inside

AUTUMN 2013

inSiDE • Vol. 11 No. 2

Innovatives Supercomputing in Deutschland

inside 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
Linda Weinmann, HLRS weinmann@hlrs.de
Antje Häusser, HLRS haeusser@hlrs.de
Sebastian Zeeden, HLRS zeeden@hlrs.de
Stefan Steinbach, HLRS steinbach@hlrs.de

Cover: Laminar-turbulent transition on a common dolphin at 1 m/s and 1 % turbulence intensity. Transition location visualized by turbulent-kinetic-energy contours combined with isolines of pressure coefficient. Streamlines indicate part of oncoming flow forced towards the lower body. (simulation: D. Riedeberger, IAG; University of Stuttgart, Germany, model: V. Pavlov, ITAW, University of Hannover, Germany)

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, we have a look at the status of PRACE, present applications, look at projects that drive Supercomputing in Germany and Europe and provide a wealth of information about German supercomputing systems, training and activities.

GCS is a strong player in PRACE and it is hence very important for GCS in which direction PRACE is headed. We open this issue with a report on the results of the last PRACE call and a report on a meeting of PRACE projects to discuss their results and challenges. PRACE also plays an important role in the interview with the new Chairman of the Board of Directors of GSC Prof. Michael Resch.

HPC is also intensively discussed in the German research community. Currently a working group of the Wissenschaftsrat (Science Council) is evaluating the existing German funding scheme for High Performance Computing and discusses possible necessary changes. We will report on the outcome of this discussion. In its general paper on the future perspectives of the German science system (published on July 12 this year) the Wissenschaftsrat put a lot of emphasis on the importance of simulation in general and on High Performance Computing in particular for the future of science & research.

With respect to hardware we expect to see new systems at GCS in 2014 and 2015 respectively. We already reported on the new systems at LRZ briefly in our last issue. HLRS has extended its contract with Cray for a second step in 2014 and plans to push the system up to more than 7 PFLOP/s late in 2015. We will report on the architecture of this system in our next issue. In terms of infrastructure HLRS is moving its plans for an HPC training center building and will start construction work in 2014.

Again, we present a rich section on applications using simulations on the GCS systems. The growing strength of GCS in HPC research is highlighted by the research project reports that show the variety of fields in which the GCS centers further push the limits of supercomputing.

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 for participate in these workshops which are now part of the PRACE Advanced Training Center and hence open to all European researchers.

Enjoy reading!

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

Contents

Contents

PRACE: Results of the 7th Regular Call

The Partnership for Advanced Computing in Europe (PRACE) is continuously offering supercomputing resources on the highest level (tier-0) 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 7th call for proposals for computing time for the allocation time period September 3, 2013 to September 2, 2014 on the above systems closed March 26, 2013. Five research projects have been awarded a total of about 178 million compute core hours on SuperMUC and six have been awarded a total of about 120 million compute core hours on Hermit. On JUQUEEN, three proposals have been awarded a total of about 52 million compute core hours, while 48 million compute core hours were given to already approved long-term projects from previous calls.

Four of the newly awarded research projects are from France, two are from Germany, Italy, and Switzerland, each, and one is from Finland, Ireland, the Netherlands, and the United Kingdom, each. The research projects awarded computing time cover again many scientific areas, from Astrophysics to Medicine and Life Sciences. More details, also on the projects granted access to the machines in France, Italy, and Spain, can be found via the PRACE web page www.prace-ri.eu/PRACE-7th-Regular-Call.

The 8th call for proposals for the allocation time period from March 4, 2014 to March 3, 2015 closed on October 15, 2013 and evaluation is still under way, as of this writing. The 9th call for proposals will open in February 2014.

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

• Walter Nadler

Jülich Supercomputing Centre (JSC)

PRACE Projects meet in Varna

Over 100 researchers from the three PRACE Implementation Phase projects convened for their annual all-hands meeting in Varna, Bulgaria, from 3 to 5 June 2013. These meetings are a unique opportunity to inform the participants in the projects about the current status and achievements of PRACE, coordinate and synchronize the work between the work packages, and plan the steps for the coming year.

Starting in 2010 the European Commission supported the implementation of the PRACE Research Infrastructure with a total of 57 million Euro. The grants were awarded to three individual, partly overlapping projects. Although the overall objective for all three is supporting the accelerated implementation of the pan-European high-end computing Research Infrastructure created and operated by PRACE, each of them has a specific focus.

Careful planning coordination is essential to avoid duplication of work in the 29 work packages and among the nearly 400 collaborators from 25 countries especially if the descriptions sound similar like dissemination, training, operations, or application support. Each of the three projects has to cover these areas and does so addressing complementary aspects or taking over after completion of a predecessor project.

Participants of the PRACE All-hands meeting in Varna, Bulgaria.

A key element to achieve this is the fact that Jülich Supercomputing Centre is the coordinator of all three PRACE projects.

On Monday, June 3, highlights of PRACE-1IP were presented. This provided a comprehensive overview of the results achieved in the first project's years setting the stage for the follow-on projects PRACE-2IP and PRACE-3IP. The majority of the work was completed in 2012 and approved by the European Commission. Until the end of 2013 the evaluation of prototypes for multi-petascale technology and the porting of applications to the emerging Intel MIC technology will complete the work. Prof. Thomas Lippert, coordinator of the project, stressed that the project has been an indispensable step for the implementation of the PRACE Research Infrastructure. Among its many achievements, he highlighted in particular the initiation and support for the update of the Scientific Case for HPC in Europe that supports the efforts of PRACE Research Infrastructure to maintain HPC as a high-priority item on the political agenda in Europe.

The remaining two days were dedicated to in-depth discussions of the future work of the PRACE-2IP and PRACE-3IP projects within the individual work package and between related work packages. The largest part of the work of PRACE-2IP will end in August; therefore, an important topic was the hand-over to PRACE-3IP of ongoing activities such as application enabling support, operation of the six PRACE Advanced Training Centres (PATC), operation of the distributed HPC infrastructure, and the Distributed European Computing Initiative (DECI) calls, which provide access to national HPC resources (Tier-1). The final arrangements for ISC'13 in Leipzig were made. This included the PRACE booth at the exhibition, the 'Dare to think the Impossible' outreach campaign to be launched at ISC'13 and the PRACE Scientific Conference 2013. Another major element of PRACE-3IP is a joint pre-commercial procurement (PCP) pilot, focusing on whole system design for energy-efficient HPC. The PCP tender will be published in October, and the meeting provided a good opportunity to continue preparing the documents for the related focused EC review.

The PRACE projects are especially grateful to Prof. Stoyan Markov and his team from our Bulgarian partner NCSA, the National Centre for Supercomputing Applications, for organizing this very productive meeting in a secluded hotel on the Black Sea near Varna. The excellent conference facilities provided a satisfying working environment that contributed greatly to the success of this all-hands meeting.

• Florian Berberich
• Dietmar Erwin

Jülich Supercomputing Centre (JSC)

Interview with Prof. Dr.-Ing. Dr. h.c. Dr. h.c. Michael M. Resch

Prof. Resch, you have been elected Chairman of the Board of the Gauss Centre for Supercomputing in March. How do you see the role for the Chairman of GCS?

First of all I feel honoured to be elected Chairman of the Board of GCS which is – we should not forget – the most powerful HPC organization in all of Europe. I follow Prof. Bachem and Prof. Hegering as Chairmen. Both are strong leaders with an outstanding scientific record and career and have given GCS stability and a reputation that is outstanding worldwide. The centres of GCS have been working together successfully for 6 years now and I see it as my duty to build on these achievements.

Basically the role of the Chairman of GCS is twofold. First, the Chairman has to coordinate and communicate internally. GCS is a collaboration of three centres but also has grown its own organization with a strong CEO. As a Chairman it is vital to make sure that the communication between the three centres and the CEO is open and well co-ordinated. We have established an excellent communication over the last years and I see it as my role to guarantee that openness.

Second, the Chairman has to communicate externally. This includes both the communication with German funding agencies and European initiatives like PRACE. On the other hand it is also important for the Chairman to make sure that GCS is communicating well with its user community. Having established a GCS steering committee – which represents our users – I will work closely with the Chairman of that committee to make sure GCS meets the needs of its users as well as possible.

You mentioned the long term collaboration of GCS – which is now in its seventh year. What are the key success stories of GCS from your point of view?

First and foremost it is a big success that we were able to bring together the three national centres into a single organization. This allows us to communicate both with funding agencies and users in a coherent way. Prior to the foundation of GCS, the three centres were only loosely coordinated, but GCS very much put this into a joint organizational and political framework.

To give on example: In the time before GCS the coordination of procurements and installations had established a certain architectural diversity and some kind of investment and installation timing which was not really coordinated. But with GCS we were able to set up an installation schedule and an architectural variety which now perfectly fits the needs of our user communities. We see further improvements in user project evaluation, user access, training, security, and public relation activities like inside.

News

Which changes should we expect from you?

There certainly are some aspects where my personal background will lead to the setting of new or additional priorities. I have worked with the industry for more than 10 years rather successfully and have identified a number of fields in which such an industrial-scientific co-operation can create a win-win situation. Over the last years, I furthermore have focussed on the integration of HPC into industrial and scientific work flows with an eye on solutions both for scientific and societal problems. These two aspects will certainly become more important in the future.

What are the future challenges?

The biggest future challenge is to make the next big step forward. The member centres of GCS are participating in an international competition that requires changing rapidly and at the same time keeping the excellent services and research activities at a very high standard. The general international trend of focussing on solutions requires that GCS and its centres adapt and improve constantly. GCS has built an ecosystem that provides everything from systems to services and research. This ecosystem needs constant improvement to adapt to the rapidly changing HPC landscape both in terms of architectures and user requirements, and it needs further funding.

GCS, however, is also part of an ecosystem. In Germany this requires a better co-operation with the Gauss Alliance. I consider the Alliance to be one of the partners for GCS and the key to a nationally coherent HPC landscape.

Last but not least GCS is faced with the many questions that arise from the PRACE projects and its activities as represented by PRACE AISBL. The first phase of PRACE has seen a substantial German contribution of 100 Million Euro provided over the last years to the European user community. A number of European users have achieved substantial breakthroughs in their respective research fields making excellent use of the German systems JUQUEEN, Hermit and SuperMUC. Any continuation of PRACE will have to honour these contributions and will have to provide a solution for the problem of financial and organizational sustainability for a truly European scientific HPC ecosystem.

What is your view of PRACE?

Europe has a variety of instruments that could help create a sustainable infrastructure. These options have to be explored openly and national contributions have to be taken into consideration. At the same time we have to find out what the best response to the true challenges of HPC is. Architectures change rapidly in HPC. Thus a one-of-a-kind project can easily turn out to be rather counterproductive. We have seen such cases worldwide over the last years. A variety of both architectures and support expertise will be absolutely necessary in the coming decade. The ability of PRACE to bring together national players will be vital to keep that variety thriving. Locality is a second important issue. While most experts assume that "The Cloud" is a solution to the HPC problem by providing cycles at request I believe that "The Cloud" is not a solution but means regress. Already 15 years ago leading HPC centres worldwide understood that raw cycles are only a small part of the solution required by the users. A bundling of cycles, operational support, services, and user support will be vital for Europe.

Many people say that the US and Japan benefit substantially from having national HPC industries. Would such a European player be an option to help European science?

It is obvious that in terms of funding national centres can easily devise an excellent strategy to get national funding if they can get the support of a national industrial partner. The huge budgets of the Japanese K-Computer and of several American centres speak for themselves. Once governments look at HPC from a purely economic point of view and take a protectionist view for national industries HPC centres can create a strong political backing. However this coin also has a flip side. Reducing HPC to a national economic problem puts science in the second place. Limiting procurements to only national players may furthermore reate a vendor-market which puts centres at a disadvantage when negotiating the best deal for the tax payer's money. Given the current market situation I personally believe that the open and competitive HPC market that distinguishes Europe both from the US and Japan is an asset which scientific HPC centres should further strive to keep.

Where do you want to be in two years from now when your term as a chairman of GCS will end?

Given the current situation I expect that GCS will have secured its funding for the coming 5-7 years. This is certainly the most challenging issue and we are on a very good path forward. Furthermore the problems of PRACE will have to be solved as the official agreements on PRACE will end in 2015. Finally I hope that GCS and its member centres will have made two steps forward by providing the right solutions to tackle the pressing problems of science, industry and society.

Prof. Resch, thank you for the interview. The interview was conducted by the inside team.

Prof. Michael M. Resch is the Chairman of the Board of Directors of GCS since May 2013. He was a cofounder of GCS. Prof. Resch is currently the director of the High Performance Computing Center Stuttgart (HLRS), the director of the Information Service Center (IZUS) of the University of Stuttgart and the director of the Institute for High Performance Computing (IHR) of the University of Stuttgart.

Simulating the Life Cycle of molecular Clouds

Star formation takes place in the densest and coldest gas in a star-forming galaxy, in so-called molecular clouds (MCs). MCs do not evolve in isolation. They are highly dynamical objects, which are born, fed, heated, and stirred from their turbulent environment into which they eventually dissolve. They form in regions where the hot or warm, ionized and atomic interstellar medium (ISM) condenses into cold (T< 300K), molecular gas. Often concentrated to the midplane of galactic disks, this process involves metallicity-dependent, non-equilibrium chemistry and molecule formation, heating and cooling, turbulence, self-gravity, and magnetic fields. Once formed, MCs further collapse to form stars and star clusters.

Observations in the Milky Way and nearby galaxies show that individual MCs convert only a few percent of their mass into stars [5]. Less than 1% of all new-born stars are more massive than 8 solar masses, but these are particularly important for galaxy formation and evolution. The life and death of massive stars differ intriguingly from those of their low-mass counterparts. Such stars affect their environment dramatically through their strong UV radiation, their energetic stellar wind, and their final explosion as supernovae. These 'feedback' processes generate turbulence in the parental molecular cloud, dissociate, ionize, and eventually destroy them from within, thereby quenching further star formation. Stellar feedback is thus thought to regulate the star formation efficiency in molecular clouds leading to a self-regulation of star formation on galactic scales (see [9] and references therein).

At the same time feedback from massive stars can drive powerful galactic outflows, removing gas from the host galaxy. In this way star formation may control the galactic gas reservoir available for star formation and regulate galaxy formation on cosmic scales. In support of this idea, high star formation rates linked to powerful galactic outflows were recently observed for gas-rich, star-forming disk galaxies at the peak of cosmic star formation at z=2-3 [12]. Galactic winds may also be of particular importance for driving the Galactic fountain [10] and removing gas from dwarf galaxies which, today, have a lower than average ratio of baryons to dark matter [11].

Understanding how feedback effects propagate from the milli-parsec scales, characterizing the sites of massive star formation, to the kilo-parsec scales of molecular cloud complexes and galactic winds is a physical and numerical challenge. In the framework of the Gauss Tier-0 project "SILCC" we model representative regions of disk galaxies using adaptive, three-dimensional simulations at unprecedented resolution and with the necessary physical complexity to follow the full life-cycle of molecular clouds. These simulations include self-gravity, magnetic fields, heating and cooling at different gas metallicities, molecule formation and dissociation, and stellar feedback in form of stellar winds and supernovae. The simulations use novel, massively parallel techniques, and are computationally very demanding. They are only feasible on a peta-scale supercomputer like SuperMUC at the Leibniz Supercomputing Centre (LRZ) in Garching near Munich.

The ultimate goal of the SILCC project is to provide a self-consistent answer as to how stellar feedback regulates the star formation efficiency of a galaxy, how molecular clouds are formed and destroyed, how galactic outflows are driven and what parameters their mass-loading depends upon. Here we discuss first results showing that the fraction of gas, which is in molecular form, is steadily increasing with the total gas surface density Σ_{gas} of the disk.

Numerical Method

The simulations are being carried out with the massively parallel, Eulerian, adaptive mesh refinement code FLASH 4.0.1 [6]. FLASH is highly modular and includes many physics units, like MHD, an external gravitational potential, or recently also a Barnes & Hut tree method for self-gravity, in the standard release. However, much more physics is required in order to carry out the proposed simulations. In preparation for this project we added new physics modules, which are not present in the standard release.

In particular, we implemented a chemical network that allows us to follow all three phases in the ISM, and which includes non-equilibrium ionization effects as well as H_2 and CO formation [7]. Also radiative heating and cooling are computed by this module. The heating and cooling terms, and also the chemical source and sink terms, are operator split from the advection and pdV work terms and are solved implicitly using the DVODE solver [2]. In cases where the chemical or cooling timesteps are much shorter than the hydrodynamic timestep, sub-cycling is used to treat the cooling and chemistry, thereby avoiding the need to constrain the global timestep. Chemical abundances are represented as tracer fields and are advected using the standard FLASH infrastructure for such fields.

In order to treat the formation of molecules, which is essential for comparing the simulations with observations, we implemented TreeCol [4] to compute shielding and self-shielding of H_2 and CO as well as the dust attenuation and dust temperature. TreeCol is a cost-efficient way to treat diffuse radiation. From every cell within the computational domain, a set of rays is casted along which the e.g. column densities are computed during the self-gravity tree-walk. The ray-casting algorithm is based on HealPix [8], an equal surface area pixilation of the sphere. Therefore, the minimum number of rays is 12, and we typically use 48 rays.

Furthermore, we implemented Supernova and stellar wind feedback. Supernovae can be placed randomly with a given Supernova rate as used in the presented simulations. Alternatively, the explosions might be launched from gas density peaks within molecular clouds or from sink particles. Each Supernova injects 10^{51} erg and a few solar masses of gas into the ISM, which we add in form of internal energy and additional mass to cells close to the explosion center, i.e. within the injection region. Before the explosion, the injection region is always refined to the finest AMR level in order to limit grid effects and over-cooling. We force a minimum number of 4 cells per radius of the injection region, but the radius may be larger if the mass within the injection region is small ($\leq 30\ M_{sun}$). With these parameters the temperature within the injection region is typically raised to ~10^7 K, and therefore the hydrodynamic timestep is recomputed and limited accordingly. Stellar winds are implemented in a similar way.

Last but not least, we use sink particles with an adequate sub-grid model to describe star clusters for simulations with lower resolution, or small N-clusters of accreting stars in case of high resolution.

We simulate representative regions of disk galaxies (500pc x 500pc x ± 2 kpc) with different gas surface densities, i.e. 10, 30, 100, and 500 M_{sun}/pc^2, to mimic the conditions of star-forming galaxies at low and high redshift. The initial density distribution of the gaseous disk is Gaussian. The boundary conditions are periodic in x- and y-, and allow for outflow in z-direction. To account for the aforementioned physical processes on spatial and temporal scales ranging from several 100 AU and 10 years to kpc-scales and 10^7 years, each run requires more than 5 million CPU hours.

The typical Milky Way Disk

Close to the solar radius (at a distance of 8.5 kpc from the Galactic Center) the Milky Way disk has an average gas surface density of 10 M_{sun}/pc^2. We present first results from simulations at an intermediate resolution, where we include turbulence driving from randomly placed Supernova explosions. The random positions are weighted with a Gaussian with a scale height of 50 pc. Assuming a certain stellar initial mass function (IMF; e.g. [3]) to estimate the number of massive stars per unit mass that is forming stars, the mean Supernova rate (SNR) for such a disk can be derived from the observed star formation rate – gas surface density relation. Observations yield nearly linear scaling of $\Sigma_{SFR} \sim \Sigma_{gas}$ [1]. Within the simulated volume, the SNR for a Milky Way-type disk is ~10 Myr^{-1}. Here we use a SNR of 14 Myr^{-1}.

The simulations show that a dynamical equilibrium is only established if the SNR does not significantly exceed or fall below the observed SNR. In Fig. 1, we show projections and slices through the (y=0)-plane. In the top row we show (a) total column density, (b) density slice, (c) temperature slice, and (d) dust temperature slice; in the bottom row we plot the column densities of (e) atomic hydrogen (HI), (e) ionized hydrogen (HII), (f) molecular hydrogen (H_2), and (g) CO. In the disk mid-plane the gas is fully molecular unless a Supernova explodes within or nearby a molecular cloud. The point-like explosions heat up the surrounding gas to 10 million Kelvin. The expanding blast wave then compresses the ISM into a denser shell, within which the dust temperature is temporarily increased. On the other hand, new molecular gas forms in very cold (< 10 K) molecular clouds within the midplane. Within the molecular clouds the gas is fully molecular and the abundances of H_2 and CO are well correlated. However, there is molecular hydrogen at higher altitudes, which is not well traced by CO. This gas, which has been driven out of the molecular clouds by Supernovae, is CO-dark and cannot be seen in observations. The reason that CO does not survive at higher altitudes is that it requires a higher column density to shield itself from the interstellar (ultraviolet) radiation field.

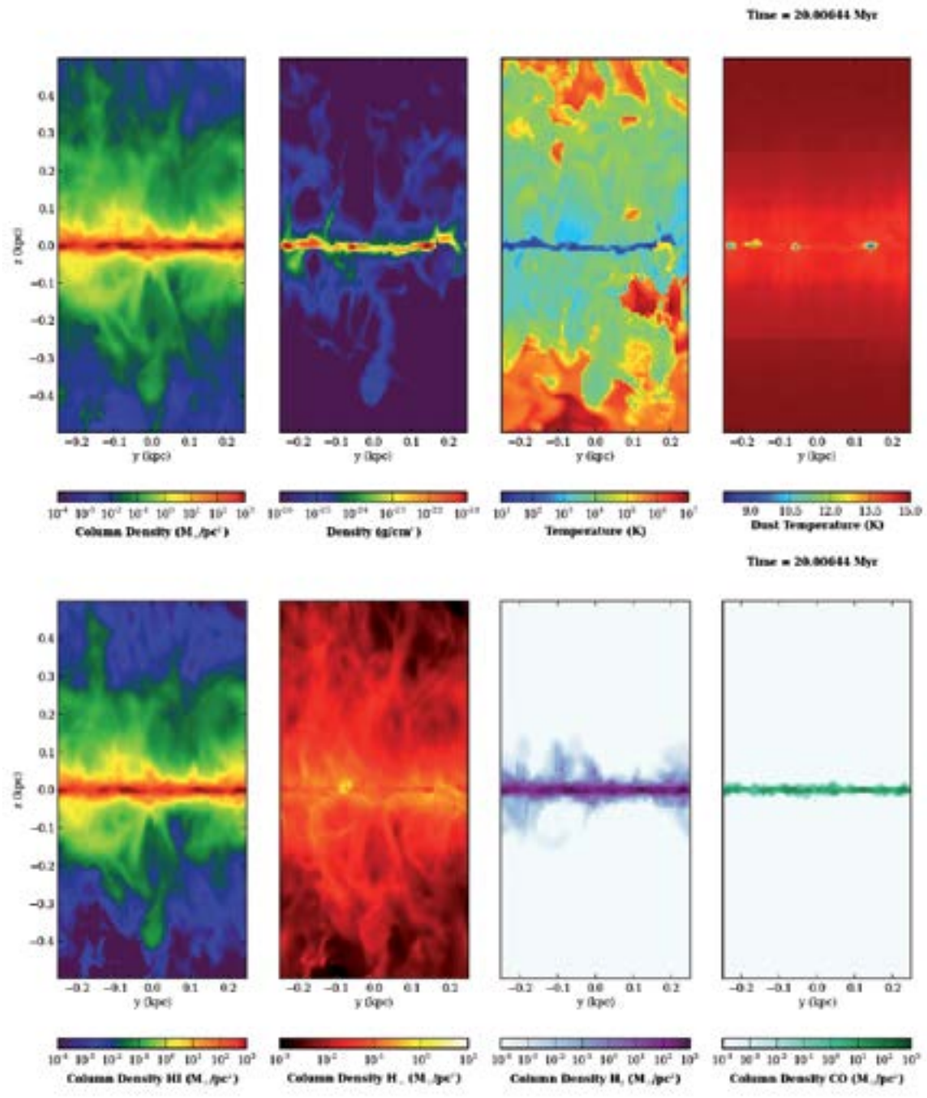

Figure 1: Projections and slices through the y=0 plane at t=20 Myr. From top left to bottom right the panels show (a) the total column density, (b) a density slice, (c) a temperature slice, (d) a dust temperature slice, and the column density of (e) atomic hydrogen (HI), (f) ionized hydrogen (HII), (g) molecular hydrogen (H_2), and (h) CO.

Disk Galaxies at different Gas Surface Densities: from low to high Redshift

Simulations of disks with higher gas surface density show that the mass fraction molecular gas is increasing with increasing surface density. This is the case even though the Supernova rate is correspondingly higher in high surface density disks, following a roughly linear scaling relation of star formation rate (hence Supernova rate) and molecular gas surface density [13].

In Fig. 2, we show the molecular gas mass fraction as a function of time for the four different initial disk surface densities: Σ_{gas} =10, 30, 100, and 500 M_{sun}/pc^2. The corresponding SNRs are 14, 42, 140, and 700 Myr^{-1}.

For the highest gas surface density of 500 M_{sun}/pc^2, the H_2 mass fraction rapidly approaches values of up to 90 %. At 100 M_{sun}/pc^2 the H_2 mass fraction saturates at 70 – 80 %. This result is in good agreement with observational estimates of the molecular gas mass fraction in centers of local galaxies as well as gas-rich high redshift disks [13]. For lower initial gas surface densities the molecular fractions are respectively lower and it takes longer to reach the saturation values. The Milky Way value of roughly 50 % (interior to the Sun's galactic orbit) agrees well with the 10 M_{sun}/pc^2 disk simulation, for which peak values of 40 % are achieved.

Even though the mass fractions are high the fractional volume filled by molecular gas is small, as it represents the densest and coldest gas component. Therefore, simulating the life-cycle of molecular clouds requires simulations with very high spatial resolution, that are currently performed on SuperMUC.

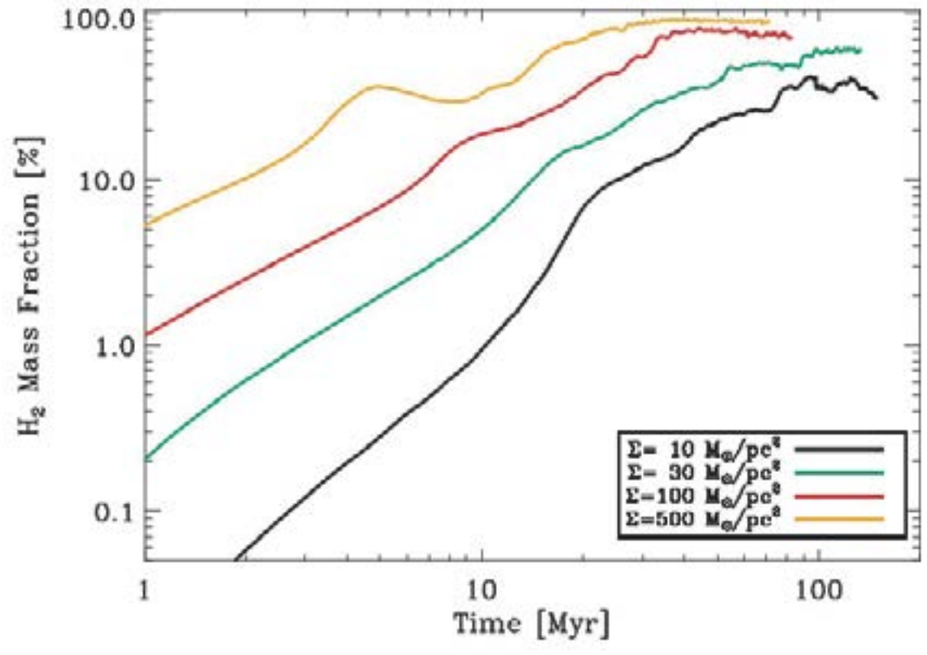

Figure 2: Mass fraction of molecular hydrogen as a function of time for four different disk surface densities.

Conclusions and Outlook

With a European team of experts from Garching, Heidelberg, Prague and Zurich we have performed the first simulations of the turbulent ISM in a galactic environment including the self-consistent formation of molecular Hydrogen and CO. In addition, the simulations include a model for energetic feedback during the late evolutionary phases of massive stars including SN explosions. With a first set of challenging simulations we have been able to demonstrate that gaseous galactic disks with properties similar to the Milky Way reach equilibrium molecular gas fractions of about 40 per cent. At higher surface densities the disks settle at higher molecular gas fractions, up to 90 per cent for 500 Msun/pc2, which is a typical value for massive galactic disks at high redshift. Moreover, we are able to formulate new constraints on the amount of CO-dark molecular gas and will be able to constrain gas outflow properties. We have already started the second phase of the SILCC project with simulations at higher spatial resolution, following the formation of small clusters of massive stars. We will resolve the relevant spatial scales to investigate the self-consistent driving of galactic winds. We will also study the important effects of magnetic fields on the turbulent multi-phase structure of galactic disks.

References

[1] **Bigiel, F., Leroy, A., Walter, F., Brinks, E., de Blok, W. J. G., Madore, B., Thornley, M. D.**
2008, AJ, 136, 2846

[2] **Brown, P. N., Byrne, G. D., Hindmarsh, A. C.**
1989, SIAM J. Sci. Stat. Comput., 10, 1038

[3] **Chabrier, G.**
2002, ApJ, 567, 304

[4] **Clark, P. C., Glover, S. C. O., Klessen, R. S.**
2012, MNRAS, 420, 745

[5] **Evans, II, N. J. and et al.**
2009, ApJS, 181, 321

[6] **Fryxell, B., et al.**
2000, ApJS, 131, 273

[7] **Glover, S. C. O., Mac Low, M.-M.**
2011, MNRAS, 412, 337

[8] **Górski, K. M., Hivon, E.**
2011, Astrophysics Source Code Library, 7018

[9] **Mac Low, M.-M., Klessen, R. S.**
2004, Reviews of Modern Physics, 76, 125

[10] **Marinacci, F., Fraternali, F., Nipoti, C., Binney, J., Ciotti, L., Londrillo, P.**
2011, MNRAS, 415, 1534

[11] **Moster, B. P., Somerville, R. S., Maulbetsch, C., van den Bosch, F. C., Macciò, A. V., Naab, T., Oser, L.**
2010, ApJ, 710, 903

[12] **Newman, S. F., et al.**
2012a, ApJ, 752, 111

[13] **Tacconi, L., et al.**
2013, ApJ, 768, 74

- Stefanie Walch[1]
- Thorsten Naab[1]
- Philipp Girichidis[1]
- Andrea Gatto[1]
- Simon Glover[2]
- Richard Wünsch[3]
- Ralf Klessen[2]
- Paul Clark[2]
- Thomas Peters[4]
- Christian Baczynski[2]

[1] Max-Planck-Institut für Astrophysik Garching

[2] Institut für Theoretische Astrophysik Heidelberg

[3] Institute, Academy of Sciences of the Czech Republic, Prague

[4] Zurich University, Zurich

Gadget3: Numerical Simulation of Structure Formation in the Universe

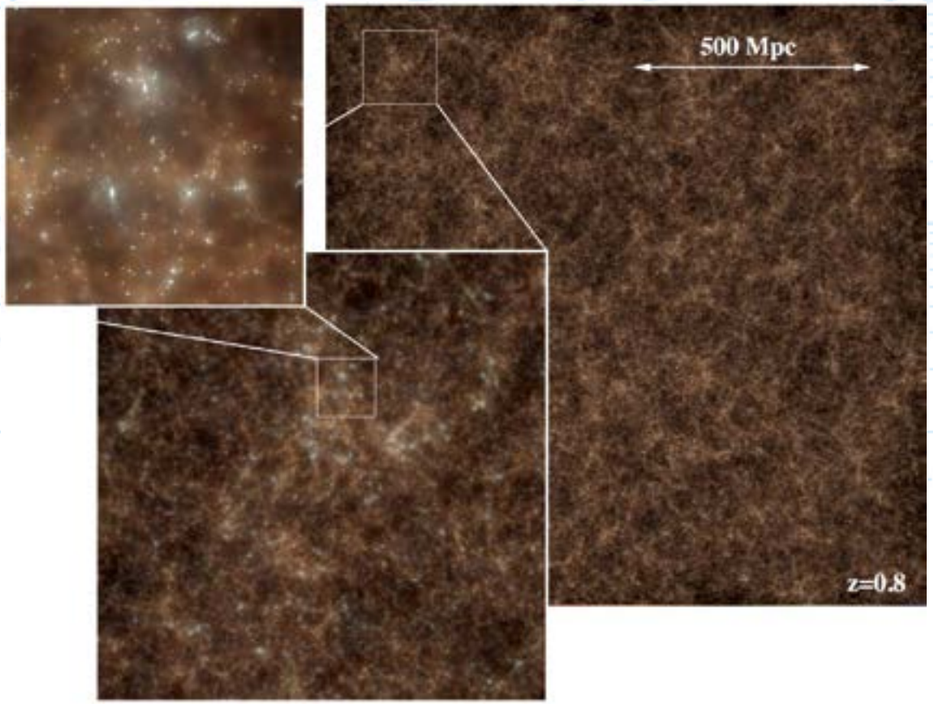

Figure 1: Visualization of a 1Gpc^3 cosmological box, showing the filamentary distribution of matter in the Universe and zooming onto a super cluster region. Simulation was performed with Gadget3.

Cosmological simulations play an important role in theoretical studies of structure formation in the Universe. Such simulations are an essential tool to accurately calculate theoretical predictions of the distribution and state of the baryonic and dark matter in the Universe, especially in the non-linear regime of gravitational dynamics and hydrodynamics, where galaxies and clusters of galaxies form out of the large scale structure. Realistic modeling must include baryonic matter as well as dark matter component and has to be able to describe physical processes such as star formation, supernovae, and AGN feedback, as well as transport processes and magnetic fields. Therefore, numerical simulations must be capable of properly resolving and following those phenomena using different numerical techniques simultaneously.

The KONWIHR-III project "Tuning a cosmo Gadget for SuperMUC" is funded by the Federal State of Bavaria. It aims at optimizing Gadget3 [1,2], a N-body Magneto-Smoothed-Particle-Hydrodynamics code for cosmological simulations, for which we have been granted computing time on SuperMUC at the Leibniz Rechenzentrum.

From the astrophysical point of view, current cosmological simulations are limited by the number of tracer particles which can be used. On one hand, large volumes need to be simulated to sample a representative part of the Universe. But on the other hand, even utilizing several 10^9 particles (as for the visualization shown in Fig. 1), the resulting resolution is still poor (~30kpc) and such simulations do not resolve properly individual galaxies. Hence, larger volumes and higher resolution are needed to produce a theoretical counterpart to interpret the data coming from current and forthcoming astronomical surveys. To achieve this, future simulations must reach 10^{11} particles and beyond following various physical processes. Being able to perform such large hydrodynamical cosmological simulations will, for the first time, allow a detailed comparison with a variety of multi-wavelength observations.

From the numerical point of view, Gadget3 is a massively parallel code, which employs various algorithms to treat different physical processes. Originally (with exception of the gravity solver) it relies on a pure MPI implementation. However, on current and future hardware configurations like

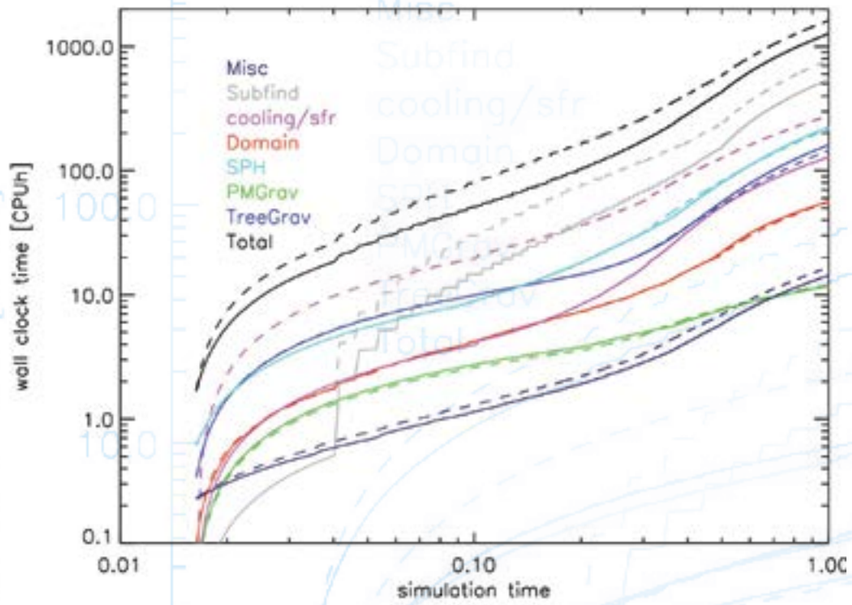

Figure 2: Cumulative CPU time as a function of simulation time given as dimensionless cosmological parameter (1 is present) broken down for the most computationally intensive parts of Gadget3. The dashed lines represent the Gadget3 version at project start and the solid lines represent the latest version after optimization. Individual parts show very large performance gain (except for the gravity tree where the correction of data race conditions slightly slows down the algorithm) and we even see an overall total improvement of almost 30%.

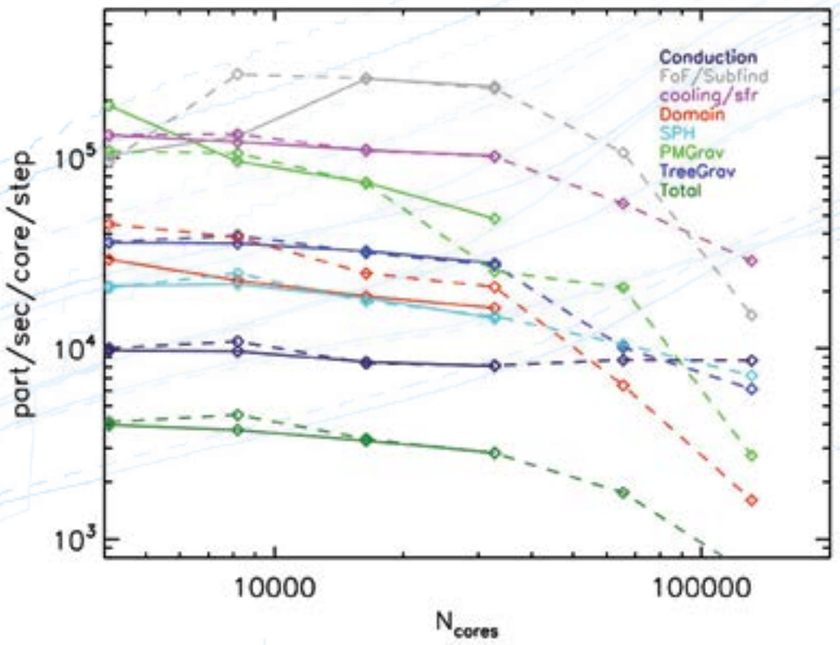

Figure 3: Strong scaling of Gadget3 on SuperMUC. The plot shows the number of particles processed per core, time step, and wall clock time, broken down for the most computationally intensive parts of the calculation and spans from 0.5 to 16 islands of SuperMUC. Solid and dashed lines indicate two different communication structures which can be used.

SuperMUC, the code is expected to run on hundreds of thousands cores, making a pure MPI implementation challenging for such a large number of MPI tasks. Additionally, previously subdominant parts of the code as well as the communication starts to strongly dominate and eventually make the code unusable on large MPI task numbers. Therefore, Gadget3 was converted to a hybrid code, making intensive use of either posix threads (as in the original gravity solver) or OpenMP (as focused in this project) for node level multithreading. We therefore aimed at identifying bottlenecks, implementing more shared memory parallelized parts of the code and optimizing Gadget3 to be usable as hybrid MPI/OpenMP code on HPC facilities like SuperMUC.

Following some of the processes in details (e.g. evolution of black holes), as well as the on-the-fly post processing (which is essential for scientific analysis of such simulations), involves finding and identifying substructures at simulation run time. Substeps, like the density estimation and gravitational unbinding check of material within substructures, are the most challenging and essential parts of the post processing. Therefore, optimization and OpenMP parallelization of those parts is a good example of the main work done within this KONWIHR-III project. Substructure identification has to be performed typically 100 times during the simulation. Within this process, density calculation (for all particles) is one of the most time consuming parts (up to 75% of the total execution time). In that case it is possible to distribute particles over multiple threads and then process many particles in parallel on each node. Additionally, various sort operations have to be performed which gain significantly using multithreaded sort algorithms. Doing so, we improved the overall performance of these parts of the algorithm by a factor 1.8 for the typical setup of 2 MPI tasks per node and 16 OpenMP threads per task which we use for large simulations on SuperMUC.

More generally, implementing multithreading in several parts of the code, especially in the most computationally intensive parts, has led so far to a total improvement in performance of almost 30% (Fig. 2).

Furthermore, our team was the first one to run Gadget3 exploiting 16 islands of SuperMUC. We successfully ran Gadget3 on 0.5, 1, 2, 4, 8, and 16 islands (i.e. 131,072 cores) using the setup given above (Fig. 3). We demonstrate that the code performs very well for a strong scaling up to a significant fraction of SuperMUC. The scaling is very close to the ideal case up to 4 islands. Beyond this point, it is encouraging to see that the improved version of Gadget3 performs reasonably well and can effectively make use of 16 thin node islands, even when various additional physical processes are switched on. This result is quite promising for future simulations on SuperMUC and for further optimizations we continue to work on.

References

[1] **Springel, V.**
2005, MNRAS, 364, 1105-1134

[2] **Springel, V., Yoshida, N., White, S.D.M.**
2001, New Astronomy, 6, 79-117

• Gurvan Bazin[1]
• Klaus Dolag[1]
• Nicolay Hammer[2]

[1] Universitätssternwarte, Ludwig-Maximilians-Universität München

[2] Leibniz-Supercomputing Centre (LRZ)

Numerical Simulation of Correlated Electron Systems

Correlated electron systems belong to the general class of complex systems whose defining property is that the whole is more than the sum of the individual parts. This emergent complexity gives rise to a rich variety of fascinating phenomena in fields ranging from biological systems to materials. At the same time, the complexity poses a challenging scientific problem: Even if we find a way to solve the corresponding equations of motion of the constituents of a many-body system, it is by no means guaranteed that we can understand the emergent collective phenomena.

In the case of crystalline materials, we can start with a quantum-mechanical description of the electronic motion in the periodic potential of the ions in terms of Bloch states. From these states, we can construct the many-body wavefunction by taking into account the Pauli exclusion principle that demands the wave function to be antisymmetric under exchange of identical fermions. These two ingredients, Bloch states and Fermi statistics, underlie band structure theory and produce a host of very interesting states, most notably insulators and metals with completely or partially filled highest occupied band, respectively. As discovered theoretically by Kane and Mele [1], band theory also permits more exotic insulating states referred to as topological insulators. These states, which arise from strong spin-orbit coupling and time-reversal symmetry, are characterized by insulating behaviour in the interior but metallic behaviour at the edge or surface of the sample. Importantly, the characteristic surface states are robust with respect to disorder, and can therefore be observed in experiments [2]. In two dimensions, the edge states are helical, that is the direction of motion of electrons is coupled to the direction of their spin [1]. In principle, this property and the topological protection make topological insulators a promising candidate for applications in the field of spintronics, which offers the vision of overcoming the problem of heating inherent to current CPU technology.

Our discussion so far has neglected electron-electron interactions. However, as students learn already in introductory courses on solid state physics, the energy scale associated with such correlations is very large, typically several electron volts. Nevertheless, in many cases, electronic correlations do not significantly alter the properties of the system – they are irrelevant in the language of Wilson's renormalization group. This remarkable fact hinges on the screening of the Coulomb repulsion, and Fermi statistics which drastically limits the phase space available for scattering due to the Coulomb repulsion, and is known as Fermi liquid theory.

Materials in which Fermi liquid theory does not hold are of particular interest. In such systems, correlations lead to exotic and/or collective states of matter. In one dimension, the concept of the Fermi liquid breaks down and correlations cause a fractionalization of electrons into independent, collective spin and charge excitations. More generally, Mott insulators are correlated materials which are incorrectly predicted by band theory to be metallic. The insulating behaviour is often attributed to the observed ordering of, for example, the spin degrees of freedom (that is, the formation of a magnetic state). However, a central question in this context is if these two phenomena are always coupled. Remarkably, the answer to this question is no. Mott insulating states with no associated symmetry breaking can be realized in a number of models of correlated electron systems. In two dimensions, these states are yet another example of exotic states of matter which support anyonic excitations with fractional statistics. Mott insulators have also received considerable attention because they produce high-temperature superconductivity upon doping with electrons or holes. Recently, the possibility of a quantum spin liquid phase has been hotly debated for the case of the Hubbard model on the honeycomb lattice [3–5]. Motivated by the discovery of topological insulators, there have been numerous studies of the honeycomb Hubbard model with additional spin-orbit coupling, see [6] for a review. The Mott transition of fermions on the honeycomb lattice has recently been realized experimentally using cold atoms in an optical lattice [7].

The Numerical Challenge

For quantum many-body problems, the Hilbert space grows exponentially with the number of electrons, or the volume of the system. A brute force diagonalization of the full problem hence requires an exponential effort. The question is whether or not we have to consider all possible states to understand the problem. Here, we choose the stochastic approach, and use the auxiliary-field quantum Monte Carlo method [8]. The method is based on a path integral formulation where the interaction is decomposed with the help of a Hubbard-Stratonovitch field that mediates the electronic correlations. The integration over the field configurations is carried out with the Monte Carlo method. A generic problem of this approach is the sign problem (the appearance of configurations with negative weights), causing an exponential increase of computer time with volume and inverse temperature to reach a given accuracy. However, for a rather large class of non-trivial models [9–12], symmetries permit us to avoid the sign problem. For such models, the numerical effort scales as $N^3 \beta$ where β is the inverse temperature and N the number of electrons. The method used here can be efficiently parallelized and run on modern architectures. Recently, various optimization schemes, in particular cache optimization, have allowed us to gain up to an order of magnitude in performance.

The Kane-Mele-Hubbard Model

Motivated by the experimental progress with graphene and related systems (such as silicene and artificial graphene),

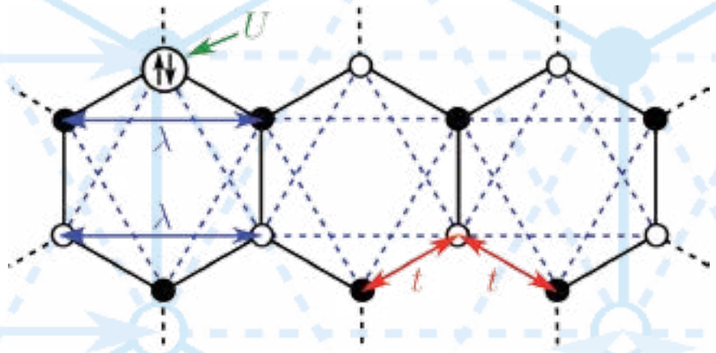

Figure 1: The Kane-Mele-Hubbard model describes electrons on the honeycomb lattice with nearest-neighbour hopping t, spin-orbit coupling λ, and an onsite repulsion U between electrons occupying the same lattice site.

we consider the paradigmatic Hubbard model for strongly correlated electrons on the honeycomb lattice [3] illustrated in Fig. 1. Electrons can hop from site to site with hopping amplitude t, and two electrons at the same lattice site experience a repulsion U that captures the essence of Coulomb repulsion. This model has time reversal, SU(2) spin, and sublattice symmetry, and the noninteracting band structure is that of massless Dirac fermions. Adding a spin-orbit term in the form of a complex second-neighbour hopping λ [1] leads to the Kane-Mele-Hubbard model for correlated topological insulators [13]. The spin-orbit term breaks sublattice symmetry and thereby opens a topological mass gap resulting in a quantum spin Hall insulator (a widely used name for two-dimensional topological insulators). At half filling (one electron per lattice site), the above symmetries permit us to apply the auxiliary-field quantum Monte Carlo method without a sign problem, and hence to investigate the role of electronic correlations accurately and without uncontrolled approximations.

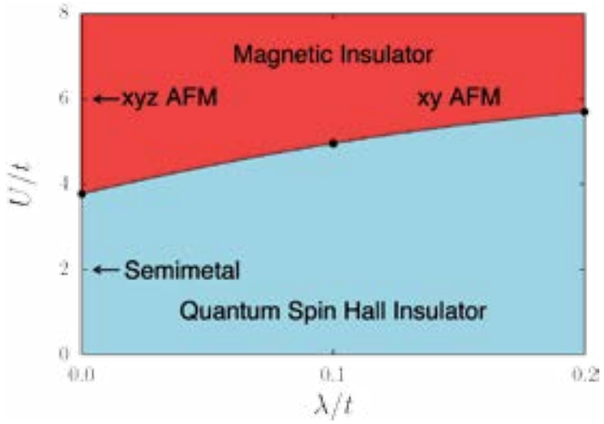

Figure 2: The zero-temperature phase diagram of the Kane-Mele-Hubbard model as obtained from quantum Monte Carlo simulations with up to 648 electrons [5, 11]. Here AFM indicates an antiferromagnetic Mott insulator with magnetic order either in the xyz or in the xy directions, respectively. The xyz AFM and semimetal phases exist only at λ = 0.

Results

In the limits of weak and strong coupling, the low-energy physics of the Kane-Mele-Hubbard model can be understood quite easily. The semimetallic state of Dirac fermions at λ = U = 0 is stable with respect to correlations and hence expected to survive at small U > 0. Similarly, at U = 0, the spin-orbit coupling establishes a gapped quantum spin Hall state, and the bulk physics remains unaffected by small interactions U. For strong coupling U/t » 1, the charge degrees of freedom become frozen (any deviations from the state with one electron per site are suppressed by U) but the spin degrees of freedom remain active. Using perturbation theory in the small parameter t/U, one can derive effective spin models. In the absence of spin orbit coupling, the SU(2) spin symmetry leads to a Heisenberg model with antiferromagnetic nearest-neighbour exchange. Since the lattice is non-frustrated, we expect long-range antiferromagnetic order at zero temperature, and a gapless Goldstone mode corresponding to spin-wave excitations. This antiferromagnetic Mott insulator (AFM) state is labelled xyz AFM in Fig. 2. Spin-orbit coupling reduces the spin symmetry from SU(2) to U(1). The resulting spin model is more complex, and includes frustrated interactions in the z direction of spin [13]. Long-range magnetic order can develop in the transverse direction, as is the case in the xy AFM phase in Fig. 2.

The validity of the above considerations can be verified using numerical simulations. Importantly, such methods also permit to study the intermediate-coupling regime and hence the evolution from small to large U/t. It is of particular interest whether or not there exists a direct transition between the weak- and strong-coupling phases, or if instead exotic intermediate phases appear. Moreover, if continuous quantum phase transitions occur, it is of great interest to identify their universality class. In the following, we summarize our present understanding of correlation effects in the Kane-Mele-Hubbard model, as it emerges from large-scale quantum Monte Carlo simulations [3,5,10,11,14].

Impact of U for λ = 0

The following two scenarios can be envisaged for the transition from the semimetal at weak coupling to the antiferromagnetic Mott insulator at strong coupling. Starting from strong coupling, and noting that the insulator to semimetal transition is numerically found to occur at values of U smaller than the electronic bandwidth W = 3t, one can construct an effective spin model for the region close to the transition where higher-order ring-exchange terms proliferate with decreasing U/t [17]. This point of view suggests that the melting of magnetic order may be independent of the metal-insulator transition. Quantum Monte Carlo simulations [3] suggested the existence of an intermediary spin liquid phase, with a single-particle gap but no long-range magnetic order, that separates the semimetal from the magnetic insulator. Similar conclusions have been drawn for the related π-flux model on the square lattice [18]. The results of [3] have been challenged by more recent studies: Entropy calculations do not favour a degenerate ground state as expected for the Z_2 spin liquid [19], and [4] demonstrates that the use of significantly larger system sizes leads to an almost complete disappearance of the purported spin-liquid from the phase diagram.

Alternatively, we can start from the weak-coupling semimetalic state. Upon increasing U/t, we eventually expect a phase transition to an insulating state with antiferromagnetic order [15,20, 21]. Because gapless, linear fermionic excitations exist below the transition, the transition is expected to fall into the Gross-Neveu universality class [15,16]. This expectation is based on Gross-Neveu-Yukawa theory [5,15,16], which describes massless Dirac fermions with a Yukawa coupling to a vector boson describing the bosonic Goldstone mode of the antiferromagnetic state. The fermionic mass (the gap of the insulating state) is generated by the condensation of the bosonic mode (antiferromagnetic order). The theory has an upper critical dimension of d = 3 where the Gaussian approximation becomes exact. Within a first-order ϵ-expansion, one can estimate the critical exponents for the present two-dimensional case. These theoretical predictions are confirmed by recent quantum Monte Carlo simulations [5]. Using the predicted critical exponents, we observe an excellent data collapse [see

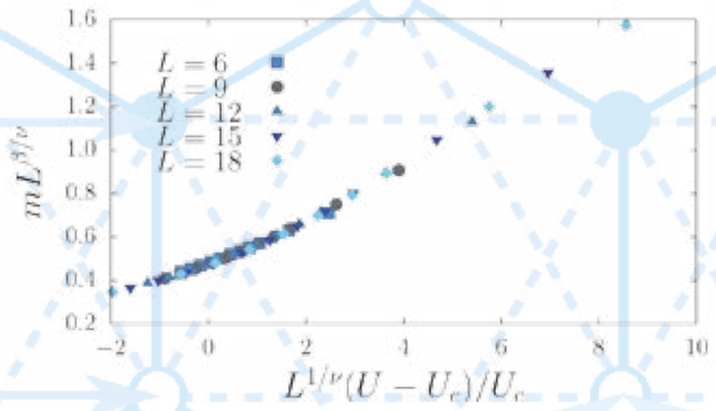

Figure 3: Data collapse of the staggered magnetic moment m at λ = 0, measured for different system sizes L. Here we have used a critical value U_c/t = 3.78 [5], and critical exponents obtained from an ε-expansion and Gross-Neveu-Yukawa theory [15,16]. Results taken from [5].

Fig. 3] and an identical finite-size scaling of both the single-particle gap and the staggered magnetization [5]. Importantly, Gross-Neveu theory predicts that the long-ranged tail of the Coulomb repulsion, relevant for possible future experiments, will not alter the nature of the transition. Numerical work to confirm this hypothesis is currently in progress.

Impact of U for $\lambda > 0$

The transition from the quantum spin Hall phase to the xy antiferromagnetic Mott insulator is an insulator to insulator transition. Because the onset of magnetic order breaks the time-reversal symmetry that protects the quantum spin Hall state, a direct transition is possible. Since the system is insulating on both sides of the transition, the latter involves only the spin degrees of freedom which order in the xy plane at the critical point. Such a transition, involving a U(1) order parameter in two spatial dimensions, is expected to fall into the same universality class as the three-dimensional XY model. Using the quantum Monte Carlo method, we have successfully verified this prediction in [11].

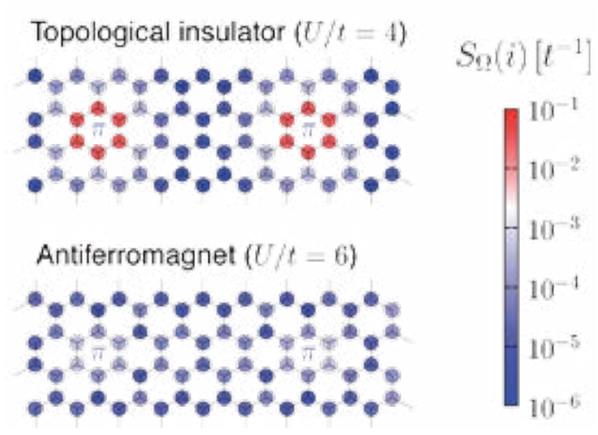

Figure 4: Integrated spectral weight of the spin fluxon states created by inserting a pair of π fluxes into the honeycomb lattice. A response occurs in the correlated topological state [see Fig. 2], but not in the topologically trivial antiferromagnetic insulator. Here, $\lambda = 0.2t$. Results taken from [14].

An important question in the context of topological insulators is how a correlated topological state can be detected numerically [6]. One attractive possibility, demonstrated in [14], is to exploit a unique property of quantum spin Hall insulators, namely the quantization of the spin Hall conductivity [1]. As a consequence, the insertion of a magnetic flux of the size of half a flux quantum (a π flux) gives rise to two types of states localized around the flux: A so-called Kramers pair of spin fluxons carrying spin ±1/2, and a pair of charge fluxons carrying charge ±e [22, 23]. The existence of these states and their quantum numbers are directly related to the topological invariant [22, 23]. Electronic interactions remove charge fluxons from the low-energy excitation spectrum of the system, leaving spin fluxons which can be detected via a Curie-law-type contribution to the bulk magnetic susceptibility [14]. Alternatively, as shown in Fig. 4, it is possible to measure the spectral weight of spin fluxon excitations in real space. The combination of π fluxes and quantum Monte Carlo simulations provides an efficient tool to identify correlated quantum spin Hall states.

Outlook

Quantum Monte Carlo simulations of the Hubbard model with additional spin-orbit coupling allow us to study the interplay between band topology and correlation effects. Our results on finite lattices are free from systematic errors and therefore also serve as benchmarks for approximate theories. The continuous algorithmic progress enables us to address exciting open questions, such as the effect of a long-range Coulomb interaction on the phase diagram and the phase transitions.

Acknowledgments

Parts of this work were done in collaboration with M. Bercx, F. Goth, I. Herbut, T. Lang, Z.Y. Meng, A. Muramatsu and S. Wessel. Funding from the DFG under the grant number AS 120/9-1, AS120/10-1, and Ho 4489/2-1 (Forschergruppe FOR 1807) is acknowledged. We thank the Jülich Supercomputing Centre for generous allocation of CPU time.

References

[1] **Kane, C.L., Mele, E.J.**
Phys. Rev. Lett. 95, 146802, 2005

[2] **König, M., Wiedmann, S., Brüne, C., Roth, A., Buhmann, H., Molenkamp, L.W., Qi, X.-L., Zhang, S.-C.**
Science 318, 766, 2007

[3] **Meng, Z.Y., Lang, T.C., Wessel, S., Assaad, F.F., Muramatsu, A.**
Nature 464, 847, 2010

[4] **Sorella, S., Otsuka, Y., Yunoki, S.**
Sci. Rep. 2, 992, 2012

[5] **Assaad, F.F., Herbut, I.F.**
Phys. Rev. X 3, 031010, 2013

[6] **Hohenadler, M., Assaad, F.F.**
J. Phys.: Condens. Matter 25, 143201, 2013

[7] **Uehlinger, T., Jotzu, G., Messer, M., Greif, D., Hofstetter, W., Bissbort, U., Esslinger, T.**
arXiv:1308.4401

[8] **Assaad, F.F., Evertz, H.G.**
Computational Many Particle Physics, Lecture Notes in Physics, Vol. 739, edited by H. Fehske, R. Schneider, and A. Weiße (Springer Verlag, Berlin, 2008) pp. 277–356

[9] **Assaad, F.F.**
Phys. Rev. Lett. 83, 796 (1999)

[10] **Hohenadler, M., Lang, T.C., Assaad, F.F.**
Phys. Rev. Lett. 106, 100403 (2011)

[11] **Hohenadler, M., Meng, Z. Y., Lang, T. C., Wessel, S., Muramatsu, A., Assaad, F. F.**
Phys. Rev. B 85, 115132, 2012

[12] **Berg, E., Metlitski, M.A., Sachdev, S.**
Science 338, 1606, 2012

[13] **Rachel, S., Le Hur, K.**
Phys. Rev. B 82, 075106, 2010

[14] **Assaad, F.F., Bercx, M., Hohenadler, M.**
Phys. Rev. X 3, 011015, 2013

[15] **Herbut, I.F.**
Phys. Rev. Lett. 97, 146401, 2006

[16] **Herbut, I.F., Juričić V., Vafek, O.**
Phys. Rev. B 80, 075432, 2009

[17] **Yang, H.-Y., Albuquerque, A.F., Capponi, S., Läuchli, A.M., Schmidt, K.P.**
New Journal of Physics 14, 115027, 2012

[18] **Chang, C.-C., Scalettar, R.T.**
Phys. Rev. Lett. 109, 026404, 2012

[19] **Clark, B.K.**
arXiv:1305.0278

[20] **Sorella, S., Tosatti, E.**
Europhys. Lett. 19, 699, 1992

[21] **Paiva, T., Scalettar, R.T., Zheng, W., Singh, R.R.P., Oitmaa, J.**
Phys. Rev. B 72, 085123, 2005

[22] **Qi, X.-L., Zhang, S.-C.**
Phys. Rev. Lett. 101, 086802, 2008

[23] **Ran, Y., Vishwanath, A., Lee, D.-H.**
Phys. Rev. Lett. 101, 086801, 2008

• Fakher F. Assaad
• Martin Hohenadler

Institut für
Theoretische Physik
und Astrophysik
Universität Würzburg

Highly-resolved numerical Simulations of bed-load Transport in a turbulent open Channel Flow

Prediction of turbulence-induced erosion and near-bed transport of sediment particles in turbulent flow is important for many processes in environmental engineering. Beyond its relation to sediment transport, the results of the present study are relevant as well for numerous industrial applications, particularly in the field of process technology, where solid particles are conveyed by a carrier flow. Traditional methods for the prediction of sediment transport are empirical and based on averaged bulk quantities. The predictive power of these formulae is low, because homogeneity of the sediment is postulated. A detailed understanding of sediment stability and the physical mechanisms involved in sediment transport is still missing due to the lack of highly-resolved data under controlled flow conditions. The presented study employs Direct Numerical Simulations (DNS) of a flow laden with a large number of particles with parameters of the disperse phase chosen similar to laboratory experiments [1, 2]. The highly-resolved simulations provide detailed and physically reliable instantaneous information on bed-load transport at medium Reynolds number, covering parameter ranges so far not reached.

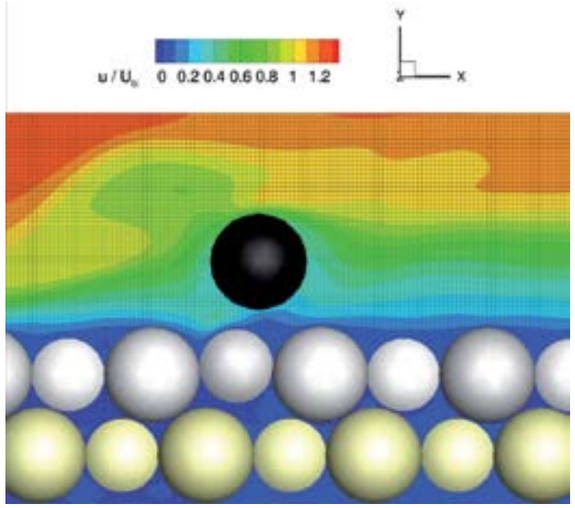

Figure 1: Particle position, background Eulerian grid and fluid velocity in the streamwise direction in a plane through the particle centre with 22 grid points per diameter.

Numerical Method

The study was conducted with the in-house code PRIME (Phase Resolving slMulation Environment) using an Euler-Lagrange approach with individual particles being geometrically resolved. The fluid is described by the unsteady three-dimensional Navier-Stokes equations discretized on a regular Cartesian grid. For each particle, the equations of motion are solved in terms of translational and angular velocity. The coupling between both phases is accomplished by an immersed boundary method [3]. For particle contact, the Adaptive Collision Model (ACM) was used [4]. It accounts for all relevant mechanisms that have to be modelled during the collision process and was validated in great detail for single and multiple simultaneous collisions.

Computational Setup

An open-channel flow is considered with a computational domain $24H \times (H+H_{sed}) \times 12H$ in streamwise, vertical, and spanwise direction, respectively, with H the channel height and H_{sed} the sediment height at the top of two layers of $N_{p,fixed}$ = 27000 mono-disperse spheres, which are arranged in a hexagonal pattern [5]. Periodic boundary conditions are applied in streamwise and spanwise direction and a free-slip condition at the top. The bulk Reynolds number based on the channel height and the bulk velocity of the fluid is 2941, i.e. slightly higher than the threshold for turbulent flow of an unladen channel. The resulting friction Reynolds number for an unladen flow (Case Fix) is Re_τ = 193 and the particle Reynolds number in wall units is D^+ = 21. The resolution of the equidistant, Cartesian grid is set to $D/\Delta x$ = 22.2 thus guaranteeing proper resolution of the viscous effects (Fig. 1). This results in a total amount of $1.4 \cdot 10^9$ grid cells, which is to the knowledge of the authors the largest grid employed so far for this kind of problem. In the reference run (case Ref) the upper layer of the sediment bed was released on top of the remaining sediment bed held fixed, which gives a total of $N_{p,mobile}$ = 13500 particles. The mobility is slightly above the threshold of initiation of motion [6]. Subsequently, additional simulations with lower mass loading (case FewPart) and lower mobility (case LowSh) were carried out to elucidate the effect of these two key parameters on bed-load transport. After initialization, the simulations were run until an equilibrium between erosion
and deposition was obtained [5]. For each simulation, statistical data was gathered for more than 240 bulk units.

Results

The reference run (Fig. 2) produces two dune-like particle clusters with their major axis in spanwise direction. The clusters have an average distance of about 12 H and travel on the surface of a layer of inactive particles. The latter are particles, which are free to move but have come to rest temporarily. The layer of inactive particles shows an increased porosity, because more than 26% of the mobile particles are moving. This allows a significant flow at this level. On top of the inactive layer, the dunes form a bed-load layer with a thickness of more than 2 D. This enhances the turbulent fluctuations in this regions substantially.

Reducing the mass loading results in inactive particle structures oriented in streamwise direction, so-called ridges (case FewPart Fig. 3). These enhance coherent vortex structures in the fluid oriented in streamwise direction. Furthermore, they generate a mean secondary flow in the troughs with

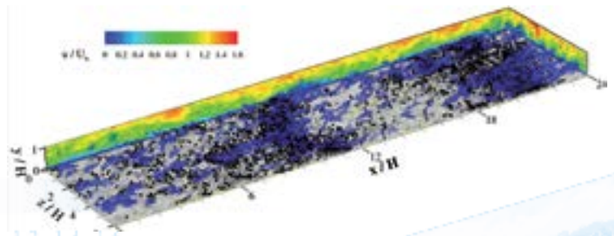

Figure 2: Instantaneous particle distribution of the reference run (Ref). Isosurfaces of fluid fluctuations blue: u'/U_b = –0.3, particles in yellow: fixed, white: $|u_p| < 1.5\ u_\tau$, black: $|u_p| \geq 1.5\ u_\tau$

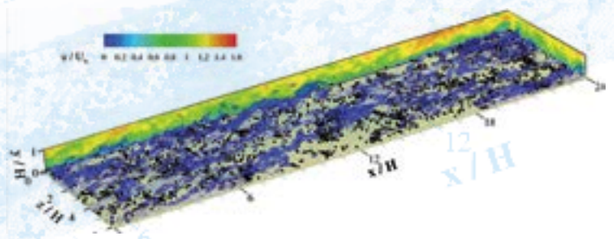

Figure 3: Same as Fig. 2, but for case FewPart.

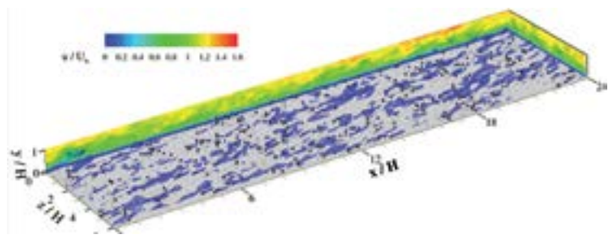

Figure 4: Same as Fig. 2, but for case LowSh.

component perpendicular to the main flow direction in turn transporting particles towards the ridges. On the other hand, fast travelling particles move in the troughs and occasionally erode the ridges at different locations. As the fixed bed is not completely covered by inactive particles, spatial heterogeneity in the spanwise direction develops that allows a significant flow with a high turbulent intensity in the sediment region. The relative amount of moving particles is reduced in comparison to Ref, as the inactive ridges are self-stabilising through hiding and shading mechanisms induced by the spatial heterogeneity. The ridges extend 12 H on average in the streamwise direction and move very slowly resulting in the separation of the fluid time scales and of the morphological time scales of the particle clusters.

In case LowSh, almost all mobile particles settle onto the fixed particles due to their increased submerged density and form an almost closed plane bed (Fig. 4). Only a small percentage of particles is eroded. The lack of hiding and shading mechanisms leads to higher translational particle velocities. In this dilute regime, the particles only have small impact on the fluid flow. Hence, the turbulent structures resemble the large-scale behaviour observed for unladen flows over rough walls with high submergence [7].

The interaction of the disperse and the continuous phase was analyzed by suitable statistical measures, such as the two-point correlation for the streamwise velocity component of the fluid in the near-wall region $R_{uu}(r_x)$, with r_x the correlation length, and the Cartesian Distribution Function (CDF) of particle pairs $G(\xi_x, \xi_z)$ [8,9] depicted in Fig. 5. Particle structures were assessed using the probability of a particle to find another particle at a given distance ξ_x, ξ_z in x- and z-direction respectively. A value of one reflects a random distribution, while values larger than one reveal particle clustering. The CDF was conditioned by the particle velocity with G^m only accounting for pairs of moving particles (Fig. 6) and G^r only pairs of two resting particles. This measure allows to distinguish between the dune-like and small scale clusters on the one hand, and the closed plane bed and the ridges on the other. To compare the two-point correlation of the fluid to the CDF of particles pairs, the values the interval 0.5 D < ξ_x < 12 H and ξ_z = 0 are shown here only. This allows to determine the relevant length scales for the two phases and to correlate them. Note that both functions $R_{uu}(r_x)$ and $G(\xi_x, \xi_z)$ are symmetric for negative values of r_x and ξ_x for large datasets.

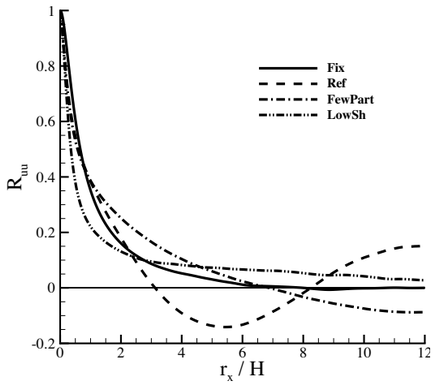

Figure 5: Two-point correlation function of the fluid in streamwise direction at y = 0.5 D. The thin horizontal line represents the value of decorrelated fluid.

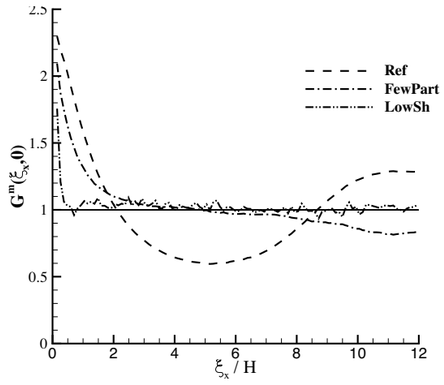

Figure 6: CDF in the streamwise direction of moving particles. The thin horizontal line represent the value of a random distribution.

The characteristic distance of the dune-like structures of case Ref becomes obvious for both the particle and the fluid statistics in Fig. 5 and 6. This suggests a strong interaction between the two phases. The phase-resolving approach of the present numerical scheme forces the fluid to flow around particle clusters and to adopt to their length scales. The same can be observed for the inactive ridges reported in case FewPart. Since the characteristic length of a ridge is up to 12 H, only one ridge fits into the computational domain driving R_{uu} below zero for r_x > 6 H. The regime of case LowSh is very dilute as only a small percentage of the particles are actually moving across the closed bed. The moving particles are randomly distributed in space and time and thus do not have the potential to alter the flow field substantially.

Conclusions

Highly-resolved simulations of three particle-laden flows across an idealized sediment bed elucidate the impact of mass loading and mobility on bed-load transport. The resulting interaction of the two phases was analyzed by suitable statistical tools to distinguish between the length scales of eroded and inactive clusters. The strong similarity of the statistical measures extracted from the two phases reveal a strong interaction. This suggest practical implications for the structural design of sediment-laden open channels, as the mobile bed ultimately alters the hydraulic resistance of the channel by introducing spatial heterogeneity that can occur in streamwise direction (dune-like clusters) or spanwise direction (ridges). It also illustrates the limits of the point-particle approach, where particles are treated as mass points without spatial extension and their motion is modelled by empirical correlations.

References

[1] **Dietrich, W.E., Kirchner, J.W., Ikeda, H., Iseya, F.**
Sediment supply and the development of the coarse surface-layer in gravel-bedded rivers. Nature, 340 (6230), 215-217, 1989

[2] **Shvidchenko, A.B., Pender, G.**
Macroturbulent structure of open-channel flow over gravel beds. Water Resour. Res., 37(3): 709–719, 2001

[3] **Kempe, T., Fröhlich, J.,**
An improved immersed boundary method with direct forcing for the simulation of particle laden flows. J. Comput. Phys. 231, 3663-3684, 2012

[4] **Kempe, T., Fröhlich, J.,**
Collision modeling for the interface-resolved simulation of spherical particles in viscous fluids. J. Fluid Mech. 709, 445–489, 2012

[5] **Vowinckel, B., Kempe, T., Fröhlich, J.,**
Particle-resolving simulations of bed load sediment transport. Paper 792 at 8th ICMF, Jeju, Korea, 2013 (electronical version only)

[6] **Shields, A.**
Anwendung der Ähnlichkeitsmechanik und der Turbulenzforschung auf die Geschiebebewegung. Ph.D. thesis, Mitteilungen der Preußischen Versuchsanstalt für Wasserbau und Schiffbau, Berlin, 1936

[7] **Yalin, M.S., Ferreira da Silva, A.M.,**
Fluvial Processes. IAHR/AIRH, Monograph, 2001

[8] **Vowinckel, B., Kempe, T., Fröhlich, J., Nikora, V.I.**
Direct numerical simulation of bed load transport fo finite-size spherical particles in a turbulent channel flow. To appear in proceedings of DLES 9, Dresden, Germany, 2013

[9] **Vowinckel, B., Kempe, T., Fröhlich, J.**
Phase-resolving simulations of bed load transport. In THESIS-2013 symposium, Chatou, France, 2013

- Bernhard Vowinckel
- Tobias Kempe
- Jochen Fröhlich

Institut für Strömungsmechanik, TU Dresden

How to fit the Local Universe into a Supercomputer?

In 1965 Arno Penzias and Robert Wilson detected the cosmic microwave background radiation (CMB). More than 13 billion years ago this radiation was imprinted on the sky, only a few 100,000 years after the Big Bang. In 1992 the COBE satellite detected anisotropies in the temperature of the CMB radiation. Meanwhile these temperature fluctuations are measured with very high precision by satellites (WMAP, Planck) as well as many ground based observations. The measured temperature fluctuations tell us that shortly after the Big Bang the Universe was almost homogeneous with tiny density fluctuations of the order of 10^{-5}. Comparing the power spectrum of measured temperature fluctuations with theoretical models, cosmologists conclude that the Universe is spatially flat and consists at present of about 68% of some unknown Dark Energy, 27% of also unknown Dark Matter and 5% of baryons.

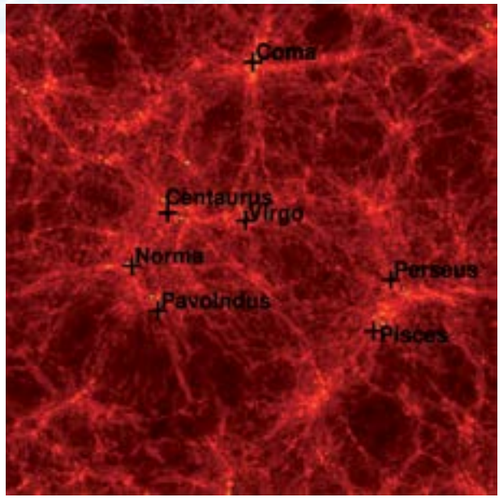

Figure 1: A slice of thickness 110 million light years through the simulation box. Some very prominent structures lie within that plane such as the galaxy clusters indicated with name. The simulated and observed positions of the clusters match very well.

In the evolved universe one can directly observe the distribution of baryons and indirectly deduce (gravitational lensing, velocity measurements) the distribution of Dark Matter. We see huge clusters of galaxies with masses up to a few 10^{15} solar masses in the knots of the cosmic web which comprises galaxies in a wide range of masses from tiny dwarf galaxies (10^{9} solar masses) to massive elliptical galaxies (10^{13} solar masses). All these structures have evolved from tiny fluctuations generated during the early inflationary phase and measured in the CMB background.

The formation of structure on large scales is well understood within the concordance model of cosmology. The initial small perturbations grow by gravitational instability and form bound objects called halos which decouple from the expansion of the universe. These bound objects grow further hierarchically by accretion of matter and merging with smaller halos. The gravitational clustering becomes increasingly nonlinear. Dark Matter is more abundant and hence most important for formation of large scale structures where gravity dominates. On smaller (galactic) scales baryons play an important role. They interact not only gravitationally but form a gas with a certain pressure and temperature. In the gravitational potential wells of the Dark Matter halos the originally hot gas cools via radiative cooling and finally stars are formed of the cooled gas. Almost all chemical elements above helium are formed in stars and then redistributed into the cosmic medium. Star formation and the feedback of the stars on the gas are the most important processes in galaxy formation. Nevertheless due to the domination of gravity at early times galaxies are closely associated with the Dark Matter halos.

The vastness of scales and the nonlinearity of gravitational clustering are the reasons why numerical simulations and the intensive use of the largest supercomputers are the only methods suited to study the gravitationally driven growth of structures down to local over-densities in halos and the formation of galaxies therein. Cosmological simulations follow the clustering of matter by numerically solving the gravitational interaction based on an N-body approach. Additionally one needs to model hydrodynamical processes, radiative cooling, star formation and the feedback of stars, in order to simulate in detail the formation of galaxies in their different environments. To this end very high mass- and spatial-resolution are necessary which imposes a strong challenge for present day computational algorithms. One important limitation of state of the art galaxy formation simulations is that galaxies are usually selected to be relatively isolated systems, to avoid extra computational costs. However it has been shown that the environment of galaxies plays an important role in their evolution. Therefore simulations of isolated galaxies can only offer limited insights into the more complex process of galaxy evolution.

One way to overcome this selection bias and to draw general conclusions about galaxy evolution is to simulate and observe galaxies in all possible environments and compare ensemble averaged observables. However this marginalization approach would require a computationally infeasible number of high-resolution simulations.

An alternative way to investigate galaxy evolution is to compare simulated and observed galaxies that reside in a similar environment. This has been the scope of the CLUES (Constrained Local UniversE Simulations) project (http://www.clues-project.org), and it is also the goal of our work - to provide initial conditions that reproduce the local environment of the Milky Way and its neighbours. The Local Universe is naturally the best observed region of the Universe where even the tiniest dwarf galaxies can be observed. Therefore, it is best suited to study the formation of structures in controlled computer experiments. To this end we need to fit the observed Local Universe into a (super)computer. In the next section we describe our way of doing this.

Recovering and Simulating Structures of the Local Universe

The initial conditions for cosmological simulations are formulated in agreement with the observed CMB power spectrum. Running the simulation over more than 13 billion years one obtains a realisation of the universe which represents the observed Universe in a statistical way, but it has nothing to do with the observed nearby universe. To simulate the formation of the observed nearby universe one has to recover the initial phase distribution of density perturbations which is responsible for the formation of the local structures. To this end observations of the nearby universe must be imposed as constraints on the initial conditions of the simulations. The resulting constrained simulations serve as a numerical laboratory of the nearby universe.

Recovering the initial phase distribution of density perturbations is challenging for several reasons. First, today's large scale structures are usually observed with galaxy surveys. Galaxies however are biased tracers of the underlying density field - like the tips of icebergs. Furthermore their receding velocities due to the expansion of the universe and their peculiar velocities (the velocity with respect to the local rest frame) are difficult to disentangle observationally. Second, the present day density field is connected to the initial density field by non-linear structure formation. Especially on small scales this makes it extremely difficult to trace it back in time.

Hence we apply an iterative Markov Chain Monte Carlo approach that approximates the evolution of the universe forward in time every iteration step. Using a combined approach consisting of this iterative reconstruction based on non-linear Lagrangian Perturbation Theory and a fully non-linear N-body simulation we are able to reconstruct the initial conditions and simulate the dark matter density in the Local Universe. The aim is to derive the density field that resembles the Local Universe as it is being observed by the 2MRS galaxy survey (www.cfa.harvard.edu/~dfabricant/huchra/2mass/) up to distances of 418 million light-years.

To this end we used a new algorithm called Augmented Lagrangian Perturbation Theory (ALPT) to model structure formation in each iteration step. ALPT complements the accuracy of Lagrangian perturbation theory on large scales with the local spherical collapse model which is more precise on small scales. This approach allows to reconstruct the initial density perturbations 13 billion years ago which are responsible for the formation of the large scale structures that are observed today. The quality of the recovered initial density field which contains the seeds of today observed structures can be assessed by direct comparison of the simulated and observed structures. To this end also the line-of-sight distortions are modelled that are inherent to modern galaxy surveys. Then the simulated structures can be directly compared to the observed positions of the 31,017 galaxies of the 2MRS that lie within the desired volume. With this comparison the sampling of the initial Gaussian density field for the next iteration step can be refined. Therefore we iteratively recover an initial density field that seeds structures as observed by the 2MRS survey. This iterative reconstruction algorithm is fast and accurate.

We then use the reconstructed initial field of density perturbations to simulate the full non-linear formation of structure using the N-body code Gadget3 (by V. Springel, http://www.mpa-garching.mpg.de/gadget/). This code follows the gravitational clustering of matter. Halos reside in the highest density peaks of the dark matter distribution and they host galaxies. Now we have a numerical realisation of the Local Universe in our computer and we can directly compare it with the real one. We have chosen a medium mass resolution of 453 million particles in the simulation box so that we achieve a mass resolution of 10 billion solar masses, i.e. a Milky Way sized galaxy is represented in the simulation by 100 particles. With this resolution we can identify galaxies in the simulation box down to the faintest ones of the observational sample in a reliable way. Our simulations resembles the Local Universe very well as we demonstrate in the two figures.

In the first figures we show the resulting density distribution in a slice of 110 million light years thickness through the center of the simulation box. This slice corresponds to the Super-galactic plane. The color coding follows the logarithm of the density field. The largest halos (with bound masses of more than 10^{14} solar masses) which presumable host galaxy clusters are marked by circles. Crosses mark the position of some nearby galaxy clusters which have been extracted from the NED database (http://ned.ipac.caltech.edu/). Our own galaxy - the Milky Way - is in the center of the box in a distance of about 50 million light years to the Virgo cluster. The second figure shows an all-sky projection of all galaxies observed in a distance of 170 to 280 million light-years from the Milky Way taken from the 2MRS survey compared to the matter density in the same shell of our simulation.

The constrained simulations of the Local Universe serve as laboratory for studying the evolution of the local density distribution, the non-linear gravitational potential and the velocity field. We can study the formation of the local cosmic web as well as the formation histories of the simulated counterparts of observed local galaxy clusters. Based on our reconstructed initial density field we are also able to re-simulate specific selected regions with higher resolution. In these so called zoomed simulations we can then re-simulate galaxies like the ones of our Local Group resolved with millions of particles including hydrodynamics and models for gas cooling, formation of stars and feedback. Being able to do so we can search for traces of the evolution in the simulated and observed galaxies and the influence of the environment on their evolution.

Figure 2: The top panel shows the sky projection of all galaxies in the 2MRS catalog at distances of 170 to 280 million light-years. The bottom panel shows the matter density in the same shell of the N-body simulations. Very prominent structures are in the left the Perseus-Pisces Supercluster connected to the lower left with the Cetus Wall and near the equator slightly to the right the Norma cluster and the Great Attractor.

Acknowledgements

The authors acknowledge the support by the Deutsche Forschungsgemeinschaft under the grant GO563/21-1. The simulations have been performed at JSC.

References

[1] **Heß, S., Kitaura, F.-S., Gottlöber, S.** Simulating Structure Formation of the Local Universe, ArXiv 1304.6565

[2] **Kitaura, F.-S.** The initial conditions of the Universe from constrained simulations, ArXiv 1203.4184

[3] **Kitaura, F.-S., Heß, S.** Cosmological Structure Formation with Augmented Lagrangian Perturbation Theory, ArXiv 1212.3514

• Steffen Heß
• Francisco Kitaura
• Stefan Gottlöber

Leibniz-Institut für Astrophysik Potsdam

A scalable hybrid DFT/PMM-MD Approach for accurately simulating Biomolecules on SuperMUC

Life sciences, from fundamental research to medical drug design, crave for detailed knowledge about the structural and dynamical properties of biomolecules. Ligand-receptor interactions, for instance, involve a complex interplay between molecules, which needs to be understood in order to selectively influence it, e.g. by medical drugs. Experiments can of course help to answer such questions, but access to the molecular world is still difficult and theoretical modeling inevitable.

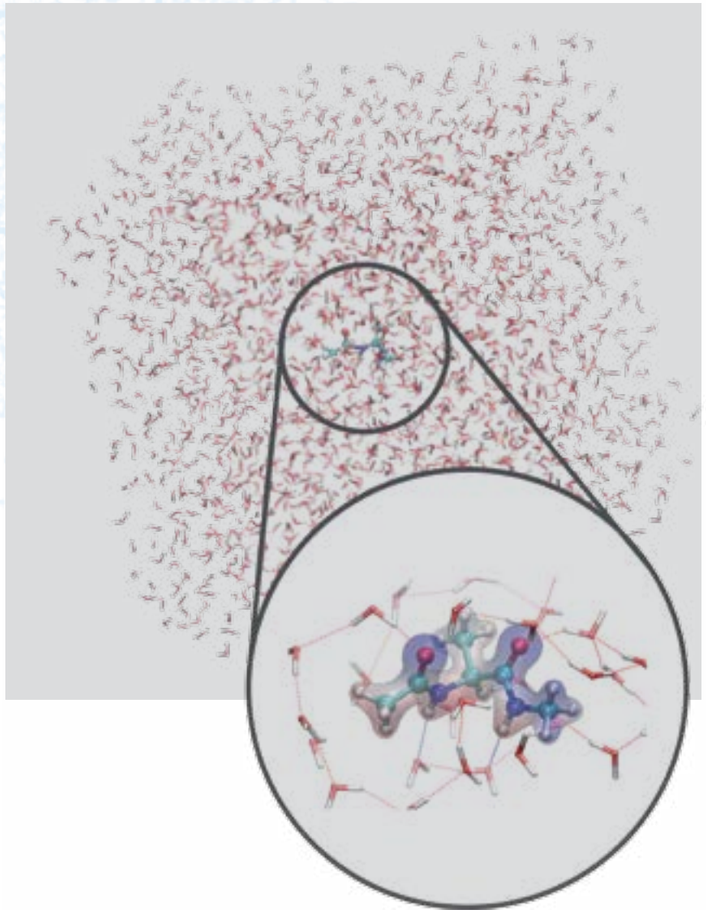

Figure 1: A typical QM/MM simulation setup consisting of an alanin-dipeptide molecule (described by QM) solvated in about 2,000 MM water molecules. The close-up shows the MM solvation shell of the solute molecule together with a sketch of its electron density, which is accessible only through the use of a QM method.

Infrared (IR) spectroscopy represents one of the few techniques to monitor the functional dynamics of biomolecules such as peptides and proteins. In order to reveal the structural changes of molecules encoded in such spectra, molecular dynamics (MD) simulations are a valuable tool that complements experimental results and helps to understand them. Molecular mechanics (MM) force fields enable MD simulations of large systems, such as a protein in solution comprising several ten thousand atoms, up to a microsecond time scale. However, such MM-MD simulations are by far not accurate enough for tasks like calculating IR spectra. In contrast, high-level quantum mechanical (QM) methods like density functional theory (DFT) provide the required accuracy, but are computationally limited to much smaller length and time scales. Thus, a sufficiently long simulation of even a small peptide in an extended aqueous environment is still beyond the scope of such QM methods.

These issues are resolved by hybrid MD approaches, which combine a QM treatment of a small subsystem with a MM description of its environment (cf. Fig. 1). This combination of an accurate, but rather slow QM description of e.g. a single (bio-)molecule with a more coarse, but efficient MM treatment of its surroundings enables MD simulations with good overall accuracy and on greater length and time scales than pure QM methods. QM/MM hybrid methods have become a standard tool in life sciences for applications like studying structure and energetics of enzyme reactions, excited-state properties or charge-transfer processes [1].

QM/MM approaches typically combine different programs for the MM and QM calculations, which need to be properly interfaced. Especially the efficient evaluation of the long-ranged electrostatic interactions between many thousand MM partial charges and the electrons treated explicitly in the QM part pose a major computational challenge.

Many current approaches are hampered by the neglect of the important polarization effects in the MM force field, and, therefore, do not yield the desired accuracy to, for example, calculate spectroscopic properties of biomolecules. Recently, a combination of a highly accurate DFT description with a polarizable MM (PMM) description of the condensed phase environment was developed in our group [2]. This new ap-

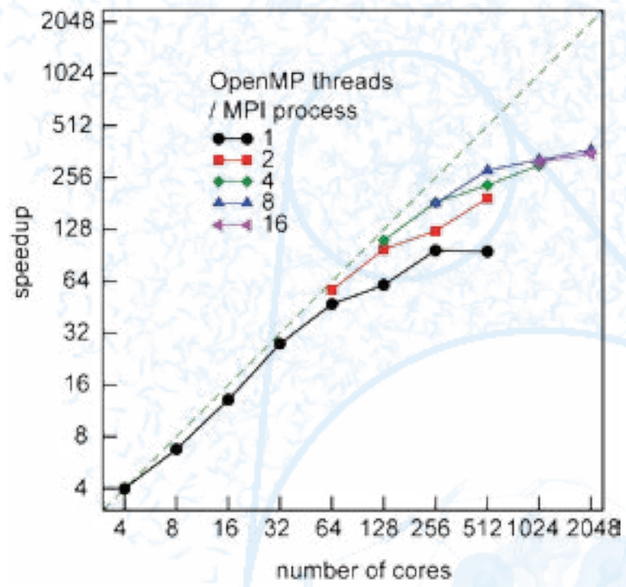

Figure 2: Scaling of IPHIGENIE/CPMD on SuperMUC for pure MPI (black) and MPI/OpenMP parallelization (colored).

proach greatly enhances the accuracy compared to conventional unpolarizable hybrid MD approaches. It treats the DFT/PMM electrostatics, which additionally covers interactions of the inducible PMM dipoles with the DFT electron density, by a linearly scaling fast multipole method (SAMM$_4$ [3]). The drastically reduced computational complexity makes studies of large DFT molecules solvated in accurately modeled extended PMM condensed phase now feasible.

We have implemented the above DFT/PMM approach by combining the well-established MPI/OpenMP-parallel DFT code CPMD [4] with our MPI/OpenMP-parallel PMM code IPHIGENIE [2, 3]. Our initial DFT/PMM implementation restricted IPHGENIE to run on a single core and used only CPMD in a parallel mode. Concomitantly the scaling of such calculations was poor, especially on HPC systems. Within a KONWIHR-III project, we completely revised the IPHIGENIE/CPMD program interface, resulting in a completely MPI/OpenMP-parallel, scalable program package. It links IPHIGENIE and CPMD together into a single executable and thereby enables direct data exchange, which is faster and more robust than the previously employed communication via file exchange. Furthermore, for each of the two program parts the respective optimal parallelization setup (number of MPI processes and OpenMP threads) can be used. Thus, this revised implementation has been carefully designed to run on HPC systems, and has been optimized for the SuperMUC petaflop system operated by the Leibniz Supercomputing Centre (LRZ). This system features 147,456 Intel Sandy Bridge cores connected by an ultrafast low-latency Infiniband network, and provides an excellent platform for our targeted large-scale DFT/PMM simulations of biomolecular systems with IPHIGENIE/CPMD.

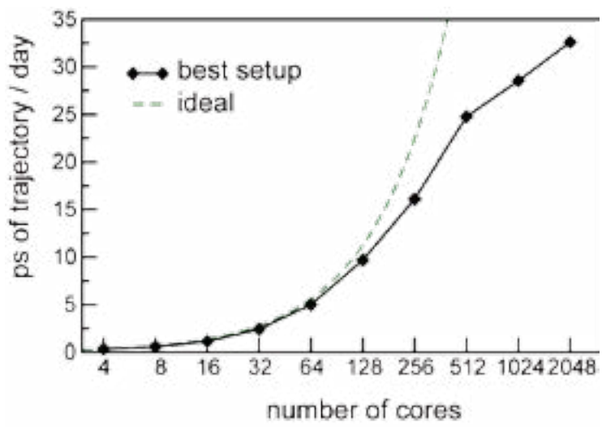

Figure 3: Absolute speed of DFT/PMM calculation involving 22 DFT atoms and 2,112 PMM waters (integration time step 0.5 fs).

We tested the strong scaling properties of our implementation on SuperMUC with a comparably small setup comprising an alanin-dipeptide molecule (22 atoms) as DFT fragment in a periodic box filled with 2,112 polarizable water molecules (PMM model TL4P), resulting in a total atom count of 8,470. Fig. 2 demonstrates that even for this small setting, an almost perfect scaling is achieved up to 512 cores, with still reasonable performance gain up to 2,048 cores. In this setup, the IPHIGENIE part scales up to 128 MPI processes, corresponding to as few as 66 atoms per MPI process. Less than 8% of the total computation time is spent on the PMM part, and the well-known excellent scaling of CPMD is not hampered by the interface to IPHIGENIE. Several tens of picoseconds of DFT/PMM trajectory per day are now feasible (cf. Fig. 3) and we are currently setting up large-scale DFT/PMM simulations of relevant biomolecules such as DNA bases on SuperMUC.

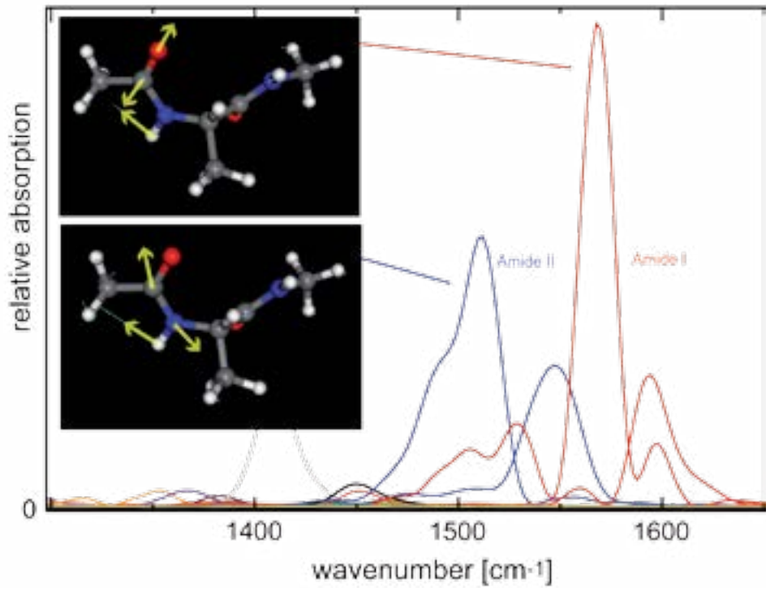

Figure 4: IR spectrum of alanin-dipeptide (treated with MT/BLYP, E_{cut}=80 Ry) solvated in PMM water. The insets show the molecular movement associated with the absorption bands.

As a sample application, we have computed two 20 ps DFT/PMM trajectories of the above system from which we calculated the IR spectrum [5] of the alanine-dipeptide in water. Fig. 4 shows the relative absorption and reveals the characteristic frequency bands Amide I and II of the left and right side of the molecule, whose positions depend on the conformation of the molecule and can thus serve to identify its current structure. As the two insets show, full insight is gained into the underlying dynamics and structure. Averaging a few of such spectra, which are rapidly calculated with our DFT/PMM approach on SuperMUC, yields accurate IR spectra of the molecule in the condensed phase helping to interpret experimental results.

Concerning future development, we are currently testing a DFT/PMM replica exchange scheme, which adds another level of parallelism to the program. The computational power of several 10,000's of cores can then easily be combined in such a calculation for rapid conformational sampling of the molecule at high accuracy.

References

[1] **Senn, H.M., Thiel, W.**
Angew. Chem., Int. Ed. 48, 1198, 2009

[2] **Schwörer, M., Breitenfeld, B., Tröster, P., Bauer, S., Lorenzen, K., Tavan, P., Mathias, G.**
J. Chem. Phys. 138, 244103, 2013

[3] **Lorenzen, K., Schwörer, M., Tröster, P., Mates, S., Tavan, P.**
J. Chem. Theory Comput. 8, 3628, 2012

[4] CPMD, http://www.cpmd.org, Copyright IBM Corp 1990-2008, Copyright MPI für Festkörperforschung Stuttgart, 1997-2001

[5] **Mathias, G., Baer, M.D.**
J. Chem. Theory Comput. 7, 2028, 2011

• Helmut Satzger [1]
• Gerald Mathias [2]
• Magnus Schwörer [2]

[1] Leibniz Supercomputing Centre

[2] Lehrstuhl für BioMolekulare Optik, LMU München

Aircraft Wake Vortex Evolution during Approach and Landing

With and without Plate Lines

As an unavoidable consequence of lift aircraft generate a pair of counter-rotating and long-lived wake vortices that may pose a potential risk to following aircraft. The empirically motivated separation standards between consecutive aircraft which were introduced in the 1970s still apply at most airports. These aircraft separations limit the capacity of congested airports in a rapidly growing aeronautical environment. The highest risk to encounter wake vortices prevails in ground proximity where the vortices cannot descend below the glide path but tend to rebound due to the interaction with the ground surface [1]. Weak crosswinds may compensate the self-induced lateral propagation of the upwind vortex such that it may hover over the runway directly in the flight path of the following aircraft. From large eddy simulation as well as from lidar field measurements it is known that wake vortices may live significantly longer than 2 min corresponding to the 5 NM separation between a leading heavy weight class aircraft and a medium follower. Consequently, most encounters are reported at flight altitudes below 300 ft. At such low flight altitudes the possibilities of the pilot to recover from a vortex encounter are limited.

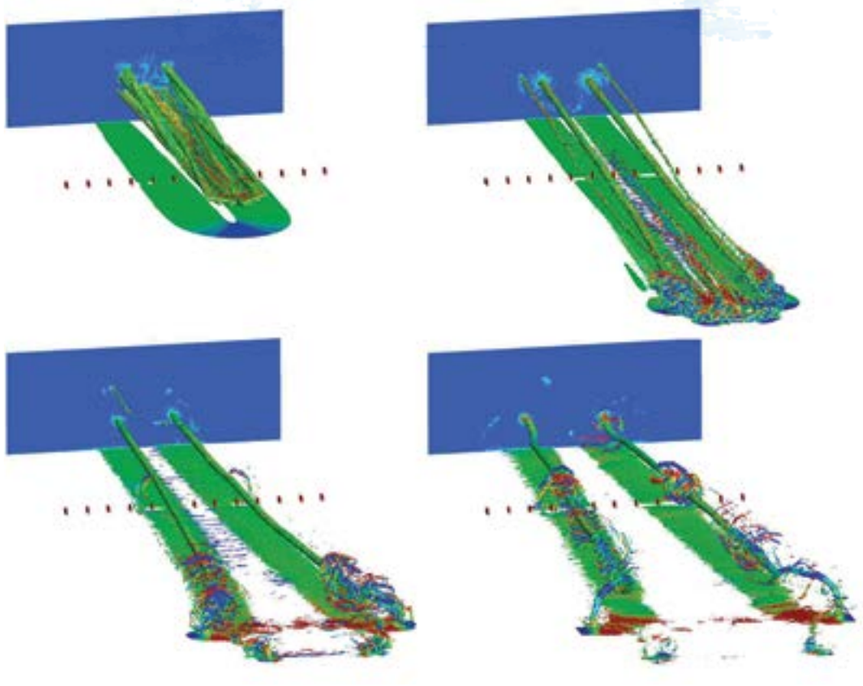

Figure 1: Simulation of wake vortex evolution during final approach of A340 aircraft in high-lift configuration. Iso-surface of vorticity magnitude colored by span-wise vorticity component.

All this suggests that comprehensive understanding of wake vortex behavior in ground proximity is of primary interest. The research activities described here are motivated by two questions:

"Why is approach and landing safe under these conditions?"

"Can we actively promote WV decay in ground proximity?"

Appropriate answers to these questions are crucial for the design of a most efficient and safe wake vortex advisory system (WVAS). Such a WVAS is conceived to adjust aircraft separations during approach and landing depending on the respective combinations of leading and following aircraft types and the prevailing meteorological conditions [2]. Although full answers to these questions are not yet tangible the numerical simulations and the related field measurement campaign described here may substantially contribute to understanding and solving the problems. Highly resolving large eddy simulations (LES) conducted on SuperMUC provide valuable insights into the physics of wake vortex behavior under various atmospheric conditions and in ground proximity. These LES also contribute indispensable guidance for the development of the real-time/fast-time wake vortex models needed in WVAS.

LES is performed using the incompressible Navier-Stokes code MGLET developed at Technische Universität München. To conduct wake vortex simulation in various atmospheric conditions, an additional equation for potential temperature is solved to take into account buoyancy effects (Boussinesq approximation).

$$\frac{\partial u_i}{\partial t} + \frac{\partial (u_i u_j)}{\partial x_j} = -\frac{1}{\rho_0}\frac{\partial p'}{\partial x_i} + (\nu + \nu_t)\frac{\partial^2 u_i}{\partial x_j^2} + g\frac{\theta'}{\theta_0}\delta_{i3},$$

$$\frac{\partial \theta'}{\partial t} + \frac{\partial (u_j \theta')}{\partial x_j} = (\kappa + \kappa_t)\frac{\partial^2 \theta'}{\partial x_j^2} + u_3 \frac{d\theta_0}{dx_3},$$

$$\frac{\partial u_j}{\partial x_j} = 0$$

where u_i, p' and θ' represent the velocity components in three spatial directions (i = 1, 2, or 3), pressure and potential temperature, respectively. The summation convention is used for the velocity components ui and δ_{ij} denotes

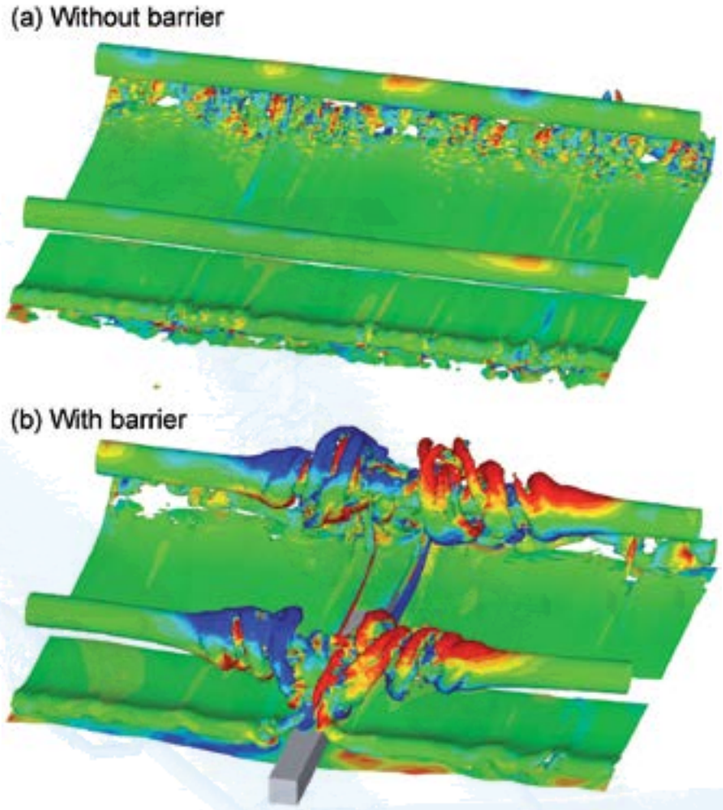

Figure 2: Wall resolving LES of wake vortex evolution in ground proximity with turbulent crosswind (a) without and (b) with a barrier on the ground surface at a vortex age of 32 s. Iso-surface of vorticity magnitude colored by span-wise vorticity component.

Kronecker's delta. The primes for pressure and potential temperature indicate that these are defined by the deviation from the reference states, hence $p = p_0+p'$, $\theta = \theta_0+\theta'$. Kinematic viscosity in the momentum equations is defined by the sum of molecular viscosity and eddy viscosity obtained from a subgrid-scale model. The corresponding diffusion coefficients in the potential temperature equation are obtained by assuming constant molecular and turbulent Prandtl numbers of 0.7 and 0.9, respectively.

The above equations are solved by a finite-volume approach with the fourth-order finite-volume compact scheme. A split-interface algorithm is used for the parallelization of a tridiagonal system for the compact scheme, which realizes good overhead time and scalability in parallel environments. In addition, a divergence free interpolation is employed for obtaining advection velocity, which ensures conservation of velocity and passive tracer fields. The pressure field is obtained by the velocity-pressure iteration method by Hirt and Cook. For time integration the third-order Runge-Kutta method is used. The Lagrangian dynamic model is employed as turbulence closure which prevents excessive eddy viscosity in the vortex cores. The closure accumulates the required averages of subgrid model coefficients along flow pathlines. This enables the Lagrangian dynamic model to distinguish the centrifugally stable vortex core regions as well as boundary layer flows from the external turbulent flow. All computations are performed in parallel by a domain decomposition approach.

Numerical simulations of wake vortices usually neglect vortex roll-up and are initialized by some analytical vortex model as for example co-rotating Lamb-Oseen vortices. A new method has been developed where an aircraft model and its surrounding flow field, obtained from high-fidelity Reynolds-averaged Navier-Stokes simulation (RANS), are swept through a ground-fixed computational domain to initialize the wake [3]. This allows the simulation of vortex evolution from the roll-up until the final decay. Close to the surface of the aircraft model the velocity field is represented by a combination of the RANS velocity field VRANS and the LES velocity field VLES with a weighting function f(y)

$$\mathbf{V} = f(y)\mathbf{V}_{LES} + [1 - f(y)]\mathbf{V}_{RANS}$$

In this study, VRANS is used as a constant forcing term of the Navier-Stokes equations solved in the LES. Since the aircraft model is swept through the computational domain, the forcing term acts as a moving boundary condition for the LES. This kind of approach might be referred to as a fortified solution algorithm or as a nudging technique as it is frequently used to assimilate meteorological measurement data into weather prediction models. In conjunction with the wake vortex evolution in ground proximity, the impact of obstacles erected on the ground surface on wake vortex evolution and decay is investigated. These obstacles are modeled by the immersed boundary method where the velocity field is modified by a forcing term such that the modified velocity field represents the obstacles on the ground.

The amount of memory used is approximately 1.5 - 3.0 GB per core. The total memory usage depends on the number of cores used, and typically is 2,048 GB for the parallel computations using 1,024 cores. For pre- and post-processing, spectral analysis requires relatively large amount of memory, e.g., 170 GB for a $1{,}024^3$ mesh case. Typical numbers of CPU cores currently used for the simulations are 512 to 2,048 cores for cases from 67 million to $1{,}024^3 \approx 1.07$ billion mesh points. The transition to SuperMUC went smoothly without larger complications.

Fig. 1 shows the simulation of the final approach of an A340 aircraft in high-lift configuration including vortex roll-up, aircraft flare and touchdown, vortex interaction with the ground and obstacles in terms of a plate line. In high-lift configuration a manifold of vortex sheets and distinct vortices are formed behind the aircraft (Fig. 1, upper left). The strongest vortices detaching form the wing tips and flap tips merge within a distance of about 10 wing spans behind the aircraft. The wake vortices as well as the bound vortex along the aircraft wing induce a vorticity layer of opposite sign at the ground surface. Shortly after touchdown the bound vortex vanishes and the free ends of the wake vortices start to interact with the vorticity layer at the ground, disturbing the wake vortices starting from the point of touchdown. This process constitutes the end effects propagating as helical disturbances along the wake vortices. Port- and starboard-vortices are no longer linked by the bound vortex and quickly diverge at the point of touchdown. Finally, we observe linking of the vortex ends with the ground. Secondary vorticity structures form at the red plate line, are wound around the primary vortices and propagate by self-induction in axial directions to either sides (lower plots). These helical disturbances finally interact with the end effects leading to rapid vortex decay.

End effects can be considered as good candidates that accelerate vortex decay close to the touchdown zone such that approach and landing can be accomplished safely. In combination with plate lines safety could be further increased

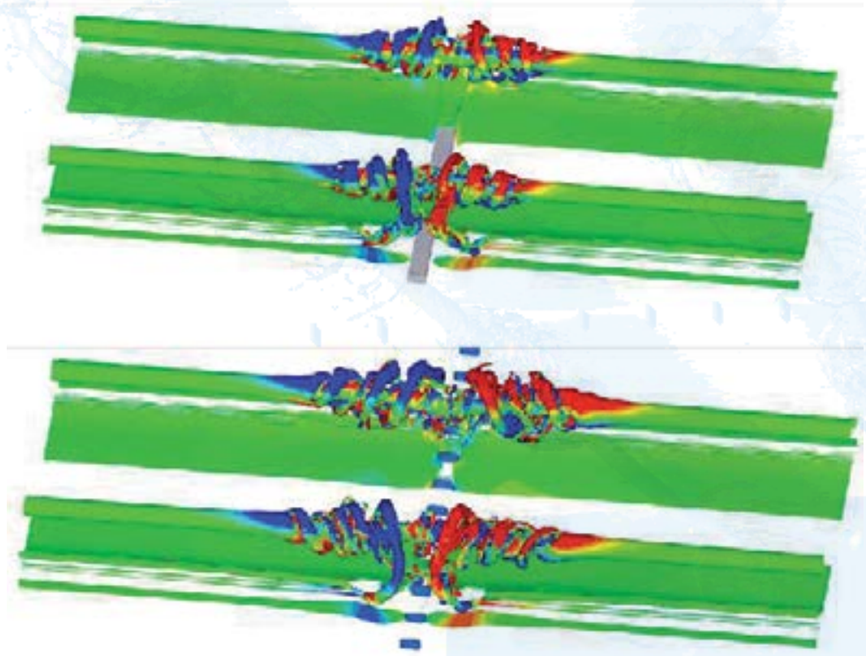

Figure 3: Wall resolving LES of wake vortex evolution in ground proximity with solid barrier (above) and plate line (below) at a vortex age of 40 s. Iso-surface of vorticity magnitude colored by spanwise vorticity component.

Figure 4: HALO research aircraft flies over plate line consisting in total of 6 wooden plates mounted perpendicular to the flight direction.

by accelerating vortex decay within the whole height range where wake vortices may rebound to the glide path.

Fig. 2a shows the interaction of the wake vortices with the turbulent structures at the ground surface generated by a crosswind blowing from left below. At a vortex age of 32 s the vorticity sheet generated by the lee (rear) vortex detaches from the ground and starts rotating around the primary vortex. Triggered by crosswind streaks the secondary vorticity sheet transforms into so-called omega loops wrapping around the primary vortices and initiating vortex decay. Under unfavorable crosswind conditions the rebounding upwind (front) vortex may hover over the runway directly in the flight corridor of a landing aircraft.

The introduction of a barrier at the ground surface may substantially accelerate vortex decay in the critical area close to the threshold where most vortex encounters occur (Fig. 2b) [4]. A respective patent entitled “Surface Structure on a Ground Surface for Accelerating Decay of Wake Turbulence in the Short Final of an Approach to a Runway” has been filed under number DE 10 2011 010 147. Such a setup specifically exploits properties of vortex dynamics to accelerate wake vortex decay in ground proximity with the following characteristics: (i) early detachment of strong omega-shaped secondary vortices, (ii) omega shape causes self-induced fast approach to the primary vortex, (iii) after the secondary vortex has looped around the primary vortex, it separates and travels both ways along the primary vortex, again driven by self induction, (iv) the artificially generated secondary vortex connects to the regular ground effect vortex and thus obtains continued supply of energy, (v) the highly intense interaction of primary and secondary vortices leads to rapid wake vortex decay independent from natural external disturbances.

Fig. 3 demonstrates that the solid barrier can be replaced by a less costly and objectionable plate line that turns out to produce similar effects. A closer look at Fig. 3 reveals that the secondary vortices are even slightly stronger

Figure 5: Overflight of HALO research aircraft and vortex roll-up in ground proximity visualized by smoke at special airport Oberpfaffenhofen.

and the propagation speed of the helical structures is even slightly higher with the plate line.

On 29 and 30 April 2013 the WakeOP field measurement campaign has been accomplished at special airport Oberpfaffenhofen with the research aircraft HALO, a modified Gulfstream G550, in order to demonstrate the functionality of the plate line to significantly accelerate vortex decay in ground proximity. During the 72 overflights of HALO at an altitude of 22 m above ground the weather impact on vortex behavior was minimized by folding away the plates alternatingly. The field experiments have successfully demonstrated the efficiency of this way to provoke premature vortex decay in the most critical flight phase prior to touch down. Already the smoke and fog visualizations documented by video and photo indicated that with the plate line the formed vortex structures are less coherent. The lidar measurements corroborate quantitatively that with the plate line vortex decay progresses faster than above flat ground at all relevant vortex ages. Further, lidar measurements in a plane with an offset of 4.5 initial vortex separations to the plate line in flight direction reveal that the disturbances travel quickly along the vortices.

References

[1] **Holzäpfel, F., Steen, M.**
Aircraft Wake-Vortex Evolution in Ground Proximity: Analysis and Parameterization, AIAA Journal, 45, 2007

[2] **Holzäpfel, F., Gerz, T., Frech, M., Tafferner, A., Köpp, F., Smalikho, I., Rahm, S., Hahn, K.-U., Schwarz, C.**
The Wake Vortex Prediction and Monitoring System WSVBS – Part I: Design, Air Traffic Control Quarterly, 17, 2009

[3] **Misaka, T., Holzäpfel, F., Gerz, T.**
Wake Evolution of High-Lift Configuration from Roll-Up to Vortex Decay, AIAA Paper 2013-0362, 2013

[4] **Stephan, A., Holzäpfel, F., Misaka, T.**
Aircraft Wake-Vortex Decay in Ground Proximity - Physical Mechanisms and Artificial Enhancement, Journal of Aircraft, DOI:10.2514/1.C032179, 2013.

Links

http://www.pa.op.dlr.de/wirbelschleppe/

• Frank Holzäpfel
• Anton Stephan

DLR,
Institut für Physik der Atmosphäre,
Oberpfaffenhofen

FORTISSIMO

Factories of the Future Resources, Technology, Infrastructure and Services for Simulation and Modelling

FORTISSIMO

Since the Cloud hype has arrived in society, the term cloud can be found nearly everywhere. Regardless of whether we look at adverts in newspapers or TV as well as at appearing add-ons in our mobiles (such as Apples iCloud or the Amazon Cloud Drive), "The Cloud" has been brought forward and is accepted as a technological standard.

However, in opposite to the initial paradigms of simplicity, ease of use and transparency, actual cloud offerings touch only a minimal subset of the desired capabilities. Data storage for the masses with limited bandwidth and adequate data security mechanisms for the industry are just a few examples concerning the current offerings. This means that the full potential of Clouds is still to be evolved and at the other hand leads to a critical and concerned view on this technology when concepts are tried to be adapted to day-to-day business of companies.

Looking at these concerns and the way the Cloud is advertised, a mentioning of Clouds in the context of supercomputing or high performance computing (HPC) applications has only led to discussions on the level of a fight of Cloud versus HPC. In a majority of cases this discussion ignored completely the potential of both technologies and thus potential synergies.

Looking at recent developments within the "classical" HPC domain there is an obvious paradigm shift from a "one supercomputer – a small number (if at all more than one) large-scale applications" relationship towards a variety of smaller applications coming from different domains to be executed on HPC systems. Especially Small and Medium-sized Enterprises (SMEs) are now discovering the potentials of simulations on supercomputers to enhance their core business and improve their competitiveness.

Thus the needs and requirements for HPC centres start to change. Whilst there was a set of rather classical HPC community members, either scientific or from major industry such as the aerospace or automotive sectors, the type of users and their expectations are about to change. In many cases SMEs that have not used simulations before, lack of expertise in the HPC technology and investments necessary for establishing this expertise internally in a company would be too risky. The main questions to be answered before there might be an investment are: "What is the benefit of using HPC?" and "What is the price?". Amongst others, these issues are discussed in detail in the PlanetHPC strategy and roadmap [1] reports.

The FORTISSIMO (Factories of the Future Resources, Technology, Infrastructure and Services for Simulation and Modelling) project has been recently started to exploit and use these synergies of High Performance Computing and Clouds. The project moves away from a plain technological viewpoint of HPC (- Clouds) up to the use of Cloud concepts empowering the provisioning of HPC services to Small and Medium Enterprises. That approach follows therefore the concepts and theories as recently published in the roadmap for cloud-computing in Horizon 2020 [2] and the Factories of the Future Roadmap of the European Factories of the Future Research Association (EFFRA [3]). Whilst the first shows clearly that developments within clouds and HPC are going in similar directions, the latter presents ICT (Information Communication Technology) as the enabling technology for the Factories of the Future and lists future needed capabilities which are definitely cloud-like.

The main focus of the FORTISSIMO project lies on the provisioning of advanced simulation in a simple and reliable way to SMEs to increase their competences in the field and to support competitiveness in Europe and beyond. As a driving force, the FORTISSIMO work plan covers the execution of a variety of various experiments, each providing requirements to the project and at the same time evaluating the results mainly from an end-users perspective. These experiments are the drivers of the FORTISSIMO evolution and will ensure focused development activities towards a tool usable by the market. Examples for those experiments are for example the simulation of high-pressure gas cylinders or the thermodynamic properties of hazardous materials. In the first year of the project, a number of 20 experiments will be active and the project foresees to announce two open calls for further experiments during its runtime to integrate new and challenging activities ensuring a proper validation of the results of FORTISSIMO and leading by that to increased maturity of the facility.

Figure 1: Complexity of using HPC.

The heart of FORTISSIMO will be the "one-stop-shop" which is a marketplace covering access to simulations and the needed hardware and software as well as to support in terms of a knowledge base and a link to experts in

Projects

different areas. Especially due to that fact, a proper execution of the experiments can be guaranteed.

In general FORTISSIMO will set up a complete ecosystem for the execution of the different experiments. This leads from the provisioning of the prototype HPC Cloud with project partners owned resources up to the evolution and realization of evolved business models (identifying the corresponding markets and pricing of offerings). In addition, the needed terms and policies for the FORTISSIMO offering and other environmental issues like license management will be addressed. An important issue here is the avoidance to start from scratch in any of those fields but to rely and base on existing software, hardware and mechanisms to build up the FORTISSIMO platform.

In order to achieve this, the primary goal of FORTISSIMO in its first months is to identify the different requirements from the experiments and resource providers (HPC centers) and transform those into a feature set that defines the functionality of FORTISSIMO's first prototype (due to March 2014).

The first version of FORTISSIMO's Marketplace will include just a minimal subset of the requested properties, such as the ability to transfer and store confidential data in a secure manner. This limitation is necessary to ensure a quick roll-out of the first release and to gain feedback from the initial 20 experiments started in parallel to the project start.

Other features that will be implemented in future versions of FORTISSIMO include single-sign-on, parallel data transfer and integrated remote visualization as well as automated suggestion of the suitable compute center regarding the special requirements. They are part of a roadmap of functionalities which will be tested and evaluated afterwards by the selected experiments of the Open Calls for experiments.

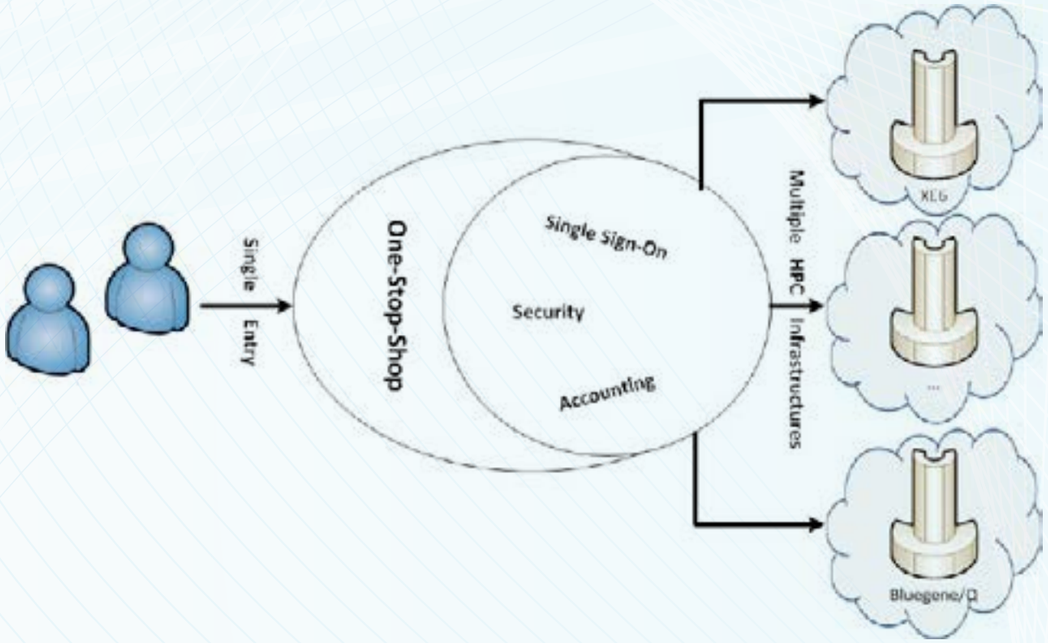

Figure 2: The FORTISSIMO Project Ecosystem.

The FORTISSIMO project started at the 1st of July 2013 and will run for 36 months with an overall budget of 21.7 Mio Euro and a funding of 16 Mio Euro provided by the European Commission. The initial consortium of 44 project partners covers both a core team of 13 partners as well as 31 experiment related consortium members. To cover all the different viewpoints on such a facility, the members of the consortium consist of hardware and software providers, ISVs, application experts and end-users.

Two Open Calls will be advertised and evaluated during the project lifetime to identify new experiments complementing the previously performed ones and to further evaluate the facility.

The role of the High Performance Computing Centre of the University of Stuttgart is thereby twofold: On one hand the operational activity of the overall FORTISSIMO facility is led by the centre, on the other hand it is deeply involved in the support of a set of experiments which are executed on the centres infrastructure.

Core Project Partners

- The University of Edinburgh (UEDIN, UK)
- SCAPOS AG (SCAPOS, D)
- Bull SAS (BULL, FR)
- Grand Equipement National de Calcul Intensif (GENCI, FR)
- Institut National de Recherche en Informatique et en Automatique (INRIA, FR)
- Universität Stuttgart (USTUTT, D)
- Consorzio Interuniversitario CINECA (CINECA, IT)
- Surfsara BV (SARA, NL)
- XLAB Razvoj Programske Opreme in Svetovanje D.O.O (XLAB, SL)
- ARCTUR Racunalniski Inzeniring D.O.O (ARCTUR, SL)
- Fundacion Centro Tecnologico de Supercomputaction de Galicia (CESGA, ES)
- Gridcore AB (GOMPUTE, SW)
- INTEL GmbH (INTEL, D)

References

[1] Sawyer, P.
A Strategy for Research and Innovation through High Performance Computing, 2012, http://www.planethpc.eu/images/stories/planethpc-strategy2.pdf

[2] Schubert, J.
Neidecker-Lutz: A Roadmap for Advanced cloud Technologies under H2020, 2012, http://cordis.europa.eu/fp7/ict/ssai/docs/cloud-expert-group/roadmap-dec2012-vfinal.pdf

[3] The European Factories of the Future Research Association (EFFRA) Factories of the Future Roadmap: http://www.effra.eu/attachments/article/335/FoFRoadmap2020_Consultation Document_ 120706_1.pdf

• Bastian Koller

University of Stuttgart, HLRS

Revisiting Dynamic Scheduling Techniques for HPC Infrastructures: The Approach of the DreamCloud Project

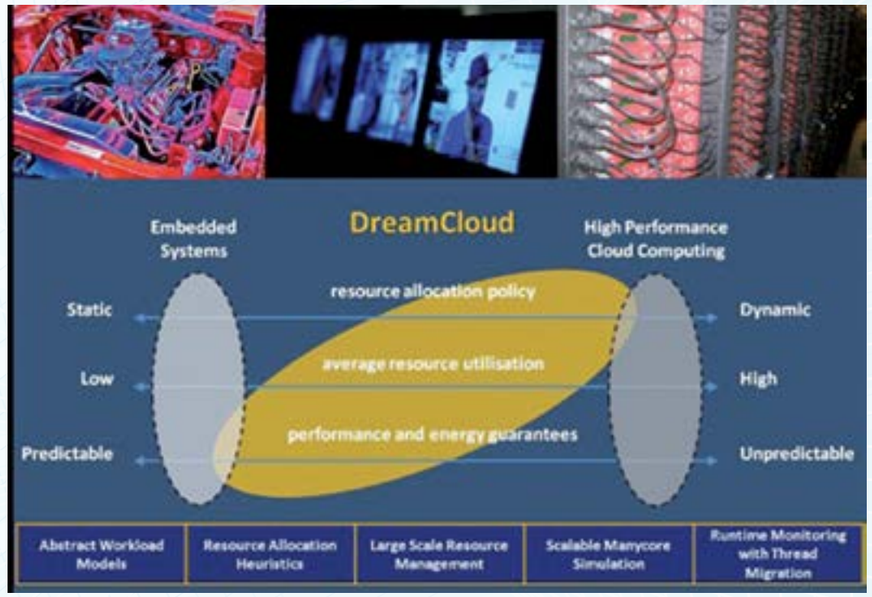

Figure 1: Revisiting Dynamic Scheduling Techniques for High Performance Computing Infrastructures: The Approach of the DreamCloud Project.

The essence of High Performance Computing (HPC) lies in sharing the large-scale hardware resources among the software applications using them. Efficient application allocation on the available compute resources is a key aspect of any HPC infrastructure functionality. The well-established HPC schedulers, such as a Portable Batch System (PBS) [1], offer effective in terms of the offered scheduling features algorithms and techniques to manage the execution of computational tasks, i.e., in the HPC terminology – batch jobs, on distributed compute nodes. However, with the emergence of high-level e-Infrastructures, such as Grid and Cloud, the traditional cluster scheduling techniques have proved useful to a limited extent only. The main reason for this is that applications running on those infrastructures require a job scheduler to offer a much more extensive set of features in terms of scalability, fault tolerance, and usability, which the traditional, static (with regard to the application) scheduling techniques are not able to meet. The execution frameworks of new-generation parallel applications, such as Hadoop/MapReduce [2], require the underlying infrastructure scheduler to be more interactive with regard to the applications, in order to enable more intelligent allocation of resources within and also beyond a batch job, i.e., the property of dynamism.

With regard to the problematic of dynamic scheduling of HPC resources, the High Performance Computing Center Stuttgart (HLRS) was engaged as a part of the University of Stuttgart in the research project “DreamCloud”, started in September 2013, partially funded by the European Commission.

The DreamCloud (Dynamic Resource Allocation in Embedded and HPC) project aims to develop novel load balancing mechanisms that can be applied during runtime in a wide range of parallel and high performance computing systems, allowing for a fine-tuning of the trade-off between performance guarantees and system efficiency according to the application needs (see Fig. 1). Such mechanisms will be organised in distinct types of cloud-like system software infrastructure that will manage the workload on different kinds of systems. Embedded Clouds will be used in systems with time-critical behaviour (such as the flight control in an aircraft), allowing for restricted load balancing and privileging strict performance guarantees. Micro Clouds will rely on novel extensions to operating systems and virtual machines, allowing for the dynamic migration of threads or full virtual machines from one core to another. Finally, High Performance Clouds will balance highly dynamic workloads, aiming for full utilisation of the underlying platform but at the same time providing performance guarantees to selected applications. A number of techniques will be explored as the underlying allocation heuristics, including bio-inspired and market-inspired techniques and control-theoretic closed loop mechanisms that rely on the monitoring capabilities of the different kinds of systems.

The HLRS is going to evaluate the major advantages of the dynamic scheduling technology for the High Performance and Cloud Computing infrastructures. For this purpose, the experimental Cloud testbed will be used, served by BonFIRE [3] and EC02Clouds [4] project infrastructures (which HLRS is also a partner of). The case study of HLRS will include two experiments representing the most common application scenarios: a workflow-based application built of loosely coupled components (plug-ins) and a massive-parallel (e.g., MPI) application. The successful conduction of this project should open a new perspective in the future of dynamic application scheduling in supercomputing environments.

Project Partners

Apart from the High Performance Computing Center Stuttgart (Germany), the following organizations participate at the project:

- X/Open Company Ltd. (United Kingdom)
- Aicas GmbH (Germany)
- University of York (United Kingdom)
- Centre National de la Recherche Scientifique (France)
- Robert Bobsch GMBH (Germany)
- Rheon Media Ltd. (United Kingdom)

References

[1] http://en.wikipedia.org/wiki/Portable_Batch_System
[2] http://hadoop.apache.org/
[3] http://www.bonfire-project.eu/
[4] http://eco2clouds.eu/

• Alexey Cheptsov
• Bastian Koller

University of Stuttgart, HLRS

SkaSim - Scalable HPC-Codes for molecular Simulation in the Chemical Industry

Molecular Dynamics (MD) and Monte-Carlo (MC) simulations allow for investigations of numerous important application scenarios in science and engineering. At their heart lie physically meaningful and quantitative models of the molecular interactions that are based on state of the art ab initio calculations and experimental data. The extreme spatial and temporal resolution (individual molecules, femtoseconds) of such simulations allows for very reliable predictions of material properties, which are particularly useful when experiments are hazardous or even impossible. However, this extreme resolution also implies a substantial computational demand in order to investigate scenarios in a timely manner. The same holds true for nanofluidics: realistic insights, not obtainable experimentally, can be captured through MD simulation. Complex phenomena as for instance phase transitions (e.g. condensation) can be investigated on the molecular level, allowing new and more fundamental understanding. However, as the dynamics of every molecule is evaluated explicitly, the number of simulated molecules needs to be large enough in order to capture the phenomenon in question. Determining experimentally elusive properties of matter is attracting increasing attention from industry. One example is process engineering, where the already highly optimized processes can only be improved through better understanding and more detailed data.

The computational power required for generating the required quantity and quality of data is significant. Thus, this demand may only be met through the efficient use of cutting-edge hardware. However, many relevant scenarios are far from trivial to simulate, e.g. interacting liquid and gas phases in a highly dynamic environment as in condensation or evaporation. The challenges to simulate such scenarios efficiently are huge.

However, the industrial development of new products and processes will experience a fundamental change in the coming years. Expensive and often potentially hazardous experiments can be replaced with safe, increasingly efficient and affordable simulations. For this transition to take place, simulations need to be yield accuracies that are comparable to high-quality experiments. In addition to the computational requirements, this calls for extremely accurate molecular models and, for complex scenarios, reliable new simulation methodologies.

The BMBF project SkaSim approaches all these challenges, namely computational performance, molecular modeling and methodological innovations. The research will build on two MD codes that were developed in a previous BMBF project: IMEMO. The code ms2 was designed to determine a large range of thermodynamic properties with high accuracy and reliability. The code ls1 mardyn was designed for large-scale nanofluidics simulations.

Performance will be tackled from two sides. On the one hand, load balancing for extremely heterogeneous simulation scenarios will be improved. The load imbalances are significant: the difference in density between a gas and a liquid is 2-3 orders of magnitude, the difference in computational load therefore is 4-6 orders of magnitude. One liquid volume element can thus be a million times more expensive to calculate than another one of the same size containing vapor. Furthermore, the load cannot be predetermined due to the highly dynamic nature of the simulations. Here, improved estimations of the load and dynamic mapping to heterogeneous hardware are the key. On the other hand, performance improvements will be achieved through an efficient use of accelerators. Accelerators have been introduced to HPC a few years ago. While they promise a huge performance boost, their utilization is far from trivial. So far in fact, accelerator nodes on supercomputers often idle due to lack of users. Here, the new pragma-based programming model OpenACC will be brought to bear. The goal of OpenACC is to simplify the usage of accelerators, thus significantly widening their usage by the application development community. OpenACC will not only be used here for efficient hybrid parallelization, but also thoroughly investigated in terms of usability and the findings will be integrated into the development process of OpenACC itself.

Molecular modeling will be addressed through new optimization strategies. Complex molecular models are difficult to parameterize. Not only do they have a high dimensionality (many degrees of freedom), a given model will also perform differently depending on the thermodynamic property in question. Thus, this can be formulated as optimization problem with multiple criteria. Determining the corresponding pareto-optimal set of solutions will allow for tuning the molecular model precisely and in essence instantly to the relevant properties from a given simulation. For the optimization of molecular models, a simulation run is the evaluation function. Because numerous evaluations are required, the performance of such simulations is essential. Furthermore, the employed optimization strategies will be parallelized.

Innovative methodologies will be investigated in the domain of phase transitions. Phase transitions are highly complex phenomena and poorly understood to date on the molecular scale. Due to the huge numbers of molecules required to capture these phenomena (even on the nanoscale), highly efficient and scalable simulations are a prerequisite. However, the actual phenomena must also be calculable. This requires, e.g. for nucleation rates, to efficiently identify, count and remove emerging clusters of molecules.

The developments described above will be brought to bear on industrially relevant simulation scenarios during the project SkaSim. The three major scenarios are: optimizing membranes for gas separation, evaluating phase transitions in reacting fluids and identifying optimal fluid mixtures for applications in the oil and gas industry. Furthermore, the integration of simulation tools with a state of the art thermodynamic database will be performed. This will allow targeting specific and relevant scenarios by simulation that are inaccessible by experiment and will thus ensure a direct route from simulations to industrial applications on a large scale.

The consortium is lead by the High Performance Computing Center Stuttgart

SkaSim

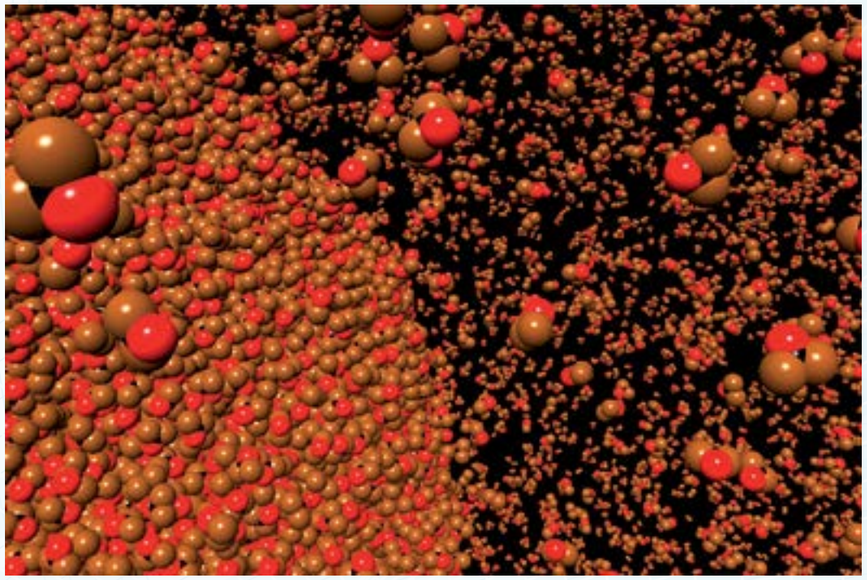

Figure 1: Simulation of the surface tension of a drop of Acetone; simulation with ls1 mardyn, one million molecules.

(HLRS), operating a petaflop system from Cray and having a history of researching high-performance molecular simulation. The consortium consists further of a leading university institute in scientific computing from TU München, the University of Applied Sciences Bonn-Rhein-Sieg and two university institutes at University of Paderborn and TU Kaiserslautern with a strong track record in molecular modeling and simulation. Moreover, two Fraunhofer institutes that routinely bridge the gap between industry and academia, i.e. ITWM in Kaiserslautern and SCAI in St. Augustin, are part of the consortium. The industrial partners come from HPC (Cray, heavily involved in the new OpenACC standard) and from applications (four enterprises involved in the chemical industry: BASF, Solvay, DDBST and Eurotechnica). BASF and Solvay are leading enterprises, whose production processes rely on thermodynamics and nanoscale phenomena. Therefore they operate large laboratories and are interested in increasingly augmenting (and on the long run replacing) experimental data with simulation data. The SME Eurotechnica specializes in supplying thermodynamic data and methods for determining them. DDBST is a leading provider of thermodynamic databases and is very experienced in interfacing material properties and applications.

Molecular modeling and simulation is a key technology for the chemical industry of the future. Having started in July 2013, the BMBF project SkaSim will run over 36 months, carrying the molecular simulation technology forward, thus enabling new products, new processes and opening up new opportunities in science and engineering.

• Colin Glass

University of Stuttgart, HLRS

POLCA: Programming Large Scale Heterogeneous Infrastructures

The limits in silicon technology have marked the end to an era denoted by constant increase of the CPUs' clock-rate following Moore's law. This has lead the main CPU manufacturers to introduce multi-core processors. But today, even the strategy of increasing performance through higher core-counts is encountering physical obstacles. For that reason, the current trend goes towards specialization of processors cores to perform well for specific recurring tasks instead of for all purposes. One implicit trend goes towards reconfigurability of the processors to allow for dynamic adaptation to the specific requirements, either only for design purposes or actually within productive systems. This implies a high degree of heterogeneity and a strong deviation on a ISA level, making programmability of such systems only possible for very experienced developers, and almost prohibits portability of the application onto different resource infrastructures.

From 2010 till 2013 the S(o)OS (Service Oriented Operating Systems) project, in which HLRS actively participated, addressed the future of distributed heterogeneous systems at the level of the operating system architecture, enabling the execution of applications across an almost unlimited number of varying resources. But just making the execution possible is not enough. Applications have to be designed to make use of parallel heterogeneous infrastructure, and for this a new programming model is needed.

Classical programming models are all designed for Turing-based sequential execution and von Neumann like memory architectures, whereas compilers typically optimize code only for one homogeneous destination infrastructure. Heterogeneous destinations typically require manual hybrid programming by the developer. There is no common current programming model that manages to address parallelization, distribution and adaptation and still remain performant.

The POLCA project will propose a programming model for minimal extensions to current models, that will allow compilers to reason over the algorithmic structure to a degree that conversion to different destination platforms and mapping to distributed, large scale environments are possible and can increase the overall execution performance significantly. What is more, these extensions will allow converting the source code to different destination platforms, including in particular FPGAs and accelerators, with a minimum of programming effort.

The POLCA Approach

In general, any program consists of a series of tasks connected by data (including state and context) passed between them. In particular in the area of High Performance Computing (HPC) the computationally demanding tasks often are strongly mathematical in nature, e.g. Molecular Dynamics (MD). In most cases, the mathematical models behind the applications are comparatively straightforward, which however does not hold true for its implementation on a given platform (i.e. the choice of an algorithm and its implementation for a given language that is able to exploit the characteristics of a specific hardware).

Pure mathematics do not imply a specific way of solving the formulas, not a specific algorithm, nor a specific platform. Though a given problem inherently favors some approaches towards the realization more than others, only the actual implementation of it specifies the concrete constraints towards the hardware and the most suitable destination platform.

$$\text{①}\quad \frac{\partial T(t,x)}{\partial t} - \alpha \cdot \nabla^2 T(t,x) = 0$$

$$\text{②}\quad \tau_i' = \begin{cases} \tau_0 & ;i = 0 \\ f(\tau_{i-1}, \tau_i, \tau_{i+1}) = \tau_i + c \cdot (\tau_{i-1} - 2\tau_i + \tau_{i+1}) & ;1 \le i \le N-1 \\ \tau_N & ;i = N \end{cases}$$

```
double heatmap[10];
double heatmap_tmp[10];
void main()
{
    initmap(*heatmap);
    for (int iter=0; iter<100; iter++)  // 100 iterations
    {
        for (int x=1; x<9; x++)
        {
            double dphi = heatmap[x-1] - 2* heatmap[x] + heatmap[x+1];
            heatmap_tmp[x] = heatmap[x] + c*dphi;
        }
        memcpy(heatmap, heatmap_tmp, 10)
    }
}
```

Figure 1: The same function of temperature development over time and space (1), its discretization (2) and its code in C language.

This does not mean that functional programming models (which are very close to mathematical formulations) are the answer to portability and performance problems. Functional languages generally have a tendency towards rather bad performance on classical (Turing-based) processing units. As opposed to that, imperative programming models, directly based on the instruction set of a classical CPU, are more aligned with the actual operation structure of the processor and thus allow for better exploitation of the hardware specifics, in particular C and FORTRAN. In other words, a purely mathematical specification alone is insufficient for an efficient execution of an application. But a mathematical declaration of the program's goal offers a fundamental benefit: the dataflow, or - more precisely - the data dependencies of the underlying algorithm can be inferred from the formula. Over recent years, the memory and communication wall have grown to be one of the most limiting factors in processing performance: a modern processor in general is able to execute operations faster than it is able to load data from storage (cache, RAM etc.). In a distributed system, and in particular in HPC parallel application, the main performance limitation is created by all exchange of data across the calculation instances. To decrease the impact of communication, application developers try to reduce the dependencies between instances to a minimum degree and try to exploit a maximum degree of concurrency (the amount of computing that can potentially be run in parallel) in the application. Most programming models however foresee that data access is immediate, i.e. comes at no cost in terms of execution performance. Though data flow models have been approached in the past, they were never really taken up by developers.

The programming model that POLCA proposes, will offer a powerful approach towards hardware-software co-design.

As explained before, mathematical specifications define all data dependencies, but due to the nature of the specification, and thus allow the dataflow to be optimized by reformulating the mathematical specification while keeping the correctness of the solution. The mathematical formulas can be restructured using typical mathematical transformations, which implicitly changes the algorithms for solving the problem and its concurrency. This way, the function (and its implementation) can be altered to improve the match of its implementation to the specific characteristics of the infrastructure, its distribution layout and its resources. The destination infrastructure is thereby not tied to a specific platform. With the mathematical notation, the resources can be equally CPUs, Accelerators (GPGPUs etc.), and FPGAs.

Programming with POLCA

With the programming model envisioned by POLCA, code will be written in a classical way, using for example C or FORTRAN, and linked through annotation with a mathematical declaration of the respective block, thereby linking input and output variables with a specific mathematical formula. In other words, POLCA allows the user to declare algorithmic blocks, "tasks", with a mathematical equivalent that can be used as additional information or even as a replacement. Especially in scientific applications, computationally demanding kernels do not change on the mathematical level, as they implement a given scientific equation.

With this specification, POLCA has direct access to all the information required to generate a dependency and dataflow graphs. Thereby, POLCA can evaluate the actual contents of the task itself and, what is more, allow for reasoning over the dependencies to reorganize this graph temporally and spatially. This means rearranging the execution order and enhancing the degree of concurrency.

As shown in Fig. 2, the dependency graph can be generated from the mathematical formula. The dependency graph provides information about the temporal behavior of the formula (y-axis) and indicates a spatial layout, without prescribing a specific clustering (x-axis). The graph is implicitly directed, in the

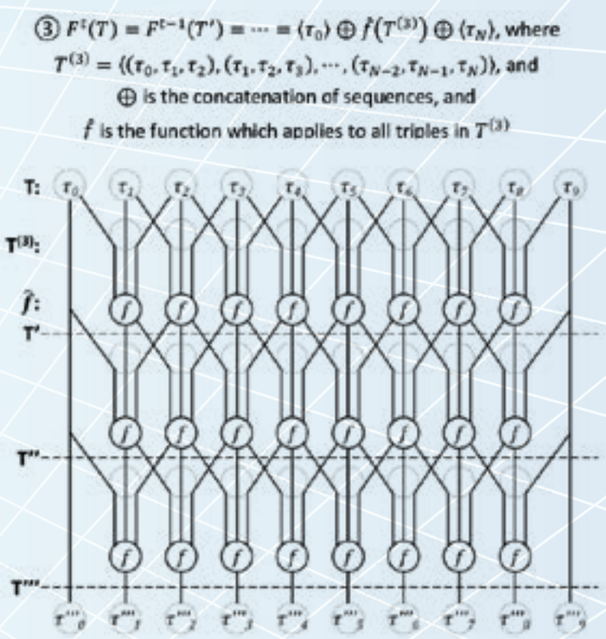

Figure 2: Simplified form of the ecuation of Figure 1 adding temporal development of the temperature and its graphical depiction for t=3 (3 timesteps).

sense that both source and destination of data are clearly identified. Reasoning over this graph allows for more freedom in the execution order and pre-calculation with partially fetched data.

Rearrangement and alteration of the code not only allows for increasing the concurrency and parallelism of a given algorithm, but also allows adaptation to a given resource infrastructure. This enables the possibility to modify the layout of the code accordingly to the communication dependencies and the actual interconnects of the destination system. Thus, the POLCA compiler and deployment tools will allow for an appropriate distribution of the executing tasks according to the infrastructure layout. Currently no programming model allows for specification of such layouts.

The capabilities and characteristics of the processing unit themselves can vary strongly. Modern processors and particularly embedded processors no longer adhere to the x86 ISA. FPGAs and GPGPUs are getting more and more common as accelerators in HPC but these hardware solutions are difficult to program individually, making the situation even more challenging in a heterogeneous environment with differing processing units combined into one system. The POLCA capability to rearrange the application layout enables the scheduler to distribute the code in a more efficient manner. The more heterogeneous the infrastructure in terms of layout, the more important the role of appropriate mapping.

GPUs are primarily vector stream processors, and hence a prime candidate for exploiting the dataflow information accessible by a mathematical model. The main problem thereby is that the memory hierarchy and the access to

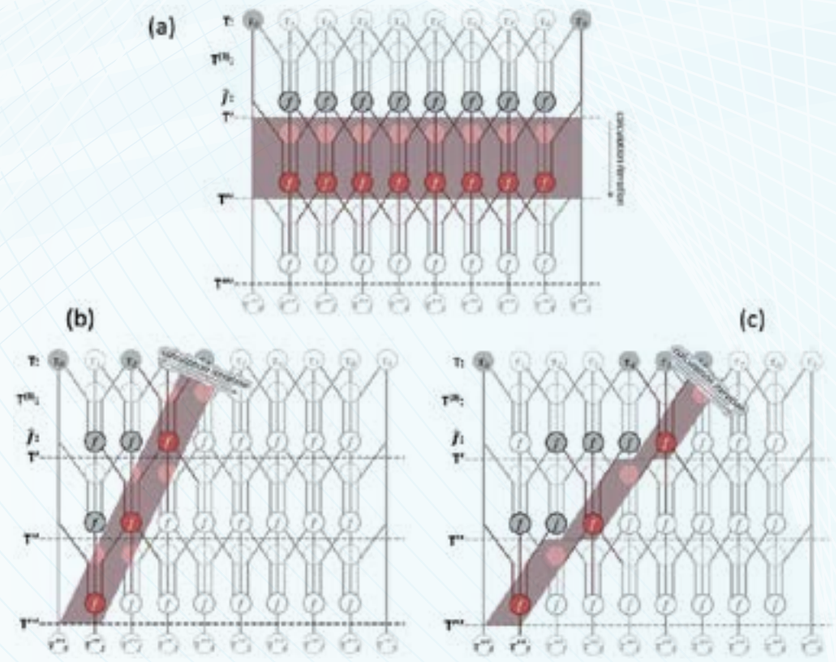

Figure 3: transformations of the calculation iterations in the dependency graph

data is complex compared to standard CPUs. A mathematical approach to programming GPUs can simplify the data and memory management. In the case of GPUs and CPUs, however, performance can be significantly improved if the core code is written in a Turing-near language, such as C. This does not necessarily hold true for FPGAs and dedicated ASICs. Mathematical equations are universal and in themselves not constrained to a specific type of processor or execution mechanism. These mathematical definitions can be converted very effectively to FPGA bitcode, making a mathematical model a very promising programming approach for such family of reconfigurable hardware. One of the goals of the POLCA project is to make FPGAs, GPUs and CPUs programmable in an accessible and hybrid fashion.

What POLCA will provide

The expected outcome of the POLCA project will be a programming model and a corresponding compiler. This compiler does not simply compile mathematical declarations into C or the other way round as this can lead to considerable impact on performance. Instead, POLCA works with a nearness model between C and its mathematical declaration that allows the tool-set and compiler to adjust the code rather than having to convert it.

POLCA tools will not compile the code in a conventional manner, but will provide a multi-phase conversion and reasoning over the whole source code of the application, delivering dedicated code segments for each resource instance, along with the according communication and data exchange mechanisms embedded into it. These code segments will thereby be ready for compilation with common, well-known and well-used compilers. In other words, POLCA will be able to incorporate common compilers into its tool-chain.

The transformed code can be represented as a graph (Fig. 3) which can be mapped against the infrastructure, which in turn can also be represented as a graph, whereby the nodes represent processing and storage units and the edges the connections between them. Different weights can represent characteristics against which the code graph can be matched. Code portability can be further increased by matching against classes of resources, rather than specific instances, reducing the last compile step to a minimum, while maintaining a maximum degree of ortability.

Who is POLCA?

POLCA is funded by the European Commission and started in September 2013. The project partners come both from academia and industry:

- University of Ulm
- University of Twente
- HLRS, University of Stuttgart
- CETIC, Centre d'Excellence en Technologies de l'Information et de la Communication
- IMDEA, Instituto Madrileño de Estudios Avanzados
- Maxeler Technologies
- Recore Systems

• Colin Glass

University of Stuttgart, HLRS

A flexible Framework for Energy and Performance Analysis

of highly parallel Applications in a Supercomputing Centre

Energy consumption is becoming the dominant cost factor in massively parallel systems. Traditionally performance properties like the number of floating point operations have been the main performance criteria for HPC systems as required by the Top500 List [1]. The aim has been to obtain the best time-to-solution of a running application. Nowadays the rising energy costs are becoming increasingly relevant and newly adopted terms such as energy-to-solution are the dominant criteria for the Green500 List [2].

The Leibniz Supercomputing Centre (LRZ) is actively involved in two projects whose topics are in the field of energy consumption of HPC systems. Both of them are sponsored by the BMBF (German Federal Ministry of Education and Research). One, the "Simulation and Optimization of Data Center energy flows from cooling networks taking into account HPC operation scenarios" (SIMOPEK, see the respective article in this issue), is a project which targets hardware components, such as the cooling infrastructure. The second BMBF project, introduced in detail in this article, is "A flexible framework for energy and performance analysis of highly parallel applications in a supercomputing centre" ("Ein flexibles Framework zur Energie- und Performanceanalyse hochparalleler Applikationen im Rechenzentrum", FEPA) which targets energy and performance optimization of applications.

Both projects are complementary and follow a global energy saving strategy of the LRZ (see Fig. 2, Page 64).

The FEPA project is conducted by the Regionales Rechenzentrum Erlangen (RRZE) in collaboration with the industrial partner NEC Deutschland GmbH. The aim of FEPA is to develop an infrastructure for detailed monitoring based on important parameters of large HPC Systems like application performance and energy efficiency.

Energy consumption is closely associated to the characteristics of the running application. For example, the number and type of instructions dispatched, cache misses, and memory utilization. The software developers need to have some background knowledge of the hardware architecture and node topology in order to do combined performance and energy optimizations within an application. HPC centres should provide assistance in this domain in order to help software developers detect energy and performance patterns. Such assistance is provided within LRZ by using an integrated monitoring tool which collects the necessary information used by the system administrators and users.

Background Software

LRZ and NEC participated in different projects from the BMBF call "HPC software for scalable parallel computing" ("HPC-Software für skalisierbare Parallelrechner") and developed independently two scalable monitoring systems. Within the project "Integrated system and application analysis for massive parallel systems" ("Integrierte System- und Anwendungsanalyse für massivparallele Rechner im Petascale-Bereich", ISAR) the PerSyst Monitoring tool was developed (PerSyst) by the LRZ. Likewise the Aggmon Monitoring System [3] was developed by NEC within the project "Tools for Intelligent System Management of Very Large Computing Systems" (TIMaCS).

The PerSyst tool [4] from LRZ runs in production mode and monitors the systemwide performance of the supercomputer. It then correlates the results with the accounting information of the running applications. The results obtained from the PerSyst tool are selectively displayed according to the permission levels of the different user groups, for instance system administrators, application support personnel, and application developers.

Aggmon, in contrast to PerSyst, is used solely by the system administrators. This tool provides the system administrators with an overview over the entire HPC system and certain events trigger system queries, which provides the monitoring information, similar to Nagios [5].

While Aggmon constitutes a generic environment for collecting and storing any arbitrary data source, the solution presented by PerSyst corresponds to the requirements of the ISAR project to monitor the performance at the application level within a Tier-1 System. Therefore, it is not yet possible for Aggmon to cover the specific requirements of PerSyst and PerSyst cannot cover the broad monitoring spectrum of Aggmon.

Different solutions for aggregation and storage were developed in both of these projects. Both monitoring systems are based on a hierarchical agent system, where the system information is collected and appropriately displayed.

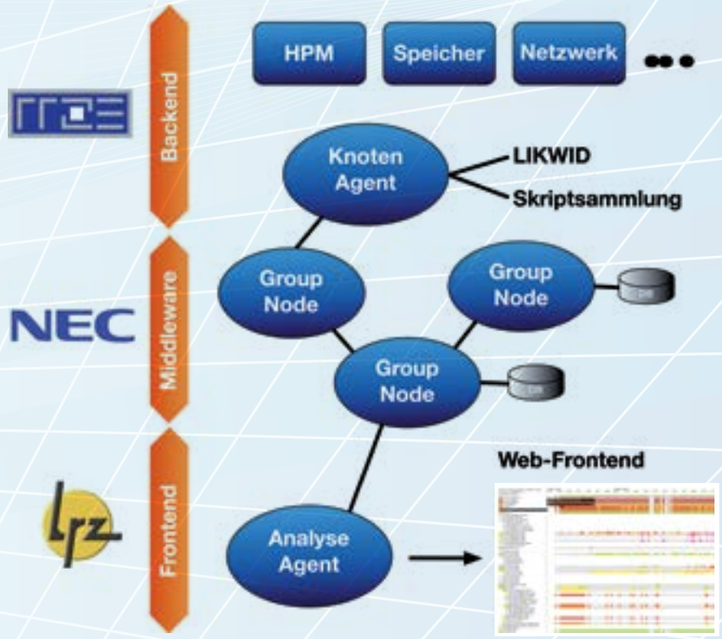

Figure 1: Software components in the FEPA Project.

The LIKWID-Tool suite [6], developed by RRZE, is a lightweight collection of performance related command line tools for the Linux operating system. LIKWID is used mainly by the end users and is designed for high flexibility and simple usage. The performance monitoring groups provide useful data for the optimization of scientific applications. Other features include the display of the node topology, the pinning of processes, and micro-benchmarking. In contrast to other existing solutions, PAPI [7] and HPCToolKit [8], LIKWID provides simple

start and stop mechanisms for measuring energy consumption, time, and performance counters. Finally, LIKWID does not use resource intensive sampling techniques, which brings the following advantages:

- LIKWID generates exact measurements with extremely low overhead.
- Given that LIKWID is not bound to any process, it is capable of measuring the uncore counters, which are part of modern architectural multicore-processors. These uncore counters contain information on the shared resources of a core. For example, bandwidth and I/O between ccNUMA domains.
- LIKWID can measure any logical thread independently from where it is running, measuring what happens on the processor, making it ideal for systemwide monitoring and application monitoring.

Both project partners LRZ and NEC had prior experience using the LIKWID tool. This tool has been used extensively by LRZ internally for performance analysis and application benchmarking. The experience attained using LIKWID and the strong cooperation with the group in Erlangen illustrates the advantages of using an internal solution for measuring performance hardware counter data.

Focus of the Project Implementation

The FEPA project is focused on three main areas: the backend components, the middleware, and the frontend.

RRZE is responsible for the development of the backend. This component provides the data acquisition mechanisms utilizing LIKWID and will be implemented as a library. It will work with the appropriate access permissions for measuring kernel data. Using simple locking mechanisms, the exclusive access to the data source will be ensured. This avoids conflict monitoring where two measurements acquired at the same time from different tools distort the results.

NEC is responsible for the middleware. The middle software layer will provide a means for the transportation, filtering, and aggregation of data. This infrastructure will provide the following functionalities:

- Perform aggregation.
- Filter and reduce the data on the application level through statistical analysis.
- The flexibility to control the resolution for data acquisition.
- Scalability without impacting the performance of the machine.
- Map the data from each node to a batch job (application context).
- Provide access to the data for the end user, as well as for the system administrator.
- Provide adequate storage, which is scalable and distributed, in order to have a fast query method for analysis algorithms.
- Provide monitoring not only for administration purposes but also for a structured performance engineering.

The frontend will be based on the web GUI from PerSyst that was developed for the ISAR project and that will be extended by LRZ. As input PerSyst will use the data obtained from the middleware. The application level as well as the system level data will be plotted.

Each component will run independently such that extensibility or exchange of one component to the other is guaranteed. There will be an interface for the interaction of each layer and, as long as the component interacts with this interface, it will be able to be coupled to the system.

Another focus of interest is the development of the performance process engineering, which should be considered by all partners in the development of the software. The performance process engineering will help an HPC centre through the process steps of detection of bottlenecks and optimization recommendations. The performance process engineering will utilize the specific hardware characteristics of the underlying HPC system to produce performance data and will be designed in collaboration with the Erlangen group who developed the LIKWID tool.

RRZE has extensive experience in the analysis of performance data for applications running on different hardware architectures. LRZ has significant experience with the PerSyst tool, which was deployed on the previous HPC systems: HLRB II and SuperMIG (the migration system of the current petaflop system SuperMUC). PerSyst was implemented with many concepts taken from the Periscope tool[9] which consider the special characteristics of the architecture while searching for bottlenecks. The emphasis is to develop an abstraction of performance patterns, which are generic and valid for different multicore platforms. Each pattern will be implemented as a function of the relevant hardware counters, in order to understand the performance characteristics of an application. In addition to the classification and analysis of the performance, the energy consumption will also be measured. This information will be used to find concrete recommendations on parameters, which will optimize the energy used by the application.

Summary

The FEPA project builds upon the TIMaCS and ISAR projects, and develops an integrated monitoring tool and a new performance process engineering. The PerSyst, Aggmon, and LIKWID tool will be integrated into the frontend, the middleware, and the backend respectively in this new integrated monitoring tool. The combination of the successful features of the three monitoring tools into a single integrated software application will ensure the continuity of the previous work done. Finally, the performance process engineering will establish the process flow for optimizing applications and will be supported by the integrated monitoring tool.

References

[1] Top500: http://www.top500.org

[2] Green500: http://www.green500.org

[3] **Focht, E., Jeutter, A.** AggMon: Scalable Hierarchical Cluster Monitoring, Proceedings of the Joint Workshop on High Performance Computing on Vector Systems, 2012

[4] **Guillen, C., Hesse, W., Brehm, M.** A scalable Monitoring Tool using Performance Properties, inside Journal, Vol. 9, No. 1, 2011

[5] Nagios: http://www.nagios.org

[6] **Treibig, J., Hager, G., Wellein, G.** LIKWID: Lightweight Performance Tools, Proceedings of an International Conference on Competence in High Performance Computing, 2010

[7] PAPI: http://icl.cs.utk.edu/papi/

[8] HPCTollKit: http://hpctoolkit.org/

[9] Periscope Project: http://www.lrr.in.tum.de/~periscop/

- Carla Guillen Carias [1]
- Wolfram Hesse [1]
- Carmen Navarrete [1]
- Jan Treibig [2]
- Matthias Brehm [1]

[1] Leibniz Supercomputing Centre

[2] Friedrich-Alexander-Universität Erlangen-Nürnberg

SIMOPEK – Simulation and Optimization of Data Center

Energy Flows from Cooling Networks taking into Account HPC Operation Scenarios

The reduction of energy consumption and the possible reuse of waste heat are becoming important issues in the design and operation of large data centers. The main objective of the SIMOPEK project [3] is to model, simulate and optimize the energy efficiency of high performance data centers, such as the Leibniz Supercomputing Centre (LRZ), using a wholistic approach. LRZ together with the Fraunhofer SCAI will develop methods and software components for modeling and simulating all energy flows of a data center for the first time. The model will take into account both the highly dynamic load behavior of the HPC system as well as new technological components (hot water cooling - IBM) and concepts for recycling the generated waste heat (adsorption cooling – SorTech AG).

The BMBF (German Federal Ministry of Education and Research) funded SIMOPEK project, together with the BMBF Project FEPA, is a critical step towards an optimal energy efficient solution.

Fig. 1 shows the vision of LRZ for an energy efficient data center using the "4 Pillar Framework for Energy Efficient HPC Data Centers [1]". The 4 Pillar Framework can be used by HPC center managers to wholistically evaluate their site, find specific focus areas, classify current research activities, and identify areas for further improvement and research. As can be seen in Fig. 2 both new BMBF projects, SIMOPEK and FEPA, together cover parts of the Modeling, Simulation and Optimization layer envisioned.

Fig. 1 shows the needed developments and connections, in and between each pillar, necessary to achieve an optimal energy efficient data center for LRZ.
In order to minimize the overall energy consumption of a data center, a tight information exchange between all four

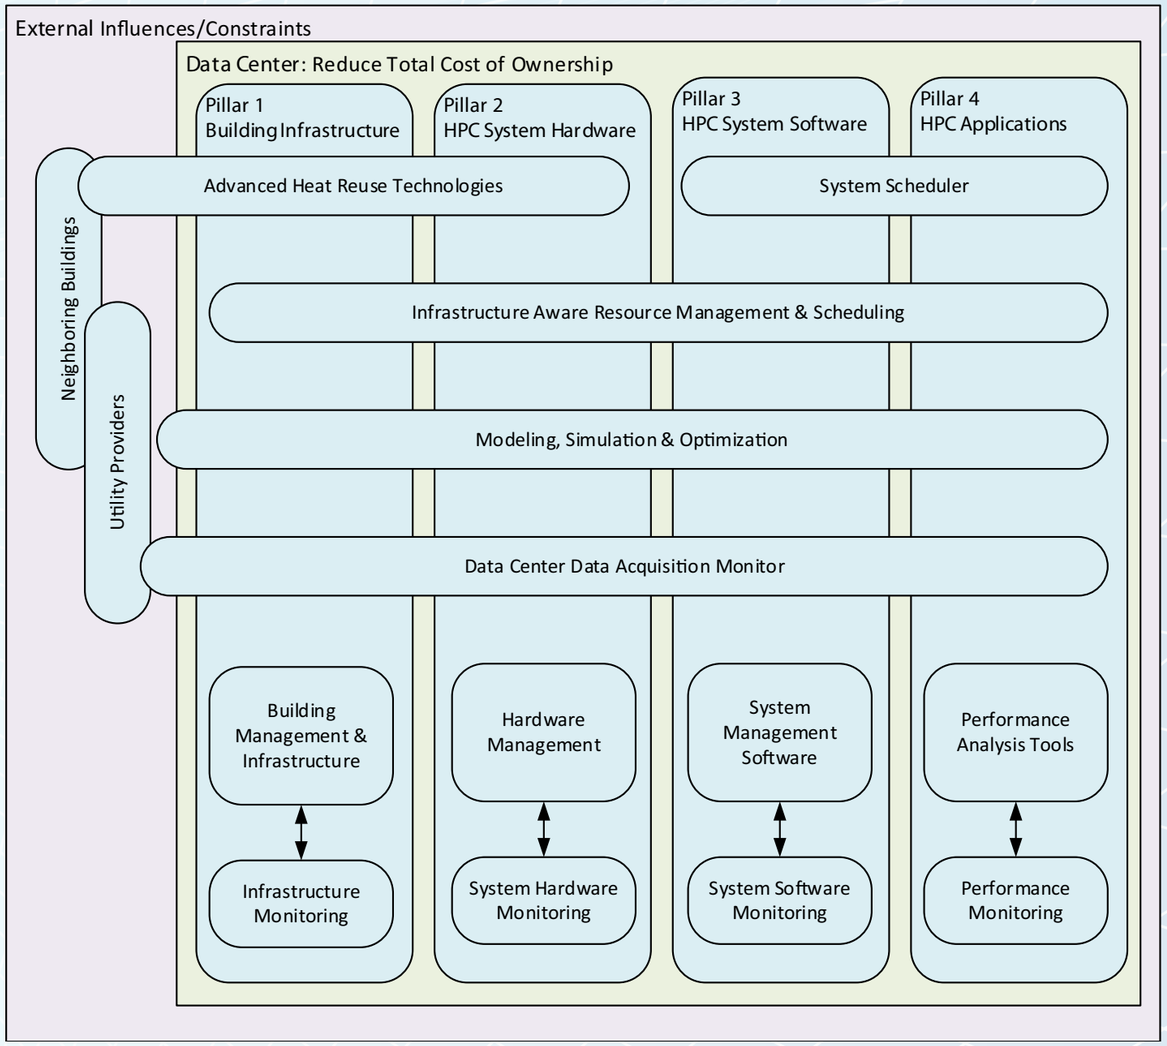

Figure 1: Vision for optimal energy efficiency solution for the BAdW-LRZ HPC data center.

pillars (Building Infrastructure, System Hardware, System Software, and Applications) is needed.

Therefore, one aim of this project is to implement a data acquisition infrastructure that is able to collect important data from pillar 1 ("Building Infrastructure"), pillar 2 ("HPC System Hardware"), and pillar 3 ("HPC System Software") as well as from the outside environment and the utility providers. This data will be fed into simulations to find energy optimized cross-pillar solutions which will be tested using the LRZ data center infrastructure. SIMOPEK will extend the current PowerDAM (a tool to calculate how much energy a user job consumed [2]) data aggregation capabilities to include the data center infrastructure and outside parameters such as air, temperature, and humidity. In addition, MYNTS (a software for gas, oil and water pipeline simulation [4]) will be extended in order to simulate and analyze energy networks of data centers (Fig. 2).

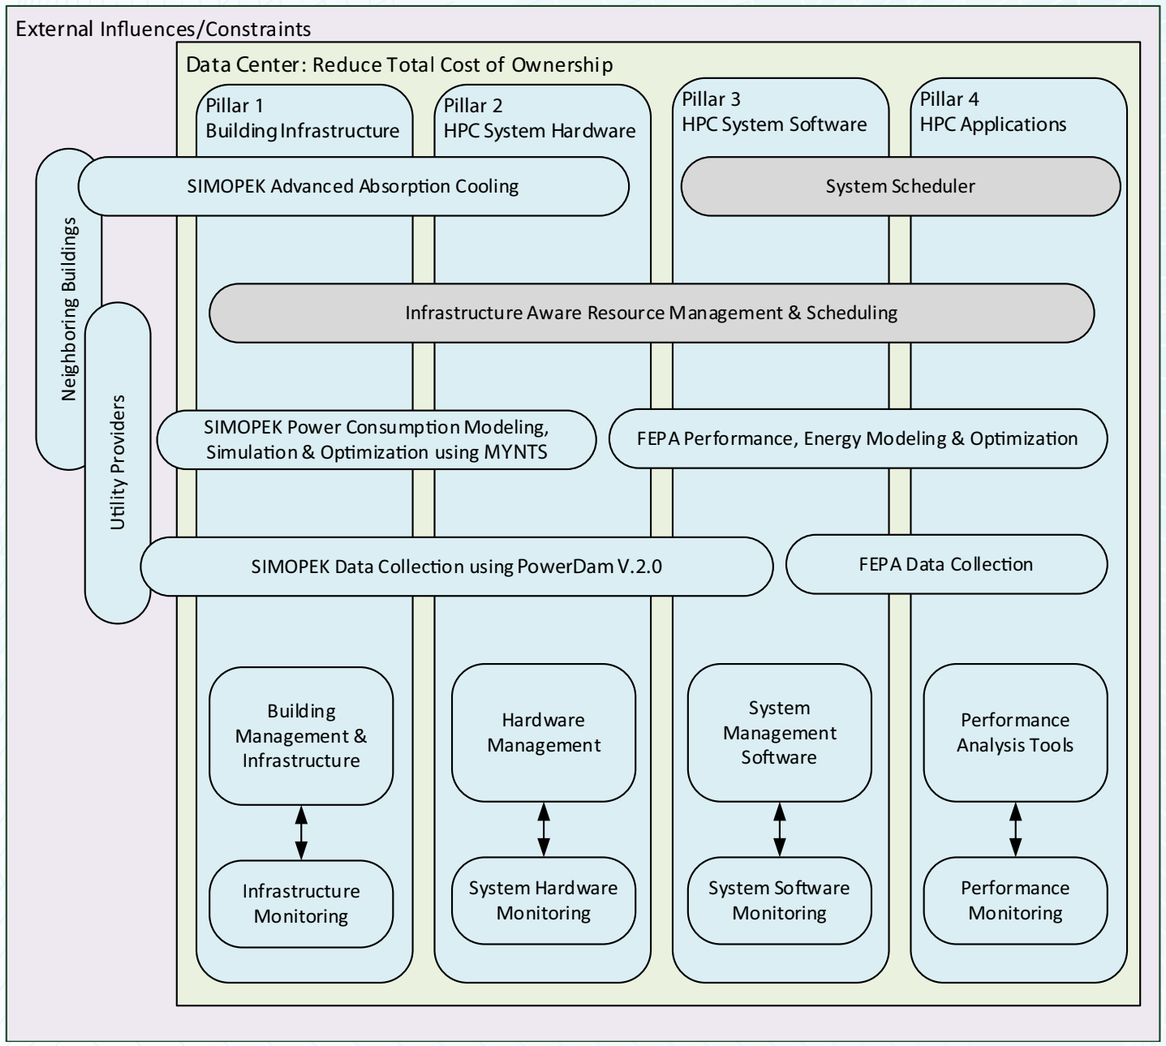

Figure 2: Current BMBF projects: SIMOPEK and FEPA, further steps to move towards the proposed vision for BAdW-LRZ HPC data center.

Another part of the SIMOPEK project is the development of concepts to allow the use of project results and software for other data centers with special focus on data centers from GCS and Gauss-Allianz. It is also envisioned that SIMOPEK will be used to support the planning of future data center infrastructures taking into account the HPC system behavior, different cooling technology options, and waste heat reuse technologies.

Facts and Figures

The SIMOPEK project started in July 2013 and will run for 3 years. The project partners are the Leibniz Supercomputing Centre (LRZ), Fraunhofer SCAI, IBM, and the SorTech AG. Associated project partners are Johnson Controls Inc., E.W. Gohl GmbH, and the Gauss Centre for Supercomputing.

References

[1] **Wilde, T., Auweter, A., Shoukourian, H.** 4 Pillar Framework for Energy Efficient HPC Data Centers, http://www.springerlink.com/openurl.asp?genre=article&id=doi:10.1007/s00450-013-0244-6

[2] **Shoukourian, H., Wilde, T., Auweter, A.** Power Data Aggregation Monitor a First Step Towards a Unified Energy Efficiency Evaluation Toolset for HPC Data Centers, ICT4S 2013: http://dx.doi.org/10.3929/ethz-a-007337628

[3] SIMOPEK Project Website: www.SIMOPEK.de

[4] MYNTS v.1.8, Fraunhofer SCAI, Sankt Augustin, Germany, 2013: www.scai.fraunhofer.de/mynts

• Tanja Cleese [1]
• Torsten Wilde [2]

[1] Fraunhofer SCAI

[2] Leibniz Rechenzentrum der Bayerischen Akademie der Wissenschaften (BAdW-LRZ)

The Catwalk Project – A quick Development Path for Performance Models

According to an old legend, the inventor of chess was asked by his king, who was thrilled by the board game, to name any reward he wanted. The inventor requested that one grain of wheat should be placed on the first square of a chessboard, two grains of wheat on the second, and so forth, doubling the number of grains for every new square. Initially, the king laughed at the inventor for asking such a low price, but later made the surprising discovery that he would not even be close to be able to pay the full reward.

Today, many HPC application developers find themselves in the situation of the king when trying to scale their code to larger numbers of processors. All of a sudden, a part of the program starts consuming an excessive amount of time. Of course, in contrast to the king in our legend, computational scientists usually possess the mathematical skills to recognize a simple geometric series. On the other hand, the laws according to which the resources needed by the code change as the number of processors increases are often much more laborious to infer and also may vary significantly across individual parts of complex modular programs. This is why analytical performance modeling is rarely attempted to predict the scaling behavior before problems manifest themselves and why this technique is still confined to a small community of experts. Unfortunately, discovering latent scalability bottlenecks through experience is painful and expensive. Removing them requires not only potentially numerous large-scale experiments to track them down, prolonged by the scalability issue at hand, but often also major code surgery in the aftermath. Not infrequently, this happens at a moment when the manpower is needed elsewhere, which is especially true for applications on the path to Exascale, which have to address numerous technical challenges simultaneously, ranging from heterogeneous computing to resilience. Since such problems usually emerge at a later stage of the development process, dependencies between their source and the rest of the code that have grown over time can make remediation even harder.

If today developers decide to model the scalability of their code, and many shy away from the effort, they first apply both intuition and tests at smaller scales to identify so-called kernels, which are those parts of the program that are expected to dominate its performance at larger scales. This step is essential because modeling a full application with hundreds of modules manually is not feasible. Then they apply reasoning in a time-consuming process to create analytical models that describe the scaling behavior of their kernels more precisely. In a way, they have to solve a chicken-and-egg problem: to find the right kernels, they require a pre-existing notion of which parts of the program will dominate its behavior at scale – basically a model of their performance. However, they do not have enough time to develop models for more than a few pre-selected candidate kernels, inevitably exposing themselves to the danger of overlooking non-scalable code.

The objective of the Catwalk project is therefore to provide a flexible set of tools to support key activities of the performance modeling process, making this powerful methodology accessible to a wider audience of HPC application developers. The tool suite will be used to study and help improve the scalability of applications from life sciences, fluid dynamics, and particle physics.

The project is coordinated by Prof. Dr. Felix Wolf of the Laboratory for Parallel Programming, German

Figure 1: Many scalability bottlenecks are almost unnoticeable at lower scales but become prohibitive once the number of processes is increased beyond a certain point.

Projects

Research School for Simulation Sciences at Aachen. Further partners are the Institute for Scientific Computing of the Technische Universität Darmstadt (Prof. Dr. Christian Bischof), the Institute of Computer Systems of the Swiss Federal Institute of Technology Zurich (Prof. Dr. Torsten Hoefler), Jülich Supercomputing Centre, Forschungszentrum Jülich (Dr.-Ing. Bernd Mohr), and the Goethe Center for Scientific Computing of Goethe University Frankfurt (Prof. Dr. Gabriel Wittum).

A first result of the Catwalk project, after running for less than one year, is a novel tool that instead of modeling only a small subset of an application program manually, generates an empirical performance model for each part of the target program automatically, significantly increasing not only the coverage of the scalability check but also its speed [1]. All it takes to search for scalability issues even in full-blown codes is to run a manageable number of small-scale performance experiments, launch the tool, and compare the extrapolated performance of the worst instances to expectations. To make this possible, we exploit several assumptions:

We take advantage of the observation that the space of the function classes underlying these models is usually small enough to be searched by a computer program. An iterative refinement process maximizes both the efficiency of the search and the accuracy of our models.

We abandon model accuracy as the primary success metric and rather focus on the binary notion of scalability bugs. Similar to a thread checker, every scalability problem we identify is a success as long as false positives that send us in a wrong direction are rare. False negatives are, of course, undesirable but acceptable as long as the number of scalability bugs we find justifies the effort.

We create requirements models alongside execution-time models. A comparison between the two can illuminate the nature of a scalability problem. Also, the generation of requirements models is less affected by performance variations.

Given that our tool relies on the standard performance-measurement infrastructures Scalasaca [2] and Score-P [3], the extra software that we developed is so lightweight that it is economically feasible to provide it in production-level quality. Finally, we generate not only a list of potential bugs but also human-readable models that can be further elaborated to conduct a variety of deeper analyses such as investigating the possibility of cache spills.

This project is part of the DFG Priority Programme 1648 Software for Exascale Computing (SPPEXA). More information can be found at http://www.vi-hps.org/projects/catwalk/

References

[1] **Calotoiu, A., Hoefler, T., Poke, M., Wolf, F.**
Using Automated Performance Modeling to Find Scalability Bugs in Complex Codes, Proceedings of the International Conference for High Performance Computing, Networking, Storage, and Analysis (SC 2013), Denver, USA, 2013

[2] **Geimer, M., Wolf, F., Wylie, B.J.N., Ábrahám, E., Becker, D., Mohr, B.**
The Scalasca performance toolset architecture, Concurrency and Computation, Practice and Experience, 22(6):702–719, April 2010

[3] **an Mey, D., Biersdorff, S., Bischof, C., Diethelm, K., Eschweiler, D., Gerndt, M., Knüpfer, A., Lorenz, D., Malony, A.D., Nagel, W.E., Oleynik, Y., Rössel, C., Saviankou, P., Schmidl, D., Shende, S.S., Wagner, M., Wesarg, B., Wolf, F.**
Score-P: A Unified Performance Measurement System for Petascale Applications, Proceedings of the CiHPC: Competence in High Performance Computing, HPC Status Konferenz der Gauß-Allianz e.V., Schwetzingen, Germany, June 2010, pages 85–97 Gauß-Allianz, Springer, 2012

• Bernd Mohr [1]
• Felix Wolf [2]
• Alexandru Calotoiu [2]
• Torsten Hoefler [3]

[1] Jülich Supercomputing Centre

[2] German Research School for Simulation Sciences

[3] Swiss Federal Institute of Technology Zurich

Projects

GROMEX – Unified Long-range Electrostatics and Flexible Ionization

The German priority program 1648 "software for exascale computing" (SPPEXA) was launched successfully in January 2013 with an initial three year funding phase. The aim of the strategic initiative of the DFG is to fund HPC software in Germany. The main topics of SPPEXA involve computational algorithms, application software, system software, programming, software tools and data management. GROMEX, one of 13 interdisciplinary research consortia within the priority program engages in the first two topics. The project brings together scientists from different fields of research like computer science, mathematics and theoretical biophysics and aims to develop a flexible and unified toolbox in the field of particle-based simulations on the exascale.

Background of the Project

Simulations of biomolecular function in atomistic detail provide insights into the inner workings of living systems that are difficult or impossible to obtain experimentally. Molecular dynamics (MD) simulation methods have been a long-standing, successful tool to tackle this problem. A particularly challenging aspect of MD simulations is the realistic modeling of electrostatic interactions. The strength and long-ranged nature of electrostatic interactions makes them important determinants of biomolecular function and properties (Fig. 1). Challenges in the computational treatment of these interactions fall into two main areas.

Usability & Scalability

The calculation of long-range electrostatic interactions is the computationally most expensive part of an MD simulation. Thus, the efficiency of this calculation is decisive for the efficiency of the whole simulation. Currently applied techniques for the treatment of long-range electrostatics like the Particle Mesh Ewald (PME) method do not scale well on large numbers of processors and will thus not be capable of harnessing the full potential of future exascale computers. With the help of the Fast Multipole Method (FMM), the scalability of the electrostatics can be significantly improved (Fig. 2). The optimal time complexity O(N), with N being the number of particles in the system, allows to efficiently utilize the available and future HPC hardware. The FMM does not suffer from the same inherent communication requirements that limit the scalability of PME-based solvers and can thus scale up to millions of cores. The implementation developed at Jülich Supercomputing Centre also provides flexibility for the user. Provided with the required accuracy for the computed electrostatic energies as only input parameter, this implementation tunes the calculation such that the optimal runtime is achieved. This user-friendly error control supports the development of a truly flexible toolbox which also can be used in other areas of research dealing with long-range interactions, namely plasma physics or astrophysics.

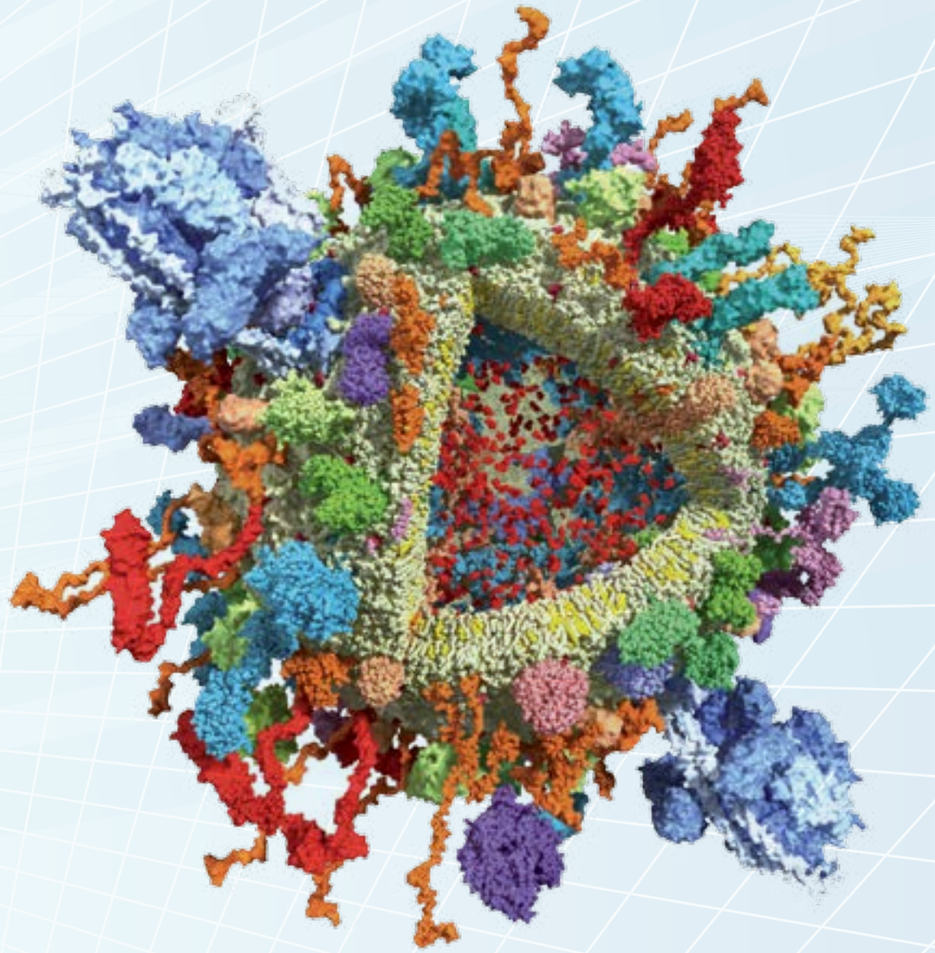

Figure 1: Functional units of biomolecular systems are often large assemblies of many different components. The simulation of these huge systems requires modern HPC computers and efficient simulation methods that take full advantage of their computational power. The figure shows an example under study at the Grubmüller department at the MPI for Biophysical Chemistry. Information between nerve cells is transmitted via messenger molecules called neurotransmitters that are transported in synaptic vesicles. The picture shows a molecular model of such a synaptic vesicle enclosed by a lipid membrane (yellow). The vesicle is filled with neurotransmitters (red spheres within the vesicle). These messenger molecules are released at the chemical synapses between nerve cells and thus transmit nerve impulses. The release of the neurotransmitters and the reloading of the vesicle involve many functional proteins, as e.g., SNARE proteins (red/orange) and V-ATPase (blue). (Figure courtesy of MPI Biophys. Chem. Jahn/Grubmüller)

Towards realistic Simulations

The treatment of biomolecular electrostatics is complicated by the fact that the charge distribution varies due to uptake and release of protons, electrons and other ions or small-molecule ligands. Here we address both tightly interlinked challenges by the development, implementation, and optimization of a unified electrostatics algorithm that will

account for realistic, dynamic ionization states (λ-dynamics) and at the same time overcome current scaling limitations. The essential idea of λ-dynamics is to enable a smooth interconversion between different protonation or other binding forms, which is crucial for MD simulations with explicit solvent molecules (Fig. 3). Improvements over previous λ-dynamics methods include that not only protonation, but also other binding reactions can be considered and that these reactions can be modeled in greater detail. The FMM allows, through its multilevel approach and spatial decomposition, to model a large variety and number of binding sites.

Additional ionization states only generate a constant computational overhead and do not require an expensive recomputation of the full interaction set. The multipole-based representation of different ionization states and sites also reduces the memory overhead and can be processed easily at very low cost.

Project Partners

- Max Planck Institute for Biophysical Chemistry, Göttingen
- Jülich Supercomputing Centre, Research Centre Jülich, Jülich
- Royal Institute of Technology, Stockholm

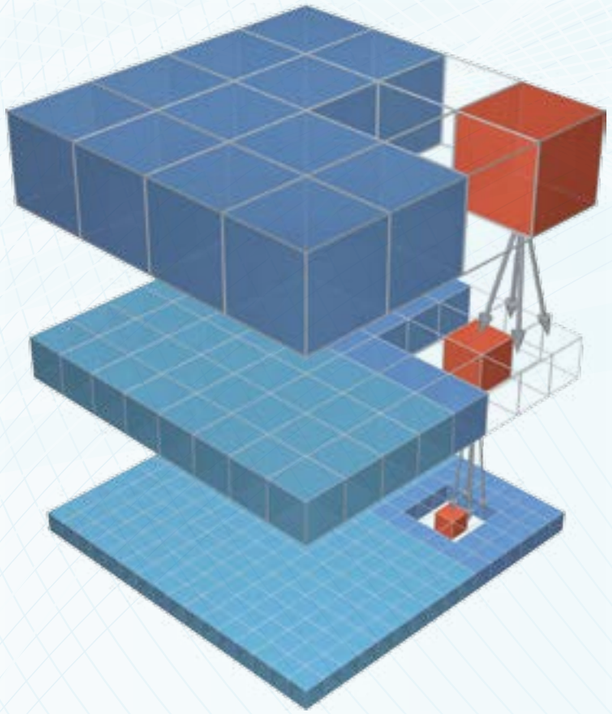

Figure 2: The figure shows the schematics of the FMM interaction set. The interaction of particles inside the depicted boxes due to their long-range forces is performed in a tree-like structure on multiple levels via box multipoles. For a given box (red) on a certain level, only a constant number of interactions (dark blue boxes) take place. The remaining interactions (light blue) are performed on a different level.

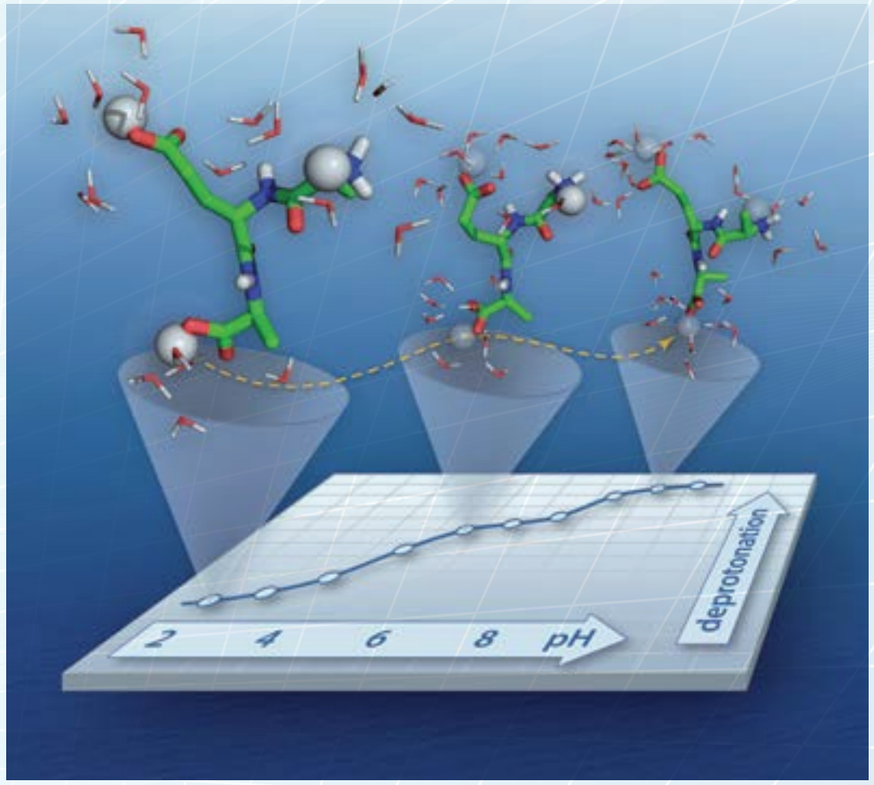

Figure 3: Binding, transfer and release of ligands, such as protons, is of central importance for the function of many biomolecules. Including these reactions in biomolecular simulations can lead to valuable insights into their influence on biomolecular function and properties. As an example, the figure shows the pH-dependent proton titration curve of a small peptide as obtained from constant-pH λ-dynamics simulations. The protonation states of the peptide's titratable groups are allowed to change dynamically during these simulations. The changes in protonation are governed by electrostatic interactions with the environment, as well as by the pH value of the surrounding solution. On top of the titration curve, the figure shows the the structures of the most probable protonation states of the peptide at selected pH values (pH 2, 6 and 10), where bound and unbound protons are depicted as opaque and faint white spheres, respectively.

References

(1) **Kabadshow, I., Dachsel, H.**
An Error Controlled Fast Multipole Method for Open and Periodic Boundary Conditions. IAS series "Fast Methods for Long-Range Interactions in Complex Systems" 6(11): 85–114, 2011

(2) **Dachsel, H.**
An Error-controlled Fast Multipole Method, J. Chem. Phys. 132(11):119901, 2010

(3) **Donnini, S., Tegeler, F., Groenhof, G., Grubmüller, H.**
Constant pH Molecular Dynamics in Explicit Solvent with λ-Dynamics. J. Chem. Theory Comput. 7, 1962-1978, 2011

(4) **Hess, B., Kutzner, C., van der Spoel, D., Lindahl, E.**
GROMACS 4: Algorithms for Highly-Efficient, Load-Balanced, and Scalable Molecular Simulation. J. Chem. Theory Comput. 4, 435-447, 2008

(5) **Ullmann, R.T., Ullmann, G.M.**
GMCT: A Monte Carlo Simulation Package for Macromolecular Receptors., J. Comput. Chem. 33, 887-900, 2012

• Ivo Kabadshow [1]
• Holger Dachsel [1]
• Carsten Kutzner [2]
• Thomas Ullmann [2]

[1] Jülich Supercomputing Centre

[2] Max Planck Institute for Biophysical Chemistry

HOPSA – A big Jump forward in HPC System and Application Monitoring

To maximize the scientific and commercial output of a High-Performance Computing system, different stakeholders pursue different strategies. While individual application developers are trying to shorten the time to solution by optimizing their codes, system administrators are tuning the configuration of the overall system to increase its throughput. Yet, the complexity of today's machines with their strong interrelationship between application and system performance demands for an integration of application and system programming.

The HOPSA project (HOlistic Performance System Analysis) therefore set out for the first time in the HPC context for combined application and system tuning developing an integrated diagnostic infrastructure. Using more powerful diagnostic tools, application developers and system administrators can easier identify the root causes of their respective bottlenecks. With the HOPSA infrastructure, it is more effective to optimize codes running on HPC systems. More efficient codes mean either getting results faster or being able to get higher quality or more results in the same time.

The work in HOPSA was carried out by two coordinated projects funded by the EU under call FP7-ICT-2011-EU-Russia and the Russian Ministry of Education and Science. Its objective was the new innovative integration of application tuning with overall system diagnosis and tuning to maximize the scientific output of our HPC infrastructures. While the Russian consortium focused on the system aspect, the EU consortium focused on the application aspect.

The HOPSA Performance Tool Workflow

One of the main results of the project was the specification and documentation of the indented usage and sequence of application of the performance tools in the form of the HOPSA performance-analysis workflow [9]. The workflow was also successfully used to structure training classes on the use of HOPSA tools, as it nicely captures the high integration of our tools set. As shown in Fig. 1, the workflow consists of three basic steps. During the first step ("Performance Screening"), we identify all those applications running on the system that may suffer from inefficiencies. This is done via system-wide job screening supported by a lightweight measurement module (LWM2) dynamically linked to every executable. The screening output identifies potential problem areas such as communication, memory, or file I/O, and issues recommendations on which diagnostic tools can be used to explore the issue further in a second step ("Performance Diagnosis"). If a more simple, profile-oriented static performance overview is not enough to pinpoint the problem, a more detailed, trace-based, dynamic performance analysis can be performed in a third step ("In-depth analysis").

The HOPSA performance tools are available as a combination of open-source offerings (the trace visualizer Paraver [6] and its measurement library Extrae and the associated performance modeling tool Dimemas [1] from BSC, the performance analysis tool Scalasca [2] and its result browser CUBE from GRS/JSC, and the community-developed performance instrumentation and measurement infrastructure Score-P [5]) and commercial products (the trace visualizer Vampir [4] from TUD and the thread and memory analyzer ThreadSpotter [3] from Rogue Wave). In the project, the individual tools have been considerably enhanced in their functionality and regarding scalability, enabling them to analyze parallel real-world applications executed with very large numbers (ten to hundred thousands) of processes and threads. Integration between the separate tool sets of the project partners also has been considerably improved. All enhancements are either already part of the latest public releases of the software packages, or at least are scheduled to be included in the next public release. Also, with the end of the project, a single unified installation package for all tools was provided [10].

Integration among the HOPSA Performance Analysis Tools

Sharing the common measurement infrastructure Score-P and its data formats and providing conversion utilities if direct sharing is not possible, the performance tools in the HOPSA environment and workflow already make it easier to switch from higher-level analyses provided by tools like Scalasca to more in-depth analyses provided by tools like Paraver or Vampir. To simplify this transition even further, the HOPSA tools are integrated in various ways. Fig. 2 gives an overview of the already implemented and envisioned tool interactions within the HOPSA tool set.

For example, with its automatic trace analysis, Scalasca locates call paths

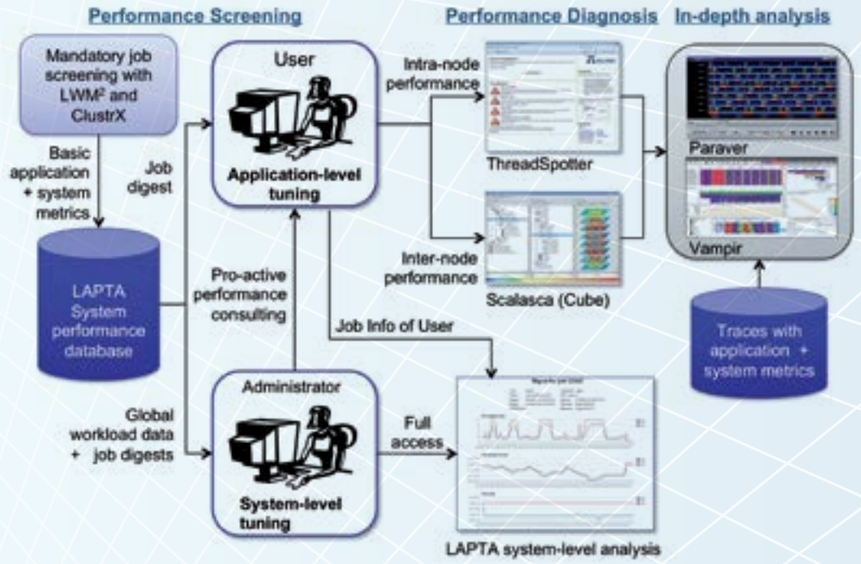

Figure 1: Overview of the performance analysis workflow.

affected by wait states caused by load or communication imbalance. However, to find and fix these problems in a user application, it is in some cases necessary to understand the spatial and temporal context leading to the inefficiency, a step naturally supported by trace visualizers like Paraver or Vampir. To make this step easier, the Scalasca analysis remembers the worst instance for each of the performance problems it recognizes. Then, the Cube result browser can launch a trace browser and zoom the timeline into the interval of the trace that corresponds to the worst instance of the recognized performance problems (see Fig. 3).

In the future, the same mechanisms will be available for a more detailed visual exploration of the results of Scalasca's root cause analysis as well as for further analyzing call paths involving user functions that take too much execution time. For the latter, ThreadSpotter will be available to investigate their memory, cache and multi-threading behaviour. If a ThreadSpotter report is available for the same executable and dataset, Cube will allow launching detailed ThreadSpotter views for each call path where data from both tools is available.

Integration of System Data and Performance Analysis Tools

The Russian ClustrX.Watch management software [7] and LAPTA system data analysis and management software [8] provides node-level sensor information that can give additional insight for performance analysis of applications with respect to the specific system they are running on. This allows populating Paraver and Vampir traces with LAPTA system information collected by Clustrx, Ganglia, and other sources (the granularity will depend on the overhead to obtain the data) and to analyze them with respect to the system-wide performance.

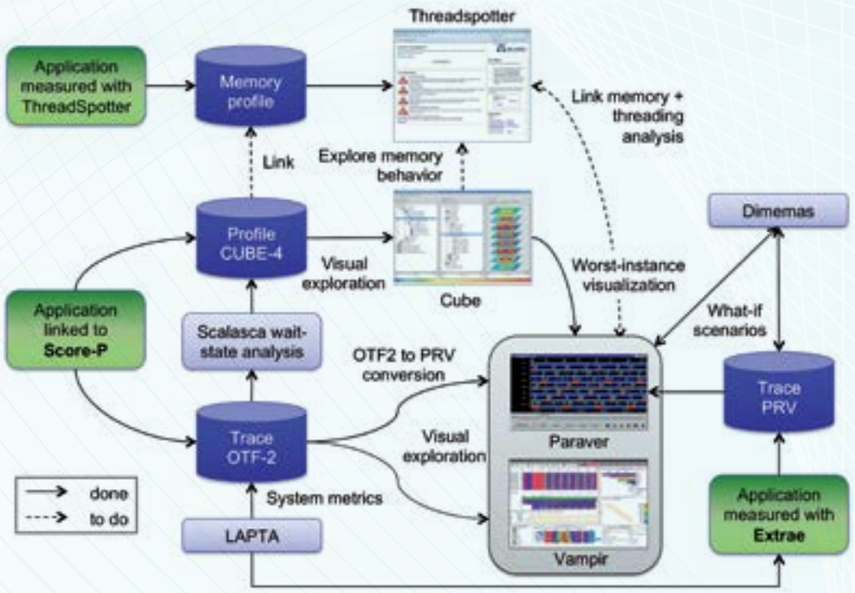

Figure 2: HOPSA Performance Tool Integration.

In the project, the Vampir team implemented a prototype Score-P adapter that enhances OTF2 traces at the end of the measurement. For evaluation, the benchmark code HPL was instrumented with Score-P. In addition to the application and MPI events, the trace was enhanced with HOPSA node-level metrics and per-process PAPI counters. Tested and working HOPSA sensors include node memory usage values and Infiniband packet counts. In the HPL code visualization (Fig. 4) one can see rising floating point operations (second timeline) resulting in a higher memory consumption per node (third timeline). Equivalent functionality was also implemented for the BSC tools Extrae and Paraver.

Conclusion

The HOPSA project delivered an innovative holistic and integrated tool suite for the optimization of HPC applications integrated with system-level monitoring. The tools are already used by the HPC support teams of project partners in their daily work. All results, documentation, and publication are available at the EU project website (http://www.hopsa-project.eu) or the Russian project website (http://hopsa.parallel.ru). For a short two-year project, dissemination was extremely successful: The project was presented at eleven events (including ISC and SC), often by multiple partners, 25 training events involving HOPSA tools have been organized, and 17 project-related publications have been published and presented at conferences.

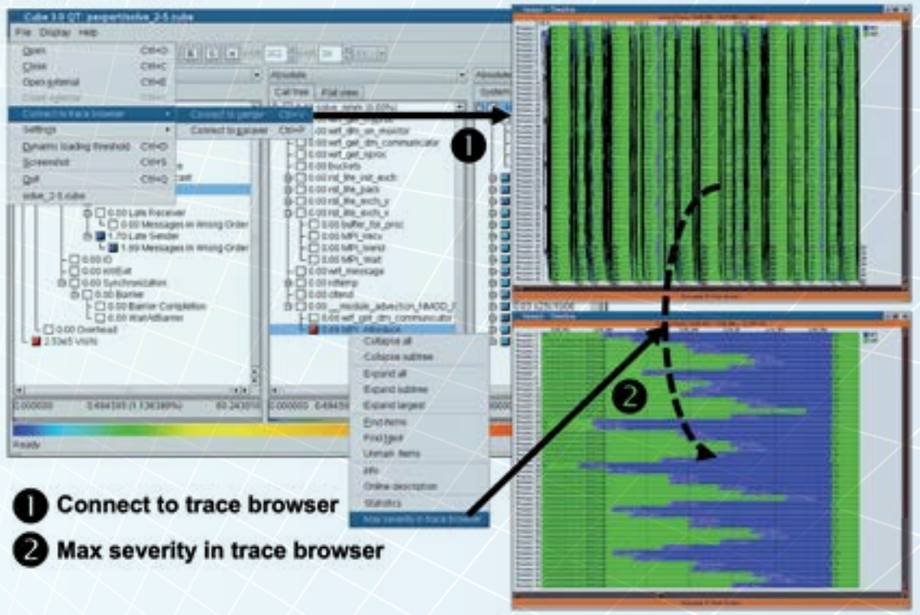

Figure 3: Scalasca →Vampir or Paraver Trace browser integration. In a 1st step, when the user requests to connect to a trace browser, the selected visualizer is automatically started and the event trace, which was previously the basis of Scalasca's trace analysis, is loaded. Now, in a 2nd step, the user can request a timeline view of the worst instance of each performance bottleneck identified by Scalasca. The trace browser view automatically zooms to the right time interval. Now the user can use the full analysis power of these tools to investigate the context of the identified performance problem.

Taking an integrated approach for the first time in an HPC context worldwide, the involved seven universities and research institutions considerably strengthened their scientific position as competence centres in HPC. Dresden University and Rogue Wave Software enriched their commercial software with unprecedented features and T-Platforms are to ship their HPC computer systems with the most advanced software offering, enabling all three of them to increase their respective market shares. Using the HOPSA tool infrastructure, the scientific output rate of a HPC cluster system can be increased in three ways: First, the enhanced tool suite leads to better optimization results, expanding the potential of the codes to which they are applied. Second, integrating the tools into an automated diagnostic workflow ensures that they are used both (i) more frequently and (ii) more effectively, further multiplying their benefit. Application programmers will ultimately benefit from higher HPC application performance by for example more accurate climate simulations or a faster market release of medication. Finally, the HOPSA holistic approach leads to a more targeted optimization of the interactions between application and system. In addition, the project resulted in a much tighter collaboration of HPC researchers from the EU and Russia.

EU Project Partners (HOPSA-EU)

- Forschungszentrum Jülich (EU Coordinator)
- Jülich Supercomputing Centre
- Barcelona Supercomputing Center
- Computer Sciences Department
- German Research School for Simulation Sciences
- Laboratory for Parallel Programming
- Rogue Wave Software AB (formerly ACUMEM)
- Technische Universität Dresden, Center for Information Services and High Performance Computing

Russian Project Partners (HOPSA-RU)

- Moscow State University (RU Coordinator)
- Research Computing Center
- T-Platforms
- Russian Academy of Sciences
- Joint Supercomputer Center
- Southern Federal University, Scientific Research Institute of Multiprocessor Computer Systems

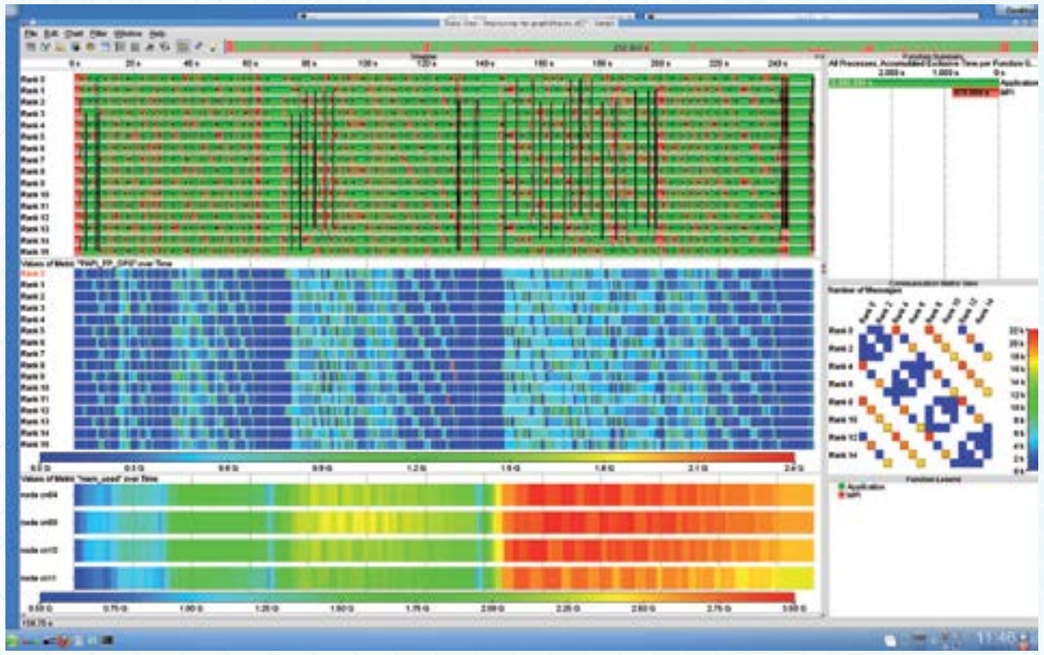

Figure 4: Vampir's Trace Visualization of the benchmark code HPL including the HOPSA node level metric "mem_used" (used memory) in the Performance Radar.

References

[1] **Labarta, J., Girona, S., Pillet, V., Cortes, T., Gregoris, L.**
DiP: A parallel program development environment, Proceedings of the 2nd International Euro- Par Conference, Lyon, France, Springer, 1996

[2] **Geimer, M., Wolf, F., Wylie, B.J.N., Abraham, E., Becker, D., Mohr, B.**
The Scalasca performance toolset architecture, Concurrency and Computation: Practice and Experience, 22(6):702–719, April 2010

[3] **Berg, E., Hagersten, E.**
StatCache: A Probabilistic Approach to Efficient and Accurate Data Locality Analysis, Proceedings of the 2004 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS-2004), Austin, Texas, USA, March 2004

[4] **Nagel, W., Weber, M., Hoppe, H.-C., Solchenbach, K.**
VAMPIR: Visualization and Analysis of MPI Resources. Supercomputer, 12(1):69–80, 1996

[5] **an Mey, D., Biersdorff, S., Bischof, C., Diethelm, K., Eschweiler, D., Gerndt, M., Knüpfer, A., Lorenz, D., Malony, A.D., Nagel, W.E., Oleynik, Y., Rössel, C., Saviankou, P., Schmidl, D., Shende, S.S., Wagner, M., Wesarg, B., Wolf, F.**
Score-P: A Unified Performance Measurement System for Petascale Applications. Competence in High Performance Computing 2010 (CiHPC), pp. 85–97. Gauß-Allianz, Springer, 2012

[6] **Servat, H., Llort, G., Giménez, J., Labarta, J.**
Detailed performance analysis using coarse grain sampling, Euro-Par 2009 - Parallel Processing Workshops, Delft, The Netherlands, August 2009, LNCS 6043, pp. 185–198. Springer, 2010

[7] T-Platforms, Moscow, Russia, Clustrx HPC Software: http://www.t-platforms.com/products/software/clustrxproductfamily.html, last accessed September 2012

[8] **Adinets, A.V., Bryzgalov, P.A., Vad, V., Voevodin, V., Zhumatiy, S.A., Nikitenko, D.A.**
About one approach to monitoring, analysis and visualization of jobs on cluster system (in Russian), Numerical Methods and Programming, 2011, Vol. 12, pp. 90–93

[9] **Mohr, B., Voevodin, V., Giménez, J., Hagersten, E., Knüpfer, A., Nikitenko, D.A., Nilsson, M., Servat, H., Shah, A., Winkler, F., Wolf, F., Zhujov, I.**
The HOPSA Workflow and Tools, Proceedings of the 6th International Parallel Tools Workshop, Stuttgart, September 2012, Springer, to appear

[10] Jülich Supercomputing Centre, Jülich, Germany, UNITE (UNiform Integrated Tool Environment): http://apps.fz-juelich.de/unite/, last accessed September 2013

• Bernd Mohr

Jülich Supercomputing Centre

End of the HPC-FF Era

The operation of the High Performance Computer for Fusion (HPC-FF), by the Jülich Supercomputing Centre (JSC), was completed in June 2013. This computer was essentially a 1/3rd partition of the JUROPA computer, consisting of Intel Xeon X5570 nodes (Nehalem-EP). Each node comprised two processors, providing a total of 8 CPU cores attached to 24 GB of main memory. Running at 2.93 GHz clock speed the 1080 nodes of the HPC-FF partition yielded a total peak performance of 101 Teraflops. The nodes were connected via Infiniband/QDR with fat-tree technology. Lustre Version 1.8 has been chosen as a parallel filesystem for both scratch and home data. The hardware was funded in-part by Forschungszentrum Jülich and in-part through the European Fusion Development Agreement (EFDA) [1], with the member countries contributing through a joint fund. Operation of HPC-FF commenced in August 2009. Allied to the physical hardware has been a High Level Support Team (HLST) [2], which has brought an important common focus to code optimisation in a multi-core environment.

Figure 1: The HPC-FF cluster at the Jülich Supercomputing Centre.

HPC-FF satisfied the need of EFDA scientists to perform large scale simulations in connection with magnetic fusion energy issues. To reach the goal of an energy producing fusion power plant (i) the plasma energy must be contained for long enough, (ii) the plasma must be macroscopically stable, (iii) the plasma facing structures must be able to handle the high heat loads coming from the plasma, (iv) in general materials are needed which survive the loads from the high energy fusion born neutrons, and (v) the components such as the blanket modules around the plasma, in which the neutron energy is extracted as heat (to generate the electricity) and the tritium to fuel the fusion reaction is bred, need to be designed. Other components such as radio frequency antennae, used to heat the plasma, also need to be designed. Studies by EFDA scientists using HPC-FF addressed aspects of all these issues.

In the tokamak, the most developed form of magnetic fusion device, the energy is usually primarily transported by short scale length turbulent instabilities (as opposed to classical collisional processes). So the understanding of turbulence, and of the means to suppress and control it, is very important in developing optimal fusion devices. Simulation of turbulence in the tokamak, and other magnetic confinement schemes (such as the stellarator), represented the largest single use of HPC-FF. Already in fluids, turbulence has been identified as one of the most challenging problems in classical physics by Nobel prize winner Richard Feynman. However, in magnetically confined fusion plasmas, the complexity is again dramatically increased as the interaction with electromagnetic fields has to be taken into account and the low collisionality means a simplified fluid description is often not applicable. A significant effort has been devoted in general, and on HPC-FF in particular by the HLST, to optimising turbulence codes.

The turbulence simulations on HPC-FF have tackled a very broad range of issues in both tokamaks and stellarators. Separately simulations have looked at turbulence on the scale lengths at which the plasma ions gyrate around the magnetic field lines and on the much smaller scale length of electron gyration, and simulations encompassing both scale lengths have also been made. Issues addressed include how flow can suppress the turbulence and how the flows themselves are generated, and in cases detailed comparisons have been made between the simulations and measurements of ion

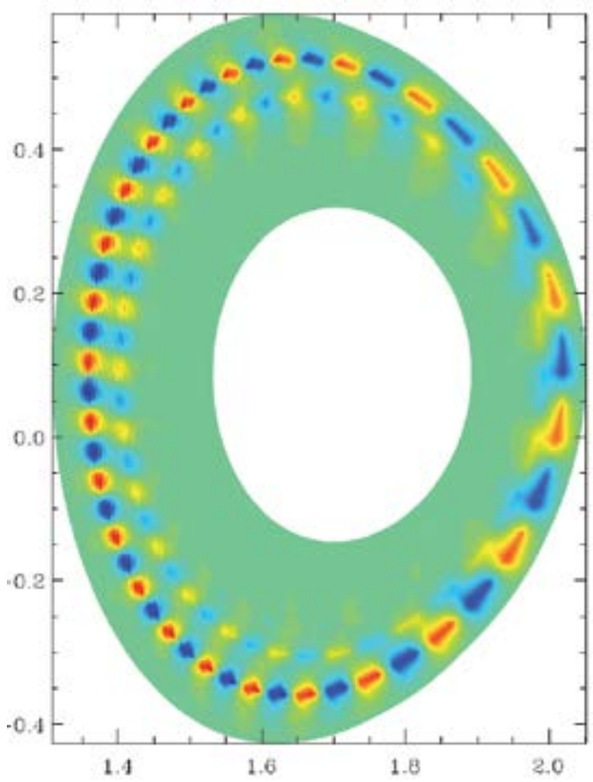

Figure 2: The colour contours (of parallel magnetic vector potential) show the linear mode structure of a microtearing instability which represents one of the most challenging turbulence types to simulate, due to its inherent multiscale character with radial structure sizes ranging from the ion-gyroradius-scale down towards the much smaller electron-gyroradius-scales. (Image courtesy of H. Doerk)

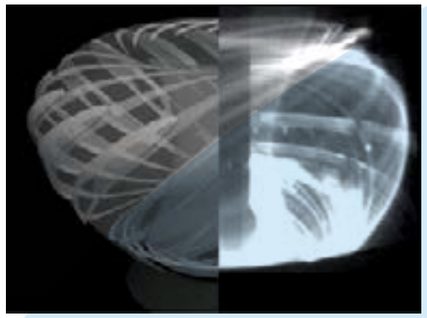

Figure 3: Comparison of the filamentation of the plasma during an ELM in the MAST tokamak (Culham UK). The left half features the simulation results; the right half features a picture from the fast camera. The simulations are based on a sophisticated magneto-hydodynamic model. (Image courtesy of S. Paméla)

scale turbulence. Not only does turbulence affect energy containment, but it also affects particle containment and simulations of the inward pinching effect on both fuel ions and impurities (which dilute the fusion reactions) have been performed. Simulations have also looked at how turbulence affects the high energy ions, resulting from the injection of neutral particle beams to heat the plasma.

Macroscopic instabilities have also been studied on HPC-FF – such instabilities may cause a complete loss of the plasma or further enhance energy/particle losses from the plasma. The largest effort on HPC-FF was on studying edge instabilities and their control. Improved confinement in the tokamak is linked with periodic instabilities known as ELMs (Edge Localised Modes) that in future tokamaks, such as ITER, are expected to impose large transient heat loads on components surrounding the plasma. HPC-FF made possible detailed ELM simulations and also allowed the study of how they may be ameliorated. Another significant area of study was on instabilities driven by high energy ions (which are generated by sources that heat the plasma or by fusion reactions).

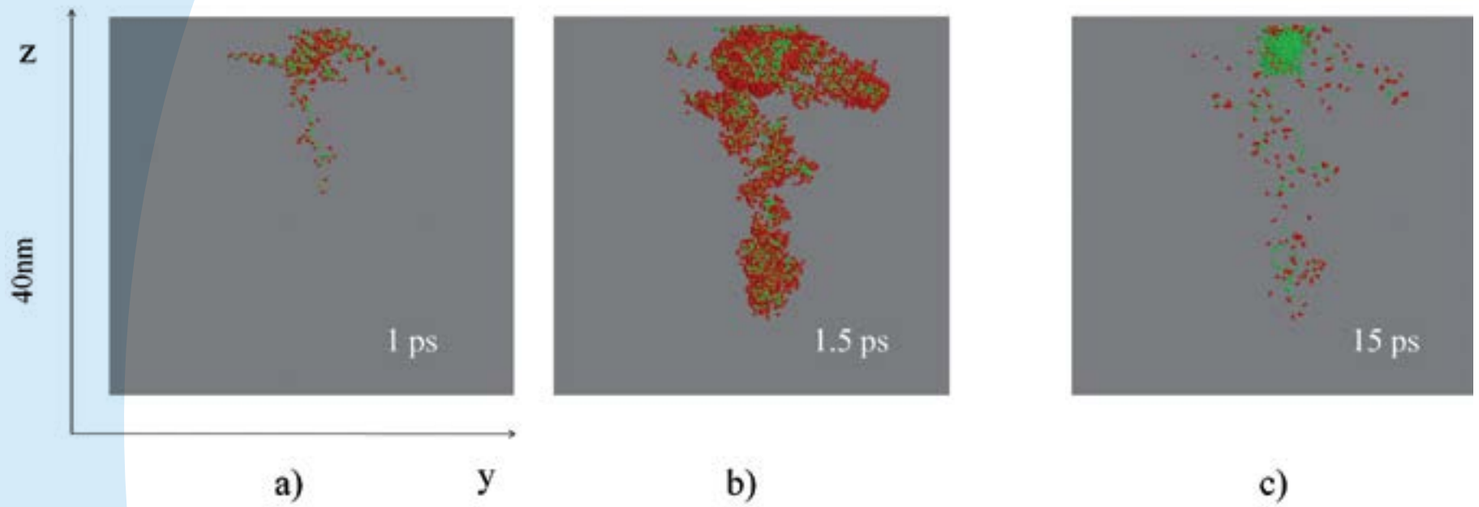

Figure 4: The figure shows three snap-shots of the time evolution of the damage produced by a high energy (150 keV) Fe ion in a 40 nm thick Fe sample. Green dots are vacancies and red dots are self-interstitials. a) 1 ps after the initiation of the recoil, b) 1.5 ps and c) 15 ps. Notice the production of a large vacancy cluster close to the surface (top of figure) showing the limitations of using small samples. (Image courtesy of M. J. Aliaga).

Density functional and molecular dynamics simulations, used to assess the effects of high energy neutrons on material properties, have also been a major user on HPC-FF. State-of-the-art ab initio electronic structure calculations have been performed by several groups to understand properties of neutron induced defects in fusion reactor materials. One of the major achievements is the study of dislocation structures in iron and tungsten, as well as the effect of alloying elements. Experimentally damage is often studied using high energy ions, as opposed to neutrons, and the use of tunnelling electron microscopes allows in-situ observations of the damage. However, it is important to understand the degree to which such experiments reproduce the real situation of neutrons impacting bulk structures, and modelling on HPC-FF has played an important role here.

Neutronics simulations (usually using Monte Carlo methods) are particularly well suited to parallelisation and important results have been obtained on HPC-FF relating to neutron shielding in ITER component designs and the associated activation of components.

Significant studies on radio frequency heating of the plasma were conducted on HPC-FF – schemes in which RF waves at frequencies that resonantly interact with ions gyrating around the magnetic field lines have been studied (this will be a primary heating means in ITER). Issues related to both the RF antenna design and how the waves couple to the plasma have been studied. Other technology related studies have focussed on blanket design and for example flows liquid metals within them.

HPC-FF has been the first shared HPC facility for the European fusion community and has allowed important progress to be made on many fronts – of which a small number of examples are presented here. Support on code development and optimisation by the High Level Support Team has also been a crucial element. Building on the success of HPC-FF, European fusion scientists now have access to a Petaflop class computer in Japan at the International Fusion Research Centre [3], under the Broader Approach [4].

References

[1] http://www.efda.org/
[2] http://www.efda-hlst.eu/
[3] http://www.iferc.org/
[4] http://fusionforenergy.europa.eu/understandingfusion/broaderapproach.aspx

• Tim Hender[1*]
• Klaus Wolkersdorfer[2**]

[1] Culham Center for Fusion Energy (CCFE), Culham, UK

[2] Jülich Supercomputing Centre (JSC)

* Tim Hender is the Fusion Programme Manager at the Culham Centre for Fusion Energy (UK) and chair of the Board that oversaw operation and CPU time allocation on HPC-FF

** Klaus Wolkersdorfer is Head of HPC systems at the Jülich Supercomputing Centre and was a permanent expert on the HPC-FF Board

JUROPA-3 - A Prototype for the Next-Generation HPC Cluster

In preparation for a future replacement of the JUROPA HPC cluster, a prototype system called JUROPA-3 has been installed at Jülich Supercomputing Centre (JSC) in April 2013. JUROPA-3 is the outcome of a cooperation of JSC, ParTec Cluster Competence Center GmbH and T-Platforms, aiming at the development of solutions for fundamental questions in large-scale cluster computing. Topics like application check-point/restart, end-to-end data integrity, network topology, failure prediction and energy efficiency are addressed in this cooperation. The development and tuning of applications for many-core processor architectures are another key aspect of the project.

Being a prototype system, the size and overall performance of JUROPA-3 are rather limited. Nevertheless, the system architecture adapts to what is expected to be relevant for a full-sized production system. It follows the same hierachical concept already successfully applied in JUROPA. Two redundant master nodes on top serve as the administrative centre of the system.

They provide fail-over functionality and house services like the batch workload manager, master LDAP and DNS servers and cluster provisioning and management functions. On the next level, a set of administration nodes provide for distributed services used by the compute nodes on level three. Typical level-two services are replica LDAP and DNS servers and DHCP. The system is complemented by front-end nodes for user login, a Lustre storage pool with OSS/MDS servers and a GPFS gateway node.

Apart from the above mentioned research topics covered in the scope of the cooperation, there are many additional issues that will be investigated on the technical level. Different from JUROPA, where SUSE SLES 11 was used as the node operating system, Scientific Linux was chosen for JUROPA-3. One task to accomplish will be the merging of administative procedures and monitoring functions with the new operating system. An expected advantage - besides cost-effectiveness - is better support of hardware and software like Infiniband and the Lustre parallel file system. New in JUROPA-3 is also, that the majority of compute nodes run in diskless mode. This has effects on installation procedures as well as run-time behaviour. Especially the handling of GPFS (General Parallel File System) has proven to be intricate in this scenario with ParaStation being used for system image administration. Despite the diskless operation, 8 nodes are equipped with local disks, which will be used for checkpoint/restart development. A scenario to be tested is the mirroring of checkpoint data to neighboring nodes, such that the data is available for restart in the case of a node failure. 16 fat nodes possess an increased amount of main memory (256 GB and 128 GB, respectively) which allow for production-sized runs of structural mechanics applications. A small storage area (2 server nodes, 60 TB disk space) is dedicated to the development and testing of future Lustre parallel file system versions and end-to-end data integrity enhancements.

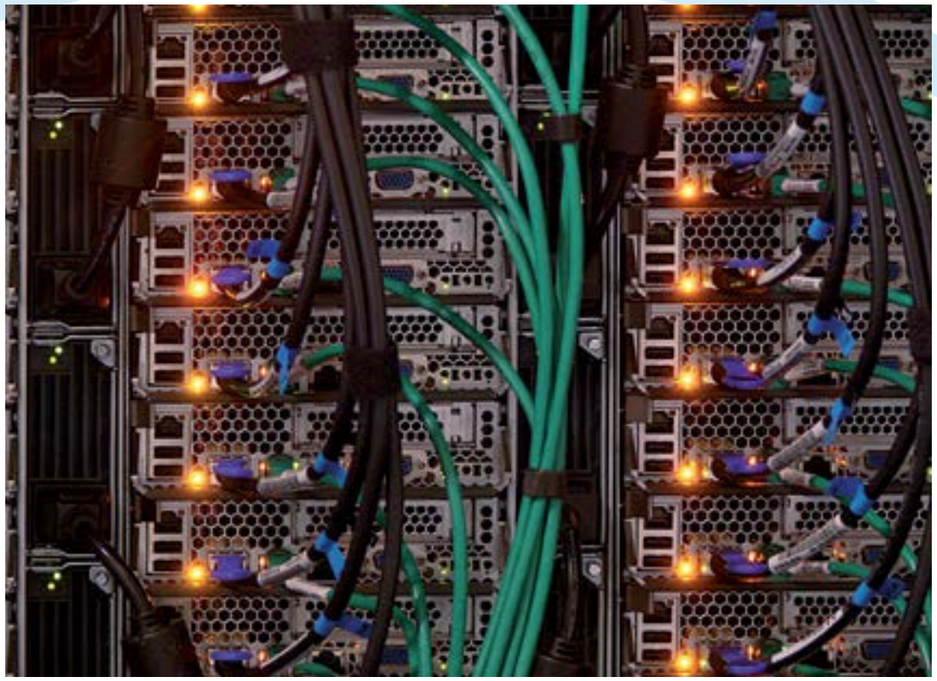

Figure 1: JUROPA-3 compute rack backview.

As mentioned above, GPFS is used as a cluster-wide file system on JUROPA-3. A GPFS gateway node with Infiniband HCA on one side and 4x10GE on the other side provides for connectivity between JUROPA-3 and FZJ's file server JUST. GPFS is mounted in addition to Lustre on all cluster nodes and serves as the main file system for home and scratch data.

Finally, the SLURM resource manager will be integrated with ParaStation and be used for job management and control. It replaces Moab and Torque used in the current JUROPA system.

System Specifications

JUROPA-3 comprises 60 compute nodes, each equipped with 2 Intel Xeon E5-2650 CPUs (Sandy Bridge-EP) providing a total of 960 processor cores with a peak performance of 15.3 teraflops. In addition, 4 compute nodes are each enhanced with 2 NVIDIA Tesla K20X GPUs and another 4 nodes with 2 Intel Xeon Phi 5110P co-processors each. The co-processors amount to an extra performance of 18.5 teraflops peak. The compute nodes, the server nodes and the storage components of the cluster are connected by a 56 Gb/s Infiniband network (FDR) with fat-tree topology providing for high communication bandwidth of parallel applications and I/O. The network topology is subject to future investigations and may be changed into hypercube or similar structure.

• Ulrich Detert

Jülich Supercomputing Centre (JSC)

Systems

First Experiences with the Intel MIC Architecture at LRZ

With the rapidly growing demand for computing power new accelerator based architectures have entered the world of high performance computing since around five years. In particular GPGPUs have recently become very popular, however programming GPGPUs using programming languages like CUDA or OpenCL is cumbersome and error-prone. Trying to overcome these difficulties, Intel developed their own Many Integrated Core (MIC) architecture which can be programmed using standard parallel programming techniques like OpenMP and MPI. In the beginning of 2013, the first production-level cards named Intel Xeon Phi came on the market. The currently fastest supercomputer in the world, Tianhe-2, heading the TOP500 list published in June 2013 uses Intel Xeon Phi coprocessors to speed up its peak performance to 34 PFlop/s. LRZ has been considered by Intel as a leading research centre for evaluating coprocessors based on the MIC architecture since 2010 under strict NDA. Since the Intel Xeon Phi is now generally available, we can share our experience on programming Intel's new MIC architecture.

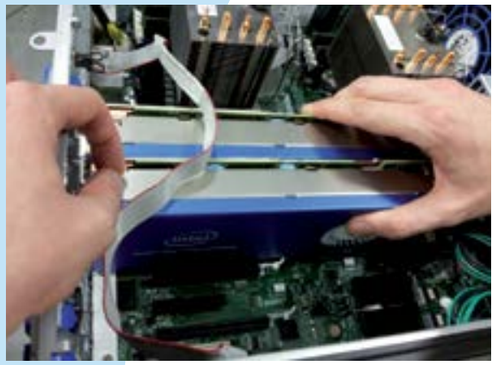

Figure 1: Installation of 2 Intel Xeon Phi cards at LRZ.

Intel MIC Architecture at LRZ

LRZ has already been selected by Intel in late 2010 as one of a few centres worldwide to evaluate the new Intel MIC architecture. Since then LRZ got various prototypes of this hardware, starting with a "Knights Ferry" prototype with 32 cores and 1 GB RAM and later various hardware stepping preproduction cards of the "Knights Corner" prototype with approximately 60 cores and up to 8 GB RAM. Initial results have been presented at the ISC '11 and under NDA in September 2012 at a workshop organised by IBM in Montpellier. LRZ is also part of the European DEEP project which develops a prototype for a potential exascale-enabling computer based on Intel Xeon Phi coprocessors and a special Terabit EXTOLL network. Currently within the PRACE-3IP project LRZ leads the subtask to create a Best Practice Guide for Intel Xeon Phi [1]. We plan to install a small Intel Xeon Phi cluster at LRZ in near future and currently prepare a MIC programming workshop in our course curriculum for 2014. Within PRACE, LRZ will analyse the scalability of the seismic community code SeisSol on up to 64 Intel Xeon Phi coprocessors of the EURORA cluster at CINECA (Bologna) as a showcase for a real MPI-based geophysics application.

Architectural Overview

The Intel Xeon Phi coprocessor is an advancement of the so called "Larrabee" chip which has never become a product. Larrabee was mainly targeting the computer graphics market and could only be programmed with lots of efforts under Windows. In contrast to that, Intel Xeon Phi is primarily targeting the HPC market and can be programmed using standard parallel programming techniques under Linux.
One of the main advantages of the coprocessor is that the programmer can directly login on the card from the host using TCP/IP based tools like ssh. This allows the user to watch and control processes running on the coprocessor using tools like top and kill and gaining the benefit of all useful information available via the Linux /proc filesystem.
An architectural overview is given in Fig. 2. A bidirectional ring interconnect connects all the cores, L2 caches and other components like the tag directories (TD), the PCIe client logic or the GDDR5 memory controllers. Details about our currently installed Intel Xeon Phi card (C0-ES2 stepping) are summarised in the following table.

Programming Models

The main advantage of the MIC architecture is the possibility to program the chip using plain C, C++ or Fortran and standard parallelisation models like OpenMP, MPI and hybrid OpenMP & MPI. The coprocessor can also be programmed using Intel Cilk Plus, Intel Threading Building Blocks, pthreads and OpenCL. OpenACC support is / will be provided by companies like CAPS or PGI. Standard math-libraries like Intel MKL are supported and last but not least the whole Intel tool chain, e.g. Intel C/C++ and Fortran compiler, debugger and Intel VTune Amplifier. It is also possible to do hardware-specific tuning using Intrinsics or Assembler. However, we would not recommend doing this (except maybe for some critical kernel routines), since MIC vector Intrinsics and Assembler instructions are incompatible with SSE or AVX instructions.

Generally speaking, two main execution modes can be distinguished: native mode and offload mode. In "native mode" the Intel compiler is instructed (through the use of the compiler-switch –mmic) to cross-compile for the MIC architecture. This is also possible for OpenMP and MPI codes. The generated executable has to be copied to the coprocessor and can be launched from within a shell running on the coprocessor.

In "offload mode" the code is instrumented with OpenMP-like pragmas in C/C++ or comments in Fortran to mark

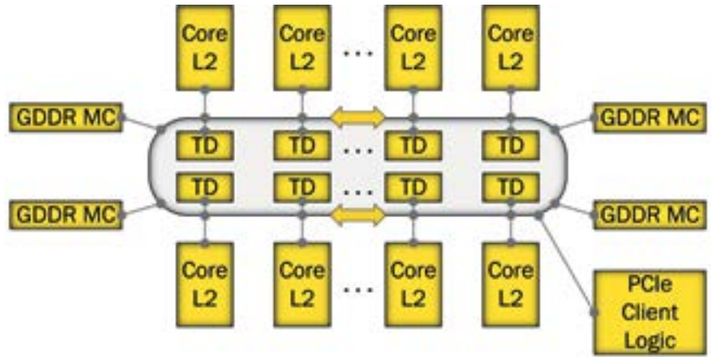

Figure 2: Overview of the Intel MIC architecture.

Number of cores	57
Frequency of cores	1.1 GHz
GDDR5 memory size	6 GB
Number of hardware threads	4
SIMD vector registere	32 (512-bit wide) per thread context
Flops/cycle	16 (DP), 32 (SP)
Theoretical peak performance	1 TFlop/s (DP), 2 TFlop/s (SP)
L2 cache per core	512 kB

regions of code that should be offloaded to the coprocessor and be executed there at runtime. The code in the marked regions can be multithreaded by using e.g. OpenMP. The generated executable must be launched from the host. This approach is quite similar to the accelerator pragmas introduced by the PGI compiler, CAPS HMPP or OpenACC to offload code to GPGPUs.

For MPI programs, MPI ranks can reside on only the coprocessor(s), on only the host(s) (possibly doing offloading), or on both the host(s) and the coprocessor(s) allowing various combinations in clusters.

Benchmarks

We have ported several mathematical and kernel benchmarks partly from the SuperMUC benchmark suite to the Intel MIC architecture. For this article we concentrate on three synthetic benchmarks from the EuroBen benchmark suite [2]: a dense matrix-matrix multiplication (MxM), a sparse matrix-vector multiplication (SpMxV) and a one-dimensional complex Fast Fourier Transformation (FFT). As representatives of three (dense linear algebra, sparse linear algebra and spectral methods) of the "seven dwarfs of HPC" they have been heavily used within PRACE to analyse 12 programming languages and

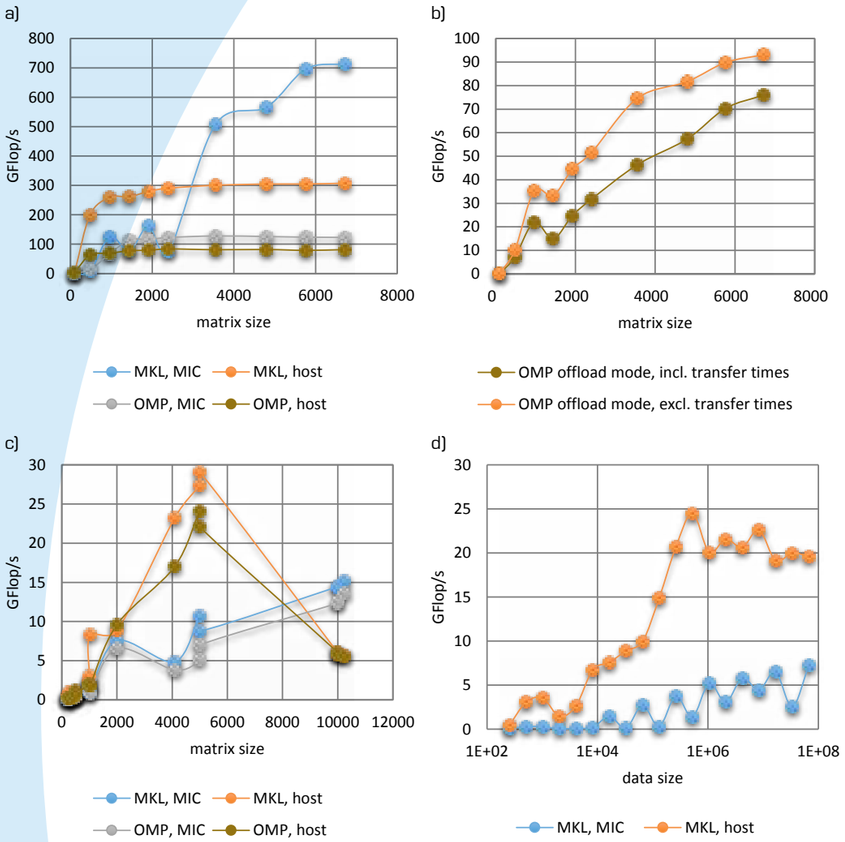

Figure 3: Performance results (double precision) for the three EuroBen kernels MxM (a) & (b), SpMxV (c) and FFT (d). Fig. (a), (c) and (d) compare the performance of MKL and/or OpenMP (OMP) based implementations on Intel Xeon Phi using approx. 228 threads in native mode with the performance on the SandyBridge-EP based host system using 16 threads. Fig. (b) shows the OpenMP performance of MxM using offload mode, both including and excluding data transfer times between the host and the coprocessor.

paradigms with respect to their performance and programmability under the leadership of LRZ [3].

Performance results for MxM (a) & (b), SpMxV (c) and FFT (d) are shown in Fig. 3. Subfig. (a), (c) and (d) compare the double precision performance on the Intel Xeon Phi coprocessor in native mode (57 cores @ 1.1 GHz with 1 TFlop/s peak) with the performance on the SandyBridge-EP based host (16 cores @ 2.6 GHz with 333 GFlop/s peak) using MKL and/or an optimised OpenMP implementation. For MxM, 714 GFlop/s could be reached using MKL on the coprocessor for the largest matrix size, which is 2.3 times faster than on the host. The optimised OpenMP implementation is 1.5 times faster for matrix sizes > 1500. Subfig. (b) displays the MxM performance in offload mode using OpenMP parallelisation in the offloaded code including and excluding data transfer times between the host and the coprocessor. For large matrix sizes data transfers become the bottleneck. In the case of SpMxV (c) the native performance on the coprocessor only excesses the host performance for the largest data sizes of the input parameter set by a factor of 2.5. The native MKL FFT performance (d) on the coprocessor is considerably below the host performance.

Concerning the ease of use and the programmability Intel Xeon Phi is a promising hardware architecture compared to other accelerators like GPGPUs, FPGAs or former CELL processors or ClearSpeed cards. Codes using MPI, OpenMP or MKL etc. can be quickly ported. Some MKL routines have been highly optimised for the MIC. Due to the large SIMD width of 64 Bytes vectorisation is even more important for the MIC architecture than for Intel Xeon based systems. It is extremely simple to get a code running on Intel Xeon Phi, but getting performance out of the chip in most cases needs hard manual tuning of the code due to failing auto-vectorisation of the Intel compiler. Using Intrinsics with manual data prefetching and efficient register usage can considerably increase the performance - but completely destroys portability. We hope that the next product of the MIC family, announced as "Knights Landing" by Intel, with integrated on-package memory and also functioning in standalone CPU mode, will deliver much better performance in combination with future releases of the Intel compiler suite.

Acknowledgements

Our work was financially supported by the KONWIHR project "OMI4papps" and by the PRACE-3IP project funded in part by the EU's 7th Framework Programme (FP7/2007-2013) under grant agreement no. RI-312763. We want to especially thank M. Klemm and M. Widmer (Intel Corp.) for continuous support and collaboration.

References

(1) **Weinberg, V., (Editor) et al.** Best Practice Guide - Intel Xeon Phi: http://www.prace-project.eu/Best-Practice-Guide-Intel-Xeon-Phi-HTML

(2) http://www.euroben.nl/

(3) PRACE deliverable D6.6, Report on petascale software libraries and programming models: http://www.prace-ri.eu/IMG/pdf/D6-6.pdf

• Volker Weinberg
• Momme Allalen

Leibniz Supercomputing Centre (LRZ)

Systems

The Extension of SuperMUC: Phase 2

On April 10, 2013, in the presence of Minister of State Dr. Wolfgang Heubisch, the contract to extend the High-Performance computer SuperMUC at LRZ was signed by Prof. Dr. Arndt Bode, Chairman of the Board of the Leibniz Supercomputing Centre (LRZ) of the Bavarian Academy of Sciences, and Martina Koederitz, General Manager, IBM Germany. The implementation of Phase 2 will start by the end of 2014 and will be finished early 2015.

SuperMUC was put into operation in July 2012 as Europe's fastest computer.

	SuperMUC – Phase 1		Phase 2
Type	Westmere-EX Intel Xeon E7-4870 10C	Sandy Bridge-EP Intel Xeon E5-2680 8C	Intel Haswell
Number of cores	8,200	147,456	74,304
Total memory (RAM)	52 TByte	288 TByte	198 TByte
Peak Performance	0.078 PFLOP/s	3.185 PFLOP/s	3.2 PFLOP/s
Size of parallel storage (GPFS)	12 PByte		9 PByte

Table 1: Summary of key parameters for SuperMUC Phase 1 and Phase 2. After the extension (Phase 2), SuperMUC will have a total of 229,960 processor cores, 538 TByte of memory, 21 PByte of storage and a peak performance of 6.4 PFLOP/s.

Figure 1: After signing the contract for SuperMUC Phase 2 (from the left): Andreas Pflieger, Director, IBM Sales Higher Education & Research, IBM Germany, Prof. Dr. Arndt Bode, Chairman of the Board of the Leibniz Supercomputing Centre, Martina Koederitz, General Manager, IBM Germany, Prof. Dr. Karl-Heinz Hoffmann, President of the Bavarian Academy of Sciences and Humanities, and Minister of State Dr. Wolfgang Heubisch. (Picture: Andreas Heddergott)

Its main characteristics are its versatility and high energy efficiency. The extension in Phase 2 will add more than 73,304 processor cores of the latest available Intel Xeon technology (Intel Haswell) to the existing 155,656 processor cores (Intel Xeon E5 – Sandy Bridge) of Phase 1. The main memory will be expanded from 340 TByte to 538 TByte, and the parallel file system (GPFS) will be extended from 12 PByte to 21 PByte with an appropriate upgrade of the bandwidth. The peak performance will be doubled to 6.4 PFLOP/s. The architecture and design of SuperMUC Phase 1 already show stable user operation and good application scaling up to more than 130,000 processor cores. This success story will be continued in Phase 2, laying the foundations for excellent scalability and the minimization of down-time despite the enormous core count.

SuperMUC Phase 2 will also include a small island based on Intels many integrated core architecture, Intel Xeon Phi. This island will be used to test and develop algorithms and applications for future supercomputer architectures.

The applications run on SuperMUC range from simulations of the evolution of the universe through the modeling of the hot interior of the earth, the propagation of seismic waves and the calculation of the flow properties of a variety of technical and natural systems to the study of biological and increasingly medical problems that directly benefit people. "In selecting the SuperMUC the decision for a computer architecture that is suitable for a wide range of scientific applications, has proved to be excellent. SuperMUC was fully occupied shortly after the start, and there are already applications that can use virtually the entire computer efficiently. Particularly in the areas of biotechnology and life sciences in the future we expect a significantly increased demand for computing power" said Prof. Dr. Arndt Bode, head of the LRZ.

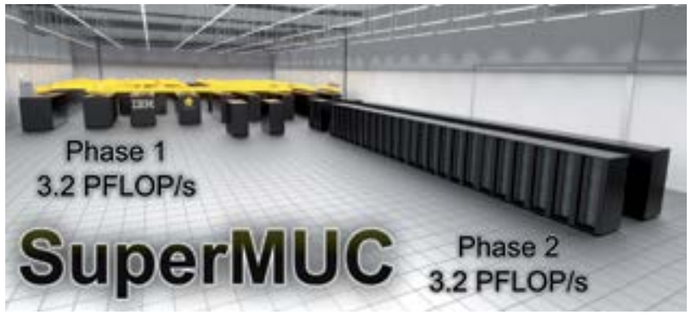

Figure 2: Illustration of SuperMUC phase 2. (Rendering: Helmut Satzger)

Figure 3: Chassis with nodes for SuperMUC phase 2. (Picture: IBM)

Innovative Water Cooling

SuperMUC Phase 2 will be cooled like the existing system using warm water (approx. 40° C). This allows on the one hand a particularly energy-efficient cooling without additional cooling machines and on the other hand the re-use of computer waste heat to heat the LRZ building. Energy-efficient system software, which was developed by IBM in collaboration with the LRZ, made it possible to optimize the power consumption and in total, SuperMUC saves more than 30% over comparable systems with a conventional cooling system.

Direct cooling of the computer chips with water also increases the durability and operational stability of the system components, as they are exposed to very little temperature variation. The chips themselves are indeed much cooler than in comparable systems with air cooling. The observed hardware failure rate, therefore, is extremely low for a system of this size.

Users from 25 European Countries

Through the Gauss Centre for Supercomputing (GCS) scientists in Bavaria, Germany and beyond use SuperMUC without modification to the existing programming concepts. The infrastructure PRACE (Partnership for Advanced Computing in Europe) opens SuperMUC's new opportunities for researchers in 25 European Member States.

Financing

The extension, including maintenance and energy, will cost 34 million Euros, and is co-financed by the German Federal Government and the state of Bavaria in an equal share. It also promotes projects such as the Bavarian Competence Network for Scientific Supercomputing KONWIHR with more than 25 applications.

References

(1) Animation of SuperMUC – Phase 1: http://www.youtube.com/watch?v=GxGrLm4ufYE

(2) Animation of SuperMUC – Phase 2: http://www.youtube.com/watch?v=qirUUIXR6XQ

(3) Video from IBM Research (SuperMUC – Phase 1): http://www.youtube.com/watch?v=LzTedSh51Tw

• Helmut Satzger
• Ludger Palm

Leibniz Supercomputing Centre (LRZ)

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

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

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

Contact:
Leibniz Supercomputing Centre

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

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

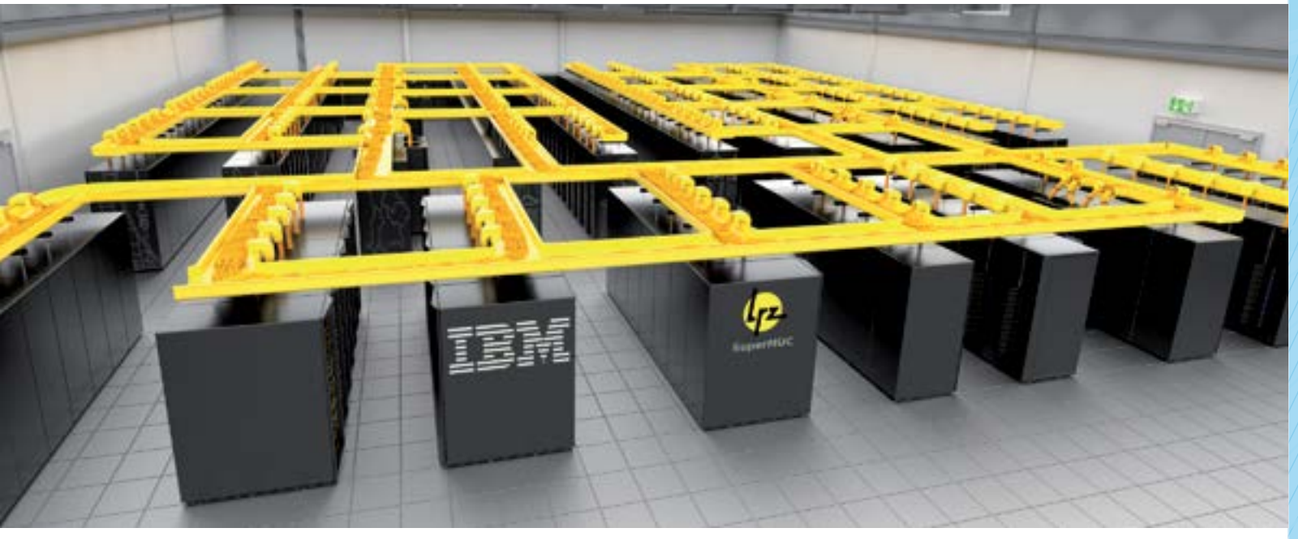

Picture of the Petascale system SuperMUC at the Leibniz Supercomputing Centre.

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

System	Size	Peak Performance (TFlop/s)	Purpose	User Community
IBM System x iDataPlex "SuperMUC"	18 islands 512 nodes with 2 Intel Sandy Bridge EP processors each 147,456 cores 288 TByte main memory	3.185	Capability Computing	German universities and research institutes, PRACE projects (Tier-0 System)
IBM BladeCenter HX5 "SuperMIG"	205 nodes with 4 Intel Westmere EX each 8,200 cores 52 TByte main memory	78	Capability Computing	German universities and research institutes, PRACE projects (Tier-1 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	Capacity 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 A dvanced Computing (MAC), Computer Science TUM

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

HLRS

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

View of the HLRS Cray XE6 "Hermit"

Compute servers currently operated by HLRS

System	Size	Peak Performance (TFlop/s)	Purpose	User Community
Cray XE6 "Hermit" (Q4 2011)	3,552 dual socket nodes with 113,664 AMD Interlagos cores	1,045	Capability Computing	European 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

JÜLICH FORSCHUNGSZENTRUM

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

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

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

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

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

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

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

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 processors (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

CECAM Tutorials at JSC

Since 2010 Forschungszentrum Jülich is an integral part of the node structure of the Centre Européen de Calcul Atomique et Moléculaire, which consists of the central organization in Lausanne and by now of 18 nodes in Europe and Israel. Since then Jülich organizes workshops, tutorials and visiting programs on a regular basis. In 2013, four events were organized, consisting of two tutorials, one workshop and a visiting programme for students. The tutorials took place at Jülich Supercomputing Centre in September.

The first tutorial, Fast Methods for Long Range Interactions in Complex Particle Systems (www.fz-juelich.de/fcs-2013), September 9 to 13 focused on providing an overview of algorithms and methods for treating long range interactions in computer simulations of particle systems. During this school, the scalable library ScaFaCoS (www.scafacos.de), which was developed in a nationwide network project funded by the Federal Ministry of Education and Research, was used to illustrate examples. Participants were encouraged to bring their own simulation codes to include the ScaFaCoS functionality. About 17 participants from 9 countries attended this event to learn about modern methods to treat long-range Coulomb interactions and reduce the numerical complexity from $O(N^2)$ to $O(N \log(N))$ or $O(N)$. Important objectives of the lectures were the introduction of modern concepts to implement these methods on parallel computers and to learn about parallel algorithms. Lecturers from the Universities of Bielefeld, Chemnitz, Stuttgart, Wuppertal and Forschungszentrum Jülich presented state-of-the-art methods, algorithms and implementations for parallel computers of various approaches to tackle the long range interactions in many-particle systems.

The second international tutorial, Multiscale modelling methods for applications in materials science (www.fz-juelich.de/mmm-2013), was jointly organized by KIT Karlsruhe and Forschungszentrum Jülich. About 30 participants from 14 countries attended this event. 12 Lecturers from 8 countries gave lectures on research including a range of mutually coupled methodologies to overcome limitations induced by various length- and time-scales in the physical problem. The tutorial was closely related to the FP7 funded project MMM@HPC. It focused on the introduction of different methods, ranging from ab initio to coarse graining techniques. It also provided insight into modern workflow design and simulation tools. Various aspects from multiscale modeling, ranging from general overview over the field to more specific topics like coarse graining techniques or electronic properties of materials were introduced. Special focus was also given to introductions to community codes, used as parts in multiscale problems, like BigDFT, DFTB+, DL_POLY or Elmer. Furthermore, the middleware UNICORE was introduced as a basic part for designing workflows in complex simulation protocols. During the afternoon sessions, the participants could get first-hand experience with tools and application codes during practical programming lessons.

The other two CECAM events, i.e. the International Guest Student Programme and the workshop on Frontiers of computational biomolecular spectroscopy and mass spectrometry organized at Forschungszentrum Jülich are still ongoing and will be reported at a later stage.

• Godehard Sutmann

Jülich Supercomputing Centre (JSC)

CHANGES Workshop

The second workshop of the CHANGES series took place from September 10 to 12, 2013 in Chicago. CHANGES stands for CHinese-AmericaN-German E-Science and cyberinfrastructure.

This second CHANGES workshop addressed Data-Driven Science: "Data Management, Analytics and Visualization".

Again about 40 well-known experts

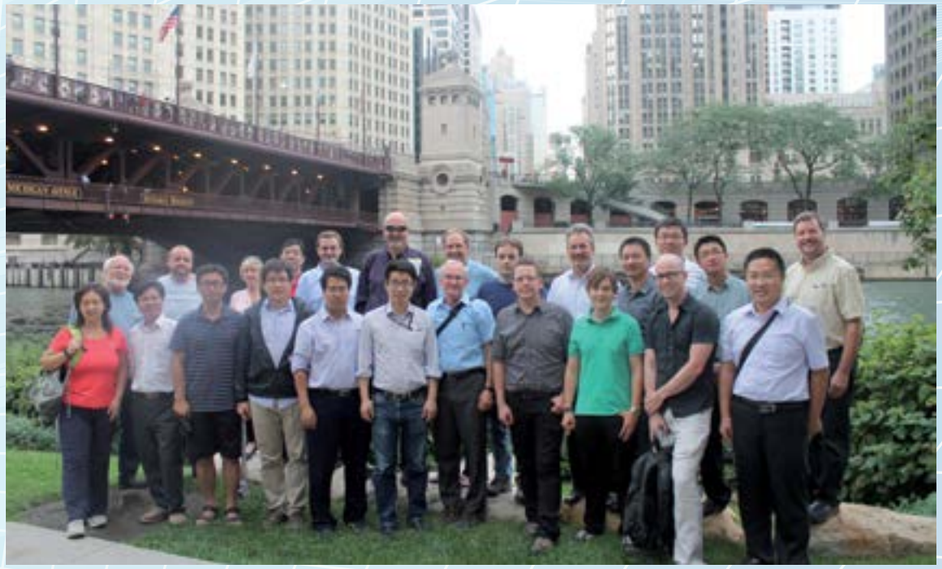

Partners are the Computer Network Information Center (CNIC) of the Chinese Academy of Sciences (CAS), the National Center for Supercomputing Applications (NCSA) at the University of Illinois at Urbana-Champaign (UIUC) and the Jülich Supercomputing Centre at Forschungszentrum Jülich. CHANGES provides a high-level platform to discuss latest trends in supercomputing, sophisticated information techniques and interdisciplinary applications. It will not only consider the issues of the partner institutions, but will also take national topics into account. The workshop series was founded in October 2011 and the first workshop took place in September 2012 in Jülich, where the topic High-Performance Computing "Performance Tools, Performance Modelling, and Algorithms" was covered.

came together by invitation and discussed latest data-related challenges in their research fields. In addition to several computer and domain scientists from JSC, Prof. Katrin Amunts from the Institute of Neuroscience and Medicine (INM-1) and Düsseldorf University, Prof. Henning Gast and Prof. Torsten Kuhlen from RWTH Aachen University were members of the German delegation. They talked about Big Data problems in their specific field of research. Besides the presentations, the workshop provided a forum for bi- and trilateral cooperations on student exchanges and mutual research projects. Special emphasis was placed on the question, how to select promising common research projects and how to secure funding for them.

• Norbert Attig

Jülich Supercomputing Centre (JSC)

Laboratory Experiments on Crowd Dynamics

At major events, accidents are rather caused by the dynamics of the crowd than external threads. Currently there are no generally accepted standards to support the actors, which are involved in the planning and realization of such events (organizers, authorities and organizations with security tasks, local authorities) to perform their tasks. The BMBF-funded project "Basigo – Bausteine für die Sicherheit von Großveranstaltungen" has set itself the goal of developing safety and security modules for major events, in order to fill this gap.

As part of this project the world's largest laboratory experiments on the dynamics in crowds took place on 19 - 22 June 2013. Over four days, more than 2,000 volunteers participated in the experiments in a 10,000 m² hall of Messe Düssedorf. Responsible for the experiments was the Jülich Supercomputing Centre at the Forschungszentrum Jülich in cooperation with the Institute for Media Research at the University of Siegen and the IBIT GmbH - International Training Centre for Crowd Safety Management.

To understand the emergence of critical conditions, in about 30 experiments with a total of 200 runs the density within the crowd was increased up to six persons per square meter. With a specially developed video technology based on a grid of 26 industrial cameras the trajectories of each person were recorded in high accuracy. The detection and distinction of individuals was carried out by Reed-Solomon codes. Each of the volunteers wore a unique cap where the cyclic, error-correcting code was printed on.

With the obtained data, e. g. the following questions will be investigated:

- When does pedestrian traffic jam occur (at narrow points, cross-ways and two-way traffic)?
- When does a jam induce a dangerous situation?
- What instruments (e. g. railings, barriers etc.) lead to an optimal use of the space?
- How can the crowd be optimally informed?

The department "Civil Security and Traffic" at the Jülich Supercomputing Centre has performed many smaller experiments on pedestrian dynamics in recent years. Based on these studies the evacuation assistant for a multi-functional arena was already developed in the BMBF project "Hermes". The core of this evacuation assistant is a real-time simulation of pedestrian flows; the modeling of the Generalized Centrifugal Force Model (GFCM) and its implementation on the parallel computer was carried out at the Jülich Supercomputing Centre. With the current experiments in BaSiGo the already examined scenarios could be upgraded at conditions of high densities of people, as they occur for example during big events or at train stations. The aim is to define reliable parameters for the planning and approval of such events and facilities, and accordingly develop the models for the simulation of pedestrian flows.

For more information:
http://www.basigo.de
http://www.fz-juelich.de/ias/jsc/cst

• Stefan Holl
• Armin Seyfried

Jülich Supercomputing Centre (JSC)

JSC Guest Student Programme on Scientific Computing 2013

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. To fulfill this mission, JSC hosts training activities and educational programmes for scientific computing on a regular basis. One of these is the Guest Student Programme (GSP) with a duration of ten weeks each summer.

The participants receive extensive training on cutting edge hardware as well as HPC-related software and algorithms. The acquired theoretical skills are turned into hands-on knowledge during supervised work on contemporary 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 also return to JSC as PhD candidates with focus on highly parallel applications.

Since the start of the GSP in 2000, a total of 146 students had the opportunity to join scientists from JSC and other institutes at Forschungszentrum Jülich. Now in it's 14th year, the GSP is more attractive to students around the world than ever. Over 50 candidates from 19 countries applied, covering a wide range of disciplines, including but not limited to mathematics, physics, chemistry, computer science, engineering, biomedicine and earth science. From this field of strong competitors thirteen students were invited to attend the programme.

This year's GSP took place from August 5th – October 11th. It was supported by CECAM (Centre Européen de Calcul Atomique et Moléculaire) and sponsored within the IBM University programme.

It started with a ten days course on parallel programming up to advanced level, covering several techniques, ranging from the usage of MPI on distributed-memory cluster systems to GPGPU programming with CUDA and OpenCL as well as OpenMP. Equipped with the required knowledge the participants were ready to focus on 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, particle-in-cell methods and safety research, as well as fundamental research in elementary particle physics and mathematical algorithms. The main platforms for code development and simulation were the multi-purpose cluster JUROPA, the GPU system JUDGE and the leadership Blue Gene/Q system JUQUEEN.

During the concluding two-day colloquium, the participants presented their accomplishments 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 4, 2014. It will be officially announced in January 2014 and is open to students from natural sciences, engineering, computer science and mathematics, who have already received the Bachelor but not yet the Master degree. The application deadline is April 25, 2014. Additional information, as well as the proceedings of the previous years, is available online at www.fz-juelich.de/jsc/gsp.

• Ivo Kabadshow
• Sven Strohmer

Jülich Supercomputing Centre (JSC)

Activities

High-Q Club – The highest scaling Codes on JUQUEEN

The trend in supercomputing towards much higher core counts seems inevitable. As a consequence, the users of supercomputing resources once again need to adapt their strategies in programming for these architectures. Whereas in the past we have for example seen vector machines come and go, the MPI standard has become established as the most widely used programming model. With multicore architectures, however, shared memory programming is making something of a comeback.

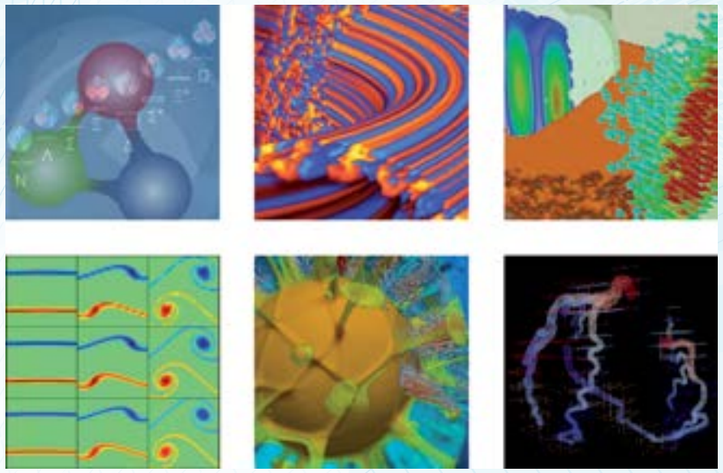

The latest supercomputer in the Jülich Supercomputing Centre, JUQUEEN, has 16 CPU cores per node and supports up to 64 hardware threads. To help our users migrate their codes and get the most performance for their algorithms, we organised a first porting and tuning workshop earlier this year [1]. Following this workshop and to promote the idea of exascale capability computing, we have established the High-Q Club, a showcase for codes able to utilise the entire 28-rack BlueGene/Q machine at JSC. The club members comprise a collection of the highest scaling codes on JUQUEEN, through which we intend to encourage other developers to invest in tuning and scaling their codes. We want our users to show that they are capable of using all 458,752 cores, and for example more than 1 million concurrent threads on JUQUEEN.

The diverse membership of the High-Q Club establishes that it is possible to scale real applications to the complete JUQUEEN using a variety of programming languages and parallelisation models, demonstrating individual approaches to reach that goal. High-Q status thus marks an important milestone in application development towards future HPC systems that envisage even higher core counts.

To qualify for membership, developers should submit evidence of the scalability of their codes across all available cores. While we currently do not set a strict minimum efficiency, we do expect the codes to profit from additional cores with an increase in speed. The benchmark used should also be as close as possible to a production scenario: trivial kernels or libraries will not be accepted.

Current members of the High-Q Club are:

dynQCD

dynQCD is a code for simulations in the field of lattice quantum chromodynamics that can be used for different fermion actions. The code is developed at the University of Wuppertal and the Simulation Laboratory for Nuclear and Particle Physics at JSC and is written in C using pthreads. The BG/Q version in particular makes use of a 4D torus and uses low level SPI calls to the network hardware instead of the MPI library.

Gysela

Gysela is a GYrokinetic SEmi-LAgrangian code for plasma turbulence simulations developed at CEA Cadarache. It is for example used in simulations of the electrostatic branch of the ion tem-perature gradient turbulence in tokamak plasmas. Gysela is written in Fortran90 and C and uses MPI, OpenMP and pthreads.

PEPC

The Pretty Efficient Parallel Coulomb solver is used for n-body simulations developed within the Simulation Laboratory for Plasma Physics at JSC. PEPC is not restricted to a specific force law or physical problem and for exampled used for beam-plasma interaction, vortex dynamics, gravitational interaction or molecular dynamics simulations. The code is written in Fortran2003 and C, making use of MPI, OpenMP and pthreads as programming models.

PMG+PFASST

PMG+PFASST combines a parallel multigrid solver with a time parallel approximation scheme to solve ODEs with linear stiff terms. The two parts have been developed at the Lawrence Berkeley National Lab (PFASST) and the University of Wuppertal (PMG) and have been coupled to one application by developers from the Cross-sectional team Mathematical Methods and Algorithms at JSC and Universita della Svizzera italiana. PMG+PFASST is written in Fortran2003 and C with MPI and pthreads.

Terra-Neo

Terra-Neo is used for modeling earth mantle dynamics and is developed specifically for the upcoming heterogeneous exascale computers by using an advanced architecture-aware co-design approach. The development team is built from members of Ludwig-Maximilians-Universität München, Universität Erlangen-Nürnberg, Regionales Rechenzentrum Erlangen and Technische Universität München. They use C++ and Fortran with MPI and OpenMP for the Terra-Neo framework.

waLBerla

waLBerla is the widely applicable Lattice Boltzmann solver from Erlangen developed by the Universität Erlangen-Nürnberg. Originally, the waLBerla framework has been centered around the Lattice-Boltzmann method for the simulation of fluid scenarios but in the meantime evolved to a code that is also suitable for a wide range of applications based on structured grids. It is developed in C++ and builds on MPI and OpenMP as well as CUDA and OpenCL for other architectures.

The updated list of members and more information is available at:

http://www.fz-juelich.de/ias/jsc/high-q-club

Reference

[1] **Brömmel, D.**
First JUQUEEN Porting and Tuning Workshop, Innovatives Supercomputing in Deutschland, Vol. 1, No. 1, 2013

• Dirk Brömmel
• Paul Gibbon

Jülich Supercomputing Centre (JSC)

Jülich Supercomputing Centre contributes to visionary "Human Brain Project"

HBP
The Human Brain Project

The goal of the Human Brain Project (HBP) [1] is to gather all existing knowledge about the human brain and to reconstruct the brain, piece by piece, in multiscale models and supercomputer-based simulations of these models. The resulting "virtual brain" offers the prospect of a fundamentally new understanding of the human brain and its diseases and of novel, brain-like computing technologies.

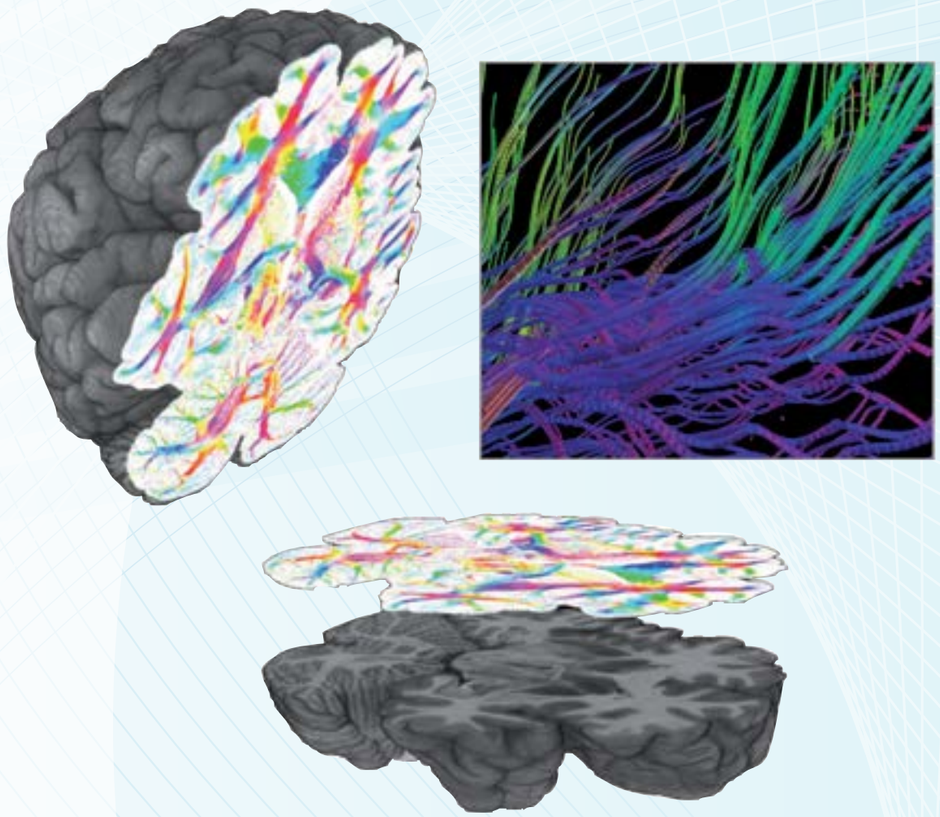

Figure 1: 3D reconstructed block face volume of a human brain with corresponding 3D fiber orientation maps obtained with Polarized Light Imaging (PLI). The pane in the upper right shows tracked fiber tubes. (Source: Axer, Amunts et al., INM, Forschungszentrum Jülich)

The HBP will develop a research infrastructure consisting of six so-called ICT (Information & Communication Technology) Platforms, dedicated respectively to Neuroinformatics, Medical Informatics, Brain Simulation, Neuromorphic Computing, Neurorobotics, and High-Performance Computing. Together, these platforms will make it possible to federate neuroscience data from all over the world, to integrate the data in unifying models and simulations of the brain, to validate the results against empirical data from biology and medicine, and to make them available to the world scientific community. The resulting knowledge on the structure and connectivity of the brain will open up new perspectives for the development of "neuromorphic" computing systems incorporating unique characteristics of the brain such as energy-efficiency, fault-tolerance and the ability to learn. The HBP's models and simulations will enable researchers to carry out in silico experiments on the human brain that cannot be done in vivo for practical or ethical reasons.

Starting from October 2013, the European Commission supports this vision through its FET (Future & Emerging Technologies) Flagship Initiative [2]. The 2.5-year ramp-up phase of the project (until March 2016) is funded by the EU's 7th Framework Programme. It will be followed by a partially overlapping operational phase under the EU's next Framework Programme, Horizon 2020. Federating more than 80 European and international research institutions, the HBP as a whole is planned to last ten years and estimated to cost one billion Euros. Additional partners will join the HBP consortium from 2014 by way of an open Competitive Call Programme [3]. The project is coordinated by Prof. Henry Markram from the Ecole Polytechnique Fédérale de Lausanne (EPFL), Switzerland.

Forschungszentrum Jülich's expertise and infrastructure in both neuroscience (e.g. in the areas of brain mapping tools, brain activity, large-scale neural network simulations) and supercomputing will make a major contribution to the project. Jülich Supercomputing Centre (JSC) leads the HBP's High-Performance Computing Platform subproject that will provide the supercomputing hard- and software capabilities necessary to simulate cellular brain models of the size of a complete human brain. JSC's central task will be to develop and host the HBP Supercomputer, the project's main production system, which will be built in stages to finally reach exascale performance.

References

[1] http://www.humanbrainproject.eu

[2] http://cordis.europa.eu/fp7/ict/programme/fet/flagship/home_en.html

[3] http://www.humanbrainproject.eu/participate/competitive-calls-programme

• Thomas Lippert
• Boris Orth

Jülich Supercomputing Centre (JSC)

Activities

Traffic and Granular Flow Conference celebrates 10th Edition by returning to Jülich

The conference Traffic and Granular Flow '13 (TGF'13) [1] brought together 105 international researchers from different fields ranging from physics to computer science and engineering to discuss the latest developments in traffic-related systems.

For its tenth edition, the TGF celebratory conference returned to the location of the very first conference held in 1995 at Forschungszentrum Jülich in Germany. Prof. Achim Bachem, one of the organizers of the first TGF and now chairman of the Board of Directors of Forschungszentrum Jülich wrote in his greetings: "I am pleased to see that after so many years the field of traffic and granular flow is still progressing and that numerous problems could be solved by new facilities." But he also points out that "we are facing plenty of new challenges in these research fields".

Originally conceived to facilitate new ideas by considering the similarities of traffic and granular flow, TGF'13 now covers a broad range of topics related to driven particles and transport systems. Besides the classical topics of granular flow and highway traffic, its scope includes data transport (Internet traffic), pedestrian and evacuation dynamics, intercellular transport, swarm behavior and collective dynamics of other biological systems. Recent progress in modelling, computer simulation and phenomenology was presented, and prospects for applications, for example to traffic control, were discussed. The conference intends to explore the interrelations between the above-mentioned fields and offers the opportunity to stimulate interdisciplinary research. This year the most prominent topic was pedestrian dynamics followed by vehicular traffic. Solely in the field of pedestrian dynamics, 47 talks were given and posters were presented. The book of abstracts can be found under [1].

Reference

[1] http://www.tgf13.de/

• Maik Boltes
• Armin Seyfried

Jülich Supercomputing Centre (JSC)

UNICORE Summit 2013

The UNICORE Summit is a unique opportunity for UNICORE [1] users, developers, administrators, researchers, service providers, and managers to meet. This year, it has been held as a satellite event at the ISC Conference in Leipzig on 18 June 2013. About 30 researchers from Germany, Italy, Switzerland, the Netherlands, Poland, Russia, Belarus, and the United States participated in the UNICORE Summit 2013 [2]. The goal of the UNICORE Summit is to exchange and share experiences, new ideas, and latest research results on all aspects of UNICORE. Since the first Summit in 2005, the organisers have received and reviewed a significant amount of distinguished contributions.

The event was opened with the keynote speech "Quo vadis EGI Clouds?" delivered by Michel Drescher, the EGI Technical Manager, who presented the EGI Federated Cloud Infrastructure. His talk was well-received by the attendees and initiated a lot of discussions about UNICORE integration with cloud infrastructures.

Recent developments in integration of applications from community projects, interoperability use cases, security aspects, virtualization techniques, performance evaluation, experiences from end users and administrators, data management, the UNICORE web portal as well as new ideas and concepts and related topics were highlighted in the following talks and demonstrations. Those selected and presented topics guaranteed lively discussions about the state-of-the art and the future of UNICORE, Grids, and distributed computing in general.

The slides to the presentations can be found on the web at http://www.unicore.eu/summit/2013/schedule.php. Accepted contributions will be published in the IAS book series of Forschungszentrum Jülich GmbH.

Figure 1: Participants of the UNICORE Summit 2013.

Figure 2: Keynote by Michel Drescher.

References

[1] UNICORE Web Page: http://www.unicore.eu
[2] UNICORE Summit 2013 web page: http://www.unicore.eu/summit/2013/

• Valentina Huber
• Daniel Mallmann

Jülich Supercomputing Centre (JSC)

3D Show at the Pharma Forum: Simulation and Visualization of the Airflow in Cleanrooms

Figure 1: Laminar Airflow in the loading side of a lyophilizer.

Knowing the details of the airflow around sterile filling lines in clean rooms is essential for the design, testing and justification of today's cutting-edge pharmaceutical filling machines. To mention just two important examples: Even in very clean rooms dust is unavoidable but it should not be transported towards the aseptic areas where the liquid and powder products are bottled. And as some of these pharmaceutical substances are toxic the contaminated air must be exhausted properly through the ventilation slots.

The OPTIMA pharma GmbH in Schwäbisch Hall has more than 25 years of experience in producing and developing sterile filling machines and wants to meet the ever increasing requirements of their customers. Dipl.-Ing. Ralph Eisenschmid, process engineer at OPTIMA pharma, run a couple of airflow simulations at the High Performance Computing Center Stuttgart (HLRS). The customer of the HLRS is excited about the available techniques: "The simulations will support the whole CAE job and minimize trial and error cycles and reconstructions. And they can replace expensive smoke studies." He thanks Dipl.-Phys. Bärbel Große-Wöhrmann, HLRS, for providing support during the start-up phase.

The CFD simulations of the airflow revealed unknown details, but how to visualize the calculated streamlines in very complex geometries? "Virtual reality techniques can be applied in many areas and help to visualize complicated processes and involved structures in three dimensions", says Dr. Uwe Wössner, Head of the Visualization Department of the HLRS. "…"

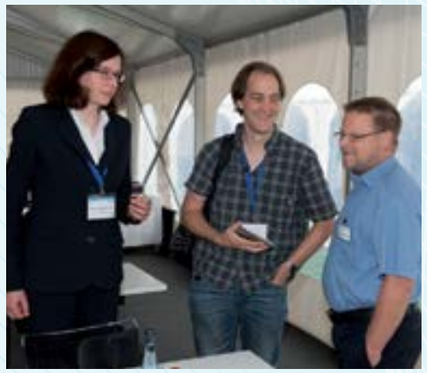

Figure 2: The project team.

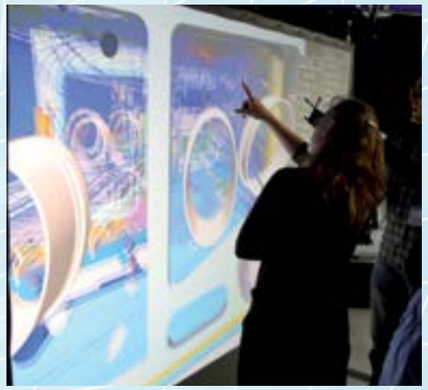

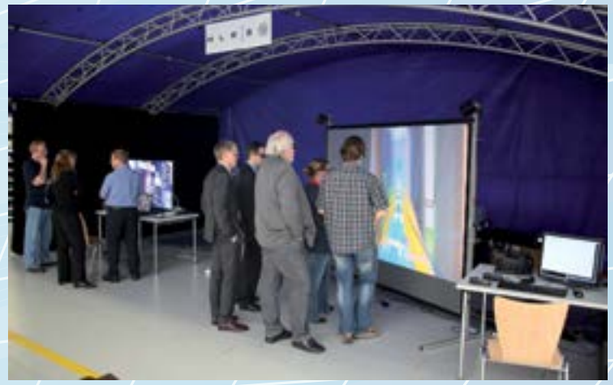

Figure 3: Virtual Reality presentation on the factory floor.

The latest developments in the pharmaceutical engine building were presented and discussed at the Pharma Forum 2013 in Schwäbisch Hall, a two-day meeting organized and hosted by OPTIMA pharma. Uwe Wössner, Ralph Eisenschmid and Bärbel Große-Wöhrmann presented the results of the collaboration between HLRS and OPTIMA pharma. An introductory presentation giving an overview over the project was followed by one of the highlights of the Pharma Forum: A 3D immersive Virtual Reality presentation of the airflow through a complete fill-finish line.

On a mobile backprojection system, the attendees where able to see a 1:1 model of the loading side, the freeze-dryer and the unloading side including animations of most relevant machine parts such as the loading table and a capper. In this model, we visualized the airflow in all machine parts and we could show how engineers can interactively place particle traces in order to find leeks or unwanted vortices anywhere in the machine.

A second installation demonstrated the use of Augmented Reality in visualizing airflow around a filling needle during the filling process. Both installations used the HLRS visualization software COVISE and it's VR component OpenCOVER.

Links

www.optima-pharma.com
www.hlrs.de

• Bärbel Große-Wöhrmann
• Uwe Wössner

University of Stuttgart, HLRS

The 17th HLRS-NEC Workshop on Sustained Simulation Performance

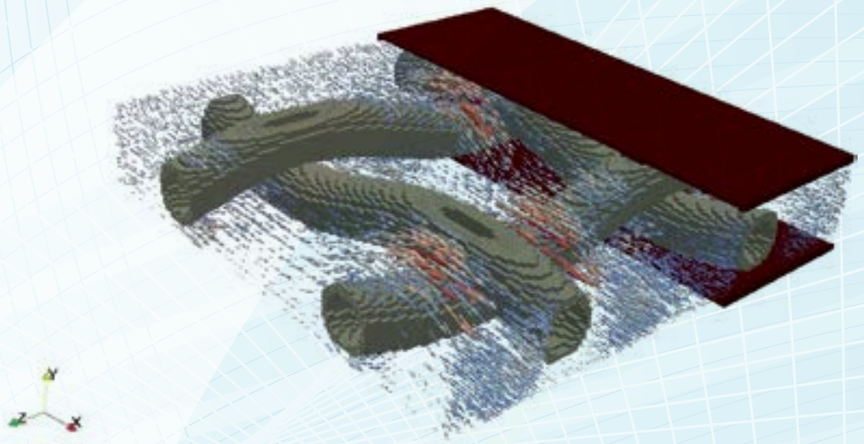

Figure 1: Simulation of a spacer.

The 17th Workshop on Sustained Simulation Performance, held on March 12th and 13th, 2013 at the NEC Corporation Headquarters in Tokyo, Japan, continued the series of workshops initiated by NEC and the HLRS in 2004. Once again, the covered topics drew a bow from leading edge supercomputing system development to the needs and results of real live applications.

The workshop was opened by Takahiro Hayashi form the Japanese Ministry of Education, Culture, Sports, Science and Technology who gave an outline of Japan's Policy on High Performance Computing. The following keynotes presented topics as diverse as an requirement analysis of HPC in engineering (Michael Resch, HLRS), the prospects of utilization of simulation in order to mitigate the damage by tsunamis and earthquakes (Fumihiko Imamura, Tohoku University), and the current state on future prospects of memory 3D DRAM technology (Jeong-Hwan Kwon, SKHynix).

The next sessions contained technical presentations on current Japanese Exascale projects, in particular on HPCI (High-Performance Computing Infrastructure), a project similar to the European PRACE presented by Kengo Nakajima (University of Tokyo), and the prospects of accelerators for Exascale computing (Mitsuhisa Sato, University of Tsukuba). The session was completed by presentations on the possible role of vector computers (Hiroaki Kobayashi, Tohoku University) and 3D memory technology (Mitsumasa Koyanagi, Tohoku University) in the post Petascale era, and an outline of NEC approach to Exascale computing by (Yukiko Hashimoto, NEC). This very intensive first day slowly wound down with a banquet with at the very top of the NEC headquarters' skyscraper. It remains unclear what left the Westerners more fascinated: the delicious Japanese food or the breath-taking view of night-time Tokyo.

The second day focussed on applications from various fields. However, it became very clear, that the 2011 Tohoku Earthquake and Tsunami let to an even stronger research effort into topics related to the usage of simulation for prediction and mitigation of such disastrous natural phenomena.

Two talks presented recent parallelization approaches to solving problems from the field of computational fluid dynamics (Sabine Roller, University of Siegen, and Satoru Yamamoto, Tohoku University). The next talks dealt with the challenges on the scaling of Exascale computing for small-scale problems (Uwe Küster, HLRS) and very large scale problems in Quantum Dynamics (Alejandro Muramatsu, University of Stuttgart).

The next sessions were devoted to usage of simulations for oceanic disaster prevention. Taro Arikawa (Port and Airport Research Institute) discussed how early warning systems combined with numerical models can predict the inundation of inhabited areas to high accuracy. Takane Hori (JAMSTEC) presented research on combination of numerical result and observations of the slip in plate boundaries for earthquake prediction. A related review on the usage of real time sensor data for numerical modelling was given by Yoshiyuki Kaneda (JAMSTEC).

The workshop was completed by presentations on performance characteristics of different CFD applications (Matthias Meinke, RWTH Aachen), the prospects of supercomputing in medical research (Joerg Bernsdolf, GRS), the benefits of usage of high-order numerical schemes on massively parallel systems (Jens Zudrop, University of Siegen), the modelling of global hydrological dynamics Keiko Takahashi (JAMSTEC), and observation data syntheses for climate studies Yoichi Ishikawa (JAMSTEC).

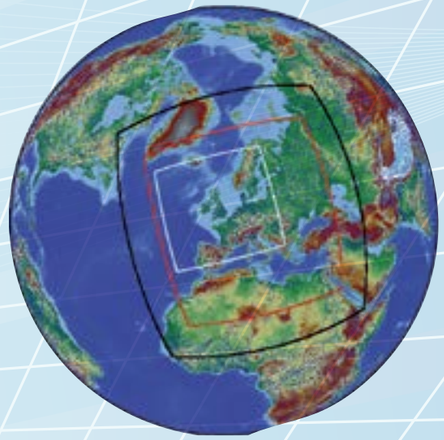

Figure 2: Illustration of the computational domain of various regional climate models with varying resolution. (Courtesy Dr. Warrach-Sagi, University Hohenheim)

We are looking forward the next instances of this workshop series at HLRS (October 2013) and Tohoku University (Spring 2014).

• José Gracia
• Ralf Schneider

University of Stuttgart, HLRS

ls1 mardyn - a Massively Parallel Molecular Simulation Code

ls1 mardyn is designed for massively parallel molecular dynamics (MD) simulations on modern supercomputers and currently holds the world record for the largest MD simulation of all times, with over four trillion particles.

In MD simulations, interactions between molecules are evaluated based on potentials. ls1 mardyn features pair potentials, point charges, dipoles, quadrupoles, Lennard-Jones sites and the Tersoff potential. Usually interaction potentials are evaluated only up to a given cut-off radius, as the contribution of far-apart pairs of molecules are relatively small and these missing contributions can be approximated with long-range correction schemes. The cut-off reduces the computational complexity of the force calculation to O(N). In order to reduce the complexity of the whole simulation to O(N), finding neighbouring molecules within the cut-off also has to be reduced to O(N). This is achieved in ls1 mardyn via the linked-cell algorithm.

However, in order to reach the extreme scales and efficiency in computing required to address length and time scales previously out of scope for simulations of highly dynamic and heterogenous systems, highly efficient methods for neighbour search and for dynamic load balancing were developed.

Neighbour Search

In highly dynamic systems neighbours have to be identified often, as the spatial arrangement changes rapidly. ls1 mardyn features an adaptive linked-cell algorithm. The basic linked-cell algorithm divides the simulation volume into equally sized cubic cells with edge length equalling the cut-off radius. Therefore, all interaction partners for any given molecule are no further than one cell away. Nonetheless, the neighbouring cells still contain molecules which are beyond the cut-off radius. Comparing the sphere within which interactions are evaluated to the volume of all neighbouring cells gives a ratio of 0.16. Thus, for a homogeneous distribution, only 16% of all molecule pairs are accepted for the force calculation.

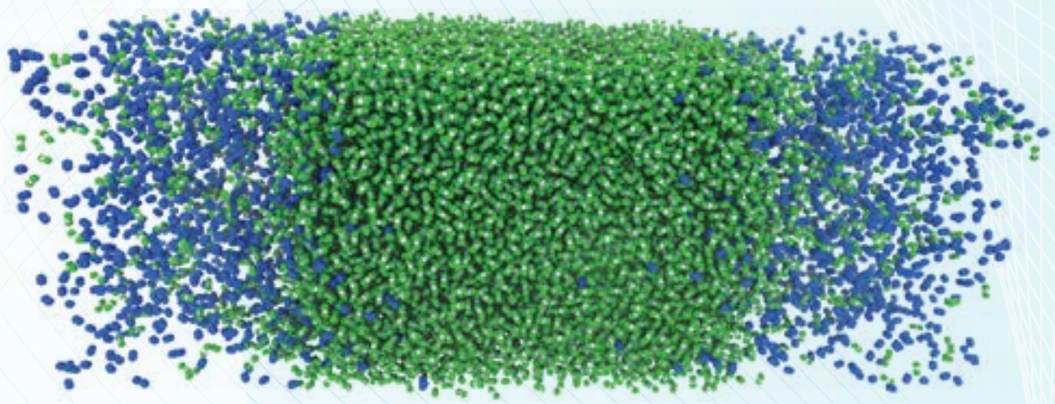

Molecular simulation of vapour-liquid phase boundaries using ls1 mardyn. Green: CO2, Blue: Oxygen, Temperature: -20° C.

Reducing the volume which needs to be searched can therefore save a lot of computing time. Using smaller cells with e.g. an edge length of half the cut-off improves the ratio to 0.27. However, this causes additional effort, as more cells need to be handled. The adaptive linked-cell of ls1 mardyn is capable of switching the cell size on the fly. For dense regions, where the time for computing the neighbour distances outweighs the cost of handling more cells, the cell size is thus reduced dynamically.

Dynamic Load Balancing

ls1 mardyn is parallelized using the domain decomposition. The simulation volume is divided into subvolumes, which are distributed to the available processing units. This method scales linearly with the number of molecules and is therefore well suited for large systems.

However, for heterogeneous scenarios (the molecules are distributed irregularly in space), the workload of equally sized subvolumes differs dramatically: it is proportional to the number of interactions and therefore grows quadratically with density. Simulations containing coexisting liquid and vapour phases, can easily vary in local density by a factor of > 100. Therefore, the workload for two subvolumes of equal size can differ by a factor of > 10,000.

In order to balance the load, the subvolumes must thus be chosen not with equal size, but equal load. In ls1 mardyn this is achieved by a kd-tree decomposition. The cells of the linked-cell algorithm are used as the basic volume units for which the load costs are determined. On the basis of the computational cost for each of the cells, the kd-tree decomposition recursively splits the simulation domain in two, along alternating dimensions, such that both sides have equal load, until the number of required subvolumes is reached. As the simulations are highly dynamic, this is repeated in regular intervals on the fly.

Scalability

Scalability studies were carried out with heterogeneous and homogeneous scenarios. Heterogeneous scenarios are very challenging - good scaling can be achieved for up to ~ 1,000 cores. Homogeneous scenarios are less challenging and show excellent scaling behaviour with ls1 mardyn. They allow for the utilization of entire state of the art supercomputers, as e.g. the Hermit system. For weak scaling, parallel efficiencies of over 90% can be reached on more than 100,000 cores.

Partners

The massively parallel simulation code ls1 mardyn is a joint development by

- TU Kaiserslautern
- University of Paderborn
- TU München
- HLRS

For more Informations:

- Martin Horsch: martin.horsch@mv.uni-kl.de
- Christoph Niethammer: niethammer@hlrs.de
- Colin Glass

University of Stuttgart, HLRS

Activities

GCS at ISC'13 – Review

Figure 1: Dr. Uwe Wössner (HLRS) gives a 3D-presentation to a delegation of international journalists.

The Gauss Centre for Supercomputing participated in the 2013 edition of ISC, held in Leipzig/Germany, yet again with a 64 m^2 booth on which its three member centres HLRS, LRZ and JSC showcased their latest research projects and presented sophisticated HPC tools. Countless like-minded HPC users, researchers, technology leaders, scientists, IT-decision makers as well as high tech media representatives stopped by to meet and talk with the directors of the three GCS centres Prof. Bode (LRZ), Prof. Lippert (JSC), the newly elected GCS Chairman of the Board, Prof. Resch of HLRS, and GCS Managing Director Dr. Müller, as well as with many other scientists of the three GCS centres.

Two GCS HPC systems amongst Top Ten of TOP500

The newly released TOP500 overview listing the most powerful supercomputers in the world once again delivered proof for Germany being a world-wide key player in HPC. GCS features the two most powerful supercomputers in all of Europe with JSC's HPC system JUQUEEN and LRZ's SuperMUC capturing positions in the Top Ten of said list. The newly coronated Number One is TH-IVB-FEP Cluster "Tianhe-2" (peak performance: 55 Petaflops) which is to be deployed at the National Supercomputer Center in Guangzho, China.

The 41th edition of the TOP500 (06/2013) revealed the following rankings for the three GCS systems:

- #7: JUQUEEN of JSC (IBM Blue Gene/Q system, peak performance: 5.9 Petaflops),
- #9: SuperMUC of LRZ (IBM iDataplex system, 3.2 Petaflops),
- #32: Hermit of HLRS (Cray XE6 system, 1.04 Petaflops).

Hermit furthermore holds second place on the list of most powerful supercomputers world wide used for industrial science and research activities (TOP500 06/2013, sub-list Industry).

GCS Booth Highlights

On the GCS booth, the three supercomputing centres presented their wide ranging HPC activities. HLRS concentrated on interactive 3D visualizations of simulation results on a tiled high-resultion 3D-display wall. Highlight of its many showcase examples was the visualization of the deformations in crystal structures. JSC exhibited a wide spectrum of scientific results obtained with its supercomputers JUGENE, JUQUEEN and JUROPA. In particular, JSC showcased LLview, the comprehensive interactive monitoring software for supercomputers developed in-house, in live demonstrations on supercomputers worldwide, while LRZ focused on highlighting their HPC-system SuperMUC and presented current science projects and research activities in 2D-videos.

Lots of attention enjoyed the Augmented Reality demo of science project CoolEmAll which aims at optimizing the energy efficiency of data centres. By moving a camera around the server, visitors could observe the temperature distribution and airflow right within a high density Christmann RECS server. They were able to see immediately where hot spots appear in operating condition and how these could be avoided through either geometry modification or adapted scheduling of workload, and they realized how this technology could be used to optimize their own servers and data centres.

ISC'13 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'13 Research Paper Sessions. This year, the award honored the paper titled „TUE, a new energy-efficiency metric applied at ORNL's Jaguar". Dr. Michael K. Patterson,

Figure 2: Gauss Award Winner Dr. Michael K. Patterson (Intel Corp.) with GCS Managing Director Dr. Claus Axel Müller (left) and GCS Chairman of the Board, Prof. Michael M. Resch.

Principal Engineer in the Intel Architecture Group of Intel Corp., had accepted the 2013 Gauss Award in the name of the Energy Efficient High Performance Computing Working Group (EE HPC WG), which was conceived and is led by Lawrence Berkeley National Laboratory to promote energy-efficient green computing best practices. (The paper can be downloaded at http://eetd.lbl.gov/sites/all/fi les/isc13_tuepaper.pdf)

• Regina Weigand

Gauss Centre for Supercomputing

Extreme Scaling Workshop at LRZ

July 9-11, 2013: Running Real World Applications on more than 130,000 Cores on SuperMUC

In July 2013, the Leibniz Supercomputing Center (LRZ) organized the first extreme scaling workshop on SuperMUC, the 3 PFLOP/s system consisting of 18 thin node islands with 147,456 Intel Sandy Bridge CPU cores. Prior to the workshop, the participants had to show that their code scales up to 4 islands (32,768 cores). Research groups from 14 international projects attained this goal and were invited to the LRZ for a three day workshop. During that time, the participants could test the scaling capabilities of their codes up to a maximum of 16 islands (two islands were continuing user operation). Application experts from the LRZ, Intel, and IBM were present during the workshop to resolve the performance optimization and tuning issues. New techniques like the fast MPI startup mechanism of large-scale special jobs were successfully executed on SuperMUC to reduce the startup time by a factor of 2-3. At the end of the third day, 6 applications were successfully running on 16 islands (131,072 cores), while the other 8 applications managed to run on 8 islands (65,536 cores).

Listed below are the name and description of the applications and the maximum number of islands the applications successfully ran (one island consists of 512 nodes with 16 physical cores each):

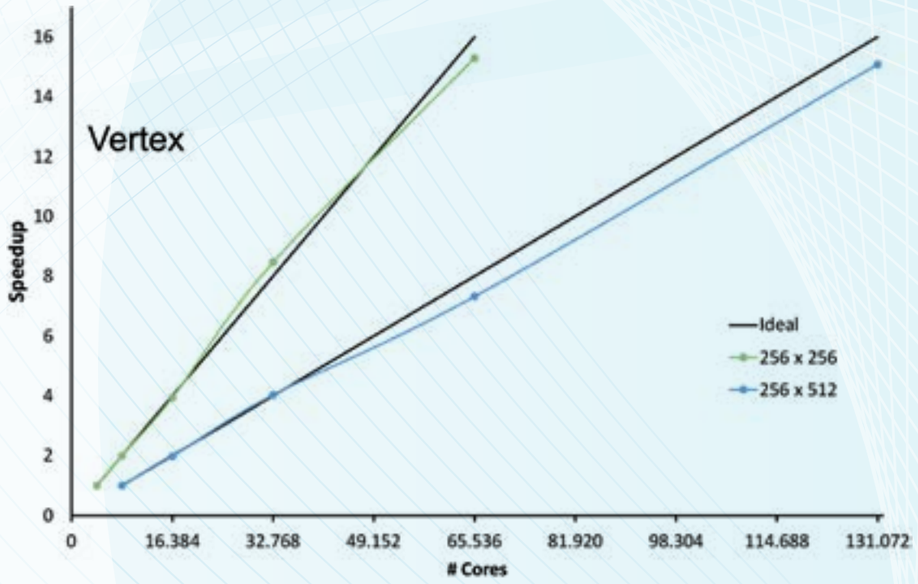

Figure 1: Scaling Plot for Vertex.

• BQCD (16 islands)

BQCD (Berlin Quantum ChromoDynamics program) is a hybrid MPI+OpenMP parallelized Monte-Carlo program for simulating lattice QCD with dynamical Wilson fermions. It allows for simulating 2 & 2 + 1 fermion flavors at a time.

• CIAO (8 islands)

CIAO solves the reacting Navier-Stokes equations in the low-Mach limit. It is a second order, semi-implicit finite difference code. It uses Crank-Nicolson time advancement and an iterative predictor corrector scheme. Spatial and temporal staggering is used to increase the accuracy of stencils. The Poisson equation for the pressure is solved by the multi-grid HYPRE solver. Momentum equations are spatially discretized with a second order scheme. Species and temperature equations are discretized with a fifth order WENO scheme.

• P-Gadget3-XXL (16 islands)

Highly optimized and fully MPI parallelized TreePM-MHD-SPH code for simulating cosmological structure formation. In its current version it also allows for an effective OpenMP parallelization within each MPI task.

• GROMACS (8 islands)

A versatile package to perform molecular dynamics, i.e. to simulate the Newtonian equations of motion for systems with hundreds to millions of particles.

• LAMMPS (16 islands)

LAMMPS (Large-scale Atomic/Molecular Massively Parallel Simulator) is a MPI parallelized numerics code for simulating molecular dynamics.

• Nyx (16 islands)

The code models dark matter as a system of Lagrangian fluid elements, or "particles," gravitationally coupled to an inviscid ideal fluid representing baryonic matter. The fluid is modeled using a finite volume representation in an Eulerian framework with block-structured AMR. The mesh structure used to evolve fluid quantities is also used to evolve the particles via a particle-mesh method. In order to more accurately treat hypersonic motions, where the kinetic energy is many orders of magnitude larger than the internal energy of the gas, Nyx uses the dual energy formulation, where both the internal and total energy equations are solved on the grid during each time step.

• Vertex (16 islands)

Neutrinoradiation hydrodynamics code, simulates from first principles the physical processes during the evolution of a supernova explosion. Therefore a MPI+OpenMP parallelized implementation of PPM hydrodynamics + a coupled ray-by-ray transport scheme is used.

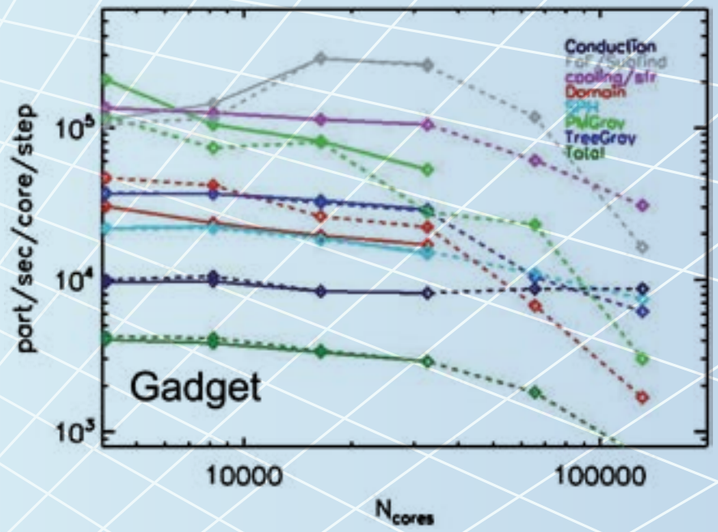

Figure 2: Scaling Plot for P-Gadget3-XXL.

• walBErla (16 islands)

A massively parallel software framework for simulating complex flows with the lattice Boltzmann method (LBM).

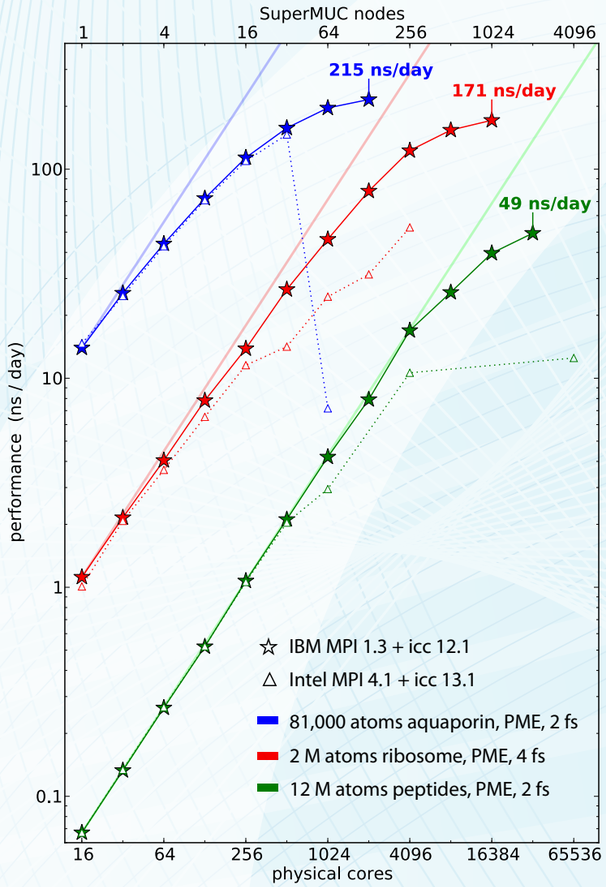

Figure 3: Scaling Plot for GROMACS.

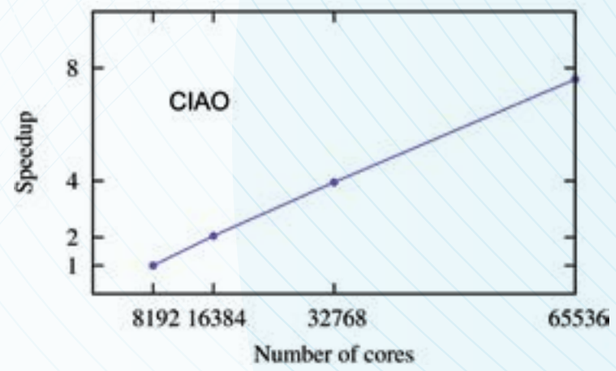

Figure 4: Scaling Plot for CIAO.

• APES (8 islands)

A suite of solvers for problems common in engineering applications. It is based on a common mesh representation library TreElM, and provides besides the solvers a mesh generation and post-processing tool. Currently there are mainly two different solvers developed within APES to implement two different numerical methods: Musubi and Ateles. Musubi implements a Lattice-Boltzmann scheme and can deal with various models. Besides the main incompressible Navier-Stokes model it is also capable of propagating passive scalars and multiple species in liquid or gas mixtures. It is mainly used for flow simulations that involve complex geometries, e.g. the flow through a channel filled by some obstacles for the simulation of electrodialysis. Another is the flow of blood through aneurysms and the simulation of the clotting effects. Ateles is a high order discontinuous Galerkin solver that is currently mainly deployed for the simulation of linear conservation laws, like the Maxwell equations.

• SeisSol (8 islands)

SeisSol is one of the leading codes for earthquake scenarios, in particular for simulating dynamic rupture processes and for problems that require discretization of very complex geometries. It allows multi-physics ground motion simulation for earthquake-engineering, including the complete dynamic rupture process and 3D seismic wave propagation with frequencies resolved beyond 5 Hz. The numerics in SeisSol are based on a higher-order discontinuous Galerkin discretization and an explicit time stepping following the arbitrary high order derivatives method. In combination with flexible unstructured tetrahedral meshes for spatial adaptivity, SeisSol shows excellent scalability and time to solution on recent supercomputing architectures.

• ExaML (4 islands)

Exascale Maximum Likelihood (ExaML) MPI application for inferring evolutionary trees of a set of species under the maximum likelihood criterion. It is an implementation of the popular RAxML search algorithm for partitioned multigene or wholegenome datasets.

• ICON (4 islands)

ICOsahedral Nonhydrostatic general circulation model is a joint development of the Max Planck Institute for Meteorology in Hamburg, and the Deutscher Wetterdienst. ICON is a next generation earth system model designed to simulate multiple scales of the atmosphere processes, enabling both climate simulations and numerical weather predictions. It provides the option to run locally nested highly refined resolutions, allowing simulations at a very fine scale. ICON is a non-hydrostatic global model with a local zoom function.

Performance Results

All projects were able to generate scaling curves up to 8 or 16 islands. From the preliminary data the following Flops rates have been obtained: 250 TFlop/s for VERTEX on 16 and 201 TFlop/s for Gromacs on 8 islands. The measured Flop rates for the complete application codes correspond to 10% or more of the peak performance of SuperMUC.

These results obtained in a short workshop can definitely compete with results reported from other Top10 supercomputers such as the K-computer and the Blue Waters system. They demonstrate the usability of SuperMUC for real world applications.

LRZ Extreme Scale Benchmark and Optimization Suite

Some of the participating projects agreed to provide their codes for an automised benchmarking and validation suite, based on the DEISA benchmark and Scalalife Validation suite (ref 1). The purpose of the package is automatic testing of the whole machine e.g. after system maintenance and identification of performance bottlenecks.

The LRZ is already planning a follow-up workshop in the near future, where the improvements and feedback from the experts will be tested.

- Momme Allalaen [1]
- Gurvan Bazin [2]
- Christoph Bernau [1]
- Arndt Bode [1]
- David Brayford [1]
- Matthias Brehm [1]
- Klaus Dolag [2]
- Jan Frederik Engels [4]
- Nicolay Hammer [1]
- Herbert Huber [1]
- Ferdinand Jamitzky [1]
- Anupam Karmakar [1]
- Carsten Kutzner [5]
- Andreas Marek [6]
- Carmen Navarrete [1]
- Helmut Satzger [1]
- Wolfram Schmidt [4]
- Philipp Trisjono [7]

[1] Leibniz Supercomputing Center

[2] Universitäts-Sternwarte, Ludwig-Maximilians Universität München, Germany

[3] Institute for Theoretical Physics, University of Zürich, Zürich, Switzerland

[4] Institute for Astrophysics, University of Göttingen, Göttingen, Germany

[5] Max Planck Institute for Biophysical Chemistry, Göttingen, Germany

[6] Rechenzentrum der Max-Planck-Gesellschaft am Max-Planck-Institut für Plasmaphysik, Garching, Germany

[7] Institut für Technische Verbrennung, RWTH Aachen, Germany

HLRS Scientific Tutorials and Workshop Report and Outlook

HLRS has installed Hermit, a Cray XE6 system with AMD Interlagos processors and 1 PFlop/s peak performance and extended with an XC30 system. We strongly encourage you to port your applications to these architectures as early as possible. To support such effort we invite current and future users to participate in the special **Cray XE6/XC30 Optimization Workshops**. With these courses, we will give all necessary information to move applications to this Petaflop system. The Cray XE6 provides our users with a new level of performance. To harvest this potential will require all our efforts. We are looking forward to working with our users on these opportunities. This four-day course in cooperation with Cray and multi-core optimization specialists is in spring 2014.

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 10 - 11, 2014.

These Cray XE6/XC30 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.

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 multi-threading with OpenMP. This course will be presented twice, in March 2014 at HLRS in Stuttgart and September 2014 at LRZ.

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 in spring 2014 in Stuttgart and in September/October in Siegen. The emphasis is placed on explicit finite volume methods for the compressible Euler equations. Moreover outlooks on implicit methods, the extension to the Navier-Stokes equations and turbulence modeling are given. Additional topics are classical numerical methods for the solution of the incompressible Navier-Stokes equations, aeroacoustics and high order numerical methods for the solution of systems of partial differential equations.

ISC and SC Tutorials

Georg Hager, Gabriele Jost, Rolf Rabenseifner: **Hybrid Parallel Programming with MPI & OpenMP.** Tutorial 9 at the International Supercomputing Conference, ISC'13, Leipzig, June 16-20. 2013.

Georg Hager, Jan Treibig, Gerhard Wellein: **Node-Level Performance Engineering.** Tutorial 2 at the International Supercomputing Conference, ISC'13, Leipzig, June 16-20. 2013.

Rolf Rabenseifner, Georg Hager, Gabriele Jost: **Hybrid MPI and OpenMP Parallel Programming.** Half-day Tutorial at Super Computing 2013, SC13, Denver, Colorado, USA, November 17-22, 2013.

Our general course on parallelization, the **Parallel Programming Workshop**, October 6 - 10, 2014 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.

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

We also continue our series of **Fortran for Scientific Computing** in December 2013 and in March 2014, 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.

With **Unified Parallel C (UPC) and Co-Array Fortran (CAF)** in April 2014, the participants will get an introduction of partitioned global address space (PGAS) languages.

In cooperation with Dr. Georg Hager from the RRZE in Erlangen and Dr. Gabriele Jost from Supersmith, the HLRS also continues with contributions on hybrid MPI & OpenMP programming with tutorials at conferences; see the box on the left page, which includes also a second tutorial with Georg Hager from RRZE.

In the table below, you can find the whole HLRS series of training courses in 2014. They are organized at HLRS and also at several other HPC institutions: LRZ Garching, NIC/ZAM (FZ Jülich), ZIH (TU Dresden), TUHH (Hamburg Harburg), and ZIMT (Siegen).

• Rolf Rabenseifner

University of Stuttgart, HLRS

2014 – Workshop Announcements
Scientific Conferences and Workshops at HLRS
12th HLRS/hww Workshop on Scalable Global Parallel File Systems (March/April 2014)
8th 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 (October 2014)
IDC International HPC User Forum (October 2014)
Parallel Programming Workshops: Training in Parallel Programming and CFD
Parallel Programming and Parallel Tools (TU Dresden, ZIH, February 24 - 27)
Introduction to Computational Fluid Dynamics (HLRS, March 31 - April 4)
Iterative Linear Solvers and Parallelization (HLRS, March 24-28)
Cray XE6/XC30 Optimization Workshops (HLRS, March 17 - 20) (PATC)
GPU Programming using CUDA (HLRS, April 7 - 9)
Open ACC Programming for Parallel Accelerated Supercomputers (HLRS, April 10 - 11) (PATC)
Unified Parallel C (UPC) and Co-Array Fortran (CAF) (HLRS, April 14 - 15) (PATC)
Scientific Visualisation (HLRS, April 16 - 17)
Parallel Programming with MPI & OpenMP (TU Hamburg-Harburg, July 28 - 30)
Iterative Linear Solvers and Parallelization (LRZ, Garching, September 15 - 19)
Introduction to Computational Fluid Dynamics (ZIMT Siegen, September/October)
Message Passing Interface (MPI) for Beginners (HLRS, October 6 - 7) (PATC)
Shared Memory Parallelization with OpenMP (HLRS, October 8) (PATC)
Advanced Topics in Parallel Programming (HLRS, October 9 - 10) (PATC)
Parallel Programming with MPI & OpenMP (FZ Jülich, JSC, December 1 - 3)
Training in Programming Languages at HLRS
Fortran for Scientific Computing (Dec 2 - 6, 2013 and Mar 10 - 14, 2014) (PATC)
URLs:
http://www.hlrs.de/events/
http://www.hlrs.de/training/course-list/
(PATC): This is a PRACE PATC course

Parallel Programming with MPI, OpenMP and PETSc

Date and Location
November 25 - 27, 2013
JSC, Forschungszentrum Jülich

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

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

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

Date and Location
November 28 - 29, 2013
JSC, Forschungszentrum Jülich

Contents
This course gives an overview of the supercomputers JUROPA and JUQUEEN. 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.

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

Node-Level Performance Engineering

(PATC course)

Date and Location
December 03 - 04, 2013
LRZ Building,
University Campus Garching,
near Munich, Boltzmannstr. 1

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:

- Introduction
- Practical performance analysis
- Microbenchmarks and the memory hierarchy
- Typical node-level software over heads
- Example problems:
  - The 3D Jacobi solver
  - The Lattice-Boltzmann Method
  - Sparse Matrix-Vector Multiplication
  - Backprojection algorithm for CT reconstruction
- Energy & Parallel Scalability.

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

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

Fortran for Scientific Computing

(PATC course)

Dates and Location
December 02 - 06, 2013 and
March 10 - 14, 2014
Stuttgart, HLRS

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.

Web Page
www.hlrs.de/training/course-list

Second JUQUEEN Porting and Tuning Workshop

(PATC course)

Date and Location
February 03 - 05, 2014
JSC, Forschungszentrum Jülich

Contents
The Blue Gene/Q petaflop supercomputer JUQUEEN marks another 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 efficiency. This course is a PATC course (PRACE Advanced Training Centres).

Web Page
http://www.fz-juelich.de/ias/jsc/events/juqueenpt14

Programming with Fortran

Dates and Locations
February 03 - 07, 2014
LRZ Building, University campus Garching near Munich.

Contents
This course is targeted at scientists with little or no knowledge of the Fortran programming language, but needing 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.

Prerequisites
Course participants should have basic UNIX/Linux knowledge (login with secure shell, shell commands, basic programming, vi or emacs editors).

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

Parallel Programming with MPI, OpenMP, and Tools

Date and Location
February 24 - 27, 2014
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.

Web Page
www.hlrs.de/training/course-list

Parallel Programming of High Performance Systems

Dates and Location
March 10 - 14, 2014
RRZE building, University campus Erlangen, Martensstr. 1: 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 Munich, Jülich and Stuttgart.

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

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

Cray XE6/XC 30 Optimization Workshop

(PATC course)

Date and Location
March 17 - 20, 2014
Stuttgart, HLRS

Contents
HLRS installed Hermit, a Cray XE6 system with AMD Interlagos processors and a performance of 1 PFlop/s. We strongly encourage you to port your applications to the new architecture as early as possible. To support such effort we invite current and future users to participate in special Cray XE6/XC30 Optimization Workshops. With this course, we will give all necessary information to move applications from the current NEC SX-9, the Nehalem cluster, or other systems to Hermit. Hermit provides our users with a new level of performance. To harvest this potential will require all our efforts. We are looking forward to working with our users on these opportunities. From Monday to Wednesday, specialists from Cray will support you in your effort porting and optimizing your application on our Cray XE6. On Thursday, Georg Hager and Jan Treibig from RRZE will present detailed information on optimizing codes on the multicore AMD Interlagos processor. Course language is English (if required).

Web Page
www.hlrs.de/training/course-list

Iterative Linear Solvers and Parallelization

Dates and Location
March 24 - 28, 2014
Stuttgart, HLRS

September 15 - 19, 2014
Garching, LRZ

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.

Web Page
www.hlrs.de/training/course-list

Eclipse: C/C++/Fortran programming

Date and Location
March 25, 2014
LRZ Building, University campus 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.

Prerequisites
Course participants should have basic knowledge of the C and/or C++/Fortran programming languages.

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

Introduction to Computational Fluids Dynamics

Date and Location
March 31 - April 04, 2014
Stuttgart, HLRS

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

Web Page
www.hlrs.de/training/course-list

Advanced Topics in High Performance Computing

(PATC course)

Date and Location
March 31 - April 03, 2014
LRZ Building, University campus 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 hands-on sessions.
Day 1
Intel tools: MPI tracing and Checking
Intel tools: OpenMP performance and correctness.
Day 2
Parallel I/O with MPI IO
Performance analysis with Scalasca.
Day 3
Tuning I/O on LRZ's HPC systems.
Portability of I/O: Binary files NetCDF HDF5.
Day 4
PGAS programming with coarray Fortran and Unified Parallel C.
PGAS hands on session.

Prerequisites
Good MPI and OpenMP knowledge as presented in the course "Parallel programming of High Performance Systems" (see above).

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

GPU Programming using CUDA

Date and Location
April 07 - 09, 2014
Stuttgart, HLRS

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.

Web Page
www.hlrs.de/training/course-list

GPU Programming

(PATC course)

Date and Location
April 07 - 09, 2014
JSC, Forschungszentrum Jülich

Contents
Many-core programming is a very dynamic research area. Many scientific applications have been ported to GPU architectures during the past four years. We will give an introduction

to CUDA, OpenCL, and multi-GPU programming using examples of increasing complexity. After introducing the basics the focus will be on optimization and tuning of scientific applications.
This course is a PATC course (PRACE Advanced Training Centres).

Web Page
http://www.fz-juelich.de/ias/jsc/events/gpu

Open ACC Programming Cray XK

(PATC course)

Date and Location
April 10 - 11, 2014
Stuttgart, HLRS

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 directive-based OpenACC programming model whose multi-vendor 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.

Web Page
www.hlrs.de/training/course-list

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

(PATC course)

Date and Location
April 14 - 15, 2014
Stuttgart, HLRS

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

Web Page
www.hlrs.de/training/course-list

Scientific Visualization

Date and Location
April 16 - 17, 2014
Stuttgart, HLRS

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.

Web Page
www.hlrs.de/training/course-list

Intel MIC&GPU Programming Workshop

(PATC course)

Date and Location
April 28 - 30, 2014
LRZ Building, University campus Garching, near Munich.

Contents
With the rapidly growing demand for computing power new accelerator based architectures have entered the world of high performance computing since around 5 years. Particularly GPGPUs have recently become very popular, however programming GPGPUs using programming languages like CUDA or OpenCL is cumbersome and error-prone. Beyond introducing the basics of GPGPU-porogramming, we mainly present OpenACC as an easier way to program GPUs using OpenMP-like pragmas. Recently Intel developed their own Many Integrated Core (MIC) architecture which can be programmed using standard parallel programming techniques like OpenMP and MPI. In the beginning of 2013, the first production-level cards named Intel Xeon Phi came on the market. The course discusses various programming techniques for Intel Xeon Phi and includes hands-on session for both MIC and GPU programming. The course is developed in collaboration with the Erlangen Regional Computing Centre (RRZE) within KONWIHR.
Each day is comprised of approximately 5 hours of lectures and 2 hours of hands-on sessions.

Prerequisites
Good working knowledge of at least one of the standard HPC languages: Fortran 95, C or C++. Basic OpenMP and MPI knowledge useful.

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

Advanced GPU Programming

Date and Location
May 05 - 06, 2014
JSC, Forschungszentrum Jülich

Contents
Today's computers are commonly equipped with multicore processors and graphics processing units. To make efficient use of these massively parallel compute resources advanced knowledge of architecture and programming models is indispensable. This course focuses on finding and eliminating bottlenecks using profiling and advanced programming techniques, optimal usage of CPUs and GPUs on a single node, and multi-GPU programming across multiple nodes.

Web Page
http://www.fz-juelich.de/ias/jsc/events/advgpu

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

Date and Location
May 19 - 20, 2014
JSC, Forschungszentrum Jülich

Contents
This course gives an overview of the supercomputers JUROPA and JUQUEEN. 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.

Web Page
http://www.fz-juelich.de/ias/jsc/events/sc-may

Parallel I/O and Portable Data Formats

(PATC course)

Date and Location
May 21 - 23, 2014
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).

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

Authors

Momme Allalen
momme.allalen@lrz.de

Fakher F. Assaad
assaad@physik.uni-wuerzburg.de

Norbert Attig
n.attig@fz-juelich.de

Christian Baczynski
baczynski@uni-hd.de

Gurvan Bazin
gurvan.bazin@gmail.com

Florian Berberich
f.berberich@fz-juelich.de

Christoph Bernau
christoph.bernau@lrz.de

Arndt Bode
arndt.bode@lrz.de

Maik Boltes
m.boltes@fz-juelich.de

David Brayford
david.brayford@lrz.de

Matthias Brehm
matthias.brehm@lrz.de

Dirk Brömmel
d.broemmel@fz-juelich.de

Alexandru Calotoiu
a.calotoiu@grs-sim.de

Alexey Cheptsov
cheptsov@hlrs.de

Paul Clark
p.clark@uni-heidelberg.de

Tanja Cleese
tanja.clees@scai.fraunhofer.de

Holger Dachsel
h.dachsel@fz-juelich.de

Ulrich Detert
u.detert@fz-juelich.de

Klaus Dolag
kdolag@mpa-garching.mpg.de

Jan Frederik Engels
mail@jfengels.de

Dietmar Erwin
d.erwin@fz-juelich.de

Jochen Fröhlich
jochen.froehlich@tu-dresden.de

Andrea Gatto
andreag@mpa-garching.mpg.de

Paul Gibbon
p.gibbon@fz-juelich.de

Philipp Girichidis
philipp@girichidis.com

Colin Glass
glass@hlrs.de

Simon Glover
glover@uni-heidelberg.de

Stefan Gottlöber
sgottloeber@aip.de

Jose Gracia
gracia@hlrs.de

Bärbel Große-Wöhrmann
woehrmann@hlrs.de

Carla Guillen Carias
carla.guillen@lrz.de

Nicolay Hammer
nicolay.hammer@lrz.de

Tim Hender
tim.hender@ccfe.ac.uk

Wolfram Hesse
wolfram.hesse@lrz.de

Steffen Heß
shess@aip.de

Torsten Hoefler
htor@inf.ethz.ch

Martin Hohenadler
martin.hohenadler@
physik.uni-wuerzburg.de

Stefan Holl
st.holl@fz-juelich.de

Frank Holzäpfel
frank.Holzaepfel@dlr.de

Valentina Huber
v.huber@fz-juelich.de

Ferdinand Jamitzky
ferdinand.jamitzky@lrz.de

Ivo Kabadshow
i.kabadshow@fz-juelich.de

Anupam Karmakar
anupam.karmakar@lrz.de

Tobias Kempe
tobias.kempe@tu-dresden.de

Francisco Kitaura
kitaura@aip.de

Ralf Klessen
klessen@uni-heidelberg.de

Bastian Koller
koller@hlrs.de

Carsten Kutzner
ckutzne@gwdg.de

Thomas Lippert
th.lippert@fz-juelich.de

Daniel Mallmann
d.mallmann@fz-juelich.de

Andreas Marek
amarek@rzg.mpg.de

Gerald Mathias
gerald.mathias@
physik.uni-muenchen.de

Bernd Mohr
b.mohr@fz-juelich.de

Thorsten Naab
naab@mpa-garching.mpg.de

Carmen Navarrete
carmen.navarrete@lrz.de

Christoph Niethammer
niethammer@hlrs.de

Boris Orth
b.orth@fz-juelich.de

Ludger Palm
ludger.palm@lrz.de

Thomas Peters
tpeters@physik.uzh.ch

Rolf Rabenseifner
rabenseifner@hlrs.de

Michael M. Resch
resch@hlrs.de

Helmut Satzger
helmut.satzger@lrz.de

Armin Seyfried
a.seyfried@fz-juelich.de

Wolfram Schmidt
schmidt@astro.physik.
uni-goettingen.de

Magnus Schwörer
magnus.schwoerer@
googlemail.com

Anton Stephan
anton.stephan@dlr.de

Sven Strohmer
s strohmer@fz-juelich.de

Godehard Sutmann
g.sutmann@fz-juelich.de

Jan Treibig
jan.treibig@rrze.fau.de

Philipp Trisjono
p.trisjono@itv.rwth-aachen.de

Thomas Ullmann
thomas.ullmann@mpibpc.de

Bernhard Vowinckel
bernhard.vowinckel@tu-dresden.de

Stefanie Walch
walch@mpa-garching.mpg.de

Regina Weigand
weigand@hlrs.de

Volker Weinberg
volker.weinberg@lrz.de

Torsten Wilde
torsten.wilde@lrz.de

Felix Wolf
f.wolf@grs-sim.de

Klaus Wolkersdorfer
k.wolkersdorfer@fz-juelich.de

Richard Wünsch
richard@wunsch.cz

If you would like to receive inside regularly,
send an email with your postal address to klank@hlrs.de
Web page: http://inside.hlrs.de/