INSIDE • VOI. 12 No.2

# Innovatives Supercomputing in Deutschland



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

# **Publishers**

Prof. Dr. A. Bode | Prof. Dr. Dr. Th. Lippert | Prof. Dr.-Ing. Dr. h.c. Dr. h.c. M. M. Resch

# Editor

F. Rainer Klank, HLRS

klank@hlrs.de

# Design

Pia Heusel, HLRS Linda Weinmann, HLRS Sebastian Zeeden, HLRS heusel@hlrs.de weinmann@hlrs.de zeeden@hlrs.de

Electrodialysis is used as an energy and cost efficient technique for sea water desalination. It uses selective ion exchange membranes to remove salt ions from sea water under the influence of an externally applied electrical field. A complex spacer structure keeps the membranes apart from each other. The figure shows the species transport in the spacer-filled channel with stream lines showing mixture flow direction. The goal of the HISEEM project, funded by BMBF, is to develop a large scale numerical simulation tool, coupling the flow around the spacer, multi-species transport and interaction, electrical fields, and membrane transport to gain a better understanding of the complex physical phenomena involved in this process.

[D. Harlacher, K. Masilamani, STS, Universität Siegen, daniel.harlacher@uni-siegen.de, kannan.masilamani@uni-siegen.de]

# Editorial

Welcome to this new issue of inside, the journal on Innovative Supercomputing in Germany published by the Gauss Centre for Supercomputing (GCS). In this issue, there is a strong focus on applications but we are also able to present a new GCS system at the HLRS at Stuttgart. The Cray XC4O code named Hornet will quadruple the peak performance and main memory of HLRS systems. A new system scheduled for 2015 at the LRZ will be presented in the next issue. The growing relevance of accelerators is reflected by the installation of a new Intel phi system at LRZ in Garching.

GCS has been a strong player in PRACE over the last years. In 2015 it will continue to provide access to leading edge systems for all European researchers through the evaluation process of PRACE. In this issue, we present deadlines and results of PRACE projects calls in 2014 and 2015. The variety of nations and projects is impressive and encouraging for the future work of PRACE. The concept of dedicated centers for research is working well both for the user community and the funding agencies. Continuation of PRACE is under discussion.

HPC continues to be an important issue in the German research community. The significance of HPC for German science is emphasized by a paper released by the Deutsche Wissenschaftsrat (German Science Council) on July 14 this year (Drs. 4032-14 at http://www. wissenschaftsrat.de/presse/pressemitteilungen/2014/nummer\_22\_ vom\_14\_juli\_2014.html) covering the relevance and future development of simulation for science in Germany. Software initiatives like the HPC Software Initiative of the Federal Ministry of Science (BMBF) or the SPPEXA (Special Program for Software for Exascale Computing) initiative of the German Research Society (DFG) start to create interesting results as shown in our project section.

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 which are now part of the PRACE Advanced Training Center and hence open to all European researchers.



- Prof. Dr. A. Bode
  (LRZ)
- Prof. Dr. Dr. Th. Lippert (JSC)
- Prof. Dr.-Ing. Dr. h.c. Dr. h.c. M. M. Resch (HLRS)

# **PRACE:** Results of the **9th Regular Call**

The Partnership for Advanced Computing in Europe (PRACE) is continuously offering supercomputing resources on the highest level (tier-O) to European researchers.

The Gauss Centre for Supercomputing (GCS) is currently dedicating shares of its IBM iDataPlex system SuperMUC in Garching, of its Cray XE6 system Hermit in Stuttgart, and of its IBM Blue Gene/Q system JUQUEEN in Jülich. France, Italy, and Spain are dedicating shares on their systems CURIE, hosted by GENCI at CEA-TGCC in Bruyères-Le-Châtel, FERMI, hosted by CINECA in Casalecchio di Reno, and MareNostrum, hosted by BSC in Barcelona.

The 9th call for proposals for computing time for a 12 month allocation time period beginning September 2014 on the above systems closed on 25 March 2014. Seven research projects have been granted a total of about 120 million compute core hours on Hermit, six have been granted a total of about 170 million compute core hours on Super-MUC, and three proposals have been granted a total of 95 million compute

PRACE

core hours on JUQUEEN. One of those research projects has been granted resources both on Hermit and Super-MUC.

Three of the newly awarded research projects are from France, two are from Germany, Italy, Spain, and Switzerland, each, and one is from Cyprus, Finland, the Netherlands, and Portugal, each. The research projects awarded computing time cover again many scientific areas, from Fundamental Physics to Medicine and Life Sciences to Universe Sciences. More details, also on the projects granted access to the machines in France, Italy, and Spain, can be found via the PRACE page www.prace-ri.eu/ PRACE-9th-Regular-Call.

The 10th call for proposals for a 12 month allocation time period starting March 2015 closed on 22 October 2014 and evaluation is still under way, as of this writing. The 11th call for proposals is exptected to open in February 2015.

Details on calls can be found on www.prace-ri.eu/Call-Announcements

# HLRS: Extending the Simulation Capacities and **Capabilities for Science and** Industry

After the installation and start of operation of its PetaFlop supercomputer Hermit in late 2011, the High Performance Computing Center Stuttgart (HLRS) recently entered installation of phase 2 of its HPC systems roadmap with the deployment of a Cray XC40 supercomputer, code named "Hornet". Hornet is destined to replace the former HLRS flagship computer Hermit, a Cray XE6 machine, which at the time of its launch featured the fastest civil and industrially used supercomputer in all of Europe.

In the configuration of this first installation step of Installation Phase Two, Hornet will deliver a peak performance of 4 PetaFlops, outperforming the maximum performance of Hermit by a factor of





Figure 1: Frontdoor Image Medicine

about four. Similar to its predecessor, Hornet has a clear focus on sustained performance of real-world applications as they are currently running on the HLRS platforms. The new HPC infrastructure will significantly extend the simulation capacities and capabilities at HLRS, supporting numerous research projects from scientific fields such as e.g. materials science, scientific engineering, life sciences, environment, energy and health as well as elementary particle physics and astrophysics.

On August 13th, 2014 the CrayXC40 system arrived at the HLRS facilities in Stuttgart and the installation commenced. Hornet had its technical operability tested in September and is now fully operational. An operational overlap



Figure 1: Frontdoor Image RECOM

phase of both HPC systems Hermit and Hornet for approximately three months is currently intended, allowing users a trouble-free migration of their applications to the new HLRS supercomputer. As HLRS had earlier installed a small transition system, the smooth move of applications is well prepared.

HLRS's new Cray XC4O system is based on the Intel Haswell processor with 12 cores each and the Cray Aries network. After completion of the first step of the current installation phase, Hornet will consist of 21 cabinets hosting a total of 3,944 compute nodes. The system's main memory capacity is 128 GB per node. This will sum up to a total of 94,656 cores with a main memory of 493 Terabyte. Hornet will provide 5.4 Petabyte of file storage to the end users with an I/O speed in the range of 150 GB/s.

In a scheduled follow-on step, which is due to be completed in 2015 ("Installation Phase 2, Step 2"), Hornet will be expanded by 2.3 Petabyte of additional file storage and by 20 more cabinets, boosting the system's expected peak performance to then over 7 PetaFlops. The installation of the new Cray supercomputer at the HLRS is carried out according to the HPC systems roadmap as defined by "Project PetaGCS", a project of the national Gauss Centre for Supercomputing (GCS) which was commenced in 2008. The goal of Project PetaGCS is to deploy petascale HPC systems in all of the three national GCS centre locations LRZ Garching, JSC Jülich, and HLRS Stuttgart. Project PetaGCS is supported by the Federal Ministry of Education and Research (BMBF) and in the case of HLRS by the Ministry of Higher Education, Research and Arts Baden-Württemberg.

contact: Bastian Koller, koller@hlrs.de

# From Dust to Planetesimals – High Resolution Simulations of Planet Formation Processes

Planet formation is a beneficial side effect of star formation. When in our milky way a large gas cloud collapses under its own weight to create a star in its center, which currently happens about once per year in our galaxy, then some of the material stays in an orbit around the newborn star. This is simply a consequence of angular momentum conservation of the turbulent gas that the star has formed from. This gas cannot fall directly onto the star, but forms a disk with the star in its center. These disks can be observationally found around most young stars that are younger than about 10 million years. The dimensions of these disks are slightly larger than our solar system and not too surprisingly we believe these disks to be the origin of planetary systems like ours. They contain 98% of Hydrogen and Helium, the two most abundant elements in the universe, and 2% heavier elements that we need



Figure 1: Dust Concentration in zonal flows. Magneto Hydro Simulation of a small patch of an accretion disk around a young star as seen from the top, thus right is the distance from the star and vertical the azimuthal stream-wise direction in the disk. Color indicates the local gas pressure where red is larger than blue. Particles are given by black dots. They preferentially concentrate in regions of higher pressure – e.g. the zonal flows.

Bastian Koller

University of Stuttgart (HLRS)





to create planets and everything that lives on them from Carbon, Nitrogen and Oxygen over Iron to even Uranium. Astronomers call all these elements heavier than Hydrogen and Helium simply metals. The initial abundance of these metals is initially on the order of 10<sup>-10</sup> kg/m<sup>3</sup> or 13 orders of magnitude lower than typical densities of a planet. Thus it becomes obvious that the planet formation process must be a very efficient mechanism that concentrates the metals locally.

The first step is the formation of dust and ice flakes out of the metals, where interestingly water ice was indeed the most abundant species of small solid objects, which are simply called dust grains by the astronomers. These grains undergo mutual collisions driven by Brownian motion, differential settling and turbulence induced relative velocities that lead to growth via sticking up to typically centimeters in size. Larger sizes are not possible because the collision velocities become so large that they rather lead to destruction of the grains than to sticking. These first growth steps from dust to planets can be detected in disks around

young stars via the observation of the spectral energy density, e.g. the color of thermal radiation emitted from the disk. Larger grains are more efficient to radiate at larger wavelength than smaller grains. What follows is unfortunately not directly observable, because meter or even kilometer sized objects are not detectable around young stars. Yet the process must be very efficient because with having more than 1.000 planets found around stars other than the sun we come to believe that basically all stars come with a smaller or larger planetary system.

# **Planetesimals – Planetary Building Bricks**

From our own solar system we know that at some stage so called planetesimals must have formed, which are 100 km sized planetary building bricks made from rock and ice. The entire solar nebula must have been full of those objects that were then merging into the cores of the gas giants like Jupiter and Saturn or into the terrestrial planets like Earth and Venus. Leftovers from this huge population of planetesimals can still be found today as asteroids and comets. Thus, we would have a relatively clear

path from dust in a disk around a young star to a planetary system, if we can explain the formation of 100 km sized planetesimals from centimeter sized gravel. As mentioned above all pure hit and stick mechanisms have failed so far. Also, a direct gravitational collapse of the dust content of the disk was ruled out, because at no time the disk was laminar enough, e.g. not-turbulent, to allow for a sufficient settling of the grains and their concentration in the mid-plane. Either the turbulence comes from large-scale magnetic fields in the disk or from hydrodynamic instabilities tapping in the vertical shear of the disk or in the radial entropy gradient. If all direct mechanisms would fail - and we know that at least some of them must work from the observation of disks - then the dust itself would trigger instabilities and turbulence in the dust and gas mixture as soon as the dust is at least as abundant as the gas. This abundance is unfortunately lower than the local concentration needed for selfgravity to take over (Roche Density) in



# **Gravoturbulent Planetesimal** Formation

In the last years we have developed a model of planetesimal formation that not only overcomes all these obstacles, it also starts to allow for the making of predictions what sizes of planetesimals should form and what the formation rate should be as function of time in the evolution of the stars and the distance from them [1]. Our models are heavily leaning on the simulations that we performed and still perform on the Jülich supercomputers JUGENE and later JUQUEEN. In these models we start with high-resolution magnetohydro-dynamical simulations of small patches of the disks around young stars and study the evolution and properties of turbulence driven from the magneto rotational instability. At the same time we simulate the dynamics of millions of test particles embedded in the turbulence that interact with the



Figure 3: A series of snapshots following the collapse of a dense particle heap. Colors indicate the particle density. The initial spherical clouds gets partly sheared apart and then collapses in a group of many planetesimals of typically 100 km in size.



a rotating and shearing system and thus no planetesimals can form.

gas via friction. The magnetic fields are too weak to have an impact on the typically only very low charged grains. The turbulence has now two effects on the grains. First, it mixes and diffuses them on the large scale and second, it concentrates them on the small scales. This can be either due to centrifugal acceleration in turbulent eddies or trapping of grains in long-lived pressure maxima. The trapping in pressure maxima is an effect of the mismatch in particle and gas dynamics. Gas feels a pressure gradient, whereas dust does not. Therefore, the turbulent gas around a star does in general not move at the Keplerian speed, but mostly slower (due to the systematic radial pressure gradient) and only sometimes faster. This leads to head- and tailwind for the dust and in order to minimize these effects the dust concentrates at the few random locations where the disk rotates at the Keplerian rate (see Fig. 1) [2]. This can occur in socalled zonal flows, a side product of the magnetic turbulence, or in large-scale

vortices like Jupiter's red spot or the high-pressure anti-cyclones in the earth atmosphere.

# A Dust-Gas Streaming Instability

Once the concentration in these dust traps reaches a dust to gas mass ratio of 1 then a streaming instability sets in. Here the dust not only feels the friction from the gas, but also the gas feels a significant back reaction. The higher the dust to gas ratio is, the closer the gas rotation velocity gets to the Keplerian rate. Thus in regions which have not been perfectly Keplerian beforehand, the dust tends to clump and reach densities that become larger than the Roche density, at which the tidal forces from the central star are no longer sufficient to prevent the direct gravitational collapse of the particle heap into a 100-1.000 km sized planitesimal (see Fig. 4). The necessary physics for this process are also incorporated into our simulations, as we simultaneously solve the Poisson



Figure 4: Artist impression of a planetary embryo within a protoplanetary disk. These so called planetesimals form out of dust by gravitational collapse and are 100 to 1.000 km in size. After their formation they grow rapidly by acceding the surrounding material.

equation for the dust and gas mixture and apply the resulting forces on gas and particles. Fig. 2 shows a snapshot of the particle distribution after collapse, where the white circle indicates the gravitational bound region around the formed planetesimals.

# Numerical Code

Our simulations use the Pencil Code [3], which contains a high order finite difference magneto hydro dynamics code that uses Runge Kutta timestepping for stability. This automatic stability was a prerequisite for us to implement the additional physics like particle feedback onto the gas, which usually results in a set of stiff differential equations. Particles are treated as individual Lagrangian point masses and self-gravity is solved via a Fast Fourier method. Our production runs use up to 512<sup>3</sup> grid-cells and 64 Mio. particles [4]. The simulations have to run for several hundreds of dynamical time scales, which is millions of computational steps before planetesimals start to form, because the concentrations develop on viscous time scales that are much longer then the dynamical time scales.

# **Future Directions**

Currently we push our simulations in two directions. On the one hand we want to understand the efficiency of the process, like knowing what the minimal concentration and size of dust might be to be converted into planetesimals, and on the other hand to derive the initial mass function of planetesimals. The first goal is treated by better and better modelling the turbulence in disks and the second by performing high-resolution studies of the collapsing and possibly fragmenting particle clouds (see Fig. 3).

If both results are in hand we can put them into semi-analytic yet global evolution codes of dusty disks and produce predictions when and where which sizes of planetesimals should form in disks around a young star, as a fundamental step in the formation of our solar system as well as for the many planetary systems around other stars.

# Acknowledgements

This project has received funding from the Deutsche Forschungsgemeinschaft within the Schwerpunktprogramm (DFG SPP) 1385 "The first ten million years of the solar system". The authors gratefully acknowledge the Gauss Centre for Supercomputing (GCS) for providing computing time for a GCS Large Scale Project on the GCS share of the supercomputer JUQUEEN at Jülich Supercomputing Centre (JSC).

# References

- [1] Johansen et al.

- 2011

contact: Hubert Klahr klahr@mpia.de



nature, 448, 1022, 2007

[2] Dittrich, K., Klahr, H., Johansen, A. Astrophysical Journal, 763, 117, 2013

[3] https://code.google.com/p/pencil-code/

[4] Johansen, A., Klahr, H., Henning, Th. Astronomy and Astrophysics, 529, 62,

- Hubert Klahr
- Andreas Schreiber

Max-Planck-Institut für Astronomie

# New Design Principles for Biomimetic Mechanical Contacts by Computer Simulation

An important function of lubricants is to keep two surfaces from touching directly so that not only friction but also undesired wear is minimal. This is not easy to achieve, because even highly polished surfaces are rough at the microscopic scale so that most fluids are squeezed out quickly when asperities from opposite surfaces approach each other. Nature and technology have so far pursued different avenues to keep a lubricating fluid within the contact [1]: commercial lubricants are based on oils that quickly increase their viscosity when the local pressure goes up. This gives oils the ability to remain between two colliding

asperities and to sustain locally a normal load over a non-negligible time. In contrast to this solution, biological lubrication is based on aqueous liquids that maintain highly fluid at all times. The lubricants yet remain in the contact, because surfaces in biological joints, specifically their cartilage, carry long, end-anchored sugar chains. These polymers are strongly hydrophilic and thus exert a large osmotic pressure up to a few 100 Mega Pascal on the aqueous lubricants giving them the ability to sustain large normal pressures in a sliding contact while remaining slippery.





Porting the principles of bio-lubrication into technological applications has been pursued for many decades [2], albeit with little success. The reason is that end-anchored polymers whose chain termini and loops penetrate the opposing "polymer brush" are prone to scission and detachment [3]. A biological system can regrow the lost endanchored chains. However, in technical systems only oligomers, i.e., relatively short chains with chemically active head groups can be dissolved as additives that replenish lost material on rubbing surfaces [4]. These molecules are not long enough to form soft, solvated polymer brushes.

To avoid chain scission and detachment in a biomimetic frictional contact, we recently proposed to decorate one surface with hydrophilic polymers and the counter surface with hydrophobic polymers [5]. If the lubricant is a mixture of two liquids so that each liquid dissolves its own favorite polymer, a fluid-fluid interface will be created. The idea is that polymers belonging to one brush cannot pass through the fluidfluid interface so that chain scission and detachment no longer occur. Here, we review computer simulations [5,6] in which we examined if the envisioned lubrication mechanism can be achieved in practical applications. We also present comparisons to experiment, which were motivated by the computer simulations.

# Simulation Method

To mimic the thermodynamic and non-equilibrium behaviour of polymers, it has become common practice to simulate coarse-grained models of polymers with molecular dynamics (MD) methods. Kremer and Grest [7] proposed one of the most successful coarse-grained polymer models more than two decades ago.



Figure 2: Friction coefficient as a function of velocity for a symmetric (orange) and an asymmetric (blue) brush contact for a tip on a surface geometry. The inset shows energy losses that occur in asperity collision mode, i.e., when the cylinders are moved in the transverse direction. Adopted from Ref. [5].



It represents polymers as bead-spring chains, in which the springs are finitely extensible and nonlinearly elastic. The parameters are designed such that crossing of two chains is extremely unlikely, even though the springs remain sufficiently flexible to allow for reasonably large MD time steps. In addition to the elastic potential between two adjacent beads in a chain, Lennard Jones (LJ) interactions act between all monomers as long as they are within a predefined interaction radius. In our simulations, LJ interactions also act between polymers and solvent particles as well as between different solvent particles. We refer to the literature for details and content ourselves with stating that the fine-tuning of LJ parameters allows one to determine if two fluids or polymer systems want to mix or to segregate. As such, the employed model has successfully reproduced the generic tribological behaviour of adsorbed [8] and end-tethered [9] hydrocarbon films.

Specific to our simulations is the use of an explicit solvent and curved surfaces. We found this to be necessary for mainly two reasons: First, the lubrication mechanism that we attempt to simulate lives from polymers not wanting to pass through a fluid-fluid interface and there would be no such repelling interface without explicit solvents. Second, in contrast to the long-held believe that friction between polymer brushes is dominated by the direct or fluid-mediated interactions occurring due to brush overlap [9], we noticed that many other dissipation mechanisms are relevant for polymers grafted to more realistic, that is, rough surfaces [10]: these are, in particular, viscoelastic deformation of the brushes in non-conformal contact geometries, squeeze-out/re-adsorption dynamics of the solvent, and

capillary hysteresis for partially solvated brushes. A setup characteristic for our simulations is shown in Fig. 1. Typical system sizes comprise 0.5 million polymer beads and a similar number of solvent particles. Simulations were run with the LAMMPS software package [11], which is a highly scalable molecular dynamics package.

# Results

In a first study [10], we demonstrated that there is no universal frictionvelocity dependence *F(v)* for polymer brush systems, although all systems showed a power-law dependence  $F \sim v^{\kappa}$ . In particular, we found the exponent  $\kappa$ to depend on the sliding direction. For example, when moving parallel to the symmetry direction in Fig. 1, the exponent turned out to be  $\kappa$ =0.57. This is in agreement with previous simulations [12] and experiments [13] employing effectively plane-plate geometries. In this case, energy losses can be rationalized as dissipation resulting from the overlap between two opposing brushes [12]. For motion in the two orthogonal directions, by which we mimic asperity collisions, we found smaller exponents, that - unlike for the motion in the symmetry direction - varied with the details of the set-up, i.e., the precise exponents depended on the degree of polymerization, the amount of solvent, and force-field parameters. The reason why exponents do not have to be universal is that different dissipation mechanisms mix and that there is no clear time-scale separation between different types of dissipation processes (viscoelastic relaxation, capillary formation, and interdigitation can all have a broad and overlapping distribution of relaxation times). Since we found a relatively small shearthinning exponent  $\kappa$  for motion in the normal and transverse directions, one



Figure 3: Set up of the experimental system in the symmetric case, (a), in which both sides are hydrophobic (acetophenone-solvated PMMA) and the asymmetric case, (b), where one side is made hydrophilic (water-immersed PNIPAM). The sphere represents a colloid, which carried the hydrophobic brush while being glued to an atomic-force microscope cantilever. Friction traces, normalized to the steady-state friction of the symmetric contact are shown in (c) and (d) for the symmetric and asymmetric system. Simulation and experiments show the same trends when represented in terms of appropriate dimensionless variables. Snapshots of the simulations are shown in (e) and (f). They reveal that polymer brushes overlap in the symmetric case, while a slick fluid-fluid interface is formed when the two brush systems are immiscible. Adopted from Ref. [5].

must expect (technically relevant) smallvelocity friction between rough surfaces to be dominated by processes other than those related to brush overlap.

It remains to be investigated if it is possible to reduce friction and wear between solids by decorating them with two mutually immiscible polymer brush systems. To answer this question we conducted simulations of two polymer brush systems, which we designed such that one surface was decorated with polymers of type "P", which was soluble in a solvent "S" but not in another solvent " $\overline{S}$ ", while the other surface carried end-anchored polymers of type  $\overline{P}$ , which were soluble in  $\overline{S}$  but not in S.

Moreover S and  $\overline{S}$  were immiscible. In all simulations that were set-up according to this scheme, we found friction to be much reduced in comparison to those where the components were mutually miscible, except at extremely large sliding velocities, where solubility no longer plays a prominent role. At the smallest velocities that we could approach in the simulations (which would roughly correspond to 1 mm/s in real-time units), the friction reduction was as big as two orders of magnitude. Results for the amount of dissipation are shown in Fig. 2. The important insight is that the exponent, with which friction becomes smaller as the sliding velocity goes down, is noticeably larger for the

asymmetric case than the one typical for overlapping brushes, revealed by the asymmetric brushes. Moreover, when going into "asperity collision mode", that is, when the cylinders are moved in the transverse direction, dissipation is again much smaller in the asymmetric than in the symmetric system. This time, however, the exponents appear to be identical. Since the time is too short for the polymers to interdigitate, even for the symmetric system, we attribute the energy losses in collision mode to be dominated by viscoelastic deformation. Interestingly, the exponent is close to that of "frictionby-interdigitation" so that it might not be possible to discriminate between the two friction mechanisms experimentally in a straightforward fashion.

To confirm the validity of our simulations, we conducted experiments mimicking the set-up of our simulations [5]. Polymers were end-anchored to a substrate as well as to a colloid, which itself was glued to an atomic-force-microscopy cantilever. The experimental system and typical force traces are presented in Fig. 3. By choosing an appropriate unit system, experimental and simulation results can almost be superimposed. This is also true for measurements not shown here explicitly, such as the normal motion of the tip after velocity inversion. We thus see it as legitimate to exploit the data obtained in MD simulations to rationalize not only the simulations themselves but also the experiments.

We can use simulations to investigate a broader range of sliding velocities and normal pressures than possible with current experimental techniques. In addition, one can investigate to what degree friction reduction persists if polymers are fully solvated and the two

solvents are miscible with each other while retaining a clear preference for one of the two brushes. A two-component miscible solvent would be preferential from a practical point of view. In Ref. [6], we demonstrate that one still has a well-defined fluid-fluid interface in those situations even if the friction is no longer reduced quite as much as in the case of partially solvated brushes, in which the two solvents are immiscible. The important aspect remains the formation of a sharp interface that prevents chain termini and loops of one side to penetrate the other side, as this prevents the scission and pulling out of polymers, thereby prolonging the lifetime of the sliding bodies. Preliminary experimental results [5] indicate a reduction of not only friction but also wear when a conventional symmetric brush system is replaced with an asymmetric one.

# Outlook

While our MD simulations have very successfully predicted that friction and wear can be strongly reduced when a conventional, i.e., symmetric brush interface (found, for example, in biological joints) is converted into one where opposing brushes are immiscible, quite a few unanswered questions persist. In particular, two issues remain challenging to address: First, there are some crucial experimental differences in the rheological response of hydrophobic and hydrophilic, in particular, zwitterion-based brush systems, whose explanation certainly requires one to go beyond generic bead-spring/Lennard-Jones models. To describe zwitterion-based brushes realistically, it might be necessary to embed the properties of real water, including its ability to participate in proton transfer reactions, into the force fields. Second, simulations of the dissipation of brushes always indicate that a power-law relation

between friction and velocity is followed by a linear response when the shear rate is decreased to extremely small values. Experimentally, friction often becomes a logarithmic function of velocity at very small v, akin of solid or Coulomb friction. To obtain such a velocity-dependence, there must be instabilities at very small length scales. The nature of such instabilities remains to be identified, before simulations can assist a material-specific design of a frictional system.

# Acknowledgement

The authors gratefully acknowledge the Gauss Centre for Supercomputing (GCS) for providing computing time through the John von Neumann Institute for Computing (NIC) on the GCS share of the supercomputer JUQUEEN at Jülich Supercomputing Centre (JSC).

# References

- [1] Lee, S., Spencer, N. D. Science 319, 575-576, 2008
- [2] Klein, J., Kumacheva, E., Mahalu, D., Perahia, D., Fetters, L. J. Nature 370, 634, 1994
- [3] Maeda, N., Chen, N., Tirrell, M., Israelachvili, J. N. Science 297, 379, 2002
- [4] Mang, T., Dresel, W. Lubricants and Lubrication, Wiley-VCH, 2007
- [5] de Beer, S. et al. Nat. Commun. 5, 3781, 2014
- [6] de Beer, S., Müser, M. H. Macromolecules (submitted)
- [7] Kremer, K., Grest, G. S. Chem. Phys. 92, 5057, 1990
- [8] He, G., Müser, M., Robbins, H. Science 284, 1650, 1999
- [9] Binder, K., Kreer, T., Milchev, A. Soft Matter 7, 7159, 2011

[11] Plimpton, S. J Comp Phys, 117, 1 1995

[12] Galuschko, A. et al.

[13] Schorr, P.A. et al. Macromolecules, 36, 389, 2003

contact: Martin H. Müser, m.mueser@fz-juelich.de Sissi de Beer, s.j.a.debeer@utwente.nl



[10] de Beer, S., Müser, M. H. Soft Matter 9, 7234, 2013

Langmuir 26, 6418, 2010

- Martin H. Müser<sup>1</sup>
- Sissi de Beer<sup>1,2</sup>
- <sup>1</sup> NIC Research Group Computational Materials **Physics** Institute for Advanced Simulation Forschungszentrum Jülich
- <sup>2</sup> Institute of Nanotechnology University of Twente

# Large-Scale Order and Small-Scale Intermittency in Turbulent Convection

Thermal convection can be found in numerous flows in Nature and technology. Examples reach from chip cooling devices and heat exchangers via air conditioning in passenger aircraft cabins to natural flow phenomena such as atmospheric clouds or astrophysical turbulence inside stars [1]. In its simplest setting the classical Rayleigh-Bénard convection setup - turbulent convection evolves in a fluid layer between two isothermal parallel plates. Turbulent Rayleigh-Bénard (RB) convection is considered as a paradigm for these applications of thermal convection. A deeper understanding of the local and global mechanisms of heat and momentum transfer, their connection to the large-scale structures which form in the flow and the small-scale statistical properties in thermal convection can be obtained from massively parallel direct numerical simulations (DNS). Such simulations fully resolve the turbulent RB convection on all scales and do not require turbulence models or subgridscale parametrizations.

Laboratory experiments and numerical simulations are carried out frequently in closed cylindrical cells with insulated side walls. Three system parameters are then relevant, the Rayleigh number Ra, a measure of the thermal driving of turbulence, the Prandtl number Pr, a measure which relates viscous and thermal diffusion, and the cell aspect ratio  $\Gamma$  which relates cell diameter to cell height. The equations of motion – the Boussinesq equations which couple the velocity and temperature fields – are discretized in space and time. We apply either finite difference (FDM) or spectral

element methods (SEM) for our DNS studies. The present spectral element simulations are based on the nek5000 software package [2] which was optimized to our particular problem. For more details and a comparison of both numerical methods we refer to [3,4].

Here we focus on two aspects of thermal convection, the large-scale flow structure (or order) and the small-scale intermittency. The first topic requires large aspect ratio cells. Since the numerical effort grows with  $\Gamma^2$ , this reduces the accessible Rayleigh number significantly. Such simulations allow us to investigate pattern formation in extended turbulent flows, a phenomenon which can be observed frequently in nature, e.g., cloud streets in the atmosphere or supergranules on the solar surface. These DNS studies are carried out with FDM. The second topic requires higher Rayleigh numbers which can only be obtained in cells with smaller aspect ratio. For example, the high-resolution DNS for the maximum Rayleigh number of Ra=10<sup>10</sup> and F=1 required two racks on the BG/Q supercomputer JUQUEEN, i.e. 65.536 MPI tasks. Spectral accuracy is necessary to study small-scale turbulence, in particular the intermittent turbulence statistics of the gradients of velocity and temperature which will be discussed later. Therefore, an SEM is applied. Fig. 1 illustrates both regimes in the parameter plane that is spanned by the Rayleigh number and the aspect ratio at a given Prandtl number.

# Large-Scale Order in Convection

Figure 1 displays the threshold at which convection sets in as a primary linear instability of the quiescent fluid layer across which heat is transferred by molecular diffusion. In a large-aspect ratio cell with solid top and bottom walls it is found at Ra, =1708 (dashed line in Fig. 1). For smaller aspect ratios this critical Rayleigh number grows as approximately indicated in the figure (see e.g. [5] for more details). A quasisteady laminar flow pattern is shown at Ra=1800 for a cell with  $\Gamma$ =50 (red framed picture). The typical convection rolls are identified in a streamline plot. For higher Ra, the system switches into a state of spiral defect chaos before the flow becomes fully timedependent and eventually turbulent at Rayleigh numbers of about 10<sup>5</sup> and

rotating concentric cylinders.



Figure 1: Large-scale flow for different setups in the parameter plane which is spanned by the Rayleigh number Ra and the aspect ratio F. In all three cases a view from the top onto the cylindrical convection cell is displayed. Red: convection slightly above the onset (Ra=1.800, Pr=0.7, F=50). Orange: soft turbulent convection regime (Ra=5×10<sup>5</sup>, Pr=0.7, Γ=50). Yellow: hard turbulent convection regime (Ra=3×10<sup>9</sup>, Pr=0.7, Γ=1). In the latter two cases, time averages over 100 convective time units are shown. The dashed line is the linear instability threshold in an infinite fluid layer at Ra\_=1.708.

beyond. The orange framed box in Fig. 1 displays a time-averaged streamline plot at Ra=5×10<sup>5</sup>, in the soft turbulent convection regime. Interestingly, it can be observed that the spiral roll patterns survive into the turbulent regime when the turbulent fluctuations are removed. A large-scale order persists far into the turbulent regime. These rolls carry a significant amount of the heat from the bottom to the top and are thus important for the global transport properties [6]. When the Rayleigh number grows further towards the hard convection regime, the temporal dynamics and evolution of these largescale patterns remains an open question. Interestingly, similar observations are reported in other systems, such as the Taylor-Couette flow between two



Figure 2: Locally seeded field lines of the convective heat current density illustrate the instantaneous pathways of heat transfer by turbulent fluid motion. The bottom contour slice displays the corresponding temperature field close to the bottom plate. Color coding stands for the vector magnitude (blue for zero, red for maximum). Field lines are colored red close to the sheet-like thermal plumes, but not necessarily exactly at the same location.

> In cells with aspect ratio of unity, one still observes a single circulation roll in the turbulent regime after timeaveraging. This is shown for a much larger Rayleigh number of Ra=3×10<sup>9</sup> (see yellow box in Fig. 1). Our DNS demonstrated that this three-dimensional large-scale flow evolves on a slow time scale [4]. It is sustained by the ongoing detachment of thermal plumes, fragments of the thin thermal boundary layers at the top and bottom plates. These plumes are the characteristic near-wall structures which carry the heat into the bulk of the cell and further up to the cold top plate. In Fig. 2 we display the lower half of a convection cell and plot field lines of the convective heat current density which is given by the product of the velocity field and the temperature

$$\vec{j}_{conv} = \vec{u} T$$

The figure displays that the method of heat transfer through the cell is rather complicated and is not necessarily synchronized with the thermal plumes which are usually identified as the hot (cold) ridges at the bottom (top) plate (see brightest areas in the contour slice of Fig. 2).

# Small-Scale Intermittency in Turbulent Convection

Convection is one particular way to generate fluid turbulence. One main question of turbulence research is if the statistical fluctuations of the turbulent fields are universal at the smaller scales. It is expected that the turbulent flow forgets the particular method of driving when the statistics at scales sufficiently below the driving scale is investigated. Such a sufficiently large scale separation usually requires large Reynolds (or Rayleigh) numbers. In [7] we suggested an alternative perspective: rather than investigating the statistics of the turbulent fields in the inertial cascade range, we study the statistics of the derivatives of the turbulent fields in the crossover to the viscous range. It turns out that these studies can be done at smaller Reynolds (or Rayleigh) numbers. This does unfortunately not mean that the simulations become less comprehensive. A correct analysis of the derivative statistics requires a very high spectral resolution which makes these studies also demanding in terms of computational resources.

The kinetic energy dissipation rate field is proportional to the magnitude of the velocity gradient tensor and the fluctuations of this kinetic energy dissipation rate are of central interest. In Fig. 3 we display a horizontal cut through this field and observe the typical spatial arrangement in the form of layers of enhanced shear.

Based on data obtained from our previous DNS and from a GCS Large Scale Project in the past year, we could show that moments of the kinetic energy dissipation rate follow the same scaling with respect to the flow Reynolds number in three turbulent flows of different complexity. This universal scaling is already detected for Reynolds numbers Re as small as few hundred and beyond. The three flows are: (i) homogeneous isotropic turbulence in a cube with periodic boundary conditions and three directions of statistical homogeneity, (ii) pressure gradient-driven channel flow turbulence with statistical homogeneity in two horizontal directions and (iii) turbulence in a closed cylindrical RB convection cell with azimuthal statistical homogeneity. This scaling with respect to Re is theoretically predicted for homogeneous isotropic turbulence and was also found in the bulk of the other two flows – a demonstration of the small-scale universality of fluid turbulence. Furthermore, we detected in the same range of Re



Figure 3: Snapshot of a horizontal slice cut through the kinetic energy dissipation rate field. The mid-plane of the cylindrical cell is displayed. Contours are plotted in logarithmic units. Spatial intermittency is visible as the inhomogeneous distribution of the high-amplitude shear layers. Data are obtained for Ra=10<sup>10</sup>, Pr=0.7 and F=1.

a transition of the derivative statistics from Gaussian or nearly-Gaussian to intermittent behavior. Current efforts focus on a better removal of effects of the large-scale flow, particularly in the RB setup (see Fig. 1).

# Outlook

Turbulent convection has been and still remains a vital field of turbulence research. We have discussed two aspects of our ongoing studies. Most experimental and numerical investigations (including our own) have been done in gases or water, i.e., in flows at Prandtl numbers around unity or above. We consider RB convection at very low Prandtl numbers as a real challenge. Applications of convection in this parameter regime reach from nuclear engineering, to the dynamo problem and planetary science. Laboratory experiments in this regime require the use of opaque liquid metals (e.g. gallium, mercury or liquid sodium) which are accessed by ultrasound or X-ray methods. High-resolution supercomputations will help to shed new light on this Terra Incognita in the RB convection map.

Finally, we wish to thank the Jülich Supercomputing Centre for their ongoing support of our research work in the field of turbulence.

# References

- [1] Chillà, F., Schumacher, J.
  - Colloquium: New perspectives in turbulent Rayleigh-Bénard convection, Eur. J. Phys. E 35, 58, 2012
- [2] http://nek5000.mcs.anl.gov
- [3] Scheel, J. D., Emran, M. S., Schumacher, J.

Resolving the fine-scale structure in turbulent Rayleigh-Bénard convection, New. J. Phys. 15, 113063, 2013

- [4] Shi, N., Emran, M. S., Schumacher, J. Boundary layer structure in turbulent Rayleigh-Bénard convection, J. Fluid Mech. 706, 5, 2012
- [5] Hébert, F., Hufschmid, R., Scheel, J. D., Ahlers. G.

Onset of Rayleigh-Bénard convection in cylindrical containers, Phys. Rev. E 81, 046318, 2010

[6] Bailon-Cuba, J., Emran, M. S., Schumacher, J.

Aspect-ratio dependence of heat transfer and large-scale flow in turbulent convection, J. Fluid Mech. 655, 152, 2010

[7] Schumacher, J., Scheel, J. D., Krasnov. D., Donzis, D. A., Yakhot, V., Sreenivasan, K. R. Small-scale universality in fluid turbulence,

Proc. Nat. Acad. Sci. USA, 111, 10961, 2014

contact: Jörg Schumacher, joerg.schumacher@tu-ilmenau.de

# Unsteady CFD for **Automotive Wheel** Aerodynamics



Figure 1: Streamlines and flow slices at the left front wheel colored by total pressure.

In the pursuit of reducing CO<sub>2</sub> emissions and extending ranges of electric vehicles automotive engineers have to minimize all driving resistances. Since aerodynamic drag, which increases with the square of velocity, becomes the dominant force at velocities higher than 70 km/h, aerodynamic optimization has gained more importance recently. In that context, the contribution of wheels and wheel houses to a passenger car's total drag amounts to approx. 25 percent, a large share considering the geometric dimensions. Consequently, industry and academia have increased their research activities in the field of automotive wheel aerodynamics. For

accessibility reasons the complex flow field around rotating wheels and inside wheel houses is, however, difficult to investigate experimentally and wind tunnel resources are usually very limited. Hence, Computational Fluid Dynamics (CFD) can be a valuable, complementary tool to improve the understanding of the flow phenomena in that region. Especially the ability to run unsteady simulations including rotating geometries has improved the physical accuracy of CFD, thus making it suitable for wheel aerodynamics.

In a joint research project of TU München and BMW Group the flow around wheels is investigated experimentally, in the

Mohammes S.

• Janet D. Scheel<sup>2</sup>

Jörg Schumacher<sup>1</sup>

<sup>2</sup> Occidental College

Los Angeles

Emran<sup>1</sup>

<sup>1</sup> Technische

Ilmenau

Universität

wind tunnel and on the road, and numerically, using the commercial Lattice-Boltzmann solver Exa PowerFLOW on the computational resources of Super-MUC at LRZ. Full vehicle simulations are conducted to assess the aerodynamics of wheel designs using a contemporary CFD approach. Beyond that, the detailed model of an isolated wheel including a tread pattern is investigated as an academic test case to develop new ways of modeling wheel and tire rotation in CFD.

# Numerical Method

The Lattice-Boltzmann method derives the macroscopic flow variables like density, momentum and energy from microscopic particle distributions following the Boltzmann kinetic theory. This method is inherently transient and discretizes the particles' motion both in velocity and direction on an equidistant, cubic lattice. The resulting equations are of

coupled, algebraic nature and macroscopic flow quantities are computed by simple summations. Furthermore, each time step is divided into a propagation step and a collision step with only the propagation step needing information from neighbor cells (compact stencil). Thus, parallelization is easier than in the case of the finite volume method which is based on coupled, partial differential equations, the Navier-Stokes equations. Since the spectrum of turbulent time and length scales is very large in the case of industrially relevant flows, a direct numerical simulation is not possible with today's computational resources. Instead, the smallest turbulent scales are modeled using an enhanced twoequation RNG k-epsilon model. The boundary layer is modeled according to the logarithmic law of the wall, additionally taking into account the effect of pressure gradients on flow separation.

# Wheel Rotation Modeling

The fact that wheels rotate at a moving vehicle is an evident observation which, however, creates challenges for the numerical simulation. Rotationally symmetrical geometries like the rim well can be handled by using a boundary condition prescribing the tangential velocity component on the surface in the form of  $u = \omega x r$ . The geometry itself does not rotate during the simulation. For the wheel spokes this boundary condition would not be appropriate since they rotate through the fluid changing their position with every time step. In the case of such a rigid body rotation the sliding mesh approach can be applied instead. The geometry is included statically in a separate, rotationally symmetrical mesh region which itself effectively rotates relatively to the global mesh. The sliding mesh region is connected to the global mesh by an interachieved in the algorithm.



Figure 2: Vortex cores visualized using a Lambda2 isosurface and total pressure distribution in the tire's wake from a transient snapshot.



face at which the interpolation between the separated domains takes place. The interpolation has, of course, implications for the solver's performance depending on the degree of parallelization that is

Handling the tire correctly is a more difficult task since it is deformed by static and centrifugal forces and is equipped with a complex tread pattern. Including the tire in a sliding mesh region is not possible because the deformed contact patch and the part of the ground intersected by the sliding mesh interface would also rotate around the axis. A possible approach is to run a Fluid Structure Interaction (FSI) simulation. The tire's deformation would be computed at every time step and then be used to update the geometry in the fluid solver. However, the material properties of the tire are usually not available for OEMs and the applied Lattice Boltzmann



Figure 4: Vortex pair forming in the upper wake of the tire, visualized from temporally averaged flow field.

> solver would need a partial rediscretization at every time step which would be extremely time consuming. Therefore, the simple approach to prescribe the tangential velocity on the tire's surface, similarly to the rim well, is chosen. In that context it is important to remove all lateral grooves from the tire's tread and to use a circumferentially averaged profile only containing the main longitudinal grooves. Otherwise, the lateral grooves would cause the flow to "trip" over the static edges leading to exaggerated separation behavior. Two flaws arise from that approach: the tangential velocity at the deformed contact patch is not correct and the shoveling effect of the lateral grooves is non-existent. In order to compensate the missing shoveling effect the tread is, therefore,

defined as a rough surface which magnifies the growth of the boundary layer.

# Influence of Wheel Designs

Both in CFD and in the wind tunnel different wheel designs have been investigated for their influence on the aerodynamic performance of the vehicle (Schnepf et al. [1]). In the numerical setup the geometry of a 2012 BMW 3 Series sedan is placed in a domain being 208 m long, 175 m wide and 138 m high. At the inlet and on the moving floor a velocity of 38.9 m/s is prescribed leading to a Reynolds number of 7 million, calculated using the wheelbase length. Using a finest cell size of 0.75 mm around the wheel spokes, 180 million volume elements and 30 million surface elements are created by the discretizer. 750.000 time steps have to be calculated for 2 s of physical time in order to achieve satisfactory convergence. In total, 32.000 core hours are consumed for one simulation run on 192 CPU cores.

The simulation results agree well with experimental data from the wind tunnel. In Fig. 1 the total pressure distribution and its downstream development are shown in three slices as an average of the last second. The most important structure is the ground vortex which evolves from the flow around the contact patch. The magnitude of total pressure loss (and thus energy loss) caused by it correlates well with the drag coefficients obtained for different wheel designs. Generally speaking, open wheel designs tend to cause more drag than closed designs with the spoke geometry being an additional factor. The same trend is observed in the wind tunnel when analyzing flow topology measurements, recorded using a five-hole probe, and force data from the balance. Altogether, the drag deltas between

different wheel designs from the simulation showed a deviation of 1% (of total vehicle's drag) from the drag deltas measured in the wind tunnel.

# **Isolated Wheel Investigations**

Focusing on the influence of a tire's tread pattern on its aerodynamic properties, Schnepf et al. [2] have investigated the flow around an isolated wheel in the wind tunnel and in CFD. For this study the tire manufacturer provided finite element analysis (FEA) deformed tires for the simulations. It was shown that the standard treatment of a tire, like in the full vehicle setup, led to a wrong separation behavior in the case of the isolated wheel. The tangential velocity boundary condition applied on the detailed tread pattern turned out to be insufficient if different tire models were to be compared. Consequently, alternative ways of modeling the wheel and, especially, tire rotation have to be investigated, starting with an attempt to rotate the whole wheel in a sliding mesh region. Since the ground must not be intersected by the interface a small gap has to be left between tire and ground. Although this is a clear error compared to the real on-road scenario it is accepted for this academic case. Using a smallest cell size of 1 mm close to the tire surface the discretizer creates 120 million volume elements. The computational effort for one simulation of 1.5 s physical time is 13.000 core hours. Evaluating the vortex structures in the tire's wake a significantly better agreement between experimental and numerical results can be achieved when including the whole tire in a rotating sliding mesh region. Although a cell size of 1 mm is not small enough to fully resolve the small grooves they affect the flow by transporting fluid upwards in the wake. Thus the flow separates earlier at the

top and the wake's height increases. This effect can only be captured using unsteady CFD as it naturally requires the geometry's rotational motion. Taking a closer look into the transient structures in Fig. 2, the complexity of the flow becomes obvious. The wake's nature is inherently unsteady including a wide spectrum of turbulent structures. Using the Lambda2 criterion small vortex cores are visualized in the transient snapshot. The evolution of vortices through the tread pattern affects the separation behavior and thus aerodynamic drag and lift.

# Summary and Outlook

software.

Unsteady CFD is a valuable tool that is already used in addition to wind tunnel testing. For the aerodynamic assessment of wheel designs it has proven to deliver reliable results. The example of a tire's tread pattern, however, shows that a new level of detail and physical modeling has to be achieved to compute the flow around tires correctly. The correct simulation model would include a contact patch at the ground and have the tread rotating at the same time. The approach to run a coupled simulation of fluid and finite element analysis solvers and to update the geometry at every time step is possible in principle, but the computational effort is beyond today's acceptable limits for an industrial application. Therefore, enhanced, alternative methods for handling moving geometries in CFD are needed. In that context, the case of an isolated wheel is a suitable academic test case to drive future developments of the simulation

# References

### [1] Schnepf, B., Tesch, G., Indinger, T.

On the Influence of Ride Height Changes on the Aerodynamic Performance of Wheel Designs. Proceedings of the JSAE Annual Spring Congress 2014 (Yokohama, Japan, May 21, 2014)

## [2] Schnepf, B., Tesch, G., Indinger, T.

Investigations on the Flow Around Wheels Using Different Road Simulation Tools. In "On Progress in Vehicle Aerodynamics and Thermal Management", Proceedings of the 9th FKFS Conference (Stuttgart, Germany, October 01-02, 2013), Expert Verlag

http://www.aer.mw.tum.de/

contact: Bastian Schnepf, bastian.schnepf@tum.de

# Bastian Schnepf

• Thomas Indinger

Institute of Aerodynamics and Fluid Mechanics, TU München

# 2nd Extreme Scaling Workshop on SuperMUC

In June 2014, LRZ organized the second installment of a four day workshop on extreme code scaling.

The extreme scaling workshop is a dedicated block operation time for selected projects which have the opportunity to test, debug and tune their software on site. Additionally, selected projects can perform unsupervised production runs during the night times of the workshop.

The goal of the workshop was to enable scaling of user implemented software on SuperMUC, the PetaFlops System at LRZ, which consists of 18 islands with 8.192 cores each. Prior to the workshop, the participants had to show that their code scales up to 4 islands (32.768 cores). Groups from 14 international projects managed to do so and came to the LRZ for a three day workshop.

Application experts from the LRZ, Intel and IBM were present during the workshop to resolve issues and assist in the performance optimization.

At the end of day three, 3 applications were successfully running on 18 islands, 5 applications on 16 islands and two applications on 8 and 12 islands respectively. During night time operation it was possible to do a run for 4 hours with a sustained PFlop/s performance.

LRZ demands physical attendance of the participants during the day time operation and provides a one-to-one support for all participants. This was possible with the help of a fine grained planning of the workshop schedule. For example on day one all participants had to measure the mpi-communication profile of their applications in order to resolve bottlenecks. Also due to a problem with the performance counters on the thin nodes, all programs were profiled on the fat node island (Flop/s rate, memory consumption, etc). Additionally to the node level optimizations the programs were tuned for better I/O utilization, which showed for some programs a dramatic increase (up to 42 GB/s for writing to a single file on 16 islands).

In summary the workshop was a huge success and the series of extreme scaling workshops will be continued in 2015.

In the following the results of selected application are presented by the authors of the programs.

contact: Ferdinand Jamitzky ferdinand.jamitzky@lrz.de

- Ferdinand Jamitzky
- Helmut Satzger

Leibniz Supercomputing Centre, Garching, Germany

# Advanced One-Sided Communication Patterns with GPI-2: Anisotropic Diffusion Filtering of Seismic Data



Figure 1: ECED filter on (artificial) seismic data. Left: before filtering. Right: after filtering, the double arrow marks two vertical fault lines.

Today modern Supercomputers like the SuperMUC at LRZ comprise a level of parallelism that was unheard of only a decade ago. This trend will continue as, within a given power envelope, it will be much easier to increase the total computing power by cloning nodes, cores, memory banks, etc. than by accelerating these very components. However, this kind of parallelism asks for new programming models.

The traditional Bulk Synchronous Processing (BSP) model that is still present in many HPC applications – and in the minds of their programmers – will fail more and more in the future for two reasons. The first reason is the tight and frequent coupling of all the components. The second reason is the growing risk of sporadic (hardware) failures in any of these components. Both challenges get worse with growing levels of parallelism and thus have to be met in future-oriented HPC applications. A general answer to this challenge is decoupling on all levels of the application. This work deals with the overlap of communication and computation and with unleashing the synchronization of the compute cores from each other. We take advantage of the features of GPI-2, a communication library that was created explicitly for one-sided communication patterns.

# **Features of GPI-2**

GPI-2 provides a concise API for distributed memory communication [1]. Exploiting the RDMA capabilities of modern high speed network interfaces it allows zero copy, one sided data transactions at wire speed which are completely detached from the CPU. This allows the application to efficiently overlap computation and communication. With GPI-2 the programmer keeps full and explicit control on all the resources. As a fully thread-safe API, it encourages a threaded model where fine-grain, task-based parallelism can be exploited. Further features are the timeout mechanism for non-local operations letting the programmer regain

control after any failed communication, to possibly fix the problem and to continue the computation.

# Anisotropic Diffusion Filtering of Seismic Data

Denoising seismic images is a common task to highlight special features, e.g. faults (see Fig. 1). The challenge here is to remove the noise without blurring the layer or fault structure. Our ECED filter stems from the class of Anisotropic Diffusion Filtering algorithms [2]. It simulates a diffusion process using an explicit Finite Differences scheme on a regular 3D grid. The Diffusion tensor comprises six components per voxel and is calculated anew for each time step. The computational effort is a mixture of eigenvector calculation to determine the diffusion tensor, a stencil scheme with constant weights for smoothing and a stencil scheme with non constant weights to perform the actual diffusion time step. In that sense the algorithm is a mixture of compute bound parts and rather memory bound parts.

# **Parallelization Strategy**

In general terms a hybrid approach is used here, i.e. one process per node is doing the communication while the computation is done SMP parallel. To build a full dynamic load balancing scheme the 3D image is logically decomposed into rods with a square front face and the full image depth. Per time step

# ECED Scaling Plot on SuperMUC



Figure 2: Scaling plot of ECED on SuperMUC.

each rod defines one workload package. Data dependencies between the packages have to be considered. The steps taken per package are reading the complete rod to the current node, performing a single time step and writing back the update. Data transfers are overlapped with computation, hiding them completely. No barriers are applied between time steps.

This approach accomplishes the goal of decoupling data transfer from computation and of decoupling the nodes from each other applying a full dynamic load balancing scheme.

# SuperMUC Adaptations and Results

Our goal is to test our ECED implementation and the parallelization scheme on a Petascale HPC system. On smaller machines its scalability has been shown before [2].

As the inter island communication bandwidth of SuperMUC is considerably lower than the intra island communication bandwidth, the workload allocation of the algorithm is adjusted. That modification keeps the execution of workloads mostly on the same island as the source of the data. The drawback here is a restriction of the dynamic load balance to the islands.

A small (artificial) seismic 3D image of 2.800x3.360x1.001 voxels (35GB) is used to demonstrate strong scaling properties. A second image of 15.000\*18.000\*1.001 voxels (1.007 GB) is used to demonstrate weak scaling properties. In both cases 10 time steps of the ECED filter are applied while the rod size is fixed to 100x100x1001 voxels.

The small data set shows perfect strong scaling up to 2k cores and reasonable scalability up to 16k cores. The big data set demonstrates weak scaling of our code up to 32k. Further improvements of the adjustments to the island structure of SuperMUC should yield weak scaling up to 130k cores.

# References

[1] http://www.gpi-site.com

[2] Bischof, C., Hegering, H.-G., Nagel, W.E., Wittum, G. (Eds.)

Competence in High Performance Computing 2010 - Proceedings of an International Conference on Competence in High Performance Computing, June 2010, Schloss Schwetzingen, Germany, chapter 9, 99-110 (2012)

contact: Martin Kühn, Martin.Kuehn@itwm.fraunhofer.de

# **Three-Dimensional Simulations** of Core-Collapse Supernovae with the VERTEX Code

Supernovae are the most spectacular cosmic explosions we know. About ten such events occur per second somewhere in the visible universe and can outshine a whole galaxy for weeks. Besides their extraordinary brightness, supernovae are vital components in the galactic cycle of matter. Essentially all elements heavier than helium were once processed in stars and eventually ejected into space by violent stellar explosions. Supernovae thus enrich the universe with heavy elements which are the raw material for new stars and planets.

A so-called core-collapse supernova signals the end of the life of a star with more than about eight solar masses. When nuclear fusion ceases in the center of such a star, the collapse of the stellar core becomes unavoidable. It implodes under its own gravity within less than one second to form an ultra-dense neutron star. A shock wave is launched and propagates outwards but stalls after roughly 100 kilometers because of the ram pressure of still infalling outer material. Huge numbers of neutrinos emitted from the hot, newly formed neutron star begin to deposit a fraction of their energy behind the stalled shock. If this energy transfer is powerful enough, the shock can be revived and eventually disrupts the whole star. The neutron star in the center cools and can even collapse to a black hole. Corecollapse supernovae are thus the birth sites of exotic objects whose interior properties are still not fully understood.

Crucial aspects of the explosion mechanism of core-collapse supernovae are still unclear. Astronomers can observe light curves and hope for neutrinos from a galactic supernova, but deducing the details of the explosion mechanism from such data is not possible. We clearly need computer simulations for looking deep into the cores of exploding stars and for a better understanding how the explosion starts.

constraints.

As described above, neutrinos play a crucial role during the explosion. They cannot be treated as a fluid component, but their propagation from the dense neutron star interior to the

Martin Kühn

Rui Machado

Germany

Fraunhofer ITWM.

Kaiserslautern,

Modeling core-collapse supernovae is challenging because a broad diversity of physics is involved, in particular the extremely complex problem of neutrino transport and neutrino reactions in dense medium. Multidimensional hydrodynamic calculations must be performed because convection and large-scale non-radial instability modes of the accretion shock are crucial for the onset of the explosion [1]. Simulations in spherical symmetry (1D) are therefore not suitable for capturing the essential aspects of the problem. Axially symmetric (2D) simulations also set artificial constraints on the fluid motions and lead to an inverse cascade of turbulent flows. Therefore, modeling the mechanism of core-collapse supernovae is only possible in full three dimensions (3D) without any symmetry

neutrino-heating region is a computationally very expensive radiation transport problem. A direct solution of all aspects of this transport problem as described by the 6+1 dimensional Boltzmann equation is computationally not feasible in current 3D simulations. Therefore a sophisticated approximation of the neutrino transport is implemented in the VERTEX supernova code [2] using the so-called "ray-byray plus" approach. For each angular "ray", spherical transport problems are solved by an iterative procedure involving a model Boltzmann equation and its first two angular moment equations. Neighboring rays are coupled by neutrino advection with moving fluid elements and neutrino pressure gradients. VERTEX provides the currently most sophisticated treatment of neutrino physics in 3D simulations.

The VERTEX code is parallelized efficiently to run large-scale simulations

of supernovae on high-performance computers. Its parallelization strategy is a hybrid OpenMP/MPI scheme where each neutrino transport ray is associated with one OpenMP thread and one MPI task runs on one compute node. This setup guarantees excellent scaling of the VERTEX code on a large variety of tested machines [3]. The code has been successfully applied in production runs for several years thanks to PRACE and GCS grants. Typical 3D simulations with an angular resolution of two degrees use nearly 16.000 cores and consume roughly 50 million core hours for one second of physical evolution. Higher grid resolution is desirable but feasible only with larger numbers of cores. The current resolution is limited by the machine partitions that are made available for long-time production runs.

In Fig. 1, we show strong scaling results of VERTEX on SuperMUC up to



Figure 1: Strong scaling on SuperMUC for two grid configurations, the spherical polar grid (top) and the Yin-Yang grid (bottom).

about 131.000 cores measured with a production setup without any special optimization. The scaling is presented both for a spherical polar grid and a so-called Yin-Yang grid [4], which we have implemented recently. The latter grid configuration consists of two lowlatitude sections of a spherical grid in overlap (see Fig. 1). Grid singularities at the polar axis and associated numerical artifacts are thus excluded and severe limitations of the hydrodynamic time step due to the grid-cell deformation near the poles are avoided. A clear disadvantage of the Yin-Yang grid is its non-trivial transformation which requires a point-to-point-like MPI communication pattern at the grid boundaries in the overlap region. However, despite this drawback we were able to achieve excellent strong scaling also with the Yin-Yang grid during the Extreme Scaling Workshop 2014 at the Leibniz Rechenzentrum.

We thus demonstrated that the VERTEX code is a perfect tool for studying core-collapse supernovae on Petaflop supercomputing platforms. Its excellent performance and nearly linear scaling will allow us to run our simulations with great efficiency also on the next generation of high-performance computers.

# References [1] Hanke, F., Müller, B., Wongwathanarat, A., Marek, A.,

- Janka, H.-T.
- [2] Rampp, M., Janka, H.-T. 396:361-392, 2002
- Janka, H.-T.
- [4] Kageyama, A., Sato, T. 2004

contact: Tobias Melson,

# Applications

SASI Activity in Three-dimensional Neutrinohydrodynamics Simulations of Supernova Cores. ApJ 770:66, 2013

Radiation hydrodynamics with neutrinos. Variable Eddington factor method for core-collapse supernova simulations. A&A

# [3] Marek, A., Rampp, M., Hanke, F.,

Towards Petaflops Capability of the VERTEX Supernova Code. arXiv:1404.1719, 2014

Yin-Yang grid: An overset grid in spherical geometry. Geochem. Geophys. Geosys. 5,

melson@mpa-garching.mpg.de

- Tobias Melson<sup>1</sup>
- Andreas Marek<sup>2</sup>
- Hans-Thomas Janka<sup>1</sup>
- Florian Hanke<sup>1</sup>
- Max Planck Institute for Astrophysics
- <sup>2</sup> Computing Centre of the Max Planck Society (RZG)

# Studying Wave-Particle-Interaction with the ACRONYM PiC Code on SuperMUC

Particle-in-Cell (PiC) Codes are a powerful numerical tool for the investigation of collisionless plasma phenomena. However, large scale simulations can be very computationally demanding. Nevertheless, due to the availability of High Performance Computing resources, such as SuperMUC at the LRZ, even big and complicated plasma simulations have become tractable. Our PiC code ACRONYM [1] has already proven to be highly suitable for such large-scale simulations.

# Wave-Particle-Interaction in Kinetic Simulations

Within project pr85li, we treat resonant interactions of charged particles and waves in a thermal plasma, which is modeled using a kinetic approach. The overall goal is to study some of the key transport mechanisms for energetic particles in the solar wind on a microscopic scale. So far, we have investigated resonant scattering of protons off low frequency waves – a process which is considered to be both an important mechanism for the transport as well as for the acceleration of relativistic protons. We are also planning to adapt our current simulation setup in order to make it suitable for studying electron scattering.

We have adopted the setup of Lange et al. [2], who approaches wave-particleinteraction with MHD simulations. We were able to reproduce the results from MHD simulations and to obtain the typical transport characteristics as



Figure 1: Scatter plots for two time steps of the same simulation. The plots show the change of each particle's direction of motion (y-axis) depending on the initial direction (x-axis). Color coding represents the (normalized) number of particles at each point in the scatter plot. The peaks in both plots represent resonant scattering, i.e. a drastic change in the particles' directions of motion. Readers who prefer a more detailed physical description are referred to [3].

predicted by quasi-linear Vlasov theory (QLT), as shown in Fig. 1.

So far, our results prove that PiC codes are capable of reproducing resonant wave-particle-interaction. This is an important result, since at the time most approaches to model the interaction of fast particles with the electromagnetic fields of waves in a thermal plasma are not self-consistent – whereas the PiC method is.

Taking the previous results as a basis, PiC codes can be used to explore waveparticle-interaction beyond the physical regime of the MHD approximation, i.e. the high frequency regime of plasma modes, where waves are dispersive or damped. In preparation for our next steps, we incorporated a model for damped waves into our code, which now allows us to drive such waves. Driving a wave means that throughout the running simulation energy is fed into the wave by adding small amounts to the electric and magnetic field strengths in the existing field vectors in each time step. Knowing the approximate damping rate for a driven wave enables us to cancel out dissipation losses and prevent the wave from decaying.

# Outlook

We see great potential for PiC simulations in this area and are planning to study electron scattering off Whistler waves, which are waves in the dispersive regime of the R-mode. The described scenario is considered to be an important process in the transport of energetic electrons within the heliosphere and the numerical modeling of the microphysical properties of the interaction can be useful for a deeper understanding of transport characteristics. Still, these simulations consume a great amount of computing time. The insights gained at the Extreme Scaling Workshop 2014 at the Leibniz-Rechenzentrum Garching allow us to better leverage High Performance Computing resources during our research.

# Results from the Extreme Scaling Workshop 2014

During the workshop, we performed a tracing run on 4.096 processor cores using the Intel Trace Collector. Analysis of the resulting trace files showed good parallelization behavior in general: Long computational sections with short communication bursts. Since the algorithm almost exclusively relies on nearestneighbor communication, good scalability could be expected. The analysis suggested, however, that a diagnostics procedure in ACRONYM might create a bottleneck in large-scale runs. This procedure was turned off as a consequence.

Output of the computational data proved to be a big problem. The use of parallel HDF5 resulted in extremely poor performance, necessitating a complete removal of output for runs with more than 32.000 cores.

We were able to complete scaling runs on up to 16 islands (131.072 cores) on SuperMUC. ACRONYM exhibited excellent, almost linear, weak scaling behavior (see Fig. 2).

More than ten billion particle updates per second (a good measure of a PiC code's performance) were reached. With an average of about 80 TFLOPS (double precision). This suggests an efficiency of about 3%. Taking previous testing into account, the code seems to be heavily limited by memory and



ACRONYM weak scaling on SuperMUC (LRZ)



Figure 1: The ACRONYM code scales nearly linearly with the number of CPU cores. A speedup by a factor of two can be achieved, if the GNU compiler is used instead of the Intel compiler. However, writing data to disk is a huge problem.

103

cache speeds. Further, smaller scale testing revealed that significant performance gains can be achieved by switching to the GNU Compiler Collection for our use case.

 $10^{2}$ 

# Acknowledgements

We would like to thank Christoph Bernau and the entire LRZ team for their help and support during the Extreme Scaling Workshop 2014.

# References

101

### [1] Kilian, P., Burkart, T., Spanier, F.

2011: The Influence of the Mass Ratio on Particle Acceleration by the Filamentation Instability. In W. E. Nagel, D. B. Kröner, M. M. Resch (Eds.), High Performance Computing in Science and Engineering '11. Springer, Berlin Heidelberg, p. 5. DOI= http://dx.doi.org/10.1007/978-3-642-23869-7

# [2] Lange, S., Spanier, F., Battarbee, M., Vainio, R., Laitinen, T.

2013: Particle scattering in turbulent plasmas with amplified wave modes. In Astronomy and Astrophysics 553, A129. D0I=http://dx.doi.org/10.1051/0004-6361/201220804

105

## [3] Schreiner, C., Spanier, F.

104

2014: Wave-particle-interaction in kinetic plasmas. In Computer Physics Communications 185 (7), p. 1981.

contact: Cedric Schreiner, cschreiner@astro.uni-wuerzburg.de

# Computational Engineering goes HPC: Thermal Comfort Assessment on Massive Parallel Systems



Figure 1: Velocity profile in a cut-plane through a test room occupied by a human manikin test subject.

# **Project Outline**

As computers tend to get faster and more powerful every year, large engineering problems such as complex buoyancy driven indoor air flow scenarios – deemed unsolvable a decade ago – can be tackled nowadays.

Within our simulation code, the main computational kernel deals with solving the Navier-Stokes equations for an incompressible Newtonian fluid flow coupled with a thermal convection-diffusion equation via the Boussinesq approximation. The applied data structure is composed of block-structured orthogonal Cartesian grids (including halos) organised in a hierarchical layout, addressing next to grid adaptation also aspects such as efficient grid distribution and migration in order to support parallel computations on several thousands of cores. For solving the Poisson equation arising in every time step of the solution of the Navier-Stokes equations, a multi-grid-like solving technique was implemented. The reader is referred to [1,2] for further details.

Fig. 1 shows a velocity profile in a cutplane through a test room, where wall temperatures were set to "cold" and the numerical manikin – coupled to a human thermoregulation model – was set to "hot". Hence, the flow is described by a natural convection phenomenon and, thus, is purely driven by buoyancy effects. One example of engineering applications comprises a thermal comfort assessment of the

• Andreas Kempf<sup>1</sup>

Felix Spanier<sup>2</sup>

<sup>1</sup> Ruhr-Universität

<sup>2</sup> North-West

University

Potchefstroom,

<sup>3</sup> Julius-Maximilians-

South Africa

Universität

Würzburg,

Germany

Bochum, Germany

• Cedric Schreiner<sup>23</sup>



0.01 128 512 2048 8192 32768 131072 amount of processes [-] Figure 2: Communication times for a total ghost layer exchange of nine independent variables per cell in seconds on two different supercomputing systems (Shaheen representing an IBM Blue Gene/P installed at KAUST) for a 3D domain with resolutions up to 4.096×4.096×4.096 on depth 8 containing a total of 78.5 billion cells.



Figure 3: Strong speedup results on two supercomputing platforms for a 3D domain with resolutions up to  $4.096 \times 4.096 \times 4.096$  on depth 8 containing a total of 78.5 billion cells for solving the pressure Poisson equation.

- Jérôme Frisch
- Ralf-Peter Mundani
- Ernst Rank

Chair for Computation in Engineering, Technische Universität München, Germany mine how comfortable an occupant would feel with respect to ambient surrounding temperatures or draught effects. Due to our simulation code, large rooms with complex geometries can now be computed using very high resolutions for a quantitative analysis.

computed results in order to deter-

# Scaling Results

During the "Extreme Scaling Workshop" held at LRZ from June 2–5, 2014, the unique possibility of testing the code on the complete SuperMUC (up to 18 islands) presented an extraordinary opportunity to investigate the code's scaling abilities and behaviour. Fig. 2 depicts the necessary halo communication time in order to exchange nine independent variables in all ghost layers of all grids. It can be observed, that the time is decreasing linearly in the double-logarithmic plot for an increasing amount of processes and at some point levelling-off-effects start to appear.

Fig. 3 shows the strong speedup for solving the pressure Poisson equation arising in every time step of solving the Navier-Stokes equations. Depth 8 contains a total of 78.5 billion cells, depth 7 has 9.8 billion cells, and depth 6 has 1.2 billion cells. Hence, in the strong speedup measurements, the domain size was kept constant, which implies a levelling-off of the performance with an increasing amount of processes at some point when the work load per process drops below a certain threshold. Nevertheless, while using the complete machine a strong scaling efficiency of 64% could be achieved while using nearly 80 billion cells.

The authors would like to thank the LRZ for the support and the opportunity to participate in the Extreme Scaling Workshop.

# References

[1] Frisch, J., Mundani, R.-P., Rank, E.,

Adaptive multi-grid methods for parallel CFD applications. Scalable Computing: Practice and Experience, 15(1), pp. 33–48, 2014

[2] Frisch, J., Mundani, R.-P., Rank, E., Parallel multi-grid like solver for the pressure Poisson equation in fluid flow applications. In Proc. of the IADIS Int. Conf. – Applied Computing, 2013

contact: Ralf-Peter Mundani, mundani@tum.de

# Hybrid Ateles on SuperMUC

Ateles is a high order discontinuous Galerkin solver and part of the Apes framework. It uses explicit time stepping and is suitable for linear and nonlinear equation systems. Its current main application fields include Maxwell equations, Acoustic wave equations and the inviscid Euler equations. The numerical scheme is based on tensorproduct polynomial representation of the solution in cubical elements. This allows an efficient dimension-bydimension approach and the usage of spectral elements. Modal and nodal representations are used as needed by the simulation and a fast transformation routine is deployed for the switch.

Using polynomial modes directly with a suitable basis, allows the fast computation of integrals, but also eases the exchange of data between differently refined elements. This scheme is therefore especially well suited for the octree meshes offered by the underlying TreEIM library, which was designed to enable large scale distributed simulations.

A big advantage of high order discontinuous Galerkin schemes is their high locality offered by the relatively small amount of data that needs to be exchanged at the surfaces of the elements. This relatively loose coupling



Figure 1: Visualization of the scattering of a electro-magnetic wave at a perfectly electrical conducting cylinder with smooth capturing of the curved boundary and sub-element features extracted by Harvester.

Applications

between the elements is nicely fitting the need to distribute computations on modern HPC systems with massively parallel resources. On the other hand, computations within elements typically require the complete data within the element accessed in different patterns. These are therefore less suitable for distributed computing, yet shared memory parallelism within nodes still can be exploited for the operations within each element.

One drawback of high order schemes could be the handling of geometric boundaries with sufficient resolution. Ateles deals with this, by using material variations within elements to represent geometrical objects in the form of polynomials, just like the solution. Another point is the visualization of the solution out of the polynomial modes. For Ateles we take care of this by the parallel post-processing tool Harvester within the Apes framework, which allows us to subsample data and create higher resolved meshes just for the visualization.

Taking advantage of hybrid parallelism, we were able to scale Ateles on up to 16 islands of SuperMUC. In the analysis of the code on a single node, we found that 2 threads per MPI process provided the best performance and utilizing all cores is beneficial. However, when distributing the simulation mesh across the network and making use of more nodes, we found that 4 threads for 4 MPI processes on each node offers highest performance. Scaling runs of the solver are therefore mainly done in this setup. To allow exploitation of shared memory parallelism and to demonstrate the actual use of high orders, we use a scheme of 64th order in the presented scaling. For this order there are more than 1.5 million degrees of freedom per element, which is



Figure 2: Weak scaling of Ateles on SuperMUC with 16 elements of 64th order and 4 OpenMP threads per MPI process.

already so large, that a single element does not fit into the cache anymore. Unfortunately it also limits the variety of different problem-sizes we are able to run, as at least one element per MPI process is required. With a fourth order Runge-Kutta time integration, each degree of freedom time update requires roughly 232 floating point operations. The largest problem we can fit onto a single process with 4 threads are 32 elements of 64th order. Due to some vectorizations over the number of elements it is beneficial to use at least 8 elements per process, leaving a rather narrow range for varying problem sizes. This is mainly due to the fact, that the scheme allows us to fine tune the discretization parameters to fit the machine as good as possible and distribute parallelism as needed.

The weak scaling with 16 elements per process shows a clear drop in the performance from a single node to two nodes and thus requiring network communication, but after this, the performance per node remains almost constant up to a single island. In the scaling beyond a single island there can again some variation be observed, however even on 16 islands of SuperMUC we achieve 88.7 % of parallel efficiency compared to a single node. Half of the drop in the parallel efficiency already takes place in the step from a single node to multiple nodes.

contact: Harald Klimach, harald.klimach@uni-siegen.de

- Harald Klimach
- Peter Vitt
- Jens Zudrop
- Sabine Roller.

Simulationstechnik & Wissenschaftliches Rechnen, Universität Siegen

# **CIAO Code - Extreme** Scaling Workshop LRZ 2014

The numerical group of the Institute for Combustion Technology at RWTH Aachen University currently uses the SuperMUC supercomputer in two different projects. One focuses on the investigation of irregular combustion phenomena in internal combustion engines and the other on improved prediction and modeling of pollutant emissions in turbulent premixed flames. Both projects use a multi-physics inhouse code called CIAO. The modular high-order, structured, finite difference code is used for simulating reactive and non-reactive flows by solving the Navier-Stokes equations. The code features efficient numerical methods for handling of complex geometry, mesh and solid body motion, and the coupling

between reaction and fluid motion. Efficient parallel scalable block decomposition is done to ensure scalability even with varying computational load in the numerical domain. Communication between grid partitions is implemented by using the Message Passing Interface (MPI).

# Investigation of Irregular **Combustion Phenomena in** SI Engines using Large-Eddy Simulations

Although spark-ignited engines have a considerable development history, the relevant flow physics are still not fully understood. One reason is the lack of The numerical group of the Institute for Combustion Technology at RWTH



Figure 1: In-cylinder flow field.

Aachen University currently uses the SuperMUC supercomputer in two different projects. One focuses on the investigation of irregular combustion phenomena in internal combustion engines and the other on improved prediction and modeling of pollutant emissions in turbulent premixed flames. Both projects use a multi-physics inhouse code called CIAO. The modular high-order, structured, finite difference code is used for simulating reactive and non-reactive flows by solving the Navier-Stokes equations. The code features efficient numerical methods for handling of complex geometry, mesh and solid body motion, and the coupling between reaction and fluid motion. Efficient parallel scalable block decomposition is done to ensure scalability even with varying computational load in the numerical domain. Communication be-

# Investigation of Irregular **Combustion Phenomena in** SI Engines using Large-Eddy Simulations

tween grid partitions is implemented by

using the Message Passing Interface (MPI).

Although spark-ignited engines have a considerable development history, the relevant flow physics are still not fully understood. One reason is the lack of experimental and numerical methods with sufficiently high resolution or capabilities of capturing stochastic phenomena. The latter aspect is of crucial importance at extreme engine operating conditions where irregular combustion may occur. More recently, Large-Eddy Simulation (LES) has been identified as a promising technique to establish a better understanding of in-cylinder flow variations. However, simulations of engine configurations are challenging due to resolution as well as modeling requirements and computational cost for these unsteady multi-physics problems.



Within this project, the CIAO Code is used for the solution of the compressible Navier-Stokes equations. Models for liquid spray injection and combustion in the premixed, partially-premixed, and non-premixed regime have been developed to be used for the simulation of complex in-cylinder processes with high accuracy.

Simulations with high spatial resolution are performed to validate the numerical framework against high-quality optical measurements for both cold and reacting flows. Multi-cycle simulations at critical operating conditions are performed at different load conditions. In addition to giving a better understanding of complex physical phenomena, the high-resolution LES project is meant to be a preliminary study for a subsequent direct numerical simulation (DNS) of a reactive engine configuration.

In order to assess the performance of the explicit low-storage Runge-Kutta time marching scheme, a compressible channel flow with scalar transport equations was tested on SuperMUC during the Extreme Scaling Workshop 2014. Momentum equations were discretized by a fourth-order accurate

### Figure 2: Speedup of compressible solver



Figure 3: Volume rendering of the scalar dissipation rate of the progress variable.

central differencing scheme while a weighted essentially non-oscillatory (WENO) scheme of the same order was used for spatial derivatives in the scalar equations. The computational grid consisted of 6.4 billion cells. Three tests on 32.768, 65.536 and 131.072 cpu cores were conducted to evaluate the parallel performance. The results indicate almost linear strong scaling behavior. In the future, the collected data will be used to optimize the code and further improve parallel efficiency. A similar test case in a more complex geometry, which is challenging with respect to mesh partitioning, will also be considered.

# Large-Scale Simulations and Modeling of Pollutant Emissions in Turbulent Premixed Flames

Within the Sonderforschungsbereich (SFB 686), we are developing LES models for turbulent premixed combustion that can be used for the design of future gas turbines. Although LES of laboratory-scale flames and devices of industrial importance have recently been carried out with great success, predictions of NO, formation are still challenging, especially for low temperature combustion, where not only thermal NO,, but also more complex formation pathways are important. The main impediment for developing high-fidelity models stems from the profound lack of high quality data that can be used to establish a better understanding of the complex NO, chemistry. Rapid advances in supercomputing make DNS a powerful tool in combustion science and enable the simultaneous simulation of turbulence and chemistry as well as the analysis of their interaction. In this project, large-scale DNS are conducted in order to understand and to model the intricate interaction of NO, chemistry and turbulence. To this end, the Navier-Stokes equations are solved in the low-Mach number limit.

In the DNS, two initially laminar lean premixed methane-air flames propagate into a temporally developing jet. The computational domain is discretized with 3 billion grid points and the methane chemistry is described with a detailed mechanism containing 32 species and 213 reactions, which results in a total of 87 billion degrees of freedom. The scalability has been assessed in three test runs on 32.768, 65.536 and 131.072 cores. The scaling plot shows excellent scaling up to 65.536 cores and reasonable scaling up to the entire machine size of 131.072 cores. The three major contributors to the cost of the flow solver are combustion (computation of reaction source terms), scalar transport, and the solution of the elliptic equation for pressure. The pressure equation requires all-to-all communication of all processors involved and therefore proves to be the limitation for better scalability, while the chemistry solver has no communication costs and enables very good scalability.

contact: Konstantin Kleinheinz, k.kleinheinz@itv.rwth-aachen.de



Figure 4: Speedup of low-Mach solver.

• Konstantin Kleinheinz

Institut für Technische Verbrennung RWTH Aachen University, Germany



47

# **Extreme Scaling of** NSCOUETTE, a Pseudospectral DNS Code



# **Motivation and Scientific** Background

In accretion disks mass flows towards a central gravitating body to accrete on it, hence losing angular-momentum. Because of angular-momentum conservation, this loss must be balanced by outward transport among gas particles, and a simple analysis of accretion time scales reveals that star and planet formation can only occur if the motion of the gas is strongly turbulent. Although rotation in accretion disks is very fast and the length scales concerned are huge, resulting in astronomical Reynolds numbers (Re), it is not clear whether one should actually expect such flows to be turbulent. In fact, Lord Rayleigh's criterion assures that laminar flows with Keplerian velocities (v ~  $1/\sqrt{r}$ ) are indeed stable to

small disturbances even in the limit of infinite Reynolds number. Nowadays, it is accepted that magnetic fields are a vigorous source of turbulence in accretion disks; nevertheless this requires disks to be ionized. Without ionization, other physical mechanisms must be responsible for generating turbulence. Arguably, the simplest explanation would be that despite stability to small disturbances, turbulence can arise if disturbances are sufficiently large.

Our aim is to elucidate whether strongly disturbed Keplerian flows become turbulent at a critical rotation velocity. Hence the key question is about the existence of turbulence and modeling strategies such as the Reynolds-averaged equations and Large-Eddy Simulation are precluded. Thus we must

resort to direct numerical simulation (DNS) of the Navier-Stokes equations, which means that all scales of motion must be resolved in space and time. In order to address the essential physics and allow comparison to laboratory experiments we focus on the canonical problem of fluid motion between two concentric cylinders, whose rotation is adjusted to generate a Keplerian profile. In addition, a negative radial temperature gradient, as expected in accretion disks, is considered. Comparison to experiment requires the simulation of flows at Re>10<sup>5</sup>, which together with the requirement for resolving all scales, implies the need for huge numerical grids and accurate

For this purpose we have developed the code NSCOUETTE (Shi et al., under review in Computers and Fluids 2014, http://arxiv.org/abs/1311.2481), which implements a hybrid-MPI-OpenMP parallel DNS method for turbulent Taylor-Couette flow. The Navier-Stokes equations are discretized in cylindrical coordinates with the spectral Fourier-Galerkin method in the axial and azimuthal directions, and high-order finite differences in the radial direction. Time is advanced by a second-order, semi-implicit projection scheme, which requires the solution

transpose (overlapped)

transpose

linear system

E FFT

res

30

25

20

15

10

weak scaling, setup A (N\*512,675,2048)

(1 MPI-task/node)

4

N SuperMUC islands (x8192 cores, x512 MPI-tasks)

non-dissipative schemes.

U. Göttingen, 2014).

# Extreme-Scaling **Experiments on SuperMUC**

Fig. 2 summarizes our measurements with NSCOUETTE for two different setups which were performed on Super-MUC during the LRZ Extreme Scaling Workshop 2014. The code was recently augmented by equations describing a temperature gradient, and for the first time we were able to test this new version at the very large scale and with significantly larger numerical resolution than before. Thanks to the hybrid parallelization, only up to 8.192 MPI tasks were required to utilize the maximum of 131.072 processor cores (16 islands) which allowed us to smoothly run the code with such high core counts.

The linear solvers (blue bars) and the fast Fourier transforms (green bars),



2

of five Helmholtz/Poisson equations. Nonlinear terms are computed with the pseudospectral method. Fig. 1 shows strong turbulent activity in an NSCOUETTE simulation which typically covers a fraction of an annulus representing an accretion disk. However, turbulence eventually decays in the course of this simulation, which was performed still without taking into account temperature gradients (cf. L. Shi, PhD thesis,

both being task-local operations parallelized with OpenMP, maintain excellent weak scalability up to 16 islands. As expected, the global transpositions of the data (yellow and orange bars) which are characteristic for pseudospectral methods eventually limit the overall scalability of the code due to collective communications of the MPI Alltoall type. Nevertheless, this communication pattern performs remarkably well up to at least 32.786 cores (4 islands) of SuperMUC, despite its 4:1 blocking factor across islands and despite the fact that message sizes decrease by a factor of two in our weak scaling setup with every doubling of the number of tasks. We expect to be able pushing scalability further, e.g., by combining the MPI\_Alltoall communications of individual arrays into a single call with larger messages.

We conclude that the NSCOUETTE

code can efficiently utilize Tier-O class

supercomputers like SuperMUC, with a

preference for high-performance net-

works like fully non-blocking InfiniBand

or other high-end fabrics. This paves

lations of Taylor-Couette flow at very

large Reynolds numbers and with high

numerical fidelity and will thus enable

contributions to long-standing hydrody-

namics problems like in star and planet

formation (Balbus, Nature, 2011).

contact: Prof. Marc Avila,

marc.avila@fau.de

the way towards direct numerical simu-

<sup>2</sup> Universitat Politècnica de Catalunya (UPC), Barcelona, Spain

Markus Rampp<sup>1</sup>

Jose-Manuel

• Liang Shi<sup>3,4</sup>

• Biörn Hof<sup>3,4</sup>

Marc Avila<sup>5</sup>

<sup>1</sup> Computing

Society,

Garching,

Germany

Center (RZG) of

the Max Planck

Lopez<sup>2</sup>

- <sup>3</sup> Max-Planck-Institute for Dynamics and Self-Organization (MPIDS), Göttingen, Germany
- <sup>4</sup> Institute of Science and Technology (IST), Klosterneuburg, Austria
- <sup>5</sup> Friedrich-Alexander-Universität (FAU), Erlangen-Nürnberg, Germany

Extreme Scaling of the PSC on SuperMUC at LRZ

# Modern Laser and Plasma **Physics**

The laser-plasma interaction provides new sources of radiation and particles like electrons and ions. Upcoming laser facilities like ELI or proof of principle experiments like the AWAKE project at CERN have budgets of about €1B and about €15M, respectively. In both cases a good understanding of the expected physics beforehand is needed and can strongly increase the return on the investment.

# **Proton Driven Wake Field** Acceleration

PdWFA is a promising candidate for GeV electron accelerators as proposed by Lotov and Caldwell (2007). The simulation of the problem is computationally demanding. The group has committed itself to contribute simulation support. Preliminary numerical investigations of the problem show a similar scaling behavior as the homogeneous case depicted in Fig. 3b.

# **Ultra-Thin Foils**

As novel radiation sources, ultra-thin foils have been intensively studied in the community over the last few years (e.g. Kiefer and Rykovanov (2009)). The simulation of thin foils in 3D is a super-computing problem with load balancing challenges (see Fig. 1).

# The PSC

Developed in the late 90s by H. Ruhl in Fortran and released under GPL the PSC (Plasma Simulation Code) is a well-tested, widely-recognized and reliable Particle-in-Cell Code that has also been fundamental to other current

codes, e.g. the EPOCH code. In 2009 Ruhl et al. and Germaschewski et al. began porting the code to a modular C framework supporting bindings to Fortran and C/CUDA, featuring selectable field and particle pushers.





Figure 1: The spacetime plot of the field energy density (below) exhibits a strong spike giving rise to an attosecond x-ray pulse (AXP). Moreover, the relative foil motion (above) displays a complicated behavior, hinting at lateral radiation emission.

The Plasma Simulation Code: A modern particle-in-cell code with load-balancing and GPU support (arXiv:1310.7866)

# **QED** with APR

Our latest findings about the simulation of quantum electrodynamics show that multiple particle weights are required to simulate cascades (Klier and Ruhl to be submitted). This can lead to a superexponential growth of particle number over time (O[e<sup>et</sup>]) that can only be managed with Adaptive-Particle-Refinement (APR) (see Fig. 2).



# **CLI and Q/A Coverage**

The modular C design makes different solvers and accelerators available via a powerful command line interface:

src/psc --psc output particles type xdmf compact\ --psc\_push\_particles\_type generic\_c\_superp --np\_z 512 --length\_z 16384e-9 --gdims\_z 262144\ --lw 1.0e-6 --pulse\_length 1.25e-6 --n0 7136e28\ --thickness 60e-12 --write checkpoint yes

To handle this flexibility thorough checks for dependencies and broken modules are necessary. They are carried out by automated build/run tests via Buildbot and LCOV.



Figure 2: Shown is the positron momentum distribution in pxpy- space from spontaneous electron-positron pair creation in the focus of 6 laser beams. On above the control of quasi particles is demonstrated.

# Further Technological Key Features

- Moving window: Multiple moving and/or dynamically growing simulation areas
- Autotools build system

- Modular I/O subsystem including several HDF5/XDMF modules for large scale parallel output
- Dynamic load and memory balancing
- MPI only or hybrid MPI/OpenMP parallelization as well as SSE/ AVX(512) micro vectorization
- CUDA and XEON PHI (MIC) acceleration
- AMR: Adaptive Mesh refinement (preliminary work, in preparation)

# Scaling

To improve scaling we moved to serial directory creation in the rank centric output module following advice by IBM. A configure flag allows to disable the code internal performance counters since they collide with LIKWID. Scaling works well up to 8 islands. A bottleneck in the initial balancing has been discussed with the LRZ application support and solved to enable a full 16 island run.

# Performance PSC internal performance counters:

******************		********	+-	*********	=== step		
	name		avg	min	Tax	total	
n.	particles	5.033164	e+07.	5.03316e+07	5.03316e+07	6.44245	e+09
nemory	max (GB)	4.73	906	4.71164	4.89679	606.599	
memory	free (GB)	94.3	443	94.0206	94.6195	12076.1	
time e	ntire step	57.8	023	57,1903	58.3213	7398.7	
time parti	cle update	52.1	686	46.989	53.323	6680.14	
time fi	eld update	0.312	526	0.272563	0.323622	40.0033	
time con	munication	5.29	058	3.86947	10.7712	677.194	
		*******					-
	t-avg	t-min	t	-max	MFP-avg	MFP-min	MFP-max
psc_step	57802.43	57190.39	5832	1.43	1051.02	1040.00	1062.00
push particles run	52188.69	46989.08	5332	3.02	1155.06	1129.00	1280.00
genc_part_xyz	203.85	183.54	20	8,28	1155,15	1129.00	1280.00
_							
xchg post	0.02	0.02		0.04	0.00	0.00	0.00
xchg photon	168.15	0.32	110	6.23	0.01	0.00	1.00
-							
Tablasted constations	min birms	(mm.) and and	4.001	. 66.66.01.	5.0		

LIKWID has measured 1.035 MFlops. The 2% difference can be attributed to the setup time.

 Inhomogeneous problem: 300-600 MFlops/s per core



Figure 3a: Scaling matrix for an inhomogeneous case from 32 up to 131072 cores. Balancing and communication is reasonable even for larger systems. The dashed red point is an interpolation to the full machine run without load balancing issue as 2 island tests indicate. Solid red is an OpenMP version without communication patterns fully respecting OpenMP parallelization.



Figure 3b: Scaling matrix for a homogeneous case up to 32768 cores. Squares refer to the time needed by the particle pusher only, while circles represent the whole step including fields, communication and administrative tasks. Close to perfect strong scaling: Efficiency Efstrong > 99%. Very good weak scaling is obtained too:  $Ef_{meak}$  (N = 512)  $\approx$  90%...

- Homogeneous problem: 1.000 MFlops/s per core
- Balanced peak performance: 1.35 GFlop/s per core

The latter is around 14% of the theoretical peak performance of a Super-MUC Fat Node Westmere EX. A good result as there is room for improvement such as more aggressive compiler optimizations, a rewrite of our SSE module, or further improvements in the OpenMP threading.

Goal: 30% peak performance

Applications



# **Conclusion and Outlook**

Modern super-computers with 10k nodes and several dozens of threads per node have 100k and more possible MPI ranks. This can be a heavy burden for load balancing as well as for MPI itself: For 4+ island runs fallback options in IBM-MPI were necessary. We suspect that the initial load balancing issue, which we observed, can be attributed to an MPI-routine. Something similar we encountered in our XEON-PHI adaption. Intra node parallelism seems necessary and delivers good results. The usage of the logical cores through OpenMP showed a 25% performance gain like the 30% we found with PHI tests. In cooperation with the LRZ application support the full machine job was redesigned. Our former OpenMP implementation was reactivated and enhanced for SuperMUC. With the modifications a full 16 island run (8.192 nodes) was possible during Block Operation in August 2014 and a great success. The peak performance of 150 billion particles/s verified the expected performance of all 131.072 cores. As the problem size was small and the communication pattern not yet fully adapted to OpenMP overall performance corresponded to an MPI-only 4 island run. With LRZ support a code version bypassing the suspicious load balancing functions was developed. A 2 island run demonstrated that our ansatz was successful, hinting at performance levels close to the one of the red dashed circle in Fig. 3a. Due to the premature end of Block Operation 16 and 18 island tests unfortunately

In conclusion, the PSC shows good scaling up to 8 islands, even with the bottlenecks that come with inhomogeneous

could not be carried out.

problems. With the application support of the LRZ and their contacts a new file layout, different approaches for initial balancing and OpenMP threading for SuperMUC have been implemented. A full machine run has been demonstrated.

The next step is a further increase of performance. For the next Block Operation it is planned to push the present 5 % - 14 % performance level of our code up to 30 % sustained performance by porting improvements in our SSE/AVX(512) implementation to SuperMUC.

# **Core Developers**

Prof. Dr. Hartmut Ruhl, Prof. Dr. Kai Germaschewski, Karl-Ulrich Bamberg, Nils Moschüring, Fabian Deutschmann, Steve Abott, Constantin Klier, Simon Jagoda

# Acknowledgement

K.-U. Bamberg and H.Ruhl acknowledge the support of the Arnold Sommerfeld Center for Theoretical Physics at the Ludwig Maximilians University as well as useful editorial suggestions from Patrick Böhl. This work was supported by Grant No. DFG, FOR1048, RU633/1-1, by SFB TR18 project B12 and by the Cluster-of-Excellence Munich-Centre for Advanced Photonics (MAP).

contact: Karl-Ulrich Bamberg, Karl-Ulrich.Bamberg@physik.unimuenchen.de

# Towards Environmental Computing on e-Infrastructures

While distributed computing infrastructures such as the Grid and Clouds have been successfully applied for a number of scientific fields (e.g. high energy physics, astrophysics), there remain many other sciences with different requirements, where todays solutions are not applicable. Yet, the observed shift to computational sciences as a third scientific pillar, in particular in disciplines that traditionally do not leverage large-scale computing facilities, is continuously introducing a spectrum of new services and tools for a wide range of science and research use cases. Rather than outlining such structures as a single integrated e-Infrastructure (Grids, Clouds), it will be far more advantageous to provide sets of well integrated core infrastructure services and a variety of facilities together with a broad set of adaptable and sustainable tools. This enables research communities to select resources, services and tools as required for their specific scientific applications.

One such multi-disciplinary scientific field with many research use cases is coined as "Environmental Computing (EC)". EC refers to a special instance of e-Science, which is both related to solving e-scientific issues in the broadest context of ecological research, and to supporting the mitigation of risks in cases of sudden hazards induced by environmental events. Typical examples of EC thus range from multi-scale simulations of climate change effects, to flash flood simulations for civil protection authorities, to assessments of changes in hydrologic cycles due to droughts, floods, or salinization of coastal aquifers.

While heatedly disputed in past years, it has recently been agreed that global warming in the earth's climate system is unequivocal [1]. Consequently, topics like climate change, water sustainability, hazard analysis, or CO<sub>2</sub> sequestration rank high on the list of challenges that must be addressed urgently. There is a common understanding that solving these challenges requires coordinated efforts in providing new computational models, new algorithms, advanced visualization technologies, and techniques to adequately cope with "big data", and new collaboration patterns between researchers of various scientific domains. Examples of the latter are virtual research environments in various flavors, the increasingly important integration of citizen scientists [2], and organizational bridges between computational sciences and IT services as for example discussed in [3].

EC can coarsely be arranged along two dimensions: The "aspects" dimension integrates various perspectives to look at EC; the "activities" dimension investigates observed phenomena from these perspectives (see Fig. 1).

A number of initiatives are beneficial for this EC vision: the publicly funded Partnership for Advanced Computing in Europe (PRACE, http://www.prace-ri.eu/); the European Grid Infrastructure (EGI, http://

• Karl-Ulrich Bamberg

• Hartmut Ruhl

and Plasma

Munich

Physics, LMU

Computational

Applications



Figure 1: Vision of Environmental Computing.

www.egi.eu/); the Extreme Science and Engineering Discovery Environment (XSEDE, https://www.xsede.org/) in the US; privately supported efforts by resource providers like the Leibniz Supercomputing Centre (LRZ, http://www. Irz.de/); projects encouraged by nonprofit organizations and governmental authorities like the United Nations Office for Disaster Risk Reduction (UNISDR, http://www.unisdr.org/), the European Geosciences Union (EGU, http://www. egu.eu/), or the FP7/Horizon 2020 programmes initiated by the European Commission. Technically speaking, "solution enabling" means providing resources, software and services that aim at overcoming (some of) the EC-challenges as for example outlined – from a more general perspective - in Europe's 2030 vision [4].

In particular, innovative solutions for model coupling, multi-scale computing, big data integration, visualization, dynamic data driven application steering - just to mention a few, are necessary. There are several service frameworks available today like the ones provided by

**Opportunities & Benefits Climate Change Responses** Local/Regional Forecasting **Project Funding Schemes** Centers of Excellence Transdisciplinary Education Market Opportunities

EU- or NSF-funded projects Multiscale Applications on European e-Infrastructures (MAPPER, http://www.mapperproject.eu), Distributed Research Infrastructure for Hydro-Meteorology (DRIHM, http://www.drihm.eu/), Virtual Earthquake and seismology Research Community in Europe e-science environment (VERCE, http://www.verce.eu/), Climate Induced Changes on the Hydrology of Mediterranean Basins (CLIMB, http:// www.climb-fp7.eu), Standards-Based Cyberinfrastructure for Hydrometeorology (SCIHM, http://scihm.org/), or the Australian water resources projects conducted by the Water Information Research and Development Alliance (WIRADA, http://www.csiro.au/). While these projects propose interesting solutions, a more general view including standardization issues, High Performance Computing in real time, dynamic data fusion and application steering, and dependability management is beyond their scope.

Fig. 1 throws light on a generic EC framework. The EC-funded project DRIHM [5] is a typical example for a specific instantiation of the framework. The project aims at configuring, implementing and operating an e-Infrastructure for hydro-meteorological research (HMR) by facilitating a transition from executing isolated models on proprietary systems to chained models on production e-Infrastructures like EGI and PRACE. Fig. 2 depicts this transition from left to right.

For orchestrating simulation chains (workflows), DRIHM offers a web-based portal. Single workflow components



Figure 2: Transition from stand-alone tool execution to a flexible chained execution on e-Infrastructures.

(jobs) are scheduled for execution on appropriate compute elements provided for the respective e-Infrastructure. Executing the jobs may require the dynamic configuration of specific tools, storage facilities for (intermediate) results, staging of (input and output) data, and transparent set up of simulation environments. If the workflow consists of several jobs, this sequence is repeated iteratively until the final results can be presented to the end user.

EC applications in general and DRIHM applications in particular, are typically

<sup>1</sup> Patent pending by Brodie (IBM): US Patent Application 20080209269.

Autumn 2014 • Vol. 12 No. 2 • inSiDE

associated with several constraints. Examples are wall-clock timing constraints for flood predictions and resource constraints for multi-scale simulations to investigate the economic damage of landslides, which requires the coupling of meteorological forecasts with hydrologic simulations in certain basins. A further non-trivial constraint is the ability of EC applications to deliver services that can justifiably be trusted, also known as dependability. The trust in mechanisms

location of schedulers, for example, can significantly be improved if they are able to exploit adequate availability information. However, DRIHM is not only interested in simple binary information, such as availability. Instead, DRIHM also analyzes the causes of erroneous behavior of components through active probing<sup>1</sup> current resource states and mining historic situations.

Workflow chains require the exchange of data between two components in a syntactically and semantically correct manner. While the former describes

Q Role E - End user T - Tool developer A - Administrator A - Datastore administrator A, - Tool administrator A, - Infrastructure administrator Ö Datastore Tool



the capability to process data provided by another model technically, the latter addresses the more difficult task of "understanding" these data. Although there are several solutions already supporting data interoperability (for storing and exposing infrastructure meta data, agentbased discovery services or the Berkeley Database Information Index (BDII) are well accepted; for describing and programming services interfaces, Web Services technologies in all flavors can be applied; the Network Common Data Form (NetCDF, http://www.unidata. ucar.edu/software/netcdf/) provides a standard for sharing array-oriented scientific data), specific EC related ontologies are still in their infancies.

Operationally, the transition from standalone model executions to integrated workflow chains has implications on the execution of the requested tools as there are numerous potential execution environments to cope with. Because e-Infrastructure resources are rarely provided exclusively, EC applications are required to be as portable as possible in the sense of executing the application on a different system without the need of (partial) modification. A special aspect of portability is performance portability, i.e., the ability to reuse code while remaining close to the theoretical peak performance of each computer. To support portability, several standards can be leveraged. Examples are the Unix Filesystem Hierarchy Standard (FHS, http://www.pathname.com/fhs/); code optimization techniques as proposed in the AutoTune project (http://www. autotune-project.eu); performance portability as discussed in [6]; or predictive modeling investigated in [7].

All in all, we hope that our efforts on environmental computing can lead to an improved framework for the research communities in this domain, and further deliver useful experiences and cookbooks for other scientific problems as well.

# Acknowledgment

The authors wish to thank the members and partners of the Munich Network Management (MNM) Team for helpful discussions and valuable comments. More information about the MNM Team is available at http://www.mnm-team. org. This work is supported in parts by the DRIHM project, funded by the European Commission under the 7th Framework Programme (G.A. no. 283568).

# References

 Stocker, T. F.
 Climate Change 2013: The Physical Science Basis, Intergovernmental Panel on Climate Change, Ed. Cambridge: Cambridge University Press, 2013

# [2] European Commission

"Green Paper on Citizen Science: Citizen Science for Europe: Towards a better society of empowered citizens and enhanced research," Brussels, Belgium, 2013

# [3] Frank, A., Jamitzky, F., Satzger, H., Kranzlmüller, D.

"In Need of Partnerships – An Essay about the Collaboration between Computational Sciences and IT Services," Proceedings of the Workshop on Bridging the HPC Talent Gap with Computational Science Research Methods (BRIDGE) at the International Conference of Computational Sciene (ICCS), Cairns, Australia, 2014

# [4] Wood, J. (Ed.)

"Riding the Wave: How Europe can gain from the rising tide of scientific data," Final Report of the High Level Expert Group on Scientific Data, 2010

# [5] Parodi, A., Rebora, N., Kranzlmüller, D., Clematis, A., Schiffers, M., Galizia, A., D'Agostino, D., Quarati, A., Cros, P., Harpham, Q., Jagers, B., Danovaro, E. Bedrina, T.

"DRIHM: Distributed Research Infrastructure for Hydro-Meteorology," in Proceedings of the 7th International Conference on System of Systems Engineering (SoSe 2012). Genoa, Italy: IEEE, July 2012

# [6] Bigot, J., Hou, Z., Perez, C., Pichon, V.

"A Low Level Component Model Enabling Performance Portability of HPC Applications" in Proceedings of the 2012 SC Companion: High Performance Computing, Networking, Storage and Analysis (SCC). IEEE Computer Society, 2012, pp. 701–710

# [7] Menzies, T., Koru, G.

"Predictive Models in Software Engineering," Empirical Software Engineering, vol. 18, no. 3, pp. 433–434, 2013

contact: Michael Schiffers, micheal.schiffers@ifi.lmu.de





# ExaStencils Advanced Stencil-Code Engineering

The German Research Foundation's German Priority Programme 1648 "Software for Exascale Computing" (SPPEXA) is nearing the end of its second of six years. 13 projects started in January 2013 to address various challenges of exascale computing. In this issue, we present project ExaStencils.

ExaStencils is being pursued by eight principal investigators of five research groups at three locations. At the University of Passau, there are the Chairs of Programming (Christian Lengauer and Armin Größlinger) and of Software Product Lines (Sven Apel), at the Friedrich-Alexander-Universität Erlangen-Nürnberg the Chairs of System Simulation (Ulrich Rüde and Harald Köstler) and of Hardware-Software-Co-Design (Jürgen Teich and Frank Hannig), and at the University of Wuppertal the Applied Computer Science Group (Matthias Bolten).

The central goal of ExaStencils is to develop a radically new software technology for applications with exascale

performance. To reach this goal, the project focuses on a comparatively narrow but very important application domain. The aim is to enable a simple and convenient formulation of problem solutions in this domain. The software technology developed in ExaStencils shall facilitate the highly automatic generation of a large variety of efficient implementations via the judicious use of domain-specific knowledge in each of a sequence of optimization steps such that, at the end, exascale performance results.

The application domain chosen is that of stencil codes, i.e., compute-intensive algorithms in which data points in a grid are redefined repeatedly as a combination of the values of neighboring points. This neighborhood pattern is called a stencil. Stencil codes are used for the solution of discrete partial differential equations and the resulting linear systems. To obtain a performance-competitive, highly automated software technology, the domain is restricted further to multigrid methods [1].



Figure 1: An F-cycle as a succession of V-cycles.



Figure 2: Successive refinement of a hierarchical hybrid grid (left to right).

Multigrid methods involve stencil computations on a hierarchy of very fine to successively coarser grids. On the coarser grids, less processing power is required and communication dominates. A multigrid method is characterized by two strategies: (1) a smoothing strategy, which is used to smooth the sampling error of the grid at hand, and (2) a coarsening strategy, which transfers data from one grid to the next coarser grid. Once one arrives at the coarsest level, one refines the grid again via some form of interpolation. This cycle of coarsening and refining is called a V-cycle. Various cycling strategies are commonly used. For instance, an F-cycle multigrid method consists of a sequence of progressively deeper V-cycles (see Fig. 1). The technology for the efficient implementation and a systematic performance engineering of parallel multigrid methods is a major current research topic [2].

ExaStencils also restricts the structure of the grids. Considered are so-called hierarchical hybrid grids: at the coarsest level, the grid is unstructured, but refinements of each segment must be homogeneous, though each segment may exhibit a different structure (see Fig. 2).



concrete solver 4 implementation

Figure 3: The ExaStencils DSL with four layers of abstraction. Layers 1-2 address the concerns of application scientists, Layers 2-3 those of mathematicians, and Layers 3-4 those of computer scientists.



Present-day stencil codes are implemented in general-purpose programming languages, such as Fortran, C, or Java, or derivates thereof, and harnesses for parallelism, such as MPI, OpenMP or OpenCL. ExaStencils favors a much more domain-specific approach with languages at several layers of abstraction, the most abstract being the mathematical formulation, the most concrete the optimized target code. At every layer, the corresponding language expresses not only computational directives but also domain knowledge of the problem and platform to be leveraged for optimization. This approach will enable a highly automated code

Continuous Domain & Continuous Model

Discrete Domain & Discrete Model

Algorithmic Components & Parameters

**Complete Program Specification** 

Hardware Description

generation at all layers and has been demonstrated successfully before in the U.S. projects FFTW [3] and SPIRAL [4] for certain linear transforms.

At the center of the project is a Scalabased code generator for a wide range of stencil codes in the domain, which is currently under development. It takes code formulated in an external domainspecific language at four different layers of abstraction (see Fig. 3). The ultimate vision is that application scientists will program at the most abstract layers, and the additional information specified at the lower layers will be generated automatically based on an analysis of the specific problem to be solved and information on the execution platform at hand. A preliminary version of the code generator in existence already demonstrates the feasibility of the ExaStencils approach [5]. It produces target code in C++ with OpenMP and CUDA. Distributed architectures are one of the main focuses of the next version of the generator framework, which is currently under development.

One major innovation in ExaStencils is that it views stencil codes not as individuals but as members of a family. The domain-specific specification pinpoints the commonalities that the code shares with the other codes of the family, and the variabilities in which it departs from the other codes. Each point of variability comes with a number of options or alternatives. The idea is that application scientists, and the ExaStencils compiler and run-time system, choose suitable options from these variabilities - and no more has to be specified to obtain a custom-optimized implementation.

In first, yet hand-coded, experiments on Jülich's BlueGene/Q JUQUEEN, we employed the Highly Scalable Multigrid Solver [6] for hierarchical hybrid grids. Commonalities and variabilities are usually specified in terms of a variability model. The variability model for the Highly Scalable Multigrid Solver is illustrated in Fig. 4. Each node denotes a configuration option - in our case, the choice of a coarse grid solver, a



Legend:	
IP_CG = In-Place Conjugate Gradient	Jac = Jacobi
IP_AMG = In-Place Algebraic multigrid	GS = Gauss-Seidel
RED_AMG = Algebraic multigrid with data reduction	RBGS = Red-Black Gauss-Seidel
GSAC = Gauss-Seidel with additional communication	BS = Block-Smoother
RBGSAC = Red-Black Gauss-Seidel with additional communication	

Figure 4: Concrete variability model for the Highly Scalable Multigrid Solver.

smoother, and pre- and post-smoothing parameter values which must satisfy the condition that their sum is greater than zero. A selection of configuration options gives rise to an executable variant of the stencil code.

Which configuration options (i.e., which choices of algorithmic components, alternatives of data structures, and parameter values) contribute to maximal performance is obvious in some cases and very surprising in others. To make this problem tractable, ExaStencils will provide a capability of recommending suitable combinations of configuration options, based on a machine-learning approach [7].

With project ExaStencils, we hope to provide proof of the application relevance of the ExaStencils paradigm of domainspecific stencil code engineering and to encourage experts of other suitable domains to take a similar approach. For up-to-date information, please visit the project's Web site at www.exastencils. org.

# Acknowledgements

We gratefully acknowledge the financial support fo the Priority Research Initiative 1648 "Software for Exascale Computing", funded by the German Research Foundation.

# References

[1] Trottenberg, U., Osterlee, C. W., Schüller, A. Multigrid, Academic Press, 2000

[2] Gmeiner, B., Köstler, H., Stürmer, M., Rüde, U. Parallel multigrid on hierarchical hybrid

grids: A performance study on current High Performance Computing clusters. Concurrency and Computation: Practice and Experience 26(1), pp.217-240, Jan. 2014

[3] Frigo, M., Johnson, S. G. Feb. 2005

Voronenko, Y. Springer, 2011

[5] Köstler, H., Schmitt, C., Kuckuk, S., Hannig, F., Teich. J., Rüde, U. A Scala Prototype to Generate Multigrid Solver Implementations for Different Problems and Target Multi-Core Platforms. Computing Research Repository (CoRR), pp.18, arXiv:1406.5369, June 2014

# Rüde, U.

A generic prototype to benchmark algorithms and data structures for hierarchical hybrid grids. In: Proc. Int. Conf. on Parallel Computing (ParCo), pp.813-822, IOS Press 2013

Jan. 2014

contact: Prof. Dr. Christian Lengauer, lengauer@fim.uni-passau.de

The design and implementation of FFTW3, Proc. IEEE 93(2), pp.216-231,

# [4] Püschel, M., Franchetti, P.,

Spiral. In: Encyclopedia of Parallel Computing, 1920-1933, Padua, D. A. et al. (eds.),

# [6] Kuckuk, S., Gmeiner, B., Köstler, H.,

### [7] Grebhahn, A., Siegmund, N., Apel, S., Kuckuk, S., Schmitt, C., Köstler, H.

Optimizing Performance of Stencil Code with SPL Conqueror. In: Proc. Int. Workshop on High-Performance Stencil Computations (HiStencils), Größlinger, A. and Köstler, H. (eds.), pp.7-14, www.epubli.de,

- Sven Apel<sup>1</sup>
- Matthias Bolten<sup>2</sup>
- Armin Größlinger<sup>1</sup>
- Frank Hannig<sup>3</sup>
- Harald Köstler<sup>3</sup>
- Christian Lengauer<sup>1</sup>
- Ulrich Rüde<sup>3</sup>
- Jürgen Teich<sup>3</sup>
- <sup>1</sup> University of Passau
- <sup>2</sup> University of Wuppertal
- <sup>3</sup> Friedrich-Alexander-Universität Erlangen-Nürnberg

# EXCESS Execution Models for Energy-Efficient Computing Systems

Information and Communication Technologies (ICT) play an important role in increasing the energy efficiency in the economy and contributing to sustainable growth in Europe. In order to achieve the ambitious goals on energy efficiency by 2020, Europe needs to improve the energy efficiency of ICT systems and use them as an enabler to improve energy efficiency across the economy. Hence, energy efficiency is becoming a leading design constraint in current and future ICT systems.

EXCESS

EXCESS [1] scientific and technological concept in addressing energy efficiency is defined by novel execution models between common High Performance Computing (HPC) infrastructures and Embedded Systems (ES). The vision of the EXCESS project is to develop energy, platform and component models that will be applicable to both embedded processors and general purpose ones. EXCESS will demonstrate this by providing resource- and energy-aware programming models, a re-targetable, generic tool chain for generating energy-optimized code, and adaptive libraries, which together can address energy efficiency issues for both classes of systems.

The EXCESS project is gathering a clear understanding of where energyperformance is wasted and will develop a continuous process to reduce the energy waste. This will lead EXCESS to achieve significant improvements in energy efficiency for computing systems. To reach the next level of energy



Figure 1: Iterative software/hardware co-design.

efficiency, the interaction of hardware and software needs to be optimized through an iterative software/hardware co-design process. EXCESS is concerned with developing new energyaware execution models that will allow holistic energy modeling and optimization of software and hardware across the whole system software stack,

ranging from the application programs via data structures and algorithms as well as libraries and the run-time system down to the actual hardware.

In order to synthesize energy optimization in embedded computing and performance optimization in HPC, we need to bridge the gaps between embedded technologies that are typically developed for small-scale customized devices and applications (e.g. smartphones and their applications) and HPC technologies which are demanded by large-scale applications and the corresponding systems. One approach to reach the objective is presented in Fig. 2 where the resource distribution of the task between HPC and the Embedded Cluster is shown. The Embedded Cluster (Movidius Cluster) is represented by the Myriad low-power, multi-core digital signal processor (DSP) architecture developed by EXCESS partner Movidius. The HPC cluster consists of standard x86 multicore CPU based servers partly extended by general purpose (Nvidia) GPUs.

EXCESS will take a holistic approach and will introduce novel programming methodologies to drastically simplify the development of energy-aware applications. These applications will be energy-portable in a wide range of computing systems while preserving relevant aspects of performance. The EXCESS project is going to be driven by the following technical components that will be developed during the lifetime of the project:

• Complete software stacks (including programming models, libraries/algorithms and runtimes) for energy efficient computing.

• Uniform, generic development methodology and prototype software tools that enable leveraging additional optimization opportunities for energy-efficient computing.

• Configurable energy-aware simulation systems for future energy-efficient architectures.

### Module server: streams algorithms sort algorithms graph computation discrete diff. operators generation vector algebra dense algebra —Network. sparse linear algebra fft

image processing

Figure 2: Execution topology on the EXECSS Cluster.

The three-year European FP7 project EXCESS, which started in September 2013, will run for 36 months with an overall budget of 3.31 million Euro and a funding of 2.5 million Euro by the European Commission.

The role of the High Performance Computing Center Stuttgart (HLRS) of the University of Stuttgart (USTUTT) is to coordinate two work packages. HLRS will be responsible for developing a runtime and monitoring framework for energy analysis, and for evaluating the outcome of the technical work pack-



65

ages with respect to their efficiency in terms of the project objectives. The Monitoring Framework will be used to provide reliable performance and energy measurements. This monitoring system will be able to collect metrics and other information necessary to assess energy consumption during the application runtime. HLRS currently builds the HPC platform and will integrate it with the embedded system contributed by partner Movidius.

Together with the industrial participant Movidius, HLRS leverage the excellent ecosystem contacts to bring in a comprehensive set of real world applications and application kernels that will guide the project. Finally, HLRS represents the conduit to actual multi- and many-core technologies that cover a wide range of hardware from embedded to high performance computing. The expertise in energy efficiency of all the partners involved in the EXCESS project is the key to make sure that EXCESS can have a sustainable impact on the state of the art in Eco-efficiency and that its results can be effectively exploited.

The rationale behind the EXCESS concept is apt to become a widely applicable solution for data centres, high performance providers, cloud computing providers as well as various other sectors, like mobile phones to optimize performance in conjunction with energy efficiency. However, as modern high performance clusters don't offer the capabilities to monitor energy consumption and performance parameters on thousands of cores, the results will be developed, evaluated and validated on rather small infrastructures.

# **Project Partners**

The EXCESS consortium unites Europe's leading experts in HPC as well as embedded computing. The consortium consists of world-class research centres, universities and companies that bring in the required expertise to accomplish the ambitious, but realistic goals of EXCESS:

- High Performance Computing Center Stuttgart (HLRS), Germany [2]
- Chalmers Tekniska Hoegskola AB (Chalmers), Sweden [3]
- Linkopings Universitet (LIU), Sweden [4]
- Movidius LTD (Movidius), Ireland [5]
- Universitetet I Tromsoe (UiT), Norway [6]

# References

- [1] "EXCESS Execution Models for Energy-Efficient Computing Systems," [Online]. Available: http://excess-project.eu/
- [2] "Home High Performance Computing Center Stuttgart (HLRS)," [Online]. Available: http://www.hlrs.de/
- [3] "Chalmers University of Technology," [Online]. Available: http://www.chalmers. se/en/Pages/default.aspx
- [4] "Linköping University," [Online]. Available: http://www.liu.se/?l=en
- [5] "Movidius" [Online]. Available: http://www. movidius.com/
- [6] "Start page UiT," [Online]. Available: http://en.uit.no/startsida

contact: Bastian Koller, koller@hlrs.de Uwe Küster, kuester@hlrs.de Yosandra Sandoval, sandoval@hlrs.de Dmitry Khabi, khabi@hlrs.de Michael Gienger, gienger@hlrs.de

# **HONK - High-Order Methods** for the Simulation of **Complex Flow Phenomena**

Numerical simulations are an essential technique for research and development in the field of engineering. They are a vital tool to ensure cost effectiveness, safety and a low ecological footprint of new products, and are essential for the competitiveness of German industry. But there remain several unsolved problems, especially in fluid dynamics. For complex simulation scenarios, e.g. problems depending on a variety of physical phenomena, excessive time-to-solution renders the integration of techniques into industrial development cycles unfeasible. For decades simulation methods have benefited from the constant increase in CPU clock frequency. Nowadays, the situation has changed. While there are still small improvements in the architecture of CPUs, the clock frequency barely increases, due to thermal and manufacturing constraints. Hence, to



push the performance of modern computers to higher levels, the number of CPUs and processing units (e.g. cores) per CPU are increased. To use the full computing capacity of these modern multi- and many-core machines, like the Cray XE6 at the High Performance Computing Center Stuttgart (HLRS), it is therefore necessary to develop and implement scalable numerical algorithms.

One of the emphases of the numerical research group of the Institute for Aerodynamics and Gas Dynamics (IAG) of the University of Stuttgart is the highly scalable discontinuous Galerkin (DG) method. The efficiency of the DG



Figure 1: Strong scaling of a code based on a discontinuous Galerkin method on the Jugene.

Bastian Koller

University of

Stuttgart (HLRS),

Uwe Küster

Germany

# HONK (S)

67

method on a state of the art peta-scale system (Jugene) is shown in Fig. 1.



Figure 2: Density profile of an exemplary 2D calculation (670.000 DOF), where gas enters the simulation volume via two valves from below. The data have been visualized using ParaView (http://www.paraview.org/) in combination with a reader plugin that was devoled within the Honk project so that FLEXI's output files can directly be evaluated with ParaView.

> The HONK project focuses on the CFD code "FLEXI", which was developed specifically for highly scalable simulations on modern supercomputers and targets the efficient simulation of highly complex transient fluid flows. FLEXI is based on a special version of the discontinuous Galerkin method; the discontinuous Galerkin Spectral Element Method (DGSEM). The efficiency of FLEXI is based on a tensor product basis of polynomials and the use of the same interpolation and numerical integration points inside the elements. With this approach FLEXI can handle general, unstructured hexahedral element meshes [1]. The parallelization is

based on MPI domain decomposition with non-blocking communication. The resulting latency hiding leads to a high parallel efficiency [2].

Today's commercial CFD software packages are mainly designed for the simulation of time averaged steady solutions with turbulence models. This results in a broad applicability of the available commercial codes. For inherently unsteady problems or simulations with high fidelity turbulence models, such solvers are not efficient. In contrast, FLEXI is designed for complex unsteady systems, using a high temporal and spatial resolution. Over the last years, the simulation of such complex fluid flows has come increasingly into the focus of industrial research. The industrialization of such a novel DG code and the application to industrial problems are the main objectives of the HONK project, which is lead by the Robert Bosch GmbH.

The goal of the HONK project is to achieve efficient simulations of pneumatic and hydraulic components that cannot be simulated satisfactorily with today's commercial CFD software tools. These complex flow scenarios include physical phenomena on different spatial and temporal scales, such as phase transitions and acoustic waves and have not yet been sufficiently studied. The project partners combine many years of expertise in the simulation and experimental study of complex flows. The latter are used for validation and exploration purposes. One of the corresponding real world applications is the injection of Compressed Natural Gas (CNG) into the intake manifold of a bi-fuel car.



Figure 3: Visualization of a high-order DG solution of the flow around a sphere with direct volume rendering.

The numerical research group of the IAG has been working for many years on the CFD code FLEXI and is actively extending the code to handle the variety of cases considered in the project: Applications like the injection of CNG in an intake manifold require a stringent representation of shocks that can occur during a high pressure injection process. Further emphasis lies on the simulation of hydraulic components. Applications in hydraulics are often characterized by the occurrence of high pressure and velocity gradients in the flow, which can cause phase transitions. The resulting two-phase mixture can only be captured when considering a compressible fluid, with a highly accurate equation of state. Thus, to simulate production-relevant processes, like, e.g. cavitation, the code will be extended to general complex equations of state for technical gases and multiphase mixtures.

These complex simulation scenarios require a modern supercomputer, like Hermit (Cray XE6) or Hornet (Cray XC40) at the High Performance Computing Center Stuttgart (HLRS), in order to yield results in a reasonable time. To benefit from the enormous computing power of these machines, sophisticated strategies for parallel execution and load balancing are required. Modern HPC architectures are characterized by a hierarchical network of computing cores. A small number of cores are closely linked on a node and share their memory. Many such nodes are connected via a network. For the parallelization of applications, there are two basic strategies: distributed and shared memory. The standard for distributed memory is MPI which is used in FLEXI. For the shared-memory concept, different competing approaches are used, e.g. OpenMP and StarSS. A hybrid parallelization strategy, which considers the hierarchical architecture explicitly, can achieve a significantly

better performance than pure distributed parallelizations. Several successful hybrid parallelizations have been performed at HLRS [3]. Different concepts for the shared-memory parallelization will be implemented in FLEXI and compared with each other. The most successful will be hybridized with the existing MPI parallelization and will be implemented in the production code. A challenge for parallelization is load imbalances, due to varying costs depending on the local occurrence of complex phenomena, e.g. shocks, which will be addressed in the hybrid approach.

The solutions based on discontinuous Galerkin methods provide new challenges also for visualization techniques. The Visualization Research Center (VISUS) of the University of Stuttgart is one of the leading visualization institutes in Europe, and has in-depth experience with flow visualization. Most current visualization approaches for DG solutions apply traditional visualization techniques by employing resampling to the DG data. However, resampling produces a substantial overhead in both computing time and memory requirements and causes loss of information. One of the research priorities at VISUS is to develop methods for direct and efficient analysis of DG solutions, like by direct volume rendering [4].

Because of the high spatiotemporal complexity and the size of the systems considered in HONK, a close integration of visualization and simulation is required, both in terms of implementation and execution, by means of in-situ visualization. This also requires new strategies for load balancing between simulation and visualization. The specific physical phenomena which are of interest in this project, e.g. cavitation, also require new specific visualization techniques that will be developed in the course of the project.

The aim of HONK is to achieve a reliable and efficient simulation process for complex fluid flows and integrated visualization of the resulting high-order DG solutions. Currently, cutting-edge methods for simulating fluid flow are implemented in scientific codes only. HONK will lift these methods into a production-ready, fully functional code, ready to address the challenges of industrial reliability requirements. To promote the sustainability of the project, the results will be launched as open source packages.

# Who is HONK?

HONK is a cooperation between academia and industry. The research institutions are from the University of Stuttgart – HLRS, IAG and VISUS. The industrial partner is the Robert Bosch GmbH. HONK started in September 2013, runs for three years and is funded by the BMBF (Bundesministerium für Bildung und Forschung, German Federal Ministry of Education and Research).

# **References:**

 Hindenlang, F., Gassner, G., Altmann, C., Beck, A., Staudenmaier, M., Munz, C.-D.
 Explicit Discontinuous Galerkin methods for unsteady problems, Computers and Fluids, 61, pp. 86-93, 2012

# [2] Altmann, C., Beck, A., Hindenlang, F., Staudenmaier, M., Gassner, G., Munz, C.-D.

An Efficient High Performance Parallelization of a Discontinuous Galerkin Spectral Element Method, (Incollection) Keller, R., Kramer, D., Weiss, J.-P. (Ed.): Facing the Multicore-Challenge III, 7686, pp. 37-47, Springer Berlin Heidelberg, 2013 [3] Gracia, J., Niethammer, C., Hasert, M., Brinkmann, S., Keller, R., Glass, C.W. Hybrid MPI/StarSs-a case study. In Proceedings of the 10th IEEE International Symposium on Parallel and Distributed Processing with Applications (ISPA), pages 48–55, 2012

# [4] Üffinger, M., Frey, S., Ertl, T. Interactive high-quality visualization of higher-order finite elements, Computer Graphics Forum, 29, 337-346, 2010

contact: Sebastian Boblest, sebastian.boblest@visus.uni-stuttgart.de Colin W. Glass, glass@hlrs.de Philipp Offenhäuser, offenhaeuser@hlrs.de Filip Sadlo, sadlo@visus.uni-stuttgart.de Malte Hoffmann, hoffmann@iag.uni-stuttgart.de

- Sebastian Boblest<sup>1</sup>
- Colin W. Glass<sup>2</sup>
- Philipp Offenhäuser<sup>2</sup>
- Filip Sadlo<sup>1</sup>
- Malte Hoffmann<sup>3</sup>
- <sup>1</sup> University of Stuttgart (Institute for Visualization VISUS)
- <sup>2</sup> University of Stuttgart (HLRS)
- <sup>3</sup> University of Stuttgart (Institute of Aerodynamics an Gas Dynamics IAG)

# Successful Technology Transfer with Siemens -The RAPID Project

# JÜLICH SIEMENS

Siemens is a true giant when it comes to software. Its products, systems, and solutions are built on billions of Euro invested in software R&D. A major part of its software portfolio was developed for single-core processors. In the future, however, there will be fewer and fewer single-core chips, with

the consequence that both the existing software portfolio and new software will need to be prepared for use on multi-core processors. An additional problem with existing software is that it has grown over time and parts of it have not been touched for long periods. These parts work, but no one is really



Figure 1: Profile data of a Pthreads application visualized in the Cube performance report explorer.



Figure 2: RAPID successfully applies HPC technology in industrial software development. © Forschungszentrum Jülich/fotolia.com.

familiar with the details anymore. In some cases, the original developers now work elsewhere or have retired.

# Optimizing Software to Work on Multi-Core Processors

If subprograms that previously worked sequentially on one processor are simply distributed across parallel CPU cores on a multi-core processor, the usual problems of parallel programs emerge: race conditions when accessing shared data and deadlocks when trying to synchronize the access to this data. A particular pitfall is that these errors often do not occur at all on single-core processors and happen in a non-deterministic way in a multicore environment. As a result, developers often build many synchronization operations into their applications. But while this eliminates race conditions, the chance for deadlocks increases. Additionally, too many synchronization operations slow down the applications – sometimes making them even slower than the single-core versions.

To remedy this situation, tools to examine the communication and synchronization aspects of the often huge software packages are needed.

In High Performance Computing (HPC), these tools already exist for a long time. Tools like the measurement and instrumentation framework Score-P [2] provides profiling and tracing data for performance analysis and visualization tools like Scalasca [3], Vampir [4], TAU [5], and Periscope [6]. However, these established tools are targeted towards the prevailing HPC programming para-

73

digms, namely MPI and OpenMP. As these tools already exist, it seems natural to contribute to and enhance the existing tools rather than to develop new tools from scratch.

# RAPID: Runtime Analysis of Parallel applications for Industrial software Development

As the Siemens software originates more from an embedded system's ecosystem than from HPC, the mentioned tools cannot be applied out-of-the-box. This is why Siemens Corporate Technology [12] and Forschungszentrum Jülich, Jülich Supercomputing Centre, collaborate in the project RAPID - Runtime Analysis of Parallel applications for Industrial software Development [1]. Goal of this project is to adapt the measurement and analysis tools Score-P and Scalasca to serve Siemens' needs. In particular, support for new threading models like POSIX threads, Windows threads, Qt threads [7], and ACE threads [8] are integrated into Score-P. In addition, support for leveraging task parallelism using MTAPI [9] is being developed. Besides supporting new programming paradigms, additional work has to be done with regards to portability. Although Score-P is already quite portable as it is running on all relevant supercomputer architectures, systems like Windows and operating systems for embedded systems have not been targeted so far. On the analysis side, new methods targeting thread-based synchronization patterns, e.g., a lockcontention analysis, are being implemented in Scalasca.

Another goal of RAPID is to assist developers in getting insight into the huge and often complex software packages by providing visual call graphs of application runs, thus overcoming the limitation of static analysis tools when it comes to the use of polymorphism or indirect calls via function pointers. To accomplish this, the established tool Cube [10] is enhanced by providing a generic plugin interface that allows various kinds of analyses on the profiling data generated during a measurement run.

The mentioned tools Score-P, Scalasca, and Cube come with the 3-clause BSD open source license [11]. All contributions to these tools developed within RAPID will be made available under the same license, which will allow the community to further improve and maintain them.

For more information please visit http://www.fz-juelich.de/ias/jsc/EN/ Research/HPCTechnology/PerformanceAnalyse/RAPID/rapid.html

# References

- [1] http://www.fz-juelich.de/ias/jsc/EN/Re search/HPCTechnology/PerformanceAnal yse/RAPID/rapid.html
- [2] http://score-p.org/
- [3] http://scalasca.org/
- [4] http://vampir.eu/
- [5] http://tau.uoregon.edu/
- [6] http://www.lrr.in.tum.de/~periscop/
- [7] http://qt-project.org/
- [8] http://www.dre.vanderbilt.edu/~schmidt/ ACE.html
- [9] http://www.multicore-association.org/ workgroup/mtapi.php
- [10] http://www.scalasca.org/software/cube-4.x/download.html
- [11] http://opensource.org/licenses/BSD-3-Clause
- [12] Siemens AG, Corporate Technology http:// www.siemens.com/innovation/en/

contact: Christian Rössel, c.roessel@fz-juelich.de

- Christian Rössel<sup>1</sup>
- Bernd Mohr<sup>1</sup>
- Markus Geimer<sup>1</sup>
- Daniel Becker<sup>2</sup>

<sup>1</sup> Jülich Supercomputing Centre (JSC)

 <sup>2</sup> Siemens AG, Corporate
 Technology, Multicore Expert Center

# The new Intel Xeon Phi based System SuperMIC at LRZ



Figure 1: The new SuperMIC system at LRZ. (Photo: Vasileios Karakasis)

Recently, LRZ has installed a new Intel Xeon Phi based cluster named SuperMIC as a prototype for future manycore technologies. SuperMIC will be partially integrated into SuperMUC and is equipped with 64 Intel Xeon Phi coprocessors based on Intel's new Many Integrated Core (MIC) architecture. Since June 2014, first selected users can login to the system and get familiar with the new architecture. First experiences with the MIC architecture at LRZ can be found in an initial evaluation report [1].

SuperMIC consists of one IBM iData-Plex rack with 32 dx360 M4 nodes. Fig. 1 shows the iDataPlex rack of SuperMIC.

Each node contains two lvy-Bridge host processors E5-2650v2 with 8 cores @ 2.6 GHz each, and two Intel Xeon Phi (MIC) coprocessors 5110P with 60 cores @ 1.053 GHz each.

	Xeon host	Xeon Phi coprocessor
Processor type	lvy-Bridge E5-2650v2	5110P
Number of nodes/coprocessors	32	64
Number of cores per node/coprocessor	2 x 8	60
Total number of cores	512	3840
Frequency of cores	2.6 GHz	1.053 GHz
Number of threads per core	2 (hyperthreading)	4 (hardware threads)
SIMD vector register width	256 bit	512 bit
SIMD instruction set	AVX	IMCI
Flops/cycle	8 (DP), 16 (SP)	16 (DP), 32 (SP)
Theoretical peak perf. per node/coproc.	332.8 GFlops/s (DP)	1.01 TFlop/s (DP)
Theoretical peak performance	10.6 TFlop/S (DP)	64.7 TFlop/s (DP)
Memory size per node/coprocessor	64 GB	8 GB
Total memory size	2048 GB	512 GB
Memory bandwidth per node/coproc.	2 x 59.7 GB/s	320 GB/s

Table 1: System parameters of the SuperMIC cluster at LRZ.

Technical details about SuperMIC are summarized in Table 1.

Linux is deployed as operating system on both the host and the coprocessors. The compute nodes are connected via Mellanox Infiniband FDR14 using OpenFabrics OFED 1.5.4.1. The connection from the host to the attached coprocessors is via PCle 2.0, which limits the host-MIC bandwidth to 6.2 GB/s. Virtualised TCP/IP and Infiniband stacks are provided over the PCle bus, allowing users to access the coprocessors like network nodes. All Intel Xeon Phi coprocessors have a unique IP address and can be directly accessed from the SuperMIC login node and the compute nodes through a virtual bridge interface. IBM LoadLeveler is deployed as batch system and used to control the access to the nodes. Once a node is reserved by the batch

> SuperMUC Login Node

SuperMIC is shown in Fig. 2.

The Intel Xeon Phi coprocessors can be programmed using the Intel C/C++or Fortran compiler and traditional HPC parallelisation techniques like (Intel) MPI and OpenMP. Furthermore, Intel Cilk Plus, Intel Threading Building Blocks, OpenCL, Pthreads and MKL are also supported.

Generally speaking, two main execution modes can be distinguished: offload mode and native mode. In "offload mode" the code is instrumented with

32 Intel Xeon Compute Nodes

SuperMIC

Login Node



Figure 2: Overview of the SuperMIC system components.



77

system, interactive login to the host as well as the attached coprocessors is granted. Every compute node has a local disk attached to it that is configured with two file systems shared by the host and its coprocessors.

An overview of the components of

### 64 Intel Xeon Phi Coprocessors

OpenMP-like pragmas in C/C++ or comments in Fortran to indicate regions of code that should be offloaded to the coprocessor and be executed there at runtime. The generated program must be executed on the host.

In "native mode" the Intel compiler is instructed to cross-compile for the MIC architecture. The generated executable must be copied to the coprocessor or the shared file system and can be launched from within a shell running on the coprocessor interactively or via ssh. In particular, this mode allows to run MPI tasks directly on the coprocessors (using ssh as the task startup mechanism). Hereby various scenarios are possible for MPI programs: MPI tasks can either be executed on up to 64 coprocessors only, on up to 32 compute nodes only (possibly doing offloading) or both on the compute nodes and the coprocessors (using an MPMD-style startup and possibly involving load balancing). Fig. 3 demonstrates the various MPI scenarios.

While on several recently installed Intel Xeon Phi based clusters the coprocessor usage is restricted to offloading, on SuperMIC users can access the Intel Xeon Phi coprocessors both in offload mode and native mode.

Currently access to SuperMIC is only granted to selected users, which already must have an active account on SuperMUC. Users interested in accessing the SuperMIC system are asked to submit a service request via the LRZ Service Desk. Further documentation on the SuperMIC system is available online [2]. In addition, LRZ organises



Figure 3: Various MPI scenarios on SuperMIC. MPI tasks are depictured by blue bars, OpenMP threads by red bars.

training workshops to assist users porting applications suitable for the new system technology.

## References

- [1] Weinberg, V., Allalen, M. First experiences with the Intel MIC architecture at LRZ, inSiDE Vol. 11 No. 2, Autumn 2013
- [2] http://www.lrz.de/services/compute/ supermuc/supermic/

contact: Volker Weinberg, Volker.Weinberg@lrz.de

Volker WeinbergMomme Allalen

Leibniz Supercomputing Centre (LRZ), Germany

79



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

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

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

# Contact:

Leibniz Supercomputing Centre

Prof. Dr. Arndt Bode Boltzmannstr. 1 85748 Garching near Munich Germany

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



Picture of the Petascale system SuperMUC at the Leibniz Supercomputing Centre.

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

	System	Size	Peak Performance (TFlop/s)	Purpose	User Community
		18 thin node islands IBM iDataPlex 512 nodes per island with 2 Intel Sandy Bridge EP processors each 147,456 cores 288 TByte main memory FDR 10 IB	3,185	Capability Computing	PRACE Projects, German universities and research institutes, PRACE projects (Tier-O System)
	IBM System x "SuperMUC"	1 fat node island with 205 nodes IBM Bladecenter HX5 4 Intel Westmere EX processors each 52 TByte main memory QDR IB	78	Capability Computing	
		32 accelerated nodes with 2 Intel Ivy Bridge EP and 2 Intel Xeon Phi each 80 GByte main memory Dual-Rail FDR 14 IB	75	Prototype System	
	Linux-Cluster	510 nodes with Intel Xeon EM64T/AMD Opteron 2-, 4-, 8-, 16-, 32-way 2.030 Cores 4.7 TByte	13.2	Capability Computing	Bavarian and Munich Universities, LCG Grid
	SGI Altix ICE	64 nodes with Intel Nehalem EP 512 Cores 1.5 TByte memory	5.2	Capacity Computing	Bavarian Universities, PRACE
	SGI Altix Ultraviolet	2 nodes with Intel Westmere EX 2,080 Cores 6.0 TByte memory	20.0	Capability Computing	Bavarian Universities, PRACE
	Megware IB-Cluster "CoolMUC"	178 nodes with AMD Magny Cours 2,848 Cores 2.8 TByte memory	22.7	Capability Computing, PRACE Prototype	Bavarian Universities, PRACE
	MAC research cluster	64 Intel Westmere Cores, 528 Intel Sandy Bridge Cores, 1,248 AMD Bulldozer Cores, 8 NVIDIA GPGPU cards, 8 ATI/AMD GPGPU cards	40.5	Testing accelerated architectures and cooling technologies	Munich Centre of Advanced Computing (MAC), Computer Science TUM

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

Centres



# **First German National Center**

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

# Service for Industry

Service provisioning for industry 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 industry always has acces to the most recent HPC technology.

# **Bundling Competencies**

In order to bundle service resources in the state of Baden-Württemberg HLRS has teamed up with the Steinbuch Center for Computing of the Karlsruhe Institute of Technology. This collaboration has been implemented in the non-profit organization SICOS BW GmbH.

# World Class Research

As one of the largest research centers for HPC HLRS takes a leading role in research. Participation in the German national initiative of excellence makes HLRS an outstanding place in the field.

# Contact:

Höchstleistungsrechenzentrum Stuttgart (HLRS) Universität Stuttgart

Prof. Dr.-Ing. Dr. h.c. Dr. h.c. Michael M. Resch Nobelstraße 19 70569 Stuttgart Germany

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



- 1114						
	System	Size	Peak Performance (TFlop/s)	Purpose	User Community	
	Cray XC40 "Hornet" (Q4 2014)	3,994 dual socket nodes with 94,656 Intel Haswell cores	3,786	Capability Computing	European (PRACE) and German Research Organizations and Industry	
	Cray XE6 "Hermit" (Q4 2011)	3,552 dual socket nodes with 113,664 AMD Interlagos cores	1,045	Capability Computing	European (PRACE) and German Research Organizations- and Industry	
	NEC Cluster (Laki, Laki2) heterogenous compunting platform of 2 independent clusters	23 TB memory 9988 cores 911 nodes	170	Laki: 120,5 TFlops Laki2: 47,2 TFlops	German Universities, Research Institutes and Industry	

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



HLRS Cray XC40 "Hornet" Frontdoor Image



View of the HLRS Cray XC40 "Hornet"

Centres



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

# **Provision of supercomputer resources**

of the highest performance class for projects in science, research and industry in the fields of modeling and computer simulation including their methods. The selection of the projects is performed by international peer-review procedures implemented by the John von Neumann Institute for Computing (NIC), GCS, and PRACE.

Supercomputer-oriented research and development in selected fields of physics and other natural sciences by research groups and in technology, e.g. by doing co-design together with leading HPC companies.

Implementation of strategic support infrastructures including communityoriented simulation laboratories and cross-sectional teams, e.g. on mathematical methods and algorithms and parallel performance tools, enabling the effective usage of the supercomputer resources.

Higher education for master and doctoral students in cooperation e.g. with the German Research School for Simulation Sciences.

# Contact:

Jülich Supercomputing Centre (JSC) Forschungszentrum Jülich

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

# **Compute servers currently operated by JSC**

	System	Size	Peak Performance (TFlop/s)	Purpose	User Community
	IBM Blue Gene/Q "JUQUEEN"	28 racks 28,672 nodes 458,752 processors IBM PowerPC® A2 448 Tbyte main memory	5,872	Capability Computing	European Universities and Research Institutes, PRACE
	Intel Linux CLuster "JUROPA"	3,288 SMT nodes with 2 Intel Nehalem-EP quad-core 2.93 GHz processors each 26,304 cores 77 TByte memory	308	Capacity and Capability Computing	European Universities, Research Institutes and Industry, PRACE
	Intel GPU Cluster "JUDGE"	206 nodes with 2 Intel Westmere 6-core 2.66 GHz processors each 412 graphic proces- sors (NVIDIA Fermi) 20.0 TByte memory	240	Capacity and Capability Computing	selected HGF Projects
	IBM Cell System "QPACE"	1,024 PowerXCell 8i processors 4 TByte memory	100	Capability Computing	QCD Applications SFB TR55, PRACE



JSC's supercomputer "JUQUEEN", an IBM Blue Gene/Q system.

Centres

# JSC Guest Student Programme 2014 -Experience Scientific Computing

As one of Europe's leading HPC centres, Jülich Supercomputing Centre (JSC) provides HPC expertise for computational scientists at German and European universities, research institutions, and in industry. Training activities and educational programmes for scientific computing are hosted by JSC on a regular basis. One of these activities is the Guest Student Programme (GSP) lasting for ten weeks each summer.

The participants receive extensive training on cutting edge hardware as well as HPC-related software and algorithms. The acquired theoretical knowledge is turned into hands-on skills by coached work on current and challenging scientific projects. For many students the programme has been the foundation for a career in HPC and the basis of fruitful long-term collaborations with their advisors. Some students even return to JSC as PhD candidates focusing on highly parallel applications.

Since the start of the GSP in 2000, a total of 157 students had the opportunity to join scientists from JSC and other institutes at Forschungszentrum Jülich. Over the course of 15 years the GSP improved continuously. This year an online application procedure was introduced. This boosted the number of applications from around 50 in the previous year to about 100. Candidates from 28 countries, covering students from mathematics, physics, chemistry, biology, and computer science, competed for open GSP positions. There were eleven students invited to participate in the programme.

This year's GSP took place from August 4th to October 10th. It was supported by CECAM (Centre Europeen de Calcul Atomique at Moleculaire) and sponsored within the IBM University programme.

It the first two weeks, courses on parallel programming up to advanced level, were run. The lectured techniques range from GPGPU programming with CUDA to the usage of MPI on distributed-memory clusters and OpenMP on shared-memory systems. Equipped with this vital knowledge the participants were ready to focus on the scientific part of the GSP. The range of scientific projects was as diverse as the user community on the hosted supercomputers, covering atmospheric science, fluid and molecular dynamics, multipole methods and safety research. Also represented was fundamental research in elementary particle physics and mathematical algorithms. In addition, this year there were also two projects supervised by the lately established Simulation Laboratory Neuroscience. This aimed at encouraging the use of supercomputers in neuroscience.

The main platforms for code development and simulation were the multipurpose cluster JUROPA, the GPU system JUDGE and the leadership Blue Gene/Q system JUQUEEN.

During the concluding two-day colloquium, the participants presented their achievements to domain experts and guests. The gained experiences were shared amongst the students, contributing to prolific discussions. Finally, as preparation for a future scientific career, the students summarized their contribution as an article.

Next year's GSP will start on August 3rd, 2015. It will be officially announced in January 2015 and is open to students from natural sciences, engineering, computer science, mathematics and the computer science related branches of neuroscience. For applicants it is mandatory, to have received the Bachelor but not yet the Master degree. The application deadline is April 24th, 2015. Additional information and the proceedings of the previous years are available online at www.fz-juelich.de/jsc/gsp.

contact: Ivo Kabadshow, i.kabadshow@fz-juelich.de Ivo Kabadshow Sven Strohmer

Jülich Supercomputing Centre (JSC), Germany

87

# Lattice Practices 2014

The 5th training workshop "Lattice Practices" was held at DESY Zeuthen March 5th to 7th this year. The scope of the Lattice Practices workshops is to provide training in state-of-the-art numerical techniques and the use of information technologies for research in lattice QCD (LQCD). Geared towards young researches, PhD students, and other interested LQCD practitioners, the workshops feature lectures on technical topics accompanied by hands on exercises with strong emphasis on practical training. Furthermore, a few very recent scientific developments are covered in order to expose the young researchers and students to potential areas of future research.

This year's workshop was organized by the Joint SimLab "Nuclear and Particle Physics" of Cyprus Institute, DESY, and JSC. Speakers from the SimLab partners and other European institutions gave technical lectures and hands-on tutorials on topics commonly dealt with in their field of research. The

topics ranged from data analysis and numerical techniques over optimization strategies and computer architecture to Higgs physics on the lattice. In the accompanying hands on sessions the participants were introduced to octave, given examples on basic techniques such as binning and error and autocorrelation analysis, but also given typical physics tasks, such as scale setting using realistic data sets. A particular emphasis was put on optimal programming, when the course of lectures and exercises went down to the silicon, to introduce the attendees to code optimization techniques and HPC architectures in general. This was completed by an introduction to numerical linear solver techniques and deepened in the accompanying exercises for both topics. The two recent developments covered in the course of lectures were: "Lattice aspects of Higgs physics" and disconnected diagrams.

A total of 29 young researches and students participated in this year's workshop, coming from institutions all over Europe, from Italy to Finland, but also from as far as the United States of America and Singapore. The opinions expressed through evaluation forms showed that the workshop was a great success and provided encouraging feedback for the next workshop, which is planned to take place during autumn of 2015. The slides of the talks and material of the hands on sessions can be found at:

https://indico.desy.de/event/LAP2014

contact: Stefan Krieg, s.krieg@fz-juelich.de



Figure 1: Participants of Lattice Practices 2014.

- Stefan Krieg<sup>1</sup>
- Dirk Pleiter<sup>1</sup>
- Rainer Sommer<sup>2</sup>
- Karl Jansen<sup>2</sup>
- Hubert Simma<sup>2</sup>
- Stefan Schäfer<sup>2</sup>
- Constantia
  Alexandrou<sup>3</sup>
- Giannis Koutsou<sup>3</sup>

# Jülich

Supercomputing Centre (JSC), Germany

# John von Neumann Institute for Computing (NIC), DESY Zeuthen, Germany

# Computationbased Science and Technology Research Center (CaSToRC), Cyprus

# **CECAM** Tutorial: Atomistic Monte Carlo Simulations of **Bio-Molecular Systems**

The CECAM tutorial "Atomistic Monte Carlo Simulations of Bio-molecular

Systems" took place at Forschungszentrum Jülich from September 15 to 19, 2014 and was attended by scientists from seven countries. The five days of the tutorial featured a range of lessons and hands-on practical sessions to provide scientists with everything necessary to apply this technique to their own research topics.

After the initial presentation by Prof. Anders Irbäck (Lund University) on the first afternoon that summarized the theory of Monte Carlo (MC) simulation and its application to biological macromolecules, the participants were introduced to the open source Monte Carlo simulation package ProFASi that served as the basis for the hands-on parts of the tutorial. ProFaSi is under active development by the organizers from the Simulation Laboratory Biology at JSC. It is a powerful alternative to molecular dynamics (MD), in particular for cases where the underlying process is too slow to be simulated by classical MD, such as in protein folding and peptide aggregation.

Two introductory sessions enabled the participants to set up, monitor, and analyze MC simulations of protein folding, peptide aggregation with ProFASi on the HPC resources of JSC. The following sessions addressed several advanced features including advanced error analysis and visualization. Wouter Boomsma (Univ. Copenhagen) demon-

strated the use of different constraints obtained from predictions and experiments in connection with MC simulations. Finally, the programming interface of ProFASi was introduced that allows for rapid development of new algorithms and simulation strategies using atomistic Monte Carlo.

The CECAM tutorial concluded with some recent research highlights using atomistic MC simulations, and a lively discussion of best practices and future developments with the participants.

contact: Sandipan Mohanty, s.mohanty@fz-juelich.de

# **UNICORE** Summit 2014

The UNICORE Summit is a unique opportunity for users, developers, administrators, researchers, service providers, and managers to meet. Its objective is to exchange and share experiences, new ideas, and latest research results on all aspects of UNICORE [1]. Since the first Summit in 2005, the organisers have received and reviewed a significant amount of distinguished contributions. Those selected and presented, complemented by invited talks, guarantee exciting Summits and lively discussions about the state-of-the art and the future of UNICORE, Grids, and distributed computing in general. The tenth edition, the UNICORE Summit 2014 [2] has been held on 24 June 2014 in Leipzig, Germany.

The invited talk "HPC Applications in Biophysics, Material Science and Biomedicine - enabled by UNICORE" by Borries Demeler, PhD, Associate Professor from University of Texas

Health Science Center at San Antonio, focused on applications of the Ultra-Scan XSEDE Science gateway [3] for high-resolution modelling of hydrodynamic experiments. The UltraScan software is used by scientists across the globe for research in biophysics, biochemistry, biomedicine, and material science to study the structure and function of biological macromolecules, investigate nanomaterials, and develop cures for diseases [4]. The keynote provided an overview of the integration of UNICORE into the gateway architecture in order to facilitate job submission and workflow management and discussed examples of science and discovery enabled by this implementation.

A second interesting use case was presented in the talk "A Workflow for Polarized Light Imaging Using UNICORE Workflow Services". The Polarized Light Imaging of brain slices is used to understand the anatomical structure of the human brain on the level of single nerve



Figure 1: Participants of UNICORE Summit 2014.

Olav Zimmermann

• Sandipan Mohanty

Supercomputing

Centre (JSC),

Germany

Jan Meinke

Jülich

fibres and is nowadays one of the most challenging tasks in neuroscience. The application of the UNICORE workflow system for this particular use case resulted in minimizing user interaction and time to completion of the scientific workflow. The next presentations highlighted current state, new ideas and concepts for the future development of the UNICORE portal [5], experiences with certificate-free user-friendly HPC access based on LDAP with UNICORE and UNITY [6], perspectives for REST services in the UNICORE environment, integration of UNICORE services in a private cloud computing platform and resource scheduling algorithms in distributed problem-oriented environments. Finally, the UNICORE roadmap and future developments were discussed by the attendees from Germany, Poland, Russia, and the United States.

The slides to the presentations can be found on the web at http://www. unicore.eu/summit/2014/schedule. php

# Valentina HuberDaniel Mallmann

Jülich Supercomputing Centre (JSC), Germany

# References

[1] UNICORE Web Page: http://www.unicore.eu

- [2] UNICORE Summit 2014 web page: http:// www.unicore.eu/summit/2014/
- [3] XSEDE Science Gateway: https://portal. xsede.org/science-gateways
- [4] UltraScan-III: http://www.ultrascan.uthscsa.edu
- [5] UNICORE Portal: http://sourceforge.net/ projects/unicore/files/Servers/Portal
- [6] UNITY Web Page: http://www.unity-idm.eu

contact: Daniel Mallmann, d.mallmann@fz-juelich.de

# "Bernstein Network -Simulation Lab Neuroscience" HPC Workshop

Neuroscience today is attacking problems of increasing complexity and scale as exemplified by projects like the Human Brain Project, which require computationally intensive simulations and the analysis of large data sets. However, many projects currently using local clusters for these purposes have not yet adapted their software and theoretical approaches to take advantage of HPC systems such as those available at the Jülich Supercomputing Centre (JSC).

The "Bernstein Network – Simulation Lab Neuroscience" HPC Workshop on June 4th and 5th at the JSC brought together Jülich computational neuroscientists and HPC experts with neuroscience domain experts from across Germany who are interested in developing petascale simulations and analyses. An important goal of this meeting was to find ways for the neuroscience community to fully exploit available JSC resources by catalyzing collaborations and adapting tools to supercomputer scales.

A total of 32 participants shared their perspectives on HPC in neuroscience. Members of the SimLab Neuroscience [1] and the JSC's HPC in Neuroscience Division delivered presentations covering a range of issues regarding the use of computing facilities at the JSC, in addition to describing work being currently done by the SimLab which leverages these resources such as structural plasticity modeling in the visual cortex using the NEST simulator [2].

> Figure 1: Participa HPC Workshop.





Figure 1: Participants of the "Bernstein Network - Simulation Lab Neuroscience"

# Other experts from the JSC and Jülich's Institute of Neuroscience and Medicine (INM) explained the computetime grant-writing process, as well as showing a variety of projects that already leverage JSC resources including large-scale neuronal network simulations on the JUQUEEN supercomputer and "Big Data" approaches to experimental electrophysiological analyses.

Fifteen external neuroscientists from the Bernstein Network [3] presented projects which they hoped to bring to the JSC supercomputers, ranging from macroscopic models of whole brain functions through neuronal network self-organization and down to ion flows in dendritic spines. Discussions regarding how to directly port these projects as well as how to further extend them so as to maximize parallelization for supercomputing architectures should lead to a new generation of neuroscience projects at the JSC.

Further details on the program are available at:

http://www.fz-juelich.de/SharedDocs/

Termine/IAS/JSC/EN/events/2014/

bernstein-hpc-2014.html

# Anne Do Lam-Ruschewski<sup>1</sup>

- Anna Lührs<sup>1</sup>
- Abigail Morrison<sup>2</sup>
- Boris Orth<sup>1</sup>
- Alexander Peyser<sup>1</sup>
- Wolfram Schenck<sup>1</sup>
- <sup>1</sup> Jülich Supercomputing Centre (JSC), Germany

# <sup>2</sup> INM-6, Forschungszentrum Jülich, Germany

# References

- [1] Simulation Laboratory Neuroscience Bernstein Facility for Simulation and Database Technology. URL http://www. fz-juelich.de/ ias/jsc/EN/Expertise/Sim Lab/slns/\_node.html
- [2] Gewaltig, M.-O., Diesmann, M.
  - NEST (NEural Simulation Tool). Scholarpedia, 2(4):1430, 2007. URL http://nestinitiative.org/Software:About\_NEST
- [3] Bernstein Network Computational Neuroscience. URL http://www.nncn.de/en

# contact: Boris Orth, b.orth@fz-juelich.de

# Jülich School on Computational Trends in Solvation and Transport in Liquids

From 23 to 27 March 2015 the IAS School on "Computational Trends in Solvation and Transport in Liquids" will take place at Jülich Supercomputing Centre. This event is part of a series of Schools in Computational Science, which are regularly organized at Jülich since more than 15 years [1]. More than 20 renowned scientists from seven countries will present lectures on modern methods and algorithms for treating solvents most efficiently on different length and time scales. The School is part of the activities of the Jülich CECAM node [2] and is cosupported by the Cluster of Excellence RESOLV [3].

"Solvation Science" is increasingly recognized as an interdisciplinary field akin to "Materials Science" or "Neuroscience" addressing a variety of different computational and simulation methods, appropriate for hierarchies of time- and length-scales, which provide a challenge to be solved. Since solvation and transport problems are apparent in a broad field, ranging from fundamental questions in chemistry or soft matter physics to industrial applications, the field is highly interdisciplinary. This calls for meetings bringing experts together from various directions and triggering exchange of ideas between disciplines. The IAS School 2015 will focus on the computational trends, multi-method approaches and modeling in this interdisciplinary field. It covers the field from

large-scale coarse grain modeling down to fully quantum-mechanical simulations of liquids at the level of electrons and nuclei. Not only bulk liquids and homogeneous solutions will be discussed, but also heterogeneous systems such as liquid/solid interfaces as well as solvated (bio)molecules. In particular, recent advances in adaptive resolution methods both in the realm of finite element modeling and of interfacing atomistic and coarse-grain descriptions of liquids will be covered. Moreover, a variety of hybrid methods, such as QM/MM approaches for solvated biomolecules, continuum solvation and lattice Boltzmann techniques, will be part of the program. Coarse graining in many distinct flavors certainly is an important approach to describe transport in liquids and solvation of increasing complex systems. This includes methods such as Brownian dynamics for biomolecular recognition, mesoscale particle methods for hydrodynamics to model particle-fluid interactions or neural network potentials to describe reactive water.

The recent revival of liquid-state integral equation methods in combination with molecular dynamics and electronic structure theory is acknowledged in the program as well as new developments in molecular density functional theory of aqueous solutions. Wellestablished techniques such as force field molecular dynamics for large-scale Activities

simulations and ab initio molecular dynamics for wet chemical reactions will be addressed as well to introduce the audience to the field. Within molecular dynamics special attention will be given to methods that are tailored to unravel solvation effects. Presentation of most recent developments in path integral simulation techniques will complement the lectures to cover the smallest scales on nuclear level, which are relevant in atomistic description of solvation.

The solution of large scale complex problems needs a direct link to high performance computing, which today includes the use of GPUs in addition to massively parallel CPU based systems. Recent developments and trends will be addressed in the School not only by providing lectures but also by including a hands-on practical tutorial on elementary GPU programming.

This IAS School is suited for highly motivated PhD students and PostDocs. Applications for participation can be sent until end of January 2015. Based on the required application documents about 50 participants will be selected by the organizers. Details about the School and the application process can be found at: http://www.fz-juelich.de/STL-2015 .

# References

- [1] For a list of past Schools, see: www.fzjuelich.de/ias/jsc/EN/Expertise/Work shops/Conferences/STL-2015/PastWork shops/\_node.html
- [2] www.cecam.org
- [3] CoE Ruhr Explores Solvation: www.ruhr-unibochum.de/solvation/

contact: Godehard Sutmann, g.sutmann@fz-juelich.de

# **3rd Workshop on Parallel-in-**Time Integration Held at JSC

At the doorstep of the Exascale era, an urging demand for improved and new numerical algorithms arises. For time-dependent problems, the idea of concurrency in the time domain attracts more and more interest in many different communities. In order to overcome the serial dependence in the time direction and to enable integration of multiple time-steps simultaneously, time-parallel methods commonly introduce a space-time hierarchy, where integrators with different accuracies and costs are coupled in an iterative fashion. Serial dependencies are shifted to the coarsest level, allowing the computationally expensive parts on finer levels to be treated in parallel. Typical examples of this concept are Parareal and the "parallel full approximation scheme in space and time" (PFASST). The space-time hierarchy used in these approaches shows strong similarities to classical multigrid structures. For example, Parareal can be interpreted as two-grid algorithm in time. PFASST uses iterative spectral deferred corrections as smoother in time and employs a full approximation scheme, thus making it conceptually similar to spatial nonlinear multigrid methods.

From May 26 to 28, 2014, the 3rd Workshop on Parallel-in-Time Integration with special focus on parallel multilevel methods in space and time was held at Jülich Supercomputing Centre. It was jointly organized by Robert Speck (Forschungszentrum Jülich), Matthias Bolten (University of Wuppertal), Rolf Krause, and Daniel Ruprecht (both USI

This workshop was the third one in a series of workshops for a fast-growing community, following the events at Università della Svizzera italiana in 2011 and at the University of Manchester in 2013. In May 2015, the 4th workshop will be held at TU Dresden and further events are already envisaged for the following years.

contact: Robert Speck, r.speck@fz-juelich.de



- Godehard Sutmann<sup>1</sup> Johannes
- Grotendorst<sup>1</sup>
- Dominik Marx<sup>2</sup>
- <sup>1</sup> Jülich Supercomputing Centre (JSC), Germany
- <sup>2</sup> Ruhr-Universität Bochum, Germany

Lugano), and was supported by DFG via SPPEXA, the German Priority Programme 1648 "Software for Exascale Computing". With 42 participants from academia, research and industry coming from eleven different countries a broad spectrum of expertise was brought together to form a great ambiance for a successful exchange of ideas. The topics ranged from applied mathematics to climate and earth science as well as engineering and software development. With sufficient time for discussions and individual meetings, new collaborations were initiated and long-lasting contacts renewed.

Robert Speck

Jülich Supercomputing Centre (JSC), Germany

97

Figure 1: Participants of the 3rd Workshop on Parallel-in-Time Integration.

# 13th HLRS/hww Workshop on Scalable Global Parallel **File System**

# SELF-MANAGING DATA

13th HLRS / hww Workshop on **Scalable Global Parallel File Systems** 

May 12-14, 2014 HLRS, Stuttgart, Germany www.hlrs.de

Representatives from science and industry working in the field of global parallel file systems and high performance storage solutions did meet at HLRS from May 12th to May 14th, 2014, for the 13th annual HLRS/hww Workshop on Scalable Global Parallel File Systems. About 75 participants did follow a total of 22 presentations that have been on the workshop agenda.

Dr. Norbert Conrad, Deputy Director HLRS, opened the workshop with an opening address on Monday morning.

In the keynote talk, Eric Barton of Intel's High Performance Data Division and former CTO of Whamcloud discussed technology developments in the Fast Forward I/O and storage program. He explained the planned I/O architecture and the different steps in the I/O workflow on large HPC systems. Sai Narasimhamurthy, Xyratex, explained the current status the vision and the roadmap of the Exascale10 I/O middle-

ware. This middleware, formerly known as EIOW will tackle the storage issues of the exascale era by providing guided I/O methods to pass more information about the I/O from the application or the middleware to the system level.

In the first presentation of the Monday afternoon session, Franz-Josef Pfreundt, FhG - ITWM, introduced BeeGFS and its connection to Big Data. BeeGFS is the new name for the well-known Fraunhofer File System FhGFS. In his talk, Sven Oehme, senior file system developer IBM, provided information about the functional enhancements in the new GPFS storage server release and he explained the influence on performance. To complete the file system session, James Coomer, DDN, illuminated Lustre Performance improvements with Large IO and solid state acceleration.

In the second Monday afternoon session, Andrew Grimshaw, University of Virginia introduced the XSEDE Global Federated File System GFFS. In his talk, he was explaining how GFFS is breaking down the barriers to secure resource sharing and he showed an online demonstration.



Thomas Schoenemeyer, Cray, provided insight to the new Cray Tiered Adaptive Storage (TAS) and showed its benefits. In the last talk of the day, Didier Gava, Netapp, discussed how flash arrays can help to solve HPC challenges.

The Tuesday storage technology talks covered technologies using fast flash memory. James Coomer, DDN, explained caching approaches for emerging, large-scale data problems and John Bent, EMC, discussed software approaches for exascale burst buffers, especially IOD, a non-POSIX interface to persistent data.

Ulrich Lechner, CTO GrauData, presented their scalable, hardware independent solutions for HSM, archiving and secure file-sharing and Johannes Reetz, RZG, concluded the data centric presentations of the day by an introduction to EUDAT the European Data Infrastructure.

The second half of the day became more network centric. The 400 GBit Testbed between TU Dresden and RZG in Garching has been shown in the talk of Eduard Beier, T-Systems while Software Defined Networking has been addressed by Torsten Omlor, IBM. Adva's Klaus Grobe went down to the hardware and explained solutions for inter-data-center 400-Gb/s WDM transport. Mondrian Nuessle, CTO and manager of Extoll introduced the brand new Tourmalet HPC Network ASIC.

On Wednesday morning, Thomas Uhl, Datera, provided insight into the companies elastic block storage. The following talks have been more research related. Andre Brinkmann, Mainz University showed the work and the results of different I/O Projects in Mainz.



The introduction of different high performance data transfer tools which have been recently developed by himself mainly at HLRS has been the topic of Frank Scheiner.

for OMPIO.

event possible.

contact: Thomas Bönisch, boenisch@hlrs.de



Activities

Michaela Zimmer from the SIOX project and University Hamburg gave an overview about the whole development the architecture and the current SIOX roadmap. Andriy Chut and Xuan Wang, both HLRS, focussed on the integration of GPFS with SIOX and a GPFS Interface

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

• Thomas Bönisch

University of Stuttgart (HLRS), Germany



# GCS @ ISC'14 in Leipzig



Figure 1: The GCS booth at ISC'14.

At ISC'14 in Leipzig, the 64 sqm large booth of the Gauss Centre for Supercomputing (GCS) was once again one of the most popular gathering points for the international HPC community. The open and inviting concept of the booth, targeted to encourage the ISC participants to stop by and interexchange with the GCS representatives, proved its merit. The GCS booth was always busy. Countless like-minded HPC users, researchers, technology leaders, scientists, IT-decision makers as well as high tech media representatives visited the GCS booth to meet and talk with the directors of the three GCS centres Prof. Arndt Bode (LRZ), Prof. Thomas Lippert (JSC), Prof. Michael M. Resch (HLRS), GCS managing director Dr. Claus Axel Müller, as well as with the present scientists of the three GCS centres.

# **TOP500**

The 43rd edition of the TOP500 delivered proof that GCS continues to play a leading role in HPC. JSC's JUQUEEN

took place 8 on the noted list and LRZ's SuperMUC occupied position 12. Hermit of HLRS, the longest serving GCS HPC system which debuted in 12/2011, still shines in the TOP500sublist for industrially used supercomputers where it holds a very strong 3rd place worldwide.

# ISC'14 Gauss Award Winner

Each year at ISC, GCS presents the Gauss Award to recognize the most outstanding paper in the field of scalable supercomputing from all papers accepted for the ISC'14 Research Paper Sessions. This year, the award honoured the paper "Exascale Radio Astronomy: Can We Ride the Technology Wave?", written by Erik Vermij, Leandro Fiorin, Christoph Hagleitner (all of IBM Research) and Koen Bertels of the Delft University of Technology.



Figure 2: The ISC'14 Gauss Award winners with Dr. Claus Axel Müller (left) and Prof. Michael M. Resch (right).

# Large Media Interest

The directors of the three GCS centres were of high demand by the international journalists. Several interviews were conducted in which the directors talked about GCS and HPC in Germany in general and during which they in particular emphasized the benefit the GCS supercomputer infrastructure delivers to science and research. They did so in pointing out several GCS large scale



Figure 3: Left to right: Professor Arndt Bode (LRZ), Professor Thomas Lippert (JSC) and Professor Michael M. Resch (HLRS) in talk with media representatives.

projects which set new world records in simulation runs and which helped obtain unprecedented details that allowed to further the until now limited knowledge about some of the most pressing scientific riddles and challenges of our time. Some examples are the SeisSol seismic simulation project by Prof. Bader (Technische Universität München), the Illustris astrophysics simulation project by Prof. Springel (Heidelberger Institut für Theoretische Studien), and the theoretical mechanochemistry project by Prof. Marx of the Ruhr-Universität Bochum, just to name a few.

# **Booth Highlights**

HLRS attracted lots of attention with their hands-on Augmented Reality demo analysing simulation results of the airflow and pressure distribution around a 3D scanned triathlete on his racing bicycle.



By moving a camera around the bicycle

or even ride the bicycle themselves, visitors could immediately observe changes to the airflow resulting from various riding positions of the triathlete. This method helps triathletes to find and verify the most efficient riding position on their racing bike and to analyse individually mounted accessories and helmets.

Apart from presenting scientific results obtained with the JSC HPC systems, **JSC** presented LLview, the in-house developed comprehensive interactive monitoring software for supercomputers, demonstrating live the operation of various supercomputers worldwide. In addition, JSC also showed the LLview monitoring components of the Eclipse PTP development environment for supercomputing applications.

LRZ focused on highlighting their HPC system SuperMUC and presented current science and research projects in 2D-videos. Additionally, LRZ ran a series of well accepted short presentations on the GCS booth on the GCS booth on subjects revolving around new and future HPC services offered by the Gaching based HPC centre.

contact: Regina Weigand, r.weigand@gauss-centre.eu Activities

Regina Weigand

Gauss Centre for Supercomputing

101

# First Intel MIC & GPU **Programming Workshop** at LRZ

In order to achieve high performance and best scaling results on heterogeneous accelerator-based systems, a three day Intel MIC & GPU programming workshop has been organized by the Leibniz Supercomputing Centre as a PRACE Advanced Training Centre (PATC) for the first time, dated 28 to 30 April 2014. The workshop attracted 25 participants, partly coming from Austria and Turkey to join the event. The goal of the workshop was to make the participants more familiar with basic GPGPU and especially Intel Xeon Phi programming techniques, and to give them the chance to work with GPGPU and – for the first time – also Intel Xeon Phi based systems. The workshop covered various highlevel programming models and optimisation techniques. While the first day focused more on GPGPU programming using CUDA, OpenACC, Python and R, the second day was devoted to Intel Xeon Phi programming using OpenMP, MPI, Offloading, Intel Cilk Plus, MKL and OpenCL. On the last day the invited

speakers Dr.-Ing. Michael Klemm from Intel Corp. and Dr.-Ing. Jan Treibig from the Regional Computing Centre Erlangen (RRZE) gave lectures about advanced MIC programming (using Intrinsics or assembly language), tuning methodologies and the new features in OpenMP 4.0.

During many hands-on sessions the participants were able to gain experience on the GPGPU cluster at LRZ and - as the very first users - also on the new Intel Xeon Phi based cluster Super-MIC at LRZ (see also the article about SuperMIC in this issue of inSiDE). In addition, the participants also had the opportunity to apply their new skills to their own codes.

Based on the very positive feedback during the workshop a similar event has recently been scheduled at LRZ for spring 2015.

Autumn 2014 • Vol. 12 No.2 • inSiDE

contact: Volker Weinberg, Volker.Weinberg@lrz.de



During its first two years of operation, SuperMUC has produced an enormous amount of results. More than 100 projects wrote reports for the proceedings of the SuperMUC Status and Results Workshop, which are now published (see ISBN and link below). From the reports, 28 interesting talks were invited to report at the workshop in Garching, from July 8-9, 2014. Users and project managers of SuperMUC projects came to discuss their scientific projects and their experiences using SuperMUC. All users participated in lively discussions, especially during the user forum, where they could bring in future requirements. Throughout the workshop, users could directly discuss their wishes with system administrators and application experts from LRZ, IBM and Intel. The LRZ also informed their users about its plans for Phase 2, which doubles the performance of SuperMUC, and introduced the new Intel Xeon Phi island (SuperMIC), as well as the new remote visualization cluster.

contact: Helmut Satzger, helmut.satzger@lrz.de



The SuperMUC proceedings can be downloaded from this URL: http://www.lrz.de/services/compute/ supermuc/magazinesbooks. The book is available as PDF and for e-book readers (epub and mobi).

ISBN 978-3-9816675-0-9

The talks from the workshop are available as PDF from: http://www.lrz.de/services/compute/ supermuc/magazinesbooks



Figure 1: Participants of the first Intel Xeon Phi & GPU programming workshop at LRZ. In the centre the main lecturers Dr.-Ing. J. Treibig (RRZE), Dr.-Ing. M. Klemm (Intel) and Dr. V. Weinberg (LRZ) holding an Intel Xeon Phi coprocessor.

Volker Weinberg

Momme Allalen

Supercomputing Centre (LRZ),

Leibniz

Garching,

Germany

Activities

5. Wagner • A. Bode • H. Satzger • M. Brehm EDITORS

# **High Performance** Computing

in Science and Engineering Garching/Munich 2014



Bayerische ademie der Wissenschaften

GCS Irz

# • Helmut Satzger

Leibniz Supercomputing Centre (LRZ), Garching, Germany

103

# Activities

# **HLRS Scientific Tutorials and** Workshop Report and Outlook

In late August 2014 we entered the next step of our HPC systems installation phase: the new Cray XC4O supercomputer is installed. It will deliver a peak performance of 3.786 Petaflops, outperforming the maximum performance of the previous system, Hermit, by a factor of about 4!

This new HPC system provides 500 TB of Main Memory and about 6 PB of disc space. It is equipped with about 100.000 computing cores and features Intel's next generation of micro processors, which are designed to optimize power savings and promise significant performance enhancements.

Rolf Rabenseifner

University of Stuttgart, HLRS, Germany

We strongly encourage you to port your applications to this architecture as early as possible. To support such efforts we invite all potential users to participate in our newly arranged series of three courses: Cray XC 40, Parallel I/O, and Optimization.

These courses will provide all necessary information to move applications to the new HPC system. We are looking forward to working with our users on this leading-edge supercomputing



technology. The next course series in cooperation with Cray specialists will take place on March 2-5, 2015. An online recording from September 2014 is also available.

Programming of Cray XK7 clusters with GPUs is taught in OpenACC Programming for Parallel Accelerated Supercomputers - an alternative to CUDA from Cray perspective on April 16-17, 2015. These Cray XC40 and XK7 courses are also presented to the European community in the framework of the PRACE Advanced Training Centre (PATC). GCS, i.e., HLRS, LRZ and the Jülich Supercomputer Centre together, serve as one of the first six PATCs in Europe. CUDA courses are also presented in April and in October 2015.

One of the flagships of our courses is the week on Iterative Solvers and Parallelization. Prof. A. Meister teaches basics and details on Krylov Subspace Methods. Lecturers from HLRS give lessons on distributed memory parallelization with the Message Passing Interface (MPI) and shared memory multithreading with OpenMP. This course will be presented twice, on March 16-20, 2015 at HLRS in Stuttgart and on October 05-09, 2015 at LRZ in Garching near Munich.

# Another highlight is the Introduction to Computational Fluid Dynamics.

This course was initiated at HLRS by Dr.-Ing. Sabine Roller. She is now a professor at the University of Siegen. It is again scheduled on February 23-27, 2015 in Siegen and in autumn 2015 in Stuttgart.

In April 2015, Performance and **Debugging Tools** are presented to assist parallel programming. In July 2015 we continue the successful series of two courses on software optimization, the Node-level Performance Engineering by Georg Hager and Jan Treibig, and User-guided **Optimization in High-Level Languages** from the Computer Graphics Lab at Saarland University. In 2015, the workshop on Scalable Global Parallel File Systems will be extended by a one day I/O Tutorial, and in June 2015, a Cluster Workshop will discuss hardware and software aspects of compute clusters.

The Visualization Courses in April and October 2015 are targeted at researchers who would like to learn how to visualize their simulation results on the desktop but also in Augmented Reality and Virtual Environments.

Our general course on parallelization, the Parallel Programming Work**shop**, September 07–11, 2015 at HLRS, will have three parts: The first two days of this course are dedicated to parallelization with the Message Passing Interface (MPI). Shared memory multi-threading is taught on the third day, and in the last two days, advanced topics are discussed. This includes MPI-2 functionality, e.g., parallel file I/O and hybrid MPI+OpenMP, as well as the upcoming MPI-3.0. As in all courses, hands-on sessions (in C and Fortran) will allow users to immediately test and understand the parallelization methods. The course language is English.

In the table below, you can find the whole HLRS series of training courses in 2015. They are organized at HLRS and also at several other HPC institutions: LRZ Garching, NIC/ZAM (FZ Jülich), ZIH (TU Dresden), and ZIMT (Siegen).

Several three and four day-courses on MPI & OpenMP will be presented at different locations all over the year.

2015 - Workshop Announcements Scientific Conferences and Workshops at HLRS 14th HLRS/hww Workshop on Scalable Global Parallel File Systems with I/O Tutorial (not yet fixed) 9th ZIH+HLRS Parallel Tools Workshop (date and location not yet fixed) High Performance Computing in Science and Engineering – The 17th Results and Review Workshop of the HPC Center Stuttgart (Autumn 2015) IDC International HPC User Forum (Autumn 2015) Parallel Programming Workshops: Training in Parallel Programming and CFD Parallel Programming and Parallel Tools (TU Dresden, ZIH, February 16-20) Industrial Services of HLRS (HLRS, February 25, July 01, and November 11) Introduction to Computational Fluid Dynamics (ZIMT, Uni Siegen, February 23-27) Cray XC4O and I/O Workshops (HLRS, March O2–O5 and October 20–23) (PATC) Iterative Linear Solvers and Parallelization (HLRS, March 16-20 / LRZ Garching, October 05-09) Tools for Parallel Programming (HLRS, April 13–15) Open ACC Programming for Parallel Accelerated Supercomputers (HLRS, April 16-17) (PATC) GPU Programming using CUDA (HLRS, April 20–22 and October 26–28) Unified Parallel C (UPC) and Co-Array Fortran (CAF) (HLRS, April 23-24) (PATC) Efficient Parallel Programming with GASPI (HLRS, April 27) Scientific Visualization (HLRS, April 28-29 and October 29-30) Cluster Workshop (HLRS, June 29-30) Node Level Performance Engineering (HLRS, July 06–07) Summer School: Modern Computational Science in Quantum Chemistry Uni Oldenburg, date not yet fixed) Introduction to Computational Fluid Dynamics (HLRS, September/October) Message Passing Interface (MPI) for Beginners (HLRS, September 07–08) (PATC) Shared Memory Parallelization with OpenMP (HLRS, September O9) (PATC) Advanced Topics in Parallel Programming (HLRS, September 10-11) Parallel Programming with MPI & OpenMP (FZ Jülich, JSC, November 30-December 02) Training in Programming Languages at HLRS Fortran for Scientific Computing [LRZ Garching, Februar 09–13 / HLRS, December 07–11, 2015] (PATC)

URLs:

Activities

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

http://www.hlrs.de/training/course-list/

(PATC): This is a PRACE PATC course

We also continue our series of Fortran for Scientific Computing on December 08–12,2014 and on March 09–13, 2015, mainly visited by PhD students from Stuttgart and other universities to learn not only the basics of programming, but also to get an insight on the principles of developing high-performance applications with Fortran.

### **ISC and SC Tutorials**

Georg Hager, Gabriele Jost, Rolf Rabenseifner: Hybrid Parallel Programming with MPI & OpenMP. Tutorial O4 at the International Supercomputing Conference, ISC'14, Leipzig, June 22-26, 2014.

Georg Hager, Jan Treibig, Gerhard Wellein: Node-Level Performance Engineering. Tutorial O1 at the International Supercomputing Conference, ISC'14, Leipzig, June 22-26, 2014.

Rolf Rabenseifner, Georg Hager: MPI+X - Hybrid Programming on Modern Compute Clusters with Multicore Processors and Accelerators. Half-day Tutorial at Supe Computing 2014, SC14, New Orleans, Louisiana, USA, November 16–21, 2014.



# GCS - High Performance Computing **Courses and Tutorials**

# Introduction

to the Programming and Usage of the Supercomputer Resources at Jülich

### **Date & Location**

November 27-28, 2014 JSC, Forschungszentrum Jülich

### Contents

This course gives an overview of the supercomputers at Jülich. Especially new users will learn how to program and use these systems efficiently. Topics discussed are: system architecture, usage model, compilers, tools, monitoring, MPI, OpenMP, performance optimization, mathematical software, and application software.

# Webpage

http://www.fz-juelich.de/ias/jsc/ events/sc-nov

# **Parallel Programming** with MPI and OpenMP

# **Date & Location**

December 1-3, 2014 JSC, Forschungszentrum Jülich

# Contents

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

## Webpage

http://www.fz-juelich.de/ias/jsc/ events/mpi

With Unified Parallel C (UPC) and

24, 2015, and an Introduction to

Co-Array Fortran (CAF) on April 23-

GASPI on April 27, 2015 the partici-

pants will get an introduction of par-

titioned global address space (PGAS)

In cooperation with Dr. Georg Hager

from the RRZE in Erlangen and Dr.

Gabriele Jost from Supersmith, the

HLRS also continues with contribu-

tions on hybrid MPI & OpenMP pro-

gramming with tutorials at confer-

ences; see the box on the left, which

includes also a second tutorial with

Georg Hager from RRZE.

languages.

# **Node-Level Performance** Engineering (PATC course)

**Date & Location** December 4-5, 2014 LRZ, Garching near Munich

## Contents

This course teaches performance engineering approaches on the compute node level. "Performance engineering" as we define it is more than employing tools to identify hotspots and bottlenecks. It is about developing a thorough understanding of the interactions between software and hardware. This process must start at the core, socket, and node level, where the code gets executed that does the actual computational work. Once the architectural requirements of a code are understood and correlated with performance measurements, the potential benefit of optimizations can often be predicted. We introduce a "holistic" node-level performance engineering strategy and apply it to different algorithms from computational science.

### Agenda:

- Introduction and Motivation
- Performance Engineering as a process
- Topology and affinity in muticore systems
- Microbenchmarking for architectural exploration
- The Roofline Model
- Basics and simple applications

# **Courses and Tutorials**

# **Programming with Fortran**

## **Date & Location**

February 9–13, 2015 LRZ, Garching near Munich

### Contents

This course is targeted at scientists with little or no knowledge of the Fortran programming language, but need it for participation in projects using a Fortran code base, for development of their own codes, and for getting acquainted with additional tools like debugger and syntax checker as well as handling of compilers and libraries. The language is for the most part treated at the level of the Fortran 95 standard; features from Fortran 2003 are limited to improvements on the elementary level. Advanced Fortran features like object-oriented programming or coarrays will be covered in a follow-on course in autumn. To consolidate the lecture material, each day's approximately 4 hours of lecture are complemented by 3 hours of hands-on sessions. Course participants should have basic

UNIX/Linux knowledge (login with secure shell, shell commands, basic programming, vi or emacs editors).

### Webpage

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

& Parallel Tools

### **Date & Location**

February 16-20, 2015 Dresden, ZIH

### 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. The last day is dedicated to tools for debugging and performance analysis of parallel applications. This course is organized by ZIH in collaboration with HLRS.

# Webpage

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

# Fortran for Scientific Computing

**Date & Location** December 8–12, 2014 (in German) HLRS, Stuttgart

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

# Contents

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

# This course is dedicated for scientists

### Webpage

- Hardware performance metrics - Typical performance patterns in

- Case study: sparse matrix-vector

- Blocking optimization for the Jacobi

Programming for optimal use of

- Single Instruction Multiple Data

- Cache-coherent Non-Uniform

Pattern-guided performance

Memory Architecture (ccNUMA)

- Simultaneous Multi-Threading (SMT)

- Case study: Jacobi smoother

• Model-guided optimization

multiplication

smoother

(SIMD)

engineering

parallel resources

- scientific computing
- Examples and best practices
- Beyond Roofline: The ECM Model
- (Optional: Energy-efficient code execution)

Between each module, there is time for Questions and Answers!

Prerequisites: Participants must have basic knowledge in programming with Fortran or C.

Third JUQUEEN Porting and

**Tuning Workshop** 

(PATC course)

**Date & Location** 

Contents

Webpage

events/juqueenpt15

February 2-4, 2015

JSC, Forschungszentrum Jülich

The new Blue Gene/Q petaflop

supercomputer JUQUEEN marks an-

other quantum leap in supercomputer

performance at JSC. In order to use

this tool efficiently, special efforts by

the users are necessary, though. The

aim of this hands-on workshop is to

support current users of JUQUEEN in

porting their software, in analyzing its

performance, and in improving its ef-

ficiency. This course is a PATC course

(PRACE Advanced Training Centres).

http://www.fz-juelich.de/ias/jsc/

# **Parallel Programming**

# **Computational Fluid Dynamics**

## **Date & Location**

February 23–27, 2015 Siegen University, ZIMT

# Contents

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

## Webpage

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

# Industrial Services of HLRS (PATC course)

**Dates & Location** February 25, and July 1, 2015 HLRS, Stuttgart

## Contents

In order to permanently assure their competitiveness, enterprises and institutions are increasingly forced to deliver highest performance. Powerful computers, among the best in the world, can reliably support them in doing so.

This information half-day is targeted towards decision makers in companies that would like to learn more about the advantages of using high performance computers in their field of business. They will be given extensive information about the properties and the capabilities of the computers as well as access methods and security aspects. In addition we present our comprehensive service offering - ranging from individual consulting via training courses to visualization. Real world examples will finally allow an interesting insight into our current activities.

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

# CRAY XC40, Parallel I/O, and Optimization courses (PATC course)

**Date & Location** Series of three Courses: March 2, 2015 March 3, 2015 March 4-5, 2015 HLRS, Stuttgart

### Contents

During August 2014 HLRS and Cray has installed "Hornet" the follow-up to Hermit. Hornet is a Cray XC40 with 3.944 compute nodes. It is based on a new Intel "Haswell" processor and the Cray Aries network. The system will have about 100.000 cores, a peak performance of 3.786 PFLOP/s and a main memory of 493 TByte. The first day of this course is targeted to Hermit users who will continue their work on Hornet. Specialists from Cray will talk about the hardware and software enhancements between the XE6 and XC4O in order to support your migration towards this new machine. The second day is dedicated

to user parallel IO at scale. The third and fourth day cover a 2 day introduction workshop about porting and optimizing your application for the Cray XC4O. The topics of this workshop include the Cray Programming Environment, scientific libraries and profiling tools.

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

# Fortran for Scientific Computing (PATC course)

**Date & Location** March 9–13, 2015 HLRS, Stuttgart

## Contents

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

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

# **Courses and Tutorials**

# **Parallel Programming of High Performance Systems**

## **Date & Location**

March 9–13, 2015 RRZE building, University campus Erlangen, via video conference at LRZ if there is sufficient interest.

### Contents

This course, a collaboration of Erlangen Regional Computing Centre (RRZE) and LRZ, is targeted at students and scientists with interest in programming modern HPC hardware, specifically the large scale parallel computing systems available in Jülich, Stuttgart and Munich.

Each day is comprised of approximately 4 hours of lectures and 3 hours of hands-on sessions.

# Day 1

- Introduction to High Performance Computing (Weinberg)
- Secure shell (Weinberg)
- Source code versioning with SVN (Guillen)
- Handling console and GUI-based interfaces (Weinberg)
- Building programs with GNU MAKE (Guillen)

### Day 2

- Basic parallel programming models: elements of MPI and OpenMP
- (Weinberg)
- Processor architectures (Hager) Day 3
- Principles of code optimization: unrolling, blocking, dependencies, C++ issues, bandwidth issues,

# performance projections (Hager)

- Basics of software engineering (Guillen)
- Advanced MPI programming (Wittmann)

### Day 4

- Advanced OpenMP programming (Weinberg)
- Performance Libraries (Weinberg)
- Parallel architectures: multi-core, multi-socket, ccNUMA, cache coherence and affinity, tools for handling memory affinity (Hager)
- Parallel algorithms: data parallelism, domain decomposition, task parallelism, master-worker, granularity, load balancing, scalability models (Hager)

# Day 5

- Intel tools: MPI tracing and Checking (Karmakar)
- Intel tools: OpenMP performance and correctness (Wittmann)
- •Performance analysis with Scalasca (Karmakar)

# Webpage

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

Autumn 2014 • Vol. 12 No. 2 • inSiDE

# Parallel I/O and Portable **Data Formats**

(PATC course)

# **Date & Location**

March 16-18, 2015 JSC, Forschungszentrum Jülich

# Contents

This course will introduce MPI parallel I/O and portable, self-describing data formats, such as HDF5 and NetCDF. Participants should have experience in parallel programming in general, and either C/C++ or Fortran in particular. This course is a PATC course (PRACE Advanced Training Centres).

# Webpage

http://www.fz-juelich.de/ias/jsc/ events/parallelio

# **Courses and Tutorials**

# **Iterative Linear Solvers and** Parallelization

# **Date & Location** March 16-20, 2015 HLRS, Stuttgart

## Contents

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

## Webpage

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

# Eclipse: C/C++/Fortran Programming

**Date & Location** March 19, 2015 LRZ, Garching near Munich

# Contents

This course is targeted at scientists who wish to be introduced to programming C/C++/Fortran with the Eclipse C/C++ Development Tools (CDT), or the Photran Plugin. Topics covered include:

- Introduction to Eclipse IDE
- Introduction to CDT
- Hands-on with CDT
- Short introduction and demo of Photran

## Webpage

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

Vectorisation and Portable Programming Using OpenCL

**Date & Location** March 19-20, 2015, 2 days JSC, Forschungszentrum Jülich

# Contents

OpenCL provides an open, portable C-based programming model for highly parallel processors. In contrast to NVIDIA's proprietary programming API CUDA, a primary goal of OpenCL is portability across a diverse set of computing devices including CPUs, GPUs, and other accelerators.

- Topics covered will include:
- Introduction to vectorization • Programming model OpenCL
- Datatypes and OpenCL vectorization featuresTools for debugging and profiling
- Tuning for architectures like CPUs, accelerators (GPUs), and co-processors (Xeon Phi)
- Heterogeneous multi-device programming

Prerequisites: Some knowledge about Linux, e.g. make, command line editor, Linux shell, experience in C.

## Webpage

http://www.fz-juelich.de/ias/jsc/ events/opencl

# Advanced Topics in High Performance Computing (PATC course)

**Date & Location** April 7–10, 2015, LRZ, Garching near Munich

## Contents

In this add-on course to the parallel programming course special topics are treated in more depth, in particular performance analysis, I/O and PGAS concepts. It is provided in collaboration of Erlangen Regional Computing Centre (RRZE) and LRZ within KONWIHR. Each day is comprised of approximately 5 hours of lectures and 2 hours of handson sessions.

## Day 1

• Processor-specific optimization strategies: compiler switches, avoiding cache thrashing, exploiting SIMD capabilities (Treibig)

## Day 2

Day 3

Day 4

Webpage

courses/

(Mendez)

- Introduction to Intel Xeon Phi (Weinberg)
- OpenMP 4.0 (Weinberg)

• Parallel I/O with MPI IO (Wittmann)

• Tuning I/O on LRZ's HPC systems

• Portability of I/O: Binary files,

NetCDF, HDF5, SIONlib (Mendez)

• PGAS programming with coarray

PGAS hands on session

Fortran and Unified Parallel C (Bader)

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

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

establishing contacts.

**Tools for Parallel** 

Programming

**Date & Location** 

April 13-15, 2015

HLRS, Stuttgart

Contents

Autumn 2014 • Vol. 12 No. 2 • inSiDE



# **Open ACC Programming for Parallel Accelerated** Supercomputers (PATC course)

**Date & Location** April 16-17, 2015 HLRS , Stuttgart

## Contents

This workshop will cover the programming environment of the Cray XK7 hybrid supercomputer, which combines multicore CPUs with GPU accelerators. Attendees will learn about the directivebased OpenACC programming model whose multivendor support allows users to portably develop applications for parallel accelerated supercomputers. The workshop will also demonstrate how to use the Cray Programming Environment tools to identify CPU application bottlenecks, facilitate the OpenACC porting, provide accelerated performance feedback and to tune the ported applications. The Cray scientific libraries for accelerators will be presented, and interoperability of OpenACC directives with these and with CUDA will be demonstrated. Through application case studies and tutorials, users will gain direct experience of using OpenACC directives in realistic applications. Users may also bring their own codes to discuss with Cray specialists or begin porting.

# Webpage

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

Developing for current and future processors will more and more require parallel programming techniques for application and library programmers. During these three days we offer the industrial and scientific user community a course on the state-of-the-art of parallel programming tools, ranging from debugging tools and performance analysis to best practices in integrated developing environments for parallel platforms. The course is focused on persons who are familiar with parallel programming. Additionally to highquality presentations, the setting and organization of the workshop leaves room for technical discussions and for

# **Courses and Tutorials**

# **GPU Programming** using CUDA

**Date & Location** April 20-22, 2015 HLRS, Stuttgart

## Contents

The course provides an introduction to the programming language CUDA, which is used to write fast numeric algorithms for NVIDIA graphics processors (GPUs). Focus is on the basic usage of the language, the exploitation of the most important features of the device (massive parallel computation, shared memory, texture memory) and efficient usage of the hardware to maximize performance. An overview of the available development tools and the advanced features of the language is given.

Webpage

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

# Introduction to GPU **Programming using CUDA** (PATC course)

**Date & Location** April 20-22, 2015, 3 days JSC, Forschungszentrum Jülich

# Contents

GPU-accelerated computing drives current scientific research. Writing fast numeric algorithms for GPUs offers high application performance by offloading compute-intensive portions of the code to an NVIDIA GPU. The course will cover basic aspects of GPU architectures and programming. Focus is on the usage of the parallel programming language CUDA-C which allows maximum control of NVIDIA GPU hardware. Examples of increasing complexity will be used to demonstrate optimization and tuning of scientific applications. This course is a PATC course (PRACE Advanced Training Centres).

# Webpage

http://www.fz-juelich.de/ias/jsc/ events/cuda

Introduction to Unified Parallel C (UPC) and **Co-Array Fortran (CAF)** (PATC course)

**Date & Location** April 23–24, 2015 HLRS , Stuttgart

## Contents

Partitioned Global Address Space (PGAS) is a new model for parallel programming. Unified Parallel C (UPC) and Co-Array Fortran (CAF) are PGAS extensions to C and Fortran. PGAS languages allow any processor to directly address memory/data on any other processors.

Parallelism can be expressed more easily compared to library-based approches as MPI. Hands-on sessions (in UPC and/or CAF) will allow users to immediately test and understand the basic constructs of PGAS lanquages.

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

# **Efficient Parallel Programming with GASPI**

**Date & Location** April 27, 2015 HLRS, Stuttgart

## Contents

In this tutorial we present an asynchronous dataflow programming model for Partitioned Global Address Spaces (PGAS) as an alternative to the programming model of MPI. GASPI, which stands for Global Address Space Programming Interface, is a partitioned global address space (PGAS) API. The GASPI API is designed as a C/C++/ Fortran library and focused on three key objectives: scalability, flexibility and fault tolerance. In order to achieve its much improved scaling behaviour GASPI aims at asynchronous dataflow with remote completion, rather than bulk-synchronous message exchanges. GASPI follows a single/multiple program multiple data (SPMD/MPMD) approach and offers a small, yet powerful API (see also http://www.gaspi.de and http://www.gpi-site.com). GASPI is successfully used in academic and industrial simulation applications. Hands-on sessions (in C and Fortran) will allow users to immediately test and understand the basic constructs of GASPI. This course provides scientific training in Computational Science, and in addition, the scientific exchange of the participants among themselves.

# Webpage

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

# Scientific Visualization

Date & Location April 28-29, 2015 HLRS, Stuttgart

# Contents

This two day course is targeted at researchers with basic knowledge in numerical simulation, who would like to learn how to visualize their simulation results on the desktop but also in Augmented Reality and Virtual Environments. It will start with a short overview of scientific visualization, following a hands-on introduction to 3D desktop visualization with COVISE. On the second day, we will discuss how to build interactive 3D Models for Virtual Environments and how to set up an Augmented Reality visualization.

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

# **High-performance Computing** with Python

### **Date & Location**

May 28–29, 2015 JSC, Forschungszentrum Jülich

# Contents

Python is being increasingly used in High Performance Computing projects such as GPAW. It can be used either as a high-level interface to existing HPC applications, as embedded interpreter, or directly. This course combines lectures and hands-on session. We will show how Python can be used on parallel architectures and how performance critical parts of the kernel can be optimized using various tools.

### Webpage

http://www.fz-juelich.de/ias/jsc/ events/hpc-python

# **Courses and Tutorials**

# **Cluster Workshop**

**Date & Location** June 29-30, 2015 HLRS, Stuttgart

### Contents

Nowadays, the IT infrastructure of many companies and institutions includes modern compute clusters. The procurement, operation and the efficient usage of such parallel systems introduces new and complex requirements.

To address these emerging requirements, the High Performance Computing Center of the University of Stuttgart (HLRS) offers a vendor-independent workshop. Topics span from the design of compute clusters to details on different hardware components, operating systems, file systems and modes of operation, touching on specific software solutions. Further, typical problems and strategies for their solution will be discussed. HLRS has been running compute clusters for many years, providing compute resources for scientific and industrial simulations. HLRS is in constant dialogue with users and hardware providers and has accumulated a large knowledge base in cluster computing.

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

# Node-level Performance Engineering

**Date & Location** July 6-7, 2015 HLRS, Stuttgart

## Contents

This course teaches performance engineering approaches on the compute node level. "Performance engineering" as we define it is more than employing tools to identify hotspots and bottlenecks. It is about developing a thorough understanding of the interactions between software and hardware. This process must start at the core, socket, and node level, where the code gets executed that does the actual computational work. Once the architectural requirements of a code are understood and correlated with performance measurements, the potential benefit of optimizations can often be predicted. We introduce a "holistic" node-level performance engineering strategy, apply it to different algorithms from computational science, and also show how an awareness of the performance features of an application may lead to notable reductions in power consumption. This course provides scientific train-

ing in Computational Science, and in addition, the scientific exchange of the participants among themselves.

- Introduction
- Practical performance analysis
- Microbenchmarks and the memory hierarchy
- Typical node-level software overheads

# • Example Problem:

- The 3D Jacobi solver
- The Lattice-Boltzmann Method (LBM)
- Sparse Matrix-Vector Multiplication
- Backprojection algorithm for CT reconstruction
- Energy & Parallel Scalability

Retween each module, there is time for Questions and Answers! Attendees are highly invited to also join the course "User Guided Optimization in High-Level Languages" held on July 8 that targets related topics!

# Webpage

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

# **User-Guided Optimization in High-Level Languages**

Data & Location July 8, 2015 HLRS, Stuttgart

### Contents

While writing code in High-level languages, HPC-programmers often have a notion of gainful optimization strategies, that should be applied to their code. Nevertheless, nowadays compilers generally use heuristics to decide on optimization strategies. To lower this gap, we will present new ideas and tools to raise the programmers control over compiler optimizations, i.e.

- 1) Noise trigger optimizations by annotations
- 2) Sierra SIMD computations with compound data types
- 3) AnyDSL create DSLs with associated optimization strategies

All of these tools will also be applied in hands-on sessions. In addition, an overview of available optimization strategies will be given by compiler experts.

Noise and Sierra are limited to C/C++, even though the underlying concepts ought to be applicable also to other imperative languages, particularly Fortran. AnyDSL uses a dialect of the Rust programming language.

Attendees are highly invited to also join the "Node-level Performance Engineering"- course held on July 6–7 that targets related topics!

## Webpage

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

# **Guest Student Programme: Education in Scientific** Computing

## **Date & Location**

August 3 - October 9, 2015 JSC, Forschungszentrum Jülich

## Contents

Guest Student Programme "Scientific Computing" to support education and training in the field of supercomputing. Students of Computational Sciences, Computer Science and Mathematics can work 10 weeks in close collaboration with a local scientific host on a subject in their field. Application deadline is April 30, 2015.

# Webpage

http://www.fz-juelich.de/ias/jsc/gsp

# Contents

News	
PRACE: Results of the 9th Regular Call	4
HLRS: Extending the Simulation Capacities and Capabilities for Science and Industry	5
Applications	
From Dust to Planetesimals –	
High Resolution Simulations of Planet Formation Processes	7
New Design Principles for Biomimetic Mechanical Contacts by Computer Simulation	12
Large-Scale Order and Small-Scale Intermittency in Turbulent Convection	18
Unsteady CFD for Automotive Wheel Aerodynamics	23
2nd Extreme Scaling Workshop on SuperMUC	29
Advanced One-Sided Communication Patterns with GPI-2:	
Anisotropic Diffusion Filtering of Seismic Data	30
Three-Dimensional Simulations of Core-Collapse Supernovae with the VERTEX Code	33
Studying Wave-Particle-Interaction with the ACRONYM PiC Code on SuperMUC	36
Computational Engineering goes HPC:	
Thermal Comfort Assessment on Massive Parallel Systems	39
Hybrid Ateles on SuperMUC	41
CIAO Code – Extreme Scaling Workshop LRZ 2014	44
Extreme Scaling of NSCOUETTE, a Pseudospectral DNS Code	48
Extreme Scaling of the PSC on SuperMUC at LRZ	51
Towards Environmental Computing on e-Infrastructures	55
ExaStencils – Advanced Stencil-Gode Engineering	60
EXCESS – Execution Models for Energy-Efficient Computing Systems	64
HUNK - High-Under Methods for the Simulation of Complex Flow Phenomena	6/
Successful Technology Transfer with Siemens – The RAPID Project	15
Systems	76
Centres	80
Activities	86
Courses	107