inSiDE

inSiDE • Vol. 5 No. 1 • Spring 2007

Innovatives Supercomputing in Deutschland

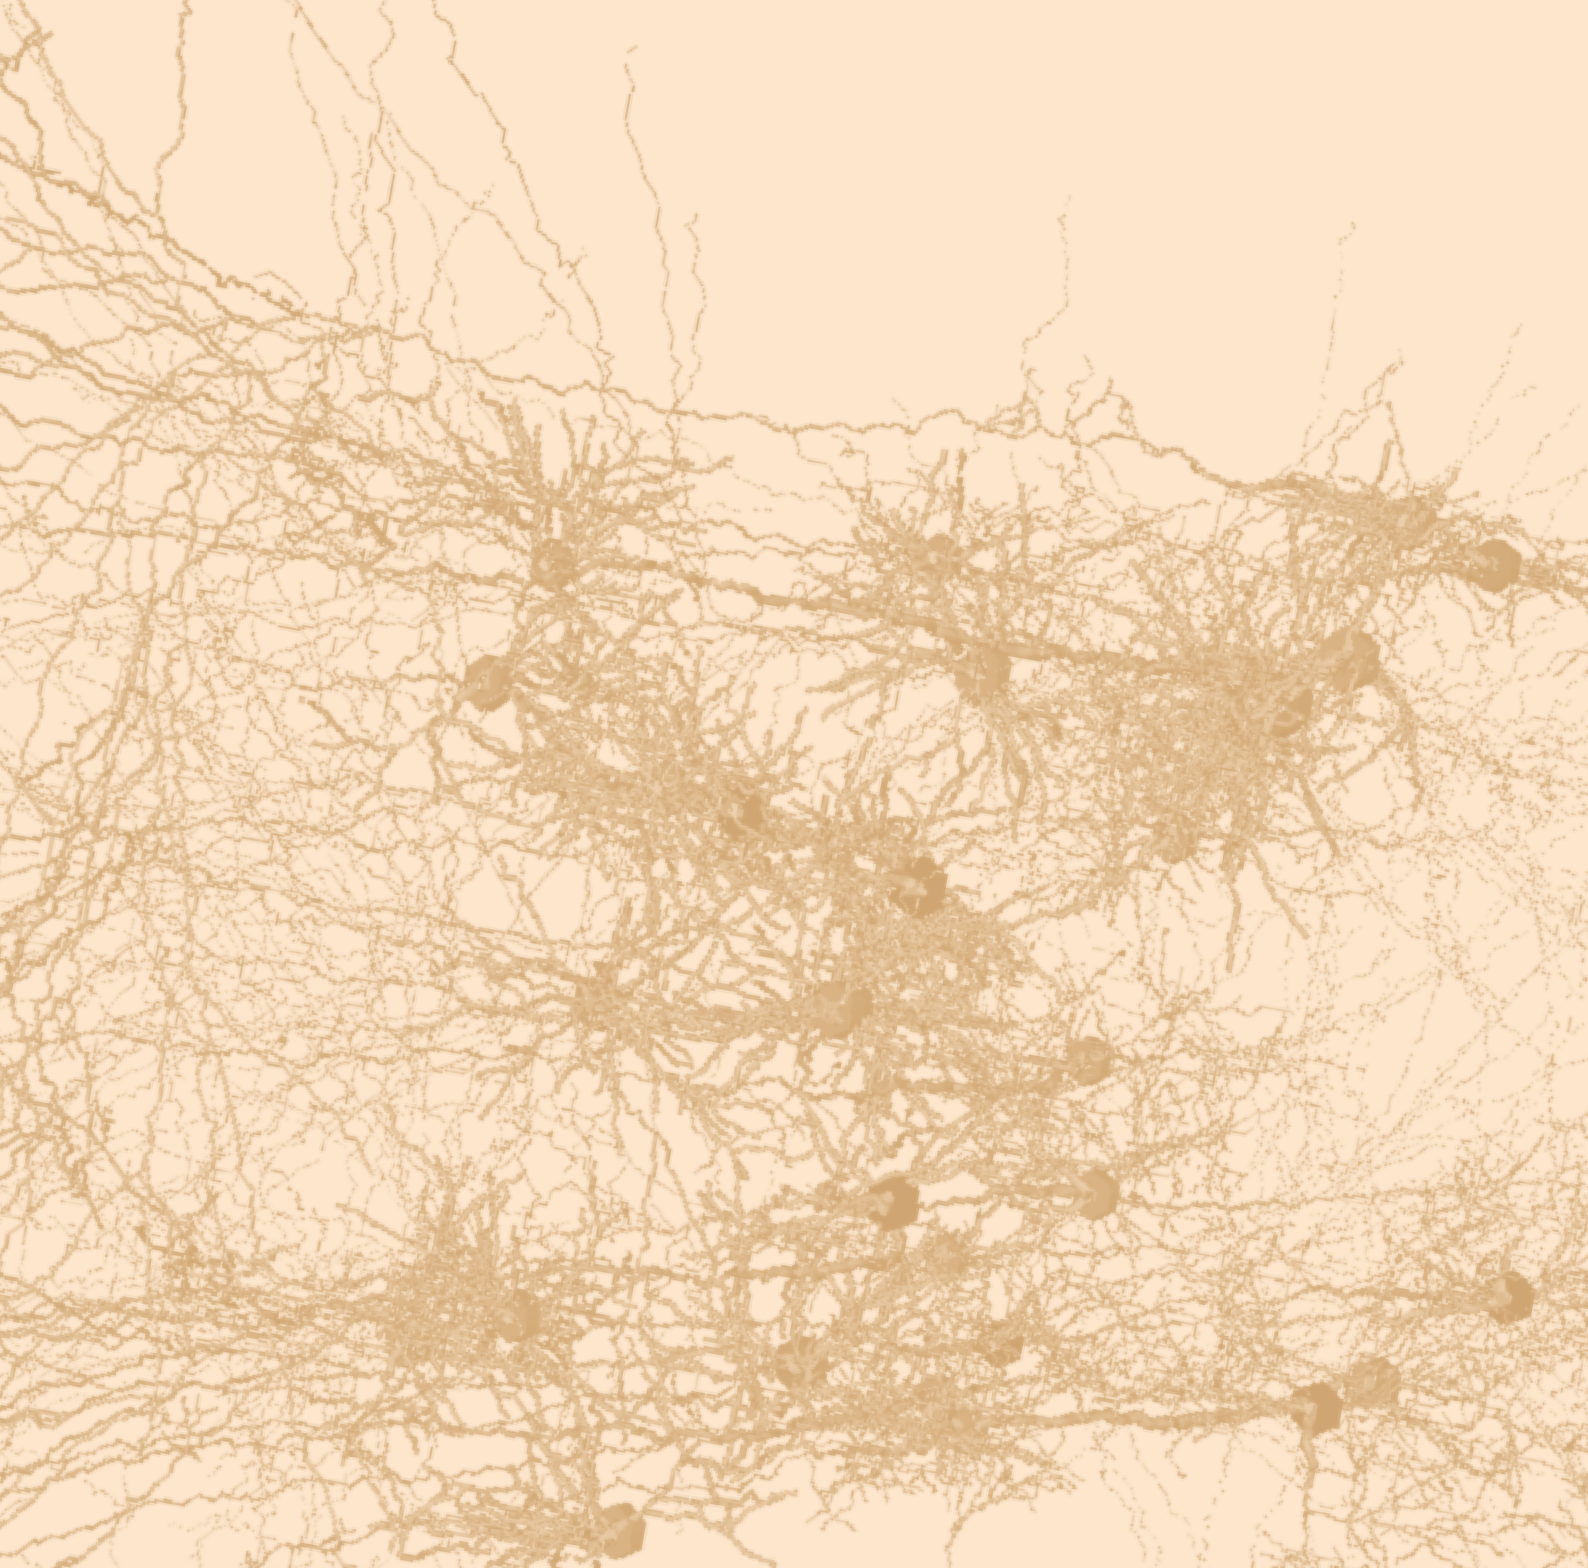

Editorial

German supercomputing has seen interesting and important changes since the last issue of InSide was published. First, the three German national supercomputing centres joined up to form the Gauss Centre for Supercomputing (GCS) with combined computing power of more than 120 TFLOP/s located in Stuttgart, Garching, and Jülich. You will find a report on the Centre in this issue. Behind these activities is the wish to bundle German HPC resources in an optimum way to provide users with the highest level of performance and service. At the same time the Gauss Centre for Supercomputing is going to represent Germany in the European HPC discussion. So, next, the Gauss Centre for Supercomputing and Partners from 14 European countries founded the consortium PACE – Partnership for Advanced Computing in Europe.

But there have also been big steps in hardware development in Germany. LRZ has installed its second phase upgrade with a peak performance of now 62 TFLOP/s, a description of which can be found in this issue. NIC only recently announced plans for its next phase installation due in autumn this year. More details will be given in the next issue. With these new hardware installations Germany continues its approach of a pyramid of performance at the national level co-ordinating its national supercomputing centres to create the best and most balanced system environment for German and European science.

In a second section a number of application papers are presented. The focus in this issue is on computational fluid dynamics. Three contributions describe problems of turbulence. With the new level of performance provided by the Gauss Centre for Supercomputing users are able to tackle problems that have been beyond their reach only 3 years ago. The contributions nicely show how increase performance and memory size can help to achieve major breakthroughs which can also contribute to the solution of flow problems in industrial environments.

A third section is devoted to projects. The ParMA project (Parallel Programming for Multi-core Architectures) aims at fully exploiting the power of multithreading on multi-core architectures for conventional HPC applications, but also for embedded applications on Multi-Processor System-on-a-Chip (MPSoC) architectures. The EU-funded project CoSpaces is aimed at developing organizational models and distributed technologies for collaborative workspaces for individuals and project teams within distributed virtual manufacturing enterprises enabling effective partnerships, innovation, improved productivity, and reduced design cycles.

Increasingly middleware and tools become important for the HPC community and German HPC centres have very early focused on research in this field. In this issue you will find a description of a Fortran binding for the GNU scientific library, a much more extensive discussion of which will be given in the August 2007 issue of the ACM SIGPLAN Fortran Forum. The UNICORE Grid system provides a seamless, secure and intuitive access to distributed Grid resources such as supercomputers, clusters, and large server farms. The most recent development of UNICORE 6 is presented here.

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

• Prof. Dr. H.-G. Hegering (LRZ)

• Prof. Dr. Dr. Th. Lippert (NIC)

• Prof. Dr.-Ing. M. M. Resch (HLRS)

Contents

Gauss and PACE: Combining Resources establishes powerful Supercomputing Organizations

Providing computing resources at the highest performance level for computer-based science and engineering in Germany and Europe: this is the common denominator behind current decisions made by German and European supercomputing facilities. First, the three German national supercomputing centres joined up to form the Gauss Centre for Supercomputing (GCS) with combined computing power of more than 120 TFLOP/s located in Stuttgart, Garching, and Jülich. Next, the Gauss Centre and Partners from 14 European countries founded the consortium PACE – Partnership for Advanced Computing in Europe.

The top computers of the Gauss Centre have a total computing power of 120 TFLOP/s at their disposal. This computing power is distributed among the partners' sites:

PACE representatives of the 15 participating countries together with the Federal Minister Dr. Annette Schavan

the High-Performance Computing Centre Stuttgart (HLRS), the Leibniz Computing Centre (LRZ) in Garching near Munich, and the John von Neumann Institute for Computing (NIC) at Research Centre Jülich. It is planned to increase the overall performance of the Gauss Centre to a value larger than 1000 TFLOP/s (1 PetaFlop/s) within the next years.

The Gauss Centre's members, who signed an agreement to found a registered association (GCS e. V.) on April 13th, will now follow a common direction in this organization. The procurement of hardware will be more closely coordinated, applications for computing time will be scientifically evaluated on a common basis, and software projects will be jointly developed. Another key area will be training. The work of specialist researchers will be supported and promoted by harmonising the services and organizing joint schools, workshops, and conferences on simulation techniques. Methodologically oriented user support is also a major concern of the Gauss Centre.

The Federal Ministry of Education and Research (BMBF), the Ministry of Innovation, Science, Research and Technology of the State of North Rhine-Westphalia, the Bavarian State Ministry of Science, Research and the Arts, and the Ministry of Science, Research and the Arts Baden-Württemberg unreservedly support the Gauss Centre for Supercomputing which is the largest national association for high-performance computing in Europe. The high-speed computer network and the scientific cooperation at the three sites are being funded by the BMBF in order to ensure that this leading international position will be maintained in future by means of optimized structures and organization. The three sites are thus making themselves visible throughout Europe and play a central role in the establishment of a European high-performance computer network. For further details see http://www.gauss-centre.eu/

Also, European scientists and engineers will be able to turn to a new resource when it comes to supercomputing. Top representatives from research institutions in 15 countries have recently laid the foundation for a leading international supercomputer infrastructure and created a European supercomputing network: PACE – Partnership for Advanced Computing in Europe. With a combined effort, Europe should stay at the top of the international competition.

The central idea behind the new partnership is a joint network of supercomputer resources with different locations, linked together by the most modern network technology. The costs were estimated by the European Strategy Forum on Research Infrastructures (ESFRI) in fall 2006 in the order of several hundred millions of Euros. They have to be covered to a large extent by national sources. The rest will be provided by the European Union through the 7th Research Framework Programme.

PACE aims at a European supercomputer eco-system that can be described in terms of a performance pyramid. At the top are a small number of leadership-class supercomputing systems, funded through national sources, with additional European funding. The middle layer of the pyramid consists of a number of national and regional supercomputers. These still should be powerful supercomputers being able to run the load below PetaFlop/s level. The bottom of the pyramid consists of local compute servers that should enable the development of a strong competence base of computational scientists.

The PACE consortium, founded formally in Berlin on April 17th by signing a corresponding Memorandum of Understanding, is made up of France, Germany, Netherlands, Spain, United Kingdom, Austria, Finland, Greece, Italy, Norway, Poland, Portugal, Sweden, Switzerland, and Turkey.

Together the members of PACE will strengthen European science, engineering and supercomputer technologies and thus secure Europe a pioneering role in the global competition.

In the coming preparatory phase of two and a half years PACE will test prototypes of Petaflop Machines and make proposals on how the funds should be efficiently deployed. The conceptual design report is the first step towards a globally competitive organizational structure for scientific computing in Europe. The principle behind this goal is using the equipment and expertise of the PACE partners, not in competition among themselves, but rather as complementing each other.

• Achim Bachem[1]
• Heinz-Gerd Hegering[2]
• Thomas Lippert[3]
• Michael Resch[4]

[1] Forschungszentrum Jülich (FZJ)
[2] Leibniz Rechenzentrum (LRZ)
[3] John von Neumann Institut für Computing (NIC)
[4] Höchstleistungsrechenzentrum Stuttgart (HLRS)

Upgrading the SGI Altix 4700 at LRZ

Migrating to Dual-core Processors

After less than a year of operation of the HLRB II, the second phase of the SGI Altix 4700 installation at LRZ has been initiated: On April 17th, 2007, the upgraded system went into user operation for the first time after an interruption of three weeks. With a peak performance of 62,3 and a LINPACK performance of 56,5 TFLOP/s, the machine is now the fastest system in Germany and one of the most powerful in Europe. This increase of capability was achieved by replacing the 4,096 Madison processors by 4,864 dual-core Montecito sockets; the increase in the number of sockets required that a fraction of the blades be equipped with two sockets ("high-density blades"). The number of memory channels as well as the NUMAlink interconnect remain essentially unchanged from phase 1.

The total memory capacity of the system has been increased from 17 to nearly 39 TBytes, and an additional 300 TBytes parallel file system has been added, for a total of 600 TBytes of scratch space.

The partitioning of the system has also been changed: It now contains 19 partitions, each with 512 cores in a single system image and 2 TBytes RAM; six of these partitions contain the high-density blades with 4 cores per blade (each blade being equipped with one memory channel and two NUMAlink connections), the remaining 13 partitions have 2 cores per blade. Hence it is possible to schedule jobs depending on their need for memory bandwidth versus computational density.

Although the memory balance of the system has worsened due to two or more cores sharing one memory channel, the aggregate benchmark performance of the machine, as defined by the benchmark suite used for the procurement process has nearly doubled to a value of 16,2 TFlop/s. First statistics taken during testing phases of part of the machine indicate that also for most user codes the per-core performance from phase 1 is preserved; a significant drop in per-core performance has up to now only been observed for the area of computational fluid dynamics, known for its comparatively low computational density. But even for these codes, 10 - 15% more than half the per-core performance is observed. The main reasons for the better-than-expected performance characteristics of the upgraded systems appear to be:

Manual Labor: preparing the SGI Altix 4700 blades for the processor upgrade in the LRZ computer room

A collection of Madison processors from phase 1

1. The Montecito processors have 9 MB level 3 cache per core, hence the size of the L3 cache for the total system has increased by a factor of more than 3. A larger percentage of applications will achieve a higher cache hit rate, and hence obtain an over-proportional performance increase.

2. The split of Montecito's L2 cache into separate instruction and data caches helps applications with high instruction density.

3. Even with tuning measures, a single Madison core does not fully exploit the memory channel; hence even a memory bound application can extract slightly more performance from a Montecito socket as compared to a Madison socket.

From the user's point of view, existing programs can be run unchanged: The machine maintains full binary compatibility even though optimal performance may require additional tuning measures. In particular, an increased use of the hybrid parallel programming model, for example using 2 threads per MPI task, may yield improved performance in many cases. During the next few months, it will be the job of HPC support staff at LRZ to help users improving the performance of their codes; a sustained performance of 5-6 TFlop/s can be expected for day-to-day operations including unoptimized programs and unused processors.

Processor cores	9,728
Peak performance	62,3 Teraflop/s
Weighted benchmark performance	16,2 Teraflop/s
Memory capacity	39 Terabytes
Memory bandwidth	34,8 Terabyte/s
Total disk space	680 Terabytes
Aggregate I/O bandwidth	45 Gigabyte/s
Latency of interconnect	1-4 microseconds
MPI bisection bandwidth	170 Gigabyte/s

SGI Altix 4700 phase 2 performance parameters

With the upgraded machine, LRZ is now well prepared to tackle the large simulations planned by the scientific communities for the next few years. While it is a rare occurrence, the integrated NUMAlink interconnect allows to use the complete system with a single program, and hence to investigate the characteristics of applications near the 10,000 core limit, opening up the route to Petaflop computing.

• Reinhold Bader

Leibniz Computing Centre

Smart Suction - An advanced Concept for laminar Flow Control of three-dimensional Boundary Layers

The list of reasons for a sustained reduction of commercial-aircraft fuel consumption is getting longer every day: significant environmental impacts of the strongly-growing world-wide air traffic, planned taxes on kerosene and emission of greenhouse gases, and the lasting rise in crude-oil prices. As fuel consumption during cruise is mainly determined by viscous drag its reduction offers the greatest potential for fuel savings. One promising candidate to reduce viscous drag of a commercial aircraft is laminar flow control (LFC) by boundary-layer suction on the wings, tailplanes, and nacelles with a fuel saving potential of 16%. (The other candidate is management of turbulent flow, e.g., on the fuselage of the aircraft, by a kind of shark-skin surface structure that however has a much lower saving potential.) Suction has been known for decades to delay the onset of the drag-increasing turbulent state of the boundary layer by significantly enhancing its laminar stability and thus pushing laminar-turbulent transition downstream. However, in case of swept aerodynamic surfaces, boundary-layer suction is not as straightforward and efficient as desired due to a crosswise flow component inherent in the three-dimensional boundary layer.

Suction aims here primarily at reducing this crossflow, and not, as on unswept wings, at influencing the wall-normal distribution of the streamwise flow component. The crossflow causes a

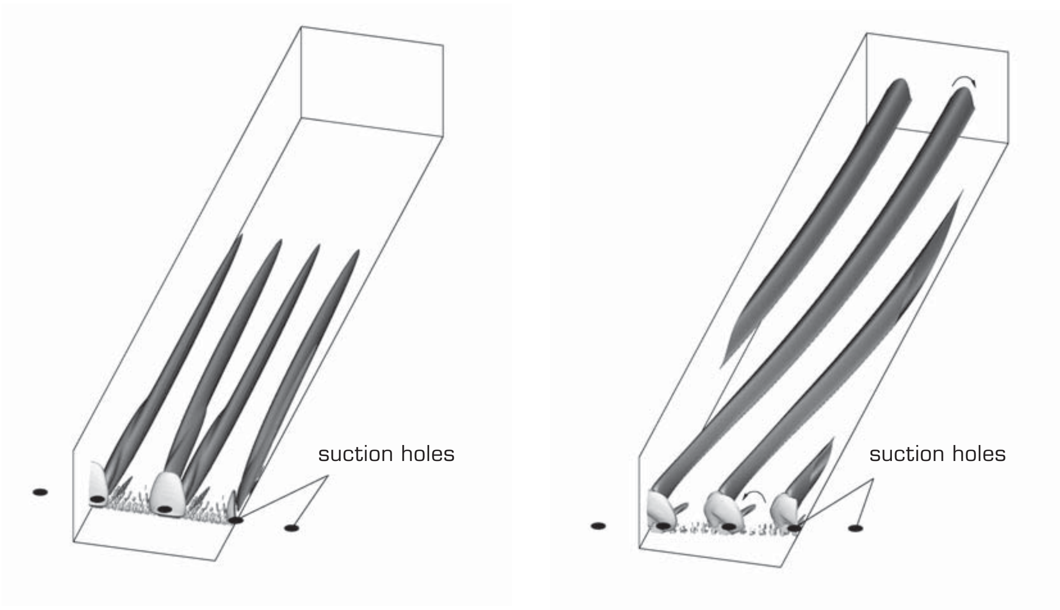

Figure 1: Visualization of vortices emanating from a single suction hole row on an unswept (left) and swept wedge (right). Arrows indicate the sense of vortex rotation. On the unswept wedge vortices are damped downstream, on the swept wegde corotating vortices are amplified due to crossflow instability

primary instability of the boundary layer with respect to three-dimensional disturbances. They can grow exponentially in downstream direction, depending on their spanwise wave number, and lead to co-rotating longitudinal vortices, called crossflow vortices (see figure 1), in the front part of, e.g., the wings. Now, on swept wings with their metallic or carbon-fibre-ceramic skins, the discrete suction through groups of micro-holes or -slots with diameters of typically 50 micrometers can excite unwanted, nocent crossflow vortices.

The grown, typically steady vortices deform the laminar boundary layer and can cause its breakdown to the turbulent state by triggering a high-frequency secondary instability, occuring now already in the front part of the wing. The onset of such an instability highly depends on the state of the crossflow vortices. A determinant parameter is the spanwise spacing of the vortices, influencing also their strength. The spacing of naturally grown vortices corresponds to the spanwise wavelength of the most amplified eigenmode disturbance of the base flow. Even in most cases of discrete suction through groups of holes or slots with relatively small spanwise and streamwise spacings such a vortex spacing appears on a suction panel as the strong growth of the most amplified disturbance always prevails. This explains why transition can also set in when "stabilizing" boundary-layer suction is applied. If the suction were perfectly homogeneous over the wall, the suction itself would not excite nocent modes, but any surface imperfections like dirt, insects and so on would.

Recently, a new strategy for laminar flow control has been proposed and experimentally [1] and numerically demonstrated [5]. At a single chordwise location artifical roughness elements are laterally placed and ordered such that they excite relatively closely-spaced, benign crossflow vortices that suppress the nocent ones by a nonlinear mechanism and do not trigger fast secondary instability. If the streamwise variation of flow conditions and stability characteristics is weak this approach has proven to impressively delay transition. A better understanding of the phyiscal background of the effectiveness of this approach has been provided by direct numerical simulations [5], who coined the term Upstream Flow Deformation (UFD).

A major shortcoming of UFD with its single excitation of benign vortices is that it works persistently only for flows with non- or weakly varying stability properties. This typically is not the case on swept wings or tail planes where the boundary-layer flow undergoes a varying acceleration. Hence, an approach combining classical suction with the Upstream-Flow-Deformation method has been proposed by the authors of this article at the Institut für Aerodynamik und Gasdynamik (IAG) at the University of Stuttgart. The research is performed by direct numerical simulations (DNS) of the problem using a computer code developed at IAG that runs on the powerful parallel vector NEC super computers, currently the SX-8, installed at the High Performance Computing Centre Stuttgart (HLRS) (for computational details see [2][3]). Its optimization has been supported by the local NEC group.

Applications

The main idea is to excite benign, closely-spaced (UFD-)vortices and to maintain them on a beneficial amplitude level by an appropriate order of suction orifices. The streamwise variation of flow conditions and stability characteristics can be taken into account by adapting the spacing of the suction orifices continuously or in discrete steps. In this way we overcome the shortcomings of the single excitation of UFD-vortices. However we note that this is not at all a trivial task because it is not clear a priori which direction the vortices follow - the flow direction depends on the wall-normal distance - and improper excitation can lead to destructive nonlinear interaction with benign vortices from upstream, or nocent vortices. For illustration of a case where the adaptation to the chordwise varying flow properties has been done in an improper way see figure 2.

If properly designed (see figure 3) the proposed method unifies the stabilizing effects of bounday-layer suction and UFD. Consequently, the new method strives for (i) securing the working of suction on swept surfaces, and (ii) an additional stabilization of the boundary-layer flow compared to classical suction alone, or, alternatively, it allows to reduce the suction rate for the same degree of stabilization. By the excitation of selected crossflow modes being exponentially amplified and finally forming crossflow vortices not triggering turbulence, the stability of the flow is enhanced as would the suction rate of a conventional suction system have been risen. The reason is that the vortices generate by nonlinear mechanisms a mean-flow distortion not unlike suction, cf. [5], influencing the stability in an equally favourable manner as suction itself. The new method is termed *smart suction* as the instability of the laminar flow is exploited to enhance stability rather than increasing the suction rate. We hope that it pushes the suction technique that has been prepared by the airliner manufacturers but not yet applied in daily operation to real operational success. Fuel and involved exhaust gas savings of about 16% can be expected using properly working suction on the wings, the tailplanes and the nacelles (see also [4]). As research must not rest on ones's laurels simulations of larger parts of the wings and investigations of exciting and maintaining benign vortices by roughness, bumps, dips, or localized suction/blowing actuators are planned. Moreover, the applicability of the technique to LFC on wind turbine rotors is scrutinized.

References

[1] Saric, W.S., Carrillo, Jr., Reibert, M.S.
Leading-Edge Roughness as a Transition Control Mechanism, (1998a), AIAA-Paper, 98-0781

[2] Messing, R., Kloker, M.
DNS study of spatial discrete suction for Laminar Flow Control, (2005), in: High Performance Computing in Science and Engineering 2004 (ed. Krause, E., Jäger, W., Resch, M.), 163-175, Springer

[3] Marxen, O., Rist, U.
DNS of non-linear transitional stages in an experimentally investigated laminar separation bubble, (2006), in: High Performance Computing in Science and Engineering 2005 (ed. Nagel, W. E. Jäger, W., Resch, M.), 103-117, Springer

[4] Schrauf, G.
Status and perspectives of laminar flow, (2005), The Aeronautical Journal (RAeS), 109, no. 1102, 639-644

[5] Wassermann, P. & Kloker, M.
Mechanisms and passive control of crossflow-vortex induced transition in a three-dimansional boundary layer, (2002), Journal of Fluid Mechanics, 456, 49-84

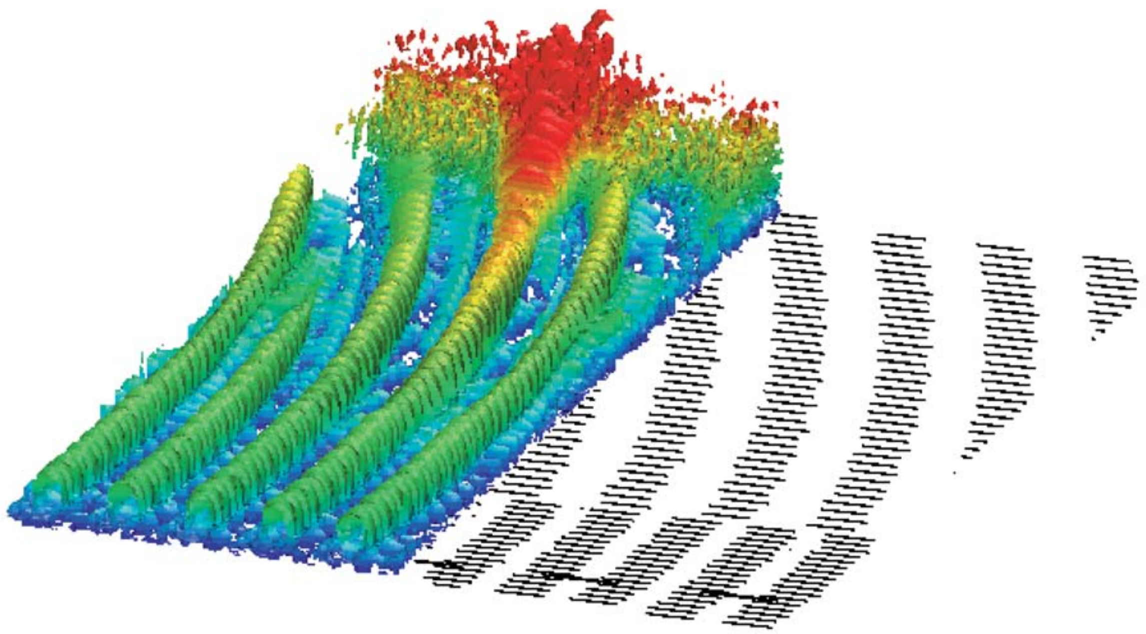

Figure 2: Visualization of vortices (left) and suction orifices/slots at the wall (right) in perspective view of the wing surface for smart suction with improper adaptation of suction orifices. Flow from bottom to top, breakdown to turbulence at the end of domain

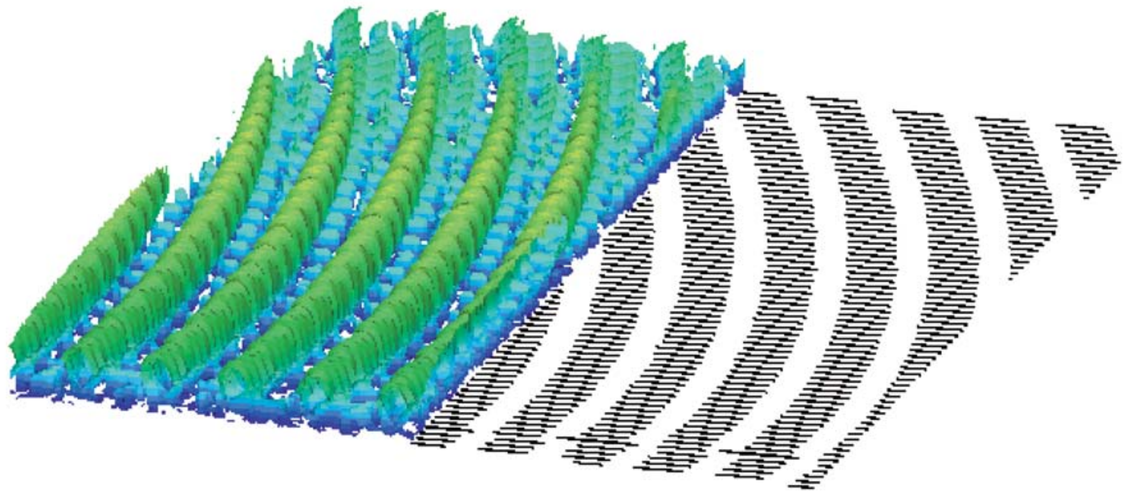

Figure 3: As for figure 2 but with proper adaptation of suction orifices for a simple case

• Ralf Messing
• Markus Kloker

Institut für Aerodynamik und Gasdynamik, Universität Stuttgart

Applications

Getting Hold of Turbulence with adaptive Methods

The common notion of turbulence is that of fluid motion changing randomly in space and time. From the theoretical point of view, randomness is caused by non-linear interactions between eddies of different size as described by the fluid dynamical equations. In the simplest picture, a cascade of turbulent eddies emerges which appears homogeneous and isotropic in a statistical sense. This is the essence of Kolmogorov's account of turbulence in 1941. However, the picture of a turbulent cascade does not fully embrace the nature of turbulence. For instance, when riding a plane, passengers might experience shaky flight stages in between calm phases. The shaking is caused by intense turbulent air motion. So, from the view point of a particular observer, turbulence can vary substantially. This property is known as intermittency.

Computationally, turbulence has frequently been treated by a method called large eddy simulation (LES) that is based on the notion of a cascade of turbulent eddies. In LES, a Grid is built up from cubic cells of given size and the dynamical equations determining the mass density, the velocity and the energy of the fluid are solved for each cell. Repeating this procedure over many time steps, a discrete numerical representation of the flow is obtained. It is an essential restriction that turbulent eddies smaller than single cells are not captured by the computation. Given the capabilities of present-day supercomputers, Grids used in LES might comprise as many as several billion cells corresponding to a ratio of about one thousand between the side length of the whole Grid and an individual cell. Hence, the different length scales, i.e. eddies of different size, span three orders of magnitude. Notwithstanding these huge numbers, a simple estimate based on the Kolmogorov turbulence theory shows that turbulent eddies in typical terrestrial or astrophysical flows vary by up to ten orders of magnitude in size. For this reason, so-called subgrid scale (SGS) models are applied in LES. These models approximate the influence of eddies smaller than the Grid cells onto the dynamics of larger eddies that is explicitly computed.

There is a long record of LES in engineering, atmospheric sciences and astrophysics. For example, thermonuclear supernova explosions have been simulated successfully by means of LES. Nevertheless, some problems in contemporary astrophysics appear to be beyond the reach of LES. A prominent example is turbulence in the interstellar medium. Other than in most terrestrial instances, the gas in the interstellar medium is in a state of supersonic turbulent motion, i.e. the flow velocity is greater than the speed of sound. Thus, shock fronts occur in addition to turbulent eddies. Moreover, regions of enhanced gas density can collapse under their own gravity. The interplay between supersonic turbulence and gravity is believed to drive the formation of stars (Mac Low and Klessen, 2004). Altogether, turbulence in the interstellar medium is highly intermittent and deviates substantially from homogeneous turbulence as envisaged by Kolmogorov. In numerical simulations, the dynamics of supersonic turbulent flow at relatively large scales as well as the gravitational collapse of highly compressed structures at typically much smaller length scales has to be tackled. With conventional LES, this task is infeasible. On the other hand, there is a method called smoothed particle hydrodynamics that works well for the treatment of self-gravitating gas, but turbulence tends to be elusive because the smoothing properties with regard to velocity fluctuations at small scales are poorly understood.

A route toward the solution to this problem was proposed by Kritsuk et al. (2006). Based on the intermittency of turbulence, they applied a numerical technique called adaptive mesh refinement (AMR). The basic idea is to work with a hierarchy of Grids dynamically adapting to the flow structure. The steep gradients of the velocity field in the vicinity of shock fronts, for instance, are covered by Grids resolving smaller scales while smoother portions of the flow are covered by a coarser Grid. In general, there are several levels of refinement, i.e. more than one step in resolution. Refined Grids might also be generated to track collapsing gas regions. Finding generally applicable criteria for refinement, however, is a highly non-trivial problem.

We made use of computational resources granted by the DEISA Extreme Computing Initiative (DECI) in order to perform highly resolved simulations of supersonic turbulence with the conventional large eddy approach and compared those with AMR runs featuring various refinement criteria. Depending on the size of the computational grid (192^3 to 768^3 cells), each simulation required between 16 and 126 CPUs of the SGI Altix supercomputer at SARA, Netherlands, and consumed up to 100,000 CPU-hrs. The total project budget was 300,000 CPU-hrs. We have not yet included self-gravity, although this is planned in upcoming applications. The basic idea was to trigger the generation of refined grids by monitoring properties of turbulence such as the vorticity and, in the case of supersonic flow, the rate of gas compression. The term vorticity is commonly used for the rotation of the velocity field in fluid dynamics. Roughly speaking, high vorticity indicates turbulent eddies.

In the case of supersonic turbulence, refinement must track the compression of gas by shock fronts as well as the formation of turbulent eddies. For this reason, we complemented the refinement based on vorticity with an additional mechanism that triggers the generation of refined Grids in regions of increasing compression. This kind of refinement was tested in simulations of supersonic turbulence, where mechanical energy was supplied by stirring and compressing the gas via random forces.

Applications

Figures 1 and 2 contrast visualizations prepared from a conventional simulation running on a static Grid and an AMR simulation, respectively. On the left hand side, isosurfaces of the vorticity are shown. Apart from thin vortex tubes, sheet-like structures stemming from shock fronts can be seen. These structures are very well reproduced by AMR. In addition, we made quantitative verifications, for example, by means of distribution functions describing the probabilities of finding certain absolute values of vorticity or other quantities. The compression and rarefaction of gas is reflected in the spatial variation of the mass density. This is illustrated by volume renderings of the mass density on the right hand side of Figure 1 and 2, respectively. Regions of rarefied gas appear bluish, whereas highly compressed gas is shown in red. The method of refinement outlined above fills in Grid patches of higher resolution if the gas compression tends to increase. We found an excellent match of the mass density distributions on the static and adaptive Grids.

The results we have obtained from the comparisons presented above corroborate the application of adaptive methods to turbulence. The postprocessing of the multi-TB data produced in the production runs with high resolution, however, has turned out to be a considerable challenge and is not yet finished. Since it is generally impossible to treat turbulence at all length scales by means of AMR only, the next step will be to combine AMR with the subgrid scale model formulated by Schmidt et al. (2006). This will link conventional AMR to LES and, therefore, we call the new method Fluid mEchanics with Adaptively Refined Large Eddy SimulationS (FEARLESS). The numerical study carried out in the framework of DECI laid the foundation. Several astrophysical applications are planned, once the complete methodology is implemented. In particular, the turbulent interstellar medium, the evolution of spiral galaxies with feedback from star formation and the dynamics of hot gas in galaxy clusters will be investigated in follow-up projects. We expect that FEARLESS will open new perspectives in astrophysics by the as yet unequalled level of sophistication in the treatment of turbulence. The calculations will be continued on the upgraded SGI Altix of Leibniz Computing Centre.

References

[1] Kritsuk, A. G., Norman, M. L., and Padoan, P. (2006)
Adaptive mesh refinement for supersonic molecular cloud turbulence, Astrophysics Journal, Letters 638, L25–L28

[2] Mac Low, M.-M. and Klessen, R. S. (2004)
Control of star formation by supersonic turbulence.
Reviews of Modern Physics 76, 125–194

[3] Schmidt, W., Niemeyer, J. C., Hillebrandt, W. (2006).
A localized sub-Grid scale model for fluid dynamical simulations in astrophysics, Theory and numerical tests, Astronomic & Astrophysic, 450, 265–281 & 283–294

• Wolfram Schmidt

Lehrstuhl für Astronomie, Institut für Theoretische Physik und Astrophysik

Universität Würzburg

Applications

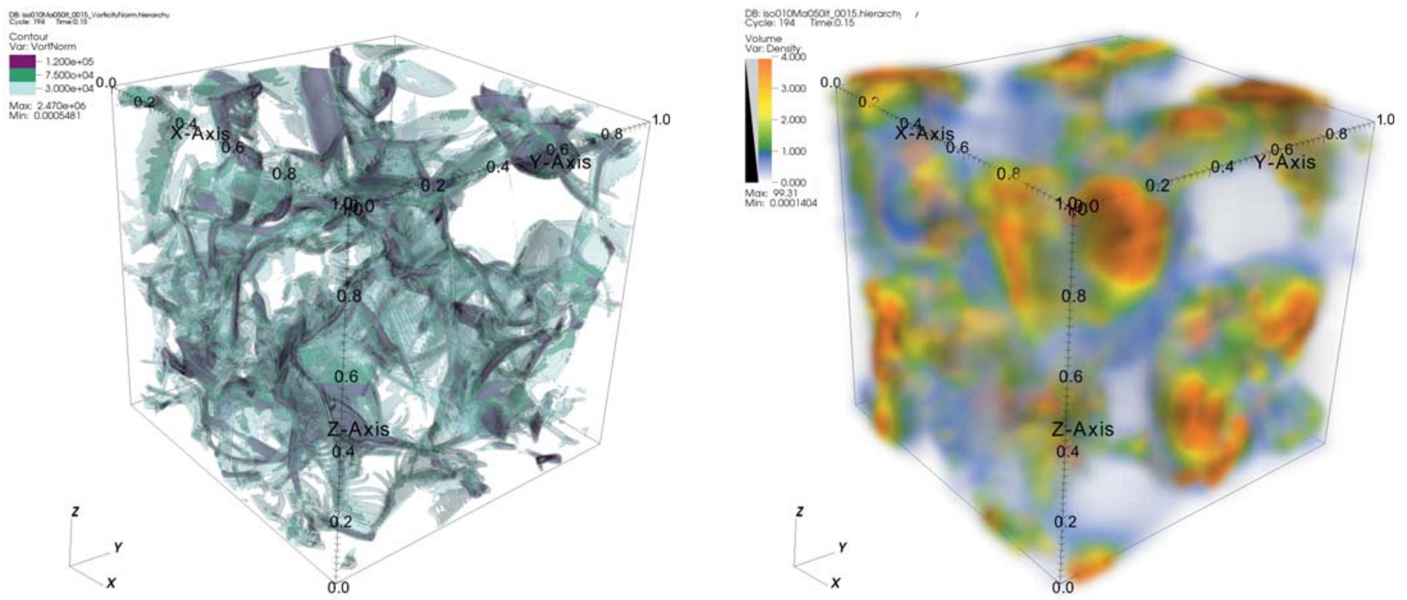

Figure 1: Isosurfaces of vorticity (left) and volume rendering of mass density (right) for a certain instant of time in a simulation of supersonic turbulence with static Grid

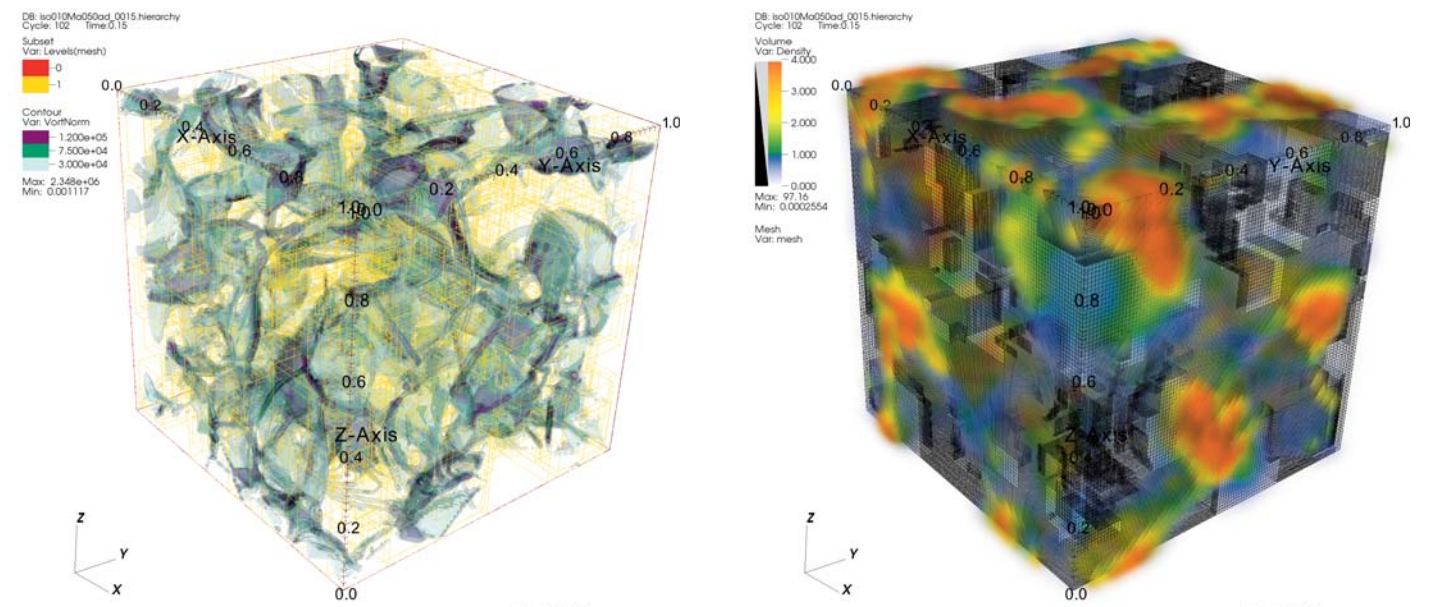

Figure 2: Same as in Figure 1 using AMR. The mesh is shown in different fashion on the left (yellowish boxes) and on the right (Grid patches with smaller cells)

Direct numerical Simulation of a Supersonic Jet and its acoustic Field

DEISA Extreme Computing Initiative

Aims of this Project

In this project, the method of direct numerical simulation will be used to compute a supersonic, three-dimensional and rectangular jet that is not perfectly expanded, as it is found at the nozzle exit of jet engines for aircraft. Numerical methods of high order of accuracy are chosen for the direct solution of the compressible Navier-Stokes equations, which gives us the possibility to compute the sound field generated by the supersonic jets directly. The aims of this study include improved modeling of sound generation by these jets. There is still lack of detailed knowledge about all the sound generation mechanisms involved in such flows, but such knowledge is needed for correct prediction of the acoustic field and for the development of control strategies to suppress it. International research makes considerable efforts to clarify questions related to sound generation by fluid flow. Co-operation with the Centre acoustique of Ecole Centrale de Lyon and LadHyX at Ecole Polytechnique, both in France, is intended on the basis of an existing DFG-CNRS project.

Physical Situation

Considering a supersonic jet, e.g. at the exit of a jet engine, in the over- or under-expanded case, a regular pattern of compression and expansion waves will be found within the supersonic part of the jet flow. A compression wave incident on the sonic line will be reflected as an expansion, and vice versa, cf. fig. 1. At the location of interaction between the compression wave and the turbulent mixing layer, acoustic waves are generated. This shock-induced noise also plays an important role in what is called jet screech. This phenomenon is caused by shock-induced acoustic waves traveling upstream and forcing the "young" shear-layer at the nozzle exit. At this point Kelvin-Helmholtz instabilities are growing to vortices, transported downstream and interacting with the shock tips which are emanating noise again and closing a feedback loop. Experimental results indicate particularly high sound pressure levels of up to 160 dB and even beyond. Prediction and reduction of shock-induced noise, produced by modern civil aircraft with jet propulsion, traveling at high subsonic or supersonic Mach numbers is a matter of particular interest. Besides pollution of the environment by the radiated sound, the latter can also lead to high dynamic loads on parts of the aircraft causing structural fatigue and even destroy them.

Methodology

The compressible Navier-Stokes equations are solved, based on a characteristic-type formulation on an orthogonal Grid stretched in both the stream-wise and the transverse directions. Along the span-wise direction, periodicity and statistical homogeneity were assumed. A low storage Runge-Kutta tree-step method is implemented for the integration in time. The computational Grid contains approx. 300 million Grid points with 2,040 points in the stream-wise direction, 1,020 points in the transverse direction and 144 points in the periodic direction. Based on the hight of the jet a Reynolds number of 30,000 can be assumed. To capture the sound generation and propagation processes, spatial discretization was done using a finite difference compact scheme of sixth order and a spectral like method in the periodic direction. A sponge region with a low-pass filter (applied spatially) and a Grid stretching in the stream-wise direction is applied to avoid large structures interacting with the outflow boundary. Additionally a boundary wave-acceleration is implemented to prevent acoustic reflections from the transverse boundaries. At the inflow the rectangular nozzle is modeled with a laminar tanh-profile. The resulting shear-layers are forced slightly downstream of the inflow boundary with a model-spectrum including the most unstable frequency of the shear-layer and its first sub-harmonic with a stochastical phase-shifting. The code is parallelized using the Massage Passing Interface (MPI). For the current setup 1,020 CPU's were used on a SGI Altix 4,700 with Itanium2 Madison 9M processors. Approx. 12 TB of data were written to disk and 0,5 TB of main memory were used.

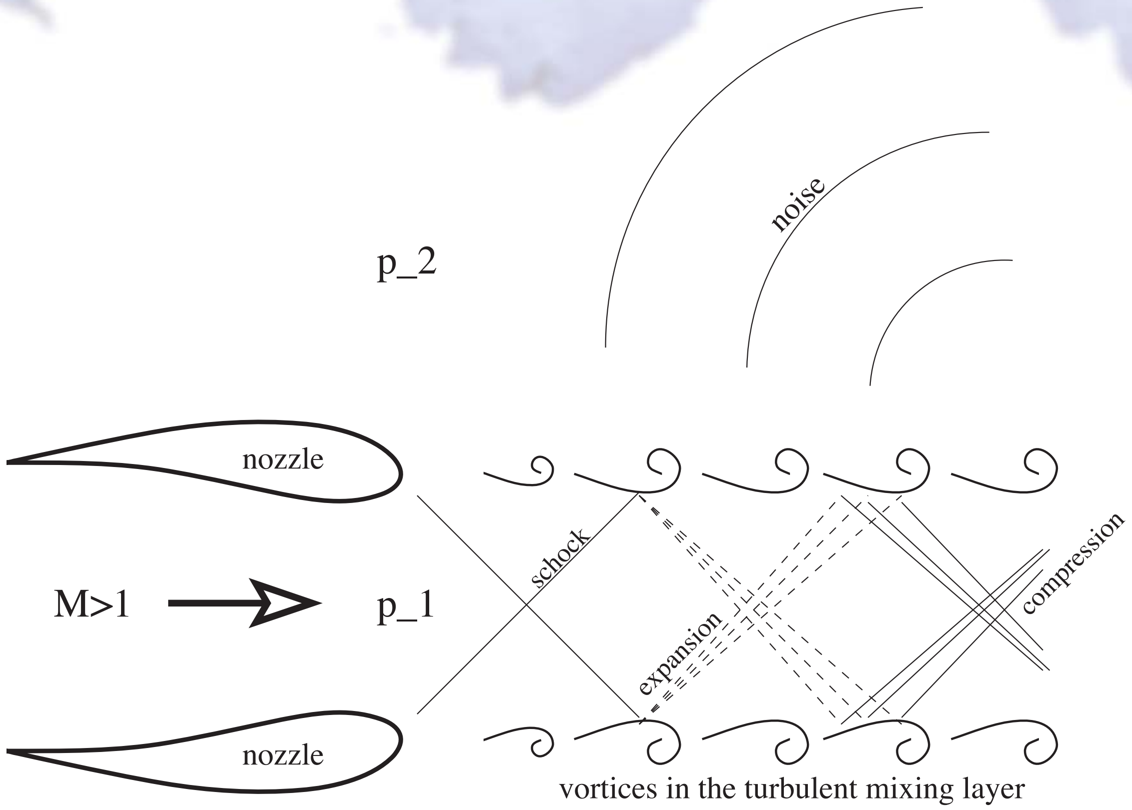

Figure 1: Schematic view of the interaction of shock and turbulent mixing layer in a jet with emanated noise. Over-expanded ⇒ $p_2 > p_1$

Applications

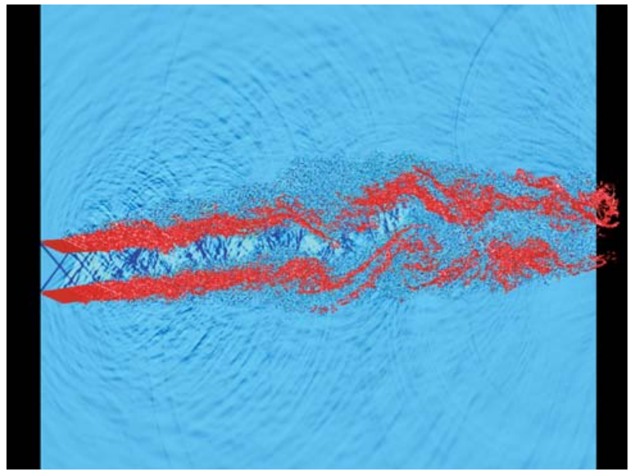

vorticity (iso-surface) and dilertation (plane)

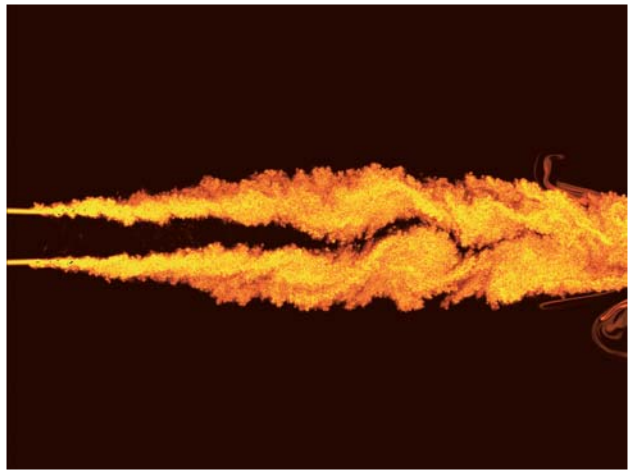

vorticity

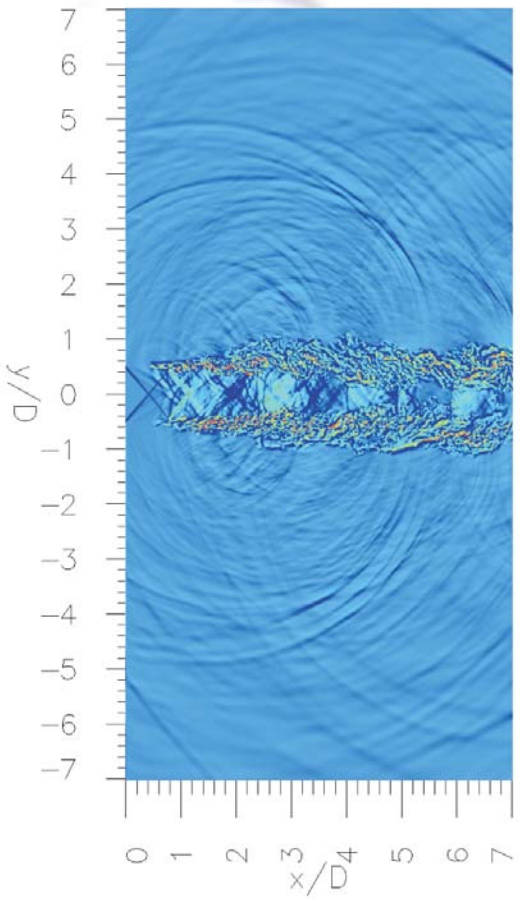

dilertation

Figure 2: Top left: 3d visualization of the jet with the vorticity as iso-surface and a plane of the dilertation. Bottom left: x-y-plane of the vorticity. Right: dilertation, D=hight of the jet

Results

Based on the high Reynolds number, the initially laminar shear-layers undergo an early transition to turbulence about 3 shear-layer thickness downstream of the forcing fringe. So the oblique shocks in the jet core interact with the turbulent mixing layers and emanate shock-induced noise. In figure (top-left) an iso-surface of the vorticity is shown in a three-dimensional sketch with a plane of the dilertation field. A dominant noise source seems to be close to the nozzle exit (between the first and second shock-cell). The early transition to turbulence and the initially laminar shear layers are clearly visible in fig. 2 (bottom-left) where a plane of the vorticity is presented. Since the jet is over-expanded (the pressure in the jet-core is below the ambient pressure; pressure ratio 0,9) the laminar shear-layers are slightly inclined towards the jet core. The last two plots in fig. 3 show the persisting shock-cell structure in the jet core with the visualization of the pressure (top) and the absolute velocity (bottom).

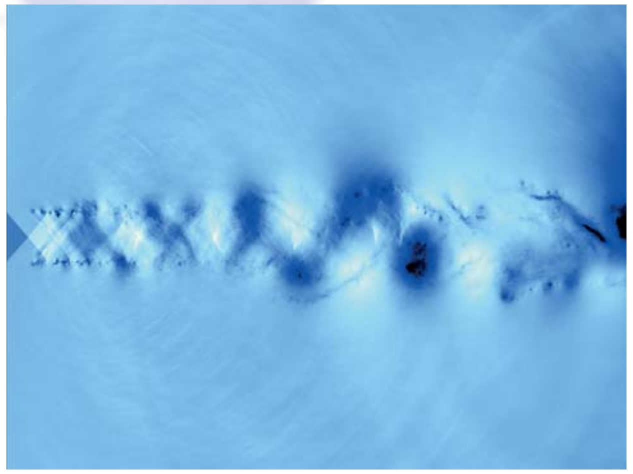

pressure field

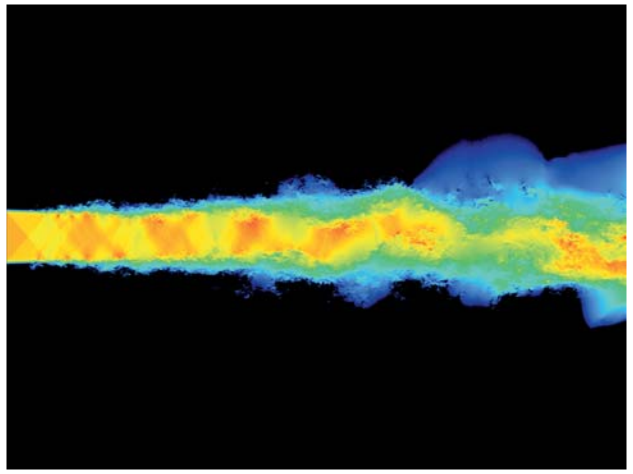

absolute velocity (|u|)

Figure 3: Top: x-y-plane of the pressure field, visualization of the shock-cells. Bottom: absolute velocity (|u|)

Acknowledgment

We thank the DEISA Consortium (co-funded by the EU, FP6 project 508830), for support within the DEISA Extreme Computing Initiative (www.deisa.org)

References

[1] Sesterhenn, J.
A characteristic-type formulation of the Navier-Stokes equations for high order upwind schemes, Computers & Fluids, Vol. 30 (1), 2001, pp. 37–67

• Jan Schulze
• Christoph Schaupp
• Jörn Sesterhenn

Fachgebiet Strömungsmechanik, TU München

Applications

Simulation of Blood Flow in a Ventricular Assist Device

Applications

Introduction

Diseases of the heart are a leading cause of death in the industrialized nations. The most reliable therapy for end-stage heart failure – heart replacement via a transplant – can be applied only in a fraction of the cases because of dramatic shortage of suitable donor hearts.

Since 1960's, attempts are being made to design a mechanical solution to heart failure; such a solution can take the form of a full replica of the heart – dual pumping chambers and complex valves – or, more commonly, of an assisting device, which pumps the blood from the existing failing ventricle into the aorta. The latter are referred to as Ventricular Assist Device, or VAD.

The Chair for Computational Analysis of Technical Systems (CATS) at the RWTH Aachen University, under the direction of Prof. Marek Behr, is specializing in CFD analysis and has been working on simulation of blood flow in VADs since 2000, with the latest analyses focusing the miniature MicroMed DeBakey VAD (see Figure 1). The main component of the DeBakey VAD is a spinning impeller propelling the fluid towards its destination and building up the required pressure head.

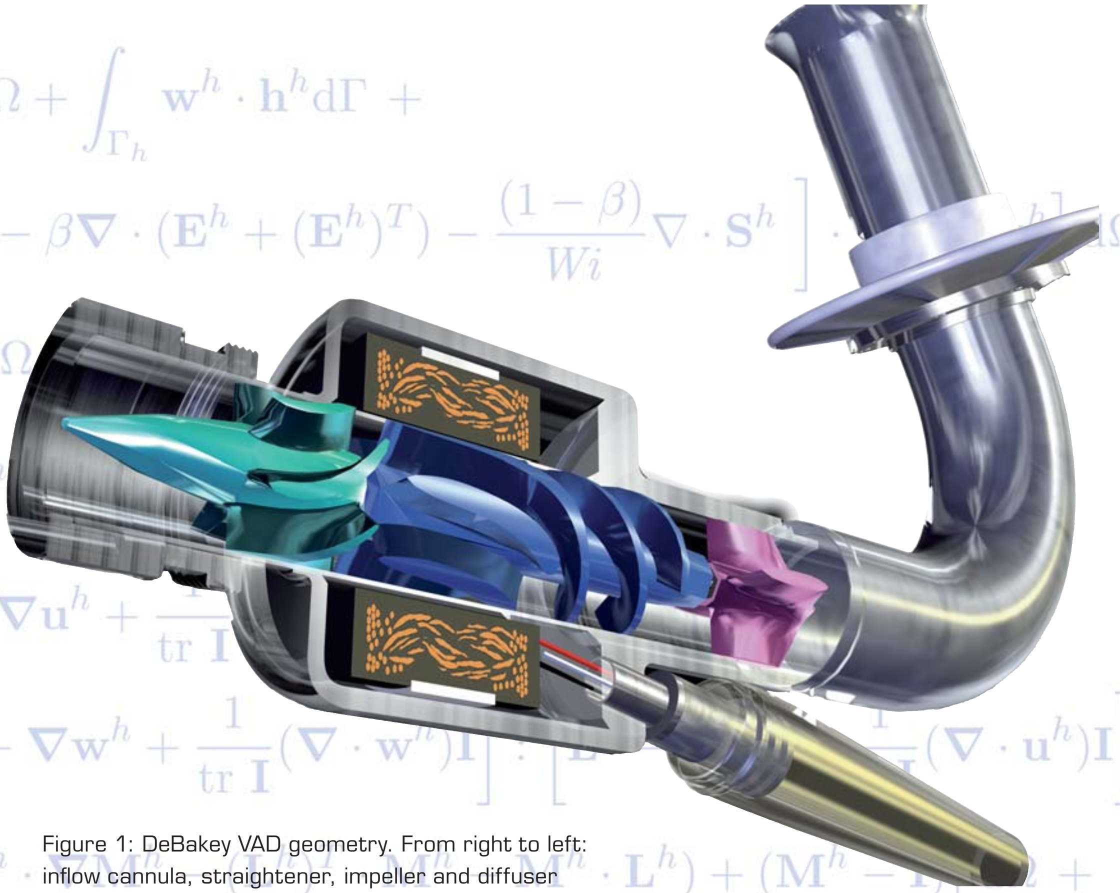

Figure 1: DeBakey VAD geometry. From right to left: inflow cannula, straightener, impeller and diffuser

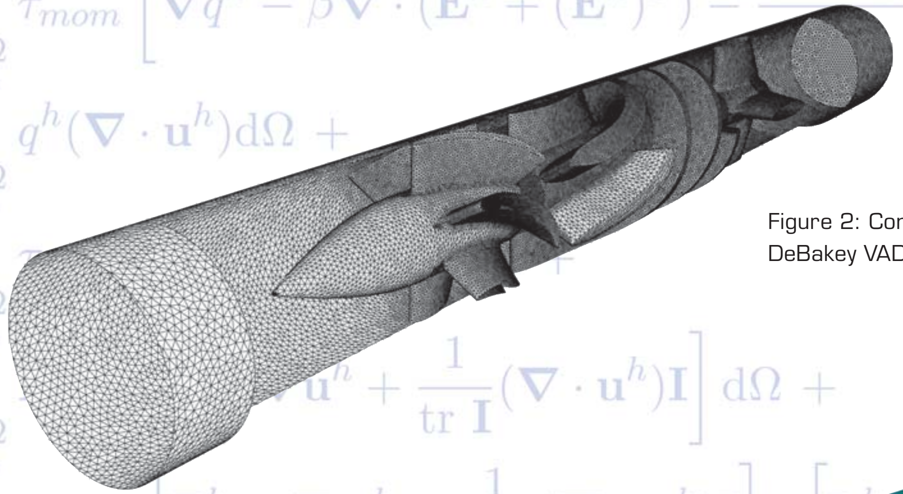

Figure 2: Computational Grid for the DeBakey VAD with 5 million elements

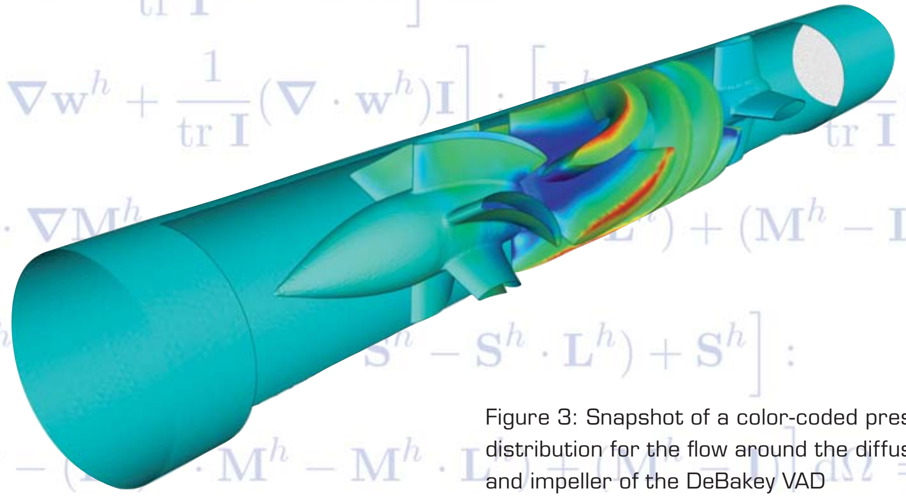

Figure 3: Snapshot of a color-coded pressure distribution for the flow around the diffuser and impeller of the DeBakey VAD

Applications

Simulation

CATS uses compute-intensive simulations to explore the potential of each design modification of the VAD, running a variety of flow profiles, flow rates, and impeller speeds to find the best way to improve the pump's biocompatibility.

Design challenges are staggering: the pumps need to be very small in order to be easily implantable, and they need to produce blood flow patterns that most closely resemble those in the body, in order to prevent hemolysis and thrombosis.

Hemolysis – the release of hemoglobin into the bloodstream – can result from damage to fragile red blood cells caused by prolonged elevated stresses imparted by the flow field. It is a potential danger to internal organs and can be life-threatening in extreme cases. In a VAD where the impeller spins at 10,000 rpm the shear stresses can be much higher than under physiological conditions.

Thrombosis – clotting of blood – can be caused by abrupt changes in the flow pattern and may lead to device malfunction or strokes.

These flow features can be predicted by simulations, but with a complex geometry such as DeBakey VAD, computational meshes in excess of 5 million computational cells are required for adequate accuracy. Thousands of discrete time intervals must be followed for simulating just a few revolutions of the impeller (see Figures 2 and 3).

Prof. Behr and his team perform computational analysis with XNS, an in-house computational fluid dynamics (CFD) code for simulations of unsteady fluid flows, including flows of micro structured liquids, in situations involving significant deformations of the computational domain. The code is based on finite element techniques using stabilized formulations, unstructured three-dimensional meshes and iterative solution strategies. Main and novel areas of XNS are: simulation of flows in the presence of rapidly translating or rotating boundaries using the shear-slip mesh update method (SSMUM); simulation of flows of micro structured (in particular viscoelastic) liquids, and simulation of free-surface flows, using a space-time discretization and staggered elevation-deformation-flow (EDF) approach.

Parallel Computing

The parallel implementation is based on message passing communication libraries, exploits graph-based mesh-partitioning techniques, and is portable across a wide range of computer architectures. The simulation for the DeBakey application is so complex that even on a large number of processors a full prediction of a quasi-stationary flow field in a DeBakey VAD may take many hours.

With simulations being repeated to proceed towards substantial design improvements, it becomes crucial to be able to exploit very large numbers of processors simultaneously, for example, the 16,384-processor IBM Blue Gene/L system operated by the Forschungszentrum Jülich.

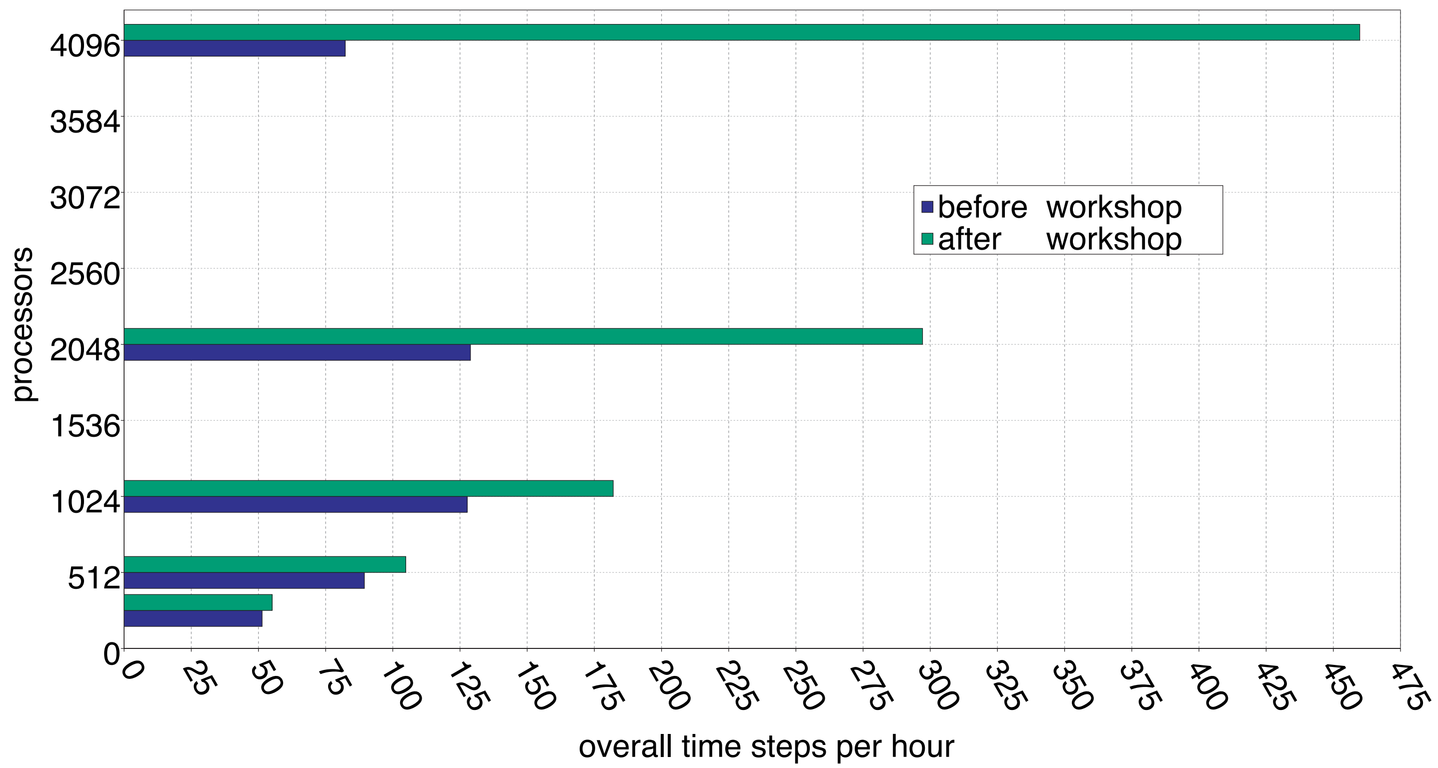

Figure 4: Overall simulation performance in time steps per hour before and after the scaling workshop

In December 2006, a scaling workshop for applications running on the 8-rack Blue Gene/L system in Jülich was organized and sponsored jointly by the John von Neumann Institute for Computing (NIC), IBM and the Blue Gene Consortium [1].

Prior to the workshop, one could observe an acceptable scaling of XNS up to 1,024 processors while there was no significant speed-up above that (see Figure 4). To find the bottlenecks in the code, the communication between the processes during the simulation runs was analyzed, both with XNS internal time measurements and the SCALASCA package [2]. After improving the communication patterns of XNS, the simulation performance could be improved remarkably up to 4,096 processes (see Figure 4). A good scaling is expected also for 8,192 processes; this is to be analyzed in future test runs.

CATS will continue its analysis of the DeBakey VAD with the objective of further improving the pump design and reducing its size, so that it could be used also for pediatric applications. The possibility to make efficient use of up to one fourth of the processors available on the Blue Gene/L at Forschungszentrum Jülich, one of Europe's most powerful computers is of great value here, because it allows generating the required data in a reasonable time.

Due to the many areas where XNS can be applied and its good scalability, we are confident that we could make efficient use of even bigger machines than the current Blue Gene/L.

References

[1] Frings, W., Hermanns, M.A., Mohr, B., Orth, B.
Report on the Jülich Blue Gene/L Scaling Workshop 2006
Technical Report FZJ-ZAM-IB-2007-02, February 2007

[2] Geimer, M., Wolf, F., Wylie, B.J.N., Mohr, B.
Scalable Parallel Trace-Based Performance Analysis
inSiDE Vol. 4, No. 2, 2006

• Mehdi Behbahani
• Mike Nicolai
• Markus Probst
• Marek Behr

Chair for Computational Analysis of Technical Systems (CATS)

RWTH Aachen University

The ParMA Project

Nowadays, programmers can no longer rely on performance improvements due to the processors' increasing clock-speeds. In contrary, all of today's common processor architectures employ multiple "slower" cores on one die. As a consequence, we are heading towards a parallel future, which is a challenging task to application developers and requires adequate programming methods and tools.

The ParMA project (Parallel Programming for Multi-core Architectures) aims at fully exploiting the power of multi-threading on multi-core architectures for conventional HPC applications, but also for embedded applications on Multi-Processor System-on-a-Chip (MPSoC) architectures. The consortium comprises 17 leading partners from four European countries, namely Bull, CAPS-Entreprise, UVSQ, DA, INT and CEA-LIST from France, HLRS, FZJ-ZAM, TUD-ZIH, GWT-TUD, RECOM Services, GNS and MAGMA from Germany, UAB, Indra and Robotiker from Spain, and Allinea from UK. Their Kick-off Meeting takes place in Dresden right before ISC 2007.

The project has been approved by ITEA2 and is being funded by the corresponding national Public Authorities, e.g. the German BMBF. The overall project will be managed by Bull, with the seven German partners being co-ordinated by HLRS.

The main achievements of ParMA will consist in:

- Evolving design and programming models to develop and restructure parallel applications in an easy and efficient way for multi-core/multi-threaded systems and MPSoCs

- Extending existing parallel tools for performance analysis and debugging (Kojak, Vampir, OPT, DDT, Marmot, Open MPI/Peruse) to support different flavours of parallelism (MPI, OpenMP, hybrid OpenMP/MPI, threads) and the latest features thereof. To offer a user-friendly and powerful environment to application developers, the integration of these tools in a common framework is envisioned

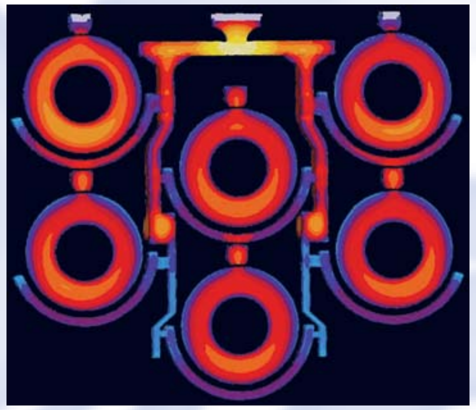

Solidification of a break disc pattern

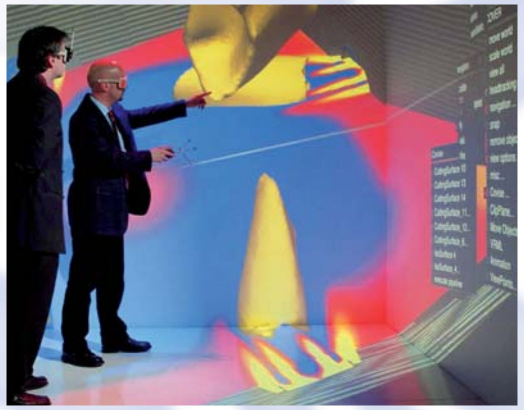

Computer-aided analysis of flame interactions in a 300 MWe combustion chamber in the Virtual Reality

- Developing and optimizing parallel applications from diverse domains, e.g. avionics applications, for multi-core architectures. Among the German partners, GNS, MAGMA, and RECOM will study metal forming or casting process simulations, to be used e.g. in the automobile industry, or combustion simulation

- Optimizing numerical libraries for multi-threaded architectures and enhancing the underlying Linux Operating System, e.g. to better utilize hierarchical NUMA systems.

HPC platforms for experimentation will be provided by Bull, while embedded systems will be provided and investigated by other French and Spanish partners.

Overall, the application developers will be supported in all steps of the development cycle. The new concepts and tools developed within ParMA will enable them to give birth to a new generation of (embedded) applications that could not be envisaged so far, to significantly speed-up their applications, and to consider much more complex models.

Contact Point at HLRS:
Bettina Krammer

Links:
http://www.parma-itea2.org
http://www.itea2.org

German Partners:
HLRS, http://www.hlrs.de
TUD-ZIH, http://www.tu-dresden.de/zih
FZJ-ZAM, http://www.fz-juelich.de/zam
RECOM Services,
http://www.recom-services.de
GWT-TUD, http://www.gwt-online.de
GNS, http://www.gns-systems.de
MAGMA, http://www.magmasoft.de

French Partners:
Bull, http://www.bull.com
CAPS-Entreprise,
http://www.caps-entreprise.com
UVSQ, http://www.uvsq.fr
INT, http://www.int-evry.fr
DA, http://www.dassault-aviation.com
CEA-LIST, http://www-list.cea.fr

Spanish Partners:
UAB, http://www.uab.es
Robotiker, http://www.robotiker.es
Indra, http://www.indra.es

UK Partners:
Allinea, http://www.allinea.com

• Bettina Krammer
• Rainer Keller

Höchstleistungs-
rechenzentrum
Stuttgart

Visualizing correctness checking results using a performance analysis framework

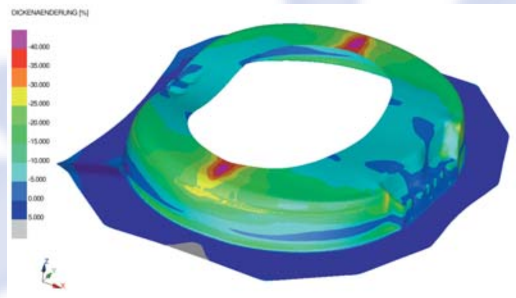

Deep drawing simulation of a wheel arch on a Linux cluster with 16 nodes

CoSpaces: Innovative Collaborative Work Environments for Individuals and Teams in Design and Engineering

CoSpaces is aimed at developing organizational models and distributed technologies for collaborative workspaces for individuals and project teams within distributed virtual manufacturing enterprises enabling effective partnerships, innovation, improved productivity, and reduced design cycles. It is a project funded by the EC running in the IST programme of the FP6. It is running from May 2006 to November 2009, within a total budget of € 11,839,600, with a total EC funding of € 8,000,000.

Introduction

Ambient Interfaces and the research domain of Ubiquitous Computing represent a third wave in computing. This paradigm shift is characterized by the integration of our increasingly complex technological landscape with the environment we occupy, while removing our perception of the computer in the tasks and activities we undertake. In practice the nature of these ambient interfaces ranges from the simple, such as devices that notify users of events by extending the computer interface into the workplace environment, through to the complex, where intuitive interface technologies such as optical tracking and gesture recognition are used to provide the user with the ability to directly manipulate highly visual information sets without perceiving the underlying technological infrastructure.

The evolution of IT over the past 20 years has led to the development of individual CAD/CAE workstations. While the computational power available to engineering professionals has grown exponentially, the collaborative dimension of the workspace has been largely underdeveloped. CoSpaces will provide an evolutionary path towards new and more collaborative work environments. Users of CoSpaces technologies in manufacturing and design, in co-operation with their suppliers, will be able to configure their own collaborative workspaces and utilize ground-breaking innovations in context-aware interfaces, natural interfaces, and "human-centric" workspaces supporting a range of collaboration scenarios and product life cycles.

Objectives

The CoSpaces project addresses three scientific and technological objectives:

- Evaluate collaboration at individual, team and enterprise levels, and develop collaboration models emphasising applications of problem solving, creativity, participatory and knowledge-based design in innovative collaborative work environments

- Create an innovative distributed software framework that will support easy creation of collaborative work environments for distributed knowledge workers and teams in collaborative design and engineering tasks

- Validate the distributed software framework for creating different classes of collaborative working styles required for collaborative design and engineering in the Aerospace, Automotive and Construction sectors.

Expected Impact

The research advances in CoSpaces will push the state-of-the-art of collaborative work environments in several key areas:

- New collaboration models addressing the development and deployment of collaborative technologies that empower workers and teams and having a strong human factors and business focus

- A practical collaborative software framework for supporting dynamic organizations that are executing complex processes in order to produce complex products

- Uncovering and addressing real-world constraints and barriers to bridge the gap between industrial engineers and scientific researchers in collaborative working

- Industry workspaces and interfaces based on deploying virtual interface technologies, leading to more practical and usable interfaces for real-world problems and users.

Projects

Projects

CoSpaces Software Framework

The CoSpaces software framework will support the easy creation of collaborative work environments for distributed workers. It includes a Foundation Platform to provide core services and high-level Collaboration Tools that support seamless and natural collaboration among distributed workers and teams.

Industry Workspaces

Three workspaces will be developed for validating industrial applications from the Automotive, Aerospace and the Construction industries:

- Distributed Design Workspace will address ad-hoc and scheduled collaboration between distributed, multi-functional design teams. Collaboration through fixed installations and mobile environments will be considered in this workspace

- Co-located Design Workspace will address how ad-hoc and planned meetings between co-located, multi-functional design teams could be supported, providing enhanced immersion, visualization, interaction, mobility and flexibility

- Mobile Service Workspace will address generic challenges in supporting mobile site workers to collaborate with remote experts during the realization or the support phase.

Roadmap Implementation

CoSpaces addresses the vision developed by the Future_Workspaces roadmap project involving over 100 key players from a multiple disciplines and defining a 10-year European vision for future collaborative working environments and scenarios. CoSpaces will initiate a programme of activities to implement this 10-year European vision for the benefit of European business, workforce, and society.

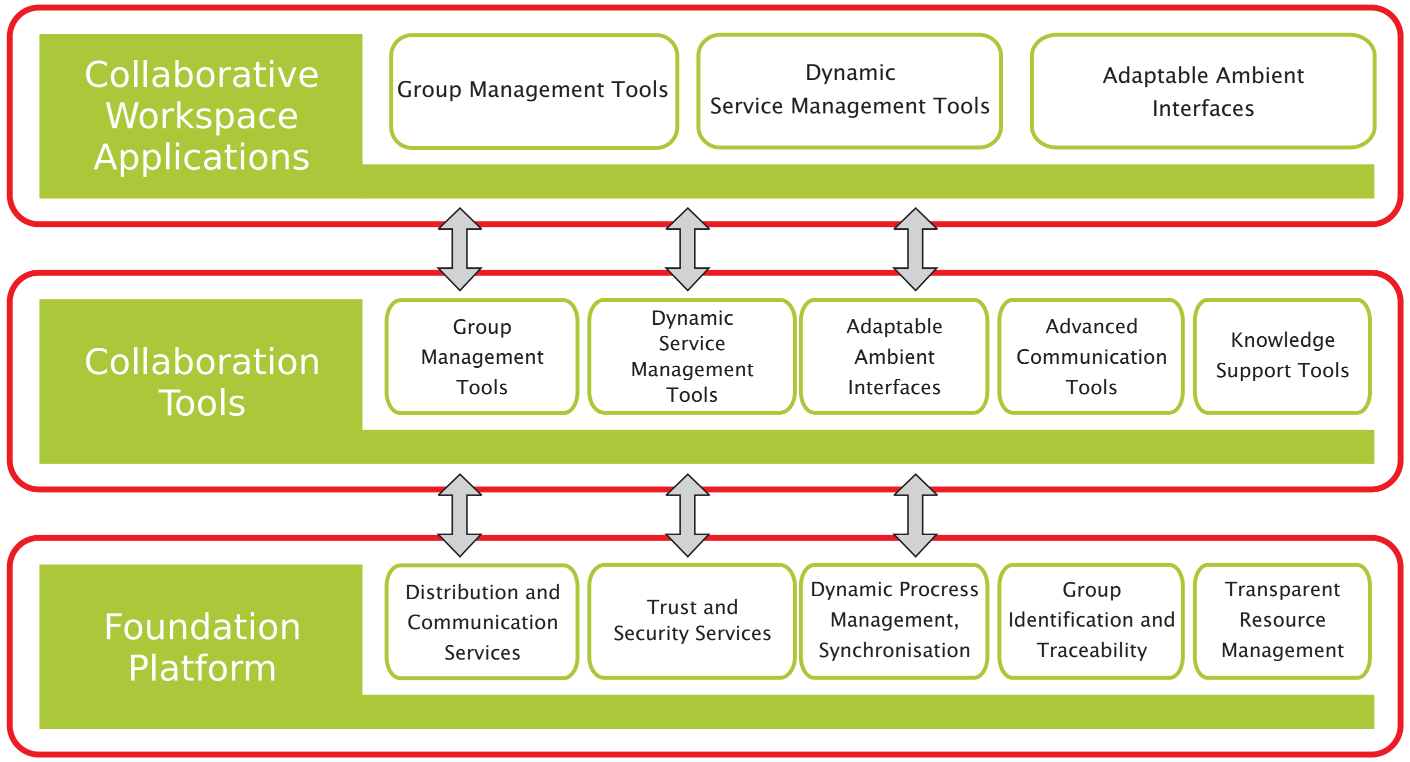

Projects

Involvement of the HLRS

The HLRS is one of the technical research partners, bringing into the project its expertise in web-services, visualization and computational simulation. Together with the other technical research partners it is creating the software foundation for the CoSpaces framework.

Project coordinator :
Scott Hansen
The Open Group

Tel: +32 2 675 1136
Fax: +32 2 675 7721

Email: s.hansen@opengroup.org
Project website: www.cospaces.org

Partners:
CIMPA (France), CERFACS (FR), CARSA (ES), COWI (DK), ESOCE Net (IT), Finanziaria Laziale di Sviluppo (IT), Fraunhofer Institute (DE), National Technical University of Athens (EL), Pragmasis (PT), Technology Application Network (UK), The Open Group (UK), UNINOVA (PT), University of Cologne (DE), University of Nottingham (UK), University of Salford (UK), University of Stuttgart (DE), Varinex Informatics (HU), Virtual Dimension Centre (DE)

• CoSpaces Consortium
• Andreas Kopecki

Höchstleistungs-rechenzentrum Stuttgart

A Fortran Binding for the GNU scientific Library

The GNU scientific library (GSL) [1] is a well-known collection of routines for scientific computing, covering many areas of numeric and semi-numeric algorithms. It has been implemented in C using object-based design principles. For this reason, accessing much of GSL's functionality from Fortran-based applications was error-prone, not easy to program, and could not be done in a portable manner. With the increasing support of the necessary Fortran 2003 features in commercial and open-source compilers, especially the C interoperation functionality, this situation is changing; most importantly, the programming effort for providing well-designed object-based Fortran interface mappings to C APIs is considerably reduced since there is no necessity any more to make use of each compiler vendor's proprietary extensions for interfacing Fortran with C. This motivated the implementation of a fully portable Fortran language binding FGSL for the GNU scientific library [2].

FGSL is for the most part implemented as a Fortran module; some auxiliary functions which require access to C specific object internals or type definition information needed to be coded in C. The module uses Fortran 90 encapsulation features to limit a client's access to the Fortran interface only; the GSL API calls are encapsulated within the FGSL module and only serve for internal use.

Usage

In order to access the abstract types' definitions as well as the explicitly interfaced API calls, a client program needs to access the module information via a **use fgsl** statement at the beginning of its scope. Apart from the larger amount of declarative code, it is then possible to proceed in a manner analogous to the C version of the client. Here an example treating numerical integration is provided, comparing the Fortran and C versions of the client. The integration routine requires a function object containing the argument function and a C-polymorphic (**void ***) parameter to be passed to it.

Fortran version

```
module mod_integration
  use fgsl
  use, intrinsic :: iso_c_binding
  implicit none
contains
  function f(x, params) bind(c)
    real(c_double), value :: x
    type(c_ptr), value :: params
    real(c_double) :: f
    real(c_double), pointer :: alpha
    call c_f_pointer(params, alpha)
    f = log(alpha * x) / sqrt(x)
  end function f
end module mod_integration
program integration
  use mod_integration
  implicit none
  real(fgsl_double), target :: alpha
  real(fgsl_double)  :: result, error
  integer(fgsl_int)    :: status
  type(c_ptr)          :: ptr
  type(fgsl_function) :: f_obj
  type(fgsl_integration_workspace) :: wk
  alpha = 1.0D0
  ptr = c_loc(alpha)
  f_obj = fgsl_function_init(f, ptr)
  wk = &
    fgsl_integration_workspace_alloc( &
    1000_fgsl_size_t)
  status = &
    fgsl_integration_qags(f_obj, &
    0.0D0, 1.0D0, 0.0_D0, 1.0D-7, &
    nmax, wk, result, error)
  write(6, fmt='(''Result : '',&
        F20.16)') result
  write(6, fmt='(''est. Error: '',&
        F20.16)') error
  call fgsl_function_free(f_obj)
  call &
    fgsl_integration_workspace_free(wk)
end program integration
```

C version

```
#include <stdio.h>
#include <math.h>
#include <gsl/gsl_integration.h>


double f (double x, void * params) {



  double alpha = *(double *) params;
  double f = log(alpha*x) / sqrt(x);
  return f;
}

int main (void) {


  double result, error;
  double alpha = 1.0;


  gsl_function f_obj;
  gsl_integration_workspace * wk;

  f_obj.function = &f;
  f_obj.params = &alpha;
  wk =
  gsl_integration_workspace_alloc(1000);


  gsl_integration_qags (&f_obj, 0, 1,
    0, 1e-7, 1000, wk, &result, &error);

  printf ("Result : .16f\n", result);
  printf ("est. error : %.16f\n",
           error);



  gsl_integration_workspace_free(wk);
  return 0;
}
```

As usual, more declarative code is needed in Fortran; the usage of the interface as well as the object definitions exactly matches the existing C API calls and typedefs. Note that while function return values can be ignored in C (cf. the call of **fgsl_integration_qags**), this is not allowed for explicit interfaces in Fortran; similarly it is seen that actual arguments given as constants are automatically cast to the correct dummy argument type in C, while Fortran, performing stricter type checking, would reject this. Since the function object components **f_obj.function** and **f_obj.params** are not directly accessible within Fortran, an additional subroutine call **fgsl_function_init** has been implemented which provides the Fortran-defined argument function as well as its parameters to the opaque object **f_obj**. Since a C function pointer is internally used, the argument function must have the **bind(c)** attribute; the function parameters need to be converted from a Fortran object (which may be non-interoperable) to a void-like pointer object using **c_loc**, and are subsequently unpacked within the argument function using **c_f_pointer**. Making use of non-interoperable argument functions and Fortran 2003-polymorphic parameters is also possible, but not yet supported by the presently available compilers; this will be implemented in a future version of FGSL. Compiling the Fortran executable requires access to the FGSL module information file (usually via a path specification given with the compilers' – I switch); linkage requires the FGSL add-on library in addition to the existing GSL libraries. Both module information file and add-on library are compiler specific and hence must be generated separately for each Fortran compiler, perhaps even for each compilers' major version changeover.

Implementation Status

The presently available version of FGSL [2] is an alpha release; furthermore not many compilers are presently able to properly build and execute the interface. However, a large subset of the version 1,8 GSL functionality is accessible. More than 1,200 API calls are available; FGSL presently consists of around 16,000 lines of Fortran and 1,000 lines of C code. Those areas where Fortran's strengths lie (complex types, vectors and arrays) are not covered, and well-established APIs like LAPACK, BLAS and FFT are also considered outside the purview of FGSL. For the case of arrays and matrices, mapping routines to Fortran arrays and array pointers of rank 1 and 2 are provided which make the GSL accessor routines superfluous; this enables the use of the abstract **fgsl_vector and fgsl_matrix** types for other areas of FGSL in conjunction with Fortran-style array processing.

A much more extensive discussion of FGSL's implementation details will be given in the August 2007 issue of the ACM SIGPLAN Fortran Forum [3].

References

[1] See the GSL home page at
http://www.gnu.org/software/gsl/

[2] The FGSL source code is available from
http://www.lrz-muenchen.de/services/software/mathematik/gsl/fortran/index.html

[3] Choose the "Newsletters" link from
http://portal.acm.org/dl.cfm?coll=portal&dl=ACM&CFID=18491455&CFTOKEN=15629419 and then select the "ACM SIGPLAN Fortran Forum" sub-entry.
A subscription is required

• Reinhold Bader

Leibniz Computing Centre

Tools

The European UNICORE 6 Grid Middleware

The UNICORE Grid system provides a seamless, secure and intuitive access to distributed Grid resources such as supercomputers, clusters, and large server farms. In recent years, UNICORE 5 is used as a well-tested Grid middleware in scientific production Grids (e.g. DEISA, D-Grid) and for business use cases (e.g. T-Systems, Philips Research). In addition, 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.

More recently, the new Web services-based UNICORE 6 has become available in beta state. It is based on common open standards that have emerged from various standardization bodies such as OASIS (Organization for the Advancement of Structured Information Standards) and OGF (Open Grid Forum).

Web and Grid Services

At the time of writing, the technologies and techniques of the Grid community follow an approach of Service Oriented Architectures (SOAs). A well-known set of service-oriented OGF standards are based on the Open Grid Services Architecture (OGSA) that emphasizes the use of so-called Grid services. In the last years, the Web services technology is used to implement these concepts of Grid services that lead initially to the Open Grid Services Infrastructure (OGSI). OGSI introduced the concepts of so-called stateful Web services that expose and manage the state (e.g. computational job status) of various kinds of Grid resources, including lifecycle management. More recently, specifications of the Web community have become increasingly important and thus the Web Services Resource Framework (WS-RF) was accepted in April 2006 as an official OASIS standard to re-factor the concepts of OGSI to exploit new Web service standards (e.g. Web Services Addressing).

UNICORE 6 Architecture

The rapid adoption of OGSA and WS-RF concepts within Grid technologies allows for an extension of UNICORE concepts through the use of common open standards. The Figure below illustrates the architecture of the new UNICORE 6 Grid middleware and its Web services-based interfaces that are conform to OGSA concepts. Among other additional interfaces, the UNICORE Atomic Services (UAS) are the main interfaces that allow for the exploitation of the core functionality by Web services-based clients. This core functionality includes job submission and control, file and data transfer as well as storage management.

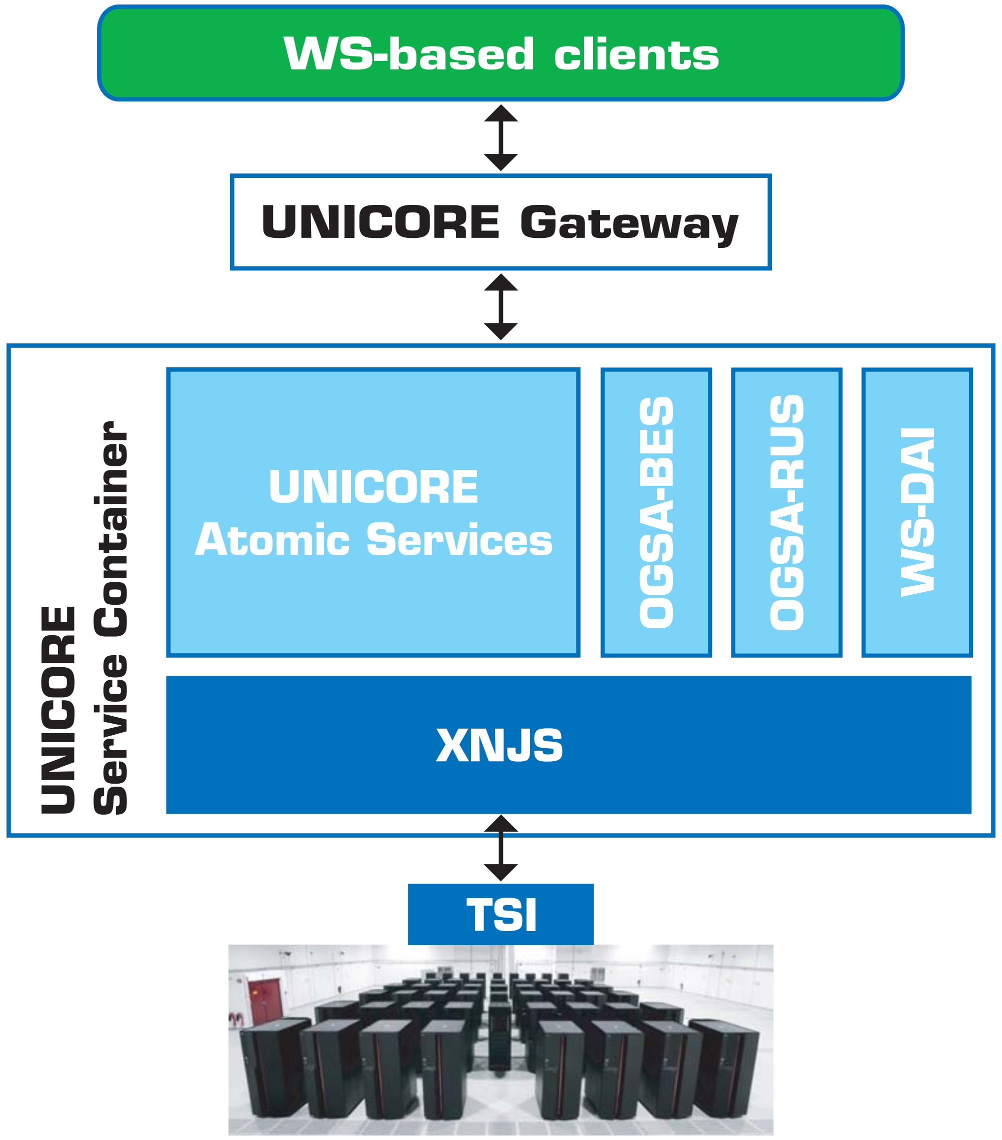

In more detail, UNICORE 6 provides a Service Container that is able to host WS-RF compliant Grid services. Authenticated end-user requests from different Web services-based clients – e.g. command line, portal, API or the GPE UNICORE Client from Intel – pass the UNICORE Gateway and initiate operations of services deployed within the UNICORE Service Container.

The UAS consists of several WS-RF standards compliant stateful Web services. First and foremost, the Target System Service (TSS) provides access to a stateful resource that models a physical computational Grid resource like a supercomputer. It exposes various pieces of information, e.g. details about the total numbers of CPUs, memory, etc. and preinstalled applications on the Grid-enabled resource. Through the TSS Grid jobs described in the OGF standard Job Submission Description Language (JSDL) are submitted to the UNICORE site. The jobs are controlled with the Job Management Service (JMS). To support data staging of JSDL jobs, the Storage Management Service (SMS) is used to access storages within Grid infrastructures. The transfer itself is realized by the File Transfer Service (FTS), which offers different solutions based on open OGF standards like e.g. Random ByteIO (RbyteIO) and Streamable ByteIO (SbyteIO).

These WS-RF compliant services expose the functionality of UNICORE 6 via common open standards. The function of the new Network Job Supervisor (XNJS) as the execution backend is to control and manage the state and persistency of jobs. Hence, one of the major tasks of the XNJS is to parse JSDL documents and turn the rather abstract job descriptions to site-specific commands by using an enhanced Incarnation DataBase (IDB). Authorization is done by using the enhanced UNICORE User Data Base (UUDB) in conjunction with extensible Access Control Markup Language (XACML) policy validations. Then all commands are forwarded to the UNCORE Target System Interface (TSI) which is directly connected to the already existing batch sub-system (e.g. LoadLeveler, Torque, LSF) running on the supercomputer.

Additional Service Interfaces

The UNICORE 6 architecture provides extensibility towards additional services being deployed in the UNICORE Service Container. These services realize additional functionalities such as the OGF standard Web services Data Base Access and Integration Service (WS-DAI) used for the access to relational or XML-based databases. In addition, the OMII-Europe project (see Inside Vol. 4 No. 2, autumn 2006) currently augments the UNICORE 6 Grid technology with other useful emerging OGF standards via services such as the OGSA-Resource Usage Service (RUS) and OGSA-Basic Execution Services (BES).

Summary

The new Web services enabled UNICORE 6.0 offers significant improvements in terms of usability for end-users, performance, extensibility for developers, interoperability with other Grid middleware, as well as open standards compliance. On top of the UNICORE Atomic Services, an integration of scientific applications via other high-level services can easily be realized and deployed. Therefore UNICORE 6.0 is well suited to be used in the future European Supercomputing infrastructure. To foster ongoing developments and the uptake in various European and international infrastructures and R&D projects, UNICORE 6 is available as open source under BSD license from SourceForge. Main contributors to UNICORE are: Fujitsu Labs of Europe, Intel, ICM Warsaw, CINECA Bologna, the University of Manchester, and Forschungszentrum Jülich.

More information and lightweight UNICORE 6 installation packages are available at http://www.unicore.eu

- Morris Riedel
- Bernd Schuller
- Achim Streit

Forschungszentrum Jülich

Tools

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

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

Contact
Leibniz-Rechenzentrum

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

Centres

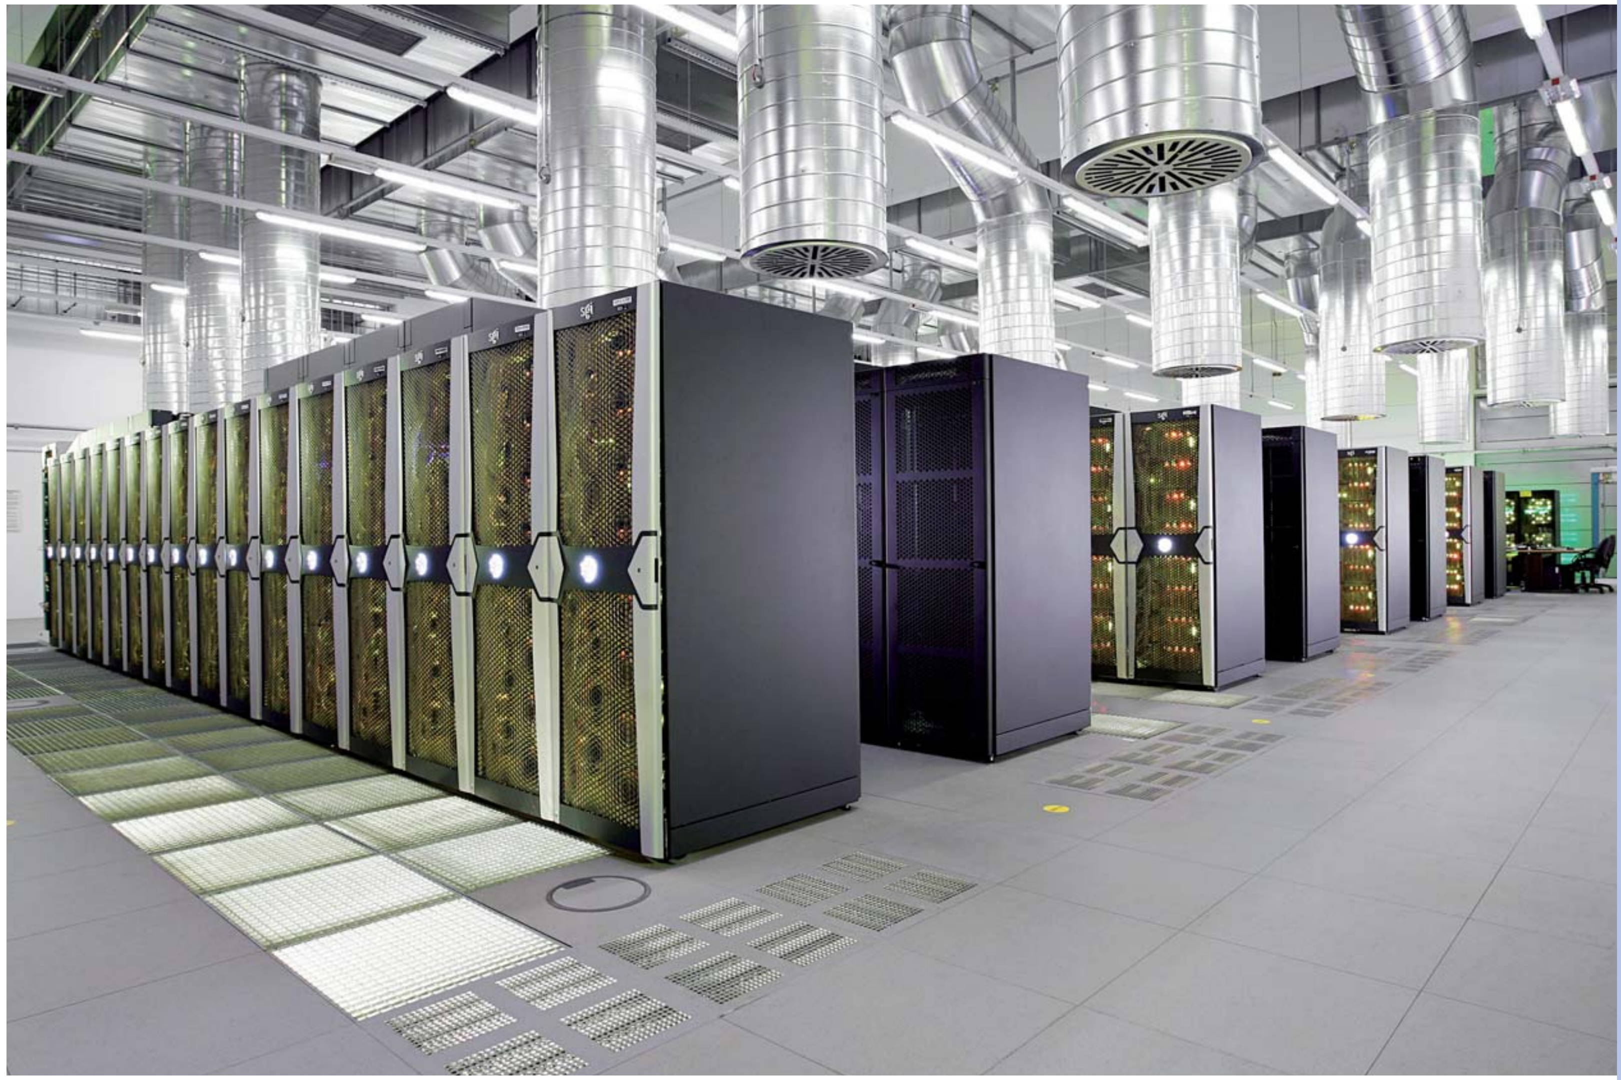

View of „Höchstleistungsrechner in Bayern HLRB II", an SGI Altix 4700
Foto: Kai Hamann, produced by gsiCom

Compute servers currently operated by LRZ are

System	Size	Peak Performance (GFlop/s)	Purpose	User Community
SGI Altix 4700 19 x 512 way	9,728 Cores 39 TByte	62,259	Capability computing	German universities and research institutes
SGI Altix 4700 256 way	256 Cores 1 TByte	1,640	Capability computing	German and Bavarian universities and research institutes
SGI Altix 3700 BX2 128-way	128 processors 512 GByte memory	820	Capability computing	Bavarian universities
Linux Cluster Intel IA64 2-way	68 nodes 136 processors 816 GByte memory	870	Capability and capacity computing	Bavarian universities
Linux Cluster Intel IA64 4- and 8-way	19 nodes 84 cores 250 GByte memory	440	Capacity computing	Munich universities
Linux cluster Intel IA32 Intel&AMD EM64T	154 nodes 192 processors 320 GByte memory	850	Capacity computing	Munich universities

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

Centres

H L R S

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

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

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

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

Contact

Höchstleistungsrechenzentrum Stuttgart (HLRS)
Universität Stuttgart

Prof. Dr.-Ing. Michael M. Resch
Nobelstraße 19
70500 Stuttgart
Germany
Phone +49-711-685-872 69
resch@hlrs.de
www.hlrs.de

View of the NEC SX-8 at HLRS

Centres

Compute servers currently operated by HLRS are

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

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

Centres

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

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

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

At present, two research groups exist: the group Elementary Particle Physics, headed by Karl Jansen and located at the DESY laboratory in Zeuthen and the group Computational Biology and Biophysics, headed by Ulrich Hansmann at the Research Centre Jülich.

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

Contact

John von Neumann -
Institut für Computing (NIC)
Zentralinstitut für
Angewandte Mathematik (ZAM)
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/nic
www.fz-juelich.de/zam

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

Compute servers currently operated by NIC are

System	Size	Peak Performance (GFlop/s)	Purpose	User Community
IBM Blue Gene/L "JUBL"	8 racks 8,192 nodes 16,384 processors PowerPC 440 4 TByte memory	45,875	Capability computing	German universities, research institutes and industry
IBM pSeries 690 Cluster 1600 "JUMP"	41 SMP nodes 1,312 processors POWER4+ 5,1 TByte memory	9,000	Capability computing	German universities, research institutes and industry
IBM BladeCentre-H "JULI"	2 racks 56 Blades 224 PowerPC 970 MP cores 224 GByte memory	2,240	Capability computing	Selected NIC projects
IBM Cell System "JUICE"	12 Blades 24 Cell processors 12 GByte memory	4,800 (single precision)	Capability computing	Selected NIC projects
AMD Linux Cluster "SoftComp"	66 compute nodes 264 AMD Opteron 2.0 GHz cores 264 GByte memory	1,000	Capability computing	EU SoftComp community
Cray XD1	60 dual SMP nodes 120 AMD Opteron 2.2 GHz processors 264 GByte memory	528	Capacity and capability computing	NIC research group "Comp. Biology and Biophysics"
apeNEXT (special purpose computer)	4 racks 2,048 processors 512 GByte memory	2,500	Capability computing	Lattice gauge theory groups at universities and research institutes
APEmille (special purpose computer)	4 racks 1,024 processors 32 GByte memory	550	Capability computing	Lattice gauge theory groups at universities and research institutes

Centres

High Performance Computing in Science and Engineering – The 9th Results and Review Workshop of the HPC Centre Stuttgart (HLRS)

The 9th Results and Review Workshop of the HLRS brought together more than 50 participants from german research institutions, the steering committee and scientific support staff of HLRS. More than 30 sophisticated talk and poster presentations were selected in advance from the steering committee out of the yearly supercomputer project reports and presented at the Results and Review Workshop. State-of-the-art scientific simulations on the supercomputer systems of HLRS again emphasized the world-class research done at the Centre obtaining outstanding results in achieving highest performance for production codes which are of particular interest for both scientists and engineers.

The presentations covered all fields of computational science and engineering ranging from CFD via computational physics and chemistry to computer science with a special emphasis on industrially relevant applications.

Every year the three most outstanding projects are honoured by the Golden Spike Award which is awarded by the steering committee of HLRS. The laureates of 2006 and their project titles are:

- Jens Harting, Institute for Computational Physics, University of Stuttgart

  "Rheological properties of binary and ternary amphiphilic fluid mixtures"

- Andreas Marek, Max-Planck Institute for Astrophysics, Garching

  "The SuperN-Project: Understanding core collapse supernovae"

- Sven Ganzenmüller, Department of Computer Engineering, University of Tübingen

  "Object-Oriented SPH-Simulations with Surface Tension".

A very interesting tutorial titled "Performance Analysis for leading-edge HPC systems" and presented by Andreas Knüpfer from the ZIH, Dresden University of Technology, completed the successful workshop.

High Performance Computing in Science and Engineering '06

This book presents the state-of-the-art in simulation on supercomputers. Leading researchers present results achieved on systems of the High Performance Computing Centre Stuttgart (HLRS) for the year 2006.

The reports cover all fields of computational science and engineering ranging from CFD via computational physics and chemistry to computer science with a special emphasis on industrially relevant applications. Presenting results for both vector-systems and micro-processor based systems the book allows to compare performance levels and 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 highest performance for production codes are of particular interest for both the scientist and the engineer. The book comes with a wealth of coloured illustrations and tables of results.

Nagel, W. E., Jäger, W., Resch, M.
High Performance Computing
in Science and Engineering '06
Springer, Berlin, Heidelberg, New York, 2006

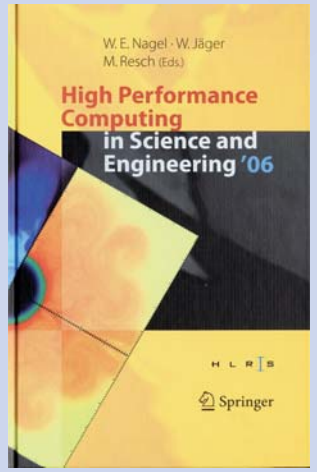

Activities

The 6th HLRS/hww Workshop on Scalable Global Parallel File Systems, HNF Europe Spring Meeting 2007 and OpenFabrics

Representatives from science and industry interested in high performance storage solutions did meet at HLRS during April 16th-18th, 2007, for the sixth annual Workshop on Scalable Global Parallel File Systems. Under the motto of "Moving the World: Global and Individual", this year's three-day event tried to leverage contributions coming from three major research and development areas: federation of OSD-based file systems for use in European HPC Grid projects, an all-optical communication and switching infrastructure being able to transport 100-Gbit/s serial signals within the framework of the German Gauss Alliance, as well as standardization efforts performed by the OpenFabrics Alliance on the InfiniBand architecture. More than 150 participants did follow a total of 45 presentations that have been on the workshop agenda.

The opening address on Monday morning was given by Prof. Dr. Michael Resch, the HLRS director. Dr. Peter Braam, founder and CEO of Cluster File Systems, delivered the keynote speech on desirable features and mechanisms needed for federated file systems. His conclusion was that OSD-based file systems, like Lustre, could very well deal with the additional control complexity. However, there would be more input needed on use patterns from the Grid community and industry. Dr. Roger Haskin, IBM Almaden, talked about research work on the IBM General Parallel File System (GPFS) with respect to wide-area operation as well as interfacing GPFS with other parallel file and HSM systems. Dr. Garth Gibson, Panasas, gave a most ingenious presentation on the evolving Parallel NFS standard, which will be part of NFS version 4.1, most likely. Tigran Mkrtchyan, DESY, demonstrated how Parallel NFS may be used as one of the preferred client interfaces to the multi-tier data distribution within the Large Hadron Collider project at CERN. Finally, Thomas Bönisch, HLRS, gave an overview on how Multi-Cluster GPFS is used to enable a globally distributed file system integration for the eleven DEISA partners in Europe which may be extended to selected TeraGrid sites in the US.

Tuesday morning saw a series of simultaneous presentations on how parallel file systems may be layered on each other in order to build multi-tier archives. Examples were GPFS/HPSS, GPFS/TSM and Quantum's StorNext as part of CASA, the Cray Advanced Storage Architecture. Exports of GPFS via Samba may enable an expanded audience for campus environments. Also OpenFabrics started with a series of interesting talks on InfiniBand converged fabrics for HPC storage, virtualization of storage hardware and software, as well as design constraints for low-latency communication platforms.

On Tuesday afternoon, HNF Europe provided the venue for the forthcoming 100-Gbit/s all-optical communication architectures in European research networks. Dr. Peter Krummrich, Siemens, gave an overview of Ex@Grid, a research project of the German Ministry of Education and Research (BMBF), which will lead to a wavelength specific switching (WSS) infrastructure for the German HPC centres. This infrastructure will enable and/or accelerate a new set of distributed scientific applications which revolve around a network-centric data management scheme. Thomas Brenner, Alcatel-Lucent, provided a profound display of the technologies needed to build up a true 100-Gbit/s serial transmission and switching system, based on experience with earlier 40-Gbit/s prototypes (as explained at last year's event). There were more talks in parallel sessions on loosely coupled TCP acceleration as well as ongoing standardization work at various IEEE 802.3 committees towards unified data Centre fabrics.

Tuesday night's agenda started with an invitation to explore the Mercedes-Benz Automobile Museum - who has contributed the "Moving the World: Global and Individual" motto. From there the bus moved on to the Stuttgart Museum of Fine Arts for a reception at o.T. Lounge and a very memorable dinner at Cube Restaurant on fourth floor.

Wednesday again saw a parallel track organized by the OpenFabrics Alliance. These sessions covered various InfiniBand issues from physical layer and kernel problems to storage and message passing. Finally, HLRS provided insight into its Centre on many levels: parallel storage systems intended for use in DGrid and Ex@Grid, the GRAU Datastorage midrange HSM, work towards massively parallel HSMs, parallel applications within the Teraflop Workbench, augmented reality and simulation steering to create hybrid prototypes, as well as private tours of the Centre.

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

Activities

The 6^{th} HLRS-NEC Teraflop Workshop

The sixth Teraflop Workshop continues the series of workshops which are held twice a year, alternatingly in Japan and at HLRS, within the Teraflop-Workbench co-operation between NEC and HLRS. The workshop provides a meeting platform for scientists, application developers, hardware designers and international HPC experts to discuss recent developments and future directions of supercomputing. This years spring event was held on March 26^{th} – 27^{th}, focusing on the applications perspective of high end computer systems. Special attention was paid to coupled systems applications combining different physical phenomena with possibly distinct requirements on computer architectures.

The opening session on Monday morning was dedicated to developments in architectures and performance achievements. Speakers Prof. Satoshi Matsuoka (Tokyo Institute of Technology), telling about "A Year with TSUBAME and its Possible Future Petaflops Extensions", and Prof. Jay Boisseau (Texas Advanced Computing Centre, The University of Texas at Austin) with his talk "The Path to Petascale" and Prof. Hiroaki Kobayashi (Information Synergy Centre, Tohoku University) with a talk on "ISC plans and update" discussed recent technology developments in the US and Japan.

Monday afternoon sessions were dedicated to applications and computational methods on leading vector and scalar platforms. Applications coming from different scientific areas, such as simulation of ionic liquids, CFD, supernovae, premixed flames and fluid-structure interaction showed the wide range of High Performance Computing areas.

The second day was focusing on coupled applications. Starting with ocean modeling and climate research over medical applications like blood flow simulations and computational modeling of the respiratory system up to aero-acoustics, ocean water plants, and aero-elastic simulations for helicopter rotors, the whole range of strongly coupled and weakly coupled applications was covered.

The success of this workshop will be followed up by the book "High performance computing on vector systems 2007", to appear in autumn. The next workshop will be held at Tohoku University, Japan, in November 2007.

Activities

20 Years John von Neumann Institute for Computing and 50^{th} Anniversary of John von Neumann's Death

In 1987 HLRZ (Höchstleistungsrechenzentrum), the first German supercomputing centre to be accessible nationwide, was founded. It opened the way for researchers to carry out scientific projects on the Cray X-MP supercomputer installed at that time in Jülich. In 1998 the HLRZ was reorganized, and the name "John von Neumann Institute for Computing" (NIC) was adopted. In 2007 the HLRZ/NIC supercomputing centre celebrates 20 years of successful operation. The celebration of the anniversary took place at Forschungszentrum Jülich on February 8^{th}, 2007. This date marks the 50^{th} anniversary of the death of John von Neumann and was therefore chosen to remember the outstanding contributions of this great scientist to the development of modern computers. At present NIC is a joint institute of Deutsches Elektronen-Synchrotron (DESY), Forschungszentrum Jülich, and Gesellschaft für Schwerionenforschung (GSI). Its supercomputer facilities support computer simulations in various fields of science for about 130 research groups at universities and national laboratories.

Celebrating John von Neumann

From Computational Biophysics to Systems Biology (NIC-Workshop)

About 120 scientists and students from Germany, Europe and the USA participated in the second annual workshop "From Computational Biophysics to Systems Biology" (CBSB07). The workshop took place from May 2^{nd} to May 4^{th}, 2007 at the Forschungszentrum Jülich and was organized by the NIC research group "Computational Biology and Biophysics" headed by Ulrich H. E. Hansmann. Leading experts from physics, chemistry, biology, and computer science discussed physics-based approaches to systems biology. An important question was how to bridge the different scales from single molecules to the entire cell that are important for a physics-based simulation of a cell. In a "town hall" meeting the participants emphasized the role of supercomputers in obtaining a detailed understanding of the working of cells. They agreed that further progress will require the supercomputer centres to provide both increased computing capacity and improved support and accessibility.

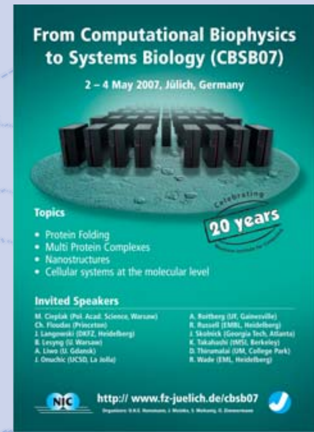

Workshop poster

Cell Cluster JUICE

At the beginning of 2007, the project JUICE (Jülich Initiative Cell cluster) was established. In the framework of this project a cluster of 12 IBM QS20 blades – 2 BladeCentre1 chassis with 24 CBE (Cell Broadband Engine) processors in total – were procured to examine the potential of Cell processors as a building block for future high-end computing systems. This multicore processor – based on a PowerPC core and eight synergistic processing elements – distinguishes itself by its extremely high performance single-precision arithmetic. A CBE running at 3.2 GHz reaches a peak performance of 204.8 GigaFlop/s, thus leading to a peak performance of nearly 4.9 TFLOP/s for the whole cluster. Due to the large performance difference between single-precision and double-precision arithmetic with a ratio of 14 it is essential to exploit 32-bit floating point arithmetic whenever feasible.

While the installation of an Infiniband network is in progress, at present, the cluster interconnect is GigaBit Ethernet. ParTec's ParaStation is used as cluster middleware. Using its MPI stack, test jobs could be set up which were able to reach more than 90% of the cluster's peak performance.

Cell Cluster Meeting at Forschungszentrum Jülich

In the context of the project JUICE a Cell Cluster Meeting took place on May 10^{th} –11^{th}. The meeting was organized by NIC/ZAM at the Forschungszentrum Jülich and was intended as a forum to discuss current research on Cell and to provide an opportunity to get in touch with other researchers working on related topics. Hardware and software developers gave an overview of the Cell roadmap. Further presentations were given about performance tools, numerical aspects when using mixed precision, an implementation of the Lanczos method, lattice QCD, and management of clusters of Cell blades.

More information on the cell cluster JUICE and this workshop can be found at http://www.fz-juelich.de/zam/juice

Virtual Institute High-Productivity Supercomputing

To advance the state-of-the-art in HPC programming tools, the John von Neumann Institute for Computing at Forschungszentrum Jülich has joined forces with the Institute for Scientific Computing at RWTH Aachen, the Centre for Information Services and High-Performance Computing at TU Dresden, and the Innovative Computing Laboratory at the University of Tennessee in Knoxville. Together they have established the Virtual Institute High-Productivity Supercomputing (VI-HPS), a collaboration that combines the expertise of some of the key players in the area of performance analysis and error detection tools for high-end computing systems.

The mission of VI-HPS is the continued development, integration, and deployment of the following major HPC programming tools, with an emphasis on scalability and ease of use:

- The automatic trace analyzer KOJAK/SCALASCA (Jülich, Knoxville)
- The visual trace browser VAMPIR (Dresden)
- The error detection tool MARMOT (Dresden in partnership with the High-Performance Computing Centre Stuttgart)
- The hardware-counter library PAPI (Knoxville)

VI-HPS spokesman Felix Wolf (Jülich) is excited about the synergies created through this partnership: "The integration of leading tool projects will provide a clear productivity benefit for the HPC user community." And Christian Bischof (Aachen), deputy spokesman of the collaboration, explains: "Besides enhancing the technical capabilities of the VI-HPS tool suite, a major focus will also be convincing more users that using these tools can save them time both during development and when the application goes into production." For this purpose, RWTH Aachen will provide substantial training and support capacities.

Virtual institutes are a well-established funding instrument of the Helmholtz Association. Helmholtz centres and universities use this instrument to concentrate research capacities and thereby create centres of excellence of international standing in key areas of research. Virtual institutes are supplemented with an annual budget of 300 K € for three years.

VI-HPS will be inaugurated during a half-day workshop held in Jülich on July 4^{th}, 2007. For more information see http://www.vi-hps.org

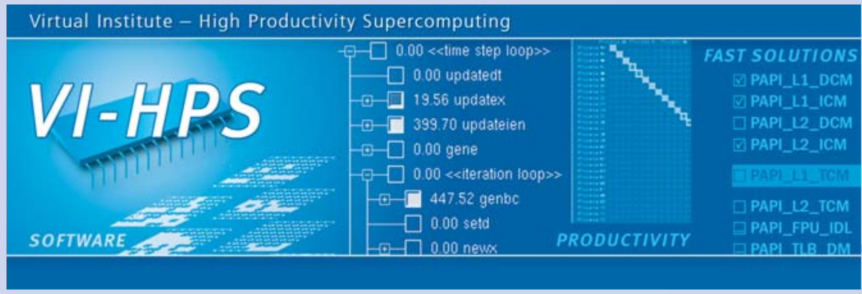

German Research School for Simulation Sciences

Computational science, based on top-level supercomputers, has established itself as the third pillar of scientific methodology alongside theory and experiment. The German Research School for Simulation Sciences (GRS) will provide excellent post-graduate education in this important interdisciplinary field of science and engineering. It will offer programs for both Master and PhD students. The studies will include all subjects relevant for simulation on high-performance computers ranging from the disciplinary sciences like physics, chemistry or biology, and the engineering sciences, to the cross-disciplinary areas of mathematics and computer science.

The GRS is founded on an equal footing by Forschungszentrum Jülich and RWTH Aachen. Both organizations have long and intense research activities in computational science and engineering, and in developing the basic methodology of the simulation sciences. Therefore they are the ideal partners for this innovative endeavour, which is supported by the Federal Ministry of Education and Research, the State of North-Rhine Westphalia and the Helmholtz Association.

The GRS will have locations both in Jülich and in Aachen; it will become operational in 2007 accepting first PhD students in autumn 2007 and Master students one year later.

Contact:
grs@grs-simulation-sciences.de

UNICORE Summit 2007

The 3rd UNICORE Summit will take place in Rennes (France) on August 28th in conjunction with the Euro-Par 2007 conference and is a unique opportunity for Grid users, developers, administrators, researchers, and service providers to meet and share experiences as well as to discuss future developments. Meet the UNICORE development community to get an inside view of the new Web Services-based and open standard compliant UNICORE 6 Grid middleware.

The Call for Papers and more details regarding the UNICORE Summit 2007 are available at
http://summit.unicore.org

UNICORE SUMMIT

Activities

ParCo 2007 in Jülich and Aachen

The well-known conference Parallel Computing (ParCo) will take place from September 3rd-7th, 2007 in Jülich and Aachen, Germany. The conference is organized by the non-profit foundation ParCo Conferences in co-operation with the Forschungszentrum Jülich and the RWTH Aachen University. ParCo 2007 marks a quarter of a century of the international conferences on parallel computing that started in Berlin in 1983. This makes ParCo the longest running series of international conferences on the development and application of high speed parallel computing technologies in Europe.

The aim of the conference is to give an overview of the state-of-the-art developments, applications and future trends in high-performance computing for all platforms. The conference addresses all aspects of parallel computing, including applications, hardware and software technologies as well as languages and development environments. Special emphasis will be placed on the role of high-performance processing to solve real-life problems in all areas, including scientific, engineering and multidisciplinary applications and on strategies, experiences and conclusions made with respect to parallel computing.

On September 3th, two tutorials on advanced parallel programming with OpenMP and MPI will open the conference. Besides presentations in the regular technical program and mini-symposia on selected topics, ParCo will feature invited talks by Barbara Chapman on "Programming in the Multicore Era", Satoshi Matsuoka on "Towards Petascale Grids as a Foundation of E-Science" and Marek Behr on "Simulation of Heart-Assist Devices".

Further information can be found at the web page
http://www.fz-juelich.de/parco2007

CSD

Computational Science & Discovery (CSD) is a new title from IOP Publishing, and is now accepting submissions. It will focus on scientific advances and discovery through computational science with a multidisciplinary breadth in physics, chemistry, biology and applied science. Within their papers, authors are encouraged to include details of the scientific advances made, their numerical methods, verification and validation of codes, and the enabling technologies they used - for example, in data management, networking and visualization, among other areas. Papers will be rigorously peer-reviewed and published electronically.

To find out more about the journal, including author guidelines, and members of the growing Editorial Board, please visit http://iop.org/journals/csd

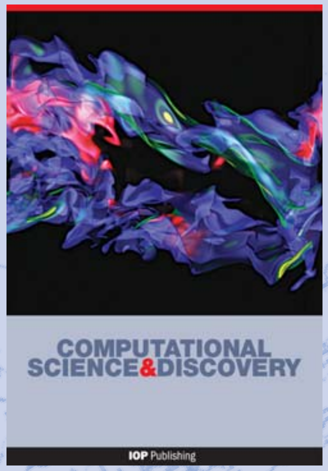

High Performance Computing Courses and Tutorials

LRZ
www.lrz.de

Multi-Core Training

Date

July 23-26, 2007

Location

LRZ building, Munich/Garching

Contents

- Munich Multi-Core Initiative
- Optimization
- SGI Montecito Multi-Core System
- Threading Concepts
- Introduction to Intel® Compilers and Intel® VTuneTM Performance Analyzer
- Programming with Threads
- Programming with OpenMP Threads
- Multithreaded Programming Methodology
- Performance Analysis and De-bugging Thread Correctness checking for Multi-Core Platforms using Intel® Threading Tools

This training is organized by Intel, the Munich Multi-Core Initiative, and LRZ.

Webpage

http://www.lrz.de/services/compute/courses/#Multi-Core_Training

Iterative Linear Solvers and Parallelization

Date

September 17-21, 2007

Location

LRZ building, Munich/Garching

Contents

The focus is on iterative and parallel solvers, the parallel programming models MPI and OpenMP, and the parallel middleware PETSc. Different modern Krylov Subspace Methods (CG, GMRES, 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 University of Kassel, HLRS, IAG, and LRZ.

Webpage

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

NIC
www.fz-juelich.de/nic

User Course NIC Guest Student Programme: Introduction to Parallel Programming with MPI and OpenMP

Date

August 7-10, 2007

Location

NIC/ZAM, Research Centre Jülich

Contents

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

Webpage

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

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

Date

August 13-14, 2007

Location

NIC/ZAM, Research Centre Jülich

Contents

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

Webpage

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

Parallel Programming with MPI, OpenMP, and PETSc

Date

November 26-28, 2007

Location

NIC/ZAM, Research Centre Jülich

Contents

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

Webpage

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

CECAM Tutorial Programming Parallel Computers

Date

January 2008 (tentative)

Location

NIC/ZAM, Research Centre Jülich

Contents

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

Webpage

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

High Performance Computing Courses and Tutorials

HLRS
www.hlrs.de

Parallel Programming with MPI, OpenMP and PETSc

Date

August 8-10, 2007

Location

Manno (CH), CSCS

Contents

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

Webpage

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

Message Passing Interface (MPI) for beginners

Date

October 8-9, 2007

Location

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/news-events/events

Shared memory parallelization with OpenMP

Date

October 10, 2007

Location

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/news-events/events

Advanced Topics in Parallel Programming

Date

October 11-12, 2007

Location

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/news-events/events

Introduction to Computational Fluids Dynamics

Date

October 15-19, 2007

Location

Stuttgart, HLRS

Contents

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

Webpage

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

Fortran for Scientific Computing

Date

October 22-26, 2007

Location

Stuttgart, HLRS

Contents

This introduction to C++ is taught with lectures and hands-on sessions. This course is organized by HLRS and Inst. 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/news-events/events

inSiDE

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

Publishers

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

Editor

F. Rainer Klank, HLRS klank@hlrs.de

Design

Stefanie Pichlmayer sp028@hdm-stuttgart.de

Authors

Prof. Dr. Achim Bachem	a.bachem@fz-juelich.de
Reinhold Bader	reinhold.bader@lrz-muenchen.de
Dr. Mehdi Behbahani	mb@cats.rwth-aachen.de
Prof. Dr. Marek Behr	behr@cats.rwth-aachen.de
Prof. Dr. Heinz-Gerd Hegering	heinz-gerd.hegering@lrz-muenchen.de
Rainer Keller	keller@hlrs.de
Dr. Markus Kloker	kloker@iag.uni-stuttgart.de
Andreas Kopecki	kopecki@hlrs.de
Bettina Krammer	krammer@hlrs.de
Prof. Dr. Dr. Thomas Lippert	th.lippert@fz-juelich.de
Dr. Ralf Messing	messing@iag.uni-stuttgart.de
Mike Nicolai	nicolai@cats.rwth-aachen.de
Markus Probst	probst@cats.rwth-aachen.de
Prof. Dr. Michael Resch	resch@hlrs.de
Morris Riedel	m.riedel@fz-juelich.de
Christoph Schaupp	christoph.schaupp@aer.mw.tum.de
Dr. Wolfram Schmidt	schmidt@astro.uni-wuerzburg.de
Dr. Bernd Schuller	b.schuller@fz-juelich.de
Jan Schulze	jan.schulze@unibw.de
Dr. Jörn Sesterhenn	joern.sesterhenn@t-online.de
Dr. Achim Streit	a.streit@fz-juelich.de

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