close look at HLRS' comprehensive disk and tape automated storage technology used to provide long-term storage for even the largest data sets (Petabyte-Archiving, Backup).

The conferences of the Storage Consortium are strongly focused on technology, planning, and application requirements of most advanced virtualization and storage architectures and are targeted at managers and operators of data centers, storage administrators, IT managers and architects, and application server administrators, but also controllers and technical buyers. Typical concerns of users, such as the planning and selection of modern storage architectures, are of particular importance. The overall motto is "from user to user," and the contents of the presentations are only partially provided by the manufacturers. This is to ensure practical relevance as much as possible which is further improved by the informal way the various participants are exchanging information between each other.

The agenda of the Storage Consortium meeting at HLRS on March 11, 2010, did include the following current topics:

- Virtualization solutions as a basis for cloud computing in modern data centers
- Solutions in the field of storage systems and file system management at HLRS
- Capacity optimized backups of virtualized server environments, data deduplication and virtualized file management
- Virtualized Data Center Architectures with standardized, virtualized IT Services
- Storage virtualization – Consistent step in protecting virtual infrastructures
- Efficient data management through file virtualization in the data center
- Storage consolidation and centralization of desktop services with WAN optimization
- Use of scenarios and reports for data deduplication
- Convergence of server, network and storage, benefits of virtualization in the storage environment and user motivations
- Solutions for cloud computing
- Virtualized storage resources as a basis for cloud computing in modern data centers.

Inquiries to the Storage Consortium, please contact: Mr. Norbert E. Deuschle, info@storageconsortium.de

See also: www.storageconsortium.de

$C^2A^2S^2E$ and JSC Co-operate to Achieve "Digital Aircraft"

Computer simulations are a key technology for the engineering as well as natural sciences. Nowadays, numerical simulations are an indispensable ingredient to all aeronautical research activities since they allow insights which would be impossible with theory and experiment alone. Thus simulations promise to reduce the development risks and to shorten the development cycles. It is expected that the next generation of aircrafts will be developed and tested completely in the computer ("digital aircraft").

The Center for Computer Applications in AeroSpace Science and Engineering ($C^2A^2S^2E$), founded by DLR and Airbus with support from the European Union and the federal state Niedersachsen, is a well-known competence centre for numerical aircraft simulations. Besides its activities on improving the modelling of the physical processes and the numerical algorithms it operates as a topical computer centre the most powerful parallel cluster for aeronautics research in Europe. Especially, the CFD code TAU is developed, maintained and applied in projects by $C^2A^2S^2E$ for its customers in the European aircraft industry.

But the computing resources available at a topical centre are not sufficient to simulate a virtual aircraft with all its multidisciplinary interactions. Therefore $C^2A^2S^2E$ needs access to capability computers at the top level of the European HPC ecosystem like JSCs JUGENE and JUROPA. Furthermore the efficiency and scalability of the simulation software, i.e. the TAU code, have to be improved to allow calculations using tens to hundreds of thousands of cores. The necessary optimization/adaptation of this research code will be supported by supercomputing experts from JSC. The formal basis for the common research activities is the co-operation agreement 'Numerical simulations for aeronautical research on supercomputers' signed by Deutsches Zentrum für Luft- und Raumfahrt (DLR) and Forschungszentrum Jülich in January 2010.

As a first step the TAU code has been implemented on JUGENE and JUROPA and successful benchmarks were performed up to several thousand cores. In order to intensify this collaboration and to achieve the "digital aircraft" JSC and $C^2A^2S^2E$ will work in future together on:

- Optimization of simulation software for up-to-date computer architectures, e.g. multi-core architectures or hybrid systems using GPUs, FPGAs, etc.
- Development of scalable CFD methods for Petaflop systems
- Evaluation of future trends in computer/cluster architectures
- Efficient usage of HPC resources
- Training of students/young scientists
- Initiation of common research projects in HPC.

Large-scale Projects at GCS@FZJ

Two calls for large-scale projects were issued by the Gauss Centre for Supercomputing (GCS) in 2009, one in April and one in July. Projects are classified as "large scale" if they require more than 5% of the available processor core cycles on each member's high-end system. In the case on the petaflop supercomputer JUGENE in Jülich this amounts to 24 rack months for a one year application period. At its meeting in June the NIC Peer Review Board – in collaboration with the review boards of HLRS and LRZ – decided to award the status of large scale project to two projects: one from the field of fluid dynamics (Prof. N. Peters, RWTH Aachen: "Geometrical Properties of Small Scale Turbulence") and one from elementary particle physics (Dr. K. Jansen, DESY Zeuthen: "QCD Simulations with Light, Strange and Charm dynamical Quark Flavours").

In October, four more projects could be awarded that status: again one from the field of fluid dynamics (Dr. J. Harting, Universität Stuttgart: "Lattice Boltzmann simulation of emulsification processes in the presence of particles and turbulence"), and three from elementary particle physics (Prof. Z. Fodor, Universität Wuppertal: "Lattice QCD with 2 plus 1 flavours at the physical mass point"; Prof. G. Schierholz, DESY: "Hadron Physics at Physical Quark Masses on Large Lattices"; Prof. S. Katz, Eötvös University, Budapest, Hungary: "QCD Simulations with Light, Strange and Charm dynamical Quark Flavours").

Altogether there are currently six large scale scientific projects working on JUGENE with an allocation of more than 40% of the available CPU cycles.

HLRB and KONWIHR Results and Review Workshop

On December 8 and 9, 2009, the Leibniz Supercomputing Centre (LRZ) hosted the "HLRB and KONWIHR Results and Review Workshop". On the one-hand side, this workshop acts as a review of the scientific research performed on the HLRB II supercomputer in the years 2008 and 2009, and on the other side, it served as a forum for the exchange of experiences between users of different projects facing similar problems as well as between users and LRZ's HPC support staff.

The workshop was attended by more than 100 participants among them scientists and researchers from all over Germany, members of the steering committee, and LRZ's HPC support staff. The workshop featured an extraordinary interdisciplinarity since the covered research areas include Computer Sciences, Mechanical Engineering (predominantly concerned with fluid-dynamics problems), Astrophysics, High-Energy Physics, Geo Sciences, Condensed Matter Physics (including solid-state physics as well as plasma physics), Chemistry, and Bio Sciences. During the two days LRZ could offer 42 projects the opportunity to present their results in talks. The talks were divided into two parallel sessions with 10 talks per day. Additionally, on each of the two days a plenary talk was given. Due to the considerable interdisciplinarity of the workshop the speakers had been asked to prepare their talks in a manner, such that even non-specialists could follow their talks to a large extend. This guidline was followed by the speakers in an outstandingly impressive way leading not only to interesting talks on a high scientific level, but also to fruitful discussions between participants from the various disciplines. The feature of interdisciplinarity was most outstanding during the plenary talks, which – following their nature – were attended by all workshop participants yielding an audience of even broader diversity in the disciplines than was faced by the speakers in the parallel sessions.

On the first day of the workshop, Bernd Brügmann (an astrophysicist fom the University of Jena) was speaking on "Black Holes and Gravitational Waves". And on the second day Sven Schmid (working at NASA for a joined project of NASA and the University of Stuttgart) gave a talk on "Characterization of the Aeroacoustic Properties of the SOFIA Cavity and its Passive Control". Both speakers earn a special thank for their outstanding talks, as they managed to give a general overview on the one-hand side, while presenting the challenging problems faced in their project on the other side.

Apart from the talks, LRZ offered the possibilty to put up posters, which was taken by another dozen of projects resulting in active discussions among the workshop participants. The workshop participants submitted seventy contributions for the workshop proceedings. Currently, these contributions are in the process of being reviewed by referees. Presumably, the proceedings will be published in spring 2010.

Since end of 2009 LRZ is actively preparing the procurement of the successor system to the HLRB II supercomputer. As the new system will provide a Peta-scale performance, the workshop also gave its participants the unique opportunity to present new challenging projects and their requirements on future hardware and software. LRZ for its part sketched its present plans for this procurement.

LRZ would like to thank the participants of the workshop for their important and interesting research performed on and justifying the HLRB II supercomputer, for their fruitful and comfortable collaboration with LRZ's HPC support staff and for their active participation in the workshop.

German-Israeli Co-operation: 24th Umbrella Symposium

JSC hosted the "Umbrella Symposium for the Development of Joint Co-operation Ideas", an annual event run jointly by Technion Haifa, RWTH Aachen University and Forschungszentrum Jülich with changing topics since 1984. This year's symposium took place from January 18 - 20, 2010, during which around 50 scientists attended more than 20 presentations and follow-up discussions on the topic of "Modelling and Simulation in Medicine, Engineering and Sciences". This wide-ranging meeting included contributions on surface science, medical engineering, biophysics, nanoelectronics and virtual reality, but with a common emphasis on numerical modelling.

An important aspect of the symposium is to give attendees the opportunity to build new cooperations in their particular field: a call for Umbrella travel grants for young scientists will be published in spring this year to support a selected number of promising ideas arising from the symposium: eight such potential partnerships were already proposed by the end of the meeting. The next symposium on the same topic will be hosted by RWTH Aachen in early 2011.

Blue Gene Extreme Scaling Workshop 2009

From October 26 - 28, JSC organized the 2009 edition of its Blue Gene Scaling Workshop. This time, the main focus was on application codes able to scale-up during the workshop to the full Blue Gene/P system JUGENE which consists of 72 racks with a total of 294,912 cores – the highest number of cores worldwide available in a single system.

Interested application teams had to submit short proposals which were evaluated with respect to the required extreme scaling, application-related constraints which had to be fulfilled by the JUGENE software infrastructure and the scientific impact that the codes could produce. Surprisingly, not only a handful of proposals were submitted, as the organizers had expected, but ten high-quality applications could be selected, among them two 2009 Gordon Bell Prize finalists. Applications came from Harvard University, Massachusetts Institute of Technology, Rensselaer Polytechnic Institute, Argonne National Laboratory, University of Edinburgh, Swiss Institute of Technology (Lausanne), Instituto Superior Técnico (Lisbon) as well as from RZG MPG, DESY Zeuthen and JSC.

During the workshop, the teams were supported by four JSC parallel application experts, the JUGENE system administrators and one IBM MPI expert; however, the participants shared a lot of expertise and knowledge, too. Every participating team succeeded to submit one or more successful full 72 rack jobs during the course of the workshop, and executing many smaller jobs to investigate the scaling properties of their applications along the way. A total of 398 jobs were launched using 135.61 rack days of the total 169.3 rack days reserved for the workshop. This is an 80% utilization which is astonishing especially considering that the system was only partially available on the third day due to necessary hardware repairs. IBM experts repaired half of the machine at a time while the other half was still used for the workshop.

Many interesting results were achieved. For example, the Gordon Bell team of Rensselaer Polytechnic Institute focused on a strong scaling study of our implicit, unstructured mesh based flow solver. An adaptive, unstructured mesh with 5 billion elements was considered in this study to model pulsatile transition to turbulence in an abdominal aorta aneurysm geometry that was acquired from image data of a diseased patient. At full system scale, they obtained an overall parallel efficiency of 95% relative to the 16 rack run which is considered as the base of their strong scaling study.

The Lisbon team investigated OSIRIS, a fully relativistic electro-dynamic particle-in-cell code written at IST and UCLA. Their goal was to determine the strong scaling results for two different test cases: the first test case ("warm") simulates the evolution of one single, hot species, while the second case ("collision") simulates two counter-streaming species. They were able to obtain the results for the strong scaling up to 294,912 cores with an efficiency of 81% (warm) and 72% (collision) compared to their baseline of 4,096 CPU cores. They prepared also a set of production runs to simulate, for the first time, a complete astrophysical plasma collision in three dimensions. Due to limited availability of CPU time, these simulations will be subject to future work. Finally, they have defined a work plan to integrate JSC's SION library in OSIRIS.

All three QCD teams (MIT, Edinburgh, DESY) reported linear scaling of their codes to the full system.

Finally, JSC's program MP2C, which couples MD and mesoscopic hydrodynamics, was scaled up to nearly full capacity of JUGENE, i.e. 262,144 cores (64 racks), as the present version requires $P=2^n$. After some adjustments of message passing protocols, a good scaling behaviour was observed. Particle numbers up to $N=8x10^{10}$ were chosen. According to memory requirements, about 10 times more particles could have been applied. Since also parallel I/O capabilities were investigated using SIONlib, benchmarking was limited to this particle number, since for one restart configuration it already produced about 5 TBytes, for which the I/O performance of 8 GB/s for writing and 16 GB/s reading data was measured.

All experiences achieved during the workshop are summarized in the technical report FZJ-JSC-IB-2010-02 available at **http://www.fz-juelich.de/jsc/docs/autoren2010/mohr1/**. On March 22 - 24, 2010, JSC organized a follow-up Blue Gene Extreme Scaling Workshop. A more detailed report about this event will be published in the next edition of inSiDE.

Activities

The 12th HLRS-NEC Teraflop Workshop

The 12th Teraflop Workshop, held on March 15th and 16th, 2009 at HLRS continued the Teraflop workshop series initiated by NEC and the HLRS in 2004. Once again, the covered topics drew a bow from leading edge operating system development to the needs and results of real live applications.

The opening talk was given by Ryusuke Egawa of Tohoku University in Sendai, Japan who gave deeper insight into the possibilities of vector computers in meta-computing environments. The following two talks addressed the SX-Linux project [1]. The first talk in the second session was given by Jochen Buchholz from HLRS who discussed the aims and first ideas of the TIMACS project. He was followed by Georg Hager form the RRZ Erlangen who showed 13 ways to sugarcode the fact that an application is not performing well on modern supercomputers.

The three talks given during the next session by Ulrich Rist from the IAG, University of Stuttgart, Hans Hasse from the University of Kaiserslautern and Matthias Meinke from the Aerodynamic Institute of RWTH Aachen targeted real live applications from CFD and molecular dynamics. The afternoon session was shared by Thomas Soddemann of the Fraunhofer SCAI, who raised the question whether GPU Computing is affordable supercomputing, and Peter Berg from IMK in Karlsruhe. Mr. Berg presented the work on ensemble-simulations done at his department. The second day started with a talk given by Wolfgang Ehlers from the Institute for Mechanics, University of Stuttgart. He presented various computational issues which have to be resolved when moving from particle dynamics to multiphase continous media. The next two talks were given by Philip Rauschenberger from the ITLR, University of Stuttgart, and Arne Nägel form the GCSC, University of Frankfurt who both gave an overview across the applications at their institutes which are related to large scale computers. In the next session, Ralf Schneider, HLRS, presented the efforts necessary when dealing with direct mechanical simulations. Another research topic which poses a similar challenge is the coupling of multi-scale, multi-physics applications which was discussed by Sabine Roller from the German Research School for Simulation Sciences, Aachen.

The Workshop closed with a session in which distributed multi-scale simulations for medical physics applications and film cooling investigated by large-eddy simulation were discussed by Jörg Bernsdorf from the German Research School for Simulation Sciences, Aachen, and Lars Gräf from the ETH Zürich, respectively.

References

[1] SX-Linux project
https://gforge.hlrs.de/projects/sxlinux/

[2] Teraflop Workbench
http://www.teraflop-workbench.de/

Publications in Computational Science and Engineering

Nagel, W. E.; Kröner, D. B.; Resch, M. M. (Eds.), High Performance Computing in Science and Engineering '09, Transactions of the High Performance Computing Center, 1st Edition., 2010, XII, 551 p., Hardcover, ISBN: 978-3-642-04664-3

This book presents the state of the art in simulation on supercomputers. Leading researchers present results achieved on systems of the HLRS for the year 2009. The reports cover all fields of computational science and engineering ranging from CFD to 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 allows comparing performance levels and the usability of various architectures. As HLRS operates the largest NEC SX-8 vector system in the world this book gives an excellent insight into the potential of vector systems. The book covers the main methods in high performance computing. Its outstanding results in achieving the highest performance for production codes are of particular interest for both scientists and engineers.

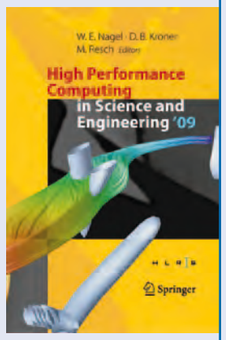

Resch, M. M.; Roller, S.; Benkert, K.; Galle, M.; Bez, W.; Kobayashi, H. (Eds.), High Performance Computing on Vector Systems 2009, 2010, XIII, 250 p., Hardcover, ISBN: 978-3-642-03912-6

The book presents the state of the art in high performance computing and simulation on modern supercomputer architectures. It covers trends in hardware and software development in general and specifically the future of vector-based systems and heterogeneous architectures. The application contributions cover computational fluid dynamics, fluid-structure interaction, physics, chemistry, astrophysics, and climate research. Innovative fields like coupled multi-physics or multi-scale simulations are presented. All papers were chosen from presentations given at the 9th Teraflop Workshop held in November 2008 at Tohoku University, Japan, and the 10th Teraflop Workshop held in April 2009 at HLRS, Germany.

Müller, M. S.; Resch, M. M.; Nagel, W. E. (Eds.), Proceedings of the 3rd International Workshop on Parallel Tools for High Performance Computing, September 2009, ZIH, Dresden, 1st Edition., 2010, 200 p., Hardcover, ISBN: 978-3-642-11260-7, to be published in June 2010

As more and more hardware platforms support parallelism, parallel programming is gaining momentum. Applications can only leverage the performance of multi-core processors or graphics processing units if they are able to split a problem into smaller ones that can be solved in parallel. The proceedings of the 3rd International Workshop on Parallel Tools for High Performance Computing provide a technical overview in order to help engineers, developers and computer scientists decide which tools are best suited to enhancing their current development processes.

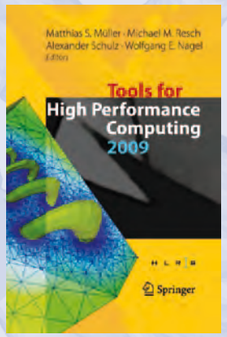

HLRS Scientific Workshops and Courses – Related to Fidelio?

In many of the previous editions of the inSiDE, I reported on past and future workshops and courses, about the new "high-productivity" programming paradigm PGAS (partitioned global address space) and related UPC and CAF courses – *but how does this fit to the pictures on these pages?* – and also about "oldstyle" efficient MPI and OpenMP programming. The courses cover also computational fluid dynamics (CFD), iterative solvers, and modern object oriented paradigms in latest (2008) Fortran standardization. After more then ten years, former participants send now their new PhD students – *still no relation to the pictures!?* – and they come from all over Germany (sometimes also from really far away) and have to stay one week in a boring/nice hotel. Scientific workshops and conferences have exciting social events, but courses? Do participants really prefer hotels' wireless LAN or a thriller in the TV? Should I cite Prof. Chr. Pfeiffer who teaches that thrillers in the evening significantly reduce the learning effort of the day, and also that men prefer thrillers? As usual in scientific workshops, we believe that also in scientific courses, after a day of intensive study and practical work in hands-on sessions[1] people want to see more than only their hotel room. We try to give them together the opportunity to experience the cities of the High Performance Computing centers. A guided walk[2] through the city can produce the necessary hunger for a comfortable dinner. In Jülich, the fortress with its several meters thick walls gives a monumental impression in the night and in Garching at a warm September evening, the "English Garden" invited several participants and lecturers to relax in a beer garden[3]. In Stuttgart after the city tour, Stuttgart's first micro brewery[4] helps to survive[5] from the stressful lectures. Some participants also experience to visit the first TVtower[6] in the world with its beautiful view over Stuttgart at day and night. Besides the city tour, always a tour through the 3D visualization in the CAVE[7] and the computing rooms[8] is included together with a special trip in the Porsche simulator[9].


12


13


9


10


11

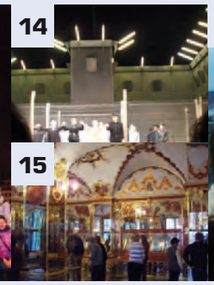
14

15


16

In Dresden, the tour through the famous historical city[10] includes a group photo in front of the "Dresdner Frauenkirche"[11]. This year, based on a special offer of the Semper Opera[12,13], half of the course enjoyed Beethoven's Fidelio with its critical, political, and now already historical stage setting[14], designed in the former "DDR". The premiere was on October 7, 1989, just 32 days before the Fall of the Wall. Some of the attendees visited also the famous art and jewels in the treasury "Grünes Gewölbe"[15] after an early arrival on Sunday. A really rare event is the coincidence of a course with Stuttgart's "Volksfest" (fair). At a ride with the largest mobile observation wheel, you can enjoy the view from 60 m height[16].

This year, the October course at LRZ starts on the last day of Munich's famous "Oktoberfest"…and on the next day the heavy work with theory and practical starts again – but more relaxed. □

Rolf Rabenseifner


1


2


3


4

5


6


7


8

2010 – Workshop Announcements
Scientific Conferences and Workshops at HLRS, 2010
12th Teraflop Workshop (March, 15 - 16)
9th HLRS/hww Workshop on Scalable Global Parallel File Systems (April 26 - 28)
4th HLRS Parallel Tools Workshop (July, not yet fixed)
17th EuroMPI 2010 (Successor of the EuroPVM/MPI Series, September 12 - 15)
High-Performance Computing in Science and Engineering - The 13th Results and Review Workshop of the HPC Center Stuttgart (October 4 - 5)
IDC International HPC User Forum (October 7 - 8)
Parallel Programming Workshops: Training in Parallel Programming and CFD
Parallel Programming and Parallel Tools (TU Dresden, ZIH, February 22 - 25)
Introduction to Computational Fluid Dynamics (HLRS, March 15 - 19)
Iterative Linear Solvers and Parallelization (HLRS, March 22 - 26)
Platforms at HLRS (HLRS, March 29 - 30)
Unified Parallel C (UPC) and Co-Array Fortran (CAF) (HLRS, May 18 - 19)
4th Parallel Tools Workshop (not yet fixed)
Parallel Programming with MPI & OpenMP (TU Hamburg-Harburg, July 21 - 23)
Parallel Programming with MPI & OpenMP (CSCS Manno, CH, August 10 - 12)
Introduction to Computational Fluid Dynamics (HLRS or RWTH Aachen, September 6 - 10)
Message Passing Interface (MPI) for Beginners (HLRS, September 27 - 28)
Shared Memory Parallelization with OpenMP (HLRS, September 29)
Advanced Topics in Parallel Programming (HLRS, September 30 - October 1)
Iterative Linear Solvers and Parallelization (LRZ, Garching, October 4 - 8)
Parallel Programming with MPI & OpenMP (FZ Jülich, JSC, November 29 - December 1)
Unified Parallel C (UPC) and Co-Array Fortran (CAF) (HLRS, December 14 - 15)
Introduction to CUDA Programming (not yet fixed)
Training in Programming Languages at HLRS
Fortran for Scientific Computing (March 1 - 5 and October 25 - 29)
URLs: http://www.hlrs.de/events/
https://fs.hlrs.de/projects/par/events/2010/parallel_prog_2010/

High Performance Computing Courses and Tutorials

GPGPU Programming

Date & Location
June 14 - 16, 2010
LRZ Building, Garching/Munich

Contents
Heterogeneous GPGPU computing promises tremendous acceleration of applications. This programming workshop includes hands-on sessions, application examples and an introduction to CUDA, CAPS, cuBLAS, cuFFT, the Portland Group Fortran Compiler, pycuda, and R.

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

Visualization of Large Data Sets on Supercomputers

Date & Location
June 29, 2010
LRZ Building, Garching/Munich

Contents
The results of supercomputing simulations are data sets which have grown considerably over the years, giving rise to a need for parallel visualization software packages.
The course focuses on the software packages Paraview, Visit and Vapor and their use to visualize and analyse large data sets generated by supercomputer applications which are typically generated in the fields of CFD, molecular modelling, astrophysics, quantum chemistry and similar. Hands-on sessions featuring example applications are given.

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

Parallel Programming with R

Date & Location
July 6, 2010
LRZ Building, Garching/Munich

Prerequisites
Participants should have some basic knowledge in programming with R.

Contents
R is known as a very powerful language for statistics, but it has also evolved into a tool for the analysis and visualization of large data sets which are typically obtained from supercomputing applications. The course teaches the use of the dynamic language R for parallel programming of supercomputers and features rapid prototyping of simple simulations. Several parallel programming models including Rmpi, snow, multicore, and gputools are presented which exploit the multiple processors that are standard on modern supercomputer architectures. Hands-on sessions with example applications are given.

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

Data Mining of XML Data Sets Using R

Date & Location
July 14, 2010
LRZ Building, Garching/Munich

Prerequisites
Participants should have some basic knowledge in programming with R, as well as basic understanding of XML.

Contents
Large data sets are often stored on the web in XML format. Automated access and analysis of these data sets presents a non-trivial task. The dynamic language R is known as a very powerful language for statistics, but it has also evolved into a tool for the analysis and visualization of large data sets.

In this course R will be used to perform and automate these tasks and visualize the results interactively and on the web. The course includes hands-on sessions.

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

Scientific High-Quality Visualization with R

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

Contents
R is a powerful software that can also easily be used for high quality visualization of large datasets of various kinds like particle data, continuous volume data, etc. The range of different plots is comprised of line plots, contour plots, surface plots up to interactive 3D opengl plotting. The course features hands-on sessions with examples.

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

Introduction to Parallel Programming with MPI and OpenMP (JSC Guest Student Programme)

Date & Location
August 3 - 5, 2010
JSC, Research Centre Jülich

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

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

Parallel Programming with MPI, OpenMP amd PETSc

Dates & Locations
August 10 - 12, 2010
Manno (CH), CSCS

November 29 - December 1, 2010
JSC, Research Centre Jülich

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

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

Introduction to Computational Fluid Dynamics

Date & Location
September 6 - 10, 2010
RWTH Aachen

Contents
Numerical methods to solve the equations of Fluid Dynamics are presented. The main focus is on explicit Finite Volume schemes for the compressible Euler equations. Hands-on sessions will manifest the content of the lectures. Participants will learn to implement the algorithms, but also to apply existing software and to interpret the solutions correctly. Methods and problems of parallelization are discussed. This course is based on a lecture and practical awarded with the "Landeslehrpreis Baden-Württemberg 2003" and organized by German Research School for Simulation Sciences.

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

Message Passing Interface (MPI) for Beginners

Date & Location
September 27 - 28, 2010
Stuttgart, HLRS

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

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

High Performance Computing Courses and Tutorials

Shared Memory Parallelization with OpenMP

Date & Location
September 29, 2010
Stuttgart, HLRS

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

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

Advanced Topics in Parallel Programming

Date & Location
September 30 - October 1, 2010
Stuttgart, HLRS

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

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

Iterative Linear Solvers and Parallelization

Date & Location
October 4 - 8, 2010,
LRZ Building, Garching/Munich

Prerequisites
Good knowledge of at least one of the programming languages C/C++ or Fortran; basic UNIX skills. Mathematical knowledge about basic linear algebra and analysis is also assumed.

Contents
The focus of this compact course 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 techniques are presented in the context of real life applications. Hands-on sessions (in C and Fortran) will allow users to immediately test and understand the basic constructs of iterative solvers, the Message Passing Interface (MPI) and the shared memory directives of OpenMP. This course is organized by LRZ, the University of Kassel, the HLRS Stuttgart and IAG.

Webpage
https://fs.hlrs.de/projects/par/events/2010/parallel_prog_2010/G.html

Advanced Fortran Topics

Date & Location
October 11 - 15, 2010
LRZ Building, Garching/Munich

Prerequisites
Course participants should have basic UNIX/Linux knowledge (login with secure shell, shell commands, simple scripts, editor vi or emacs). Good knowledge of the Fortran 95 standard is also necessary, such as covered in the previous winter semester course.

Contents
This course is targeted at scientists who wish to extend their knowledge of Fortran beyond what is provided in the Fortran 95 standard. Some other tools relevant for software engineering are also discussed. Topics covered include

- object oriented features, submodules
- design patterns
- handling of shared libraries
- mixed language programming
- standardized IEEE arithmetic
- I/O extensions from Fortran 2003
- parallel programming with coarrays
- the Eclipse development environment
- source code versioning system.

To consolidate the lecture material, each day's approximately 4 hours of lecture are complemented by 3 hours of hands-on sessions.

Webpage
http://www.lrz.de/services/compute/courses/#TOC16

Parallel Performance Analysis with VAMPIR

Date & Location
October 18, 2010
LRZ Building, Garching/Munich

Prerequisites
Participants are required to have good knowledge of C, C++ or Fortran, as well of the message passing concepts (MPI) needed for parallel programming.

Contents
Running parallel codes on large-scale systems with thousands of MPI tasks typically requires tuning measures in order to reduce MPI overhead, load balancing problems and serialized execution phases. VAMPIR is a tool that allows to identify and locate these scalability issues by generating trace files from program runs that can afterwards be analyzed using a GUI.

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

Parallel I/O and Portable Data Formats

Date & Location
October 25 - 27, 2010
JSC, Research Centre Jülich

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

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

Fortran for Scientific Computing

Date & Location
October 25 - 29, 2010
Stuttgart, HLRS

Contents
This introduction to C++ is taught with lectures and hands-on sessions. This course is organized by HLRS and Institute for Computational Physics. 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/events/

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

Date & Location
November 25 - 26, 2010
JSC, Research Centre Jülich

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

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

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

Date & Location
December 14 - 15, 2010
Stuttgart, HLRS

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

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

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

Publishers

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

Editor & Design

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

Authors

Christian Alig alig@usm.lmu.de
Momme Allalen allalen@lrz.de
Martin Breuer breuerm@earth.uni-muenster.de
Andreas Burkert aburkert@mpe.mpg.de
Cezary Czaplewski czarek@chem.univ.gda.pl
Volker Gaibler vgaibler@mpe.mpg.de
Leonhard Gantner leonhard.gantner@imk.fzk.de
Georg Hager Georg.Hager@rrze.uni-erlangen.de
Ulrich Hansen hansen@earth.uni-muenster.de
Helmut Harder harder@earth.uni-muenster.de
Dawid Jagiela lightnir@gmail.com
Ferdinand Jamitzky jamitzky@lrz.de
Sarah Jones sarah.jones@kit.edu
Norbert Kalthoff norbert.kalthoff@imk.fzk.de
Martin Krause krause@mpe.mpg.de
Sven Krüger krueger@theochem.tu-muenchen.de
Adam Liwo adam@chem.univ.gda.pl
Michael Meier michael.meier@rrze.uni-erlangen.de
Stanislaw Ołdziej stan@scheraga2.chem.cornell.edu
Morris Riedel m.riedel@fz-juelich.de
Ulrich Rist rist@iag.uni-stuttgart.de
Notker Rösch roesch@theochem.tu-muenchen.de
Helmut Satzger satzger@lrz.de
Marc Schartmann schartmann@mpe.mpg.de
Harold A. Scheraga has5@cornell.edu
Steffen J. Schmidt schmidt@flm.mw.tu-muenchen.de
Günter H. Schnerr schnerr@flm.mw.tu-muenchen.de
Bernd Schuller b.schuller@fz-juelich.de
Juliane Schwendike juliane.schwendike@kit.edu
Björn Selent selent@iag.uni-stuttgart.de
Achim Streit a.streit@fz-juelich.de
Matthias Thalhamer thalhamer@flm.mw.tu-muenchen.de
Jan Treibig jan.treibig@rrze.uni-erlangen.de
Tobias Trümper tobit@earth.uni-muenster.de
Gerhard Wellein gerhard.wellein@rrze.uni-erlangen.de

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