


inside

SPRING 2015

inSiDE • Vol. 13 No.1

**Innovatives Supercomputing
in Deutschland**



Publishers

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

Editor

F. Rainer Klank, HLRS klank@hls.de

Design

Pia Heusel, HLRS heusel@hls.de

Fine dust particle distribution in the human lung during inspiration. The particles, tracked by a Lagrangian model, are transported down the trachea and distribute in the left and right main bronchi before they deposit in higher lung generations where they can cause severe damage to the respiratory system potentially leading to lung cancer.

Gonzalo Brito Gadeschi¹, Lennart Schneiders¹, Christoph Siewert^{1, 2} Andreas Lintermann¹, ³ Matthias Meinke¹, Wolfgang Schröder^{1, 3}

¹ Institute of Aerodynamics, RWTH Aachen University, Germany, ² Laboratoire Lagrange, OCA, CNRS, Nice, France, ³ SimLab "Highly Scalable Fluids & Solids Engineering", JARA-HPC, RWTH Aachen University, 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 introduce the new Chairman of the Board of GCS, Prof. Thomas Lippert. Prof. Lippert took over from Prof. Michael Resch in April and talks about his ideas and views on the future of High Performance Computing in an interview with inside. Unfortunately we also have to report sad news. Walter Nadler of the Jülich Supercomputing Centre passed away recently having served Gauss Centre for Supercomputing in a variety of functions over the last years. An obituary reflects on a life dedicated to High Performance Computing.

With a selection of applications we present the use of GCS supercomputers over the last months. The variety of applications shows that GCS systems are able to cover a wide spectrum of research topics and basically make supercomputing available for all scientific fields that rely on compute performance to break new barriers and achieve new results.

The section "Projects", on the other hand, shows how tightly GCS and its centres are enmeshed in the European and German scientific community when it comes to research in the field of High Performance Computing. We find a variety of project reports that range from basic research to applied research and industrial development. Most outstanding is the research conducted to guarantee a leading role for German and European research in the field of Exascale Computing.

High Performance Computing continues to be an important issue in the German research community. The significance of HPC for German science is emphasized by a new paper released by the Deutsche Wissenschaftsrat (German Science Council) on April 24 this year (Drs. 4488-15 at www.wissenschaftsrat.de/download/archiv/4488-15.pdf) addressing the question of the funding of HPC in Germany.

As usual, this issue includes information about events in supercomputing in Germany over the last months and gives an outlook of workshops in the field. Readers are invited to participate in these workshops which are now part of the PRACE Advanced Training Centres (PATCs) and hence are open to all European researchers.

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

Interview with Prof. Dr. Dr. Thomas Lippert



Prof. Lippert, you were elected chairman of the Board of the Gauss Centre for Supercomputing in April. How do you see the role of the Chairman of GCS?

Since its foundation in 2007, the Gauss Centre for Supercomputing (GCS) has become the most powerful national HPC infrastructure in Europe. My predecessors as chairmen of the GCS, Prof. Achim Bachem, Prof. Heinz-Gerd Hegering, and Prof. Michael Resch all have

significant merits as far as the success of this association is concerned. They have given GCS maximal stability and a reputation that is outstanding worldwide. Furthermore, the centres of the GCS have been working together ever closer within the last eight years; I am very proud to follow the role model of my colleagues as the Chair by building on their virtues and these past achievements.

Thanks to the formidable engagement of the German ministry for education and research (BMBF) on the one hand and the ministries in Baden-Württemberg, Bayern, and North Rhine-Westphalia on the other hand, between 2010 and 2016 the GCS was enabled (and still is) to provide Germany's HPC infrastructure of the highest national performance class, called Tier-1. Evidently, GCS has achieved one of its major objectives. In addition, GCS has almost completed its delivery of supercomputing capacity worth Euro 100 Mio. to the European association Partnership for Advanced Computing in Europe (PRACE) on behalf of the BMBF, so far the most significant contribution to Europe's Tier-0 HPC infrastructure. On top of it, GCS not only is commissioned to provide the most powerful HPC infrastructure in Germany and Europe,

its mission also comprises to serve a broad range of research and industrial activities in a variety of disciplines.

At this stage, I see it as a most important responsibility of the chairman to ensure that the GCS success story will continue and even can become a greater success in the future. This goal, on the one hand, requires leading the GCS safely into its next phase, starting in 2016, by deepening the deep trust GCS enjoys from its ministries. On the other hand, GCS has always been, as truly European partner, crucial for the progress of PRACE; the chairman must contribute his best to ensure that GCS stays being the stabilizing element in PRACE.

You mentioned the long-term collaboration between the GCS partners, which is now in its ninth year. What is the key success story of GCS from your point of view?

What was considered a tremendous challenge eight years ago – bringing together the three national centres to form a single organization – has turned out to become a most conducive research infrastructure of supercomputing for Germany and Europe that sub-

stantially has enhanced Germany's position and visibility in the field of computational science and engineering worldwide. It is quite interesting to observe the reason for such a successful long-term stability.

My view as to a GCS key success story more strongly reflects my computational scientist background rather than my computer science history – I am a practicing computational elementary particle physicist – and I think, the crucial step was a scientific one, namely the establishment of a rigorous GCS peer review process, joining the virtues of established procedures at the three centres. In fact, it is remarkable that GCS achieved to promote its peer-review procedure to become the template for the European PRACE peer review. On the organizational level, GCS has established a Joint Steering Executive Committee (Lenkungsausschuss) led by experts from science and engineering of highest reputation and supported by a team from the three GCS centres. This executive committee is responsible for the peer review process and the practical coordination between the centres. Our GCS evaluation practice promotes a continuously increasing level of quality of the proposals submitted, and, most importantly,

it attracts those scientific and engineering groups to GCS facilities with highest scientific claims.

Which changes should we expect from you?

We have gained lot of experience in the last years on data intensive supercomputing, being increasingly requested in a variety of fields like for instance terrestrial systems, neuroscience, climatology and engineering. Given the growing importance of Scientific Big Data Analytics (SBDA) for science and engineering, I assume that SBDA in science and industry will become a major component of high-end HPC activities worldwide, and both fields, data intensive simulations and SBDA, will grow. I am convinced that the GCS has the ability to become the trailblazer for peer reviewed SBDA in Germany and Europe and I will strongly try encouraging the GCS to engage in this field.

What are future challenges?

The HPC landscape is known to transform very rapidly. Every 3 to 5 years a new generation of HPC technology appears. The past and future challenge for GCS is to take our users with us, together with their highly evolved application codes, as we move towards exascale compute performances and data storages becoming established around 2025. We have to substantially expand our training efforts and our support activities

in order to help our communities keeping up with the exponentially growing demands on parallelism, scalability, hierarchical memory layers and interactivity.

Many people say that the US and Japan benefit substantially from having national HPC industries. Would such a European player enhance European science?

It was obvious in the early times of supercomputing that US scientists and engineers had some competitive advantage as early adopters of their homemade technologies. I think that since then the rest of the world, in all scientific and engineering fields, has been able to strongly improve its position due to excellence in application software, algorithms and development of tools. As example, materials sciences groups in Europe are probably in the lead compared to US due to their high expertise in molecular dynamics and ab-initio simulation codes. Of course, it is natural and important for the HPC centres to engage in co-design projects with German and European technology providers of HPC technology; a most welcome effect is that this engagement will foster the field not only in Europe but also as a whole, not least because high-end HPC technology producers and vendors more and more become globalized companies. On all accounts, GCS must be open to the best technology available worldwide, procured on the basis of accepted rules, to ensure its continued ability to offer

world-leading supercomputing systems and software to its users.

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

I hope that GCS II will have been launched by the same winning team as in the first round of GCS, that we can manage to realize the vision of an interacting HPC pyramid with the Gauss Alliance, and that GCS will have succeeded to rejuvenate its leadership in the next round of PRACE.

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

Obituary: Dr. Walter Nadler



On June 9, 2015, Dr. Walter Nadler passed away at the age of 61. Even to those close to him this came as an unexpected shock.

Dr. Walter Nadler was not only an established expert in simulating complex biological systems, he also excelled as a science manager, efficiently and effectively coordinating peer review procedures which had been established to grant Tier-1 and Tier-2 supercomputer resources to different user alliances. He was the driving force in implementing the GCS governance which

was agreed upon by the GCS member centres in 2012/13, where he acted as indispensable advisor to the chairmen of numerous compute time granting commissions.

Walter received his PhD in Physics in 1985 at Technische Universität München, supervised by Prof. Klaus Schulten. In the following years he worked at different institutions, including the California Institute of Technology, Wuppertal University and Michigan Technological University. In October 1996 he joined Forschungszentrum Jülich for the first time and became for two years a member of the HLRZ (later NIC) research group Complex Systems, led by Prof. Peter Grassberger. He collaborated scientifically with many groups, among them the Complexity Science Group at Calgary University and the group Cardiac MRT and Biophysics at Würzburg University. Results of the latter collaboration were honoured 2003 with the Helmholtz prize. In 2007 he joined Forschungszentrum Jülich for a second time and became member of the NIC research group "Computational Biology and Biophysics", headed by Prof. Ulrich Hansmann. Mid 2008 he was appointed head of the newly established NIC coordination office. Since then he focussed on introducing innovative enhancements to the NIC peer review and FZJ-internal allocation procedures. In 2012 an additional regional peer-review process for researchers from Forschungszentrum Jülich and RWTH Aachen Uni-

versity was approved (JARA-HPC partition), an extraordinarily complex undertaking which succeeded largely thanks to the skill and leadership demonstrated by Walter. One year later, the peer-review process for the newly established GCS Large-Scale Projects - defined by the GCS Governance - had to be integrated and implemented, again done by Walter. In the last few months he provided most valuable input to the redesign of JSC's application and peer-review server software.

A role that Walter will also be fondly remembered for both inside and outside JSC is as an enthusiastic, unflappable and fearsomely efficient booth manager: in this capacity he organized and coordinated the exhibition of the Jülich Supercomputing Centre at the European International Supercomputing Conference (ISC) and the US Supercomputing Conference (SC) every year since 2008. Under his supervision these exhibitions became a cornerstone with respect to the international visibility of the Jülich Supercomputing Centre.

The Jülich Supercomputing Centre and its partner institutions in the Gauss Centre for Supercomputing, the John von Neumann Institute for Computing, the JARA-HPC Vergabegremium and the Vergabekommission für Supercomputer-Ressourcen at Forschungszentrum Jülich will all sorely miss him.

- Thomas Lippert¹
- Dietrich Wolf²

¹ Jülich
Supercomputing
Centre, Germany

² Universität
Duisburg-Essen,
Germany

Minister of Science of Baden-Württemberg visits HLRS

On April 13th 2015 Teresia Bauer, Minister of Science of the State of Baden-Württemberg visited the HLRS. Having declared High Performance Computing one of the key technologies for the state of Baden-Württemberg Minister Bauer came to see the center and the new supercomputer "Hornet" that is operational since January 2015. In a guided tour through the facilities Prof. Michael Resch – Director of HLRS – explained the potential and power of the Cray XC40 and pointed to the challenges for both HLRS and its users that come with the ever growing level of parallelism. Minister Bauer was deeply impressed both by the Cray system and the technical facilities at HLRS.

After the technical tour Minister Bauer was shown to the Virtual Reality environment of HLRS. Dr. Wössner – Head of Visualization at HLRS – gave a presentation of how visualization can be used to make simulation results not only visible but also understandable for both the users of HLRS and the general public. Showing the three-dimensional view of a planned water power plant Dr. Wössner gave an example for how the results of visualization had been



made available to citizens affected by the construction and operation of the power plant. Minister Bauer got a very good impression of how simulation and visualization can be used in a dialogue between politics and citizens to discuss better solutions especially for large scale projects.

In a final discussion Minister Bauer and Prof. Resch touched on a number of further important issues related to simulation. Both agreed that the role of simulation in the political decision making process will be more important in the future. They also agreed that training and education in simulation technology and high performance computing are vital issues especially for a high-tech region like the state of Baden-Württemberg.

- inside team

University of
Stuttgart, HLRS,
Germany

Cray CEO visits HLRS



On January 20th 2015 Peter J. Ungaro President and CEO of Cray visited the High Performance Computing Center Stuttgart (HLRS) to discuss future strategies and collaborations in HPC. In a face-to-face meeting with the Director of HLRS Prof. Michael Resch strategic issues in the development of the worldwide HPC community as well as the collaboration between Cray and HLRS were discussed. After a tour through the facilities of HLRS and a visit to the new flagship system "Hornet" – a Cray XC40 with a peak performance of 3,8 POFLOP/s – Peter Ungaro sat down for an intensive discussion with the management and technical leadership of HLRS.

During the discussion Mr. Ungaro presented the future product strategy of Cray discussing technical options for an engineering focused center like HLRS. Both sides agreed that sustained performance is going to be the key to further success in HPC and that solutions rather than pure flop counting are the driving force for the further collaboration between HLRS and Cray. HLRS and Cray also agreed that data issues are growing in importance and have become a key driving factor for supercomputers requiring solutions that go way beyond traditional data storage and management. Data analytics and data discovery are considered one important aspect. However, beyond the classical issues of Big Data a combination of data and simulation will become ever more important in the future. It will therefore be vital for an HPC center to be able to provide solutions that integrate both worlds in a single solution for its users.

- inside team

University of
Stuttgart, HLRS,
Germany



Member of Executive Board of Porsche visits HLRS

On May 13th 2015 Wolfgang Hatz, Member of the Executive Board – Research and Development of Porsche, visited the HLRS. Together with his team in charge of virtual car development Wolfgang Hatz came to get a view of the new Cray system and to discuss with HLRS the long term relation of Porsche and HLRS in the field of high performance technical computing. Prof. Michael Resch – Director of HLRS – gave an overview of the center and the history of the collaboration with Porsche. In his talk he highlighted the role of Porsche as a partner for HLRS but also the potential of HPC and especially of the new Cray XC40 system. Wolfgang Hatz was impressed by the variety of issues and the potential of the new system for technical simulation.



After a tour through the computer room and the facilities Wolfgang Hatz was shown to the Virtual Reality environment of HLRS. Dr. Wössner – Head of Visualization at HLRS – gave a presentation of how visualization can be used to make simulation results not only visible but also understandable for both the users of HLRS and the general public. Wolfgang Hatz was invited to take a ride in the Porsche driving simulator at HLRS and was impressed by the potential of such a simulator.

- inside team

University of
Stuttgart, HLRS,
Germany



During the presentation Porsche and HLRS discussed the potential of Augmented Reality in the design and development process of cars. Both sides agreed that there is a lot of potential for the extension of the collaboration not only in terms of computing but also in the integration of visualization and simulation.

Jülich's Bernd Mohr to Chair SC17



Jülich Superscomputing Centre is extremely happy to announce that the SC Steering Committee elected Jülich scientist Dr.-Ing. Bernd Mohr as General Chair of the SC17 Conference in Denver, Colorado. It is the first time ever that the most important conference for High Performance Computing, networking, storage and analysis – the Supercomputing Conference (SC) in the USA – will be chaired by a non-American. Shortly after the announcement, the portal HPCwire listed Mohr as one of the people to watch in 2015. See: www.hpcwire.com/peoplewatch-2015/

Bernd Mohr started to design and develop tools for performance analysis of parallel programs already with his diploma thesis (1987) at the University of Erlangen in Germany, and continued this in his Ph.D. work (1987 to 1992). During his PostDoc position at the University of Oregon, he designed and implemented the original TAU performance analysis framework. Since 1996 he has been a senior scientist at Forschungszentrum Jülich, Germany's largest multidisciplinary research center. He now acts as team leader for the group "Programming Environments and Performance Optimization" and serves as deputy

head for the JSC division "Application support". Besides being responsible for user support and training in regard to performance tools at the Jülich Supercomputing Centre (JSC), he is leading the Scalasca performance tools effort in collaboration with Prof. Dr. Felix Wolf, now at TU Darmstadt.

His first visits to SC were 1993 in Portland and 1994 in Washington, D.C., during his PostDoc days at the University of Oregon in Eugene. Later, in 1999, he was among a small team which was responsible for setting up and staffing the research exhibits booth of the Jülich Supercomputing Centre for a few years. Dr.-Ing. Mohr also gave 11 SC tutorials between 1999 and 2009. Finally, he got involved in helping organize the conference as a research paper reviewer for SCO3 and worked his way up serving various roles in the technical program committee. In 2011, he became the first European to be elected to the SC Steering Committee.

Established in 1988, the annual SC conference attracts scientists and users from all over the world who come together to discuss current developments. More than 10,000 people now participate in SC each year. The conference is sponsored by the Association for Computing Machinery (ACM) and the Institute of Electrical and Electronics Engineers (IEEE).

The work for SC17 is already well under way. One of the most important tasks is to build his core organization team, around which some 500 volunteers will gather in the coming year. He also started to think about fine-tuning the topics for 2017: Besides an even stronger international focus for the conference, potential directions will be linking HPC more closely to topics such as big data analytics and pre- and post-processing including visualization.

I am personally very proud of working with Bernd at Jülich Supercomputing Centre. He was the person promoting the international reputation of JSC's computer science research most profoundly. His personality and enthusiasm have encouraged many international researchers to cooperate with JSC and have helped to create numerous personal relationships between JSC and international experts.

- Thomas Lippert

Jülich
Supercomputing
Centre (JSC),
Germany

HLRS Supercomputer successfully executed Extreme-scale Simulation Projects

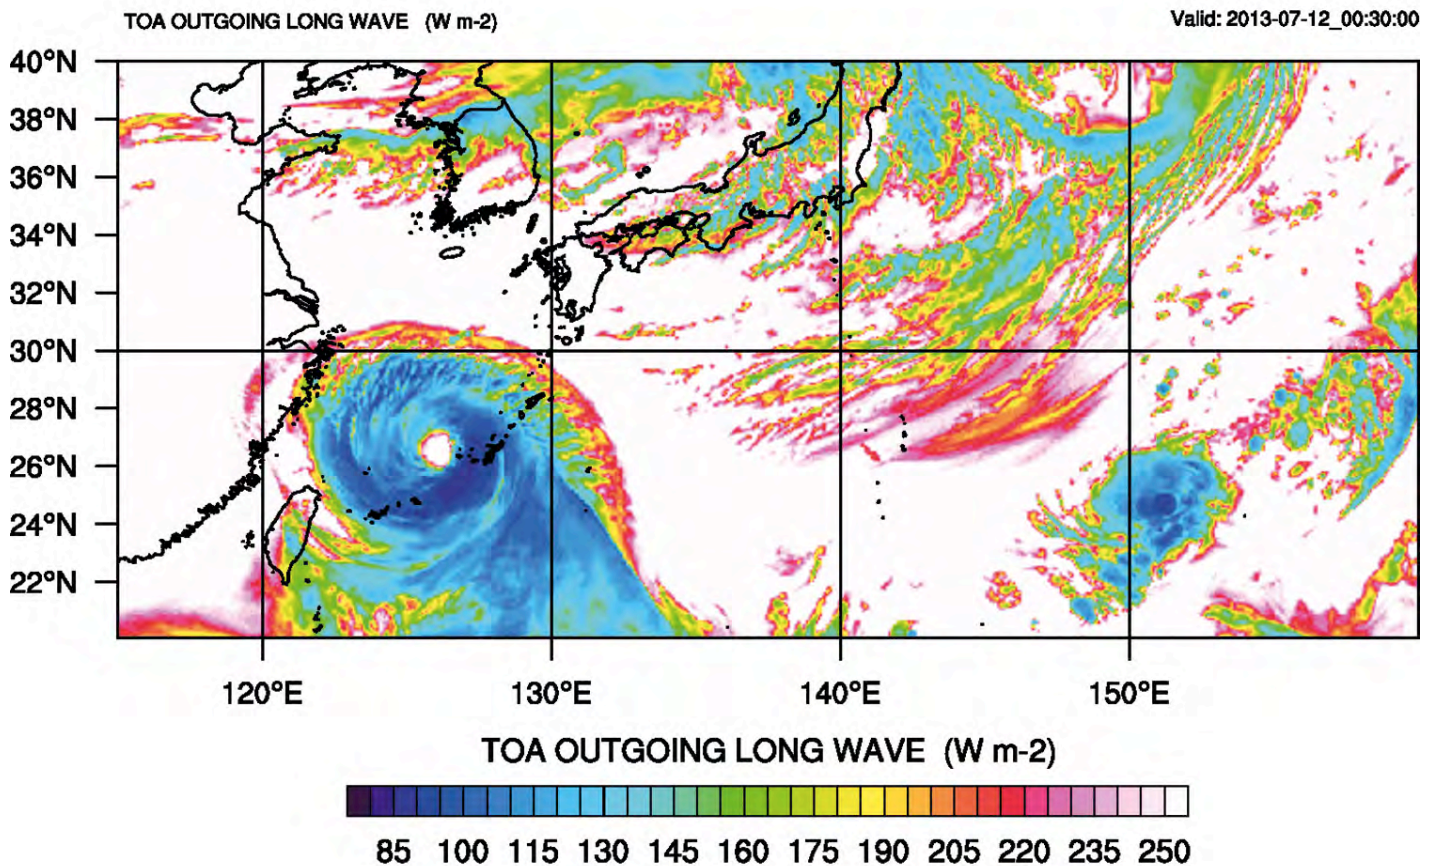


Figure 1: Fine scale cloud structures of typhoon Soulik depicted by the outgoing long wave radiation. Blue areas indicate dense high clouds whereas white areas show cloud-free regions. Because of the high resolution, the eyewall of the typhoon with the strongest precipitation is clearly seen by the dark blue area southwest of Okinawa.

© Institute of Physics and Meteorology, Universität Hohenheim

The newly installed HPC system Hornet of HLRS successfully finished extensive simulation projects that by far exceeded the calibre of previously performed simulation runs at HLRS: Six so called XXL-Projects from computationally demanding scientific fields such as planetary research, climatology, environmental chemistry, aerospace, and scientific engineering were applied on the HLRS supercom-

puter in early January. With each application scaling up to all of Hornet's available 94,646 compute cores, the machine was put through a demanding endurance test. The achieved results more than satisfied the HLRS HPC experts as well as the scientific users: Hornet lived up to the challenge and passed the simulation "burn-in runs" with flying colors.

The six XXL-Projects implemented on the HLRS supercomputer, are:

Convection permitting Channel Simulation

Institute of Physics and Meteorology, Universität Hohenheim, Wulfmeyer, V., Warrach-Sagi, K., Schwitalla, T.

- 84,000 compute cores used
84 machine hours
- 330 TB of data + 120 TB for post-processing.

Current high-resolution weather and climate models are operated over a domain which is centred over the region of interest. This configuration suffers from a deterioration of large-scale features like low pressure systems when propagating in the inner domain. This feature is strongly limiting the quality of the simulation of extreme weather events and climate statistics. A solution can be a latitude belt simulation around the Earth at a resolution of a few km. By the XXL project, a corresponding simulation became possible for a time period

long enough to cover various extreme events on the Northern hemisphere and to study the model performance. The results confirm an extraordinary quality with the respect to the simulation of extreme weather events such as the taifun Soulik in the Pacific from July 10-12, 2013.

The storage capabilities of the new Hornet system allowed the scientists to run the simulation without any interruptions for more than three days. Using the combination of MPI+OpenMPI including PNetCDF libraries, the performance turned out to be excellent. Another finding was that not the computing time but the I/O performance became the limiting factor for the duration of the model run.




Figure 2: λ_2 -visualization of the turbulent structures along the flat plate.
© IAG, Universität Stuttgart

Direct Numerical Simulation of a Spatially-developing Turbulent Boundary Layer Along a Flat Plate

Institute of Aerodynamics and Gas Dynamics (IAG), Universität Stuttgart, Atak, M., Munz, C. D.

scaling attributes being a necessity to satisfy the computational demand in a sustainable and efficient way. The outcome of this work allowed the researchers to establish a database which can be used for further complex investigations such as shock / wave boundary layer interactions.





Figure 3: Vorticity contours with mapped on Mach number
© Institute of Aerodynamics, RWTH Aachen University

- 93,840 compute cores used
70 machine hours
- 30 TB of data

The intake flow of hypersonic air-breathing propulsion systems is characterized by laminar and turbulent boundary layers and their interaction with impinging shock waves. The objective of this work was to conduct a direct numerical simulation of the complete transition of a boundary layer flow to fully-developed turbulence along a flat plate up to high Reynolds numbers. The scientists applied a high-order discontinuous Galerkin spectral element method which inherently involves excellent

Prediction of the turbulent Flow Field around a ducted axial Fan

Institute of Aerodynamics, RWTH Aachen University, Pogorelov, A., Meinke, M., Schröder, W.

- 92,000 compute cores used
110 machine hours
- 80 TB of data

The turbulent low Mach number flow through a ducted axial fan is investigated by large-eddy simulations using an unstructured hierarchical Cartesian method. It is the purpose of this computation to understand the development of vortical flow structures and

the turbulence intensity in the tip-gap region. To achieve this objective a resolution in the range of 1 billion cells is necessary. This defines a computational problem that can only be tackled on a Tier-0 machine.

Large-Eddy Simulation of a Helicopter Engine Jet

Institute of Aerodynamics, RWTH Aachen University, Cetin, M. O., Meinke, M., Schröder, W.

- 94,646 compute cores used
300 machine hours
- 120 TB of data

The impact of internal perturbations due to geometric variations, which are generally neglected, on the flow field and the acoustic field of a helicopter engine jet was analyzed by highly resolved large-eddy simulations based on hierarchically refined Cartesian meshes up to 1 billion cells. The intricacy of the flow structure requires such a detailed resolution which could only be realized on an architecture like featured by Hornet.

Ion Transport by Convection and Diffusion

Institute of Simulation Techniques and Scientific Computing, Universität Siegen, Masilamani, K., Klimach, H., Roller, S.

- 94,080 compute cores used
5 machine hours
- 1.1 TB of data

The goal of this computation was a more detailed simulation of the boundary layer effects in electro-dialysis processes, used for seawater desalination. This simulation involves




Figure 4: Flow structure downstream of the centre body
© Institute of Aerodynamics, RWTH Aachen University

the simultaneous consideration of multiple effects like flow through a complex geometry, mass transport due to diffusion and electrodynamic forces. The behavior in the boundary layer has a big influence on the overall process, but is not well understood. Only large computing resources offered by Petascale systems such as Hornet enable the consideration of all involved physical effects enabling a more realistic simulation than ever before.

Large Scale Numerical Simulation of Planetary Interiors

German Aerospace Center / Technische Universität Berlin, Hüttig, C., Plesa, A. C., Tosi, N., Breuer, D.

- 54,000 compute cores used
3 machine hours
- 2 TB of data

The interior of planets and planet-like objects have a very hot interior that causes heat-driven motion. To study the effect of this kind of motion on the evolution of a planet large-scale computing facilities are necessary to understand the evolving patterns under realistic parameters and compare them with observations. The goal is to understand how the surface




Figure 5: Species transport in the spacer filled flow channel

© Institute of Simulation Techniques and Scientific Computing, Universität Siegen

is influenced, how conditions for life are maintained, how plate-tectonics work and how quickly a planet can cool.

With the Hornet project, scientists were able, for the first time, to study the flow under realistic assumptions in full 3D and capture valuable information such as surface heat-flow and stresses.

Demand for HPC on the Rise

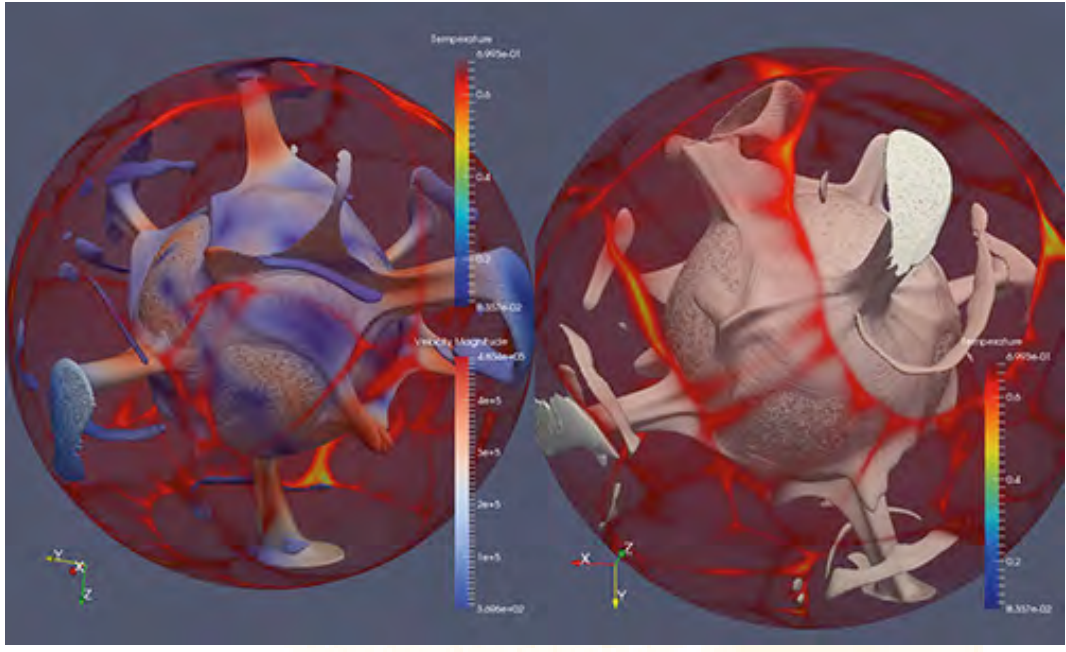
Demand for High Performance Computing is unbroken. Scientists continue to crave for ever increasing computing power. They are eagerly awaiting the availability of even faster systems and better scalable software enabling them to attack and puzzle out the most challenging scientific and engineering problems. "Supply generates demand", states Prof. Michael M. Resch, Director of HLRS. "With the ability of ultra-fast machines like Hornet both industry and researchers are quickly realizing that fully leveraging the vast capabilities of such a supercomputer opens unprecedented opportunities and helps them to deliver results previously impossible to obtain.

We are positive that our HPC infrastructure will be leveraged to its full extent. Hornet will be an invaluable tool in supporting researchers in their pursuit for answers to the most pressing subjects of today's time, leading to scientific findings and knowledge of great and enduring value," adds Prof. Michael M. Resch.

Outlook

Following its ambitious technology roadmap, HLRS is currently striving to implement a planned system expansion which is scheduled to be completed by the end of 2015. The HLRS supercomputing infrastructure will then deliver a peak performance of more than seven PetaFlops (quadrillion mathematical calculations per second) and feature 2.3 petabytes of additional file system storage.

contact: Uwe Küster,
kuester@hls.de



- Regina Weigand
Gauss Centre of Supercomputing

Figure 6: Bottom and internally heated high-Rayleigh number convection with strong viscosity contrasts in a 3D spherical shell to study patterns and stresses relevant for earth.
© German Aerospace Center, Berlin

Discontinuous Galerkin for High Performance Computational Fluid Dynamics




Figure 1: Artist's concept of the X-51A WaveRider (Image credit: www.nasa.gov)

Towards Scramjets using Discontinuous Galerkin Spectral Element Methods

A space shuttle taking off from the ground is an impressive moment and the images broadcasted on TV are still attracting our attention. But as fascinating as they are, the future of space flight may belong to air-breathing supersonic and hypersonic vehicles. Particularly scramjets (supersonic combustion ramjets) are considered

as an ambitious candidate for classical rocket driven systems. In contrast to today's space transportation technologies which need to carry tons of oxygen supplies, the scramjet inhales the atmospheric air to gain the oxygen required for combustion. Thus, the aircraft becomes lighter, faster and eventually more economic since the payload can be significantly increased. The idea of an air-breathing vehicle is not new though, scramjets having been

a concept since the 1950s and in the 1960s the first scramjet engines were tested in labs. However, after more than 50 years of scramjet research a breakthrough in aviation history was achieved in May 2013 when the U.S. Air Force successfully launched the X-51A WaveRider (Fig. 1) reaching a top speed of more than five times the speed of sound for a record-breaking flight time of four minutes.

Despite extensive studies over decades and first encouraging flight tests, the scramjet still merges and challenges various research fields like aerodynamics, combustion and material science. Especially the intake of a scramjet plays a key role for the proper functionality of air-breathing vehicles as it does not compress the incoming air via moving parts like compressors, but through a series of shock waves generated by the specific shape of the intake and the high flight velocity (see Fig. 2). As the supply of compressed air is of crucial importance for the subsequent efficient combustion of the fuel-air mixture to produce thrust, the intake flow also determines the operability limits of the whole system. Along with the previous two flight tests of the X-51A WaveRider which failed due to inlet malfunctions, the scientists attached great importance to the detailed study of the

intake flow. Since experimental and flight data of hypersonic air-breathing vehicles are difficult and utterly expensive to obtain, numerical methods are applied to enhance our understanding of the involved complex physical phenomena.

The Discontinuous Galerkin Method

The intake flow of a scramjet is characterized by laminar and turbulent boundary layers and their interaction with shock waves, yielding a three-dimensional unsteady complex flow pattern. Particularly shock wave / boundary layer interactions are of fundamental importance as they may cause intense heat loads leading to serious aircraft damages. A reliable numerical investigation of the occurring phenomena within the intake can be accomplished by means of high fidelity simulations, as direct numerical simulations (DNS) and large eddy simulations (LES) offer, which in turn require large computational resources. Hence, the applied numerical method has to satisfy many requirements, like for instance high spatial resolution and low dissipation errors, to accurately resolve the turbulent features of the flow. Moreover, the method should also allow an efficient usage of HPC systems to reduce the computational time. Due to its excellent dispersion




Figure 2: Schematic of a scramjet

and dissipation attributes as well as its highly performant parallelization, the discontinuous Galerkin (DG) method is recently considered as a promising approach to conduct such high-fidelity simulations in an efficient way. Beyond that, the DG method also combines arbitrary high-order resolution in space with geometrical flexibility, even enabling the computation of complex geometries.

In the present work, however, we apply a special collocation type formulation of the DG method for unstructured hexahedral element meshes, namely the discontinuous Galerkin spectral element method (DGSEM) [1], whose main advantage is based on its HPC capability. In contrast to other high-order schemes (e.g. finite difference and finite volume methods), the DGSEM algorithm with explicit time discretization is inherently parallel since all elements communicate only with direct neighbors. Furthermore, the DGSEM operator itself can be split into two parts, namely the volume part, which solely relies on local data, and the surface part, for which neighbor information is required. This feature of the DG method can be exploited by hiding communication latencies and the nega-

tive influence of data transfer can be reduced to a minimum. It is therefore possible to send surface data while simultaneously performing volume data operations. Hence, the DGSEM facilitates a lean parallelization strategy, where except for direct neighbor communication no additional operations are introduced, being important for an efficiently scalable algorithm. The DGSEM algorithm has been implemented into our code framework FLEXI [2], which is fully MPI-parallelized and exhibits excellent parallel efficiency and scaling attributes. In this context, Figs. 3 and 4 show strong scaling results of the DGSEM code FLEXI on the HLRS Cray XE6 supercomputer cluster for different polynomial degrees N and different mesh sizes, respectively. Given a fixed number of elements we increased the number of processors by a factor of two in each step and, thus, decreased the load per processor up to the one-element-per-processor limit. Here, we see that for the highest polynomial degree $N=8$, corresponding to an order of accuracy of 9, the code permanently yields a so-called super-scaling, i.e., scaling efficiency higher than 100%, owing to low data communication and memory consumption such that caching effects can be exploited. It is also




Figure 3: Strong scaling for different polynomial degrees N on the same mesh

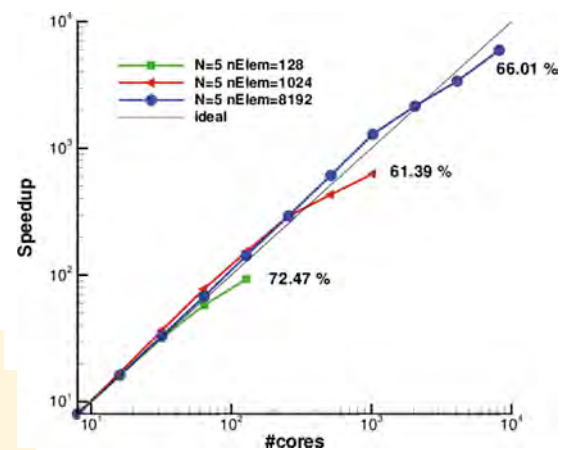


Figure 4: Strong scaling for a fixed polynomial degree $N=5$ on different meshes




Figure 5: Visualization of the turbulent structures using the λ_2 -criterion and colored by the streamwise velocity

possible to achieve super-scaling for lower polynomial degrees by increasing the load per processor to optimize the volume-to-surface data ratio (4 elements per processor for $N=5$ and 8 elements per processor for $N=3$). Hence, the scaling results prove that the DG method is very well suited for massively parallel high performance computing.

DNS of a Turbulent Boundary Layer Along a Flat Plate

Transition describes the process of a laminar flow becoming turbulent. It also represents the initial phenomenon that occurs along the intake walls before shock waves interact with the emerging boundary layers. Thus, we will first investigate the most canonical form of transition, namely the DNS of a spatially-developing supersonic turbulent boundary layer on a flat plate, in order to demonstrate the high potential of DG schemes for compressible wall-bounded turbulent flows. The free-stream Mach number, temperature and pressure of the boundary layer flow are given by $M = 2.67$, $T = 564$ K and $p = 14890$ Pa, respectively. We follow the approach of developing turbulence,

which involves the simulation of the complete laminar-turbulent transition process. Fig. 5 visualizes the streamwise development of turbulent structures showing how the initially laminar flow experiences transition and eventually becomes turbulent.

We choose a polynomial degree of $N=5$ to enable a direct comparison with Keller & Kloker [3] who used a sixth order compact finite difference (cFD) code for a similar simulation setup. As




Figure 6: Van Driest transformed streamwise velocity profiles at the most downstream location




Figure 7: Spanwise averaged Reynolds-stress profiles

displayed in Fig. 6, the van Driest transformed velocity profile at the most downstream position matches both the compressible reference solution and the incompressible scaling laws. In Fig. 7 we compare the Reynolds-stresses, again, at the most downstream position where the momentum thickness based Reynolds number reaches a value of $Re_\theta=1361$. In addition, the Reynolds-stresses are density-weighted to compare the results with Spalart's incompressible turbulent boundary layer at $Re_\theta=1410$ [4]. Except for the immediate near-wall region, where slight differences to the incompressible Reynolds-stresses can be observed, the DG results are in a very good agreement with both the compressible reference case and Spalart's incompressible boundary layer.

The computational efficiency of the DGSEM scheme can be assessed by the so-called performance index (PID). The PID is a convenient measure to judge the computational efficiency and it expresses the computational time needed to update one degree of freedom (DOF) for one time step: $PID = (\text{Wall-clock-time} \times \#\text{cores}) / (\#\text{DOF} \times \#\text{timesteps})$. The comparison of the PID reveals that the DG method ($PID_{\text{DGSEM}}=15$) is able to compete with

state-of-the-art compact FD codes ($PID_{\text{cFD}}=17$) and emphasizes the strong efficiency of the DGSEM code at performing high-fidelity simulations of wall-bounded compressible turbulent flows. Based on the encouraging results, the outlook of the present study is to conduct a DNS of a shock wave / boundary layer interaction to gain a deeper insight into the complex physics of the intake flow.

Acknowledgments

This research has been supported by the Deutsche Forschungsgemeinschaft through the research training group "Aero-Thermodynamic Design of a Scramjet Propulsion System for Future Space Transportation Systems". The simulations were conducted on the HLRS Cray XE6 supercomputer system.

References

- [1] Kopriva, D. A. Implementing Spectral Methods for Partial Differential Equations: Algorithms for Scientists and Engineers. Springer Publishing Company, Incorporated, 2009.
- [2] Hindenlang, F., Gassner, G., Altmann, C., Beck, A., Staudenmaier, M., Munz, C.-D. Explicit discontinuous Galerkin methods for unsteady problems. *Computers & Fluids*, 61 pp. 86-93, 2012.
- [3] Keller, M. and Kloker, M. DNS of effusion cooling in a supersonic boundary-layer flow: Influence of turbulence. *AIAA Paper 2013-2897*, 2013.
- [4] Spalart, P. R. Direct simulation of a turbulent boundary layer up to $Re_\theta=1410$. *Journal of Fluid Mechanics*, 187:61-98, 1988.

contact: Muhammed Atak,
 atak@iag.uni-stuttgart.de
 Claus-Dieter Munz,
 munz@iag.uni-stuttgart.de

- Muhammed Atak
- Claus-Dieter Munz

Institute of
 Aerodynamics
 and Gas Dynamics,
 University of
 Stuttgart, Germany

Golden Spike Award by the HLRS Steering Committee in 2014

Large-scale Simulations of Particle-laden Flows

The Computational Fluid Dynamics group of the Institute of Aerodynamics and Chair of Fluid Mechanics at the RWTH Aachen University uses the HLRS Supercomputers to study particle-laden flows by both Lagrangian models as well as fully resolved particles methods. The latter are particularly challenging due to the wide range of time and length scales involved. As a consequence a high mesh resolution is required to capture the fluid motion around the particles, responsible for the particle motion,

and the large scales of the flow. The interaction of these phenomena is, e.g., responsible for the particle concentration within the fluid domain.

Particle-laden flows play a major role in many engineering applications and natural sciences, which include fields as diverse as biomedical applications, weather forecasting, manufacturing and combustion processes. Studying the deposition of inhaled aerosols and particles in the human lung is a biomedical application, which aims to help

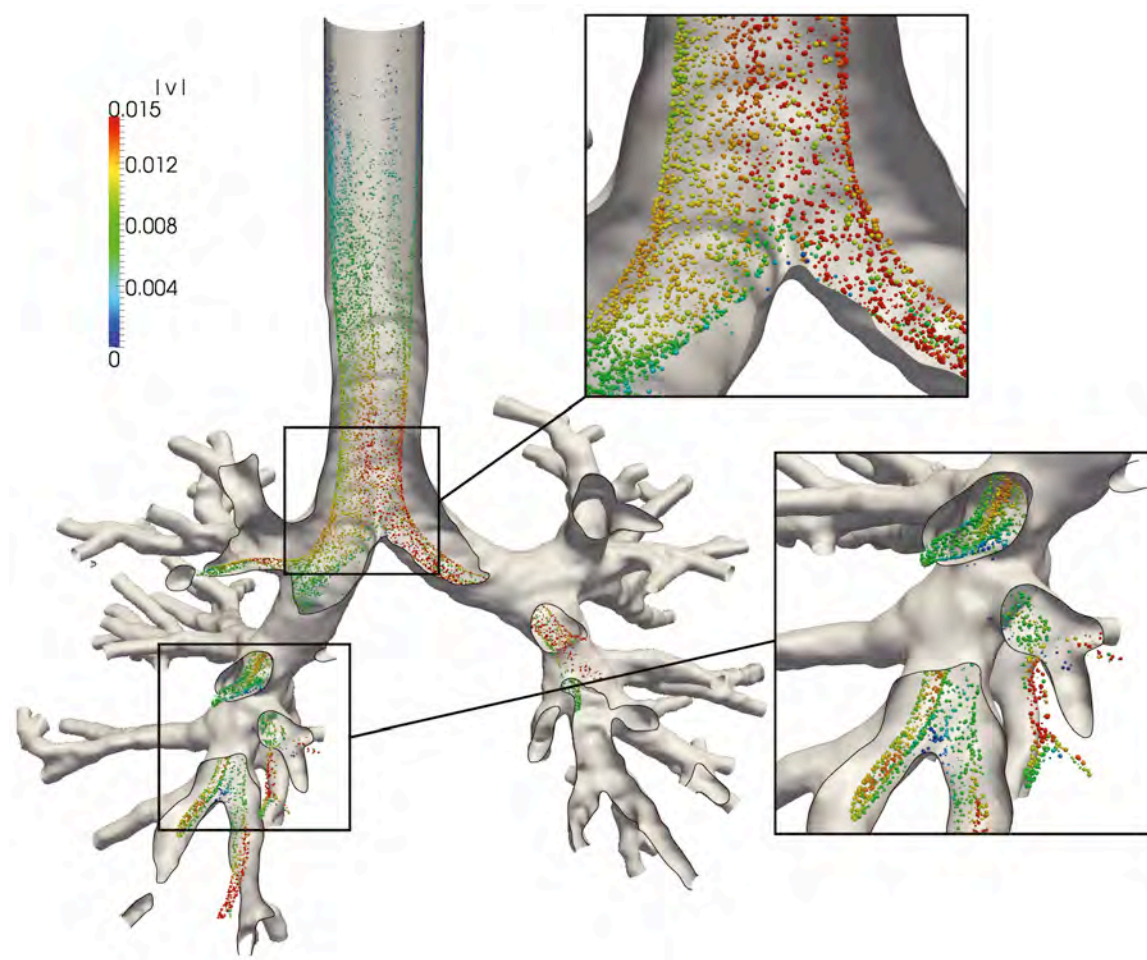


Figure 1: Fine dust particle distribution in the human lung during inspiration. The particles, tracked by a Lagrangian model, are transported down the trachea and distribute in the left and right main bronchi before they deposit in higher lung generations where they can cause severe damage to the respiratory system potentially leading to lung cancer [2].

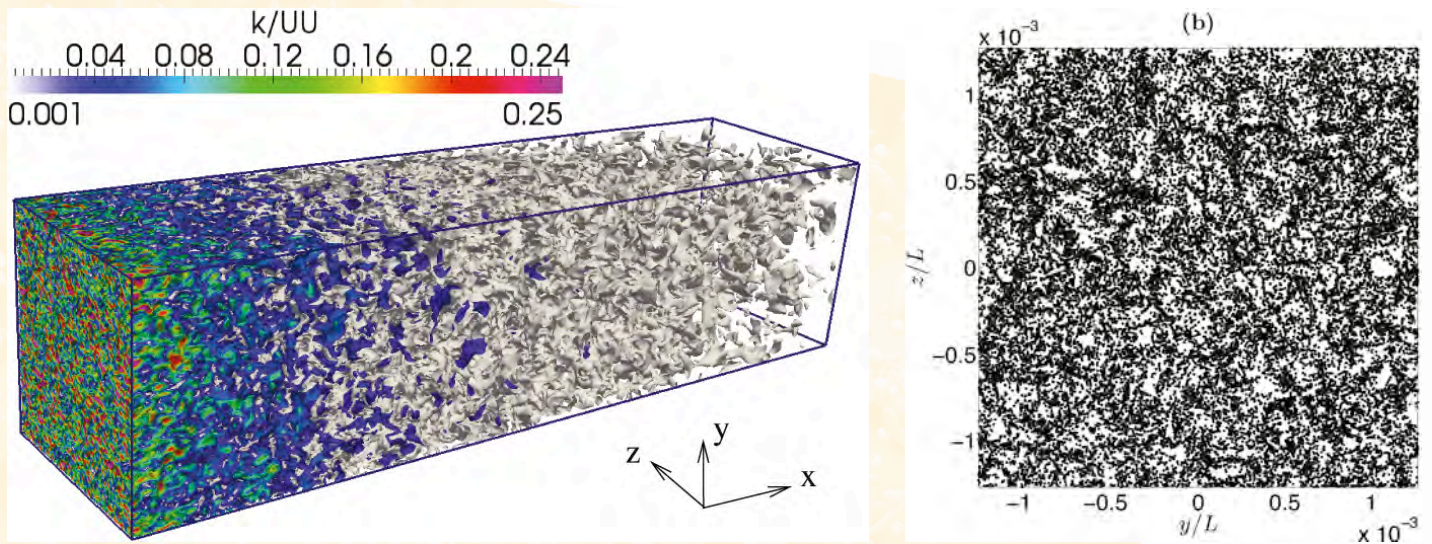


Figure 2: (Left) At the inflow plane of the computational domains (at the left side) synthetic eddies are generated to generate an isotropic turbulence field. It is shown how the turbulent intensity decays due to the dissipation as the turbulence is convected downstream (in positive x direction). (Right) Rain particle distribution determined by Lagrangian particle-tracking for a Stokes number of ~ 1 in an isotropic, decaying turbulence field. The particles concentrate in regions of high strain and low vorticity leading to a higher collision probability.

to understand associated pathologies (see Fig. 1). Within the high-priority research program SPP1276 METSTROEM the growth of water droplets within vapor clouds is studied to increase our understanding of the rain formation process [1]. The range of scales in this process spans from the size of aerosol particles (micrometers) to the cloud's length (kilometers), see Fig. 2. The manufacturing quality of the electrical discharge machining process is known to be sensitive to the debris removal. In the SFB/TRR136 (Process Signatures) the physics of the debris deposition process for multiple removal strategies is studied to better understand the underlying physical phenomena of the cleaning and characterize the signature of the pro-

cess. Finally, oxy-fuel combustion is an emerging combustion technology that significantly reduces the CO_2 emissions of coal power plants. Detailed numerical simulations of the oxy-fuel combustion process including the burning coal particles are performed within the trans-regional collaborative research center SFB/TRR129 (Oxyflame), see Fig. 3.




Figure 3: Fully-resolved particle simulation involving 1,000 particles in a Taylor Green Vortex at Reynolds' number 1,600 on 1,800 cores. The computational grid is a periodic cube. On its surfaces the vorticity distribution is shown (heat map). Vortical structures are visualized with the contours of the Q-criterion (green to blue).

With this objective in mind, the multi-physics in-house code ZFS (Zonal Flow Solver) has been developed at the Institute of Aerodynamics over the past decade. It is based on the concept of hierarchical Cartesian meshes [3], which allows adaptive mesh refinement techniques to automatically invest the computational effort on those regions of the domain where the solution changes more rapidly, while reducing the grid resolution where the solution is smooth. This allows tracking the small scales of fluid flows where they exist, i.e., maintaining a fine grid e.g. around the particles as they move through the domain. A space-filling curve is used to partition the domain such that the communication between processors is minimized. The concept of space filling curves also allows to redistribute the load dynamically as the particles move from one processor to another, enabling efficient large-scale adaptive simulations. Finally, the mathematical models of the fluid problems are solved on these hierarchical adaptive Cartesian grids by numerical methods optimized for the underlying physical phenomena characteristic of different flow regimes. ZFS provides a

Lattice-Boltzmann method for incompressible flows, finite-volume methods for compressible flow and combustion problems, and high-order discontinuous Galerkin methods for acoustic problems. These are complemented with particle solvers that allow simulating millions of point particles using Lagrangian models, and level-set methods for performing fully-resolved particle simulations involving thousands of particles.

The parallel efficiency of ZFS is demonstrated by means of a so-called strong-scaling experiment, in which the size of the problem is kept constant while the number of computing cores is continuously increased. For an ideally efficient solver the product of the number of computing cores time the required computing time would stay constant. In Fig. 4 a study of the parallel efficiency of the grid generation process for a grid containing 9.82×10^9 cells is shown. A parallel efficiency close to 100% is maintained till 32,768 cores. As the number of cells to be generated per core further decreases the parallel efficiency also decreases due to a lacking amount of work per core, finally achieving a parallel efficiency of ~75% on 131,072 cores.

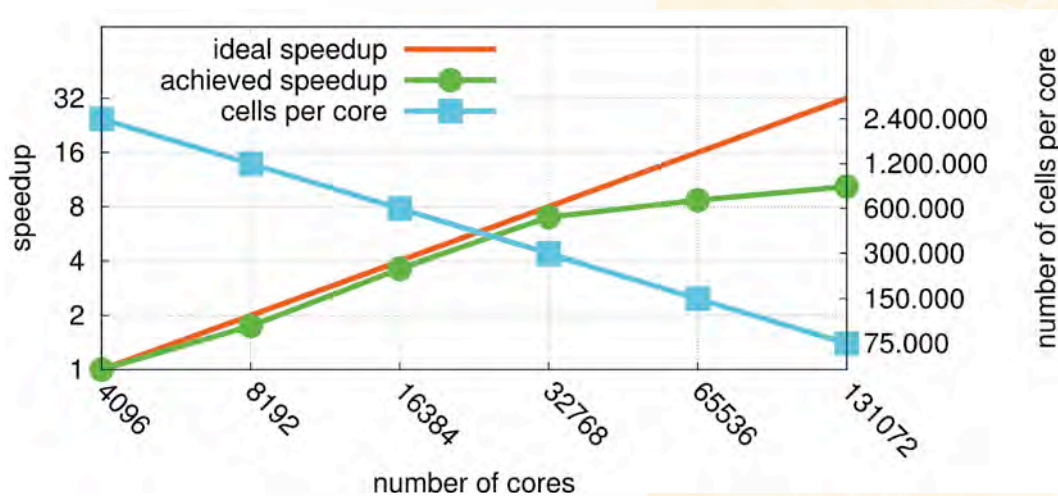


Figure 4: Results of the strong scaling experiment for the generation of a grid consisting of 9.82×10^9 cells. The ideal speed up (red), the actual achieved speedup (green), as well as the number of cells generated per core (blue) are shown [4].

Acknowledgements

This work has been financed by the research cluster Fuel production with renewable raw materials (BrenaRo) at RWTH Aachen University as well as by the German Research Foundation (DFG) within the framework of the SFB/Transregio 129, SFB/Transregio 136, the transregional collaborative research center SFB 686, and within the framework of the priority program SPP 1276 METSTROEM under grant number SCHR 309/39. The support is gratefully acknowledged. The authors thank for the provided computing time at the IBM BlueGene/Q system (JUQUEEN) at the Jülich Supercomputing Center (JSC) and the CRAY XE6 system (HERMIT) at the High Performance Computing Center Stuttgart (HLRS).

- Gonzalo Brito Gadeschi¹
- Lennart Schneiders¹
- Christoph Siewert^{1, 2}
- Andreas Lintermann^{1, 3}
- Matthias Meinke¹
- Wolfgang Schröder^{1, 3}

¹ Institute of Aerodynamics, RWTH Aachen University, Germany

² Laboratoire Lagrange, OCA, CNRS, Nice, France

³ SimLab "Highly Scalable Fluids & Solids Engineering", JARA-HPC, RWTH Aachen University, Germany

References

- [1] **Siewert, C., Kunnen, R.P.J., and Schröder, W.**
Collision Rates of Small Ellipsoids Settling in Turbulence, *Journal of Fluid Mechanics*, 758, 686–701, 2014.
- [2] **Lintermann, A. and Schröder, W.**
Simulation of Aerosol Particle Deposition in the Lower Human Respiratory System, submitted to the *Journal of Aerosol Science*, 2015.
- [3] **Lintermann, A., Schlimpert, S., Grimm, J. H., Günther, C., Meinke, M., and Schröder, W.**
Massively parallel grid generation on HPC systems, *Computer Methods in Applied Mechanics and Engineering*, 277, 131–153. doi:10.1016/j.cma.2014.04.009, 2014.
- [4] **Brito Gadeschi, G., Siewert, C., Lintermann, A., Meinke, M., and Schröder, W.**
Towards large multi-scale particle simulations with conjugate heat transfer on heterogeneous super computers, *High Performance Computing in Science and Engineering '14*, 2015.

contact: Gonzalo Brito Gadeschi,
g.brito@aia.rwth-aachen.de
Dr. Matthias Meinke,
m.meinke@aia.rwth-aachen.de

Quarks and Hadrons - and the Spectrum in Between

The mass of ordinary matter, due to protons and neutrons, is described with a precision of about five percent [1] by a theory void of any parameters – (massless) Quantum Chromodynamics (QCD). The remaining percent are due to other parts of the Standard Model of Elementary Particle Physics: the Higgs-Mechanism, which explains the existence of quark masses, and a per mil effect due to the presence of Electromagnetism. Only at this stage, free parameters are introduced¹. Mass is thus generated in QCD without the need of quark masses, an effect termed “mass without mass” [2].

This mass generation is due to the dynamics of the strongly interacting quarks and gluons, the latter mediating the force like the photon in the case of Quantum Electrodynamics. However, in contrast to the photons, which are themselves uncharged and thus do not interact (directly) with one another, the gluons themselves carry the strong (color) charge. This is the essential complication in the dynamics of the strong interaction, since it renders the interaction non-linear. Furthermore, due to the uncertainty relation of quantum physics and relativity, even the number of quarks inside the proton or neutron is not fixed: quark anti-quark pairs are created and annihilated constantly and contribute significantly to the overall mass of the particle. So far, the only known way to analyze QCD at small energies, where its interaction strength is large, is through simulations. These simulations are based on the discretized theory, called

Lattice QCD, which traditionally uses a space-time lattice with quarks on the lattice sites and the gluons located on the lattice links. In the so-called continuum limit, the lattice spacing is sent to zero and QCD and the continuous space-time are recovered. Simulations of Lattice QCD require significant computational resources. As a matter of fact, they have been a driving force of supercomputer development. A large range of special-purpose Lattice QCD machines, e.g. the computers of the APE family, QPACE, and QCDOC among others, have been developed, the QCDOC being the ancestor of the IBM Blue Gene family of supercomputers. Over the last decade, these simulations have matured substantially. In 2008, the first fully controlled calculation of the particle spectrum [1] became available and simpler quantities such as masses and decay constants

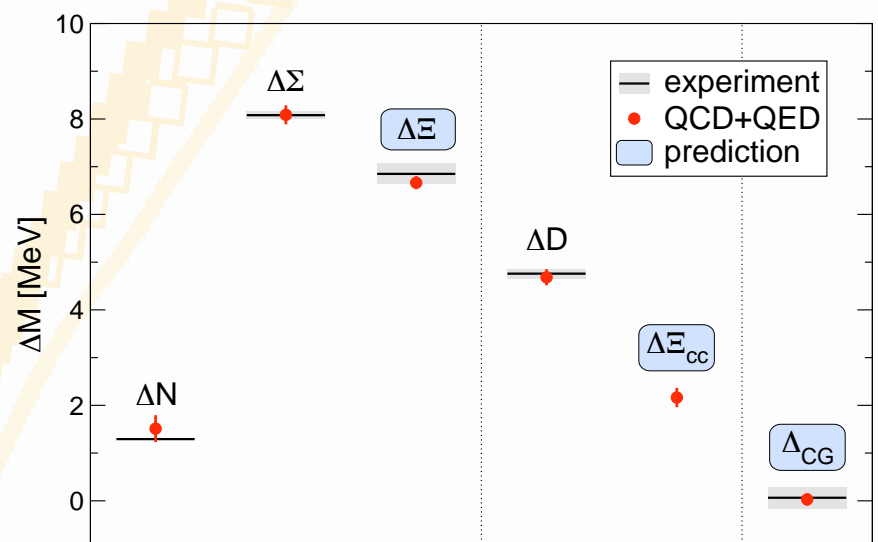


Figure 1: Mass splittings [3]. The horizontal lines are the experimental values and the grey shaded regions represent the experimental errors than those of our predictions. Our results are shown by red dots with their uncertainties. A blue shaded region indicates that no experimental results are available or that these results have larger errors.

¹ These parameters are a coupling (electromagnetic charge) and the quark masses

are now routinely computed to percent precision.

Neutron, Proton, and the Stability of Matter

In order to increase the precision of the calculations further, one has to address the largest sources of uncertainties, which, in case of the spectrum, are due to Electrodynamics and the difference between the up- and down-quark masses. Once these effects are properly included in the simulations, one can calculate per mil effects of the particle spectrum of the Standard Model, e.g. the difference between the proton and the neutron mass.

For equal light quark masses, Electrodynamics renders the proton slightly heavier, due to energy stored in the electromagnetic field that surrounds it. The light quark mass splitting, conversely, increases the neutron mass, since it contains two of the heavier down-quarks compared to one in case of the proton. The interplay between these effects has significant implications for the stability of matter. If the neutron-proton mass splitting was about a third of the 0.14% found in nature, hydrogen atoms would undergo inverse beta decay, leaving predominantly neutrons. Even with a value somewhat larger than 0.05%, Big Bang Nucleosynthesis (BBN) would have produced much more helium-4 and far less hydrogen than it did in our universe. As a result, stars would not have ignited in the way they did. On the other hand, a value considerably larger than 0.14% would result in a much faster beta decay for neutrons. This would lead to far less neutrons at the end of the BBN epoch and would make the burning of hydrogen in stars and the synthesis of heavy elements more difficult.

Including these effects is, however, non-trivial. The biggest obstacle turns out to be the long ranged nature of electro-dynamics. Whereas the strong force is essentially confined inside its "bound-states", the so-called hadrons, the electromagnetic force, falling off according to the well-known $1/r^2$ rule, is still felt at large distances. This introduces significant finite size effects, which are typically larger than the mass splittings one is interested in. In our recent calculation, the correct theoretical framework for a treatment of these effects was established and the finite-size corrections were calculated analytically. A new simulation algorithm for the Electrodynamics part of the calculations was developed, which reduced the autocorrelation by three orders of magnitude. Combined with other advanced methods, such as the latest multi-level solvers, this allowed us to compute the particle splittings using the presently available resources of the Gauss Centre for Supercomputing², JUQUEEN at JSC, Hermit at HLRS, and SuperMUC at LRZ (Fig. 1).

From Hadrons to Quark Soup

In the case of the proton and the neutron, quarks and gluons are confined to the hadron. If we, however, increase the temperature of the system sufficiently, both particles will "melt" and quarks and gluons behave as free particles, forming an exotic state of matter called quark-gluon plasma. This (rapid) transition from the quark-gluon to the hadronic phase occurred when the early universe evolved from the "quark epoch", lasting from 10^{-12} to 10^{-6} seconds after the Big Bang, to the following hadron epoch, which ended when the universe was about

² see acknowledgements

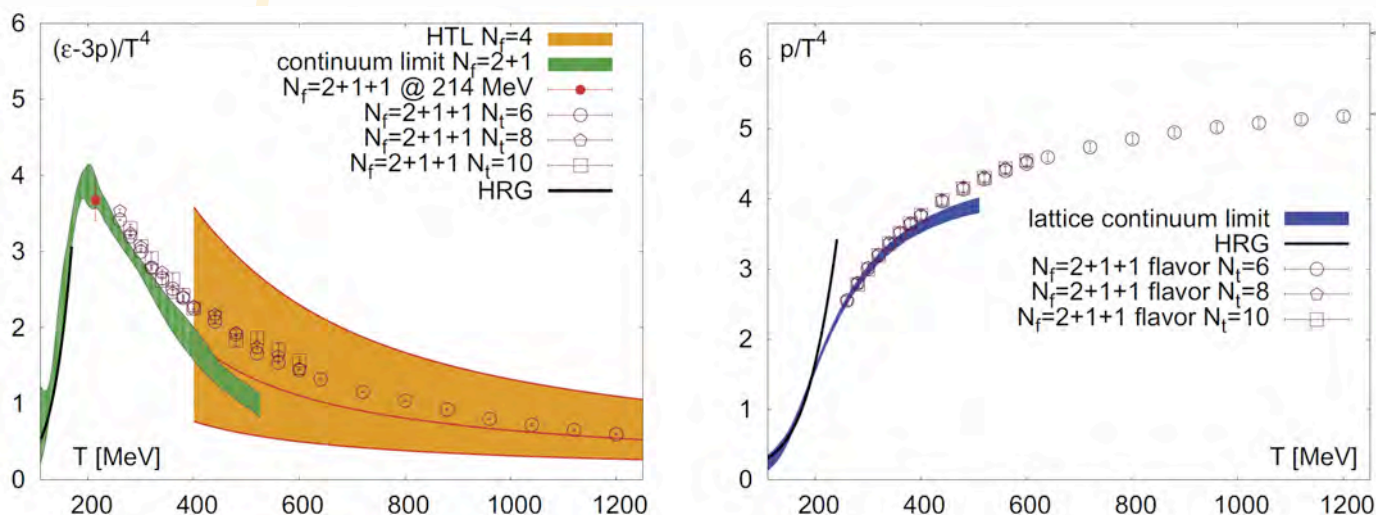


Figure 2: (Left) Preliminary results for the charmed EoS. For comparison, we show the HRG result, the $N_f = 2 + 1$ band, and, at high Temperatures, the HTL result, where the central line marks the HTL expectation for the EoS with the band resulting from (large) variations of the renormalization scale. (Right) Preliminary results for the pressure, errors indicate the Stefan-Boltzmann value. All errors are statistical only.

one second old. Present heavy-ion experiments (LHC@CERN, RHIC@BNL, and the upcoming FAIR@GSI) create, for a brief moment when two heavy nuclei collide, the extreme conditions of the early universe, allowing us to study this transition and the properties of the quark-gluon plasma some 13 Billion years later.

Considerable theoretical effort is invested attempting to describe these experiments, from collision to detector signals. Here, the Equation of State (EoS) of QCD [4] is a central ingredient for a complete understanding of the experimental findings. At low temperatures, the EoS can be calculated using the so-called "Hadron Resonance Gas" (HRG) model. At high temperatures, perturbative analyses of QCD become possible (e.g. "Hard Thermal Loop" (HTL) perturbation theory). The intermediate region, from ca. 100 MeV to 1 GeV, can be covered systematically through simulations of Lattice QCD.

Presently available Lattice QCD results for the EoS neglect the effects of the

charm quark, which restrict their region of applicability to temperatures below about 400 MeV. In order to reach larger temperatures, we have set up new simulations which take the charm quark into account, using an improved formulation of Lattice QCD. Our preliminary results illustrate the impact of the charm quark for temperatures larger than 400 MeV (Fig. 2), and make contact with both HRG at low and HTL at large temperatures. The EoS is thus becoming available for the whole temperature region.

Computational Aspects

Simulations of Lattice QCD generally happen in three main phases. In the first phase, an ensemble is generated through a Markov process. This phase is thus usually scaled to a large number of cores, minimizing "wall-clock" time. We have, so far, scaled our production code used for the ensemble generation up to 1.8 million parallel threads, running at a sustained performance of over 1.6 PFlop/s (Fig. 3). The second production stage then analyzes the individual "configurations" that

constitute an ensemble one by one. Since an ensemble can contain 1,000 configurations and more, this greatly reduces the need for scaling to a large number of cores. Therefore, we can optimize production at this stage for efficiency (which reaches up to 70% of the hardware peak flop rate) and queue throughput. Physics results are then extracted in the final step of the calculation, which, with our involved blind analysis procedure [1,3], requires a small compute cluster on its own.

Outlook

Simulations of Lattice Quantum Chromodynamics have reached per mil level precision. By now, we are able to reproduce even intricate details of the particle spectrum, such as the neutron-proton and other mass splittings at high precision. The inclusion of Quantum Electrodynamics was essential to archive this level of accuracy, rendering calculations of the combined theory possible in cases where this is needed in the future. Beyond conceptual advances, the correct reproduction of the mass splittings found in nature provides further strong evidence that Quantum Chromodynamics correctly accounts for the properties of strongly interacting matter. Moving beyond the mass spectrum, we can now calculate the properties of the

early universe transition between, and the properties of matter in, the quark and the hadron epoch, starting 10^{-12} seconds after the Big Bang.

Acknowledgements

We thank all the Members of the Budapest-Marseille-Wuppertal and Wuppertal-Budapest collaborations for the fruitful and enjoyable cooperation.

Our simulations require substantial computational resources. We are indebted to the infrastructure and computing time made available to us by the Gauss Centre for Supercomputing and the John von Neumann Institute for Computing.

Further support for these projects was provided by the DFG grant SFB/TR55, the PRACE initiative, the ERC grant (FP7/2007-2013/ERC No 208740), the Lendület program of HAS (LP2012-44/2012), the OCEVU

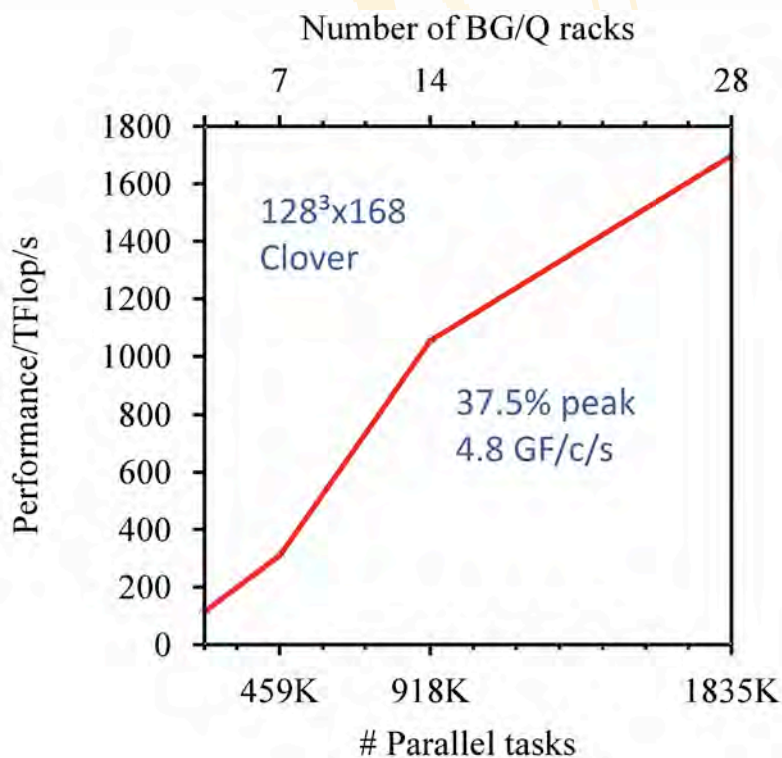


Figure 3: Strong scaling analysis of the Lattice QCD simulation code, reaching a sustained performance of over 1.6 Flop/s on the Blue Gene/Q at JSC, Jülich.

Labex (ANR-11-LABX-0060), the A*MIDEX project (ANR-11-IDEX-0001-O) and the GENCI-IDRIS Grand Challenge grant 2012 "StabMat" as well as grant No. 52275. The computations were performed on JUQUEEN and JUROPA at Forschungszentrum Jülich, on Turing at the Institute for Development and Resources in Intensive Scientific Computing in Orsay, on SuperMUC at Leibniz Supercomputing Centre in München, on Hermit and Hornet at the High Performance Computing Center in Stuttgart, and on local machines in Wuppertal and Budapest. Computing time on the JARA-HPC Partition is gratefully acknowledged.

References

[1] Dürr, S. et al.

Ab-Initio Determination of Light Hadron Masses, *Science* 322, 1224-1227, 2008

[2] Wilczek, F.

Mass Without Mass I: Most of Matter, *Physics Today* 52, 11, 1999

[3] Borsanyi, S., et al.

Ab initio calculation of the neutron-proton mass difference, *Science* 347, 1452-1455, 2015

[4] Borsanyi S, et al.

The QCD equation of state with dynamical quarks, *JHEP* 1011, 077, 2010

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

- Stefan Krieg^{1, 2}
- Zoltan Fodor^{1, 2, 3}

¹ JSC & JARA-HPC, Forschungszentrum Jülich, Germany

² University of Wuppertal, Germany

³ Eötvös Lorand University, Hungary

Advance Visualisation of Seismic Wave Propagation and Speed Model

With the induction of the Virtual Reality and Visualisation Centre (V2C) at Leibniz Supercomputing Centre (LRZ), many domain specialists have approached LRZ to leverage on the immersive projection technology. Large datasets can now be stereoscopically displayed and specialists can interact with their complex datasets intuitively. Seismologists is one group of domain specialists that have benefited from the use of this virtual reality (VR) technology. To allow a deep insight into the simulated data, the seismologists make use of VR installations like a 5 sided projection installation based on the concepts of a Carolina Cruz-Neira's CAVE Automated Virtual Environment (CAVE)¹. In this article, the CAVE like installation at LRZ will be referred to as the CAVE for convenience.

Seismology

Seismology is a field of science where earthquakes and propagation of seismic waves that move through the Earth and on its surface are scientifically studied. The geological structure and physical properties of

the Earth have a significant impact on how the waves propagate. Depending on the material properties, e.g. density and elasticity of the medium, the speed at which the seismic waves propagate differs. As seismic waves travel through different materials, they can be reflected, refracted, dispersed, diffracted and/or attenuated. A wave speed model is used to approximate and represent the different materials that compose the Earth structure and the speed at which the waves travel through each of these regions. Having a wave speed model that accurately represents a region implies understanding more about the Earth's interior and also the possibility to determine the earthquake location more accurately.

Computation

A forward simulation of a North Italian earthquake was computed to analyse the wave speed model and how its features affect the propagation of seismic waves. In particular, the considered event was the mainshock of a seismic sequence that struck North Italy in




Figure 1: The Wave Speed Model

¹ CAVETM, a registered trademark of the University of Illinois' Board of Trustees




Figure 2: Slices and Contours of the Wave Speed Model

2012. It occurred on May 20, 2012, at 02:03 UTC with a local magnitude of 5.9, e.g. [5], and a hypocentre location of 44.98°N, 11.23°E and 6.3 km depth relocated by the Istituto Nazionale di Geofisica e Vulcanologia (INGV). The source parameters were given by the centroid moment tensor solution calculated using the Time Domain Moment Tensor technique implemented at INGV [6]. The considered region in North Italy is characterised by a large sedimentary basin, a significant presence of fluid and strong heterogeneities, leading to remarkable site effects and liquefaction phenomena. To run the simulation, this region was discretised by constructing a conforming, unstructured mesh of hexahedral elements that covered a volume of ~350 km in longitude, ~230 km in latitude and 60 km in depth. The mesh was composed of ~1.6 million of hexahedral elements and honoured the free surface topography. Then, to represent the geological structure of the region, a wave speed model was constructed that was based on a 3D tomographic model for the Italian peninsula derived by the inversion of P-wave travel time measurements [2]. The P-wave wave speed model was scaled into the speed of S-waves and the density using a scaling relation. The model included the signature of the 8 km thickness sedimentary basin. The combination of the described mesh and wave speed

model allows to produce simulations that resolve a minimum period of about 4 seconds. The forward simulation was performed using SPEC-FEM3D_Cartesian [3], a very popular wave propagation simulation code, which employed the continuous Galerkin spectral-element method with arbitrary unstructured hexahedral meshes. The computation was submitted to SuperMUC at LRZ, utilising 500 cores to generate 1 minute of seismograms with 114 seismic stations. Since this application was well parallelised, it managed to take good advantage of the parallel architecture of SuperMUC. The computation, including the generation of additional output files for visualisation purposes, completed in less than 30 minutes. The corresponding wave speed model in Fig. 1, the slices and contours of this model in Fig. 2 and the animated surface propagation of seismic waves in Fig. 3 were then ported and displayed in the CAVE. These figures are renderings for illustration purposes of the actual displays in the CAVE. The interactive CAVE displays are particularly useful to allow the visualisation of the contours and slices of the wave speed model. It allows the scientists to literally step into the model, intuitively visualise the critical layers and associate them to the simulated speed of the wave propagation. The computed synthetic seismograms can be compared to the observed seis-

mograms to calculate the discrepancy quantified by a misfit function. The misfit is then used to correct the wave speed model. With each such iteration, an increasingly accurate wave speed model can be represented and bring the seismologists one step closer to understanding the Earth's interior.

Immersive Visualisation

A visualisation prototype has been developed from scratch which smoothly integrates into the seismologists' workflow. The prototype is based on OpenSG [4], an open source scene graph, which supports multi-display installations. Additionally VRPN [7] is used as an abstraction layer for a large set of input devices. Thus it is possible to display the wave speed model, contours of it and a sequence of surface propagation on powerwalls, CAVE-like installations and Head-Mounted Displays (HMDs). Providing stereoscopic real-time interaction turning the model allows building a mind map of the displayed structures. The advantage of the HMDs and surround displays lies in the intuitive data analysis. By providing head tracking it is easily possible to change the perspective on the data in a fluent and natural way fostering exploration of the displayed data discovering unexpected connections inside the data.

Future Work

The benefits of using CAVE to analyse their models has motivated the seismologists to integrate their existing tools to create data products that are useful for CAVE displays. The seismologists aim to include the CAVE visualisation items as an additional end-product of their scientific workflow. This will allow them to produce multiple types of representation, highlighting features and properties depending

from their interests. This level of automation could allow a dynamic annotation of the visualisation elements to enrich the description and the metadata associated to the specific workflow run, experiment or study. From the perspective of advance visualisation, the data exploration is currently focused on a single user setup but it is easily possible to integrate network communication and allow collaborative data analysis. This is also something to explore in the future to allow the seismologists not only to analyse the models independently but also collaboratively with their fellow seismologists.

Acknowledgement

This work was conducted in cooperation with the team from the EU project, Virtual Earthquake and seismology Research Community in Europe e-science environment (VERCE). The output data required for the visualisation was generated by the execution of scientific workflows that were built around SPECFEM3D_Cartesian. These workflows were submitted and controlled within a virtual research environment (VRE), which is available to seismologists via the VERCE project as the VERCE platform. The platform offers ease and homogeneous access to the data and workflow management systems by connecting numerous distributed computational resources. These services are fully integrated within an interactive science gateway that is tailored to the needs of the seismologists. It enables them to experiment with different wave speed models and simulation parameters in an assisted environment.

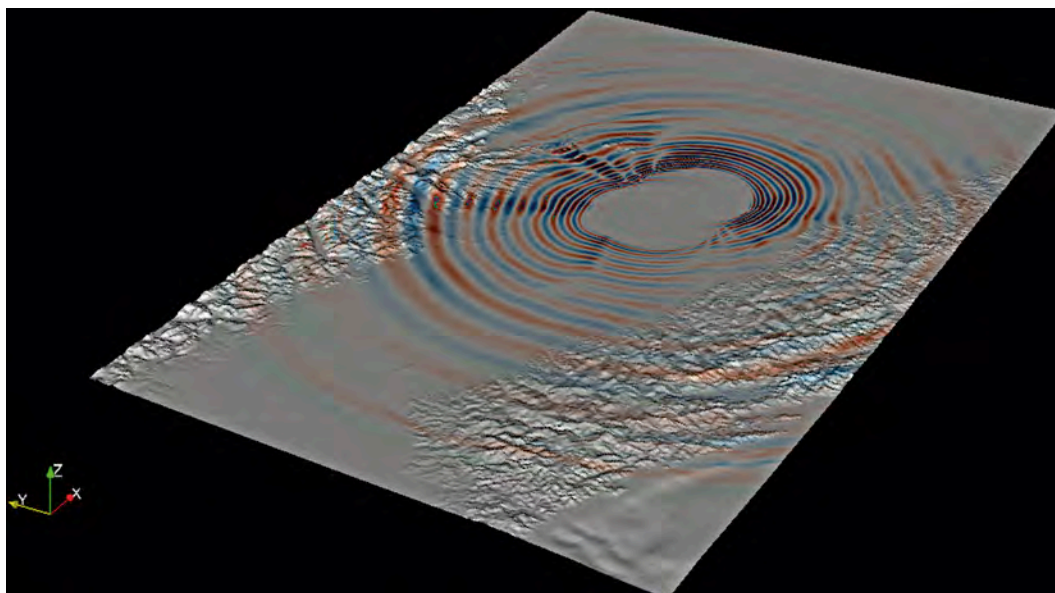


Figure 3: Surface Propagation of Seismic Waves

References

- [1] Cruz-Neira, C., Sandin, D. J., Defanti, T. A., Kenyon, R. V., and Hart, J. C. The cave: Audio visual experience automatic virtual environment, *Communications of the ACM* 35, 6, 64–72., 1992.
- [2] Stefano, R. D., Kissling, E., Chiarabba, C., Amato, A., and Giardini, D. Shallow subduction beneath Italy: three-dimensional images of the Adriatic-European-Tyrrhenian lithosphere system based on high-quality P wave arrival times., 114, B05305, 2009.
- [3] Peter, D., Komatitsch, D., Luo, Y., Martin, R., Le Goff, N., Casarotti, E., Le Loher, P., Magnoni, F., Liu, Q., Blitz, C., Nissen-Meyer, T., Basini, P., and Tromp, J. Forward and adjoint simulations of seismic wave propagation on fully unstructured hexahedral meshes, *Geophysical Journal International* 186, 721–739, 2011.
- [4] Reiners, D. Opensg: A scene graph system for flexible and efficient realtime rendering for virtual and augmented reality applications, Ph.D. thesis, Technische Universität Darmstadt, Mai 2002.
- [5] Scognamiglio, L., Margheriti, L., Mele, F. M., Tinti, E., Bono, A., De Gori, P., Lauciani, V., Lucente, F. P., Mandiello, A. G., Marcocci, C., Mazza, S., Pintore, S., and Quintiliani, M. The 2012 Pianura Padana Emiliana seismic sequence: locations, moment tensors and magnitudes, *Annals of Geophysics* 55, 4, 2012.
- [6] Scognamiglio, L., Tinti, E., and Michelini, A. Real-Time Determination of Seismic Moment Tensor for the Italian Region, *Bulletin of the Seismological Society of America* 99, 4, 2223–2242, 2009.
- [7] Taylor, R. M. II, Hudson, T. C., Seeger, A., Weber, H., Juliano, J., and Helser, A. T. Vrpn: a device-independent, network-transparent vr peripheral system, *Proceedings of the ACM symposium on Virtual reality software and technology* (New York, NY, USA), VRST '01, ACM, pp. 55–61, 2001.
- [8] VERCE - <http://www.verce.eu>
- contact: Siew Hoon Leong,
siew-hoon.leong@lrz.de
Christoph Anthes,
christoph.athanes@lrz.de
Federica Magnoni,
federica.magnoni@ingv.it
Alessandro Spinuso,
xspinuso@knmi.nl
Emanuele Casarotti,
emanuele.casarotti@ingv.it
- Siew Hoon Leong^{1,2}
 - Christoph Anthes^{1,2}
 - Federica Magnoni³
 - Alessandro Spinuso⁴
 - Emanuele Casarotti³
- ¹ Leibniz Supercomputing Centre, Germany
- ² Ludwig-Maximilians-Universität München, Germany
- ³ Istituto Nazionale di Geofisica e Vulcanologia, Italy
- ⁴ Royal Netherlands Meteorological Institute, The Netherlands

A new Neutrino-Emission Asymmetry in forming Neutron Stars

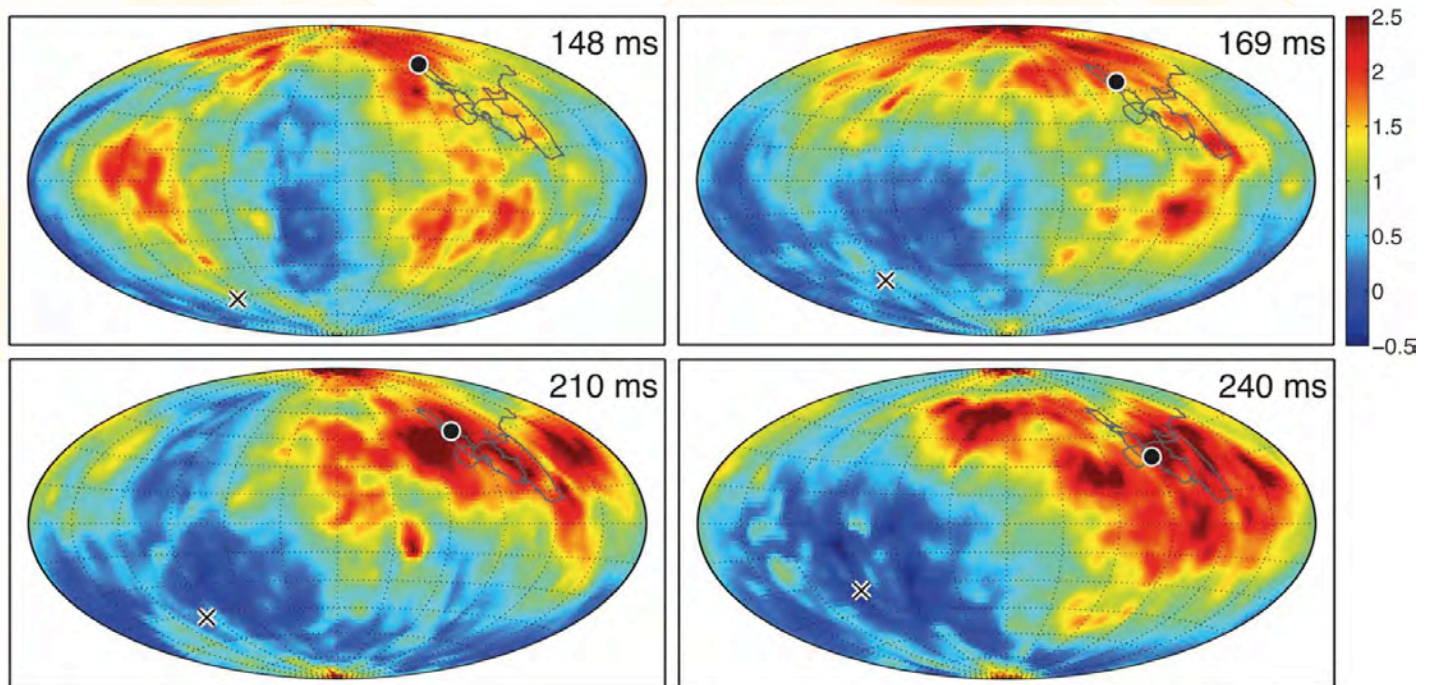


Figure 1: Evolution of the neutrino emission asymmetry in a collapsing star of 11.2 solar masses. The ellipses represent the whole surface of the nascent neutron star (analog to world maps as planar projections of the Earth's surface). Red and yellow mean a large excess of electron neutrinos compared to electron antineutrinos, normalized to the average, blue means a low excess or even deficit of electron neutrinos. The images show the merging of smaller patches that are present at 0.148 seconds (upper left panel) to a clear hemispheric (dipolar) anisotropy at 0.240 seconds (lower right panel). The dot and cross indicate the emission maximum and minimum, the dark-grey line marks the path described by a slow motion of the dipole direction.

© Author and The American Astronomical Society

Supernovae are the spectacular explosions that terminate the lives of stars more massive than about nine times our sun. They belong to the most energetic and brightest phenomena in the universe and can outshine a whole galaxy for weeks. They are important cosmic sources of chemical elements like carbon, oxygen, silicon, and iron, which are disseminated in the circumstellar space by the blast wave of the explosion. Supernovae are also the birth places of the most exotic celestial objects: neutron stars and black holes.

Neutron stars contain about 1.5 times the mass of the sun, compressed into

a sphere with the diameter of Munich. Their central density exceeds that in atomic nuclei, gigantic 300 million tons (the weight of a mountain) in the volume of a sugar cube. Neutron stars are formed as extremely hot and dense objects when the central core of the highly evolved, massive star undergoes a catastrophic collapse because it cannot withstand its own gravitational weight any longer. Newly born neutron stars cool by the intense emission of neutrinos, ghostly elementary particles that hardly interact with matter on earth but that are produced in gigantic numbers at extreme temperatures and densities. These neutrinos are thought to trigger the violent disruption of the

dying star in the supernova if even only one percent of their huge total energy can be tapped to heat the stellar mantle that surrounds the forming neutron star.

Because neither experiments nor direct observations can reveal the processes at the centers of exploding stars, highly complex numerical simulations are indispensable to develop a deeper and quantitative understanding of this hypothetical "neutrino-driven explosion mechanism", whose solid theoretical foundation is still missing. The computational modeling must be done in three dimensions (3D), simulating the whole star, because turbulent flows as well as large-scale deformation play a crucial role in enhancing the neutrino-matter interactions. This requires the solution not only of the fluid dynamics problem in a strong-gravity environment including a description of the properties of neutron-star matter and of nuclear reactions. In particular the neutrino propagation and processes pose a grand computational challenge, because besides the three spatial dimensions there are additional three dimensions for neutrino energy and direction of motion. Not even the biggest existing supercomputers can solve such a six-dimensional, time-dependent transport problem in perfect generality.

Simulation Methods

In this project the Supernova Simulation Group at the Max Planck Institute for Astrophysics (MPA) simulates the gravitational collapse of stellar cores to neutron stars and the onset of the supernova explosion with the PROMETHEUS-VERTEX code for multi-dimensional hydrodynamical simulations including a highly sophisticated description of three-flavor neutrino

transport and neutrino-matter interactions with full energy dependence. While the former is treated with an explicit, higher-order Godunov-type scheme, the latter is solved by an implicit integrator of the neutrino energy and momentum equations, supplemented by a state-of-the-art set of neutrino-reaction kernels and a closure relation computed from a simplified model-Boltzmann equation. Consistent with the basically spherical geometry the equations are discretized on a polar coordinate grid, and the computational efficiency is enhanced by the use of an axis-free Yin-Yang implementation and a time- and space-variable radial grid. The thermodynamics and changing chemical composition of the stellar medium are determined by high-dimensional equation-of-state tables and nuclear burning at non-equilibrium conditions.

Employing mixed MPI-OpenMP parallelization, our ray-by-ray-plus approximation of multidimensional transport allows for essentially linear scaling up to tested processor-core numbers of more than 130,000. In production applications we are granted access to up to 16,000 processor cores (allowing for two degrees angular resolution) and the code typically reaches 10–15% of the peak performance. The computational load is strongly dominated by the complexity of the neutrino transport. Despite remaining approximations, a single supernova run for an evolution period of roughly half a second takes several months of uninterrupted computing and up to 50 million core hours, producing more than hundred Tbytes of data. PRACE and GAUSS contingents on the SuperMUC infrastructure of LRZ have enabled the successful execution of this important and timely project of fundamental research in theoretical stellar astrophysics.




Figure 2: Bubbles of "boiling" gas surrounding the nascent neutron star (invisible at the center). Despite the highly time-variable and dynamical pattern of plumes of hot, rising gas and inflows of cooler matter, the neutrino emission develops a hemispheric asymmetry (see Fig. 1) that remains stable for periods of time much longer than the life time of individual bubbles.

© Author and The American Astronomical Society

Results

On the way to producing the first-ever 3D explosion models with a highly sophisticated treatment of the neutrino physics, the MPA team made a stunning and unexpected discovery [1]: The neutrino emission develops a strong dipolar asymmetry (Fig. 1). Neutrinos and their anti-particles are not radiated equally in all directions but with largely different numbers on opposite hemispheres of the neutron star. As expected, the neutrino emission starts out to be basically spherical except for smaller variations over the surface (see Fig. 1, upper left panel). These variations correspond to higher and lower temperatures associated with violent "boiling" of hot matter inside and around the newly formed neutron star, by which bubbles of hot matter rise outward and flows of cooler material move inward (Fig. 2). After a short

while, however, the neutrino emission develops clear differences in two hemispheres. The initially small patches merge to larger areas of warmer and cooler medium until the two hemispheres begin to radiate neutrinos unequally. A stable dipolar pattern is established, which means that on one side more neutrinos leave the neutron star than on the other side. Observers in different directions thus receive different neutrino signals. While the directional variation of the summed emission of all kinds of neutrinos is only some per cent (Fig. 3a), the individual neutrino types (for example electron neutrinos or electron antineutrinos) show considerable contrast between the two hemispheres with up to about 20 per cent deviations from the average (Figs. 3b,c). The directional variations are particularly pronounced in the difference between electron neutrino

and antineutrino fluxes (Fig. 1, lower right panel), the so-called lepton number emission.

The possibility of such a global anisotropy in the neutrino emission was not predicted and its finding in the first-ever detailed three-dimensional simulations of dynamical neutron-star formation comes completely unexpectedly. The team of astrophysicists named this new phenomenon "LESA" for Lepton-Emission Self-sustained Asymmetry [1], because the emission dipole seems to stabilize and maintain itself through complicated feedback effects despite the violent bubbling motions of the "boiling" hot and cooler gas, which lead to rapidly changing structures in the flow around and inside the neutron star (Fig. 2).

Outlook

The new, stunning neutrino-hydrodynamical instability that manifests itself in the LESA phenomenon is not yet well understood. Much more research is needed to ensure that it is not an artifact produced by the highly complex numerical simulations. If it is physical reality, this novel effect would be a discovery truly based on the use of modern supercomputing possibilities and not anticipated by previous theoretical considerations.

If LESA happens in collapsing stellar cores, it will have important consequences for observable phenomena connected to supernova explosions. A directional variation between electron neutrino and antineutrino emission will lead to differences of the chemical element production in the supernova ejecta in different directions.

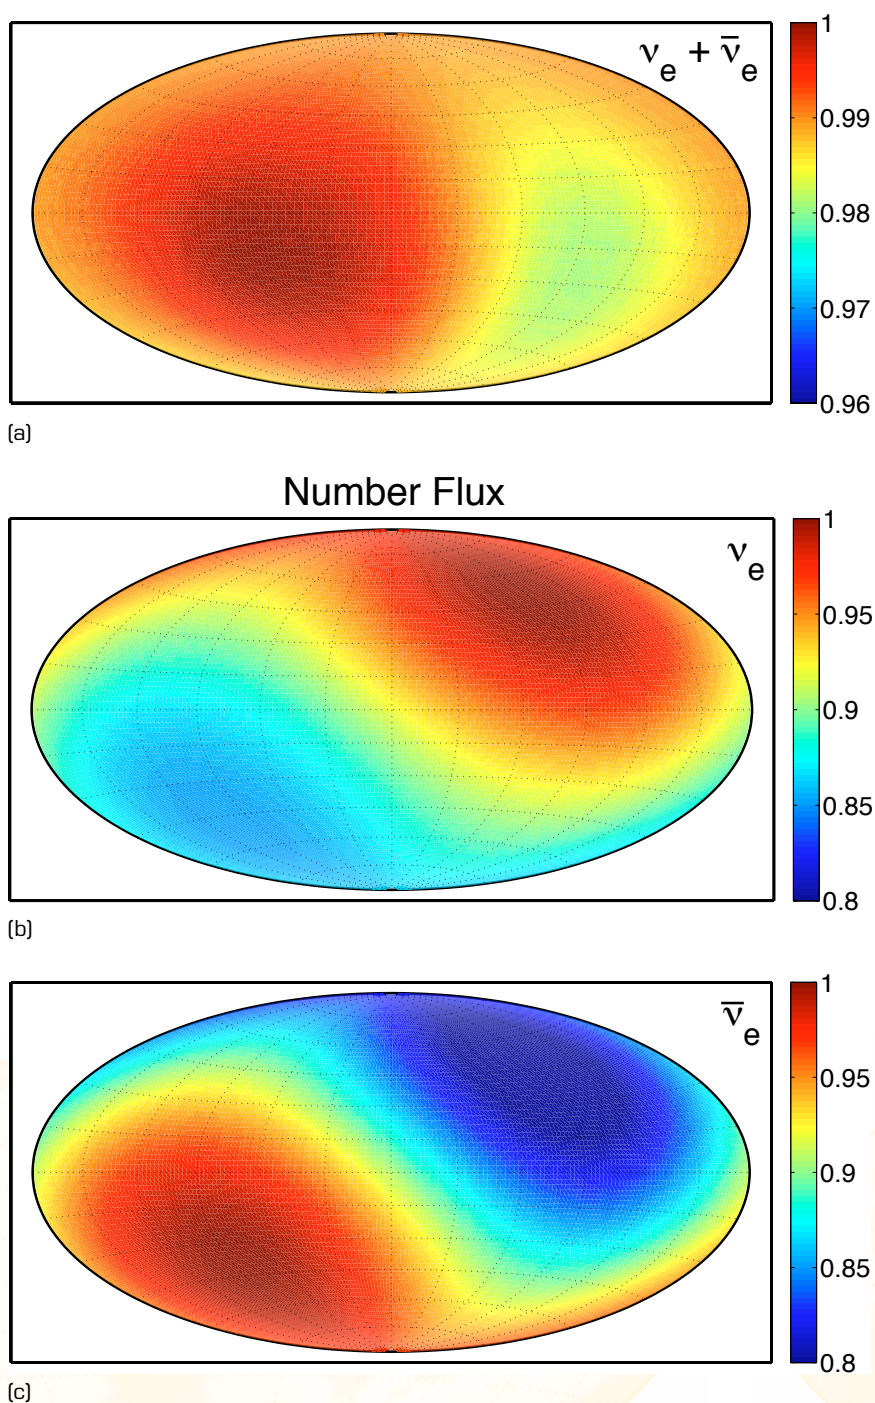


Figure 3: Neutrino emission asymmetry as observable over a period of 0.1 seconds. Analog to Fig. 1 the ellipses show all possible viewing directions. Observers in the red regions see the highest emission, whereas those in the blue areas receive lower emission. While the hemispheric difference of the emission of electron neutrinos plus antineutrinos is only one to two per cent (a), the differences of electron neutrinos (b) and antineutrinos (c) individually amount to up to 20 per cent of the maximum values with extrema on opposite sides.

© Author and The American Astronomical Society

Moreover, the global dipolar anisotropy of the neutrino emission carries away momentum and thus imparts a kick to the nascent neutron star in the opposite direction. Also the neutrino signal arriving at Earth from the next supernova event in our Milky Way must be expected to depend on the angle from which we observe the supernova [2,3].

Acknowledgments

The author is grateful for GAUSS and PRACE computing time on LRZ's SuperMUC and for project funding by the European Research Council through grant ERC-AdG No. 341157-COCO-2CASA.

References

- [1] **Tamborra, I., Hanke, F., Janka, H. Th., Müller, B., Raffelt, G. G., Marek, A.**
Self-sustained asymmetry of lepton-number emission: A new phenomenon during the supernova shock-accretion phase in three dimensions, *Astrophysical Journal* 792, 96, 2014, <http://arxiv.org/abs/1402.5418>
- [2] **Tamborra, I., Hanke, F., Müller, B., Janka, H. Th., Raffelt, G.**
Neutrino signature of supernova hydrodynamical instabilities in three dimensions *Physical Review Letters* 111, 121104, 2013, <http://arxiv.org/abs/1307.7936>
- [3] **Tamborra, I., Raffelt, G., Hanke, F., Janka, H. Th., Müller, B.**
Neutrino emission characteristics and detection opportunities based on three-dimensional supernova simulations *Physical Review D* 90, 045032, 2014, <http://arxiv.org/abs/1406.0006>

• Hans-Thomas
Janka

Max Planck
Institute for
Astrophysics
Garching,
Germany

Internet Links

Institute web page:
www.mpa-garching.mpg.de

Core-collapse supernova data archive:
www.mpa-garching.mpg.de/ccsnarchive

contact: Hans-Thomas Janka,
thj@mpa-garching.mpg.de

Precision Physics from Simulations of Lattice Quantum Chromodynamics

In the last decade, simulations of Lattice Quantum Chromodynamics (QCD) have reached a new level of precision. By now we are able to compute per mill effects of the particle mass spectrum of QCD, by combining Lattice QCD with Lattice Quantum Electrodynamics (QED). This dramatic increase of precision was made possible by the combination of new and more powerful machines and new lattice actions and simulation algorithms.

With these new methods at hand, we have computed the neutron-proton and other mass splittings from first principles [1,2], studied the range of applicability of chiral perturbation theory [3] and extracted the corresponding low energy constants [4]. Furthermore, we have computed the Equation of State of QCD [5], the freeze-out parameters of the cooling quark-gluon-plasma produced in heavy ion experiments [6,7], including the flavor dependency of the freeze-out temperature [8].

In the following, we will briefly discuss these different results.

Precision Spectrum of the Standard Model

Moving beyond the per cent level precision of our 2008 calculation of the light hadron spectrum [9] (see also Fig. 1), requires the addition of two formally per cent level effects in the calculation: Quantum Electrodynamics (QED), due to the magnitude of the fine structure constant $\alpha \approx 1/137$, and the up/down




Figure 1: Ab initio calculation of the light hadron spectrum of QCD [9]. Blue circles indicate experimental input required to fix the quark masses and overall scale of the calculation. Red points indicate predictions of QCD. Experimental widths are grey, with the central value indicated by a black line.

quark mass splitting $(m_d - m_u)/\Lambda_{\text{QCD}} \approx < 0.01$ (which are of almost the same magnitude). In particular including QED poses a range of conceptual issues [10].

Incorporating these effects allows us to calculate the neutron-proton and other mass splittings. This implies calculating a per mill level difference between particle masses that were previously available at per cent precision.

However, by making use of statistical correlations between (e.g., neutron and proton) propagators, it is possible to calculate the splitting directly.

A first step in this direction was our result [2] using quenched QED, i.e. neglecting the effect of QED in quark loops. Such a calculation cannot be fully satisfactory. For our recent result




Figure 2: Mass splittings. Shown are individual masses from Figure 1, with the results [1] on the mass splittings shown in 40x magnification.

on the mass splittings [1], we took the next step and included these unquenching effects, which required us to develop new simulation techniques and to calculate finite volume effects analytically [10]. The precision of our results is illustrated in Fig. 2.

Confronting Chiral Perturbation Theory

Chiral perturbation (XPT) theory is an effective theory that is used to compute low-energy properties of QCD. It does not describe the dynamics of quarks and gluons, but rather the dynamics of hadrons. Since it is an expansion around zero quark masses (and momenta), it is in principle unclear if it does in deed apply to the physical

world, where the quark masses are non-zero. XPT has also been used in the past to extrapolate from heavier than physical quark masses (due to the costs of simulating with physical parameters) down to the physical mass point.

In 2010, we were the first to compute the quark masses [11] using ensembles with physical quark mass parameters. With these new ensembles that now include the physical point, we could check the applicability of XPT for different observables. Furthermore, we could calculate several (low energy) constants (LECs) of XPT that are not fixed by the theory. This is illustrated in Fig. 3 for one particular LEC, where also the limited applicability of XPT for heavy pion masses is clearly visible.

Lattice QCD for heavy Ion Experiments

Heavy ion experiments (such as RHIC at BNL, LHC at CERN or the upcoming FAIR at GSI) require the Equation of State of QCD (EoS) as a central ingredient for the understanding of the evolution of the quark-gluon plasma generated in collisions of heavy nuclei. The only tool available that allows for a calculation of the EoS from first principles is Lattice QCD.




Figure 3: Deviation of the low energy constant B in the RGI scheme from the value at physical mass parameters. Clearly, for simulated pion masses much larger than 300 MeV, the applicability of XPT (SU(2) in this case) to (N_f=2+1) QCD is limited.

In 2013 [5] we presented the first full calculation of this central quantity, by carefully balancing and controlling the different uncertainties and taking the continuum limit (vanishing lattice spacing). Recently, our findings were corroborated by the independent hotQCD collaboration, as shown in Fig. 4. This settled a long standing discrepancy between the collaborations, dating back beyond our first continuum estimate of 2010 and became possible with the most recent results from hotQCD.

By studying so-called generalized susceptibilities (derivatives of the partition function with respect to the chemical potentials), it is also possible to directly match Lattice QCD results at finite temperature and chemical potentials to experimental data.

In heavy-ion experiments, the nuclei rarely hit head on, and the "cross-section volume" is not constant on an event-by-event basis. This implies that otherwise conserved quantities like the total electric charge or the baryon number, as measured from the collision products, fluctuate as well. If additional cuts on the experimental data are applied, this experimental setting can be described using a grand-canonical ensemble and thus be simulated using Lattice QCD. The moments of the distributions of the conserved charges found in experiment can then be matched to Lattice QCD results, when appropriate ratios of moments are taken, such that the interaction volume cancels out. In this way it is possible to extract the experimental freeze-out parameters, i.e. the temperature and chemical potential at "chemical freeze-out" (last inelastic scattering of hadrons before detection).




Figure 4: Comparison of independent results for the EoS of QCD (taken from Phys. Rev. D90 (2014) 094503 by hotQCD). Shown are our results ("stout") and the results of the hotQCD collaboration ("HISQ").

We calculated these parameters [6] using preliminary data from the STAR collaboration. Using their latest experimental results, our findings for the freeze-out parameters are consistent [7], independent whether electrical charge or baryon number fluctuations were used for the matching. Furthermore, we see indications that hadrons with different quark flavors may have different freeze-out parameters [8], a finding that could be verified experimentally at LHC.

Acknowledgements

We thank all the Members of the Budapest-Marseille-Wuppertal and Wuppertal-Budapest collaborations for the fruitful and enjoyable cooperation.

Our simulations require substantial computational resources. We are indebted to the infrastructure and computing time made available to us by the Gauss Centre for Supercomputing and the John von Neumann Institute for Computing.

Further support for these projects was provided by the DFG grant SFB/TR55, the PRACE initiative, the ERC grant (FP7/2007-2013/ERC No 208740),

the Lendület program of HAS (LP2012-44/2012), the OCEVU Labex (ANR-11-LABX-0060), the A*MIDEX project (ANR-11-IDEX-0001-O) and the GENCI-IDRIS Grand Challenge grant 2012 "StabMat" as well as grant No. 52275. The computations were performed on JUQUEEN and JUROPA at Forschungszentrum Jülich, on Turing at the Institute for Development and Resources in Intensive Scientific Computing in Orsay, on SuperMUC at Leibniz Supercomputing Centre in München, on Hermit and Hornet at the High Performance Computing Center in Stuttgart, and on local machines in Wuppertal and Budapest. Computing time on the JARA-HPC Partition is gratefully acknowledged.

References

- [1] **Borsanyi, S., et al.**
Ab initio calculation of the neutron-proton mass difference, *Science* 347, 1452-1455, 2015.
- [2] **Borsanyi, S., et al.**
Isospin splittings in the light baryon octet from lattice QCD and QED, *Phys. Rev. Lett.* 111, 252001, 2013.
- [3] **Dürr, S., et al.**
Lattice QCD at the physical point meets SU(2) chiral perturbation theory, *Phys. Rev. D* 90, 114504, 2014.
- [4] **Borsanyi, S., et al.**
SU(2) chiral perturbation theory low-energy constants from 2+1 flavor staggered lattice simulations, *Phys. Rev. D* 88, 014513, 2013.
- [5] **Borsanyi, S., et al.**
Full result for the QCD equation of state with 2+1 flavors, *Phys. Lett. B* 730, 99, 2014.
- [6] **Borsanyi, S., et al.**
Freeze-out parameters: lattice meets experiment, *Phys. Rev. Lett.* 111, 062005, 2013.
- [7] **Borsanyi, S., et al.**
Freeze-out parameters from electric charge and baryon number fluctuations: is there consistency?, *Phys. Rev. Lett.* 113, 052301, 2014.
- [8] **Borsanyi, S., et al.**
Is there a flavor hierarchy in the deconfinement transition of QCD?, *Phys. Rev. Lett.* 111, 202302, 2013.
- [9] **Dürr, S., et al.**
Ab-Initio Determination of Light Hadron Masses, *Science* 322, 1224-1227, 2008.
- [10] See our contribution "Quarks and Hadrons - and the spectrum in between" to this inside issue.
- [11] **Dürr, S., et al.**
Lattice QCD at the physical point: Simulation and analysis details, *JHEP* 1108 (2011) 148, and Lattice QCD at the physical point: light quark masses, *Phys. Lett. B* 701, 265-268, 2011.

contact: Zoltan Fodor,
fodor@physik.uni-wuppertal.de
Stefan Krieg,
s.krieg@fz-juelich.de

• Zoltan Fodor^{1, 2, 3}

• Stefan Krieg^{1, 2}

¹ JSC & JARA-HPC
Forschungs-
zentrum Jülich,
Germany

² University of
Wuppertal,
Germany

³ Eötvös Lorand
University,
Hungary

Turbulence in the Planetary Boundary Layer

The Planetary Boundary Layer (PBL) is the lowest part of the atmosphere, the part that is in contact with the surface and that feels the cycle of day and night. The PBL is important not only for being the place where we live, but also because this relatively shallow layer – about 2 kilometers deep or less – regulates the exchange of mass, momentum and energy between the rest of the atmosphere and the land and the oceans. This exchange across the PBL strongly depends on how turbulence mixes the air. The correct representation of turbulence is then key in weather prediction and climate research, but turbulence remains poorly understood in some relevant PBL regions and regimes. By providing a faith-

ful description of turbulence across all relevant scales, without any turbulence model, Direct Numerical Simulation (DNS) is opening new avenues to advance this understanding. In this brief article, we provide two examples that illustrate this new development.

Turbulence Collapse in the Stable Boundary Layer

One common PBL regime is a wind-driven boundary layer with a stable density stratification. Such a situation develops for instance at night or when a mass of air is advected over a relatively cold surface – layers of heavy fluid form below layers of lighter fluid due to the heat loss towards the earth's surface. Stable stratification weakens the




Figure 1: Horizontal cross-sections near the surface showing the magnitude of the angular velocity of the fluid particles (this angular velocity is one of the defining properties of turbulence), obtained from a DNS of an Ekman layer using a grid size 3072x6144x512. In a neutrally stratified case (left panel), we appreciate turbulence homogeneously distributed everywhere. In contrast, during turbulence collapse in a stably stratified case (right panel), the morphology is completely different: turbulence is reduced to turbulence pockets that are embedded in an otherwise smooth variation of flow properties – turbulence is intermittent in space.

turbulence as heavy fluid particles are lifted and kinetic energy is converted into potential energy. If strong enough, stable stratification can even induce a turbulence collapse, the flow becomes laminar (smooth, without turbulent fluctuations) and the mixing rates across the PBL drop significantly. The understanding and the correct representation in atmospheric models of such a process is a long-standing issue.

So far, simulations could not reach into the strongly stratified regime and the analysis of turbulence collapse has been based on field measurements, to a large extent [1]. This paradigm has changed. Using DNS and a stably stratified Ekman layer as a simplified physical model of the PBL, we have succeeded to reproduce, in a single configuration, for the first time, the three stratification regimes found in nature: weakly, intermediately and strongly stratified [2].

An Ekman layer is a boundary-layer type of flow that results from the balance between the Coriolis force and the pressure force, along with the boundary condition that the flow at the surface has the same velocity as the surface itself (the so-called no-slip boundary condition). The density difference between the surface and the free flow above the boundary layer can also be controlled. This physical model is fully characterized by only two non-dimensional parameters: the Reynolds number, which measures the relative strength between the inertia forces and the viscous forces, and the Richardson number, which measures the relative strength between the buoyancy forces, caused by the density variations, and the inertia forces. By systematically varying this second parameter, we can reproduce the different regimes found in nature (see Fig. 1). This work has demonstrated that DNS has become a suitable tool to study turbulence collapse in the PBL under controlled conditions.




Figure 2: Vertical cross-sections showing the liquid content (mass of the liquid divided by the total mass of the fluid mixture) and the buoyancy, obtained from a DNS of a cloud-top mixing layer using a grid size $3072 \times 3072 \times 2216$. The numbers within parentheses correspond to the first research flight of the DYCOMS-II field campaign, used as reference in the DNS study. The turbulence is driven by the convective instability caused by the evaporation of the droplets near the cloud interface.

As a first application, we have shown that turbulence collapse need not be an on-off process in time but can occur intermittently in space without the need of external factors that induce such an intermittency, like surface heterogeneity – intermittency is intrinsic to stably stratified turbulence. It suffices that wave-like, large-scale structures (several tens of times the boundary-layer depth) have space and time to develop (see Fig. 1). This result helps explain the difficulty to obtain spatial intermittency in simulations, because we need to retain these large scales and, simultaneously, we need to resolve the small-scale turbulence inside the turbulence regions.

Catalytic Effect of Wind Shear at the Cloud Top

Another example of the relevance of turbulence in the PBL is the dependence of marine-boundary clouds on the representation of turbulent mixing: Differences among climate models in the representation of this dependence explain about half of the model spread in the surface temperature response to increasing CO₂ [3]. Part of the difficulty lies in representing correctly entrainment, or, more generally, small-scale turbulence at the PBL top. Part of the difficulty consists as well in understanding better the role of turbulence in clouds [4].

DNS allows us to address several aspects of this problem. For instance, we can disentangle key aspects of the role of the stratocumulus-top cooling that is caused by droplet evaporation. This process works as follows (see also Fig. 2). The cloud interface lies in a relatively thin layer across which the buoyancy rapidly increases from the in-cloud value (in white, corresponding

to relatively cold fluid) to the tropospheric value (in red, corresponding to relatively warm fluid). As the cloud mixes with the dry air above it, droplets evaporate and tend to cool down the resulting mixture, as indicated in Fig. 2 by the blue colors in the buoyancy field. For some thermodynamic conditions, like in this example, this evaporative cooling is strong enough to render parcels of fluid colder, and thus heavier, than the in-cloud value. This condition is referred to as buoyancy reversal and leads to convective instability – heavier fluid on top of lighter fluid – and hence turbulence. The key question is: how important is this instability for the cloud-top dynamics?

By deriving an explicit parametrization of the mixing rates from the analysis of DNS data, it was possible to show that, although necessary, buoyancy reversal is not a sufficient condition for a rapid desiccation of the cloud. The reason is that the turbulence generated by buoyancy reversal alone is limited by molecular transport and that is a very slow process compared with other processes acting at the cloud-top [5]. This work settled a three-decades-long discussion on the topic.

Still, evaporative cooling can become important when enhanced by other turbulence sources. DNS analysis has demonstrated that wind shear localized at the cloud-top (a vertical gradient of the horizontal velocity), which is often found in nature, can serve as such a source [6]. The reason is that turbulence generated by buoyancy reversal can locally thin the entrainment zone and thereby enhance the local shear, which in turn enhances mixing and the evaporation of more droplets, creating a positive feedback. The implication of

this finding is twofold. First, evaporative cooling needs indeed to be retained in the analysis of stratocumulus-topped PBLs but always in combination with other processes. Second, wind shear needs to be added to that analysis.

Outlook

We have presented two examples of how DNS is helping us to advance our understanding of turbulence inside the planetary boundary layer. By faithfully representing the flow properties across all the relevant spatial and temporal scales, DNS has led, in some cases, to new insights in the mechanisms that control the PBL dynamics. In some other cases, it has settled long-standing discussions. The idea of using DNS as a research tool to investigate this type of problems was already explored in the early seventies [7], but the computational resources were not yet there – now they are here. The possibilities that are emerging with the current generation of supercomputers, and that will further emerge during the coming decades, are opening new avenues in weather and climate research.

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

References

- [1] **Mahrt, L.**
Stably stratified atmospheric boundary layers. *Annu. Rev. Fluid Mech.* 46, 23-45, 2014.
- [2] **Ansorge, C., Mellado, J. P.**
Global intermittency and collapsing turbulence in the stratified planetary boundary layer. *Boundary-Layer Meteorol.*, 153, 89-116, 2014.
- [3] **Sherwood, S. C., Bony, S., Dufresne, J. L.**
Spread in model climate sensitivity traced to atmospheric convective mixing. *Nature* 505, 37-42, 2014.
- [4] **Bodenschatz, E., Malinowski, S., Shaw, R., Stratmann, F.**
Can we understand clouds without turbulence? *Science* 327, 970-971, 2010.
- [5] **Mellado, J. P.**
The evaporatively driven cloud-top mixing layer. *J. Fluid Mech.*, 660, 5-36, 2010.
- [6] **Mellado, J. P., Stevens, B., Schmidt, H.**
Wind shear and buoyancy reversal at the top of stratocumulus. *J. Atmos. Sci.*, 71, 1040-1057, 2014.
- [7] **Fox, D. G., Lilly, D. K.**
Numerical simulation of turbulence flows. *Rev. Geophys. Space Phys.*, 10, 51-72, 1972.

contact: Juan Pedro Mellado,
juan-pedro.mellado@mpimet.mpg.de

• Juan Pedro Mellado

Max-Planck-Institute for Meteorology, Germany

The strong Interaction at Neutron-rich Extremes

The microscopic understanding of atomic nuclei and of high-density matter is a very challenging task. Powerful many-body simulations are required to connect the observations made in the laboratory to the underlying strong interactions between neutrons and protons, which govern the properties of nuclei and of strongly interacting matter in the universe. Renewed interest in the physics of nuclei is driven by discoveries at rare isotope beam facilities worldwide, which open the way to new regions of exotic, neutron-rich nuclei, and by astrophysical observations and simulations of neutron stars and supernovae, which require controlled constraints on the equation of state of high-density matter. Fig. 1 shows the substantial region of exotic nuclei that will be explored at the future FAIR facility in Darmstadt. The nuclear many-body problem involves two major challenges. The first concerns the derivation of the strong interaction between nucleons (neutrons and protons), which is the starting point of few- and many-body ab initio calculations. Since nucleons are not elementary particles, but composed of quarks and gluons, the strong interaction has a very complex structure. Although it is becoming possible to study systems of few nucleons directly based on quarks and gluons, the fundamental degrees of freedom of Quantum Chromodynamics (QCD), high-precision calculations of nuclei based on quarks and gluons will not be feasible in the foreseeable future. As a systematic approach, chiral effective field theory (EFT) allows to derive nuclear forces in terms

of low-energy degrees of freedom, nucleons and pions, based on the symmetries of QCD [1]. The chiral EFT framework provides a systematically improvable Hamiltonian and explains the hierarchy of two-, three-, and higher-body forces. The presence of such many-body forces is an immediate consequence of strong interactions. In particular, the computation and inclusion of three-nucleon (3N) forces in many-body calculations is one of the current frontiers [2].

The second challenge concerns the practical solution of the many-body problem based on nuclear forces. Since the computational complexity grows significantly with the number of particles, up to about 10 years ago the scope of ab initio calculations was limited to light nuclei up to around carbon (with nucleon number $A=12$). Due to advances on several fronts and also due to rapidly increasing




Figure 1: Chart of atomic nuclei. The black points mark the stable nuclei, while the grey region denotes all nuclei known to date. Of the latter, the points in purple have been studied with mass measurements at the GSI Helmholtzzentrum für Schwerionenforschung Darmstadt. The blue points mark the exotic nuclei that will be explored at the future FAIR facility. Figure courtesy Y. Litvinov (GSI).




Figure 2: Contour plots of RG-evolved 3N interactions as a function of three-body hypermomenta (see Ref. [5] for details). Shown is the systematic decoupling of low- and high-momentum parts in the interaction as the RG resolution scale λ is lowered from left to right. (Figure taken from Ref. [5]).

computing power, this limitation has nowadays been pushed to much heavier systems (see, e.g., the recent work of Ref. [3]). One key step was the development of powerful renormalization group (RG) methods that allow to systematically change the resolution scale of nuclear forces [4]. Such RG transformations lead to less correlated wave functions at low resolution and the many-body problem becomes more perturbative and tractable.

Our work focuses on the derivation of RG evolved interactions and electroweak operators, the inclusion of chiral 3N forces in many-body calculations of extreme neutron-rich nuclei, and on the development of Quantum Monte Carlo simulations with chiral EFT interactions, which open up nonperturbative benchmarks for high-density matter. These calculations enable us to explore the formation of structure in exotic nuclei, the properties of neutron-rich nuclei and matter that play a key role in the synthesis of heavy elements in the universe, as well as the nuclear physics involved in applications to fundamental symmetries, e.g., for the nuclear matrix elements of neutrinoless double-beta decay that probes the nature and mass scale of the neutrinos.

RG Evolution of nuclear Interactions

The convergence behavior and the required computational resources of many-body calculations for a given nucleus are governed by the properties of the employed nuclear forces. It is convenient to visualize nuclear interactions as a function in momentum space, where low momenta correspond to large interparticle distances and high momenta to short-range correlations. In general, a strong coupling of low- and high-momentum parts in nuclear interactions induce strong virtual excitations of particles and imply a poor perturbative convergence and large required basis spaces for the solution of the many-body Schrödinger equation. The similarity renormalization group (SRG) allows to systematically decouple low-momentum physics from high-momentum details via a continuous sequence of unitary transformations that suppress off-diagonal matrix elements, driving the Hamiltonian towards a band-diagonal form [4]. This decoupling is illustrated in Fig. 2 on the basis of a representative chiral 3N interaction.

Computationally, the SRG evolution of NN interactions is straightforward and can be performed on a local computer. However, when evolving nuclear interactions to lower resolution, it is inevitable that higher-body interactions

are induced even if initially absent. This might be considered unnatural, if nuclei could be accurately calculated based on only NN interactions. However, chiral EFT reveals the natural scale and hierarchy of many-body forces, which dictates their inclusion in calculations of nuclei and nuclear matter. In fact, the importance of 3N interactions has been demonstrated in many different calculations [2]. The RG evolution of 3N forces is computationally challenging since typical dimensions of interaction matrices in a momentum-space partial-wave representation can reach about $10^4 - 10^5$. This means that the required memory for storing a single interaction matrix in double precision can reach about 40 GB. For the solution of the RG flow equations it is necessary to evaluate efficiently matrix products of this dimension. Since numerical solvers of differential equations need several copies of the solution vector for a stable and efficient evolution, a distributed storage of all matrices and vectors is mandatory. For an efficient evaluation of large matrix products we have employed a hybrid OpenMP/MPI strategy for our implementation.

3N Forces and Neutron-rich Nuclei

Nuclei with a certain number of protons and neutrons are observed to be particularly well-bound. These closed-shell or "magic" nuclei form the basis of the nuclear shell model, which is a key computational method in nuclear physics. Exploring the formation of shell structure and how these magic configurations evolve with nucleon number towards the limits of the nuclear chart is a frontier in the physics of nuclei, and the microscopic understanding from nuclear forces represents a major challenge. The theoretical shortcomings in predicting shell structure are




Figure 3: Two-neutron separation energy S_{2n} of the neutron-rich calcium isotopes as a function of neutron number N (see Ref. [7] for details). Shown are the new experimental results from the ISOLTRAP collaboration (red squares) in comparison with our predictions based on chiral NN+3N interactions (blue line), as well as shell-model (KB3G and GXPF1A) and coupled cluster (CC) calculations. (Figure taken from Ref. [7]).

particularly evident in the calcium isotopes. While microscopic calculations with well-established NN forces reproduce the standard magic numbers $N = 2, 8, 20$, they do not predict ^{48}Ca as a doubly-magic nucleus when neutrons are added to ^{40}Ca . As a result, phenomenological forces have been adjusted to render ^{48}Ca doubly magic, and it has been argued that the need for these phenomenological adjustments may be largely due to neglected 3N forces. In recent work, we have shown that 3N forces play a decisive role in medium-mass nuclei and are crucial for the magic number $N = 28$ [6]. For the calcium isotopes, the predicted behavior of the two-neutron separation energy S_{2n} up to ^{54}Ca is in remarkable agreement with precision mass measurements of the ISOLTRAP collaboration at ISOLDE/CERN using a new multi-reflection time-of-flight mass spectrometer, as shown in Fig. 3. The new $^{53,54}\text{Ca}$ masses are in excellent agreement with our NN+3N

predictions and unambiguously establish $N = 32$ as a shell closure. This work was published with the ISOLTRAP collaboration in Nature [7]. Since it is not possible to solve the many-body problem exactly for general medium-mass nuclei, valence-space methods utilize a factorization of nuclei into a core and valence nucleons that occupy a truncated single-particle space above the core. The interactions of particles in this valence-space are computed microscopically in many-body perturbation theory (MBPT), whereas the primary computation lies in the self-consistent evaluation of a large number of one- and two-body diagrams. The resulting effective Hamiltonian can then be diagonalized exactly, and within certain limits reproduces the exact eigenvalues.

Quantum Monte Carlo Simulations with chiral EFT Interactions

Quantum Monte Carlo (QMC) methods have been proven to be a very powerful tool for studying light nuclei and neutron matter [8]. In Refs. [9,10,11], we have presented first QMC calculations

based on chiral NN interactions. This was not possible before due to nonlocalities in chiral EFT interactions. However, it is possible to remove all sources of nonlocality in nuclear forces up to next-to-next-to-leading order (N^2LO) in the chiral expansion. This enables us to perform auxiliary-field diffusion Monte Carlo (AFDMC) calculations for the neutron matter equation of state up to nuclear saturation density based on local leading-order (LO), next-to-leading order (NLO), and N^2LO NN interactions. Our results exhibit a systematic order-by-order convergence in chiral EFT and provide nonperturbative benchmarks with theoretical uncertainties. For the softer interactions, perturbative calculations are in excellent agreement with the AFDMC results, as shown in Fig. 4. These advances also opened up first Green's Function Monte Carlo calculations of light nuclei based on chiral NN interactions [11]. Presently, we are working on the implementation of the leading 3N forces in QMC simulations. This paves the way for QMC calculations with systematic chiral EFT interactions for nuclei and nuclear matter, for testing the perturbativeness of




Figure 4: Neutron matter energy per particle E/N as a function of density n based on AFDMC and MBPT calculations (see Ref. [10] for details). The different panels show results using local N^2LO NN interactions with different resolution scale $R_0 = 1.0 - 1.2$ fm (the latter being the softest interaction). For the MBPT results, we show the Hartree-Fock energies as well as the energy at second and third order. The width of the bands provide a measure of the theoretical uncertainty. For the softer interactions (right panel), the perturbative calculations are in excellent agreement with the AFDMC results. (Figure taken from Ref. [10]).

different orders, and also allows for matching to lattice QCD results in a finite volume.

The QMC methods we use in our calculations treat the Schrödinger equation as a diffusion equation in imaginary time and project out the ground-state wave function from a trial wave function by evolving to large imaginary times. GFMC performs, in addition to a stochastic integration over the particle coordinates, explicit summations in spin-isospin space, and is thus very accurate but computationally very costly, so that one can only access particle numbers with $A \leq 12$. In contrast, AFDMC also stochastically evaluates summations in spin-isospin space and shows a better scaling behavior at the cost of less accuracy. We can thus simulate 66 fermions in our neutron matter calculations. For our QMC simulations we typically average over 5-10k walkers for several thousand time steps. Since we use independent walkers, the code is easily parallelizable and shows an excellent and almost linear scaling behavior with the number of cores. Typically, we use 200-400 cores. In contrast to SRG transformations, we have only moderate memory requirements of typically 1GB per core.

Acknowledgement

Our results could not have been achieved without an allocation of computing resources at the Jülich Supercomputing Centre. We are grateful to the John von Neumann Institute for Computing for selecting the present project as 2014 Excellence Project, and thank our collaborators on the JUROPA project, in particular Alexandros Gezerlis, Jason Holt, Javier Menéndez, and Johannes Simonis. This work was supported in part by the ERC Grant No. 307986 STRONGINT.

References

- [1] **Epelbaum, E., Hammer, H. W., Meißner, U.-G.**
Modern theory of nuclear forces, Rev. Mod. Phys. 81, 1773, 2009.
- [2] **Hammer, H. W., Nogga, A., Schwenk, A.**
Three-body forces: from cold atoms to nuclei, Rev. Mod. Phys. 85, 197, 2013.
- [3] **Binder, S., Langhammer, J., Calci, A., Roth, R.**
Ab initio path to heavy nuclei, Phys. Lett. B 736, 119, 2014.
- [4] **Furnstahl, R. J., Hebeler, K.**
New applications of renormalization group methods in nuclear physics, Rept. Prog. Phys. 76, 126301, 2013.
- [5] **Hebeler, K.**
Momentum space evolution of chiral 3N interactions, Phys. Rev. C 85, 021002, 2012.
- [6] **Holt, J. D., Otsuka T., Schwenk, A., Suzuki, T.**
Three-body forces and shell structure in calcium isotopes, J. Phys. G 39, 085111, 2012.
- [7] **Wienholtz, F., et al.**
Masses of exotic calcium isotopes pin down nuclear forces, Nature 498, 346, 2013.
- [8] **Carlson, J., et al.**
Quantum Monte Carlo methods for nuclear physics, arXiv:1412.3081.
- [9] **Gezerlis, A., et al.**
Quantum Monte Carlo calculations with chiral effective field theory interactions, Phys. Rev. Lett. 111, 032501, 2013.
- [10] **Gezerlis, A., et al.**
Local chiral effective field theory interactions and quantum Monte Carlo applications, Phys. Rev. C 90, 054323, 2014.
- [11] **Lynn, J. E., et al.**
Quantum Monte Carlo calculations of light nuclei using chiral potentials, Phys. Rev. Lett. 113, 192501, 2014.

contact: Achim Schwenk,
schwenk@physik.tu-darmstadt.de

- Kai Hebeler
- Achim Schwenk
- Ingo Tews

Institut für
Kernphysik,
Technische
Universität
Darmstadt,
ExtreMe Matter
Institute EMMI,
GSI, Germany

Lattice QCD as Tool for Discoveries

With the discovery of the Higgs particle the Standard Model of particle physics has once again passed a crucial experimental test. Already before this discovery the gauge sector of the theory, which describes the electromagnetic, weak and strong interactions was tested far beyond any reasonable doubt. Now that also the Standard Model mechanism for electroweak symmetry breaking (through the Higgs mechanism) has been proven to be correctly described, the Standard Model is established even further as one of the great cultural achievements of the 20th century. Consequently, the task of the 21st century is to clarify the physics Beyond the Standard Model (BSM) which is already known to exist (e.g., cosmologically relevant entities like dark matter, dark energy, inflatons etc. are not part of the Standard Model) but for which the fundamental theory responsible has not been determined. For this new endeavor the task is no longer to test the Standard Model but to achieve the highest possible accuracy in Standard Model calculations and thus to maximize the potential to uncover BSM physics in, e.g., the LHC experiments at CERN. In fact, success depends equally crucially on experimental progress as on a continuous improvement of theory. As the physics of quarks and gluons (QCD: Quantum Chromodynamics) is by far the most difficult part of the Standard Model, progress in QCD is especially important. Fig. 1 shows an attempt to illustrate the complexity of this task. What is shown is a schematic illustration of a proton. The properties of a proton are

dominated by quantum effects. Therefore, it is a much too strong simplification to say that proton is built up from three quarks. Instead, its properties depend very strongly on effects like the production of virtual quark-antiquark pairs and gluons. The resulting complex quantum state has definitive mass, charge and spin but how these quantities split into individual contributions of quarks and gluons can only be answered by taking the statistical average over all possible quantum states. At the LHC protons are collided with extremely high energy such that due to relativistic time dilatation the result of an individual collision depends on what one might want to call a snapshot of this complex state, making the task to fully understand the "underlying event" a formidable one.

Many highly sophisticated QCD techniques have been developed over the last decades for this purpose and many aspects of perturbative and non-perturbative nature can be treated satisfactorily by analytic means. However, there exist also many important non-perturbative quantities which can only be calculated (at least up to now) by computationally expensive numerical simulations using Lattice QCD (LQCD). This is the field we are working in.

Linking QCD to purely statistical Calculations

The fundamental observation, which is the basis of LQCD is that the analytic continuation of time to imaginary time allows one to map quantum field theory onto thermodynamics and statistics.




Figure 1: Attempt to illustrate the complex quantum state of a proton. Figure taken from [1].

This allows for a purely statistical treatment of otherwise untreatably complicated QCD problems. Discretization of space-time, i.e. introducing a "lattice" of space-time points, reduces the number of degrees of freedom of the corresponding statistical problems to a very large but finite number. Finally, Monte Carlo techniques allow one to generate ensembles of representative gauge field configurations. Calculating physical observables in QCD is then reduced to calculating suitable expectation values on these ensembles. For large ensembles with realistic quarks masses, which allow to reach the needed precision, the last two steps need super-computer resources like SuperMUC. To describe any given experiment one always needs a combination of analytic and numerical techniques. However, the required analytic tools have reached such a high precision, that the information needed from LQCD became the main source of uncertainty. The

LQCD community has promised this vital information already many years ago but fulfilling this promise has again and again turned out to be much more difficult than anticipated. We have now added a new chapter to this epic story.

Hadronic Structure

As argued above, a hadron like the proton, i.e. a quark-gluon bound state, has a truly mind-boggling complexity, see Fig. 1. QCD combines relativistic quantum field theory and strongly-coupled nonlinear dynamics. The task of deducing the incredibly complex internal structure of a hadron from the debris of a high energy particle collision "long" (relatively speaking) after the collision is over sounds like a completely hopeless enterprise. Still, the standards now reached by modern hadron physics means that this is done successfully every day. Over the years an ever longer list of quantities which




Figure 2: Our new results (dark green circles) in comparison to results from other collaborations and experimental data (black symbols). Please note the reduced parameter range plotted, resulting in optically large error bars. Figure taken from [2].




Figure 3: Our new results for g_A (red squares) and the ratio g_A divided by the pion decay constant labeled as F_π . Please note the reduced parameter range plotted, resulting in optically large error bars. Figure taken from [3].

provide information about the internal structure of hadrons in terms of their fundamental constituents (quarks and gluons) have been introduced and studied in great detail by the international physics community. For these quantities the crucially needed lattice input can be organised in the form of matrix elements of the type

$$\langle \text{hadron 1} | \text{specific quark gluon operator} | \text{hadron 2} \rangle$$

where hadron 1 and hadron 2 can be identical. The matrix element describes the transition from hadron 1 into hadron 2 via an interaction represented by the quark-gluon operator. The task of LQCD is to determine a large number of such matrix elements. Some of these are known experimentally to high precision. These provide valuable test cases which can be used to demonstrate how well the systematics involved in a lattice calculation are under control. This is used to guide estimates of the precision that can be obtained for those matrix elements which are not yet known experimentally. Unfortunately, so far LQCD has failed to reproduce the experimental results for two such test cases: the second moment of the isovector parton distribution function of the nucleon, denoted $\langle x(u-d) \rangle$, which gives information about the fraction of the nucleon momentum carried by the individual quarks and the isovector axial-vector coupling of the nucleon denoted g_A , which is associated with the beta decay of a neutron into a proton. These quantities are typical examples of hadron structure observables such that failing to calculate them reliably is a most serious reason for concern.

Our new results for $\langle x(u-d) \rangle$ and g_A

Lattice results are assessed in terms of the associated (Monte Carlo) statistical and systematic errors. The goal of LQCD is to simulate QCD using quarks with physical masses on a lattice with a large enough physical volume and small enough lattice spacing that finite volume and discretisation effects are not significant. However, the numerical cost of the simulations grows very rapidly as the quark mass is decreased towards the physical value while keeping the physical volume large enough. This leaves any lattice collaboration with basically two options. One can simulate with unphysically large quark masses and then extrapolate to the physical ones. The advantage of this option is that for a given amount of computing time one gets much smaller statistical errors. The disadvantage is that there are systematic uncertainties introduced by the mass extrapolation. Alternatively, one can simulate at or close to the physical masses but then has to live with (other) large systematic uncertainties unless one has access to very large computer resources. Most collaborations do not have the possibility to choose the second approach. However, thanks to QPACE, our home-made supercomputer, and SuperMUC, we had the computing resources required. Our results for the momentum fraction $\langle x(u-d) \rangle$ and the coupling g_A are shown in Fig.1 and 2. The results are plotted against the square of the pion mass which is proportional to the u/d quark mass.

The message of both figures is clear: as the quark mass is decreased the experimental results can be reproduced with an error of around 10%. However, to achieve greater precision more work has to be done.

Our new data at low pion mass achieves an overall uncertainty of less than 5% but significantly disagrees with the phenomenological values. For $\langle x(u-d) \rangle$ precision results are likely to require simulations with smaller lattice spacings to have better control of the continuum limit. Unfortunately with the standard lattice formulation we used so far, simulations on finer lattices are unacceptably expensive such that one needs a new approach. The only promising formulation proposed so far is that used by the CLS collaboration. We have joined this collaboration and already started large scale simulations. For g_A the remaining discrepancy is caused by finite volume effects which can be corrected for with more computer time by simulating on larger space-time lattices.

While we obviously would have preferred to determine many different physics quantities already now with high precision, this story provides a nice illustration of the fundamentals of science research: One should never fall prey to wishful thinking but should always work on improving the precision achieved and be open to identify problems which were previously not apparent.

References

- [1] http://uni-hamburg.cfel.de/images/content/e8/e76/index_eng.html
- [2] **Bali, G. S., et al.**
Physical Review D 91, 5, 054501, ArXiv: 1412.7336, 2015.
- [3] **Bali, G. S., et al.**
Physical Review D 90, 7, 074510, ArXiv: 1408.6850, 2014.

contact: Andreas Schäfer,
andreas.schaefer@physik.uni-regensburg.de

• Andreas Schäfer

Institute for
Theoretical
Physics,
Regensburg
University,
Germany

LIKWID Performance Tools

LIKWID is a set of performance-related open source command line tools targeting X86 processors. It was initially presented in late 2009. In inside spring edition 1, 2010 we presented it to a wider audience. Since that time a lot has happened and we think it is time to report on the latest LIKWID developments. LIKWID is developed at the Erlangen Regional Computing Centre (RRZE). Since July, 2013 LIKWID is officially funded as part of the BMBF FEPA project. FEPA is a joint effort together with LRZ-BADW and NEC Deutschland to enable system-wide application profiling on large-scale HPC cluster systems. LIKWID was downloaded several thousand times, and it has found its place in the tooling community as a set of simple to use command line tools with a unique feature set. Many HPC centers offer LIKWID as part of their standard tool set. With regard to functionality LIKWID enables to get information about compute nodes, measure various profiling data sources on processors (e.g., HPM data, RAPL energy counters), control thread/core affinity and perform microbenchmarking. The current stable release is version 3.1.3. We are currently in the process to finish the next major release LIKWID 4, which includes major changes in the internal software architecture of LIKWID. The command line applications in LIKWID 4 are implemented in the Lua scripting language and are based on a common C library API. This API is designed to also be used by other tools building on LIKWID functionality.

In its initial version, LIKWID consisted of three command line tools:

- likwid-topology – shows all topology information on a compute node that is relevant for software developers.
- likwid-pin – controls thread to core affinity.
- likwid-perfctr – enables measurements of hardware performance monitoring data on X86 processors.

Since that time several other tools were added:

- likwid-bench – a microbenchmarking framework and application allowing rapid prototyping of threaded assembly kernels.
- likwid-mpirun – enables simple and flexible pinning of MPI and MPI/threaded hybrid applications with integrated likwid-perfctr support.
- likwid-powermeter – tool for accessing RAPL counters (power and energy consumption) and query Turbo mode steps on Intel processors.
- likwid-memsweeper – tool to eliminate file system buffer cache from ccNUMA domains.
- likwid-setFrequencies – tool to set specific processor clock frequencies.

What sets LIKWID apart from many other tools in this area is that it implements all functionality on its own in user space instead of relying on the Linux perf_event interface like most other tools in the Hardware Performance Monitoring (HPM) sector. This allows us to support new processors quickly without the need to install a new Linux kernel. As a consequence, LIKWID was one of the first tools offering a usable access to the Uncore HPM units on newer Intel processors. For a long time LIKWID used the "cpuid" instruction to query node topology information. Starting with LIKWID 4 we decided to rely on "hwloc" as an alternative backend for topology information. This enables a more robust affinity interface and also allows us

to port LIKWID to non-X86 architectures more easily in the future. In the following we want to focus on a few features that let LIKWID stand out against other tools.

Thread Group Concept

The of processor IDs as provided by the Linux OS are tedious to use in practice, since there are no fixed rules for how they should be mapped to the node topology. To alleviate this problem LIKWID provides "logical" processor IDs and the concept of "thread groups". A thread group is a topological entity shared by several processors. The thread group syntax is used throughout LIKWID to specify processor lists. In LIKWID a thread group is indicated by single capital letters: N for node, S for socket, M for memory domain, C for last level cache. There are two syntax variants for specifying thread groups, a list-based one and an expression-based one. Here is an example using the list-based syntax to pin an OpenMP application:

```
likwid-pin -c S0:0-3@S1:0-3 ./a.out
```

Above command will set OMP_NUM_THREADS to 8 in case it is not yet set from outside and place the first four threads on the first four physical cores on socket 0 and the second four threads on the first four physical cores on socket 1. Using the list-based syntax LIKWID deals with SMT threads using a "physical cores first" policy, which makes it easy to ignore the SMT feature. This means that on a socket with four cores having two SMT threads per core the IDs 0-3 are hardware threads on the first four distinct physical cores and IDs 4-7 are their counterparts. As can be seen in the example, multiple thread group expressions can be chained using the "@" character.

As an alternative LIKWID also supports an expression-based syntax. The following could be used, e.g., on the Intel Xeon Phi:

```
likwid-pin -c E:N:30:2:4 ./a.out
```

This command will set OMP_NUM_THREADS to 30 using two threads per core (out of four possible SMT threads). Such a placement would be difficult to realize using the list-based syntax. It is also more convenient to use in benchmarking scripts. LIKWID thread groups offer a consistent and flexible interface for expressing node topology domains using a future-proof unified command line syntax.

Performance Groups

One reason why hardware performance monitoring (HPM) is difficult to use is that the things a software developer is interested in can usually not be directly measured. Instead, a collection of HPM events must be configured to compute the resulting derived metric. There is no accepted standard for naming HPM events, and due to poor documentation it is difficult for the application developer to pick the right event sets for a specific purpose. Moreover at least on X86 processors HPM events are not an integral part of the product for processor vendors. Consequently, the documentation may be outdated or wrong, and the events are sometimes unreliable. LIKWID provides "performance groups", which bundle event sets and derived metrics computed from them. This relieves the developer from finding meaningful event sets on every new processor generation. We are currently working on validating the supported performance groups with respect to accuracy. Another feature of LIKWID is its easy extension. INCE the performance groups are defined in plain text files it is simple to alter existing groups or add new ones. In the current stable release LIKWID must be recompiled to be aware of changed or additional performance groups. In the upcoming LIKWID 4 release performance groups are dynamically interpreted without the need to recompile.

As an example, the following command measures the memory bandwidth on an IvyBridge-EP system:

```
likwid-perfctr -C S0:0-1@S1:0-1 -g
MEM -m ./a.out
```

This will run the thread-parallel application with four threads, two on socket 0 and two on socket 1. The `-m` option indicates that the application was instrumented using the "marker API". The output could look as follows (shortened):

```
-----
CPU type: Intel Core IvyBridge EP processor
CPU clock: 2.20 GHz
Measuring group MEM
-----

Region: copy
-----
+-----+-----+-----+-----+-----+
| Region Info | core 0 | core 1 | core 10 | core 11 |
+-----+-----+-----+-----+-----+
| RDTSC Runtime [s] | 0.623158 | 0.624219 | 0.623072 | 0.624195 |
| call count | 10 | 10 | 10 | 10 |
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
| Event | core 0 | core 1 | core 10 | core 11 |
+-----+-----+-----+-----+-----+
| INSTR_RETIRED_ANY | 1.02895e+09 | 1.02998e+09 | 2.17321e+08 | 2.16647e+08 |
| CPU_CLK_UNHALTED_CORE | 1.7788e+09 | 1.78486e+09 | 1.80275e+09 | 1.80704e+09 |
| CPU_CLK_UNHALTED_REF | 1.34993e+09 | 1.35452e+09 | 1.36762e+09 | 1.37088e+09 |
| CAS_COUNT_RD | 2.93971e+07 | 0 | 15068 | 0 |
| CAS_COUNT_WR | 2.04592e+07 | 0 | 3583 | 0 |
| CAS_COUNT_RD | 2.95764e+07 | 0 | 13294 | 0 |
| CAS_COUNT_WR | 2.0555e+07 | 0 | 1603 | 0 |
| CAS_COUNT_RD | 2.94727e+07 | 0 | 346550 | 0 |
| CAS_COUNT_WR | 2.04713e+07 | 0 | 334941 | 0 |
| CAS_COUNT_RD | 2.94429e+07 | 0 | 14053 | 0 |
| CAS_COUNT_WR | 2.04603e+07 | 0 | 2346 | 0 |
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
| Metric | core 0 | core 1 | core 10 | core 11 |
+-----+-----+-----+-----+-----+
| Runtime (RDTSC) [s] | 0.623158 | 0.623158 | 0.623158 | 0.623158 |
| Runtime unhaltd [s] | 0.80853 | 0.811283 | 0.819414 | 0.821364 |
| Clock [MHz] | 2899 | 2899 | 2900.02 | 2900.01 |
| CPI | 1.72876 | 1.73291 | 8.29531 | 8.34091 |
| Memory Read BW [MBytes/s] | 12107.5 | 0 | 39.9477 | 0 |
| Memory Write BW [MBytes/s] | 8416.06 | 0 | 35.1729 | 0 |
| Memory BW [MBytes/s] | 20523.6 | 0 | 75.1206 | 0 |
| Memory data volume [GBytes] | 12.7894 | 0 | 0.046812 | 0 |
+-----+-----+-----+-----+-----+
```

The first output box shows the total run-time spent in a region ("copy" here) and how many times it was executed (10 times). The next table shows the raw events as measured per core on the hardware. The last box displays derived metrics computed from the raw event counts. The reported clock speed of 2.9GHz indicates that Turbo mode is enabled on this system (the nominal clock is 2.2 GHz as shown in the header). It can also be seen that almost all data volume is served by just one of the two sockets. The application may thus have a ccNUMA -related problem. There is a final statistics table for threaded measurements providing MIN, MAX, MEAN and SUM for the derived metrics, which was omitted in this output.

LIKWID Libraries

Likwid-perfctr performs measurements on the specified cores, but it does not know about the code that was executed during the measurement. Measurements are connected to code by affinity enforcement. All the pinning facilities of likwid-pin are therefore also available in likwid-perfctr. To get a meaningful profiling result it is usually necessary to instrument regions in the code for measurement. LIKWID provides a simple marker API, which allows to tag and name regions in the code. This API only consists of six calls and supports a Fortran 90 and a C interface.

The following code snippet illustrates the marker API:

```
#include <likwid.h>
LIKWID_MARKER_INIT; // must be called from serial region
LIKWID_MARKER_THREADINIT; // must be called from parallel region

LIKWID_MARKER_START("Compute");
<code>
LIKWID_MARKER_STOP("Compute");

LIKWID_MARKER_CLOSE; // must be called from serial region
```

The upcoming LIKWID 4 release will provide is an extensive library API which makes it easier to implement tools on top of LIKWID. The C library consists of 72 functions.

Outlook

LIKWID 4 will provide a new software architecture and enhanced functionality. In the context of the BMBF FEPA project LIKWID is established as part of a unified application profiling framework on the system level. Measurements have indicated that the current interface to the LINUX OS generates significant overhead [4]. Hence, we will provide a dedicated Linux kernel module enabling low overhead measurements. Porting of LIKWID to new microarchitectures is a constant effort. Depending on demand, LIKWID is also ported to non-HPC processors like, e.g., various flavors of Intel Atom. Currently LIKWID is hosted on Google Code. LIKWID depends on its users to report errors and guide its development by feedback and feature requests. If you want to try LIKWID we would be happy to receive your feedback on the LIKWID user mailing list.

References

- [1] **Treibig, J., Hager, G., and Wellein, G.**
LIKWID: A lightweight performance-oriented tool suite for x86 multicore environments. Proceedings of PSTI2010, the First International Workshop on Parallel Software Tools and Tool Infrastructures, San Diego CA, September 13, 2010. DOI: 10.1109/ICPPW.2010.38
- [2] **Treibig, J., Hager, G., Meier, M., and Wellein, G.**
LIKWID performance tools. InSiDE 8(1), 50-53, Spring 2010.
- [3] **Treibig, J., Hager, G., and Wellein, G.**
LIKWID performance tools. In: C. Bischof et al. (eds.), Competence in High Performance Computing 2010. Springer, ISBN 978-3-642-24025-6 (2012), 165-175. DOI: 10.1007/978-3-642-24025-6_14.
- [4] **Röhl, T., Treibig, J., Hager, G., and Wellein, G.**
Overhead analysis of performance counter measurements. Proceedings of PSTI2014, the Fifth International Workshop on Parallel Software Tools and Tool Infrastructures, Minneapolis MN, September 12, 2014.

How to get it

The LIKWID performance tools are Open Source. You can download them at:

<http://tiny.cc/LIKWID>

LIKWID is is funded as part of the FEPA project.



This work was funded by KONWIHR.



contact: Jan Eitzinger,
jan.eitzinger@fau.de

- Jan Eitzinger
- Thomas Röhl
- Georg Hager
- Gerhard Wellein

Friedrich-Alexander-Universität
Erlangen-Nürnberg
Regionales RechenZentrum
Erlangen (RRZE),
Germany

ECO₂Clouds - Experimental Awareness of CO₂ in Federated Cloud Sourcing



Project Outline

Nowadays, data centers consume a high amount of energy and in addition, they are responsible for the generation of a significant portion of CO₂ emissions. Furthermore, when thinking of cloud computing the increasing adoption of cloud based developments has a big impact on the environment. The energy consumption and the resulting emissions are growing dramatically. The ECO₂Clouds project target was to analyse and provide solutions for ecological implications of cloud-based IT infrastructures (Fig. 1) and bridging the critical gap between latest state-of-the-art in research and business. It enfolded three geographically distributed testbeds across Europe offering heterogeneous cloud resources including the possibility to measure precisely the energy consumption of each physical host. ECO₂Clouds entire infrastructure providers are also part of the BonFIRE project [1], the resources can be accessed seamlessly by describing an application energy profile which is parsed by the ECO₂Clouds scheduler in order to submit the extended experiment descriptor to the BonFIRE control mechanisms which are mainly based on the Open Cloud Computing Interface (OCCI) [6]. Within the ECO₂Clouds

project strategies for enhancing cloud infrastructures by considering the eco-awareness with the aim to reduce the costs and protecting the environment by additionally tracing the carbon footprint and supporting the reduction of CO₂ emissions from current cloud infrastructures were developed. For achieving the ECO₂Clouds goals the project did not only focus on the infrastructure layer but as well on the virtual and the application layer to achieve the following project vision:

- Improving an effective application deployment on federated cloud infrastructures.
- Reducing the energy consumption and thus the arising costs and the CO₂ emissions.
- Optimizing key assets such as deploying virtual machines, applications and databases.

In the scope of the ECO₂Clouds project, eco-efficient data were collected at the physical cloud infrastructure and virtual machine level. Further, required quality and cost parameters for deploying virtual machines in multi-cloud environments were identified and additionally, evaluation mechanisms and optimizations for algorithms to assess different parameter configurations and their influence in energy-efficient cloud sourcing and application deployment strategies were identified.

The ECO₂Clouds project has achieved its goals including infrastructure sup-

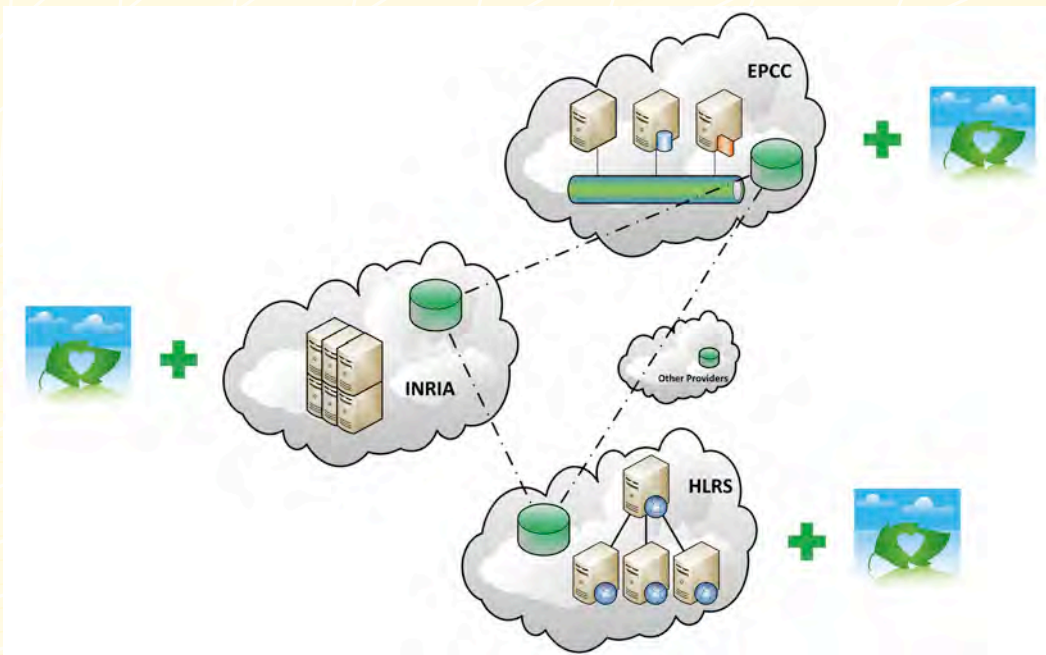


Figure 1: Schematic infrastructure of ECO₂Clouds

port for energy efficiency by taking steps towards the provision of necessary information through monitoring metrics. Thus, ECO₂Clouds monitoring metrics quantify the energy consumption and the environmental impact arising through application execution. The information regarding the energy consumption were used to identify the produced grams of CO₂. Furthermore, an application strategy focusing on the environmental impact for deploying applications on multi or federated clouds was developed. In this scope, key model parameters such as ecological impact, quality and cost dimensions were focused. ECO₂Clouds produced the following results:

- Methods for collecting and exposing the carbon footprint at the infrastructure and the virtual machine level.
- An approach for incorporating the carbon footprint into a federated cloud deployment strategy.

- Mechanisms for an optimized resource utilization of federated clouds.
- An approach for adapting changes to a running application based on the energy consumption.

Through implementing the above mentioned issues, the ECO₂Clouds project has highlighted that monitoring and collecting knowledge about the execution of applications in a federated cloud environment enables a "green optimization".

Work of HLRS

Within the two year duration of ECO₂Clouds, HLRS was responsible for various tasks, starting with providing an efficient and scalable monitoring infrastructure for all the required data and supporting the data mining approach needed for making assumptions regarding eco-efficiency, offering two case studies for the improvement of the ECO₂Clouds scheduler as well as providing a dedicated cloud infrastructure to the project.

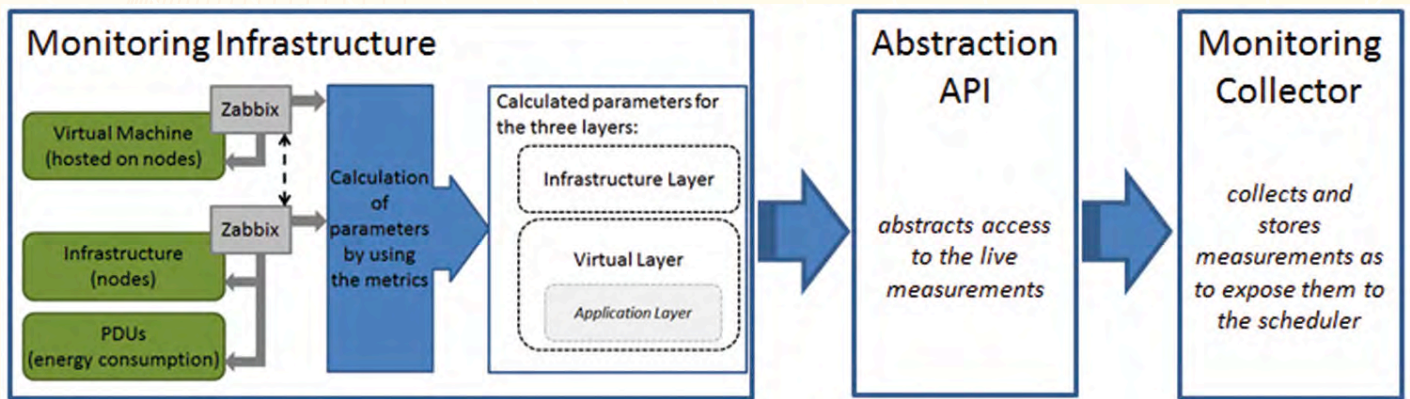


Figure 2: The monitoring infrastructure

The project goals were achieved through implementing the monitoring framework (Fig. 2) taking into account the data mining approach for collecting and elaborating the available data and further, capturing long term data. In the scope of the monitoring framework development, HLRS considered various previously developed monitoring solutions such as GAMES [2] and OPTIMIS [3] projects. These solutions mainly rely on the monitoring tool Nagios [4] whereas BonFIRE already provides a complete Zabbix [5] monitoring framework. Hence, the monitoring solutions were adapted in order to overcome different monitoring approaches as well as potential monitoring overhead.

lector, collects the required monitoring data from the power distribution units (PDUs), physical hosts and virtual machines. The abstraction API facilitates the access to the live measurements so that the monitoring collector can collect and store those measurements in the accounting database. The data collection of the monitoring infrastructure is performed through the monitoring metrics applicable for the PDUs, physical hosts, virtual machines and applications. A central feature of the monitoring system is the possibility to estimate the energy consumption of virtual machines (Fig. 3) as represented by the following simplified formula:

$$\frac{\text{IDLE energy of the host}}{\text{Number of VMs}} + \text{Dynamic energy of the host} \times \frac{\text{CPU used by the VM}}{\text{CPU used by all VMs running on the host}}$$

The monitoring infrastructure, including the three layers infrastructure, virtualization and application layer, and the related monitoring components, the abstraction API and the monitoring col-

Based on the ECO₂Clouds monitoring framework, the ECO₂Clouds data mining service makes use of the data collected through the monitoring metrics. The ECO₂Clouds data mining architecture consists of mainly two major components, the accounting service, to perform transfer of non-reduced data to a remote data storage (DM storage), and the DM storage, which gathers the non-reduced data and performs statistical analysis over them, e.g., correlation analysis over a large




Figure 3: Energy consumption per VM

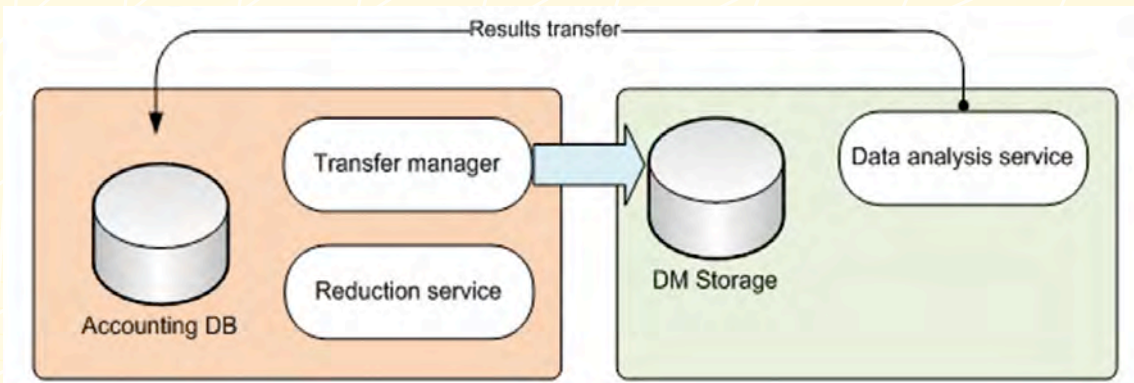


Figure 4: Data mining architecture

enough portion of data. The data reduction service is a crucial component of the accounting service to avoid an uncontrollable growth of the accounting DB, which could cause overloading with the old data while the new data could not be stored and used.

Further implementations were performed for the two case studies and additionally, GAMES project case studies were the foundation of the case studies dealing with Finite Element (FEM) simulations for malicious bones which represent a High Performance Computing (HPC) application. The second case study is dealing with modern e-Business applications: various benchmarks like the Linpack as well as I/O intensive operations were executed to simulate a common e-Business application behaviour. Moreover the mandatory cloud infrastructure was provided to enable the deployment execution and obtain detailed data regarding the energy consumption and thus, facilitating the calculation of the carbon footprint. For that purpose, the already offered infrastructure of BonFIRE was used and extended with PDUs. HLRS was providing the host for the central access points of ECO₂Clouds, an application portal and the scheduler.

The ECO₂Clouds results can be applied and reused to other cloud infrastructures and research projects dealing with federated clouds and eco-efficiency, focusing on reducing the energy consumption and CO₂ emissions.

References

- [1] BonFIRE Project Website: <http://www.bonfire-project.eu/>, Last visited: March 16, 2015.
- [2] GAMES Project Website: <http://www.green-datacenters.eu/>, Last visited: March 16, 2015.
- [3] OPTIMIS Project Website: <http://www.optimis-project.eu/>, Last visited: March 16, 2015.
- [4] Nagios Website: <http://www.nagios.org/>, Last visited: March 16, 2015.
- [5] Zabbix Website: <http://www.zabbix.com/>, Last visited March: 16, 2015.
- [6] OCCI working group Website: <http://occi-wg.org/>, Last visited March: 16, 2015.

- Axel Tenschert
- Michael Gienger

University of
Stuttgart, HLRS,
Germany

contact: Axel Tenschert,
tenschert@hls.de
Michael Gienger,
gienger@hls.de

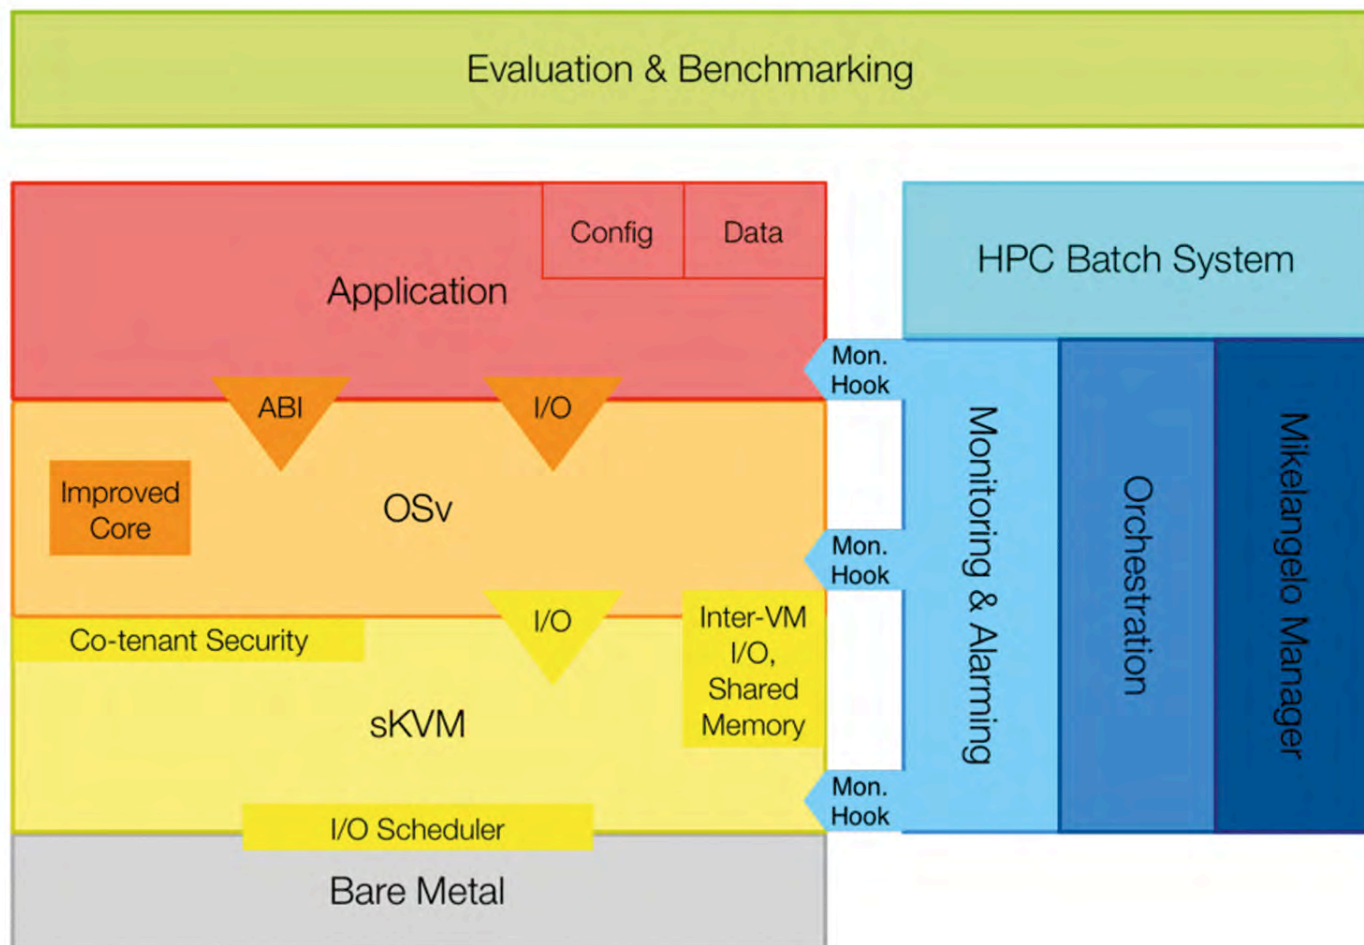
MIKELANGELO - Micro Kernel Virtualization for High Performance Cloud and HPC Systems



MIKELANGELO

MIKELANGELO is a project, targeted to advance the core technologies of Cloud computing, enabling even bigger uptake of virtualization mechanisms, especially in the HPC and Big Data domains. HPC in the Cloud and Big Data technologies under one umbrella. The vision of MIKELANGELO is to improve responsiveness, agility and secu-

rity of the virtual infrastructure through packaged applications, using the lean guest operating system OS^v [1] and the newly developed superfast hypervisor sKVM. In short, the work will concentrate on significant performance improvements of virtual I/O in KVM [2], using enhanced virtual I/O expertise and technologies, integrated and optimized in conjunction with the lightweight operating system OS^v. Thus, it will provide the




private and public cloud communities technologies for fast, agile and secure Cloud application deployments in manifold hardware infrastructures and environments.

With respect to the proposed initial architecture above, an approach and the accompanying software stack that will disrupt the traditional HPC and Private Cloud fields will be developed, tested and validated. The key enablers for achieving this ambitious goal are the use and enhancements of the state-of-the-art hypervisor software KVM, improvements in Virtio [3], a virtualization standard for network and disk device drivers as well as mechanisms for cross-layer optimization in conjunction with the guest operating system OS^v. In order to validate the achieved results, the solution will be integrated in one of the well-accepted Cloud and HPC management systems (such as OpenStack and Torque).

Use Cases

The MIKELANGELO project will implement various use cases, covering different application domains with different requirements for Cloud and High Performance Computing environments in order to demonstrate the functionality and flexibility of the developed and optimized components. Therefore, all use cases are built to show the capabilities and agility of the MIKELANGELO virtualization stack, as well as to measure and compare their performance with traditional approaches.



Big Data


The Big Data use case will leverage the MIKELANGELO software stack in order to provide flexible Big Data services enabling high performance processing. In the context of this use case, flexibility needs to cover on-demand access and elasticity. To achieve this special kind of flexibility, a dedicated Big Data software stack will be directly integrated within a Cloud middleware.

Cloud Bursting

In a Cloud environment, MIKELANGELO's software stack will showcase how to solve the problem of Cloud Bursting efficiently. Cloud bursts are a web-

phenomenon, which leads many users to suddenly use a particular service, outstripping its capacity. MIKELANGELO is used to optimize and enhance unprecedentedly reactive cloud elasticity.

MIKELANGELO software stacks are targeted. Flexible and scalable workload based services to simulate and analyse computational fluid dynamics for new aircraft designs in the Cloud will be enabled.



Simulation using OpenFOAM

Within this use case, high performance simulation of plane wings using the OpenFOAM [4] and



Cancellous Bones Simulation

This use case will implement mechanisms to determine material behaviour of cancellous bones in order to improve bone implant systems. As high performance for both, processing and I/O are required, this use case is executed at the edge of performance capabilities of modern Cloud systems and will strongly benefit of the MIKELANGELO developments.

ROLE of HLRS in the Project

The role of the High Performance Computing Center Stuttgart is to coordinate the Workpackage on use cases and architecture analysis, which is concerned with the definition and evaluation of the use cases and the overall architecture of the project. Furthermore, HLRS will integrate the developed MIKELANGELO components into its Research Cloud and a HPC test system in order to support the execution and validation of the use cases. Finally, HLRS will develop and optimize the Cancellous Bone Simulation software kernel for distributed Cloud execution.

Partners

The MIKELANGELO consortium brings together 9 partners, providing comprehensive research, technical and business expertise in hypervisor I/O research and development (IBM, Huawei), guest operating systems development and kernels (Cloudeus), security in co-tenant environments (BGU), management of Cloud and HPC environments (GWVG, XLAB, HLRS), monitoring and performance evaluation of systems and

applications (GWVG, HLRS, Intel) and finally, a strong business context for exploiting the results (Intel, IBM, XLAB, PIPSTREL).

Key Facts

MIKELANGELO is a large-scale project funded by the European Commission within Horizon 2020 Research and Innovation Programme. It has started at 01.01.2015 and will run until 31.12.2017.

Webpage:

<http://www.mikelangelo-project.eu>

Twitter:

twitter.com/mikelangelo_eu

References

- [1] OS⁹, <http://osv.io/>
- [2] KVM, http://www.linux-kvm.org/page/Main_Page
- [3] Virtio, <http://www.linux-kvm.org/page/Virtio>
- [4] OpenFOAM, <http://www.openfoam.com/>

- Michael Gienger

University of
Stuttgart, HLRS,
Germany

contact: Michael Gienger,
gienger@hlrs.de
Bastian Koller,
koller@hlrs.de

FFMK: A Fast and Fault-tolerant Microkernel-based Operating System for Exascale Computing

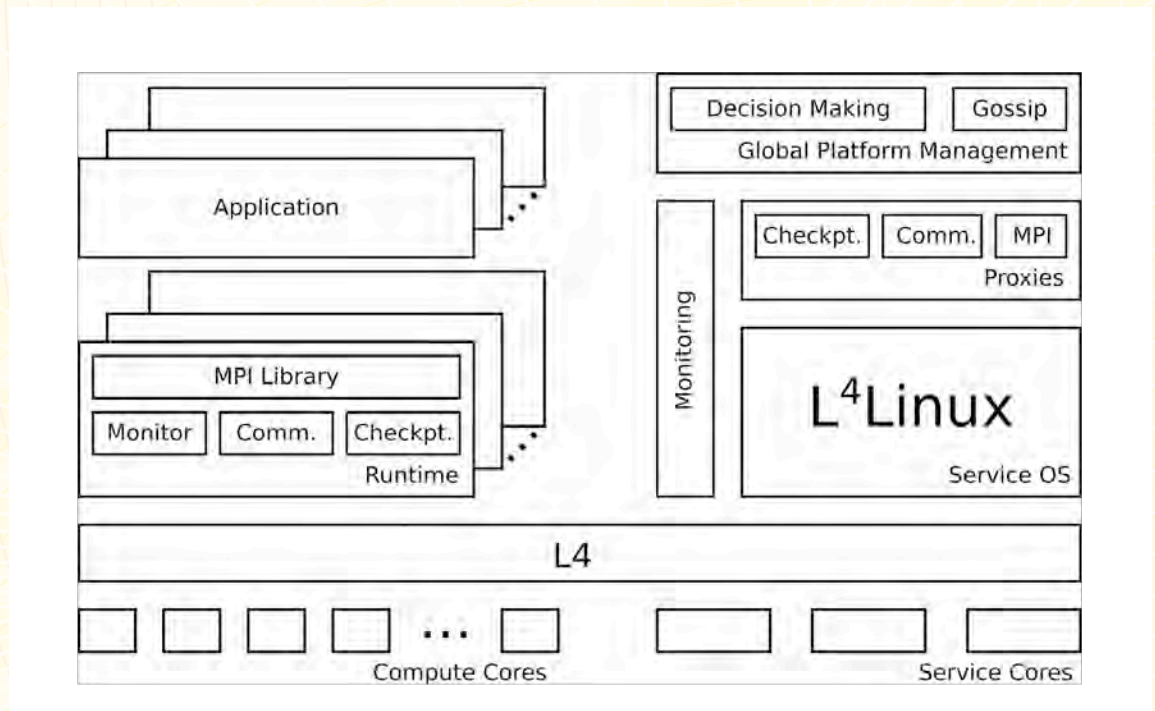


Figure 1: FFMK software architecture: Compute processes with performance-critical parts of (MPI) runtime and communication driver execute directly on L4 microkernel; non-critical functionality is split out into proxy processes on Linux, which also hosts global platform management.

Three out of five research directions of the priority program "Software for Exascale Computing" (SPPEXA) are named system software, computational algorithms, and application software. The project "FFMK: A Fast and Fault-Tolerant Microkernel-based Operating System for Exascale Computing" connects these research directions in an international project.

The FFMK project [1] is driven by Prof. Hermann Härtig (Computer Science, TU Dresden), Prof. Alexander Reinefeld (Zuse Institute Berlin), Prof. Amnon Barak (Computer Science, Hebrew University of Jerusalem), and

Prof. Wolfgang E. Nagel (Center for Information Services and High Performance Computing, TU Dresden). The starting point for the design of the FFMK platform is the expectation that the following major challenges have to be addressed by systems software for exascale computers.

Current high-end HPC systems are tailored towards extremely well-tuned applications, which are assigned fixed partitions of hardware resources by a scheduler such as SLURM. Tuning of these applications includes significant load balancing efforts. We believe a major part of this effort will have to




Figure 2: Run time of COSMO-SPECS+FD4 weather code on BlueGene/Q system JUQUEEN with gossip algorithm running in parallel: Gossip messages sent at different intervals; baseline is same benchmark without running gossip algorithm. Increasing the number of nodes causes more overhead at same intervals due to larger gossip messages describing more nodes. Shorter bars are better; orange part of the bars indicate the time spent communicating via MPI.

shift from application programmers to operating systems and runtimes (OS/Rs) because of the complexity and dynamics of future applications. A number of runtime systems already addresses that challenge, notably Charm++ and X10. We further believe that an exascale OS/R must accommodate more dynamic applications that extend and shrink during their runtime, for example, when computational demand explodes as particle density in a simulation increases.

In addition to imbalances within applications, we expect future exascale hardware to have less consistent performance due to fabrication tolerances and thermal concerns. This is a radical change from the general-purpose CPUs (like x86-64) and accelerators (e.g., GPGPUs) that are used today, which are assumed and selected to perform very regularly. We also assume that not all compute elements can be active at all times due to power and heat issues (dark silicon). Thus, hardware will add to the unbalanced execution of applications in a way that cannot be predicted upfront.

Furthermore, the sheer size of exascale computers with their unprecedented number of components will have significant impact on fault rates. Some OS/Rs already address this concern in part by enabling incremental and application-specific checkpoint/recovery and by using on-node memory to store checkpoint data. However, for exascale machines, we expect more types of memory that differ in aspects like persistence, energy requirements, fault tolerance, and speed. Important examples are on-node non-volatile memory (phase-change memory, flash, etc.) and stacked DRAM. A checkpoint store that can keep up with increased fault rates and the vastly enlarged application-state on exascale machines requires an integrated architecture that uses all these different types of memory.

We believe that a systems-software design for exascale machines that addresses the just described challenges must be based on a coordinated approach across all layers, including applications. The platform architecture as shown in Fig. 1 uses an L4 microkernel [2] as the




Figure 3: Run time of MPI-FFT benchmark on BlueGene/Q system JUQUEEN with gossip algorithm running in parallel: Gossip messages sent at different intervals; baseline is same benchmark without running any gossip algorithm. Increasing the number of nodes causes more overhead at same intervals due to larger gossip messages describing more nodes. Shorter bars are better; orange part of the bars indicate the time spent communicating via MPI.

light-weight kernel (LWK) that runs on each node. All cores run under this minimal common foundation consisting of the microkernel itself and a few extra services that provide higher-level OS functionality. Additionally, an instance of a full commodity OS (Linux in our case) runs on top of it, but only on a few dedicated cores that we refer to as service cores.

Application programming paradigms are supported by small dedicated library operating systems running directly on L4. These library OSes and their supporting components are split into performance-critical parts running directly on L4 and uncritical parts running as mostly idle proxy processes on Linux; we refer to both parts as a runtime.

In the presence of frequent component failures, hardware heterogeneity, and dynamic demands, applications can no longer assume that compute resources are assigned statically. We envision the platform as a whole will be managed by load distribution components. Monitoring and decision making is done at three levels: (1) on each multi-core node

(i. e., local scheduling of threads/processes), (2) per application/partition among nodes, and (3) based on a global view of a master management node. Node-local schedulers take care of (1); scalable gossip algorithms support (2) and (3).

Using gossip, the nodes build up a distributed, inherently fault tolerant, and scalable bulletin board that provides information on the status of the system. Nodes have partial knowledge of the whole system: they know about only a subset of the other nodes, but enough of them in order to make decisions on how to balance load and how to react to failures in decentralized way. The global view over all nodes is available to the master node, which receives gossip messages from some nodes. It makes global decisions such as where to put processes of a newly started application. We demonstrated that this two-layer gossip algorithm matches predictions of an analytical model and simulations for the average age (i.e., quality) of the gossiped information; we further simulated the effect of node failures [3]. We found that — at reasonable

gossip intervals of a few hundred milliseconds — there is no noticeable overhead on application performance when the gossip algorithm runs on the same nodes and interconnect as the application (see Fig. 2 and Fig. 3 from [4]).

To handle hardware faults, a fast checkpointing module takes checkpoint state from applications and distributes and stores it redundantly in memory across several nodes. Our prototype is currently based on XtreamFS [5], which uses erasure coding to provide an efficient, redundant checkpoint store backed by RAM disks. We demonstrated checkpoint/restart of a production code on a Cray XC30, where a node holding checkpoint state was taken offline to simulate failure.

The prototype provides a split MPI runtime and communication driver (Infiniband). The Linux-based components provide management-related functionality and are not on the critical path, whereas performance-critical functionality of MPI and the Infiniband driver runs directly on the LWK. Management functionality of Linux-hosted proxy processes includes forwarding messages to and from the MPI process manager during startup and shutdown. They also interface with the Mellanox Infiniband kernel driver in Linux to manage allocation of message transfer buffers and to take care of the rather complicated connection handling. We currently use MPI as an application runtime also to demonstrate extending and shrinking of existing applications. We therefore create more MPI processes than there are cores avail-

able. The process-level "taskification" of MPI also enables the FFMK OS on each node to migrate MPI processes among cores to balance load locally; cross-node migration is currently not part of our prototype, as the previously described platform management is not yet integrated.

At the time of this writing, after the second year of the project run time, the basic prototype (L4 microkernel, L4Linux, MPI, Infiniband) has been tested on a small partition of an HPC system installed at TU Dresden's "Center for Information Services and High-Performance Computing". The weather code COSMO-SPECS+FD4 ran on 17 nodes connected with an FDR Infiniband network (240 cores in total). Another large production code (CP2K), small test codes, and the Intel MPI benchmarks have been ported to our OS prototype. They comprise in the order of millions of lines of C, C++, and FORTRAN source code, which we did not have to modify.

While source-level compatibility to existing platforms is a great advantage, we believe that efficient management of the platform at exascale requires cooperation from applications. In particular, application-level hints can improve efficiency for fault-tolerance mechanisms (e.g., optimize checkpoint/restart or avoid it altogether, if algorithms can compensate for failing nodes). They can also enable load-balancing components predict what resources will be needed next (e.g., CPU or memory required by a process to compute the next time step) and whether migration is worthwhile.

We develop in cooperation with application partners novel interfaces such that an HPC program can pass hints to the OS/R about what it is doing and what resources it needs. The ultimate goal of the project is to build and integrate mechanisms of a self-organizing, fault-tolerant HPC platform that frees application developers from micromanaging the dynamic exascale hardware we expect in the future.

- Hermann Härtig
- Carsten Weinhold

Department
of Computer
Science, TU
Dresden,
Germany

References

- [1] FFMK project website: <http://ffmk.tudos.org>
- [2] Härtig, H., Hohmuth, M., Liedtke, J., Schönberg, S., Wolter, J.
The performance of μ -kernel-based systems. In: Proceedings of the 16th ACM Symposium on Operating System Principles (SOSP), pages 66–77, Saint-Malo, France, 1997.
- [3] Barak, A., Drezner, Z., Barak, A., Levy, E., and Shiloh, A.
Resilient gossip algorithms for collecting online management information in exascale clusters, *Concurrency and Computation: Practice and Experience*, 2015.
- [4] Levy, E., Barak, A., Shiloh, A., Lieber, M., Weinhold, C., Härtig, H.
Overhead of a Decentralized Gossip Algorithm on the Performance of HPC Applications, *ROSS 2014*, June 2014, Munich, Germany.
- [5] XtreamFS website: <http://www.xtreamfs.org>

Acknowledgments

This project is supported by the priority program 1648 "Software for Exascale Computing" funded by the German Research Foundation (DFG). The authors acknowledge the Jülich Supercomputing Centre, the Gauss Centre for Supercomputing, and the John von Neumann Institute for Computing for providing compute time on the JUQUEEN supercomputer.

Contact: Carsten Weinhold,
carsten.weinhold@tu-dresden.de

EUDAT2020

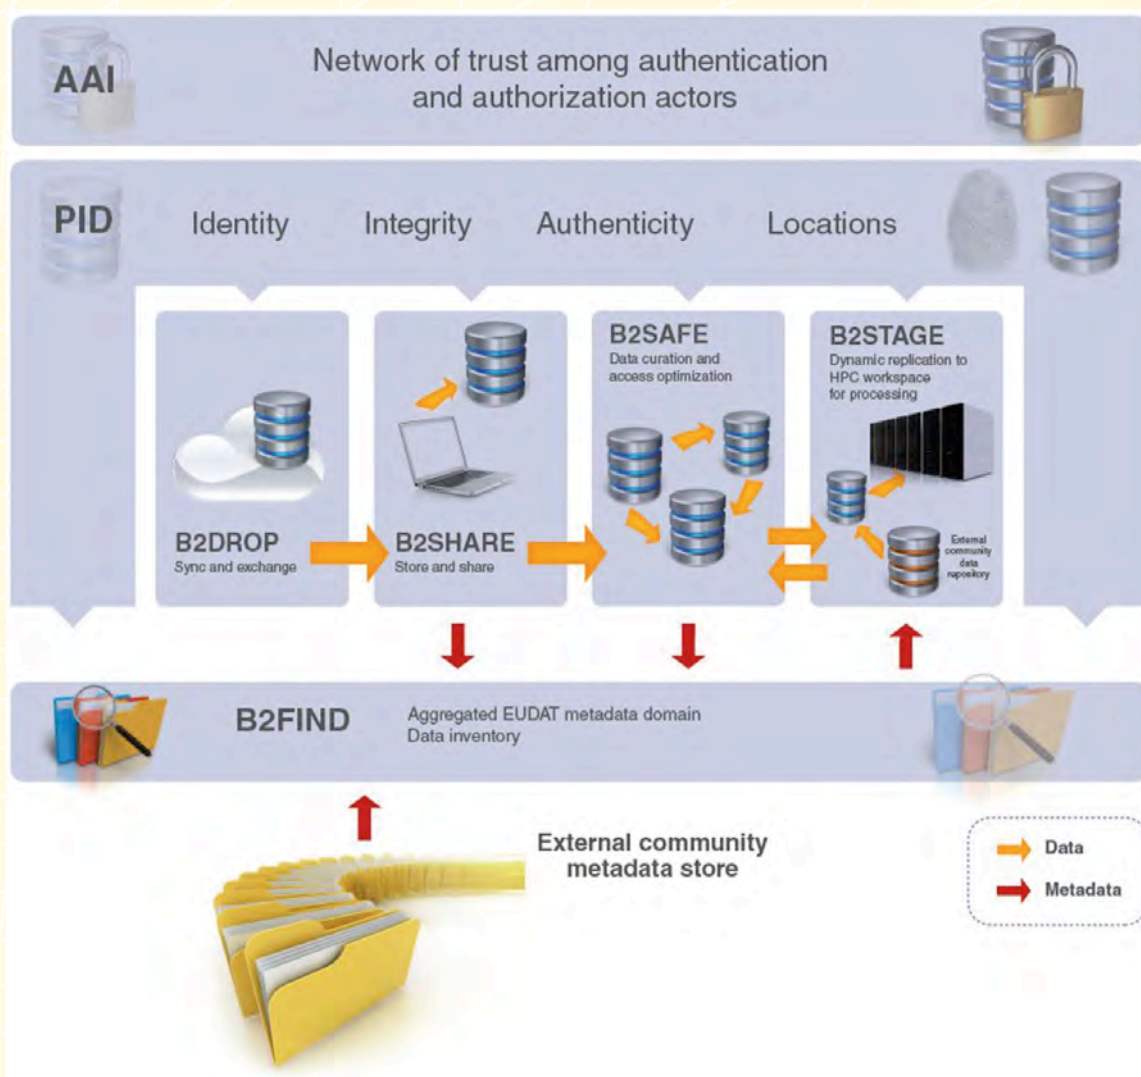


Figure 1: Interaction of EUDAT services

EUDAT2020 brings together a unique consortium of e-infrastructure providers, research infrastructure operators, and researchers from a wide range of scientific disciplines, working together to address the new data challenge. In most research communities, there is a growing awareness that the "rising tide of data" will require new approaches to data management and that data preservation, access and sharing should be supported in a much better way. Data, and in particular Big Data, is an issue touching all research infrastructures.

EUDAT2020's vision is to enable European researchers and practitioners from any research discipline to preserve, find, access, and process data in a trusted environment, as part of a Collaborative Data Infrastructure (CDI) conceived as a network of collaborating, cooperating centres. The CDI combines the richness of numerous community-specific data repositories with the permanence and persistence of some of Europe's largest scientific data centres.

EUDAT2020 builds on the foundations laid by the first EUDAT project (Fig. 1), strengthening the links between the CDI and expanding its functionalities and remit. Covering both access and deposit, from informal data sharing to long-term archiving, and addressing identification, discoverability and computability of both long-tail and big data, EUDAT2020's services will address the full lifecycle of research data.

One of the main ambitions of EUDAT2020 is to bridge the gap between research infrastructures and e-Infrastructures through an active engagement strategy, using the communities that are in the consortium as EUDAT beacons and integrating others through innovative partnerships.

During its three-year funded life, EUDAT2020 will evolve the CDI into a healthy and vibrant data-infrastructure for Europe, and position EUDAT as a sustainable infrastructure within which the future, changing requirements of a wide range of research communities are addressed.

EUDAT Service Suite

Research communities from different disciplines have different needs for data management and analysis, but they all have requirements for basic data services.

EUDAT offers common data services (Fig. 2) that support multiple research communities through a geographically distributed network of general-purpose data centres and community-specific data repositories.

B2DROP is a secure and trusted data exchange service for researchers and scientists to keep their research data synchronized and up-to-date and to exchange it with other researchers.

B2SHARE is a user-friendly, reliable and trustworthy service for researchers, scientific communities and citizen scientists to store and share small-scale research data from diverse contexts. B2SHARE offers the data owner to define access policies, add community-specific metadata and assigns persistent identifiers to data sets.

B2SAFE is a robust, safe and highly available service that allows community repositories to implement data management policies on their research data and replicate data



Figure 2: EUDAT service suite

across multiple administrative domains in a trustworthy manner.

B2STAGE is a reliable, efficient, and lightweight service to transfer research data sets between EUDAT storage resources and High Performance Computing (HPC) workspaces.

B2FIND is a simple, user-friendly metadata catalogue of research data collections stored in EUDAT data centres and other repositories.

• Daniel Mallmann

Jülich
Supercomputing
Centre (JSC),
Germany

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

ORPHEUS – Fire Safety in the Underground

ORPHEUS

Public transport systems form a fundamental part of our cities and play a major role in our society. Thus, civil safety aspects of underground transport facilities are of crucial importance.

Risks and danger manifest in large disasters that take many innocent lives and therefore draw large attention (like 1987 in London, Great Britain or 2003 in Daegu, South Korea, to name only a few). However, there are many more incidents that do not trigger strong medial attention (like recently in Washington, USA, and London, Great Britain, 2015).

The demand of safety research in this context is represented by recent projects such as the Swedish project METRO [1] or the German projects OrGaMIR / OrGaMIRPLUS [2]. Both research projects covered selected safety aspects of underground metro stations.

So far, there exist only limited concepts to construct smoke extraction systems in complex underground facilities. Such systems play a major role for the personal safety during fire hazards and therefore smoke extraction systems represent the main aim of the ORPHEUS project – abbreviated **"Optimierung der Rauchableitung und Personenführung in U-Bahnhöfen: Experimente und Simulationen"**. Thereby, the main obstacle in designing safety concepts for underground sta-

tions is the thermally driven movement of hot smoke or toxic gases. The gas dynamics may become very intensive and hard to control. This puts the passengers in a very dangerous situation, since smoke has a very negative impact on the rescue process, as it limits the mobility and visibility of the occupants and the rescue teams. Hence, most fatalities occur due to the toxic nature of smoke. In contrast to common stations or street tunnels, underground stations often have a very low ceiling height and therefore cannot sustain a smoke free layer, which is large enough to prevent suffocation.

As a consequence, fires in underground stations massively challenge the rescue operation and demand an



Figure 1: Typical metro stations are complex and compact structures. Hot smoke and passengers have in general the same movement paths to the surface. The figures illustrate the metro station "Osloer Strasse" in Berlin.

effective inter-organisational crisis management. Considering the high relevance, the presented project investigates concepts and strategies to improve the safety of underground stations during a fire with an intensive smoke yield. This covers technical as well as inter-organisational aspects. February 2015 saw the start of the ORPHEUS project, which is funded by the Federal Ministry of Education and Research (BMBF) for 36 months [3]. The consortium is coordinated by JSC's division "Civil Safety and Traffic" [4] and consists of the following funded partners:

- Bundesanstalt für Materialforschung und -prüfung
- IBIT GmbH
- Imtech Deutschland GmbH
- Institut für Industrieaerodynamik GmbH
- Ruhr-Universität Bochum, Lehrstuhl für Höhlen- und U-Bahnklimatologie

and the associated partners respectively subcontractors

- Berliner Feuerwehr
- Berliner Verkehrsbetriebe
- Deutsche Bahn Station&Service AG
- Hekatron Vertriebs GmbH
- Karstadt GmbH
- NVIDIA GmbH
- Team HF PartG

All of them met at the project kick-off meeting in Jülich on 12 March 2015. The proposed research plan is divided into three main parts:

- 1) Fire experiments in existing stations and assessment of the current state-of-the-art of technical fire safety engineering systems and methods, safety concepts, detection systems as well as the validity of numerical tools.
- 2) Technologies and concepts for personal and operational safety in case of fire with emphasis on people with special requirements and handicaps covering physical and numerical modelling of smoke extraction and guiding systems as well as numerical pedestrian simulations.
- 3) Reactive situation detection and support systems for rescue teams with integration of the project's results in inter-organisational crisis management systems.



Figure 2: Participants of the kick-off meeting in Jülich

Fire Experiments

The first of the three main parts consists of the execution and validation of experiments in an operative metro station in Berlin. Outside of operating hours, real fire experiments will be carried out in a fully monitored station. The target station (Osloer Strasse) has a complex structure with three levels. Multiple experiments will be carried out with real fires. These technical fires produce no soot or toxic gases and are only used as a dynamic heat source to model thermally driven flows. The diagnostic is based on tracer gases (e.g. SF₆) that are added to the fire plume. The obtained data (e.g. velocity, temperature, and tracer gas concentrations) measured at many locations will be used to validate small scale and numerical models. These experiments are accompanied by long-term underground climate measurements and will apply novel techniques that make use of existing communication cables in the tunnels. The collected climate data will also be considered during the concept phase.

Preventive Fire Safety Concepts

The second part aims to investigate concepts for novel smoke control and extraction systems as well as evacuation aspects. During the initial phase of the project, the fire safety objectives will be defined and used to compare different safety systems and concepts. Small-scale



Figure 3: A snapshot of a smoke spread simulation in an underground station

physical experiments (1:5 and 1:20) and numerical simulations of smoke and heat propagation in underground stations will form the basis for the project studies.

Besides the application of existing CFD (Computational Fluid Dynamics) models, new concepts for mesh-adaptive methods will be developed by JSC. They will allow refining the numerical mesh in regions of high dynamics and gradients, while passive regions will be resolved by a coarse mesh. These techniques are crucial to provide sufficient numerical resolution




Figure 4: First steps toward mesh adaptive simulations. Here, the structure of a hot plume is well resolved, while the surrounding may stay coarsely represented.

inside small subdomains of a large and complex structure. An accompanying goal is to develop a fire simulation model that can make efficient use of HPC systems.

Special emphasis is put on the availability of evacuation paths and routes used by fire fighters. Therefore, pedestrian simulations will accompany the work on preventive technical systems. These simulations will be coupled to fire simulations and will therefore allow evaluating the impact of smoke on the evacuation (e.g. visibility or toxic affec-

tion). Additionally, the pedestrian simulation model JuPedSim [5], developed in-house at JSC, will be expanded to include psychological decision models.



Figure 5: Mobility and vision handicapped people pose special requirements, which are considered in this project.

Humans with special requirements, like sight and/or mobility handicapped people (e.g. elderly), will play a major role in the analysis and optimisation of evacuation routes.

Support During Rescue Operations

The third part covers the interaction between the operators, emergency services, and third parties, like shopkeepers. In the case of fire, processes on the surface are essential for rescue attempts. Revision of past events and interviews with organisations involved will allow a communication pattern

analysis – also with respect to warnings and alarms. Besides the individual organisation analysis, the cooperation and interaction between affected organisations before and during an incident is studied. The investigation covers the transfer of knowledge and information between the emergency services and the communication to infrastructure users and concerned residents.

To tackle the complex interaction scheme, an inter-organisational map and interfaces will be developed. The interfaces include communication paths as well as interaction points in the individual crisis management. This work will allow identifying inconsistent processes and potential for optimisation. Other important communicative aspects are warnings and alarms. Thereby, investigated topics are the contents and forms of the messages sent to the occupants.

To support rescue teams and technical systems, JSC will also examine a real-time smoke propagation prognosis model. This CFD based model is supposed to run on GPGPU systems to provide a short time-to-solution and should be coupled to sensors to include live parameters such as the position or the intensity of the fire as well as climate data.




Figure 6: Surface events like markets may be located directly at the metro's exits and therefore amplify the complexity of a rescue operation.

Conclusion

All concepts created in this project focus on underground metro stations. However, during the final phase of the project the outcome will be evaluated for the application to other transport infrastructures, like airports or car parks. The technical and organisational results form the scientific basis to develop production systems and tools for fire safety engineering.

Funding

BMBF, call "Zivile Sicherheit – Schutz und Rettung bei komplexen Einsatzlagen"

Funding amount

Euro 3.2 million

Project duration

February 2015 – January 2018

Project website

<http://www.orpheus-projekt.de>

Links

- [1] <http://www.metroproject.se>
- [2] <http://www.bmbf.de/pubRD/OrGaMIR.pdf>
- [3] http://www.bmbf.de/pubRD/Projektumriss_ORPHEUS.pdf
- [4] <http://www.fz-juelich.de/ias/jsc/cst>
- [5] <http://www.jupedsim.org/>

- Lukas Arnold

Jülich
Supercomputing
Centre (JSC),
Germany

contact: Lukas Arnold,
l.arnold@fz-juelich.de

Seamless Continuation: 4th PRACE Implementation Phase Project

In response to the European Commission's call for proposal in 2014 within the new European framework programme Horizon 2020 PRACE partners from 25 countries submitted a successful proposal and started the 4th Implementation Phase project (PRACE-4IP) on 1 February 2015. The project will support the transition from the initial five-year period (2010-2015) of the Research Infrastructure established by the Partnership for Advanced Computing in Europe (PRACE) to PRACE 2.0.

Key objectives of PRACE-4IP are:

Ensure long-term sustainability of the infrastructure. The project will assist the PRACE Research Infrastructure (PRACE RI) in managing the transition from the business model used in the Initial Period. It demonstrated the case for a European HPC research infrastructure relying on the strong engagement of four hosting partners (BSC representing Spain, CINECA representing Italy, GCS representing Germany and GENCI representing France) who funded and deployed the petaflop/s systems used by PRACE RI.

Promote Europe's leadership in HPC applications. Scientific and engineering modelling and simulation require capabilities of supercomputers. The project will enable application codes for PRACE leadership platforms and prepare for future systems, notably those with architectural innovation embodied in accelerators

or co-processors, by investigating new programming tools and developing suitable benchmarks.

Increase European human resources skilled in HPC and HPC applications.

The project will contribute by organizing highly visible events, enhancing the state of the art training provided by the PRACE Advanced Training Centres (PATCs), targeting both the academic and industrial domains. On-line training will be improved and a pilot deployment will assess a Massively open online Course (MooC).

Support a balanced eco-system of HPC resources for Europe's researchers.

The project will contribute to this objective through tasks addressing: a) the improvement of PRACE operations; b) the prototyping of new services including "urgent computing", the visualization of extreme size computational data, and the provision of repositories for open source scientific libraries. Links will be established with other e-infrastructures, the Centres of Excellence which will be created in Horizon 2020 and the existing international collaborations will be extended.

Evaluate new technologies and define Europe's path for using ExaFlop/s resources.

The project will extend its market watch and evaluation based on user requirements, study best practices for energy-efficiency and lower environmental impact throughout the life cycle of large

HPC infrastructures and define best practices for prototype planning and evaluation. This will contribute to solve a wide range of technological, architectural and programming challenges for the exaflop/s area.

Disseminate effectively the PRACE results. This targets engaging European scientists and engineers in the wider utilisation of high end HPC. The project will continue to organise well known events like PRACEdays, Summer of HPC or the International HPC Summer School in order to promote and support innovative scientific approaches in modelling, simulation and data-analysis. With the extended presence at conferences (e.g. SC, ISC or ICT) the project is seeking wider support of the general

public for investment in HPC, in particular by illustrating success stories and raising awareness in the potential for development.

PRACE-4IP is again coordinated and managed by by Forschungszentrum Jülich. It has a budget of nearly Euro 16,5 Mio including an EC contribution of Euro 15 Mio. The duration will be 27 month.

Over 250 researchers collaborate in PRACE from 49 organisations¹ in 25 countries. 106 collaborators met in Ostrava from 28-29 April 2015, for the PRACE-4IP kick-off meeting. The meeting was organised by IT4Innovation-VSB and held at the Technical University of Ostrava, Czech Republic.



Figure 1: PRACE collaborators at the PRACE-4IP kick-off meeting

Synopsis of the PRACE Projects

The European Commission supported the creation and implementation of PRACE through five projects with a total EC funding of Euro 82 Mio. The partners co-funded the projects with more than Euro 33 Mio in addition to the commitment of Euro 400 Mio for the initial period by the hosting members to procure and operate Tier-0 systems and the in-kind contribution of Tier-1 resources on systems at presently twenty partner sites. The following table gives an overview of the PRACE projects.

Project Id Grant Number	Number Partners	Euro Budget Mio.	Euro EC Funds Mio.	Duration	Status
PRACE-PP RI-211528	16	18.9	10.0	1.1.2008 - 30.6.2010	completed
PRACE-1IP RI-261557	21	28.5	20.0	1.7.2010 - 31.8.2013	completed
PRACE-2IP RI-283493	22	25.4	18.0	1.9.2011 - 31.8.2013	completed
PRACE-3IP RI-312763	26	26.8	19.0	1.7.2012 - 30.6.2014	completed PCP ² to end June 30, 2016
PRACE-4IP 653838	26	16.5	15.0	01.02.2015 - 30.4.2017	started
(total)		116.0	82.0		



• Florian Berberich

Jülich Supercomputing Centre (JSC), Germany

¹ This includes 26 Beneficiaries and 23 linked Third Parties from associated universities or centres.

² The Pre Commercial Procurement Pilot (PCP) is part of the PRACE-3IP project, however with a duration of 48 month.

Intel® Parallel Computing Center (IPCC) at



The Intel® Parallel Computing Centers (IPCC), are novel initiatives established and financially supported by Intel within universities, institutions, and labs that are leaders in their own fields, and focus on modernizing scientific applications to increase parallelization efficiency and scalability through optimizations that leverage cores, caches, threads, and vectorization capabilities of microprocessors and coprocessors of Intel [1]. The main target architectures for these centers are the Intel MIC (Many Integrated Cores), the Intel Xeon Phi coprocessor [1].

The IPCC at Garching Research campus, consisting of two closely collaborating partners, LRZ and the Computer Science Department, Technische Universität München (TUM) [2], was started in July, 2014 and is funded for two years. The Garching IPCC focuses on the optimization of four different highly acclaimed applications from different areas of science and engineering: earthquake simulation and seismic wave propagation with SeisSol [3,4], simulation of cosmological structure formation using Gadget3 [5], the molecular dynamics code Is1 mardyn [6] developed for applications in chemical engineering, and the software framework SG++ [7] to tackle high-dimensional problems in data mining or financial mathematics (using sparse

grids). All these codes have already demonstrated very high scalability on SuperMUC (with performance up to a Petaflop/s, see SeisSol ACM Gordon Bell paper [4]), but are in different stages of development, e.g. with respect to running on the Intel MIC architecture. While particularly targeting the new Intel architectures, e.g. Xeon Phi (MIC) coprocessors, the project also simultaneously tackles some of the fundamental challenges that are relevant for most supercomputing architectures – such as parallelism on multiple levels (nodes, cores, hardware threads per core, data parallelism) or compute cores that offer strong SIMD capabilities with increasing vector register width, e.g. Intel Haswell architecture [2].

Within this project, LRZ contributes on optimizing Gadget3, a numerical simulation code for cosmological structure formation, for the Intel Haswell and the MIC architectures. Gadget3 has already established itself as a community code and been already scaled fairly efficiently on the entire SuperMUC but Xeon Phi was basically an unexplored territory; no MIC version of Gadget3 existed at the project start. The Garching IPCC at the same time will also provide the opportunity to draw out the best possible performance with such codes from the Su-

perMUC Haswell extension or potential SuperMUC successors. The far-reaching goal of this project is to establish a model process and collection of experiences as well as best practices for similar codes in the computational sciences. To prepare simulation software for this new platform and to identify possible roadblocks at an early stage, one has to tackle two expected major challenges: (1) achieving a high fraction of the available node-level performance on (shared-memory) compute nodes and (2) scaling this performance up to the range of around 10,000 compute nodes. Concerning node-level performance, the Garching IPCC considers compute nodes with one or several Intel Xeon Phi (co-) processors. A respective small evaluation platform (32 nodes, each with 2 MIC co-processors and powerful Infiniband fabric between the nodes), known as SuperMIC, is in operation at LRZ since June 2014 [8]. Scalability on large supercomputers is studied on SuperMUC, which in its upgrade phase in 2015 is extended by a second 3-PetaFlop/s partition based on the new Intel Haswell architecture. All successful code optimizations and improvements done within the IPCC works will be integrated in the regular software releases of the four research codes and will therefore provide feedback to the scientific user communities. The initial results of the code optimization and re-engineering were presented in the first IPCC User Forum Meeting in Berlin (Sep, 2014). More results on performance and MIC related improvements were showcased in the second IPCC User Forum Meeting in Dublin (Feb, 2015). To widen the community outreach of the IPCC works LRZ, together with its funding partner Intel, has also taken initiatives to spread the MIC know-how

and best known methods to its interested HPC users. A workshop, as part of the LRZ activities as an Intel Parallel Computing Center, was organized together with Intel during Oct 20-21, 2014 at LRZ [9]. The main objective of this workshop was to learn and discuss advanced Xeon Phi programming and analysis methodologies. Case Studies of real world codes have been presented as hints for planning the optimal strategy of Xeon Phi Coprocessor usage. Intel engineers were also available supporting this effort. The major topics covered under this workshop are Xeon Phi Coprocessor Architecture, Coprocessor Software Infrastructure and Programming (e.g. Vectorization, OpenMP and Offloading). Furthermore, Intel MPI best practices for native and symmetric usage in combination with offloading were discussed. To analyze user applications Intel engineers introduced the Intel Trace Analyzer and Collector (ITAC) which can profile and visualize Intel MPI based parallel application behavior, the Intel VTune Amplifier XE for advanced sequential and threading analysis e.g. Vectorization analysis, Debugging for Xeon Phi. Case Studies for symmetric MPI and for Offload in combination with MPI on




Figure 1: Intel® Xeon Phi™ Coprocessor (image courtesy ©Intel Inc.)

the host only were demonstrated to the audience. Special emphasis was given on "hands-on" sessions on the four IPCC codes. There were 21 participants from different German universities and research institutes. Due to its overall success and extremely high demand from the users the IPCC is planning a follow-up workshops in the near future.

Project Partners

Principle collaborators

1. Leibniz Supercomputing Centre of Bavarian Academy of Sciences and Humanities

PI: Arndt Bode

Responsible for Gadget3

2. Technische Universität München, Department of Informatics

PI: Michael Bader

Responsible for Seissol.

PI: Hans-Joachim Bungartz

Responsible for ls1 mardyn and SG++

Associated Partners

1. Ludwig-Maximilians-Universität, Geo and Environmental Sciences

Heiner Igel and Christian Pelties

2. Ludwig-Maximilians-Universität, USM, Computational Astrophysics

Klaus Dolag

3. Intel Corporation USA, Intel Labs, Parallel Computing Lab

Alexander Heinecke

References

[1] Intel Parallel Computing Center
<https://software.intel.com/en-us/ipcc>

[2] IPCC at LRZ and TUM
<https://software.intel.com/en-us/articles/intel-parallel-computing-center-at-leibniz-supercomputing-centre-and-technische-universit-t>

[3] **Dumbser, M., Käser, M.**
 An arbitrary high order discontinuous Galerkin method for elastic waves on unstructured meshes II: The three-dimensional case. *Geophys. J. Int.* 167(1), pp. 319–336, 2006. Dumbser, M., Käser, M., Toro, E.F. An arbitrary high order discontinuous Galerkin method for elastic waves on unstructured meshes V: Local time stepping and p-adaptivity. *Geophys. J. Int.* 171(2), pp. 695–717, 2007.

[4] **Heinecke, A., Breuer, A., Rettenberger, S., Bader, M., Gabriel, A., Pelties, C., Bode, A., Barth, W., Liao, X., Vaidyanathan, K., Smelyanskiy, M., and Dubey, P.**
 Petascale high order dynamic rupture earthquake simulations on heterogeneous supercomputers. *Proc. Int. Conf. High Performance Computing, Networking, Storage and Analysis (SC '14)*. IEEE Press, Piscataway, NJ, USA, 3-14, 2014 (DOI=10.1109/SC.2014.6 <http://dx.doi.org/10.1109/SC.2014.6>)

[5] **Bazin, G., Dolag, K., Hammer, N.**
 Gadget3: Numerical Simulation of Structure Formation in the Universe, inside 11, 2, 2013.

[6] **Niethammer, C., Becker, S., Bernreuther, M., Buchholz, M., Eckhardt, W., Heinecke, Werth, A. S., Bungartz., H.-J., Glass, C. W., Hasse, H., Vrabec, J., and Horsch, M.**
 Ls1 mardyn: The massively parallel molecular dynamics code for large systems., *Journal of Chemical Theory and Computation* 10 (10): pp. 4455-4464, 2014.

[7] **Pflüger, D.**
 Spatially Adaptive Sparse Grids for Higher-Dimensional Problems. Dissertation, Verlag Dr. Hut, München, 2010. ISBN 9-783-868-53555-6.

[8] **Weiberg, V., Allalen, M.**
 The new Intel Xeon Phi based System SuperMIC at LRZ, inside 12, 2, Autumn 2014.

[9] IPCC Intel Xeon Phi Coprocessor Workshop
http://www.lrz.de/services/compute/courses/archive/2014/2014-10-20_hmic1w14/

- Anupam Karmakar
- Nicolay Hammer
- Luigi Iapichino
- Vasileios Karakasis

Leibniz
 Supercomputing
 Centre (LRZ),
 Germany



contact: Anupam Karmakar,
Anupam.Karmakar@lrz.de
Nicolay Hammer,
Nicolay.Hammer@lrz.de
Luigi Iapichino,
Luigi.Iapichino@lrz.de
Vasileios Karakasis,
Vasileios.Karakasis@lrz.de



Figure 2: Participants of the IPCC Intel Xeon Phi Coprocessor Workshop during October 20 - 21, 2014 at LRZ

JURECA: Jülich Research on Exascale Cluster Architectures



Figure 1: (Left) Front view of a V-class enclosure. (Right) Rear view of a V-class enclosure hosting ten V210S compute nodes (slim or fat nodes of type 1 in JURECA). The same chassis can host five V210F accelerated compute nodes occupying twice the space of V210S node.

According to its dual architecture strategy, Forschungszentrum Jülich is offering access to a leadership-class capability system and to a general-purpose supercomputing system. The latter meets the users' need for mixed capacity and capability computing time. Since 2009, the JUROPA (Jülich Research on Petaflop Architectures) cluster, based on Intel Nehalem CPUs and quad data rate InfiniBand networking technology, has taken up this challenge and has enabled outstanding science by researchers from around the world. Now, Jülich Supercomputing Center (JSC) has started to install the next-generation general-purpose cluster JURECA (Jülich Research on Exascale Cluster Architectures) which will supersede JUROPA.

JURECA will be based on latest-generation Intel Haswell CPUs and provide a peak floating point

performance of roughly 1.8 petaflops per second – a six-fold increase over JUROPA. With its high speed connection of about 100 GiBps to the center-wide exported GPFS file systems, JURECA will not only serve the widest variety of user communities from the traditional computational science disciplines but will also be a welcoming home to data-intensive science and big data projects.

JURECA will be built from V-class blade servers of the Russian supercomputing vendor T-Platforms. Once fully installed, the system will consist of more than 1,800 compute nodes with two Intel Haswell E5-2680 v3 12 core CPUs per node. About 1,680 compute nodes will be equipped with 128 GiB DDR4 main memory, i.e., more than 5 GiB per core. In support for workflows requiring more memory per core, additional 128 nodes with 256 GiB and 64 nodes with

512 GiB DDR4 RAM will be available. 75 nodes will be equipped with two NVIDIA K80 cards each, providing an additional flop rate of 430 teraflops per second for accelerator-capable applications. 12 Login nodes with 256 GiB per node will be available for workflow- and data-management as well the convenient execution of short running pre- and post-processing operations. Additionally, 12 visualization nodes with 512 GiB (10 nodes) and 1 TiB main memory (2 nodes) and two NVIDIA K40 GPUs are available for advanced visualization purposes. JURECA's visualization partition will replace the older JUVIS visualization cluster at JSC and move the analysis of data closer to its source. An overview of the different node types in JURECA is shown in Table 1. Fig. 1 and Fig. 2 show the V-class chassis and blade servers employed in JURECA.

The Haswell CPUs in JURECA support the AVX 2.0 instruction set architecture extension and can perform two 256-bit (i.e., two times 4 double precision floating point numbers) wide multiply-add operations per cycle. Due to the increased core count and improved microarchitecture of the Haswell CPUs,

the peak floating point capabilities of compute nodes in JURECA are 10 times higher than that of a JUROPA node. From the user perspective, however, this performance improvement does not come for free but requires code optimizations and potentially refactoring in order to take advantage of the wider single instruction multiple data (SIMD) units of the Haswell CPUs. Since October 2014, JSC is providing JUROPA users with access to the 70 TFlops/s Haswell cluster JUROPATEST to foster such efforts.

JURECA will be interconnected with a cutting-edge Mellanox EDR (extended data rate) InfiniBand network. EDR InfiniBand with four lanes per direction achieves a unidirectional point-to-point bandwidth of 100 Gbps – a significant bandwidth improvement compared to the 4x QDR InfiniBand technology in use by JUROPA. As in JUROPA, the JURECA components will be organized in a fully non-blocking fat-tree topology. The core of the fabric is constituted by four 648-port EDR director switches connected to more than a hundred 36-port leaf switches located in the compute node racks.



Figure 2: A V210S compute node as used in the JURECA cluster. Each node hosts two CPU sockets. The V201F version, which has twice the width of a V201S module, can additionally host two PCIe-based accelerator cards.

JURECA will not be equipped with a system-private (global) file systems like JUROPA, but will be connected to JSC's storage cluster JUST4 and mount the work and home file systems from there. For users with access to multiple computing systems at JSC this consolidation of the storage landscape simplifies the data management and reduces data movement across the center. The connection to JUST4 will be based on InfiniBand and Ethernet technologies. By utilizing the high-bandwidth InfiniBand network for storage, JURECA will not only feature a high accumulated storage bandwidth but individual nodes will be able to sustain a significant portion of this total bandwidth. This design choice ensures JURECA's ability to service the widest variety of users' requirements from capacity to capability computing and traditional computational to emerging big data sciences. To bridge between the high speed InfiniBand network and JSC's Ethernet-based storage backbone network, Mellanox gateway switches are employed in an active/active mode pro-

viding high bandwidth and reliability through failover mechanisms. The employed gateway switches will each route traffic between eighteen FDR (Fourteen Data Rate) InfiniBand links on one side and eighteen 40GigE links on the other side. While JURECA is built around EDR InfiniBand technology, the storage connection is realized with lower performing FDR links due to the available market offerings.

JURECA's cutting edge hardware setup is matched by a state-of-the-art software stack. The system will be launched with a CentOS 7 Enterprise-Linux installation featuring a 3.10 Linux kernel. The main MPI (Message Passing Interface) implementation will be ParaStation MPI which, in the newest version available on JURECA, supports MPI-3.0. Additionally, Intel MPI will be supported on JURECA. JURECA will be the first large-scale system at JSC on which the open-source Slurm workload manager will be employed. In the context of the JUROPA collaboration a new plugin for the ParaStation resource management

Node Type	Number	Characteristics
Standard/Slim	1.605	2x Haswell E5-2680 v3, 128 GiB DDR4 RAM
Fat Type 1	128	2x Haswell E5-2680 v3, 256 GiB DDR4 RAM
Fat Type 2	64	2x Haswell E5-2680 v3, 512 GiB DDR4 RAM
Accelerated	75	2x Haswell E5-2680 v3, 128 GiB DDR4 RAM, 2x K80 GPUs
Login	12	2x Haswell E5-2680 v3, 256 GiB DDR4 RAM
Visualization Type 1	10	2x Haswell E5-2680 v3, 512 GiB DDR4 RAM, 2x K40 GPUs
Visualization Type 2	2	2x Haswell E5-2680 v3, 1 TiB DDR4 RAM, 2x K40 GPUs

Table 1: Overview of the different node types in the JURECA system.

daemon has been developed by ParTec together with JSC. This plugin allows for using the Slurm batch system in combination with the ParaStation resource management – the resource management of choice on JSC's clusters – without requiring additional daemons on the compute nodes that would inject spurious jitter.

Building on the JUROPA experience – where the successful collaboration between JSC and the hardware and software vendors has contributed to the exceptional life span of the system – JSC, T-Platforms and the software provider ParTec are engaging in a collaborative project to further develop and augment the JURECA system after installation and to approach urgent research questions in the scalability of large-scale cluster systems.

The installation of JURECA will proceed in two phases so as to minimize the service interruption for the users by allowing to install the first JURECA phase while JUROPA continues to operate. The first phase of JURECA consists of only six racks but delivers a performance equivalent to the current JUROPA system. It thus allows users to continue working while the JUROPA system is dismantled to free up space for the remaining 28 JURECA racks. The installation of this second phase will be done during production and only a short offline maintenance window is required to integrate phase one and two in a single fat-tree InfiniBand network and perform benchmarks as part of the acceptance procedure.

contact: Dorian Krause,
d.krause@fz-juelich.de

• Dorian Krause

Jülich
Supercomputing
Centre (JSC),
Germany



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

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

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

Contact:

Leibniz Supercomputing Centre

Prof. Dr. Arndt Bode
Boltzmannstr. 1
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/Lenovo System x Cluster "SuperMUC Phase 1"	18 thin node islands IBM/Lenovo iDataPlex 512 nodes per island with 2 Intel Sandy Bridge EP processors each 147,456 cores 288 TByte main memory FDR 10 IB	3,185	Capability Computing	German universities and research institutes, PRACE projects (Tier-O System)
	1 fat node island with 205 nodes IBM/Lenovo Bladecenter HX5 4 Intel Westmere EX processors each 52 TByte main memory QDR IB	78	Capability Computing	
	32 accelerated nodes with 2 Intel Ivy Bridge EP and 2 Intel Xeon Phi each 72 GByte main memory Dual-Rail FDR 14 IB	75	Prototype System	
Lenovo/IBM Nextscale Cluster "SuperMUC Phase 2"	6 medium node islands Lenovo/IBM Nextscale 512 nodes per island with 2 Intel Haswell EP processors each 86,016 cores 197 TByte main memory FDR 14 IB	3,303	Capability Computing	German universities and research institutes, PRACE projects (Tier-O System)
Linux-Cluster	510 nodes with Intel Xeon EM64T/AMD Opteron 2-, 4-, 8-, 16-, 32-way 2,030 Cores 4.7 TByte	13.2	Capability Computing	Bavarian and Munich Universities, LCG Grid
SGI Altix 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

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



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 access 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. Hon. Prof.
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 XC40 "Hornet"

Compute servers currently operated by HLRS

System	Size	Peak Performance (TFlop/s)	Purpose	User Community
Cray XC40 "Hornet"	3,944 dual socket nodes with 94,656 Intel Haswell cores	3,786	Capability Computing	European (PRACE) and German Research Organizations and Industry
NEC Cluster (Laki, Laki2) heterogenous computing platform of 2 independent	23 TB memory 9,988 cores 911 nodes	170	Laki: 120,5 TFlops Laki2: 47,2 TFlops	German Universities, Research Institutes and Industry
NEC SX-ACE	64 compute nodes with 256 cores 4 TB memory	16	Vector Computing	German Universities, Research Institutes and Industry



HLRS Cray XC40 "Hornet" frontdoor image



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 international peer-review procedures implemented by the John von Neumann Institute for Computing (NIC), GCS, and PRACE.

Supercomputer-oriented research and development in selected fields of physics and other natural sciences by

research groups and in technology, e.g. by doing co-design together with leading HPC companies.

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

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

Contact:

Jülich Supercomputing Centre (JSC)
Forschungszentrum Jülich

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




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 "JURECA"	1,884 SMT nodes with 2 Intel Haswell 12-core 2.5 GHz processors each 150 graphics processors (NVIDIA K80) 281 TByte memory	2,245	Capacity and Capability Computing	European Universities, Research Institutes and Industry
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

New Books in HPC

High Performance Computing in Science and Engineering '14



Transactions of the High Performance Computing Center Stuttgart (HLRS), 2014


Editors: Nagel, W. E., Kröner, D. H., Resch, M. M.

This book presents the state-of-the-art in supercomputer simulation. It includes the latest findings from leading researchers using systems from the High Performance Computing Center Stuttgart (HLRS). The reports cover all fields of computational science and engineering ranging from CFD to computational physics and from chemistry to computer science with a special emphasis on industrially relevant applications. Presenting findings of one of Europe's leading systems, this

volume covers a wide variety of applications that deliver a high level of sustained performance.

The book covers the main methods in HPC. Its outstanding results in achieving the best performance for production codes are of particular interest for both scientists and engineers. The book comes with a wealth of color illustrations and tables of results.

Sustained Simulation Performance 2014




Proceedings of the Joint Workshop on Sustained Simulation Performance, University of Stuttgart (HLRS) and Tohoku University, 2014

Editors: Resch, M. M., Bez, W., Focht, E., Kobayashi, H., Patel, N.

This book presents the state-of-the-art in High Performance Computing and simulation on modern supercomputer architectures. It covers trends in hardware and software development in general and the future of high-performance systems and heterogeneous architectures in particular. The application-related contributions cover computational fluid dynamics, material science, medical applications and climate research; innovative fields such as coupled multi-physics and multi-scale simulations are highlighted. All papers were chosen from presentations given at the 18th Workshop on Sustained Simulation Performance held at the HLRS, University of Stuttgart, Germany in October 2013 and subsequent Workshop of the same name held at Tohoku University in March 2014.

Tools for High Performance Computing 2014



Proceedings of the 8th International Workshop on Parallel Tools for HPC, 2014

Editors: Knüpfer, A., Gracia, J., Nagel, W. E., Resch, M. M.

Current advances in High Performance Computing increasingly impact efficient software development workflows. Programmers for HPC applications need to consider trends such as increased core counts, multiple levels of parallelism, reduced memory per core, and I/O system challenges in order to derive well performing and highly scalable codes. At the same time, the increasing complexity adds further sources of program defects. While novel programming paradigms and advanced system libraries provide solutions for some of these challenges, appropriate supporting tools are indispensable. Such tools aid application developers in debugging, performance analysis, or code optimization and therefore make a major contribution to the development of robust and efficient parallel software. This book introduces a selection of the tools presented and discussed at the 8th International Parallel Tools Workshop, held in Stuttgart, Germany, October 1-2, 2014.

Third JUQUEEN Porting and Tuning Workshop




Figure 1: Participants of the Third JUQUEEN Porting and Tuning Workshop (© Forschungszentrum Jülich GmbH).

Jülich Supercomputing Centre (JSC) continued its successful series of JUQUEEN Porting and Tuning Workshops from 2nd to 4th of February this year. The PRACE Advanced Training Centre (PATC) course attracted over 20 participants from various institutions in three European countries. The workshop familiarised the participants with the Blue Gene/Q supercomputer installed at JSC, including the provided toolchain from compilers and libraries to debuggers and performance analysis tools. The participants received help porting their codes, analysing execution performance and scalability, and

in improving the efficiency of their applications and workflow. Each year the workshop also focuses on one special user group, this time inviting users from the Earth system modelling community. This activity was supported by the Simulation Laboratory Climate Science and the Simulation Laboratory Terrestrial Systems.

The programme of the course started with overviews of JUQUEEN's hardware and software environment. A summary of the best practices was then followed by talks on performance analysis tools, debugging, and efficient I/O. The

workshop concluded with tips on very specific hardware features of the Blue Gene/Q architecture (QPX and TM/SE). At the heart of the workshop were hands-on sessions with the participants' codes, supervised by members of staff from JSC's Simulation Laboratories and cross-sectional teams (Application Optimisation, Performance Analysis, Mathematical Methods and Algorithms) as well as IBM. The general programme was accompanied by sessions for the Earth system modelling groups with short talks by the participants and members of the two Simulation Laboratories involved. Those covered a large variety of codes and applications from process and sensitivity studies to numerical weather prediction and climate change projections. Examples shown included ensemble approaches with multiple concurrent realisations, parallel data assimilation frameworks, and innovative variable model grids. In general, the challenge here is to optimise the often large legacy codes used in geosciences with their multiphysics models (clouds and precipitation, convection, chemistry, radiation, groundwater, sediment transport) and multiple spatial and temporal scales from riverbed water percolation and fine sediment movement to global-scale climate simulations.

The slides of the talks can be found on the web at <http://www.fz-juelich.de/ias/jsc/jqws15>

contact: Dirk Brömmel,
d.broemmel@fz-juelich.de
Klaus Görgen, k.goergen@fz-juelich.de
Lars Hoffmann,
l.hoffmann@fz-juelich.de

- Dirk Brömmel¹
- Klaus Görgen^{1, 2}
- Lars Hoffmann¹

¹ Jülich
Supercomputing
Centre (JSC),
Germany

² University of
Bonn, Germany

Helmholtz Portfolio Theme "Supercomputing and Modelling for the Human Brain" - General Assembly 2015

A total of 92 scientists from Forschungszentrum Jülich and several partner institutions collaborating within the Helmholtz Portfolio Theme "Supercomputing and Modelling for the Human Brain" (SMHB) [1] met for their annual General Assembly on March 30th and 31st at the Jülich Supercomputing Centre (JSC), to present and discuss the work progress achieved in 2014 and to agree on the next steps in the project.

The SMHB started in January 2013 as a Portfolio Theme [2] of the Helmholtz Association. Its overarching goal is to better understand the organization and functioning of the human brain by developing a realistic model of the brain. To meet this grand scientific challenge, an appropriate infrastructure for High Performance Computing (HPC) in the exascale range and Big Data analytics needs to be built.

The work plan of the SMHB therefore integrates a wide range of knowledge and expertise from fundamental neuroscience, brain modelling and simulation, simulation software technology, High Performance Computing, large-scale data management, scientific workflows, and interactive visualization and analysis. The SMHB collaborates with the JSC's Exascale Labs for the co-design of neuroscience applications and HPC technology.

The SMHB is embedded into the allied partnership of the two new Helmholtz Programmes "Decoding the Human Brain" and "Supercomputing and Big Data". Both were successfully reviewed in 2014 and started in January 2015 with the third period of the Programme-oriented Funding of the Helmholtz Association. The SMHB was also conceived as part of the Helmholtz contribution to the European Future and Emerging Technology (FET) Flagship "Human Brain Project" [3].

The SMHB General Assembly 2015 was opened by Prof. Wolfgang Marquardt, Chairman of the Board of Directors of Forschungszentrum Jülich, and by the two project speakers, Prof. Katrin Amunts (INM-1) and Prof. Thomas Lippert (JSC). In the course of this meeting, the vivid collaboration in the SMHB became once more evident through many examples of fruitful interactions between the SMHB work packages and tasks. The work plan was also updated in order to adapt the work to upcoming needs and to further strengthen existing or newly established links. For instance, two new tasks "3D cellular architecture" and "multimodal modelling of structure, function and connectivity" were added to the work plan.

Dr. Moritz Helmstaedter from the Max Planck Institute for Brain Research gave an invited and very well received keynote talk on his research topic of Connectomics, in which he introduced the audience to the dense reconstruction of neuronal circuits.

A further highlight of this year's meeting was the SMHB young scientists presenting their work in a spotlight talk session and discussing it afterwards with colleagues and members of the SMHB Scientific Advisory Board with expertise in several fields during the poster session.

References

- [1] <http://www.fz-juelich.de/JuBrain/EN>
- [2] <http://www.helmholtz.de/en/research>
- [3] <https://www.humanbrainproject.eu>

contact: Anne Do Lam-Ruschewski,
a.dolam@fz-juelich.de

- Anne Do Lam-Ruschewski
- Anna Lührs
- Boris Orth

Jülich
Supercomputing
Centre (JSC),
Germany



6th Blue Gene Extreme Scaling Workshop



As an optional appendix to this year's JUQUEEN Blue Gene/Q Porting and Tuning Workshop at Jülich Supercomputing Centre (JSC), two additional days (5-6 February) were offered for select code-teams to (im)prove their applications' scalability to the entire 458,752 cores. This continued the tradition of the initial initial 2006 workshop using the JUBL Blue Gene/L, 2008 workshop using the JUGENE Blue Gene/P and three subsequent workshops dedicated to extreme scaling which attracted participants from around the world. [1,2,3,4]

Seven international application code-teams took up this offer: **CoreNeuron** brain activity simulator (EPFL Blue Brain Project), **FE2TI** scale-bridging incorporating micro-mechanics in macroscopic simulations of multi-phase steels (University of Cologne and TU Freiberg), **FEMPAR** finite-element code for multi-physics problems (from UPC-CIMNE), **ICON** icosahedral non-hydrostatic atmospheric model (DKRZ), **MPAS-A** multi-scale non-hydrostatic atmospheric model for global, convection-resolving climate simulations (KIT and NCAR), **psOpen** direct numerical simulation of fine-scale turbulence (RWTH-ITV and JARA), and **SHOCK** structured high-order finite-difference kernel for compressible flows (RWTH-SWL).

JSC Simulation Laboratories for Climate Science, Fluids & Solids Engineering and Neuroscience assisted the code-teams, along with JSC Cross-sectional Teams, JUQUEEN and IBM technical support.

12 million core-hours were used, covering a 30 hour period with the full 28 racks reserved, and all seven codes managed their first successful execution using all 28 racks within the first 24 hours of access. Figs. 1 & 2 show that the codes demonstrated excellent strong and/or weak scalability, six using 1.8 million MPI processes or OpenMP threads, which improved the existing High-Q Club entry for FEMPAR and qualified five new members for the High-Q Club. MPAS-A unfortunately was not accepted as its scaling was limited to only 24 racks (393,216 cores) with its 1.2 TB 3-km dataset of 65 million grid points.

Detailed workshop reports provided by each code-team, and additional comparative analysis to the other 16 High-Q Club member codes, are available in a technical report and expected to be published in the proceedings of a ParCo conference minisymposium later this year. The workshop surpassed our expectations and completely achieved its goal, with all participants finding it to have been extremely useful as on-hand support made it possible to quickly overcome issues. To follow-up a workshop is being organised at this year's ISC-HPC conference to compare JSC application extreme-scaling experience with that of leading supercomputing centres. [5]

contact: Brian Wylie,
b.wylie@fz-juelich.de




Figure 1: Strong scaling of workshop application codes compared to existing High-Q Club members.




Figure 2: Weak scaling of workshop application codes compared to existing High-Q Club members.

- Dirk Brömmel
- Wolfgang Frings
- Brian Wylie

Jülich
Supercomputing
Centre (JSC),
Germany

References

- [1] Mohr, B., Frings, W. Jülich Blue Gene/P Porting, Tuning & Scaling Workshop 2008, inside 6(2), 2008
- [2] Mohr, B., Frings, W. (Eds.) Jülich Blue Gene/P Extreme Scaling Workshops 2009,2010,2011, Tech. Reports FZJ-JSC-IB-2010-02, FZJ-JSC-IB-2010-03 & FZJ-JSC-IB-2011-02.
- [3] The High-Q Club <http://www.fz-juelich.de/ias/jsc/high-q-club>
- [4] Brömmel, D., Frings, W., Wylie, B. J. N. 2015 JUQUEEN Blue Gene/Q Extreme Scaling Workshop, Tech. Report FZJ-JSC-IB-2015-01 <http://juser.fz-juelich.de/record/188191>
- [5] ISC-HPC Workshop on Application Extreme-scaling Experience of Leading Supercomputing Centres (Frankfurt, 16 July 2015) <http://www.fz-juelich.de/ias/jsc/aXXLs/>

Retrospective of International Jülich CECAM School on Computational Trends in Solvation and Transport in Liquids

The Jülich CECAM School [1], organized this year at JSC continued a series of Schools which started in 2000 and which takes place in bi- or triennial periods [2]. It has entered the CECAM [3] event calendar since 2012 and addresses a larger audience interested in computational science and parallelization of algorithms in various fields of application. This year the focus of the School was on Computational Trends in Solvation and Transport in Liquids, partially motivated by the Cluster of Excellence RESOLV, coordinated at Ruhr-University Bochum, which also co-sponsored the School financially [4].

The participants, 43 young scientists from 8 countries (Europe and overseas), proved a broad interest in the topic, which covers both Soft Matter Physics and Solvation Chemistry. The well balanced topics of the lectures made it possible to present both topics as coupled together. The 23 lecturers from 7 countries (also Europe and overseas) were highly renowned scientists who presented latest developments of research embedded in a broad overview over the field.



The topics of the School were grouped into several blocks, covering Atomistic Methods, Coarse Grain and Continuum Methods, Hybrid Methods, Mesoscale Fluid Methods as well as Numerical Methods and Hardware. It is the mix of topical blocks, which has proven to be most efficient in addressing the audience, having sometimes a very diverse background but profiting highly of the introductory character of the lectures spiced with advanced topics which gives a perspective of research in the various fields.

An essential part of the School is the provision of Lecture Notes to the audience on the first day of the School, which makes it even more accessible to follow the courses and refresh the knowledge during and after the School. The Lecture Notes have been published as volume 28 of the IAS series, which is available either as hard copy or as PDF file on the web [5]. Although the topics were well structured in blocks in the Lecture Notes, during the School topics were intentionally scattered over the whole week to provide a broad spectrum of topics to the whole audience and to avoid topic selected days, which would have reduced the number of participants during School days. As a practical aspect of the School, a hands-on session for computing on graphical processing units (GPU) was organized, providing a code oriented approach for first steps in GPU computing.

To stimulate discussion between participants and lecturers, 4 evening events were organized, containing a get-together during the first evening, 2 poster sessions and an excursion of the whole group to the brown coal open cast mining in the close neighbor-

hood of Jülich together with a dinner. The positive feedback from participants and lecturers provides strong motivation to further organize Computational Science oriented Schools in a similar scheme.

References

- [1] www.fz-juelich.de/stl-2015
- [2] For a list of past Schools, see: www.fz-juelich.de/ias/jsc/EN/Expertise/Workshops/Conferences/STL-2015/PastWorkshops/_node.html
- [3] www.cecarn.org
- [4] CoE Ruhr Explores Solvation: www.ruhr-uni-bochum.de/solvation/
- [5] IAS Series, Vol. 28: <https://juser.fz-juelich.de/record/188877/files/IAS-Band-28-V2.pdf>

- Godehard Sutmann¹
- Johannes Gritendorst¹
- Gerhard Gompper
- Dominik Marx²

¹ Jülich Supercomputing Centre (JSC), Germany

² Ruhr-University Bonn, Germany

Workshop on Force-Field Development (Forces 2014)

Many processes in nature and technology can be rationalized with computer simulations of a few thousands to a few hundred millions of atoms. Examples range from protein folding via plastic deformation to the dynamics of explosions. Simulating such processes in a meaningful fashion is outside the scope of electronic-structure based techniques and will remain so for the next few decades, even if computers and algorithms keep improving at the current rate. Gaining useful insight into the various processes then often requires one to use low-cost force fields that still properly reflect the intricate quantum mechanics responsible for interatomic bonding and repulsion. For instance, the failure mechanism of a specific material can only be unraveled in simulations that accurately account for defect energetics. This cannot be achieved with generic two-body potentials. To advance the field, a workshop on the development of force fields was held in Jülich from November 3–5, 2014.

Seventy participants from four continents attended the workshop, in which eleven invited talks, seven contributed talks, as well as twenty posters were presented. Topics included, amongst others, the systematic and fitting-free bottom-up design of force fields from first principles, machine-learning strategies, force fields for non-equilibrium (excited electronic states) to potential repositories and force-field standardization. The IOP journal Modeling and Simulation in Materials Science and Engineering (MSMSE) will dedicate a special issue Force Fields: From Atoms to Materials to this very successful workshop. Ten invited contributions to the proceedings volume are currently under review and expected to be published within the next six months. More details on the workshop, including a link to the special MSMSE link in the near future, can be found at <http://www.fz-juelich.de/ias/jsc/ForceFields2014>

- Martin Müser

Jülich
Supercomputing
Centre (JSC),
Germany



Figure 1: Participants of the NIC workshop Force Fields: From Atoms to Materials in front of the Jülich Supercomputing Centre.

JSC becomes full Member of JLESC



The Joint Laboratory brings together researchers from the Institut National de Recherche en Informatique et en Automatique (France), National Center for Supercomputing Applications (USA), Argonne National Laboratory (USA), Barcelona Supercomputing Center (Spain), RIKEN Advanced Institute for Computational Science (Japan) and JSC. The Joint Laboratory brings together researchers from the Institut National de Recherche en Informatique et en Automatique (Inria, France), the National Center for Supercomputing Applications (NCSA, USA), Argonne National Laboratory (ANL, USA), Barcelona Supercomputing Center (BSC, Spain), RIKEN (Japan) and JSC.

The objectives of JLESC are to initiate and facilitate international collaborations on state-of-the-art research related to computational and data focused simulation and analytics at extreme scales. JLESC promotes original ideas, research and publications, as well as open source software, aimed to address the most critical issues in advancing from petascale to extreme scale computing. Research topics include parallel programming models, numerical algorithms, parallel I/O and storage systems, data analytics, heterogeneous computing, resilience and performance analysis. JLESC involves scientists and engineers from many different disciplines as well as from industry

to ensure that the research supported by the laboratory addresses science and engineering's most critical needs and takes advantage of the continuing evolution of computing technologies.

The collaborative work within the joint laboratory is organized in projects between two or more partners. This includes mutual research visits, joint publications and software releases. The results of these projects are discussed during biannual workshops, where also new ideas and collaboration opportunities are presented. The next event will take place in Barcelona from June 29 to July 1, 2015, followed by a two-day summer school on big data. In December 2015, the event will be organized by JSC.

JSC is involved in many new developments related to key extreme-scale applications. This expertise nicely complements JLESC's current work and will bring new problems to the applied mathematics and computer science communities. As this is truly a high-ranking consortium in supercomputing, JSC felt very much honored being invited to join as full partner. JSC is represented in JLESC by the Steering Committee member Prof. Thomas Lippert and the JSC Executive Director Dr. Robert Speck. For more information visit the official JLESC website under <http://publish.illinois.edu/jointlab-esc>

• Robert Speck

Jülich
Supercomputing
Centre (JSC),
Germany

JuPedSim: Framework for simulating and analyzing the Dynamics of Pedestrians




Figure 1: Screenshot of the simulation of a building evacuation with three floors connected with two staircases. Pedestrians are represented with cylinders and their velocities is color coded (red for slow and green for fast (ca. 1.34m/s)). After 4 seconds, there are still 36 pedestrians in the building.

The growing complexity of modern buildings and pedestrians' outdoor facilities make the use of simulation software inescapable. Simulations are used not only in the design phase of new structures but also later in the preparation or monitoring of large-scale events. To obtain reliable results, it is paramount to understand and reproduce the underlying phenomena which rule the dynamics of the pedestrians. There exist numerous software tools, mostly commercial (or with proprietary licenses), for simulating pedestrians. They usually implement a single model and are only of limited use for academic purposes, where the aim is generally a rapid prototyping framework for implementing and testing new concepts or models. There exist open source tools as well. But they are mainly designed for applications and do not offer easy access for model testing or extension. In addition the state of the art in the modeling of pedestri-

ans is characterized by a high dynamic and the zoo of models is still growing. However, it lacks a common basis for model comparison and benchmarking. From the scientific and academic point of view, it is often crucial to understand how a model has been implemented since the mathematical description and the computer implementation sometimes differ. This also raises issues about the validation of the models, especially against empirical data, which is often neglected in many software systems.

The Jülich Pedestrian Simulator, JuPedSim, is an extensible framework for simulating and analyzing pedestrians' motion at a microscopic level. It consists at the moment of three modules which are loosely coupled and can be used independently. JuPedSim implements state of the art models and analysis methods.

The module JPSScore computes the trajectories of the pedestrians given a geometry and an initial configuration. The start configuration includes the desired destinations, speeds, route choices and other demographic parameters about the pedestrians. Three-dimensional geometries are also supported. Two models at the operational level (locomotion system, collision avoidance) and three models at the tactical level (route choice, short term decisions) are actually implemented in the framework [1,2]. Other models are in the process of being integrated. Further models can be incorporated by third parties without much effort. Other behavioral

parameters are implemented such the possibility to share information about closed doors with other agents and the ability to explore an unknown environment looking for an exit.

The second module JPSvis visualizes the geometry and the trajectories, either from files or streamed from a network connection. High resolution videos can also be recorded directly from the module interface.

The module JPSreport analyses the results from the simulation or any other source for instance experiments and generates different type of plots. The reporting tool integrates four different, measurement methods [3]. Possible analyses include densities, velocities, flows and profiles of pedestrians in a given geometry.

Planned features for the framework include a graphical user interface for editing the geometry, which will also include import capabilities for various CAD formats and the connection of the pedestrian simulation with a fire simulator. In contrast to other simulation packages, an emphasis is set on the validation of the implemented models. The empirical data used for the validation come from numerous experiments that have been conducted in different geometries during the past years. All inputs and output files are XML based. JuPedSim is platform independent, released under the LGPL License and written in C++. All information including documentation, source code and experimental data are available at www.jupedsim.org. JuPedSim has been used for instance in a real time evacuation assistant for arenas [1] and will be used to simulate the Berlin underground station Osloer Straße in the ORPHEUS project [5].

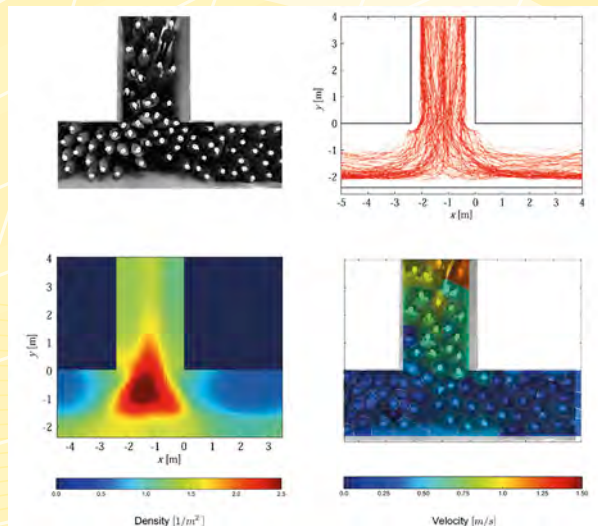


Figure 2: Workflow for analyzing trajectories using JPSreport. Top left: experiment involving two pedestrian streams coming from left and right and merging in a T-Junction. Top right: automatically extracted trajectories using PeTrack [4]. Bottom left: density profiles computed with JPSreport using the Voronoi method [3]. Bottom right: instantaneous velocities of the pedestrians using the Voronoi method [3].

References


- [1] **Kemloh Wagoum, A. U.**
Route choice modelling and runtime optimisation for simulation of building evacuation, Dissertation, Schriften des Forschungszentrums Jülich: IAS Series 17, XVIII, 122 p., 2013.
- [2] **Chraibi, M.**
Validated force-based modeling of pedestrian dynamics, Dissertation, Schriften des Forschungszentrums Jülich: IAS Series 13, 2012.
- [3] **Zhang, J.**
Pedestrian fundamental diagrams: Comparative analysis of experiments in different geometries, Dissertation, Schriften des Forschungszentrums Jülich: IAS Series 14, 103 p., 2012.
- [4] **Boltes, M.**
Automatische Erfassung präziser Trajektorien in Personenströmen hoher Dichte, Dissertation, Schriften des Forschungszentrums Jülich: IAS Series 27, xii, 308 p., 2015.
- [5] **Arnold, L.**
ORPHEUS - Fire Safety in the Underground, inside Spring 2015.

- Armel Ulrich
- Kemloh Wagoum
- Armin Seyfried

Jülich
Supercomputing
Centre (JSC),
Germany

contact: Armel Ulrich Kemloh Wagoum,
u.kemloh.wagoum@fz-juelich.de

HLRS at SC14 in New Orleans, Louisiana



Belonging to the steady circle of participants at the annual Supercomputing Conference (SC), HLRS had returned to SC14 held in New Orleans, Louisiana/ USA (November 16-21, 2014). In addition to HLRS representatives participating in a number of workshops, birds-of-a-feather sessions, tutorials and other SC events, HLRS also hosted a booth in the SC exhibition hall where they presented details about the large number of simulation projects of breakthrough caliber the HLRS supercomputing infrastructure is being used for.

Highlight of the HLRS's exhibits was a hands-on Augmented Reality demo which analyzed simulation results of the airflow and the pressure distribution around a 3D scanned triathlete on her racing bike. By moving a camera around the bicycle or even riding the

bicycle themselves, visitors could immediately observe changes to the airflow resulting from various riding positions of the triathlete. This method helps triathletes to find and verify the most efficient riding position on their racing bikes and to analyze individual mounted accessories and helmets.

Apart from HLRS researchers and scientists inter-exchanging with the like-minded HPC experts, visitors to the HLRS booth were also given the opportunity to have some fun. The HLRS staff had organized a special competition with the opportunity for the participants to win real trophies: The HLRS HPC Awards (High Performance Cycling awards). Those who dared to meet the challenge of the 5 km long bike trail were given the opportunity to enjoy the (simulated) charming scenery of the

beautiful Black Forest. However, they also had to fight the elevation of the German mountain range: up- and down-hill passages had to be mastered while at the same time "racing against the clock" and trying to beat trail records was in order.

What at first seemed like a feasible task proved for some biking fans to be a too big challenge. Not all courageous racers eventually crossed the "finishing line".



• Regina Weigand

University of
Stuttgart, HLRS,
Germany

contact: Uwe Wössner,
woessner@hlrs.de

14th HLRS/hww Workshop on Scalable Global Parallel File Systems



From April 27th to April 29th, 2015, representatives from science and industry working in the field of Global Parallel File Systems and High Performance Storage did meet at HLRS for the fourteenth annual HLRS/hww Workshop on Scalable Global Parallel File Systems "The Non-Volatile Challenge". About 75 participants did follow a total of 22 presentations which have been on the workshop agenda.

Prof. Michael Resch, Director of HLRS, opened the workshop with an opening address on Monday morning.

In the keynote talk, Eric Barton, CTO of Intel's High Performance Data Division, discussed "A new storage paradigm for NVRAM and integrated fabrics". He explained emerging trends and upcoming technologies which might significantly change the storage landscape. Peter Braam, Braam Research and University of Cambridge, explained the exa-scale data requirements and issues of the square kilometre array which is currently under development.

In the first presentation of the Monday afternoon session, Torben Kling-Petersen, Seagate, discussed the Lustre Storage Enterprise HPC technology of Seagate which is enabling energy efficient, extreme performance storage solutions. Following new approaches for large HPC systems, Wilfried Oed, Cray, explained the Cray Data Warp solution which can already be deployed in today's Cray XC30 and XC40 systems. Afterwards, Alexander Menck, NEC, gave an overview of the NEC storage portfolio.

In the second afternoon session, James Coomer, DDN, introduced the IME Burst Buffer technology and he showed, how real world applications can profit from its usage. Franz-Josef Pfreundt, FhG – ITWM, showed BeeOND, which is BeeGFS (formerly known as Fraunhofer File System) on demand. He explained how a file system can be setup and provided automatically on nodes which have been reserved for a user job, e.g. by a batch system. In the last presentation of the day, Mellanox' Oren Duer gave an overview on Mellanox efforts in the storage field including the different technologies they are providing in the field.

In the first presentation on Tuesday morning, Akram Hassan, EMC, provided EMC's view to elastic Cloud and object storage. He showed how the object storage solution fits the needs of today's applications. The following talk have been more research oriented.

Tim Süß, University Mainz presented results of his studies on the potential of data deduplication in checkpoint data of scientific simulations.

The second session started with a presentation of IBM, given by Olaf Weiser, explaining new developments in GPFS, especially GPFS native RAID. For half a year, Lenovo is now a new and active player on the HPC market. Michael Hennecke provided an overview of Lenovos HPC storage solutions. In the last talk of the morning sessions, Thomas Uhl introduced a new highly performant HSM solution provided by GrauData which is especially working together with Lustre.

The second half of the day is traditionally reserved for network oriented presentations. This year a focus have been on the opportunities of Software Defined Networking. Yaron Ekshtein, presented in detail the Pica8 approach for a switch operating system providing an open networking approach for improving cost/performance and scalability of compute clusters. This was followed by Edge Core's Lukasz Lukowski, who showed its bare metal switches as underlying hardware for SDN and how they can accelerate open networking. Klaus Grobe, Adva Optical, went even more down to the hardware and explained solutions for inter-data-center 400-Gb/s WDM transport.

As a preview for the Wednesday session, Radu Tudoran, Huawei, introduced the three storage musketeers of the Big Data era: Reliable, Integrated and Intelligent and Technology-Convergent.



The other talks touching also the big data arena followed on Wednesday morning. Mario Vosschmidt, NetApp, discussed the benefits of declustered RAID technologies and especially covered Hadoop file systems with this. Alexey Cheptsov, HLRS, introduced the Juniper project and how HPC can be used for data-intensive applications with Hadoop and OpenMPI.

The last session has been again more future and research oriented. The Seagate perspective to the future of HPC Storage has been given by Torben Kling-Petersen. Michael Kuhn, Hamburg University, followed with a new approach to make future storage systems aware of I/O semantics and finally Thomas Bönisch, HLRS, showed the results of the project SIOX, Scalable IO Extensions.

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

- Thomas Bönisch
University of
Stuttgart, HLRS,
Germany

HLRS Scientific Tutorials and Workshop Report and Outlook

The **Cray XC40 Optimization and Parallel I/O** courses will provide all necessary information to move applications to the new Cray XC40 Hornet and future Hazel Hen HPC system. We are looking forward to working with our users on this leading-edge supercomputing technology. The next course series in cooperation with Cray specialists will take place on October 20-23, 2015. An online recording from September 2014 is also available.

One of the flagships of our courses is the week on **Iterative Solvers and Parallelization**. Prof. A. Meister teaches basics and details on Krylov Subspace Methods. Lecturers from HLRS give lessons on distributed memory parallelization with the Message Passing Interface (MPI) and shared memory multithreading with OpenMP. This course will be presented twice, on March 16-20, 2015 at HLRS in Stuttgart and on September 07-11, 2015 at LRZ in Garching near Munich.

- Rolf Rabenseifner

University of
Stuttgart, HLRS,
Germany

ISC and SC Tutorials

Georg Hager, Gabriele Jost, Rolf Rabenseifner: **MPI+X - Hybrid Programming on Modern Compute Clusters with Multicore Processors and Accelerators**. Tutorial O2 at the International Supercomputing Conference, ISC'15, Frankfurt, July 12-16, 2015.

Georg Hager, Jan Treibig, Gerhard Wellein: **Node-Level Performance Engineering**. Full-day Tutorial at Super Computing 2014, SC14, New Orleans, Louisiana, November 16-21, 2014.

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

Another highlight is the **Introduction to Computational Fluid Dynamics**. This course was initiated at HLRS by Prof. Sabine Roller. It is again scheduled on February 23-27, 2015 in Siegen and on September 14-18, 2015 in Stuttgart.

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

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

Our general course on parallelization, the **Parallel Programming Workshop**, September 07-11, 2015 at HLRS, will have three parts: The first two days of this course are dedicated to parallelization with the Message Passing Interface (MPI). Shared memory multi-threading is taught on the third day, and in the last two days, advanced topics are discussed. This includes MPI-2 functionality, e.g., parallel file I/O and hybrid MPI+OpenMP, as well as MPI-3.0 (and upcoming 3.1) with its extensions to the one-sided communication, a new shared memory programming model and a new Fortran interface. As in all courses, hands-on sessions (in C and Fortran) will allow users to im-



diately test and understand the parallelization methods. The course language is English.

Three and four day-courses on **MPI & OpenMP** are presented at other locations all over the year.

We also continue our series of **Fortran for Scientific Computing** on March 09–13, 2015 and December 07–11, 2015, mainly visited by PhD students from Stuttgart and other universities to learn not only the basics of programming, but also to get an insight on the principles of developing high-performance applications with Fortran.

With **Unified Parallel C (UPC) and Co-Array Fortran (CAF)** and an **Introduction to GASPI** each year in spring, the participants will get an introduction of partitioned global address space (PGAS) languages.

GPU programming is taught in **OpenACC Programming for Parallel Accelerated Supercomputers – an alternative to CUDA from Cray perspective** in April and in two **CUDA** courses in April and in October 2015.

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 on the right hand side, you can find the whole HLRS series of training courses in 2015. They are organized at HLRS and also at several other HPC institutions: LRZ Garching, JSC (FZ Jülich), ZIH (TU Dresden), and ZIMT (Siegen). GCS serves as a PRACE Advanced Training Centre (PATC). PATC courses are marked in the table.



2015 – Workshop Announcements

Scientific Conferences and Workshops at HLRS

14th HLRS/hww Workshop on Scalable Global Parallel File Systems (April 27–30)

High Performance Computing in Science and Engineering – The 17th Results and Review Workshop of the HPC Center Stuttgart (October 05–06)

IDC International HPC User Forum (Autumn 2015)

Parallel Programming Workshops: Training in Parallel Programming and CFD

Parallel Programming and Parallel Tools (TU Dresden, ZIH, February 16–19)

Industrial Services of HLRS (HLRS, February 25, July 01, and November 11)

VI-HPS Tuning Workshop (HLRS, February 23–27) (PATC)

Introduction to Computational Fluid Dynamics (ZIMT, Uni Siegen, February 23–27)

Cray XC40 and I/O Workshops (HLRS, March 02–05 and October 20–23) (PATC)

Iterative Linear Solvers and Parallelization (HLRS, March 16–20 / LRZ Garching, September 07–11)

Tools for Parallel Programming with OmpSs (HLRS, April 13–15)

OpenACC Programming for Parallel Accelerated Supercomputers (HLRS, April 16–17) (PATC)

GPU Programming using CUDA (HLRS, April 20–22 and October 26–28)

Unified Parallel C (UPC) and Co-Array Fortran (CAF) (HLRS, April 23–24) (PATC)

Efficient Parallel Programming with GASPI (HLRS, April 27)

Scientific Visualization (HLRS, April 28–29 and October 29–30)

Cluster Workshop (HLRS, June 29–30)

Node Level Performance Engineering (HLRS, July 06–07)

Introduction to Computational Fluid Dynamics (HLRS, September 14–18)

Message Passing Interface (MPI) for Beginners (HLRS, September 28–29)

Shared Memory Parallelization with OpenMP (HLRS, September 30)

Advanced Topics in Parallel Programming (HLRS, October 01–02) (PATC)

Advanced Parallel Programming with MPI & OpenMP (FZ Jülich, JSC, November 30–December 02)

Training in Programming Languages at HLRS

Fortran for Scientific Computing (HLRS, Februar 09–13 / December 07–11) (PATC)

URLs

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

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

(PATC): This is a PRACE PATC course

GCS - High Performance Computing

Node-Level Performance Engineering

(PATC course)

Date & Location

July 6-7, 2015

HLSR, Stuttgart

and December 10-11, 2015

LRZ Garching near Munich

Contents

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

- Introduction and Motivation
- Performance Engineering as a process
- Topology and affinity in multicore systems

- Microbenchmarking for architectural exploration
- The Roofline Model
 - Basics and simple applications
 - Case study: sparse matrix-vector multiplication
 - Case study: Jacobi smoother
- Model-guided optimization
 - Blocking optimization for the Jacobi smoother
- Programming for optimal use of parallel resources
 - Single Instruction Multiple Data (SIMD)
 - Cache-coherent Non-Uniform Memory Architecture (ccNUMA)
 - Simultaneous Multi-Threading (SMT)
- Pattern-guided performance engineering
 - Hardware performance metrics
 - Typical performance patterns in scientific computing
 - Examples and best practices
- Beyond Roofline: The ECM Model
- Optional: Energy-efficient code execution

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

Webpage

www.hlsr.de/training/course-list
and

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

Introduction to Parallel Programming with MPI and OpenMP

Date & Location

August 11-14, 2015

JSC, Forschungszentrum Jülich

Contents

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

Webpage

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

Courses and Tutorials

Iterative Linear Solvers and Parallelization

Date & Location

September 7-11, 2015
LRZ, Garching near Munich

Contents

The focus is on iterative and parallel solvers, the parallel programming models MPI and OpenMP, and the parallel middleware PETSc. Different modern Krylov Subspace Methods (CG, GMRES, BiCGSTAB ...) as well as highly efficient preconditioning techniques are presented in the context of real life applications.

Hands-on sessions (in C and Fortran) will allow users to immediately test and understand the basic constructs of iterative solvers, the Message Passing Interface (MPI) and the shared memory directives of OpenMP. This course is organized by LRZ, University of Kassel, HLRS, and IAG.

Webpage

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

Advanced Fortran Topics

Date & Location

September 14-18, 2015
LRZ, Garching near Munich

Contents

This course, partly a PRACE Advanced Training Center (PATC) course, is targeted at scientists who wish to extend their knowledge of Fortran beyond the Fortran 95 standard. Days 1-3 (LRZ course): best practices (global objects & interfaces, abstract interfaces, import statement, object-based programming), object-oriented programming (type extension, polymorphism, inheritance, type- & object bound procedures, generic type-bound procedures, abstract types & deferred bindings), IEEE features & floating point exceptions, Fortran 2003 I/O extensions. Days 4-5 (PATC course): object-oriented design-patterns (polymorphic objects & function arguments, interacting objects, submodules & plugins), interoperability with C (mixed language programming), coarrays (basics, dynamic entities, advanced synchronization, programming patterns, collectives, events, teams, atomic subroutines, performance aspects). To consolidate the lecture material, each day's approximately 4 hours of lecture are complemented by 3 hours of hands-on sessions. The last 2 days of the course are a PATC event.

Prerequisites

- good knowledge of Fortran 95

Webpage

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

Introduction to Computational Fluid Dynamics in High Performance Computing

Date & Location

September 14-18, 2015
HLRS, Stuttgart

Contents

The course deals with current numerical methods for Computational Fluid Dynamics in the context of high performance computing. An emphasis is placed on explicit methods for compressible flows, but classical numerical methods for incompressible Navier-Stokes equations are also covered. A brief introduction to turbulence modelling is also provided by the course. Additional topics are high order numerical methods for the solution of systems of partial differential equations. The last day is dedicated to parallelization.

Hands-on sessions will manifest the contents of the lectures. In most of these sessions, the application Framework APES will be used. They cover grid generation using Seeder, visualization with ParaView and the usage of the parallel CFD solver Ateles on the local HPC system.

The course is organized by HLRS, IAG (University of Stuttgart) and STS, ZIMT (University of Siegen).

Webpage

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

GCS - High Performance Computing

Message Passing Interface (MPI) for Beginners

Date & Location

September 28-29, 2015
HLRS, Stuttgart

Contents

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

Webpage

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

Shared Memory Parallelization with OpenMP

Date & Location

September 30, 2015
HLRS, Stuttgart

Contents

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

Webpage

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

Advanced Topics in Parallel Programming

(PATC course)

Date & Location

October 1-2, 2015
HLRS, Stuttgart

Contents

Topics are MPI-2 parallel file I/O, hybrid mixed model MPI+OpenMP parallelization, MPI-3.0 parallelization of explicit and implicit solvers and of particle based applications, parallel numerics and libraries, and parallelization with PETSc. Hands-on sessions are included. Course language is English (if required).

Webpage

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

Courses and Tutorials

Introduction to GPU programming using OpenACC

Date & Location

October 19-20, 2015
JSC, Forschungszentrum Jülich

Contents

GPU-accelerated computing drives current scientific research. Writing fast numeric algorithms for GPUs offers high application performance by offloading compute-intensive portions of the code to the GPU. The course will cover basic aspects of GPU architectures and programming. Focus is on the usage of the directive-based OpenACC programming model which allows for portable application development. Examples of increasing complexity will be used to demonstrate optimization and tuning of scientific applications.

Webpage

<http://www.fz-juelich.de/ias/jsc/events/openacc>

Cray XC40 Workshop and Parallel I/O Courses

Date & Location

October 20-23, 2015
HLRS, Stuttgart

Contents

This course consist of two parts.

A three-day **optimization course** covers the main topics of porting applications for the Cray XC40 system Hornet, application profiling, application debugging, and parallel I/O. Tools provided on the system are presented in talks and live demos. Furthermore a series of hands-on sessions give the opportunity to exercise also with own applications.

An increasing issue in parallel computing is the efficient handling of reading and writing files. The one day workshop **Efficient parallel I/O** offers users the ability to gain insight on how to optimize the I/O of an application. It shows methods and tools for parallel I/O including MPI-IO, HDF5, NetCDF. Thus the time spent in reading and writing files can be significantly reduced, increasing the total efficiency of an application.

The workshop does include some Cray specific information about tools but is still useful for I/O on other systems as well.

Webpage

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

GPU Programming using CUDA

Date & Location

October 26-28, 2015
HLRS, Stuttgart

Contents

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

Webpage

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

GCS - High Performance Computing

Scientific Visualization

Date & Location

November 5-6, 2015
HLRS, Stuttgart

Contents

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

Webpage

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

Data Analysis and Data Mining with Python

Date & Location

November 9-11, 2015
JSC, Forschungszentrum Jülich

Contents

Pandas, matplotlib, and scikit-learn make Python a powerful tool for data analysis, data mining, and visualization. All of these packages and many more can be combined with IPython to provide an interactive extensible environment. In this course, we will explore matplotlib for visualization, pandas for time series analysis, and scikit-learn for data mining. We will use IPython to show how these and other tools can be used to facilitate interactive data analysis and exploration.

Webpage

<http://www.fz-juelich.de/ias/jsc/events/python-data-mining>

Industrial Services of the National HPC Centre Stuttgart (HLRS)

Date & Location

November 11, 2015
HLRS, Stuttgart

Contents

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

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

Webpage

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

Courses and Tutorials

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

Date & Location

November 26-27, 2015
JSC, Forschungszentrum Jülich

Contents

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

Webpage

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

Advanced Parallel Programming with MPI and OpenMP

Date & Location

November 30 - December 2, 2015
JSC, Forschungszentrum Jülich

Contents

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

Webpage

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

Fortran for Scientific Computing

Date & Location

December 7-11, 2015
HLRS, Stuttgart

Contents

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

Webpage

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

Contents

News

Interview with Prof. Dr. Dr. Thomas Lippert	4
Obituary: Dr. Walter Nadler	7
Minister of Science of Baden-Württemberg visits HLRS	9
Cray CEO visits HLRS	10
Member of Executive Board of Porsche visits HLRS	11
Jülich's Bernd Mohr to Chair SC17	12

Applications

HLRS Supercomputer successfully executed Extreme-scale Simulation Projects	14
Discontinuous Galerkin for High Performance Computational Fluid Dynamics	20
Large-scale Simulations of Particle-laden Flows	25
Quarks and Hadrons – and the Spectrum in Between	29
Advance Visualisation of Seismic Wave Propagation and Speed Model	34
A new Neutrino-Emission Asymmetry in forming Neutron Stars	38
Precision Physics from Simulations of Lattice Quantum Chromodynamics	43
Turbulence in the Planetary Boundary Layer	47
The strong Interaction at Neutron-rich Extremes	51
Lattice QCD as Tool for Discoveries	56
LIKWID Performance Tools	60

Projects

ECO ₂ Clouds – Experimental Awareness of CO ₂ in Federated Cloud Sourcing	64
MIKELANGELO – Micro Kernel Virtualization for High Performance Cloud and HPC Systems	68
FFMK: A Fast and Fault-tolerant Microkernel-based Operating System for Exascale Computing	72
EUDAT2020	77
ORPHEUS — Fire Safety in the Underground	80
Seamless Continuation: 4th PRACE Implementation Phase Project	85
Intel® Parallel Computing Center (IPCC) at LRZ	88

Systems

92

Centres

96

Activities

102

Courses

122