

in Deutschland

Contents

News

HLRS@GCS Hazel Hen Europe's fastest system	4
Leibniz Extreme Scaling Award goes to VERTEX	6
Start of Production on the JURECA System	8
Europes biggest Training Facilities for High Performance Computing	10
New Team for Coordination of Computing Time Allocation located at JSC	/ 11
20 Years of High Performance Computing for Science and Economy	12
"Industry Relations" Team paves the Way to HPC Expertise	15
Reallabor Stadt:quartiere 4.0	16

Applications

Aeroacoustic Simulation of complete Helicopter Simulation using CFD	18
Large Scale numerical Simulations of planetary Interiors	23
Theory of dynamical Processes in Semiconductor Nanostructures via	
large Scale ab initio Calculations	28
Simulating Quarks and Gluons in physical Conditions	32
High Resolution gravity Field Modeling	37
Power Data Aggregation Monitor	40
4D City – Space-time urban Infrastructure Mapping by Multi-sensor	
Fusion and Visualization	46
Simulation of Interactions at the LHC	50

54

56

58

62

64

70

75

76

78

82

Projects BEAM-ME Mont Blanc 3 Dynamic Scheduling of reconfigurable scientific Workflows – The DreamCloud approach to Energy-efficient HPC Infrastructures ParaPhase: Towards HPC for adaptive Phase-Field Modeling DEEP Project successfully completed Scientific Big Data Analytics by HPC Improved Safety for disabled People New Czech-Bavarian Competence Centre for Supercomputing Applications established TerraNeo - Integrated Co-Design of an Exascale´earth Mantle Modeling Framework

Activities			88
Courses			107



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

Publishers

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

Editor

F. Rainer Klank, HLRS

klank@hlrs.de

Design

Pia Heusel, HLRS

heusel@hlrs.de

Mantle convection in the interior of the terrestrial planets of the Solar System.

A.-C. Plesa¹, C. Hüttig¹, M. Maurice¹, N. Tosi^{1,2}, D. Breuer¹

¹Department of Planetary Physics, German Aerospace Center (DLR), Berlin, Germany, ²Department of Astronomy and Astrophysics, Technische Universität Berlin, 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 present Europe's fastest PRACE system Hazelhen that was installed at HLRS in September 2016. The system already won an award for its industrial usage at SC'15. At the same time we present a new generalpurpose system at JSC called JURECA. So, GCS is moving forward with its successful national HPC strategy supported by both the federal government and the state governments of Bavaria, Baden-Württemberg and Northrhein-Westfalia.

The growing capabilities of GCS systems are reflected in a growing number of applications that make use of our systems achiev-ing breakthroughs in their fields. This issue again covers a wide variety of activities and again gives proof that high performance computing is an important if not vital tool for a large number of applications.

Again we present a large number of research projects and activities in high performance computing. We present a selection of HPC related projects that are funded at the European, the national and the state level. This is proof of how deep and well GCS is connected with the research communities that embrace and further drive high performance computing. Key aspects like Ex-ascale Computing, Big Data or Cloud Computing and more are covered by the projects conducted at each of the three GCS cen-ters. This shows how much the three national German high performance computing centers

have turned from operational sites into research centers of their own.

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

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

HLRS@GCS Hazel Hen Europe's fastest system



In August 2015, HLRS installed its currently figures as the fastest second phase of a large scale Cray XC40 system. It was code named acceptance in September 2015. During October HLRS ran a numthe system and routinely added the HPL benchmark and the HPCG benchmark to the tests. The results were surprising for the HLRS team.

In the TOP500 list Hazel Hen was able to take the 8th place slightly ahead of the Shaheen II system of the King Abdullah University of Science. In this friendly competition with the colleagues in Saudi Arabia Hazel Hen seemed to have found a sweet spot for the HPL benchmark the fastest system in Europe. since first results had not been very convincing. Ranked number When designing the concept for 8, Hazel Hen was ahead of all

open European system accessible via the Pan-European PRACE con-Hazel Hen and went through cept. The only European system faster than Hazel Hen is the Swiss Piz Daint system which is based on ber of large scale simulations on the intensive use of accelerators - a concept which did not show convincing price/performance for the real world applications of HLRS and wastherefore ruled out during the negotiation for the second phase of installation.

> Even more surprising was the fact that for the HPCG benchmark Hazel Hen showed an impressive performance taking 6th position worldwide. For this more realistic benchmark Hazel Hen is currently

Hazel Hen, HLRS was focused on other PRACE systems and hence the sustained performance of user

codes. Peak performance was not relevant and the ranking in any benchmarks was also not an issue. HLRS was all the more surprised to see that sustained application performance can result in a high ranking in benchmark lists. With the ongoing discussions about the relevance of existing benchmarks HLRS found that HPCG reflects much better the performance level that can be expected from a user's code than HPL. This does not come as a surprise. CG methods are widely used in a variety of engineering codes that are at the forefront among the HLRS users. It is also obvious that a single benchmark will never be able to give a true picture of the performance an HPC system is capable of. Especially when it comes to the use of accelerators application performance can vary extremely when using different algorithms for different applications.

From the point of view of a general purpose usage center Hazel Hen has proven to be an excellent choice. Large scale simulations conducted during the acceptance phase and the first few weeks of operation showed excellent results for a variety of applications. Even industrial users benefit from the balanced design of the system and one care manufacturer was able to do more than thousand crash simulations within 24 hours on Hazelhen - earning HLRS an HPC-Wire Award for "Best Use of HPC in Automotive" at supercomputing 2016.

contact: Michael M. Resch, resch@hlrs.de



- Michael M. Resch
- University of Stuttgart, HLRS, Germany

Leibniz Extreme Scaling Award goes to VERTEX



Figure 1: Prof. Bode and the four finalists from VERTEX, wallberla, PSC and INDEXA

It was the fourth "Extreme Scaling Workshop" organized by the Leibniz Supercomputing Centre (LRZ) to optimize the most demanding programs running on LRZ's supercomputer SuperMUC. As 2016 is the "Year of Leibniz", Arndt Bode, Chairman of the LRZ, handed the new "Leibniz Extreme Scaling Award" to the VERTEX team, that made the most remarkable progress in optimising its code.

Gottfried Wilhelm Leibniz, who died on November 14th, 1716, was the first to construct a mechanical computing machine that was not only able to add, subtract and multiply, but also to divide floating point numbers. SuperMUC is able to perform 3,000,000,000,000,000 of these floating point operations in a single second delivering 3 Petaflops of peak performance. And LRZ has two of them, the SuperMUC Phase 1 from 2012 and the new SuperMUC Phase 2 that was put to operation in 2015.

VERTEX needs more than several Exallop, 10 to the 18th floating point operations, during one single run of an hour. Nevertheless, it was the fastest program during the Extreme Scaling Workshop 2016 running on SuperMUC. 13 teams from international projects gathered in Garching near Munich to optimize their programs for SuperMUC together with the experts from LRZ and the hardware and software vendors. This achievement was honored with the newly created Leibniz Extreme Scaling Award. VERTEX and the other teams succeeded in optimizing their codes to extreme scaling so that they now can exploit SuperMUC's huge compute power to the extreme with optimal efficiency.

Applications from **Divers Research Fields:** Supernovae, Weather Forecast, Medicine

VERTEX is used to simulate supernovae, those gigantic explosions ending the existence of massive stars that can shine brighter even than galaxies for weeks thereby producing neutron stars, black holes and all the atoms from which the earth and men are composed. And there should also be generated measurable amounts of neutrinos and signals of gravitational waves. Of course, the scientists behind VERTEX from the Max-Planck-Institute for Astrophysics are on the hunt for these waves that have been detected for the first time recently.

The other programs carry names like HemeLB, waLBerla, LISO, MPAS,

SWIFT, GADGET or GHOST. No ghosts but very real humans are the scientists writing them. They originate from many different countries all over Europe and the world. The projects are hosted at Munich and Erlangen, Darmstadt, Greifswald, Dresden, London, Durham and Valencia. Also the areas of application are very divers.

MPAS uses a great mesh of worldwide weather and climate prediction and fine mesh simulations to create a model to forecast the local weather even of single valleys in the alps with excellent precision.

HemeLB and waLBerla are Lattice Boltzmann Codes. While sounding abstract it may soon be saving human lifes. The blood flow is simulated using these programs, especially in aneurysms, those very dangerous balloon like bulges in the walls of blood vessels, that might rupture and lead to life threatening bleeding. Physicians have to decide whether it is better to remove the aneurysm or to leave it. This is a very delicate decision and simulations help to make them on a better basis. Today, they need a supercomputer for these simulations. But not for too long these computations will be possible at the hospital and the images can be inspected immediately.

Extreme Scaling Workshop: Function Check for SuperMUC

Since 2011 LRZ organises the Extreme Scaling Workshops to optimize the programs that already run well on its supercomputer so that they run with extreme scaling. The objective is to run realistic, huge use cases with as many processor cores as possible as efficiently as possible.

And it is worth the effort. In an earlier workshop the geophysical simulation SeisSol from Ludwig-Maximilians-Universität in Munich could achieve excellent performance and made it to one of the five finalists of the "Gordon Bell Award" at the Supercomputing Conference 2014.

The workshop also is a thorough function check. During the workshop SuperMUC is available exclusively for the participants. The programs are deeply understood and problems can be investigated with precision not possible in normal operation to detect any faults in hardware, software, communication or file systems. This helps to improve normal operation with many different applications running at the same time.

And the effort is worth the money. Improving the efficiency of the simulations saves budget. The optimized versions of the huge programs using the most of SuperMUC at once need only one fourth of time and energy after the optimization. Thus, taking into account the enormous costs of operating such a big computer there will well be saved some Million cpu cycles that may now be spent on other research projects. This is true especially for Bavaria that pays half of SuperMUC's costs and supports most of the projects from Bavarian universities participating in the workshop through the competence network for scientific supercomputing in Bavaria KONWIHR.

contact: Ludger Palm, ludger.palm@lrz.de • Ludger Palm

Start of Production on the JURECA System

On November 2, 2015 the JURECA general-purpose supercomputer started production at full scale at Jülich Supercomputing Centre in Forschungszentrum Jülich. JURECA, which stands for "Jülich Research on Exascale Cluster Architectures", succeeds the popular JUROPA cluster which served computing time to a wide variety of European research and industry projects from 2009 to 2015. As its predecessor, the system has been co-designed by JSC together with the project partners, in this case T-Platforms and ParTec, as a versatile and balanced system for compute- and data-intense scientific applications. JURECA's peak performance of 1.8 Petaflop/s is six times higher than that of JUROPA while energy consumption has dropped

the system components and configuration see, [1].

In order to minimize the downtime for the users during the system installation as much as possible, the buildup of JURECA had been organized in two phases. This was necessary since a parallel operation of JUROPA and the complete 34 rack JURECA system was not feasible due to a lack of available floor space. The first phase, consisting of 160 compute nodes in six racks, replaced the JUROPA system; this first phase offered the same peak performance at a strongly reduced energy and floor footprint. In this way large parts of the full JURECA system could be set up and tested while production continued. In October a downtime was



Figure 1: Side view of the JURECA system at Jülich Supercomputing Centre in Forschungszentrum Jülich.

by approximately one third. With the addition of 150 NVIDIA K80 GPUs in Mid-December 2015, the peak performance was further boosted to 2.2 Petaflop/s. For technical details about tion tests.

required to finalize the top-level Infini-Band cabling, while integrating the six racks from the first phase during the process, and to perform final stabiliza-

With a High-Performance Linpack (HPL) performance of 1.42 Petaflop/s on 1,764 compute nodes JURECA entered on position 50 in the TOP500 November 2015 list and placed among the five fastest systems in Germany. On the High-Performance Conjugate Gradient (HPCG) benchmark, JURECA achieved an impressive 68.3 Teraflop/s corresponding to place 18 on the HPCG November 2015 list.

From the start, JURECA has shown a very satisfactory stability despite the cutting edge software and hardware technology employed in the system, such as the brand-new Mellanox Extended Data Rate (EDR) InfiniBand network. During the setup and the early operation of JURECA the lessons learned from the JUROPA project have proven invaluable. Building on this foundation, JSC, T-Platforms and ParTec, will collaboratively improve the system setup and software stack to establish a similarly solid foundation for nextgeneration cluster systems.

JURECA has been warmly welcomed by users. Between autumn 2014 and summer 2015, JSC provided access to a Haswell test cluster that allowed users to adapt their codes and workflows early on to the software stack on JURECA. This opportunity has been widely used, as could be seen by the nearly optimal usage of the system shortly after the start of production on the first phase.

JURECA is available for all eligible scientists. Computing time is granted by the John von Neumann Institute for Computing (NIC) as well as through the Jülich Aachen Research Alliance (JARA-HPC/VSR).

References

2016

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

Jülich Supercomputing Centre, JURECA: General-purpose supercomputer at Jülich Supercomputing Centre. Journal of large-scale research facilities, 2, A62. http://dx.doi.org/10.17815/jlsrf-2-121,

• Dorian Krause

Europes biggest Training **Facilities for High Performance** Computing



• Norbert Conrad

ing, the topping-out ceremony for the new Education and Training Center at the University of Stuttgart (HLRS) could be celebrated on October 23, 2015.

With the new Education and Training Center the HLRS made space for the integration of research, development, production and teaching as well as further training of its users all around the world. "The HLRS is the biggest training center for High Performance Computing with around 800 participants", according to Prof. Michael Resch, Director of the High Performance Computer

tion with a gross floor area of 2.003 square meters creates a new training room of 254 square meters with cutting-edge IT equipment in a fully integrated surrounding. The training participants have interactive access to the training system directly next to the High Performance Computer systems. The upstream forum offers the possibility for relaxed conversations of the participants and allows for an efficient service "on-site". The new center is supposed to come on stream in October 2016. The cost of construction of almost 6 million Euro will be borne by the University of Stuttgart.

New Team for Coordination of Computing Time Allocation located at JSC

Two new members have joined the team coordinating computing time allocation PRACE Preparatory Access Calls. at Jülich Supercomputing Centre.

The new head of the Coordination Office for the Allocation of Computing Time is Dr. Florian Janetzko. Florian Janetzko studied Chemistry at the University of Hannover. He received his PhD in Theoretical Chemistry in 2003. After post-doctoral stays at CINVESTAV (Mexico City, Mexico), University of Calgary (Canada) and University of Bonn (Germany), he joined JSC's Application Support Division in 2008 as a member of the Cross-Sectional Team Application Optimization. He has been involved, among other things, in support and teaching activities as well as in the European Projects DEISA and PRACE [1] as the person locally responsible for the DECI calls.

Dr. Alexander Schnurpfeil is the new Scientific Secretary of the John von Neumann Institute for Computing (NIC) [2]. He is the successor of Dr. Manfred Kremer, who retired from this position end of March 2016. Alexander Schnurpfeil studied Chemistry at the Universities of Leipzig, Bonn and Siegen and received his PhD in Theoretical Chemistry from the University of Cologne in 2006. After a post-doctoral stay at the University of Cologne, he worked as a software developer and teacher. He joined JSC in 2008, also in the Application Support division, where he was, among other duties, responsible for user support for the HPC-FF



Dr. Florian Janetzko is the new Head of the Coordination Office for the Allocation of Computing Time.

References

contacts: Florian Janetzko, Alexander Schnurpfeil, nic@fz-juelich.de

cluster and was local coordinator of the



The new Scientific Secretary of the John von Neumann Institute for Computing, Dr. Alexander Schnurpfeil.

[1] http://www.prace-ri.eu/

[2] http://www.john-von-neumann-institut. de/nic/EN/Home/home_node.html

• Alexander • Florian Janetzko

Centre (JSC),

20 Years of High Performance Computing for Science and Economy



Thus, an anniversary colloquium was held in the Motorworld in Böblingen on December 1. This location was not chosen randomly as it was this location where the closeness to the automobile was substantiated in the form of the most beautiful vintage cars and most rakish sports cars, which hww has shaped from the beginning. Hence, the combination of maximum performance and a glance into the past, which is perfect for the hww-anniversary. The close proximity of the Daimler Plant in Sindelfingen tops this choice off.

The High Performance Computer for Science and Economy GmbH (hww) runs so reliable and experienced that a round birthday was almost missed: In July 1995, hww was founded as a supercomputer-joint-venture of the public shareholder University of Stuttgart and the county Baden-Württemberg, and on the other hand of Porsche AG and debis Systemhaus GmbH, a subsidiary of Daimler Benz Inter-Services AG.

After more than 20 years, the social structure has changed a little: instead of debis, T-Systems is now involved and on the public side the Karlsruhe Institute of Technology came along. Of course there were also a lot of changes on the technical side. This was reason enough to celebrate the anniversary with a small ceremony and to invite all old and new companions. The current Managing Directors of hww, Dr. Alfred Geiger and Dr. Andreas Wierse welcomed the distinguished group of guests, among them the current representatives of the shareholders but also numerous companions, who in the meantime are enjoying their well-deserved retirement. At this point the founding Director and Leader of the Computer Center Prof. Roland Rühle was commemorated. He significantly shaped the foundation of the hww and sadly died two years ago.

With the first lecture, the second founding director Mr. Michael Heib, formerly debis, could lead the guests back to the very beginnings. Together with Mr. Wolfgang Peters, who accompanied the foundation on the side of the Ministry of Economics, offered an overview over the computer landscape and told some high-spirited anecdotes from this time. Even back then this Joint-Venture could only succeed because of numerous important supporters.

At this point let us make a comparison between the computer landscape from the past and the present flagship of the hww, the Cray XC40 Hazel Hen. The computing power of the old Cray T94 from debis was illustrated in an article of the debis company magazine: "If a person could multiply 2 14-digit numbers per second, he would need almost 55 years in order to achieve what one single processor of T94 can do in 1 second", Professor Resch, Director of the HLRS recently described the computing power of the current Cray thus: "What Hazel Hen can calculate in one second would take all humans, so about 7 billion inhabitants, 55 years given an eight hour working day." To say this in other numbers: The computing

power of a hww super computer has

multiplied millionfold from 5 gigaflops in the year 1995 to 7 Petaflops in 2015.

A similar tendency can be found with power consumption, thankfully the factors are not quite as high: Here,an increase of power of approximately 64 KW in the past faces 3 Megawatt today. Worth mentioning is also that the Cray T94 needed 60 % less (!) than it's previous system with a simultaneous increase of computing power.

Prof. Norbert Kroll from DLR highlighted the perspective of the user in the second colloquium. In his presentation "Perspectives of High Performance Computing in aircraft development", he impressively showed how this tremendous computing power can be used nowadays. Today, the flow simulation is a highly developed technical tool and you can't imagine airplane development without it anymore;



News



• Andreas Wierse

the combination of wind tunnel and flight tests with the flow simulation is what in many cases allows a good understanding of the aerodynamic processes; the integration of structuremechanical calculations allows for a view of the complete airplane and the interaction of deformation and aerodynamic.

Prof. Kroll also took this chance to point out the problems which are still ahead of us, which challenges we still need to handle. The number of simulations (more than 2,000 computational grids per year, more than 1,000 simulations per week with the airbus!) as well as the size of the result data increases permanently, transient calculations are much greater than stationary ones. The aim of a complete virtual airplane is despite all progress still not achieved concerning the hardware as well as in the software. With help of a timeline, he located the multi-disciplinary optimization towards the end of the decade. the CFD-based inflight flow simulation in real time was found in the year 2030. If you looked at this point of the timeline at the available computing capacity, you could see: 1 Zeta-flops!!

It is out of the question that the topic of High Performance Computing will be at the top of the agenda of airplane development in the foreseeable future, a combination of research and industry, or better yet of science and economy will be exceptionally important in the foreseeable future. At the subsequent dinner memories of the past as well as expectations of the future were extensively discussed. On one thing all anniversary guests agreed: Even if nobody bet that hww would celebrate its 20th birthday when it was established, the prospect of another 20 years of cooperation of science and economy with High Performance Computing in Baden-Württemberg is extremely good.

"Industry Relations" Team paves the Way to HPC **Expertise**

Since January 2016, JSC has provided a new central contact point for enterprises that wish to use HPC cluster systems for their R&D activities, manufacturing, or their products. The work of the "Industry Relations" Team is inspired by various existing cooperations between JSC and industrial companies and by a generally increasing interest in and need for parallel computing in various fields of engineering and technology.

JSC therefore offers its expertise in HPC to potential customers from industry in the form of "simulation consulting", which ranges from software analysis and optimization techniques targeting small systems up to elaborate means of running highly parallel code on compute clusters with thousands of cores. In this context, the Simulation Laboratories of JSC and JARA-HPC deliver core competences in the application and HPC field to establish cooperations on different levels, varying from short-term core consulting services up to large-scale project partnerships with intensive code treatment and performance enhancement. Last but not least, the Industry Relations team enables industry partners and SMEs to access Jülich know-how and experts in the rapidly emerging field of big data analytics.

This essential new interface provided by JSC emerged naturally from existing services and capabilities and aims to

facilitate the transfer of specific knowledge in the field of parallel computing to a wider industrial community and to further areas of applied research and development. It offers a tremendous opportunity for cooperations that can leverage the vast potential HPC is able to deliver to IT-based R&D. This opportunity is even extended by the work of the Simulation Laboratories of JSC and JARA-HPC that allows to boost projects by using cutting edge research knowledge.

Further details on the new central contact point can be found at http:// www.fz-juelich.de/ias/jsc/industryrelations.

contact: Hartmut Fischer, ha.fischer@fz-juelich.de

Industry Relations

• Hartmut Fischer

Germany

Reallabor Stadt:quartiere 4.0



The University of Stuttgart has received project budget over Euro 1,195,000 for the research project "Reallabor Stadt:quartiere 4.0". The assessment of support was given within the context of the announcement "Wissenschaft für Nachhaltigkeit: Reallabor Stadt" of MWK. The festive ceremony took place on December 2, 2015 at HLRS. In the focus of the research project are participatory and digital planning procedures. City districts should turn into "Real Labors of Change" through this.

The digitalization of economy and society is one of the most comprising revolutions of our time. Apart from the industrial transformation of an industry 4.0, the change of living environments and the working worlds has to be shaped significantly in the future. At the same time, the impact of digitalization on cities as our central

habitat is hardly known and there is no public understanding due to the complexity and the interlinking of topics. Planning and decision-making processes are often conducted with very conventional and highly regulated methods and tools which can hardly keep pace with the current challenges. However, the fields of innovation of digitalization and planning and decisionmaking processes are supposed to be viewed in combination with more and more complex demands under consideration of all relevant players as well as the central location for living in urban space.

Objectives

We examine which digital participation formats, visual tools and simulations can be used in order to optimize planning processes as well as to include citizens early and constructively in the city-district-planning.

We work out in which phase which digital procedures, formats and simulations with which participation are suitable for the application of planning processes on a district level.

We formulate a planning guide for cities and communities. This will give us an indication of what digital instruments and what participation formats are suitable for which phases of the planning process.



Approach

With this project, which is set up for three years, the Research Association relies on a strategy of co-design of research and of the common knowledge acquisition with citizens. With the help of a series of real-world experiments and inventions, innovative and future-oriented procedures for city district planning are supposed to be developed with early shaped citizen participation and tested in cooperation with the cities Stuttgart and Heidelberg.

Project Partners

Institute for Human Factors and Technology Management (IAT), High Performance Computing Center Stuttgart (HLRS), Institute of Urban Planning and Design (SI), Research Center for Interdisciplinary Risk and Innovation Studies (ZIRIUS) at the University of Stuttgart, Fraunhofer Institute for Labour Economics and Organisation (IAO)

- Wolfgang Schotte
- University of Stuttgart, HLRS, Germany

Aeroacoustic Simulation of complete Helicopter Simulation using CFD



Figure 1: λ_p -visualization of the flow field around the helicopter in a BVI relevant descent flight

In the past years, the aeroacoustic noise emission of helicopters became one of the most important, but also challenging issues in helicopter development. The Blade Vortex Interaction (BVI) phenomenon is one of the dominant phenomena characterizing the helicopter's aeroacoustic footprint. Everybody who witnessed a helicopter's landing approach will remember a strong impulsive noise emission for quite a long time during the whole maneuver. In particular, people living in the entry lane of a heliport or a hospital suffer from increased noise exposure. Often, helicopters are deployed in search and rescue, and security (such as police) missions, which require flight

routes above populated areas, and are in conflict with night time banning.

As of technical progress, the mission profile of helicopters is continuously extending. But the society's noise sensitivity increased within the recent past as well. The acoustic admission requirements for all means of transportation including helicopters are continuously tightened. On this account, it is compulsory to consider helicopter's noise emission already in early design phases.

BVI noise results from the interaction of the rotor blades with the blade tip vortices of the preceding blades. As shown in Figure 1, in forward descent flight the blade tip vortices move through the rotor disk and are hit by succeeding rotor blades, which leads to strong pressure fluctuations on their surface. Inside the cabin, these fluctuations can be felt like a ride over cobblestone. Outside the cabin, the fluctuations can be heard as strong impulsive noise.

Addressing the aeroacoustic behavior of rotor systems in early design phase in a conventional-experimenta-manner is connected with high costs and lot of time. For example, every variation of the rotor blade geometry causes a rebuilt of all rotor blades and a new cost and time intensive wind tunnel tests. In this way, wide-range aeroacoustic examination of rotor systems during design time is impossible.

For efficient analysis and optimization of the aeroacoustic and aerodynamic behavior of helicopters, a virtual wind tunnel is required, which allows detailed, reliable, and very accurate predictions of the flow field in all flight situations. Particularly, aeroacoustic phenomena are often insufficiently predicted by low fidelity computational methods. For high fidelity noise prediction of a helicopter configuration, a multidisciplinary CFD-CSD-CAA tool chain has been established at the Institute of Aerodynamics and Gas Dynamics of the University of Stuttgart (IAG).

For detailed simulation of the aerodynamic, the block structured finite volume Reynolds-averaged Navier-Stokes (RANS) CFD code FLOWer, originally developed by the German Aerospace Center (DLR) [1], is used. It can handle rotating coordinate systems, overlapping meshes by using the CHIMERA method, and aero-elastic components by applying deformations derived from CSD-tools. The accurate flight state including the deflection of the rotor blades due to the acting aerodynamic forces is determined by the flight mechanic tool CAMRAD II [2], which provides the information required for deforming the volume meshes within FLOWer. To achieve a realistic (force and moment free) flight state of the helicopter, the flight mechanics of the helicopter is taken into account. Therefore, a free-flight helicopter trim is performed including six degrees of freedom. Beside the three main rotor controls, collective and cyclic pitch angles, the fuselage pitch and roll orientation as well as the tail rotor collective are taken into account.

Limited computational resources and too high numerical dissipation of the vortex structures long prevented the successful prediction of BVI using CFD methods. Within the past years, several important helicopter related features have been implemented into FLOWer by IAG [3,4]. The CFD solver was extended with different methods of fifth order spatial WENO schemes [5] to facilitate a detailed conservation of the flow field and especially vortices.

For tool chain validation a full helicopter configuration in a BVI flight case was simulated and compared to experimental flight test data. For resolving the aeroacoustic phenomenology, a very high discretization density of the CFD volume meshes is required. As a CFD simulation of the whole space between helicopter and the aeroacoustic observers (e.g. microphones on the ground), which preserves acoustic

pressure fluctuations, would go far beyond the scope of today's available computational power, only the area of aerodynamic relevance is represented by volume meshes.

For transportation of the acoustics to distant observers, the CAA-Tool



Figure 2: Radiation of BVI assignable noise emission of the helicopter with a slice 4.5 rotor radii under the helicopter



Figure 3: Difference of BVI assignable noise 4.5 radii under the helicopter between full helicopter and rotor only evaluation (Grey areas excluded due to low overall noise level)

ACCO [6] developed by the IAG is used, which uses the Ffowcs-Williams-Hawkings equations for acoustic modelling. Applying ACCO on converged time resolved 3D CFD results allows aeroacoustic examination of the flow field at any arbitrary location.

The final CFD solution was computed on a setup consisting of 59 separate volume meshes with ca. 200 million cells using a higher order WENO scheme of fifth order. (This setup is comparable to a conventional second order simulation with 1.6 billion cells.) The fifth order WENO scheme represents the optimum with regard to the efficiency. Schemes with even higher order require more computation power to achieve the same physical results. For this simulation, focus was set on vortex convection and mapping of vortex-structure interactions using the higher order methods. Since the method works best for Cartesian meshes, the body meshes are embedded in a Cartesian background mesh with refinements in the vicinity of the helicopter's surface. The usage of hanging grid nodes enables a coarsening towards the far field to ensure a reasonable grid size. With this approach 50% of the overall cells are located in the Cartesian background mesh to guarantee best higher order results.

Besides aeroelasticity of the rotor blades, the mass flow through the engines is considered by prescribed pressures at the inlet and exhaust. The exhaust flux is furthermore prescribed with the average exhaust temperature of the specific helicopter type. Especially the mass flux influences the fuselage wake significantly in terms of occurrence and extent of separation areas.

Already in 2014 a main rotor only simulation (under neglect of the fuselage, the tailboom, and the tail rotor system) by Kowarsch et. al [7] showed good agreement to the experimental flight test data regarding BVI noise emissions. But some minor differences regarding the direction of radiation were found. The presented simulation of a

is found upstream. Tracing back its radiation to the helicopter, the source area is found at the advancing blade side of the main rotor. A further BVI associated noise emission downstream is found with lower strength and expan-



Figure 4: Comparison of PNLT time history of flight test and simulation results

complete helicopter configuration performed on the HLRS Hazel Hen system and improves the results to excellent agreement to the flight test data. The high grid resolution in the near field area of the helicopter in combination with the higher order WENO method yields a very detailed resolution of the vortex structure. As shown in Figure 1, the blade tip vortices are preserved compact during their entire convection downstream, which is a necessary condition for an accurate representation of the blade vortex interactions. In the wake of the skids a typical Kármánlike vortex street is visible produced by the cylindrically shaped skids.

Figure 2 shows the BVI related noise footprint 4.5 radii under the helicopter. The strongest BVI noise emission with its radiation marked by the isosurface mission decreases [8].

sion. The source for this area can be identified in the retreating blade side area of the main rotor. The nonlinear distortions of the isosurface arise from reflection and diffraction caused by the fuselage. The difference of the BVI associated noise emission between the evaluations of the main rotor solely and the complete setup including fuselage and tail rotor is shown in Figure 3. The effects of reflection and shading are clearly visible. At the advancing blade side the noise immission increases whereas at the retreating side the im-

Figure 4 shows the PNLT values for the certification relevant approach, which represent the noise level how it is perceived by humans. During the whole flyover, the achieved results match the experimental data extremely well.

Acknowledgements

The authors would like to thank Airbus Helicopters Deutschland GmbH for the esteemed cooperation, as well as providing us experimental data to enable this investigation. We would like to express our thanks to the Federal Ministry of Economics and Technology for providing us the resources to realize this project. Furthermore, the investigation is based on the longstanding cooperation with the German Aerospace Center (DLR) making us their CFD code FLOWer available for advancements and research purpose, which we would like to thank for.

Further acknowledgement is made to the High Performance Computing Center in Stuttgart who provided us with support and service to perform the computations on their High Performance Computing system Hazel Hen.

• Patrick Kranzinger

- Ulrich Kowarsch
- Manuel Keßler
- Ewald Krämer

Aerodynamik und Gasdynamik, University of Stuttgart, Germany

References

- [1] Kroll, N., Eisfeld, B. and Bleeke, H.M., "The Navier-Stokes code FLOWer," Notes on Numerical Fluid Mechanics, pp. 58-71, 1999.
- [2] Johnson, W.,

CAMRAD II Comprehensive analytical model of rotorcraft aerodynamics and dynamics, edition 4.8 ed., 2009.

[3] Kranzinger, P., Kowarsch, U., Schuff,

M., Keßler, M. and Krämer, E., "Advances in Parallelization and Highfidelity Simulation of Helicopter Phenomena," in High Performance Computing in

Science and Engineering '15, Heidelberg, Springer Verlag, 2015.

[4] Schuff, M., Kranzinger, P., Keßler, M. and Krämer, E.,

"Advanced CFD-CSD coupling - Generalized, high performant, radial basis function based volume mesh deformation algorithm for structured, unstructured, and overlapping meshes," in 40th European Rotorcraft Forum, Southampton, GB, 2014.

[5] Kowarsch, U., Keßler, M. and Krämer, E.,

"High Order CFD-Simulation of the Rotor-Fuselage Interaction," in 39th European Rotorcraft Forum, Moscow, 2013.

[6] Keßler, M. and Wagner, S.,

"Source-Time Dominant Aeroacoustics," in Computer & Fluids, Vol. 33, 2004, pp. 791-800.

[7] Kowarsch, U., Lippert, D., Schneider, S., Keßler, M. and Krämer, E., "Aeroacoustic Simulation of an EC145T2

Rotor in Descent Flight," in 71st American Helicopter Society Annual Forum, Virginia Beach, VA, 2015.

[8] Kowarsch, U., Öhrle, C., Keßler, M. and Krämer, E.,

"Aeroacoustic Simulation of a Complete H145 Helicopter in Descent Flight," in 41st European Rotocraft Forum, Munich, 2015.

contact: Patrick Kranzinger, kranzinger@iag.uni-stuttgart.de

Golden Spike Award by the HLRS Steering Committee in 2015

Large Scale numerical Simulations of Planetary Interiors

Mantle convection describes the slow creep of rocky materials caused by temperature-induced density variations and compositional heterogeneities inside planetary bodies. This process is ultimately responsible for the heat transport from the deep interior and for the large-scale dynamics inside the Earth and other terrestrial planets thus influencing surface geological

structures like volcanoes, rifts and tectonic plates as well as the magnetic field generation.

Different styles of mantle convection are reflected in the surface regimes of the terrestrial bodies of our Solar System. Thus, the surface of the Earth, covered by seven major plates, is characterized by the generation of



Figure 1: Mantle convection in the interior of the terrestrial planets of the Solar System

new oceanic crust and lithosphere occurring at spreading centers and its subsequent subduction into the mantle at convergent plate margins. On all other terrestrial bodies of the Solar System, the so-called one-plate bodies like Mercury, Venus, the Moon or Mars the outermost layers comprise an immobile lid – the stagnant lid – across which heat is transported by conduction, while convection takes place beneath [e.g., 1]. With the tremendous increase of computational power over the past decades, numerical models of planetary interiors have become state-of-the-art tools for understanding the processes that shape the interior of terrestrial bodies across our Solar System and beyond. By considering sophisticated initial conditions, accounting for complex phase transitions and employing a non-linear rheology, these models aim to reconstruct the entire thermo-



Cores



Over the past decades, numerous space missions to the terrestrial planets of the Solar System have returned large amounts of data revealing spectacular surface features and offering important constraints on the evolution of their interior. However, space mission data offer only a present-day snapshot of the thermo-chemical history and, even though the lack of plate tectonics on bodies like Mercury, the Moon or Mars has maintained surface traces of events throughout the planetary history, old surfaces are sometimes buried by younger features and a time reconstruction of the past events is difficult.

chemical history of a terrestrial body and match observations from space mission data.

Numerical Method

To this end, in our project, MAntle THErmo-chemical COnvection Simulations (MATHECO) we perform large scale numerical simulations of mantle convection using the finite-volume code Gaia [2,3]. The code is a fluid flow solver which employes a fixed mesh in arbitrary geometries to solve the conservation of mass, momentum and energy. Gaia uses various methods to decompose the mesh into equal volume domains which are then distributed on the computational cores. An efficient load balancing ensures a good performance and guarantees code scalability. While for 3D Cartesian box grids, due to their regular nature, an optimal domain decomposition can be easily achieved, the problem becomes highly complex when 3D spherical shells are involved. In Gaia we use the so-called Thomson-points to laterally decompose the sphere by distributing points, curate track compositional variations within the mantle, we employ tracer particles. Massless particles move according to the velocity field computed on the fixed mesh. They transport material properties which are used to update various compositional fields at grid resolution by interpolating values stored on particles back to the center of the grid cells. The method has the advantage over classical grid-based

Figure 3: Thermal convection at high Rayleigh numbers. The ISO-surface shows thermal instabilities which develop in a system heated both from below and from within at high Rayleigh numbers $Ra \approx 10^{11}$ including a temperature dependent viscosity with a viscosity contrast of 10^6 . The simulation uses a computational grid with 54 million computational points (275 million unknowns) and was performed on 54×10^3 computational cores at the HLRS supercomputing center.

all assumed with equivalent potential energy, on the surface of the sphere and minimizing the global potential field energy [4]. In Figure 2 we show the code performance obtained with up to 54 x 10^3 computational cores and a typical domain decomposition for a 3D spherical shell grid with boundary refinement at the top surface.

The Gaia code is used for modeling Stokes-flow with strongly varying viscosity, a scenario often met in geophysics and in particularly in mantle convection simulations. Further, beside thermally driven convection, compositional changes may be the source of additional buoyancy. To acmethods of being essentially free of numerical diffusion and of enabling the natural advection of an arbitrary large number of different properties with limited computational effort [5,6]. The ability to use such technique in various geometries makes Gaia unique among mantle convection codes. The code is written in C++ and does not need additional libraries, thus being easily portable on various systems.

Results

Numerical simulations of planetary interiors often use simplified approaches and limited parameter ranges in order to reduce computational costs. The convection vigor expressed in terms of



the Rayleigh number is usually reduced and often a simplified mantle rheology is used. With our models we investigate the behavior of systems using high Rayleigh numbers (vigorous convection) as appropriate for the Earth's mantle and complex rheologies, i.e. viscoplastic rheology [e.g.,7], to model the subduction of tectonic plates consistent with the Earth's mobile surface (Figure 3).

We also take into account complex starting conditions by considering the existence of a liquid magma ocean during the early stages of planetary evolution [8]. The subsequent cooling of such magma ocean can result in a unstable density configuration which may induce a global mantle overturn (Figure 4). The style of the overturn as well as the subsequent configuration of the density profile are ultimately important for the subsequent thermochemical evolution of a planetary body and may have a major impact on the

later surface tectonics and volcanic history [9,10].

Conclusions and Outlook

Our simulations, using highly resolved meshes with up to 54 million points necessary for the parameter ranges studied, have been performed on the new Cray XC40 (Hornet) machine at the HLRS supercomputing center in Stuttgart using up to 54,000 computational cores. During our project, we have largely improved our code, making Gaia one of the few codes world-wide able to efficiently run large scale simulations of interior dynamics. High-performance computational centers like HLRS with its massive parallel computing resources make possible to address complex fluid dynamical problems of planetary interiors and help improve our understanding of physical processes active in the interior of terrestrial bodies using reallistic parameter values.



Figure 4: Overturn of an unstable density gradient. The ISO-volume shows the dense material sinking and accumuating at the coremantle boundary. The simulation uses a 12 million points computational grid (60 million unknowns) and additional 240 million tracer particles to accurately track compositional gradients. It was performed using 13 x 10³ computational cores on the HLRS platform.

Numerical simulations of planetary interiors will continuously improve to account for more realistic conditions. On the one hand, the large amount of data available from various space missions offers a valuable mean to validate the models and on the other hand the highperformance computational resources will allow to consider e.g., complex mineralogical phase transitions, more realistic rheological behavior and sophisticated starting conditions.

Acknowledgements

This work has been supported by the Helmholtz Association through the research alliance "Planetary Evolution and Life" and through the grant VG-NG-1017, by the Deutsche Forschungsgemeinschaft (grant TO 704/1-1), and by the Interuniversity Attraction Poles Programme initiated by the Belgian Science Policy Office through the Planet Topers alliance. Computational time has been provided by the High-Performance Computing Center Stuttgart (HLRS) through the project Mantle Thermal and Compositional Simulations (MATHECO).

References [1] Schubert, G., Turcotte, D.L.,

Olson, P., (2001)

- Plesa, A.-C.,

[4] Thomson, J.J., Cambridge Philosophical Magazine Series 6 7, 237 (1904)

- 4(4), 8302, (2003)
- Sciences, 2012
- [8] Solomatov, V.S.,
- 118, 1-17 (2013)

[10] Plesa, A.-C., Tosi, N., Breuer, D., Earth and Planetary Science Letters 403, 225-235 (2014)

contact: Ana-Catalina Plesa ana.plesa@dlr.de

Cambridge University Press, Cambridge

Hüttig, C., Tosi, N., Moore, W.B.,

Physics of the Earth and Planetary Interiors 220, 11-18 (2013)

INFOCOMP 2011: The First International Conference on Advanced Communications and Computation, 2011

Tackley, P.J., King, S.D.,

Geochemistry, Geophysics, Geosystems,

Plesa, A.-C., Tosi, N. and Hüttig, C.,

Integrated Information and Computing Systems for Natural, Spatial and Social

[7] Ismail-Zadeh, A., Tackley, P.J.,

Cambridge University Press, (2010)

Origin of the Earth and Moon. University of Arizona Press, Tucson (2000)

[9] Tosi, N., Plesa, A.-C., Breuer, D.,

Journal of Geophysical Research Planets

- A.-C. Plesa¹
- C. Hüttig¹
- M. Maurice¹
- N. Tosi^{1,2}
- D. Breuer¹

¹Department of Planetary Physics, German Aerospace Germany

Astronomy and Astrophysics, Germany

Theory of dynamical Processes in Semiconductor Nanostructures via large Scale ab initio Calculations



Figure 1: Atomistic models of InAs-InP core-shell nanocluster with diameter of 3 nm.

Motivation

To unravel the effects of temperature (lattice vibrations) on the electronic and optical properties of semiconductor nanostructures known as quantum dots, we perform large scale abinitio density functional theory (DFT) calculations on realis- tic structures with diameters rang- ing from 2 to 3

nm, which contain up to one thousand atoms and several thousand electrons. The scientific interest resides not only on the understanding of fundamental physics ("how does matter behave at the nanometer scale?") but also on a reliable assessment of the process of carrier relaxation – how do excited electrons relax to their ground state, and how fast – which is most relevant for semiconductor nanodevices. Specifically for the design and control of nanomaterials in the fields of optoelectronics, spintronics, photovoltaic, biolabeling, and the next generation of displays.

The task is computationally most chal-

Computations

percomputer Hazel Hen at the HLRS using well parallelized ab initio codes.

Results

The electronic wave functions of the highest occupied molecular orbital (HOMO) and the lowest unoccupied molecular orbital (LUMO) states of semiconductor nanocluster, which



Figure 2: Electronic wave functions of (a) HOMO and (b) LUMO states for a silicon nanocluster with a diameter of 2.4 nm.

lenging. On one hand the chemistry is complicated due to, e.g., surface effects, which requires a computationally expensive quantum mechanical ab-initio treatment. On the other hand the structures, although small in size, are large in the number of atoms and number of electrons that need to be considered. A typical nanostructure studied in this project (an InAs-InP core-shell nanocluster with diameter of 3 nm), which is composed of more than 1,000 atoms, is presented in Figure 1. In this project, all of the DFT calculations are performed on the suare used to calculate the magnitude of the electron-phonon (e-ph) coupling [1,2], are plotted in Figure 2 (a) and (b), respectively. With the calculated e-ph coupling, a size dependent carrier relaxation time of semiconductor nanoclusters induced by lattice vibrations is obtained (see Figure 3 for e-ph coupling, electronic energy level spacing, and carrier lifetime as a function of size for an InAs quantum dot [2].) We also perform large-scale ab-initio DFT calculations to study the lattice strain and the vibrational properties of colloidal semiconductor core-shell nanoclus-



Figure 3: (a) Electron-phonon coupling strength, (b) electronic energy level spacing, and (c) carrier lifetime of InAs nanoclusters as a function of nanocluster size.

ters. Our results lead to a different interpretation of the frequency shifts of recent Raman experiments, while they confirm the speculated interface nature of the low-frequency shoulder of the high frequency Raman peak [3] (see Figure 4 for calculated and measured CdSe-CdS nanoclusters).

This project contributed to several important ingredients in the fundamental understanding of the effects of size confinement (the effect of the reduced

dimensionality) on the material properties, with a special emphasis on the involvement of lattice vibrations. One of the key results was the understanding of the process how excited electrons relax to their ground states by successively emitting quantum of vibrations (called vibrons). We showed, for instance, that this relaxation process is highly dependent on the size of the quantum dots and the material. The quantitative quality of the results allows us to give precise numbers and not only trends. We showed that within certain ranges of radii (depending on material) the relaxation is inhibited (called a "phonon bottleneck") while in others the relaxation is very fast, in the femtosecond range. The results achieved in this project helped to establish a physical picture of ultrafast electronic processes in colloidal semiconductor nanostructures, which benefits further developments and applications of colloidal semiconductor optoelectronic devices.

Reference

[1] Han, P. and Bester, G.,

First-principle calculation of the electronphonon interaction in semiconductor nanoclusters, Phys. Rev. B 85, 235422 (2012)

[2] Han, P. and Bester, G.,

Carrier relaxation in colloidal nanocrystals: Bridging large electronic energy gaps by low-energy vibrations, Phys. Rev. B 91, 085305 (2015)

[3] Han, P. and Bester, G.,

Heavy strain conditions in colloidal core-shell quantum dots and their consequences on the vibrational properties from ab initio calculations, Phys. Rev. B 92, 125438 (2015)

contact: Gabriel Bester, gabriel.bester@uni-hamburg.de



Figure 4: (a) Calculated Raman spectrum for CdSe-CdS core-shell r CdSe-CdS core-shell nanocluster with R= 3 nm

Applications

Figure 4: (a) Calculated Raman spectrum for CdSe-CdS core-shell nanocluster with R=1.5 nm. (b) Measured Raman spectrum of

• Gabriel Bester,

University of Hamburg and The Hamburg Centre for Ultrafast Imaging, Germany

Simulating Quarks and Gluons in physical Conditions



Figure 1: The speedup of our code on JUQUEEN relative to 128 nodes with problem size fixed (strong scaling). The dashed line represents linear scaling.

The standard model (SM) of elementary particle physics is extremely successful with the latest triumph of the Higgs boson discovery at the LHC. The known interactions in the SM are the electromagnetic, weak and strong forces. While the electromagnetic and weak forces can be analyzed soundly by perturbation theory, because the interaction strength is weak, this is not possible for the strong force at large distances where the constituent particles, the quarks and gluons, interact strongly.

However, the strong force between quarks and gluons is responsible for the form of matter surrounding us, namely protons and neutrons: only

a few percent of the mass of proton and neutron stems from the guark masses (gluons are massless), the predominant part comes from the strong interaction itself. In fact, all nuclear matter is governed by the strong force and hence its understanding is of crucial importance to understand nuclear reactions. In addition, at the early stage of the universe quarks and gluons were not bound in hadrons but interacted within a plasma of particles. To understand the transition of this plasma to our observed universe today is of fundamental importance to reveal the origin of our world and hence of our sheer existence. Also, important nuclear processes like the triple-alpha process (nuclear fision of three Helium nuclei) are ultimately determined by the strong force. Finally, there are many physical processes detected in world-wide experiments where the strong force plays a crucial role and a theoretical understanding of the contributions from the strong force is essential to interpret the experiments. Below, we will provide one example of such a process with the muon anomalous magnetic moment.

On the theoretical side, we have a very elegant model, called quantum chromodynamics, which is supposed to describe the interaction between quarks and gluons. Still, since this interaction is of non-perturbative nature, in many cases there exists up to now no analytical tool to solve the underlying equations. The way out is to formulate the theory on a Euclidean 4-dimensional space-time grid. Then it becomes pos-

sible to implement this lattice quantum chromodynamics (LQCD) on a computer and solve the relevant equations numerically.

Clearly, to this end, one needs to introduce a non-zero lattice spacing and the 4-dimensional box has to be finite. Furthermore, the employed algorithms scale with some power of the inverse quark mass and it turns out that the physical value of the quark mass is so small that - at least in the past numerical simulations directly at the physical quark mass were prohibitively expensive. Thus, LQCD calculations have to include extrapolations to zero lattice spacing, infinite volume and also to the physical value of the quark mass - a clearly demanding undertaking.

However, the last years have seen a dramatic improvement in algorithms. In addition, the supercomputer architectures have developed tremendously, the JUQUEEN (JSC), Hazel Hen (HLRS) and SuperMuc (LRZ) machines at the Gauss center being prime examples. In addition, the LQCD application is highly appropriate for massively parallel supercomputer architectures. In Figure 1 we show a strong scaling plot for our particular code. We show the speedup of the iterative solver, the so-called conjugate gradient solver, used in our software relative to 128 nodes of JUQUEEN for fixed problem size. The most important driver routines are improved using XLC intrinsics for QPX and low level communication routines leading to more than 30% of peak performance. We observe super-linear scaling up to 1,024 nodes, followed by some flattening due to too small local problem size. The dashed line represents linear scaling.

These developments led to a situation unthinkable a few years ago: simulations directly at the physical value of the guark mass are feasible. Our collaboration uses a particular form of LQCD called twisted mass fermions. With this formulation, which has the great advantage that physical quantities scale with a rate proportional to the square of the lattice spacing and therefore faster towards the continuum limit, we were indeed able to reach the physical value of the quark mass and perform simulations directly in physical conditions.

Such calculations are still very demanding and we have, therefore, built a team of physicists from universities and institutes in Cyprus, France, Germany, Italy and Switzerland working together in the European Twisted

M_N	
M_{D_s}	F
f_{D_s}	
f_K	
f_K/f_π	
M_K/f_K	
f_{D_s}/f_{π}	
f_{D_s}/f_K	

respectively.



 $Q_{\rm lat}/Q_{\rm phys}$

Figure 2: Ratios of basic physical results Q_{iat} computed by ETMC using LQCD and phenomenological results Q_{nhue} published by the particle data group. The error bars represent the statistical and systematical LQCD error added in quadrature. The second, outer error bar for the dimensionful quantities M_s, and f_{ν} comes from the estimate of the lattice spacing. The red bars indicate the errors on the phenomenological quantities separately. The dotted and dashed lines indicate per-mille and per-cent deviations from 1.0.

Mass Collaboration (ETMC) which has taken up the challenge to carry out the demanding calculations directly at the physical point.

And, our collaboration has been successful. In Figure 2 we compare a number of physical quantities computed by our collaboration to those published by the particle data group (PDG). The PDG collects the most up to date experimental values for physical quantities within the framework of high energy physics. Figure 2 demonstrates clearly that our LQCD calculations reproduce the physical quantities measured in nature which is a very reassuring result for the appropriateness of the LQCD approach to understand the strong force between quarks and gluons.

The quantities shown in Figure 2 are rather basic physical observables to be computed in LQCD. But, the agreement of LQCD calculations with experiment opens now the door for more complicated and demanding quantities. A number of such quantities, in particular hadronic form factors have already been computed and results are published in renowned international journals.

Here, we want to give one other example which is related to the muon anomalous magnetic moment. The muon is a so-called lepton which has very similar properties as the electron, for example the charge and the spin, but it is about 200 times heavier. The spin of this muon is 1/2 on the classical level, see left panel of Figure 3. However, an elementary particle such as the muon is described by a quantum field theory which allows spontaneous generations of particle anti-particle pairs out of the vacuum. If we con-

sider, for example, the generation of electron-positron pairs (the positron being the anti-particle of the electron) inside the muon, we will find a deformation of the muon and, therefore, a deviation from the classical value of the spin from 1/2. This is illustrated by the right panel of Figure 3.



Figure 3: The spin of the muon. Left: on the classical level, the spin of the muon is 1/2. Right: when quantum corrections appear (here the generation of electron positron pairs), the muon is deformed and a deviation of the spin from the value 1/2 can be computed theoretically to a very high accuracy. This deviation has been indeed detected experimentally to a precision that matches the theoretical calculation.

What makes this observation very attractive is the fact that the spin of the muon or the electron can be measured extremely precisely with 7 (in case of the muon) to 10 (in case of the electron) significant digits. This

precision opens a window to detect possible particle anti-particle pair creations of particles that are not within the standard model and come from some new, so far completely unknown new physics. Since, theoretically, the chance to detect such new particles is proportional to the mass square of the underlying lepton and since the mass of the muon is about 200 times larger



than the one of the electron, the muon is an ideal laboratory to search for new physics beyond the standard model.

It is clear, however, that in order to observe deviations from the SM, the SM prediction itself - here for the true spin of the muon - must match the experimental precision with all contributions from the SM forces well controlled. And, it is here where LQCD comes into the game since it is only through ab-initio, non-perturbative lattice calculations that the non-perturbative strong force contribution can be computed. Thus, experiments eagerly await results from LQCD for the nonperturbative strong force contribution to the spin of the muon.

However, the fact that in the past simulations were performed at unphysically large values of the quark masses introduced a hard to control systematic error from the extrapolation to the physical value of the quark mass. Although improved extrapolation techniques have been designed by our collaboration, the extrapolations to the physical point led to large errors for the lattice results. Therefore, the simulations directly at the physical quark mass value carried out in this project constitutes great progress, avoiding the difficulty to control the extrapolation in the quark mass.

agreement.

As a reassuring result, we could reproduce the values of the strong force contribution to the muon spin. This is usually measured as a deviation from the value 1/2 and is called the muon anomalous magnetic moment, amuon. The value of amuon would, hence, be zero at the classical level and a non-vanishing value indicates the size of quantum corrections. For the strong force contribution we find a____=5.52(39)10⁻⁸ from our calculation directly at the physical point and a____=5.72(16)10⁻⁸ from our earlier work extrapolated in the quark mass. The error of the latter is still smaller since there more independent data points were used to compute the value. Our results can be compared to a dispersive analysis of experimental data which gives a_{max}^{exp} = 5.66(5)10⁻⁸ and we find a good

We consider the here presented work as a significant step forward to understand the strong force through LQCD calculations and, thus, to unravel the mysteries of nuclear matter and the early universe. However, many more steps are to be done. So far, the calculations are at only one value of the [2] Abdel-Rehim, A. et al., lattice spacing and one volume. Moreover, other, heavier quarks have to be taken into account as active degrees of freedom. Finally, the effect of the slight difference between the masses of the up and down quarks and electromagnetism need to be addressed. The ETM collaboration is working on these issues and is pursuing a long-term and broad research programme towards the goal to solve quantum chromodynamics - at least numerically.

Acknowledgements

The authors gratefully acknowledge the Gauss Centre for Supercomputing (GCS) for providing computing time for a GCS Large-Scale Project on the GCS share of the supercomputer JUQUEEN [5] at Jülich Supercomputing Centre (JSC). We would like to thank all members of the European Twisted Mass Collaboration (ETMC) in which this work is embedded for a most enjoyable and fruitful collaboration. Without this common effort it would have not been possible to obtain the interesting and important results described in this report.

Karl Jansen²

- Carsten Urbach¹
- für Strahlen- und Kernphysik Bethe Center for
- Neumann Germany

References

Abdel-Rehim, A., Alexandrou, C., Constantinou, M., Hadjiyiannakou, K., Jansen, K., Kallidonis, C., Koutsou, G. and Aviles-Casco, A. V., Direct evaluation of the guark content of

the nucleon from lattice QCD at the physical point, arXiv:1601.01624 [hep-lat]

[ETM Collaboration], Simulating QCD at the Physical Point with N f=2 Wilson Twisted Mass Fermions at Maximal Twist, arXiv:1507.05068 [hep-lat]

[3] Abdel-Rehim, A. et al.,

Nucleon and pion structure with lattice QCD simulations at physical value of the pion mass, Phys. Rev. D 92 (2015) 11, 114513 [Phys. Rev. D 93 (2016) 3, 039904] doi:10.1103/ PhysRevD.92.114513, 10.1103/Phys-RevD.93.039904 [arXiv:1507.04936 [hep-lat]]

[4] Jansen, K. and Urbach, C.,

tmLQCD: A program suite to simulate Wilson Twisted Mass lattice QCD, Comput. Phys. Commun. 180 (2009) 2717 doi:10.1016/j.cpc.2009.05.016 [arXiv:0905.3331 [hep-lat]]

Jülich Supercomputing Centre. (2015). JUQUEEN: IBM Blue Gene/Q Supercomputer System at the Jülich Supercomputing Centre. Journal of largescale research facilities, 1, A1. http://dx.doi.org/10.17815/jlsrf-1-18

contacts: Karl Jansen, karl.jansen@desy.de, Carsten Urbach, urbach@hiskp.uni-bonn.de

High Resolution Gravity Field Modeling



Figure 1: The gravity field of the Earth in terms of gravity anomalies [mgal].

The static gravity field of the Earth is one of the key parameters for the observation and measurement of a number of processes and flows in the dynamic system of the living planet Earth. Its knowledge is of importance for various scientific disciplines, such as geodesy, geophysics and oceanography. For geophysics the gravity field gives insight into the Earth's interior, while by defining the physical shape of the Earth it provides an important reference surface for oceanographic applications, such as the determination of sea level rise or modelling of oceanographic currents.

Moreover this reference surface is a key parameter on the way to a globally unified height system.

The scientific goal is to estimate the static gravity field as precise and detailed as possible. As the gravity field in general is represented by a spherical harmonic series, the parameters to be estimated in gravity field modelling are spherical harmonic coefficients. There exist various techniques to observe the gravity field, which have different advantages and complement each other. The observation of the Earth gravity field

from dedicated satellite missions delivers high accurate and globally homogenous gravity field information for the long to medium wavelengths of the spherical harmonic spectrum (corresponding to spatial resolutions down to roughly 100 km). However, due to the large distance between the satellite and the Earth's surface, the gravity field signal is damped in satellite height. Therefore short wavelengths of the spherical harmonic spectrum (smaller than 100 km) cannot be observed from space. To complement the satellite information, terrestrial gravity field measurements over land and satellite altimeter observations over the oceans, which need to be converted to gravity field quantities, are used as additional data. As these observations are taken at the Earth surface (land and ocean) they contain the full undamped signal. The scientific challenge is to combine the different types of gravity field observations in the way that all data types keep their specific strengths and are not degraded by the combination with other information in specific spherical harmonic wavelength regimes. As mentioned, this procedure shall result in a set of spherical harmonic coefficients representing the global Earth gravity field up to highest possible resolution.

Strict full normal equation approach and the need for supercomputing

Our approach to determine the spherical harmonic coefficients is based on a strict least squares adjustment with a Gauß-Markov model. By that, the different data types can be combined optimally on normal equation base. The approach enables the ideal relative weighting between different data

sets as well as the individual weighting of every single observation, what is important as especially the quality of terrestrial measurement data differs significantly (observations are in general guite accurate in regions such as Europe or North America, whereas they are worse in Africa or South America). Due to the high correlation of the unknown spherical harmonic coefficients, the corresponding normal equation system is a dense matrix. As the number of unknowns increases quadratic with the spherical harmonic degree, and again the number of elements of the full normal equation matrix is quadratic to the number of unknowns, full normal equation systems become quite large. This is a computational challenge, which can only be solved with supercomputers such as SuperMUC. With the available data we currently estimate the coefficients of spherical harmonic expansion up to degree and order 720, which corresponds to more than 500,000 unknowns and a normal equation system of 2 TByte. With availability of denser ground data coverage further extensions to even higher resolutions are planned for the near future.

Results

Below the benefit of the full normal equation approach is demonstrated exemplary for a gravity field solution in the area of South America. As observations data sets from the US-German satellite mission GRACE (Gravity and Climate Recovery Experiment), the satellite mission GOCE (Gravity Field and Steady State Ocean Circulation Explorer) of the European Space Agency (ESA), and terrestrial observations provided by the National Geospatial Intelligence Agency (NGA) are available. The quality of the terrestrial data differs significantly, e.g. it is quite high for example for coastal areas in Brazil and low for areas in the Amazon area. The major challenge is to find the optimal combination of the different data sets.

as that proofs, that the good performance of the satellite information was not degraded by the combination with terrestrial data. It is clearly visible, that the differences are quite high for the



Figure 2: Accuracy map of terrestrial observations in South America [mgal] (left); difference between the block-diagonal (middle) and full normal (right) equation solution to the satellite information [mgal].

For that purpose, we derive first an accuracy map of the terrestrial observations by comparison with the satellite information in the low to medium wavelengths (Figure 2 left). In the following, we calculate two different gravity field solutions. One is based on a reduced block-diagonal normal equation approach, which requires strong data constraints like equal data weights, but can be executed on a single node computer system. The second approach is our full normal equation approach applying individual weights according to a predetermined accuracy map and using supercomputing facilities. One can investigate the difference between the two results and the original satellite information in the low to medium wavelengths, which delivers a measure of how good the estimation procedure is able to recover the field in this frequency range. The parameter estimation process can be regarded as good, when the differences are small,

Details about the procedure and results can be found in [1].

References

Fecher, T., Globale kombinierte Schwerefeldmodellierung auf Basis voller Normalgleichungssysteme; Dissertation, Ingenieurfakultät Bau Geo Umwelt (BGU), TU München, 2015

contact: Thomas Fecher,

reduced block diagonal approach (Figure 2 middle). Bad data in the Amazon area and in the Andes affected the combination solution. With full normal equations and supercomputing approach, nearly no differences are visible (Figure 2 right). The combination is ideal for that case, what demonstrates the power of using full normal equation systems and supercomputing faciltiies for this purpose.

fecher@bv.tum.de

- Th. Fecher
- T. Gruber
- R. Pail

München, Germany

Power Data Aggregation Monitor

The power consumption of modern High Performance Computing (HPC) systems is rapidly increasing. For example, Tianhe-2 (which translates as Milky Way-2), the supercomputer developed by China's National University of Defense Technology and currently the fastest high-end system in the world according to TOP500 November 2015 rankings [http://top500.org/], delivers around 33.86 PetaFLOPS (Floating Point Operations per Second) performance at the cost of 17.8 MW power (and that excluding the 6.2 MW cooling overhead [1]) which is already a sufficient amount of power for sustaining a small city. These high power consumption rates do not only convert into high operational costs (which already start to be a dominant factor for the Total Cost of Ownership of the target high-end system [2]), but they also transform into high carbon emissions and serve as a limiting factor for hosting new generation supercomputers. The latter issue can already be



Figure 1: Overview of the 4 pillar framework

seen with Exascale computing systems (i.e. systems representing the next thirtyfold increase in computing capabilities beyond currently existing Multi-Petascale systems) where the goal, set by the US department of Energy, is to have a maximum of 20 MW power consumption [3]. These issues once again show that the overall reduction of power/energy consumption is an important milestone towards a sustainable, reliable, and cost-efficient large scale modern data center.

In order to comprehensively asses the current energy efficiency status of the target HPC data center as well as to understand the success rate of various optimization solutions applied for the reduction of power/energy usage, a mechanism is needed that will measure, collect, and correlate the power/ energy consumption data from all the relevant aspects of the data center. Figure 1 shows the 4 pillar framework [4] which represents all energy relevant aspects of a data center from building infrastructure (pillar 1) to the running application (pillar 4). A mechanism for covering a wide range of power consumption analysis capabilities (global optimization strategy) is needed - from the complete data center view (e.g. ability to report on metrics such as PUE, ERE, WUE, etc. [http://www.thegreengrid.org/]] over the individual systems (e.g. ability to report on system power consumption, utilization ratio, etc.) and deployed heat reuse technologies (e.g. assessing the COP of deployed adsorption chiller, i.e. a machine that removes heat from

a liquid via a vapor-compression or refrigeration cycle) to the user applications (e.g. ability to report on aggregated energy consumption for a given HPC workload by correlating the information from the pillars 2 and 3).

There are several monitoring and management toolsets that can be applied for each of these mentioned components (i.e. building infrastructure, deployed high-end systems, etc.). Although diverse in their nature, they are all scope-based, e.g. the current data center cooling infrastructure management and control systems do not consider the actual behavior of the target supercomputing systems residing in the data center and do not provide a control-side coupling between building infrastructure and the present high-end systems. Therefore the cooling infrastructure is typically designed

for the operation at maximum power draws (10 MW at Leibniz Supercomputing Centre [http://www.lrz.de/]) of deployed HPC systems.

Figure 2 shows the annual power consumption profile of Leibniz Supercomputing Centre for 2014. The topmost curve, labeled as P(DC) and colored in violet, shows the complete power consumption of the data center; the second curve (counted from the top), labeled as P(SuperMUC_IT) and colored in blue, shows the power consumption profile of the deployed SuperMUC supercomputer (currently the 23rd fastest supercomputer in the world according to TOP500 November 2015 rankings) [https://www.lrz.de/ services/compute/supermuc/]; and the bottom curve, labeled as P(DC_ Cooling) and colored in red, shows the power consumption of the cooling infrastructure.





Figure 3: PowerDAM overview

and analysis.

Figure 2: Power consumption profiles of LRZ in 2014

As can be seen the overall power profile of the data center, i.e. P(DC), is affected by the power consumption variation of SuperMUC (mainly due to the energy efficiency features of the new generation CPU's where the power consumption of processors varies strongly with the workload [5]) as well as by the power profile of the cooling infrastructure, once again showing the importance of in depth holistic data collection, correlation,

Power Data Aggregation Monitor (PowerDAM) [6,7] was specifically designed and developed at LRZ for addressing this need. It is a unified energy measuring and monitoring toolset that is aimed towards collecting and correlating energy/power consumption relevant data from all the aspects of the HPC data center. Currently Power-DAM covers pillar 1, 2, and 3 of the 4 pillar framework illustrated in Figure 2.

Figure 3 provides an overview of the developed toolset. As can be seen, PowerDAM is using an agent-based data communication model for retrieving the actual sensor data from the remote entities/systems. The agents are maintained on the remote entity side, and are configured to have a permission for using entity/system vendor specific commands to access the point of interest data and push this data back over the network to PowerDAM. PowerDAM on its turn maintains plug-ins for communicating with these agents. This approach



Figure 4: PowerDAM detailed EtS report for a user on a given HPC system

makes PowerDAM loosely coupled with system agents and easies the extension process of monitored entity set.

PowerDAM provides a separate plugin interface for defining customized data anaylsis reports. Figure 4 shows one of these options – the detailed EtS report of all submitted and executed applications of a given user on a given HPC system [7] - in this case the MPP CoolMUC Linux cluster (shown in Figure 4 and briefly described below) abbreviated as mpp1. The y-axis represents the EtS values of the user submitted applications in kWh. The xaxis represents a table where the first row presents the system scheduler assigned application ids, and the second, third, and fourth rows present the computation, cooling, and networking consumption percentages respectively. For example, the application with the scheduler assigned 18,339 id (upper "zoomed-in" table-fragment) has consumed 18.9 kWh energy, of which 78.8 % is attributed to computation, 11.4 % to networking, and 9.7 % to cooling.

CoolMUC, also referred to as Massively Parallel Processing (MPP) Linux cluster system, shown in Figure 5, is a direct warm water cooled AMD processor based Linux cluster built by MEGWARE [http://www.megware. com/]. It has 178 compute nodes (8 nodes of which are dedicated to interactive sessions) within 3 compute racks. Each of the compute nodes is equipped with two AMD Opteron 6,128 HE 8 core processors.

The classification maps of the dynamic change of compute node sensor data through a color mapping is another reporting option provided by PowerDAM. These maps update automatically after a specified amount of time (by default the update frequency is set to the sensor data collection frequency). Figure 6 illustrates the example classification of average CPU temperature heat map data for all compute nodes of CoolMUC. The three columns in this figure show the three compute racks of the CoolMUC prototype system, and each cell represents a compute node. Each cell in Figure 5 contains the compute node name (e.g. lxa27), the power value (e.g. 224 W), and the utilization/ load rate (e.g. 99.13%). outside temperature. As can be seen, at warmer outside temperatures the adsorption chiller generates less cold water power as expected.

Figure 8 illustrates the Power Usage Effectiveness (PUE) [6] of SuperMUC Phase 1 [8] system. SuperMUC is the 23rd fastest supercomputer in the world according to TOP500 November 2015 rankings. It is a Gauss Center for Supercomputing (GCS) [http://www. gauss-centre.eu] infrastructure system and one of the Partnership for Advanced Computing in Europe (PRACE)



Figure 5: Overview of CoolMUC Linux cluster

Figure 7 illustrates the correlation between Coefficient of Performance (i.e. the COP defined as the fraction of the cold water power that goes out of the adsorption chiller to the hot water power that goes in to the adsorption chiller) of the adsorption chiller (Figure 5), average inlet temperature, and the [http://www.prace-ri.eu/] TierO systems and is built by IBM. SuperMUC consists of 18 thin node islands and 1 fat node island, which is referred to as SuperMIG, and also used as a migration system. All compute nodes within an individual island are connected via a fully non-blocking Infiniband network



SuperMUC has an average PUE of 1.23 for the complete month of May 2015. The spikes in the PUE can be explained by the system maintenance windows.

As can be seen. PowerDAM can be used for the evaluation of energy/power consumption of various HPC applications as well as for the energy efficiency assessment of the entire HPC system, reusetechnologies, and the complete data center. The collected data is also used to predict the energy and power consumption of large-scale HPC applications for a given resource configuration [7].

Figure 6: PowerDAM temperature heat map for CoolMUC Linux cluster



Figure 7: Analyzing the efficiency of the adsorption chiller with PowerDAM

FDR10 for the thin nodes and QDR for the fat nodes. The 18 thin islands in total contain 147,456 processor cores in 9,216 compute nodes. Each node is equipped with 2 x 8 core processors. Each processor is of Sandy Bridge-EP Xeon E5 - 2680 8C type. SuperMUC has a peak performance of 3.2 PetaFLOPS (=3.2 × 10¹⁵ Floating Point Operations per Second (FLOPS)). As can be seen,

PowerDAM is and will be the underlying framework for all further investigations regarding the reduction of energy consumption and the improvement of data center energy efficiency at LRZ, by collecting and analyzing the full set of influencing parameters: building and cooling infrastructure, supercomputer hardware, application and algorithms, systems software and tools.

Acknowledgments

The work presented here has been carried out within the SIMOPEK [http:// simopek.de] project which has received funding from the German Federal Ministry of Education and Research (BMBF) under grand agreement no. 01IH13007A. The work was achieved using the GSC (Gauss Center for Supercomputing) resources at BAdW-LRZ with support of the State of Bavaria, Germany. The authors would like to thank Jeanette Wilde for her valuable suggestions and comments.

contact: Hayk Shoukourian, hayk.shoukourian@lrz.de



Figure 8: PowerDAM PUE report for SuperMUC Phase 1

References

[1] Dongarra, J.,

Visit to the National University for Defense Technology Changsha, China. Technical report, University of Tennessee, 06 2013

[2] Koomey, J. G., Belady, C., Patterson, M. Santos, A. and Lange, K.-D.,

Implications of recent trends in performance, costs, and energy use for servers. The Green Computing Book: Tackling Energy Efficiency at Large Scale, page 297, 2014

[3] Bergman, K., Borkar, S., Campbell, D., Carlson, W., Dally, W., Denneau, M., Franzon, P., Harrod, W., Hiller, J., Karp, S., Keckler, S., Klein, D., Lucas, R., Richards, M., Scarpelli, A., Scott, S., Snavely, A., Sterling, Th., Williams, R. S., Yelick, K., Bergman, K., Borkar, S., Campbell, D., Carlson, W., Dally, W., Denneau, M., Franzon, P., Harrod, W., Hiller, J., Keckler, S., Klein, D., Kogge, P., Williams, R. S., and Yelick, K.,

ExaScale Computing study: Technology Challenges in Achieving Exascale Systems, Peter Kogge, Editor & Study Lead, 2008

[4] Wilde, T., Auweter, A., Shoukourian, H.,

The 4 Pillar Framework for energy efficient HPC data centers. Computer Science - R&D, Springer, Vol 29, pages 241 - 251, 2013. DOI: http://dx.doi. org/10.1007/s00450-013-0244-6

[5] Hackenberg, D., Schöne, R., Ilsche, T., Molka, D., Schuchart, J., Geyer, R.,

[6] Shoukourian, H., Wilde, T., Auweter, A., and Bode, A., Monitoring Power Data: A first step of Environmental Modelling & Software, Elsevier, Vol 56, pages 13 - 26, 2013 ISSN: 1364 - 8152. DOI: http://dx.doi. org/10.1016/j.envsoft.2013.11.011

[7] Shoukourian, H.,

Adviser for Energy Consumption Management: Green Energy Conservation. Doctoral dissertation, Technical University of Munich, Munich, Germany, 2015

An energy efficiency feature survey of the intel Haswell processor, 2015

towards a unified energy efficiency evaluation toolset for HPC data centers. Journal

- Hayk Shoukourian
- Torsten Wilde
- Axel Auweter
- Arndt Bode

4D City - Space-time **Urban Infrastructure** Mapping by Multi-sensor **Fusion and Visualization**

Static 3D city models are well established for many applications such as architecture, urban planning, navigation, tourism, and disaster management. However, they do not represent the dynamic behavior of the buildings and other infrastructure (e.g. dams, bridges, railway lines). Very high resolution spaceborne Synthetic Aperture Radar (SAR) Earth observation satellites, like the German TerraSAR-X, provide for the first time the possibility to derive both shape and deformation parameters of urban infrastructure on a continuous basis.

Therefore, this project is aimed to generate 4D (space-time) city models and their user specific visualizations to reveal not only the 3D shape of urban infrastructures but also their deformation patterns and motion.

a very high level of detail. The deformation of different parts of individual buildings will be accessible for different users (geologists, civil engineers, decision makers, etc.) to support city monitoring and management and risk assessment.

Results and Methods

The main technique employed in the project is the so-called differential SAR tomography (TomoSAR). For retrieving the 3D position and the deformation parameters of one pixel, we solve an inversion problem with typical dimension of 100x1,000,000 (the forward model matrix). Considering the large size of spaceborne SAR images, the computation over a large area is very time consuming.

Dataset	# of images	Image size	Resolution
Las Vegas	180	11kx6k	0.6x1.1m ²
Berlin	550	11kx6k	0.6x1.1m ²
Shanghai	29	25kx55k	1.2x3.3m ²

Table 1: Processed datasets and size

The research envisioned in this project will lead to a new kind of city models for monitoring and visualization of the dynamics of urban infrastructure in

With the resource of LRZ, we are so far the only team in the world that can produce the 3D reconstruction

and deformation in city-scale using TomoSAR. In total, this project consumed 6 mio core-hours and 10 TB of storage. For each processing, over 500 cores are required. So far, we have processed the datasets listed in Table 1.

In the following content, some representative results are shown.

Las Vegas





Figure 1: Upper: TerraSAR-X high resolution spotlight image of Las Vegas, and lower: 3D point cloud of Las Vegas reconstructed using our algorithm. Color represents the height [2]-[3].

The upper subfigure of Figure 1 is one of the input TerraSAR-X images of Las Vegas. By applying the TomoSAR algorithm on tens of such images, a 3D point cloud was reconstructed (lower subfigure). This point cloud contains around 10 million points. Most importantly, each point contains not only the 3D position information, but also its deformation information with an accuracy of better than millimeter per year (so-called 4D).



3cm/y.

Berlin



Figure 3: Fusion of two reconstructed 3D point clouds. The combined point cloud contains over 40 million points [5].

By fusing two point clouds from different viewing angles, we obtain a complete coverage over an entire city. Figure 3 is the example of Berlin. As always, each point is associated

Figure 2: Deformation estimates of an area around the Las Vegas Convention Center: linear deformation velocity (unit: mm/y)[4].

Figure 2 shows an example of motion estimation. Since July 2009, an area in the city of Las Vegas is undergoing a pronounced subsidence centered at the Convention Center. The estimated linear deformation velocity is up to

with its movement information. The combined point cloud contains about 40 million points. The number of points exceeds 100 million, if we combine all six reconstructed point clouds of Berlin.



Figure 4: Reconstructed 4D building façade model, amplitude of seasonal motion is color-coded. [unit: mm]

4D Building Model

Figure 4 illustrates the potential of these point clouds. It depicts the façade model of the Bellagio Hotel in Las Vegas reconstructed from the TomoSAR point cloud [6]. The points are overlaid onto the model with the color indicating the estimated motion parameter (here is the amplitude of seasonal motion caused by thermal dilation). This information can be used for developing dynamic building models from spaceborne SAR data that can help to monitor individual buildings and even the whole city.

On-going Research / Outlook

Based on the achieved results from previous years, we would like to conduct further research on large area object reconstruction from the resulting TomoSAR point clouds exemplified in Figure 1 and Figure 3, as well as SAR image filtering for improving the quality of existing global digital elevation model (DEM) in terms of resolution and noise reduction.



Figure 5: Jülich city area: Optical image ©Google (left), the standard 12m TanDEM-X DEM (middle) and improved nonlocal TanDEM-X DEM with 6m resolution (right).

This would require computational resource for:

- Object reconstruction algorithm developments and tests based on the TomoSAR point clouds
- Scientific visualization of the reconstructed dynamic city models
- Non-local filtering of SAR images.

Figure 5 is a comparison of the standard 12m TanDEM-X DEM and the nonlocal filtered TanDEM-X DEM [7].

References

- [1] http://www.sipeo.bgu.tum.de/
- [2] Zhu, X.,

Very High Resolution Tomographic SAR Inversion for Urban Infrastructure Monitoring: A Sparse and Nonlinear Tour, Deutsche Geodätische Kommission, 2011

[3] Zhu, X., Wang, Y., Gernhardt, S. and Bamler, R.,

Tomo-GENESIS: DLR's Tomographic SAR Processing System. In Urban Remote Sensing Event, 2013 Joint. 159-162, 2013

[4] Zhu, X., Bamler, R.,

Let's Do the Time Warp: Multicomponent Nonlinear Motion Estimation in Differential SAR Tomography. IEEE Geoscience and Remote Sensing Letters 8 (4), 735-739, 2011

[5] Wang, Y., Zhu, X.,

Automatic Feature-based Geometric Fusion of Multi-view TomoSAR Point Clouds in Urban Area. IEEE Journal of Selected Topics in Applied Earth Observation and Remote Sensing 8 (3), 953-965, 2014

[6] Zhu, X. and Shahzad, M.,

Façade Reconstruction Using Multiview Spaceborne TomoSAR Point Clouds. IEEE Trans. Geosci. Remote Sens. 52, 6, 3541-3552, 2014

Germany, 2014

contact: Xiaoxiang Zhu, xiao.zhu@dlr.de



[7] Zhu, X., Bamler, R., Lachaise, M., Adam, F., Shi, Y., Eineder, M., Improving TanDEM-X DEMs by Non-local InSAR Filtering, in: EUSAR 2014, Berlin,



- Xiaoxiang Zhu
- Yuanyuan Wang
- Gerald Baier
- Yilei Shi

und Raumfahrt Methodik der

Simulation of Interactions at the LHC



Figure 1: Example of a proton-proton collision event in ATLAS LHC Run-2. This event was collected in September 2015 and characterized by two particular high energetic jets which have an invariant mass of 8.8 TeV

Our project (ATLMUC) runs simulations of high energy proton-proton collisions in the large hadron collider (LHC) at CERN combined with a simulation of the ATLAS detector response on SuperMUC. After the successful LHC run-1 from 2010-2012 where protonproton collisions at a centre-of-mass energy of 7/8 TeV were studied, LHC started its 2nd phase (LHC run-2) in 2015 operating at a centre-of-mass energy of 13 TeV. The ATLAS experiment [1] is one of two multi-purpose experiments at the LHC designed to record large numbers of these protonproton collision events. Figure 1 shows

an example event from the recent data taking. The ATLAS collaboration has already published more than 400 journal articles including the celebrated discovery of the Higgs Boson. In searches for new phenomena, as well as for precise measurements, simulations of protonproton collisions, based on theoretical predictions, combined with a detailed simulation of the detector response are indispensable. These simulations are computationally expensive, e.g. the complete simulation of a complex collision event takes up to 1,000 seconds on a single CPU core.

The ATLAS experiment records about 10 billion collision events per year. The detailed analysis of this data requires at least the same amount of simulated events for the standard processes in order to perform the baseline optimizations and background corrections. Detailed searches for contributions from "New Physics" processes - the key focus of the LHC program – require additional samples of simulated events for these processes, typically for multiple settings of parameters specific for these models. These measurements provide stringent constraints on theoretical models beyond the Standard Model of particle physics, as illustrated in Figure 2 for the case of Supersymmetry [2], one possible extension of the Standard Model.

In many cases the scientific output of the ATLAS collaboration is not limited by the capacity to process and reduce the data but by the capacity to produce sufficient simulated data. Therefore using CPU resources at HPC systems such as SuperMUC/LRZ is a crucial extension of the worldwide LHC computing grid resources which primarily focus on data storage and reconstruction of LHC events.

Our project is a cooperation between groups at LMU Munich and Max-Planck-Institut für Physik, which are both members of the ATLAS collaboration. We also did benefit from the C2PAP cluster at LRZ which provided both a testbed for the batch system and basic infrastructure services.

Results and Methods

SuperMUC was integrated into the ATLAS production system to run the cpu-intensive part of the Monte Carlo simulation of LHC events in the ATLAS detector. The integration required a gateway service to receive job requests, stage-in input data, submit into the batch system, and stage-out the output data. Due to the large number of jobs submitted, automatized submission and data distribution procedures are required. The gateway is provided by an ARC CE [3] running on a remote node with key-based ssh access to the SuperMUC login nodes. Submission into the batch system and subsequent monitoring proceeds using commands run via ssh. The GPFS files systems are fuse-mounted (sshfs) and therefore available for stage-in and -out of data. Although this gateway solution serves the current purpose and is remarkably stable, a similarly capable gateway native to SuperMUC would be the optimal solution. Submission is strictly controlled via X509 certificates. The workloads are a well-defined subset of ATLAS central production workflows, namely detector simulation based on Geant4 [4]. Geant4 is a toolkit for simulation of the passage of particles through matter. It is cpu-limited and dominated by integer arithmetic operations. Although the passage of particles through matter is serial by nature, ATLAS developed a means to efficiently utilize multiple CPU cores. After initialization the process forks to N sub-processes making use of copy-on-write shared memory, thereby significantly reducing the memory requirement per core. Each process then processes a stream of independent events, before merging at the end. This enables the efficient use of whole-nodes and thereby fulfills the basic SuperMUC requirement.

The workloads are deliberately defined to be short (<4hrs) in order to maximize backfill potential. The project

was accepted on the basis of backfill with pre-emptable jobs. In lieu of checkpointing, upon which we are working, the short jobs ensure little work is lost In order to cache the input files GPFS in case of pre-emption.

The initial 10M core-hours allocation was consumed by October 2015, after which 10M core-hours were added. The main problem encountered was the limited GPFS client performance on phase-1 compute nodes. It leads to a halving of the CPU efficiency, due to delays in file access. On the MPCDF Hydra system this issue could be solved by some GPFS client reconfiguration. Also SuperMUC phase-2 nodes do not show this problem. For the phase-1 nodes a partial solution was found by using Parrot-cvmfs [5] for the software access. The cymfs part caches file metadata and leads to reduced GPFS lookup.

scratch space is used. The GPFS scratch space is also used for the work directories of active jobs. The GPFS work space is used to store the software, in a format required by the cymfs client. The 11 TB guota is adequate for the foreseeable needs.

On-going Research / Outlook

The LHC run-2 is planned to continue until end of 2018 and should increase the data volume by at least a factor 5 compared to run-1. A corresponding increase of the simulated data volume is required in order to analyze and interpret the recorded data. This will allow



Figure 2: Example of ATLAS results on searches for super-symmetric particles predicted in extensions of the Standard Model theory. The plot shows 95% confidence level exclusion limits for the production of charginos and neutralinos [2].

us to determine with much better precision the properties of the Higgs Boson and either find new particles as predicted by "New Physics" theories or further increase the constraints on these models. Using SuperMUC to simulate events will be a crucial component to reach these goals.

Active development of the simulation software is ongoing in order to make the workflow more flexible and better parallelizable for smaller work-units. An MPI based application to distribute events across nodes in a multi-node job, should be available soon. This would allow us to run jobs on multiple nodes and flexibly adjust to available slots per SuperMUC island for backfill. The same development will allow us to store the events simulated before the system preempts a job, thereby only losing the events processed at that moment. This provides an effective high-level check-pointing - a pre-requisite for the effective exploitation of preemptable backfill capacity.

A longer-term goal is to adapt the software for Intel/Mic architectures, though presumably this will be ready only after LHC run-2 (Run-3 is planned to start in 2021). We would welcome if "SuperMUC Next Generation" includes Intel/Mic architecture extensions.

References

- [3] Advanced Resource Connector
- (2003) 250-303
- filesystem/hpc

contact: Günter Duckeck, Guenter.Duckeck@physik.uni-muenchen.de



[1] The ATLAS experiment, http://atlas.ch/

Search for the electroweak production of super-symmetric particles in $\sqrt{s=8}$ TeV pp collisions with the ATLAS detector, ATLAS collaboration, Phys. Rev. D 93, 052002 (2016)

middleware for lightweight computational Grids. M.Ellert et al., Future Generation Computer Systems 23 (2007) 219-240

Geant4, Nuclear Instruments and Methods in Physics Research A 506

http://cernvm.cern.ch/portal/

- Günter Duckeck
- Rodney Walker
- S. Kluth

Maximilians-München,

BEAM-ME

BEAM-ME

The liberalization of the energy sector and the increasing decentralization of country wide electrical power supplies cause a high complexity of techno-economical systems in terms of production, distribution and storage of electrical power when realizing renewable energy technologies. Against the background of the desired creation of a European internal energy market and the on-going transformation of the renewable energy - energy-supply-system (RE-ESS) in the cause of Germany's renewable energy strategy, the application of system analytical models gains more and more significance since questions have to be answered arising from this development. These questions especially concern the identification of sustainable investment strategies to ensure supply safety, in connection with the constantly increasing share of renewable energy production.

During the past years it became obvious that the application of energy-system models for focused investment- and deployment-optimizations requires more computational power than even a modern workstation can deliver. This problem becomes even more evident when models with high temporal and special resolution are applied to enable the optimization of decentralized RE-ESSs.

Facing these challenges the target of this research project is to harvest the potential that parallelized computations on High Performance Computing systems provide for the high resolution optimization strategies of the energy-system analysis.

To reach this target developments in the area of solution algorithms as well as in the implementation of new computing strategies for energy-system models have to be done. These developments are only possible with an approach that integrates the required expertise of applied mathematics and informatics. A proof of this concept will

be done by means of the energysystem model ReMix, developed over the past 10 years at the DLR, providing the needed level of detail in all sectors of the European energy-systems. The portability of the newly developed optimizing strategies to other energy-system models will additionally be analyzed by means of external modeling partners, to ensure that the developed methods are of general kind. Another important aspect, is the portability of the developed methods to different (High Performance) Computing systems, which is why two of Germany's biggest Computing centers are direct partners of the project.

contact: Ralf Schneider, schneider@hlrs.de



Gefördert durch:

Bundesministerium für Wirtschaft und Energie

aufgrund eines Beschlusses des Deutschen Bundestages

Ralf Schneider

Mont Blanc 3



Back in October 2011, the Mont Blanc consortium launched the first phase of a project aimed at exploring an energy-efficient alternative to current supercomputers, based on low-power mobile processors, with the ambition of setting future HPC standards for the Exascale era. Four years later, the Mont-Blanc project has indeed given birth to a fully functional prototype that allowed project members to demonstrate the viability of using European embedded commodity technology for High Performance Computing. The

project also defined a set of developer tools and ported real scientific applications to this new environment. More globally, the project has given good visibility to the concept of using mobile technology for HPC.

The third phase of the Mont Blanc project is coordinated by Bull, the Atos brand for technology products and software, and has a budget of 7.9 million Euros, funded by the European Commission under the Horizon2020 program. It was launched in October 2015 and will run for 3 years until autumn of 2018. Mont Blanc 3 adopts a co-design approach to ensure that hardware and system innovations are readily translated into benefits for HPC applications. It aims at designing a new high-end HPC platform that is able to deliver a new level of performance / energy ratio when executing real applications.

HLRS' Perspective on the Project

The High Performance Computing Center Stuttgart (HLRS) joined the Mont Blanc consortium in phase two of the project with the goal to develop a graphical debugging tool for the project's task-based parallel programming model OmpSs. In the third phase, the research focus of HLRS will shift from tools to hybrid programming models and evaluation through applications.

With increasing number of concurrent processes it becomes increasingly important reduce idle times due to MPI communication. One way to do that is hiding communication latencies by overlapping message transfers with computations. HLRS develops a hybrid programming model that will allow



encapsulating MPI operations inside OmpSs tasks in order to progress MPI communication and application code concurrently with minimal effort from the developer.

• BSC, Spain



CINIS

н

UC

In addition, HLRS will contribute to the evaluation of the project's prototype systems and it's the developed software stack by porting various relevant simulation codes taken from the engineering application field.



The Mont Blanc 3 consortium consists of 3 industrial partners and 7 academic partners from 6 countries.

Industrial partners:

• Bull/Atos (coordinator), France



• ARM, United Kingdom







contact:

Josè Gracia, gracia@hlrs.de

Academic Partners:

CNRS, France

HLRS, Germany



• ETH Zürich, Switzerland,

ETH zürich

• Universidad de Cantabria, Spain

Université de Versailles, France



• Universität Graz, Austria

• Josè Gracia

Dynamic Scheduling of reconfigurable scientific Workflows - The DreamCloud approach to Energy-efficient **HPC** Infrastructures

The major challenges for the HPC infrastructures of the current, "exascale-aiming" generation are:

- steadily growing scale of applications and
- corresponding energy consumption increase.

Whereas the parallelization technologies like MPI or PGAS offer some (however, imperfect) solutions to the application scaling problem, all of them remain absolutely agnostic of the energy-efficiency of the underlying infrastructure. Starting from 90 nm chip manufacturing technology, the amount of power lost due to the static power dissipation has been

taking over the power lost due to dynamic one. Power usage constraints, along with traditional speedup and scaling factors, became one of the important optimization objectives for computation-intensive applications and for the HPC providers. On the other hand, explicit underutilization of the existing resources defeats the original purpose of owning and running the computing resources.

The mainstream HPC schedulers assume the infrastructure resources are homogeneous and the applications workflows are static. However, this is no longer so. When looking through the recent TOP500 [1] list of the world's most powerful HPC infrastructures, it is not hard to note that



Figure 1: Example of dynamic workflow scheduling

the amount of heterogeneous hardware has considerably increased. Accelerators like GPU and FPGA are not novel to the HPC world any more. The modern applications take advantages of the virtualization (aka Cloud) technologies and component-based execution (such as microservices), which allow them to run on any resource of the heterogeneous infrastructure. Moreover, the typical HPC application workflows are getting increasingly dynamic, which means they can be scaled (e.g. by replicating workflow components, see Fig. 1) to fit the amount of computation/data to be performed/analyzed.

Parallel applications are largely influenced by a multitude of run-time events including network congestion, hardware failure, or contention for compute resources. As a consequence, executing the same application on the same infrastructure may result in significantly varying execution times. The challenge of scheduling dynamic application workflows to the heterogeneous infrastructures lies in reconsideration of the canonical HPC scheduling algorithms, which employ simplified allocation policies by targeting only homogeneous infrastructures. In particular, the following essential requirements have to be covered:

• the user should be able to expose non-functional requirements such as the desired duration of the application execution or the allowed energy consumption per execution to the scheduler

 the duration of tasks that are combined in dynamic workflows might not be known in advance

• the scheduler should be predictive with regard to the expected application execution time and energy consumption (along with the other nonfunctional requirements) and offer the resources of the type and in capacity as required for meeting them

 the scheduler should enable nesting of different tasks on the resource aiming to increase the tradeoff between the infrastructure utilization and the application characteristics (QoS guarantees) for oversubscribed HPC infrastructures, i.e. when the application demand is much higher than the available resources.

The EU-funded project DreamCloud [2], started in December 2014 and finishing in August 2016, addresses the problems of the static scheduling by elaborating new, dynamic techniques of resource allocation. DreamCloud proposed a new, adaptive approach to application scheduling in HPC environments that takes into account the dynamic power management strategies. The approach aims to enable the awareness of the scheduling middleware of the application behavior in heterogeneous environments and thus become predictive with regard to the i) application performance and ii) infrastructure utilization. DreamCloud leverages heuristics to analyze the application behavior (e.g. performance and energy consumption) in different modes and conditions that the infrastructure might operate. The main innovation of DreamCloud is that it seeks to cross-fertilize the allocation techniques used in the embedded (with a well-established set of the resource allocation and management solutions) and in the HPC domains.

Adaptive Scheduling of **Reconfigurable Applications**

DreamCloud's dynamic approach (see Fig. 2) specifically targets scalable workflow-based applications, complementary to the traditional techniques established for HPC. The primary goal is to enable a trade-off between the optimal resource utilization and the required application performance by adhering to the required energyefficiency standards. The chosen way is to allocate computational tasks to resources in a way that addresses both energy consumption and execution time (an "NP-complete problem"). This challenge is tackled by the DreamCloud platform by elaborating scheduling heuristics [3]. The heuristics leverage efficient machine learning algorithms in order to react to any changes in the infrastructure, workload, or user preferences as well as to come up with a scheduling strategy that selects the most opti-



Figure 2: Dynamic scheduling approach.

mal system configuration as a function of energy, power, and workflow characteristics.

The dynamic approach improves the functionality of the native scheduler by:

 suggesting a specification for dynamic application workflows that enables a number of optimization options that specifically target heterogeneous infrastructures

- reducing the energy consumption by enabling dynamic power management strategies for the hardware
- leveraging heuristics to optimally control the application performance and energy consumption in different modes of the hardware functioning

 using the monitoring platform to collect performance and energy profiles of the dynamic application workloads.

DreamCloud Platform

The conceptual architecture of the dynamic scheduling platform is shown in Figure 3. The key components of the architecture are following:

- Workflow Manager responsible for processing incoming workflow submission requests, initiating task execution, managing task status
- Resource Manager manages the nodes status according to the selected energy-efficiency policy
- Scheduling Advisor generates deployment plan for a task, enables interfacing with Heuristics Manager



Figure 3: The dynamic scheduling platform (numbers represent the interaction flow between the components during task scheduling)

 Heuristics Manager - implements the dynamic scheduling heuristics

 Monitoring Platform - is responsible for generation, collection, and analysis of the communication and power efficiency profiles which are required by the Heuristics Manager.

Outlook

Leverage of the heuristics in order to determine the dynamic properties of the reconfigurable scientific HPC workflows offers a promising strategy to improvement of the balance between the application's nonfunctional requirements and the infrastructure's operation policies. The dynamic scheduling platform, that was proposed by DreamCloud project as an extension of the native scheduling middleware, has a potential to achieve a performance and infrastructure utilization improvement, which was assessed up to the level of 30% (cf. [3]). We will be keeping the HPC community informed on the final evaluation results, which are planned to be provided by the middle of 2016.

References

Indrusiak, L.S.

contact: Alexey Cheptsov, cheptsov@hlrs.de

[1] TOP500 list of the most powerful HPC infrastructures. Available at: http://top500.org/lists/2015/11/, retrieved on 4.05.2016

> DreamCloud project website, available at http://www.dreamcloud-project.org/, retrieved on 4.05.2016

[3] Singh, A.K., Dziurzanski, P., and

Market-inspired Dynamic Resource Allocation in Many-core High Performance Computing Systems, Proceedings of the IEEE International Conference on High Performance Computing & Simulation, 2015, pp. 413-420

Alexey Cheptsov

ParaPhase: Towards HPC for adaptive Phase-Field Modeling



Figure 1: Phase-field simulation of drying soil (Miehe & Mauthe, Univ. Stuttgart)

Phase-field models are an important class of mathematical techniques for the description of a multitude of industryrelevant physical and technical processes. Examples are the modeling of cracks and fracture propagation in solid media like ceramics or dry soil (see Figure and [1]), the representation of liquid phase epitaxy [2] for solar cells, semi-conductors or LEDs as well as melting and solidification processes of alloys.

The price for the broad applicability and mathematical elegance of this approach is the significant computing cost required for the simulation of phase-field equations at large scales. Solutions of these equations 1. Heterogeneous parallelization in typically contain sharp interfaces moving through the domain. Such structures can only be resolved with carefully tuned, adaptive discretization schemes in space and time. Even worse, many key phenomena start to emerge only when the simulation 2. Innovative parallelization in time. For domain is large and the simulation time is long enough. For example, in order to simulate micro cracks leading to fatigue failure of a piece of machinery, the domain must contain a certain number of these cracks. For epitaxy, in turn, structures are normally described on nano-scales, while the specimen sizes are on the order of centimeters. Thus, the enormous number of

degrees-of-freedom for the discretization in space and time as well as the significant complexity of the simulation demand the use of modern HPC architectures.

The goal of the BMBF project "ParaPhase - space-time parallel adaptive simulation of phase-field models on HPC architectures" (FKZ 01IH15005A, BMBF program "IKT 2020 - Forschung für Innovation") is the development of algorithms and methods that allow for highly efficient space-time parallel and adaptive simulations of phasefield problems. Three key aspects will be addressed in the course of the project:

- space. The adaptive phase-field multigrid algorithm TNNMG [3] will be parallelized using load-balanced decomposition techniques and GPUbased acceleration of the smoother.
- optimal parallel performance even on extreme scale platforms, novel approaches like Parareal and the "parallel full approximation scheme in space and time" [4] for the parallelization in the temporal direction will be used, exploiting the hierarchical structures of spatial discretization and solver.

3. High-order methods in space and time. To increase the arithmetic intensity, i.e., the ratio between computation and memory access, flexible highorder methods in space (using the Discontinuous Galerkin approach) and time (using spectral deferred corrections) will be implemented and combined.

This project brings together scientists from applications, numerical mathematics and HPC. The consortium is lead by JSC and consists of six German partners: Prof. Dr.-Ing. Heike Emmerich, Chair for Materials and Process Simulations at Universität Bayreuth, JProf. Dr. Carsten Gräser, Department of Mathematics and Computer Science at Freie Universität Berlin, Prof. Dr.-Ing. Christian Miehe, Chair of Material Theory at Universität Stuttgart, Prof. Dr. Oliver Sander, Institute of Numerical Mathematics at Technische Universität Dresden, Dr. Robert Speck, Jülich Supercomputing Centre at Forschungszentrum Jülich GmbH, and Jiri Kraus at NVIDIA GmbH.

With two partners with a long-standing experience in the particular field of applications (the groups of Heike Emmerich and Christian Miehe) as well as four partners with a strong background in methods, algorithms and HPC (the groups of Carsten Gräser, Oliver Sander, Robert Speck and Jiri Kraus), one of the key goals of this project is the mutual extension of competences: While the application scientists will ultimately be able to work with unprecedented resolutions in space and time on modern HPC platforms, the HPC experts and method developers will have the unique opportunity to test and tune their ideas for industry-relevant, highly complex problems. The algorithms developed in this project will be primarily used for studying fracture propagation

and liquid phase epitaxy, but these problem classes already represent a wide range of challenges in industrial applications. Based on the open source software Dune, the "Distributed and Unified Numerics Environment" [5], the resulting algorithms will help to make large-scale HPC simulations accessible for researchers in these fields

The Kickoff Meeting is planned for the end of May 2016, marking the official start of this three-years project. With two postdocs and three PhD students fully funded by BMBF, this will also contribute to the educational effort in the field of computational science and engineering, enabling young scientists to address challenges of real-world applications with HPC-ready methods and algorithms.

References

Hofacker, M.,

Eck, C., Emmerich, H.,

A two-scale model for liquid-phase epitaxy. Mathematical Methods in the Applied Sciences 2009;32: 12-40

Gräser, C., Sack, U., Sander, O.,

Truncated Nonsmooth Newton Multiarid Methods for Convex Minimization Problems. In: Domain Decomposition Methods in Science and Engineering XVIII (eds. Bercovier, M., Gander, M., Kornhuber, R., Widlund, O.), Volume 70 of LNCSE, pages 129-136

- Gander, M. J.,

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

1] Miehe, C., Welschinger, F.,

Thermodynamically consistent phase-field models of fracture: Variational principles and multi-field FE implementations. International Journal for Numerical Methods in Engineering 2010;83: 1273-1311

50 years of Time Parallel Time Integration. In: Multiple Shooting and Time Domain Decomposition, Springer, 2015

http://www.dune-project.org

Robert Speck¹ Oliver Sander²

Numerische Mathematik,

DEEP Project successfully completed



The DEEP (Dynamical Exascale Entry Platform) project [1] has come to an end. In its final review, which took place in Jülich on the 2nd October 2015, DEEP was evaluated as a very successful project with a lasting impact in the European HPC landscape.

DEEP ran from December 2011 until August 2015. It is one of the first "European Exascale projects" selected by the European Commission through the FP7 funding programme, and received over 8 million Euro in funding. Led by the Jülich Supercomputing Centre (JSC), the consortium of 16 partners from European HPC industry also and academia included the Leibniz Supercomputing Centre (LRZ).

The Cluster-Booster Architecture

Striving at pushing the applications' scalability to the limits, DEEP proposed a new approach to heterogeneous computing that best matches the different intrinsic concurrency levels in large simulation codes. The Cluster-Booster architecture [2] combines two distinct hardware components in a single platform:

• the Cluster – equipped with fast, general-purpose processors that show highest single thread performance but with a limited total number of (expensive) cores and less energy efficient

• and the Booster – composed of many-core Intel® Xeon Phi™ coprocessors connected by the EXTOLL [3] network, all together most



Figure 1: Cluster-Booster architecture in DEEP

energy efficient, highly scalable and massively parallel.

In DEEP, code parts of an application that can only be parallelized up to a limited concurrency level run on the Cluster profiting from its high singlethread performance, while the highly parallelizable parts of the simulation exploit the high scalability and energy efficiency of the Booster.

Hardware Prototypes

To demonstrate the Cluster-Booster concept, the DEEP project built hardware prototypes that fully leverage leading-edge multi-core and many-core processors, interconnects, packaging and cooling methods and monitoring/control approaches.

The DEEP Cluster consists of 128 dual-socket blade nodes integrated in Eurotech's proven and highly efficient off-the-shelf Aurora technology. Its direct liquid cooling technology enables year-round chiller-less cooling of the system.

The DEEP Booster, on the other hand, has been designed and built from scratch within DEEP and is also based on Eurotech's Aurora technology. With 384 first-generation Intel® Xeon Phi™ coprocessors, the DEEP Booster is one of the largest Xeon Phi based systems in Europe, with a peak performance of up to 460 TFLOP/s. But even more importantly, this prototype is different from anything seen in the HPC landscape until now: it is the only platform world-wide in which the Xeon Phi processors do operate autonomously without being attached to a host. This provides full flexibility in configuring the right combination of Cluster and Booster nodes, to opti-



Figure 2: DEEP System. On the left side the Cluster; on the right side (with blue cabling): the Booster

mize the use application.

Additionally, two smaller prototypes have been constructed in DEEP:

• The Energy Efficiency Evaluator: a smaller DEEP System (4 Cluster nodes + 8 Booster nodes) installed at LRZ to perform energy efficiency experiments

• The GreenICE Booster: a 32-node (Intel Xeon Phi) Booster prototype built by University of Heidelberg and Megware to test the ASIC implementation of the EXTOLL interconnect, and explore immersive cooling.

Projects

mize the use of the hardware for each

Energy Efficiency

The DEEP prototype is based on technologies that reduce the system's energy consumption and help users optimize and tune the system according to their needs. Partner LRZ developed in DEEP comprehensive monitoring capabilities that integrate system and infrastructure sensors, enabling an in-depth analysis of the machine's operating conditions.

The DEEP monitoring system delivers a wealth of voltage, current and temperature data for all system components at high frequency. The innovative monitoring and control hardware and software infrastructure prototyped in DEEP has created substantial progress in the field, showing how high-frequency sensor data can be collected and processed in a scalable way, and how it can effectively interact with the firmware of the system components to ensure safe and efficient operation.

Software Developments

Programming a heterogeneous system like DEEP is a challenging task for developers of HPC applications. To minimize the effort of porting existing applications to the Cluster-Booster architecture, special emphasis was put in the DEEP project on developing a programming model that gives as much support as possible to the users. To achieve it, the team performed a tremendous co-design effort in collaborating with the hardware and the applications team.

While traditional supercomputers are either totally homogeneous or heterogeneous only at the node level, the DEEP system is heterogeneous at the cluster level, combining different

kinds of nodes and networks. In order to hide this complexity from the application developers, the software stack implements two abstraction layers:

 ParTec's ParaStation MPI serves as the basic parallelization layer and was extended into a Global MPI covering both Cluster and Booster

 OmpSs from the Barcelona Supercomputing Centre (BSC) was chosen as programming model and extended to provide flexible and powerful offload features.

The DEEP system software and programming model allow computing tasks to be distributed dynamically to the most appropriate parts of the hardware to achieve highest computational efficiency. Based on existing standards and product quality solutions, extensions were made where necessary to make the unique DEEP features available or enhance the ease of programming. Today, it provides a solid base for increasing the circle of applications optimized for heterogeneous architectures in general, and in particular for the DEEP-ER project.

Furthermore, proven performance analysis and modelling tools from JSC and BSC were extended in the DEEP project to fully support the programming models; they were also used to predict the performance of scaled-up systems, establishing a precedent for full system performance projection in the scaling dimension without the need to first create analytical application models [4].

Applications

Six pilot applications were selected to investigate and demonstrate the benefits of combining hardware, system software and the programming model to leap beyond the limits of Amdahl's Law. During the project they were highly optimized, acted as drivers for co-design leading to the final realiza-

of the DEEP hardware and software architecture. For instance, reverse offloading [6] (Booster to Cluster), I/O offloading and dynamic offloading of discrete tasks are all possible on a DEEP machine, and can easily be ported to other systems.

Across all applications it could be



Figure 3: Application running on the DEEP programming environment

tion of hardware and software in the project, and served to identify the main features of applications that most benefit from the DEEP concept [5].

Every application is different and therefore needs to be considered as a different use case. However, the project delivered impressive evidence of the number of ways HPC applications can benefit from the flexibility

shown that only a limited amount of changes are necessary to benefit from the Cluster-Booster architecture. Furthermore, since the DEEP software interfaces are based on standards, the DEEP-enabled codes continue to run on conventional architectures, sometimes showing surprising performance and efficiency improvements compared to their old formulation.



Figure 4: Project Manager Dr. Estela Suarez standing between the two DEEP Boosters. (Copyright: Forschungszentrum Jülich/Sascha Kreklau)

Even more importantly, the experience gathered in the application analysis and adaptation was distilled into "best-known methods", resulting in a playbook for tackling a wide range of additional applications and preparing them for DEEP-class systems. It is our hope that this will have a profound beneficial effect on the entire application ecosystem.

Impact

With its key achievements and the large body of expertise created, the DEEP project is poised to have a significant and lasting impact. Besides opening up new avenues for the architecture of efficient HPC systems, it has materially increased Europe's indigenous capabilities in HPC system design and production, and has produced a complete system software stack together with a programming environment for heterogeneous platforms. Six relevant applications in

critical fields of the European Research Arena have been remodeled and adapted, and what is more, best-known methods have been established that will enable many more codes to reap the benefits of the DEEP software and hardware architecture.

The DEEP system has proven that the Cluster-Booster architecture concept of dynamically associating different kinds of computing resources to best match workload needs can be implemented with state-of-the-art multicore and many-core technology, and that such a system can indeed provide a superior combination of scalability and efficiency. It has thereby opened up a new avenue towards affordable, highly efficient and adaptable extreme scale systems (up to Exascale-class), merging the hitherto separate lines of massively parallel and commodity Cluster systems. The sibling project DEEP-ER is already carrying the flag further by integrating novel memory and storage concepts and providing scalable I/O and resiliency capabilities [7].

References

[1] The DEEP project, http://www.deepproject.eu

[2] Eicker, N., Lippert, T., Moschny, T., Suarez, E.

The DEEP Project An alternative approach to heterogeneous cluster-computing in the many-core era, Concurrency and computation (2015) DOI: 10.1002/cpe.3562

[3] Neuwirth, S., Frey, D., Nuessle, M., Bruening, U.,

> Scalable communication architecture for network-attached accelerators, High Performance Computer Architecture (HPCA), 2015 IEEE 21st International Symposium on; 7-11 February 2015, IEEE, 2015, pp. 627-638, DOI: 10.1109/ HPCA.2015.7056068

[4] Rosas, C., Giménez, J., Labarta, J., Scalability prediction for fundamental performance factors, Supercomputing frontiers and innovations 1 (2) 2014, 4-19, DOI: 10.14529/jsfi140201

[5] Alvarez Mallon, D., Lippert, T., Beltran, V., Affinito, F., Jaure, S., Merx, H., Labarta, J., Staffelbach, G., Suarez, E., Eicker, N.,

> Programming Model and Application Porting to the Dynamical Exascale Entry Platform (DEEP), Proceedings of the Exascale Applications and Software Conference, Edinburgh, Scotland, UK. 9-11 April 2013

[6] Jakobs, A., Zitz, A., Eicker, N., Lapenta, G.,

Particle-in-Cell algorithms on DEEP: The iPiC3D case study, Proceedings of the 26. GI/ITG Workshop Parallele Algorithmen, Rechnerstrukturen und Systemsoftware (PARS), Potsdam, 7-8 May 2015, PARS-Mitteilungen 32 (2015), p. 38-48

[7] The DEEP-ER Project, http://www.deep-er.eu

contact: Estela Suarez, e.suarez@fz-juelich.de • Estela Suarez

Scientific Big Data Analytics by HPC



Figure 1: 'Scientific Big Data Analytics' by HPC and the role of large-scale research

Managing, sharing, curating and especially analyzing the ever increasing large quantities of data face an immense visibility and importance in industry and economy as well as in science and research. Commercial applications exploit "Big Data" for predictive analysis, to increase the efficiency of infrastructures, customer segmentation, and tailored services. In contrast, scientific big data allows for addressing problems with complexities that were impossible to deal with so far. This article offers the reader a view of how HPC addresses some of the exponentially growing data challenges that are evident in all areas of science with a particular focus on "Big Data Analytics".

Scientific Big Data Analytics

There is a wide variety of parallel and scalable "Big Data Analytics" technologies and approaches in the field. Researchers are constantly distracted from exploring new and partly innovative technologies and in too many cases lack a sound infrastructure for developing solutions beyond smallscale technology testbeds. The notion "Scientific Big Data Analytics (SBDA)" comprises the work of researchers and engineers, that is based on the creation and analysis of big data sets, and thus relies on the availability of appropriately sized infrastructures as highlighted in red in Figure 1. This is required in order to be competitive in the respective scientific domain.

It is necessary to provide large-scale infrastructures to scientists and engineers of university and research institutes, who perform projects with highest demands on storing, transfer, and processing of data sets. We therefore foresee the importance of establishing a systematic provisioning process in the field of SBDA. This provisioning should be done similar to the provisioning of HPC resources as done for the simulation sciences. This constitutes a key element in the SBDA building block in our framework shown in Figure 1 and explained by Lippert et al. in [2].

Importance of Peer-Review

In order to guarantee that the data analysis achieves the highest scientific quality, it is necessary to apply the principle of autonomous controlling of resource allocation by science in a competitive peer-review process, like it is common practice for international large-scale infrastructures. In addition, the scientific controlling of resource allocations will allow to focus on problem areas that are highly

relevant for science and society. The

steering process prevents that sci-

ence gets lost in the details of this industry-driven topical area, with many technologies that are highly relevant only for commercial applications (e.g. recommender engines, shopping basket association rule mining, etc.). Instead, new approaches will be developed and, subsequently have to be translated to economy and industry. Scientific technologies and approaches in this field will mature, leading eventually to community-approved codes to tackle an ever increasing amount of research questions. Hence, the effective usage of the SBDA infrastructure as illustrated in Figure 1 is ensured by a scientific peer-review. These scienceled processes not only guarantee the most beneficial usage of the infrastructure, but also steer their evolution and focus through the involvement of research communities in key areas of science and engineering.

Role of the NIC

The John von Neumann Institute for Computing (NIC) has established a scientific peer-review process for the provision of supercomputing cycles at the Jülich Supercomputing Centre (JSC).



Figure 2: The full productive loop of scientific big data analytics by HPC

Scientists and researchers who apply for computing time on infrastructure resources are supported by a continuously growing number of domain-specific simulation laboratories (SimLabs) at JSC. The SimLabs offer support on advancement, it is natural to apply the NIC provisioning concept to data infrastructures and analysis resources as well. This kind of provisioning together with the activities in the Helmholtz programme "Supercomput-





Figure 3: AVIRIS Indian Pines (a) image cube representation and (b) ground reference

a high level and push forward research and development in co-design together with specific scientific communities that take advantage of HPC techniques in parallel applications.

The NIC allocation principles served also as a blueprint for the allocation of computing time in the Gauss Centre for Supercomputing (GCS). Given NIC's strong experience and its scientific ing & Big Data", will promote research and development for supercomputing and data infrastructures. It is the heart of truly innovative SBDA and ensures scientific advancement.

In order to gain better insights into the demand by communities and requirements, the JSC has performed an initial step towards implementing principles as proposed in this article. The importance of data analytics, management, sharing, and preservation of very big, heterogeneous or distributed data sets from experiments, observations and simulations is of increasing significance for science, research and industry. This has been recognized by many research institutions, among them leading HPC centers. They want to advance their support for researchers and engineers using SBDA by HPC.

First Experience revealed

For the first time NIC invited Expressions of Interest (Eol) for SBDA projects using HPC to identify and analyze the needs of scientific communities. The goal is to extend and optimize the HPC and data infrastructures and services in order to provide optimal methodological support. As described above, a peer-review was performed on the Eol submissions in a similar manner as known from HPC calls. First experiences were presented as posters at the last NIC symposium in Jülich [1]. They clearly demonstrate in which areas SBDA by HPC is of major importance.

Eol submissions have been received from the field of Biology (e.g. mining of molecular structure data) and Neurosciences (e.g. deep learning analysis of brain images), from the field of statistical turbulence research and from traditional HPC simulation communities in the context of earth sciences (e.g. SimLab Climate and SimLab Terrestrial Systems). One identified common scheme we refer to as the "full loop of SBDA" is illustrated in Figure 2. Solving a wide variety of inverse problems can actually lead to better algorithms for simulation sciences that in turn then deliver more accurate models to understand our world. In order to

establish this strong "productive loop" between HPC simulations and more data-intensive applications, a strong foundational infrastructure is required.

One Example of HPC Impact

Recent advances in remote sensor and computer technology are substituting the traditional sources and collection methods of data, by revolutionizing the way remotely sensed data are acquired, managed, and analyzed. The term "remote sensing" refers to the science of measuring, analyzing, and interpreting information about a scene acquired by sensors mounted on board of different platforms for Earth and planetary observation.

Our motivation is driven by the needs of a specific remote sensing application dataset as shown in Figure 3 based on data from AVIRIS. It raises the demand for technologies that are scalable with respect to big data and thus this application represents one example of a SBDA project that requires HPC. One of the main purposes of satellite remote sensing is to interpret the observed data and classify meaningful features or classes of land-cover types. In hyperspectral remote sensing, images are acquired with hundreds of channels over contiguous wavelength bands, providing measurements that we consider as concrete big data in this example. The reasoning includes not only a large data volume but also a large number of dimensions (i.e., spectral bands). A deeper introduction to the field and its application to HPC using parallel and scalable support vector machines on HPC resources are in G. Cavallaro et al. in [3].

- Morris Riedel^{1,2}
- Thomas Lippert¹
- Daniel Mallmann¹
- Gabriele Cavallaro²

Iceland, Reykjavik

References

We highlight in this article the impact

of HPC since parallelization techniques

lead to significant speed-ups for the

cross-validation and for each training

and testing process of the aforemen-

tioned dataset. To provide an example

during cross-validation, HPC SBDA

techniques have been able to reduce

the time from 14.41 minutes (Matlab)

to 1.02 minutes using just the best parameter set. Since this "best param-

eter" needs to be searched a typical

grid search is performed consisting

of many runs with different parameter

sets. In this context, parallel and scal-

able SBDA by HPC methods have been

able to reduce the time to solution

More notably, this is achieved by

maintaining the same accuracy as achieved when performing the

processes with serial tools (e.g. Matlab, R, libSVM). In the majority

of cases, the minimal training and

testing time was around 1 min

that still can be considered as an

interactive experience thus enabling

remote sensing scientists to perform

easier and faster experiment with

different techniques (e.g., applying quick parameter variations of feature-

extraction techniques). The parallel

and scalable community code for

support vector machines and its

feature engineering approach based

on mathematical morphology are now

maintained by JSC and starting to be

used for other application areas in

science (e.g. neuroscience images) and

industry (e.g. welding image analysis).

from roughly 9 hours to 35 minutes.

http://www.john-von-neumann-institut. de/nic/EN/News/Symposium/NIC-Symposium-2016/PosterSession/_node. html

Lippert, Th., Mallmann, D., Riedel, M.,

Scientific Big Data Analytics by HPC, in: NIC Symposium 2016, Publication Series of the John von Neumann Institute for Computing (NIC), NIC Series 48, ISBN 978-3-95806-109- 5, pp. 1 - 10

[3] Cavallaro, G., Riedel, M., Richerzhagen, M., Benediktsson, J.A., Plaza, A.,

On Understanding Big Data Impacts in Remotely Sensed Image Classification Using Support Vector Machine Methods IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing, Issue 99, pp. 1-13, 2015

contact: Morris Riedel, m.riedel@fz-juelich.de

Improved Safety for disabled People

In order to help disabled people gain easier access, many buildings, e.g. sheltered workshops, as well as other venues are now accessible barrier-free. But how can such places be evacuated effectively if disabled people are involved? How can stakeholders be trained, evacuation plans be adjusted, and facilities be designed to make these sites safer? To answer these questions, the joint project "Safety for people with physical, mental or age-related disabilities" (SiME) has been funded in the context of the "Research Programme for Civil Security" by the German Federal Ministry of 2. Collective movement of people with Education and Research (BMBF). The three-year project began in February 2016 and is coordinated by the Federal Institute for Materials Research and Testing (BAM). Other partners are Ottovon-Guericke University Magdeburg, Hochschule Niederrhein, Werkstatt Lebenshilfe, PTV Transport Consult GmbH, and Forschungszentrum Jülich.

The simulation of the process of evacuating a building enables the identification of bottlenecks or lack of assistance and the calculation of the evacuation time. For this purpose, parameters of the realistic movement of persons involved are needed, but such data are as yet only available for people with unrestricted mobility. In SiME, the team from JSC will execute parameter studies for mixed traffic, i.e. for people with and without disabilities, and also analyse the process of movement of disabled people, e.g. transfer from a wheelchair to an evacuation chair, during an evacuation process.



The intended parameter studies have two focuses:

- large public events.

For the extraction of trajectories of individual participants methods for visual sensors developed during previous projects like Hermes and BaSiGo can be used. But people with a low height inside a dense gathering like persons using wheel chairs will often be occluded so that sensors have to be examined and new methods be developed to track also covered people. After sensor fusion the movement of every single person will be available for further analysis.

With the collected data, more reliable models could be developed to simulate the evacuation of sheltered workshops or homes for people with disabilities. A simulated forecast of the dynamic inside gatherings including people with limited mobility will be more realistic.

contacts: Maik Boltes, m.boltes@fz-juelich.de, Stefan Holl, st.holl@fz-juelich.de



1. Evacuation of disabled people from their daily environment like sheltered workshops or residences

and without disabilities in venues for

 Maik Boltes Stefan Holl

New Czech-Bavarian Competence Centre for Supercomputing **Applications established**



Figure 1: Participants of the "Intel MIC Programming Workshop" at IT4Innovations, Ostrava, February 3-4, 2016

Supercomputing resources are a critical infrastructure for many areas of computational science and engineering: diverse "grand challenges" such as cosmological simulations, hazard assessment of large earthquakes or floods (tsunamis, storm surges, raininduced floods) as well as evaluating risks of long-term storage of CO, or nuclear wastes are strongly relying on large-scale simulations.

In July 2015, the new 2 PFlop/s supercomputer Salomon was put into operation at the IT4Innovations Czech national supercomputing centre in Ostrava, Czech Republic, based on the latest Intel CPU and Intel Xeon Phi coprocessor technology. It is currently the largest machine of this architectural type in Europe and thus serves as a role model for machines on the road to exascale. To ensure the efficient

use of such a critical infrastructure for advanced computing in Europe, it is imperative to provide optimized simulation software and well-trained experts in computational science.

LRZ and the Department of Informatics of TUM in cooperation with IT4Innovations have thus recently established a Czech-Bavarian Competence Centre for Supercomputing Applications (CzeBaCCA) [1] to address the following major goals:

(1) Foster Czech-German collaboration of computational scientists focusing on cutting-edge supercomputing in various fields of high scientific and societal impact

(2) Establish communities of computational scientist that are well-trained on the latest supercomputing architectures

(3) Improve the efficiency of simulation software on latest supercomputing architectures to maximise the utilisation of supercomputers as critical research infrastructure.

Besides their joint research programme around simulation software and tools for Salomon, TUM will organize a series of scientific workshops on topics like optimization of simulation codes, and LRZ will offer a training program specifically designed for Intel Xeon Phi based systems.

The first "Intel MIC Programming Workshop" [2] in this series took place at IT4Innovations, Ostrava on February 3-4, 2016 and attracted over 40 participants. The training covered various Intel Xeon Phi programming models and optimization techniques. During many hands-on sessions the participants were able to gain experience on the new Salomon machine. The training was followed by a one-day workshop "SeisMIC - Seismic Simulation on Current and Future Supercomputers" [3] on February 5, 2016. 13 international speakers contributed talks about the demands and desired features of seismic simulation software, software and performance engineering and parallelization for current and novel HPC platforms, as well as establishing scalable simulation workflows on supercomputing environments.

The Czech-Bavarian Competence Centre for Supercomputing Applications is funded by the Federal Ministry of Education and Research.

References

- supercomputers
 - courses/



Figure 2: Participants of the scientific workshop "SeisMIC - Seismic Simulation on Current and Future Supercomputers" at IT4Innovations, Ostrava, February 5, 2016

The next MIC programming workshop [4] will take place at LRZ on June 27-29, 2016 and will be complemented with a scientific workshop on "High Performance Computing for Water Related Hazards" [4] on June 29 - July 1, 2016 to create synergy effects between the different communities.

contact: Volker Weinberg,

Acknowledgements

1 http://www.lrz.de/forschung/projekte/ forschung-hpc/CzeBaCCA/

http://training.it4i.cz/en/SMIC-02-2016

http://training.it4i.cz/en/seismicseismic-simulation-on-current-and-future-

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

Volker.Weinberg@lrz.de

- Michael Bader¹
- Arndt Bode²
- Volker Weinberg²
- Momme Allalen²
- Branislav Jansík³
- Martin Palkovic³
- Vít Vondrák³
- Ondřej Jakl³

of Munich (TUM),

TerraNeo - Integrated Co-**Design of an Exascale Earth** Mantle Modeling Framework



Figure 1: Velocity streamlines below Iceland, illustrating the different velocity directions due to separation of the Eurasian- and North-American plate



Figure 2: Velocity streamlines below the Indian plate, illustrating a temperature dependent viscosity and an astenosphere with diameter 410 km

Much of what one refers to as geological activity of Earth owes to the fact that heat is transported from our planet's interior to the surface in a planet-wide solid-state convection of Earth's rocky mantle. This convection causes plate tectonics, mountain and ocean building, volcanism, and the accumulation of stresses in tectonic plates leading to earthquakes. Modeling the convective process (see Fig. 1 and Fig. 2 for velocity streamlines and Fig. 3 for a temperature distribution) and linking it to surface manifestations is seen widely as among the most fundamental Earth Sciences problems. Not surprisingly a better understanding of mantle convection belongs to the "10 Grand Research Questions in the Earth Sciences" identified by the U.S. National Research Council.

The sheer magnitude of the computational challenge of mantle flow simulations on modern supercomputers is grasped easily from the dramatic growth of disciplinary complexity. Dominating roadblocks in Geophysics are model complexity and uncertainty in parameters and data, e.g., rheology and seismically imaged mantle structure. Equally challenging are the enormous space and time scales one must resolve, with Earth's interior discretized by \$10^12\$ grid points and simulations spanning \$10^9\$ years. Exascale Computing is needed, but the disruptive transition from modest concurrency to billions of threads leaves the software development trailing behind.

The TerraNeo project set out in 2013 to address these challenges. The project's ultimate goal is the delivery of a new community code based on a carefully designed multi-scale finite element discretization using block-wise refinement for simulating the nonlinear mantle flow transport processes. TerraNeo is part of the SPP 1648 "Software for Exascale Computing" (SPPEXA) of the German Research Foundation, and recently entered its second three-year funding period. The project consortium comprises the groups of Barbara Wohlmuth (Numerical Mathematics, TUM), Ulrich Röde (System Simulation, FAU) and Hans-Peter Bunge (Geophysics, LMU).

Realistic mantle convection simulations require solutions of extremely large systems of equations, and thus the first challenge is to design optimal algorithms that permit a scalable implementation. In TerraNeo we employ multigrid methods, since only this algorithm class has yet been demonstrated to solve indefinite systems with \$10^{13}\$ degrees of freedom on current supercomputers. For reaching problems of this magnitude, scalability is a necessary but no sufficient condition. On modern architectures, the real-life performance is increasingly determined by the intra-node and intra-core efficiency. Therefore it is essential that the algorithms and their implementation also exploit the available instruction level parallelism and the CPU micro-architecture in the best possible way.

Here we employ performance engineering in a co-design process together with the algorithmic development. We quantify the efficiency of a parallel algorithm by generalizing Achi Brandt's notion of textbook multigrid efficiency (TME). The classical definition of TME postulates that a partial differential equation (PDE) must be solved with an effort that is at most the cost of ten operator evaluations. As an extension of this paradigm, we quantify the cost of



applying the discretized operator by analytic architecture-specific performance models. These analytical performance predictions lead to nonfunctional design goals that guide the implementation of the algorithmic kernels and these in turn eventually determine theefficiency of the TerraNeo software.

Figure 3: Temperature plumes in a synthetic coupled mult-physics setting

Among the many other challenges in designing an exascale-ready new framework, we next report on two: resilience and uncertainty. The future era of exascale computing will necessitate that fault tolerance is supported not merely on the level of hardware or operating system, but also algorithmically. Effectively the algorithms themselves must be made error tolerant and augmented with intelligent adaptive strategies that detect and compensate for faults.

On exascale systems, traditional redundancy approaches, such as global checkpointing, might become too costly. With multigrid methods, the dynamic loss of data due to a hard fault can be compensated for by solving a local boundary value problem using multigrid cycles, see Figure 4. This local recovery can benefit from a special superman acceleration. Specifically, the additional compute resources for it are of the size of the faulty



Figure 4: Influence of a dynamic data loss in case of a fault on the multigrid convergence.



Figure 5: Redundant data structure on lower dimensional primitives.

> processing unit and can be provided by additional parallelization, assuming that the use of more processors for the faulty subdomain leads to a significant speed up. We propose here, that e.g.

one assigns a full shared memory node to perform the local recovery for a domain previously handled by a single core. In our prototype this has been realized by a tailored OpenMP parallelization for the recovery process.

The recovery from one or multiple faults becomes especially efficient, when it is suitably decoupled from the global solution process, and both are continued asynchronously in parallel. Several variants of fault tolerant multigrid strategies such as a Dirichlet-Dirichlet (DD) and a Dirichlet-Neumann (DN) approach show excellent parallel performance. During \$n_F\$ multigrid cycles the faulty and the healthy domain are separated and asynchronous sub-jobs take place. For (DD), we perform \$n_F\$ multigrid cycles on the healthy domain with Dirichlet data at the interface to the faulty domain while for (DN), Neumann data are specified. Figure 5 illustrates the halo data structure. We run \$n F \eta s\$ local multigrid cycles with Dirichlet data on the faulty domain. Here \$\eta_s\$ stands for the superman speed up. After \$n_F\$ global steps both sub-jobs are interconnected.

Beyond deterministic forward computations it is crucial to account for uncertainty. Multi-level Monte-Carlo methods are widely applied. But their efficient realization on peta-scale systems is far from trivial. Since the number of samples is often small, their parallel execution alone cannot exploit modern parallel systems to their full potential. This means that multiple layers of parallelism must be identified. Three natural layers exist: level, sample, and solver parallelism, see Figure 6 for a heterogeneous three layer scheduling. The loops over the levels and over the samples are inherently parallel, except for the a

posteriori computation of the statistical values of interest. Nevertheless, it is neither obvious how to handle the load balancing between the different levels of parallelism nor how to schedule the solvers for each sample. The scheduling algorithms must be flexible with respect to the scalability window of the solver, robust with respect to the number of processors and at the same time cheap in light of the NP-completeness of the general scheduling problem.

Promising scheduling strategies may impose additional assumptions and constraints, such that an exact solution of the resulting simplified optimization problem can be computed. Or alternatively they may use meta-heuristic search algorithms such as, e.g., simulated annealing (SA). Three main sources of parallel inefficiencies exist in parallel MLMC algorithms resulting from the non-optimal weak scalability properties, from the machine being partly idle due to large run-time variations of the sample scheduled in parallel or due to over-sampling, i.e., more samples than required are scheduled on the machine. Minimizing the run time as single goal,

leads to a rather flat objective function

in large parts of the search domain. To

improve the performance of the SA, the

number of idle processors is introduced

as second auxiliary objective. When two

candidates have the same MLMC run-

time then the one with the higher num-

ber of idle processors is selected. This

choice is motivated by the observation

that the probability to find a shorter run-

time in the neighborhood of a candidate

is higher if more processors are idle.

-C168, 2015

Waluga, C., and Wohlmuth, B., Towards textbook efficiency for parallel multigrid., Numer. Math. Theory Methods Appl., 8, 2015



and Wohlmuth, B.



Figure 6: Three-layers for parallelism for a

Acknowledgements

We gratefully acknowledge the financial support of the Priority Programme 1648 "Software for Exascale Computing (SPPEXA) of the German Research Foundation as well as 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) and computing time on the GCS Supercomputer SuperMUC at Leibniz Supercomputing Centre.

References

[1] Huber, M., Gmeiner, B., R.de, U., Resilience for multigrid software at the extreme scale., submitted

Gmeiner, B., Röde, U., Stengel, H., Waluga, C., and Wohlmuth, B.,

Performance and Scalability of Hierarchical Hybrid Multigrid Solvers for Stokes Systems., SIAM J. Sci. Comp., 37(2):C143-

Gmeiner, B., Röde, U., Stengel, H.,

- Ulrich Röde¹
- Björn Gmeiner¹
- Dominik Bartuschat¹
- Barbara Wohlmuth²
- Christian Waluga²
- Markus Huber²
- Lorenz John²
- Hans-Peter Bunge³
- Jens Weismüller³
- Simon Baumann³
- Marcus Mohr³

Nürnberg,

München,

Maximilians-München,

lrz

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

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

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

Contact:

Leibniz Supercomputing Centre

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

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



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

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

System	Size	Peak Performance (TFlop/s)	Purpose	User Community
	18 thin node islands IBM iDataPlex 512 nodes, 147,456 cores Sandy Bridge EP 288 TByte, FDR 10 IB	3,185	Capability Computing	German universi- ties and research institutes, PRACE (Tier-O System)
IBM/Lenovo System x Cluster "SuperMUC Phase 1"	1 fat node island with 205 nodes, 8,200 cores IBM Bladecenter HX5 Westmere EX 52 TByte, QDR IB	78	Capability Computing	
	32 accelerated nodes Intel Ivy Bridge EP, Intel Xeon Phi 76 GByte, Dual-Rail FDR 14 IB	100	Prototype System	
Lenovo/IBM Nextscale Cluster "SuperMUC Phase 2"	6 medium node islands Lenovo/IBM Nextscale 512 nodes, 86,016 cores Haswell EP 197 TByte, FDR 14 IB	3,580	Capability computing	German universi- ties and research institutes, PRACE (Tier-O System)
Lenovo/IBM Nextscale Cluster "CooLMUC2"	252 nodes, 7,056 cores Haswell EP 16.1 TByte FDR 14 IB	270	Capability computing	Bavarian Universities
Megware IB-Cluster "CoolMUC"	178 nodes, 2,848 Cores AMD Magny Cours 2.8 TByte QDR IB	22.7	Capability Computing, PRACE Prototype	Bavarian Universities
SGI Altix Ultraviolet	2 partitions, 2,080 cores Intel Westmere EX 6.0 TByte main memory	20.0	Capability Computing	Bavarian Universities
Linux-Cluster	510 nodes, 2,030 cores Intel Xeon EM64T/ AMD Opteron 2-, 4-, 8-, 16-, 32-way 4.7 TByte main memory	13.2	Capability Computing	Bavarian and Munich Universities, LCG Grid

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

Centers

R S **H**

First German National Center

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

Service for Industry

Service provisioning for industry is done together with T-Systems, T-Systems sfr, and Porsche in the public-private joint venture hww (Höchstleistungsrechner für Wissenschaft und Wirtschaft). Through this co-operation industry always has acces to the most recent HPC technology.

Bundling Competencies

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

World Class Research

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

Contact:

Höchstleistungsrechenzentrum Stuttgart (HLRS) Universität Stuttgart

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

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



64 nodes

256 cores

4 TB memory

independent

NEC SX-ACE

clusters

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

16



View of the HLRS Cray XC40 "Hazel Hen"

Centers

Purpose	User Community
Capability Computing	European (PRACE) and German Research Organizations and Industry
Laki: 120,5 TFlops Laki2: 47,2 TFlops	German Universities, Research Institutes and Industry
Vector Computing	German Universities, Research Institutes and Industry



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.

The Jülich Supercomputing Centre (JSC) at Forschungszentrum Jülich is committed to enabling scientists and engineers explore some of the most complex grand challenges facing science and society. Our research is performed through collaborative infrastructures exploiting extreme-scale supercomputing and federated data services.

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

Implementation of strategic support

Higher education for master and doctoral students in close cooperation with neighbouring universities.

Contact:

Jülich Supercomputing Centre (JSC) Forschungszentrum Jülich

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

Compute servers currently operated by JSC

	System	Size	Peak Performance (TFlop/s)	Purpose	User Community
	IBM Blue Gene/Q "JUQUEEN"	28 racks 28,672 nodes 458,752 processors IBM PowerPC® A2 448 Tbyte main memory	5,872	Capability Computing	European Universities and Research Institutes, PRACE
	Intel Linux CLuster "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
	IBM Cell System "GPACE"	1,024 PowerXCell 8i processors 4 TByte memory	100	Capability Computing	QCD Applications SFB TR55, PRACE



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

JSC hosted the 4th JLESC Workshop

From December 2 to 4, the 4th JLESC Workshop took place at the Gustav-Stresemann-Institut in Bonn-Bad Godesberg, organized this time by JSC. This event was the second in 2015 of the biannual meetings of the Joint Laboratory on Extreme Scale Computing (JLESC) and the first one hosted by 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 AICS (Japan) and JSC.

The key objective of JLESC is to foster international collaborations on state-of-the-art research related to computational and data focused simulation and analytics at extreme scales. Within JLESC, scientists from many different disciplines as well as from industry address the most critical issues in advancing from petascale to extreme scale computing. The collaborative work is organized in projects between two or more partners. This includes mutual research visits, joint publications and software releases. Every six months, all JLESC partners meet during a workshop to discuss the most recent results and to exchange ideas for further collaborations.

With more than 100 scientists and students from the six JLESC partners, the meeting in Bad Godesberg covered a broad range of topics crucial for to-

day's and tomorrow's supercomputing. Together with the other participants, scientists and PhD students from JSC and German partner universities could catch up on cutting-edge research from the fields of resilience, I/O and programming models as well as numerical methods, applications, big data and performance tools. Besides their talks, the participants had time for fruitful discussions about their on-going and future research during project meetings, open microphone sessions and a social event. In addition, PhD students and postdocs could attend the first JLESC "young scientist dinner" to exchange ideas in a less formal (and very well-received) setting.

Organized by Inria, the next JLESC workshop will take place at ENS Lyon from June 27 to 29, 2016 continuing this successful series of internationally recognized and valued meetings. This workshop is followed by a two-day summer school on resilience. For more information on JLESC, the workshops and JSC's participation, visit http:// www.fz-juelich.de/ias/jsc/jlesc

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

7th Blue Gene Extreme Scaling Workshop

Feedback from last year's very successful workshop motivated the organization of a three-day extreme scaling workshop February 1-3, 2016 at Jülich Supercomputing Centre (JSC), continuing the series started in 2006. The entire 28-rack JUQUEEN Blue Gene/Q with 458,752 cores was reserved for over 50 hours to allow eight selected code teams to investigate and improve their application scalability, assisted by staff of JSC Simulation Laboratories and Cross-Sectional Teams. Scalasca/ Score-P and Darshan were employed to profile application execution performance, focussing on MPI/OpenMP communication and synchronisation as well as file I/O.

Code Saturne from STFC (Daresbury) and Seven-League Hydro from HITS (Heidelberg) were both able to display good strong scalability and thereby become candidates for High-Q Club membership. Both used 4 OpenMP threads per MPI process, over 1.8 million threads in total. Existing members, CIAO from RWTH-ITV and iFETI from University of Cologne and TU Freiberg, were able to show that they had additional solvers which also scaled acceptably using purely MPI, and in-situ visualization was demonstrated with a CIAO/JUSITU/Visit simulation using 458,752 MPI processes running on 28 racks.

Two adaptive mesh refinement libraries, **p4est** from University of Bonn and **IciMesh** from Ecole Centrale de Nantes, showed that they could respectively run with 917,504 and

Spring 2016 Vol. 14 No. 1 InSiDE

dynQCD JuSPIC KKRnano PMG+PFASST CoreNeuron ICON pSOpen SHOCK PP-Code

k 3

High-Q Club members

....

IMD Gysela JURASSIC MP2C muPhi OpenTBL PEPC TERRA-NEO CoreNeuron FE2TI FEMPAR SHOCK

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

458,752 MPI ranks, but both encountered problems loading large meshes. Parallel file I/O limitations also prevented large-scale executions of the FZJ IEK-6/Amphos²¹ **PFLOTRAN** subsurface flow and reactive transport code, however, a **NEST-import** HDF5 module developed by the EPFL Blue Brain Project could be optimized to use collective MPI file reading calls to load and connect 1.9 TB of neuron and synapse data and enable large-scale

Robert Speck

Activities



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





- Dirk Brömmel
- Wolfgang Frings
- Brian J.N. Wylie

data-driven neuronal network simulations with 458,752 threads.

Detailed workshop reports provided by each code-team, and additional comparative analysis to the 25 High-Q Club member codes, are available in a technical report. Figures 1 and 2 show the strong and/or weak scalability of this year's workshop codes. Despite more mixed results than the previous workshop, we learnt more about application file I/O limitations and inefficiencies which continue to be the primary inhibitor to large-scale simulations, and all of the participants found the workshop to have been very valuable.

References

[1] The High-Q Club, http://www.fz-juelich. de/ias/jsc/high-q-club

2] Brömmel, D., Frings, W., Wylie, B.J.N.

JUQUEEN Extreme Scaling Workshop 2015, Tech. Report FZJ-JSC-IB-2015-01, https://juser.fz-juelich.de/ record/188191

Brömmel, D., Frings, W., Wylie, B.J.N.

JUQUEEN Extreme Scaling Workshop 2016, Tech. Report FZJ-JSC-IB-2016-01, https://juser.fz-juelich.de/ record/283461

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



Figure 3: Visualization client showing CIAO/JUSITU/Vislt with 458,752 MPI processes on JUQUEEN

JURECA Porting and **Tuning Workshop**

In the past couple of years, Jülich Supercomputing Centre (JSC) has been running Porting and Tuning Workshops on its highly scalable BlueGene/Q system JUQUEEN. These workshops attracted up to 47 participants from our wide, European user base and focussed on the specialized hardware and software of BlueGene/Q.

The expected trend for the near future is a continued increase in complexity for the supercomputers to come. A precursor of this trend is the latest system installed in Jülich, the general purpose cluster JURECA, the "Jülich Research on Exascale Cluster Architectures" system. While it is comprised of standard Intel processors and thus appears to be a off-the-shelf computer, with a high-speed Mellanox EDR InfiniBand network, nodes with varying amounts of memory and additional 75 GPU nodes with 2 NVIDIA K80 accelerators each it is not. The complexity starts with the Intel Xeon chips supporting simultaneous multithreading with up to 48 threads per node, each of them being able to use FMA instructions and wide SIMD vectors. The latter being essential to reach peak performance. Ideally, this can be combined with the available GPUs, adding another way to parallelize codes and increasing the complexity when it comes to coordinating distributed and shared memory parallelization with executing kernels on the accelerators and transferring the necessary data.

The workshop will take place June 6-8., its goal will be making the users aware of the increasing effort to reach peak performance and suggesting possible routes for this task. The topics will include best practices for JURECA, possibilities for visualization, and scientific big data analytics. We will also cover efficient I/O and ways to achieve multi-threading and vectorization. To facilitate GPU programming, we will compare OpenCL, OpenACC and CUDA. Special focus will be put on node-level performance and performance analysis.

Since this is a very broad range of topics, we will only have time to introduce the ideas and highlight their applicability. At the heart of the workshop will be extensive hands-on sessions with the participants' codes, aimed at helping with porting applications to JU-**RECA** and understanding performance bottlenecks. This will be supervised by members of staff from JSC's Simulation Laboratories and cross-sectional teams Application Optimization, Performance Analysis, and Mathematical Methods and Algorithms. At the end of the workshop the participants should have their codes running on JURECA and have a clear picture on how to improve the performance.

contact: Dirk Brömmel, d.broemmel@fz-juelich.de

Dirk Brömmel

Michael M. Resch

23rd Workshop on sustained Simulation Performance

On March 22 and 23, about 60 researchers met for the 23rd time to discuss the sustained performance of simulations on supercomputers. The workshop originated from a collaboration between HLRS and the Cyberscience center of Tohoku University in Sendai, Japan. Meanwhile the workshop brings together a number of organisations among them the Japan Agency for Marine-Earth Science and Technology (JAMSTEC). Since 2005, two workshops are held every year - one in spring in Japan and one in autumn in Stuttgart. The proceedings of these workshops are published at Springer.

The 23rd workshop on sustained simulation performance was giving insight into the planning both of the Japanese government and of HLRS for future HPC systems and into a variety of applications currently running on systems in Stuttgart and Japan.

Katsuyuki Kudo - a representative of MEXT - gave an interesting talk about Japanese HPC policies and about a follow-on system for the K-Computer to be expected for the year 2020. In his talk he pointed out that it is scientific and societal questions that drive the project for the development of an exascale system. This perfectly matched with the talk of Prof. Michael Resch from HLRS who presented ideas on the future development of HPC and their implementation at HLRS. Prof. Vladimir Voevodin from Moscow State University gave an insight into the educational needs for parallel programming - an issue that will become even more important in the future with ever larger numbers of cores in HPC systems.

Prof. Hiroaki Kobayashi from the Cyberscience Center of the Tohoku University gave an impressive talk on the benefits of vector architectures and especially the NEC SX-ACE system. He showed how vector technology outperforms any other architecture by far when it comes to memory bound applications. Prof. Toshimitsu Yokobori from Tohoku University showed how HPC can help to better understand blood vessel behavior and the treatment of diseased blood vessels.

The workshop dinner was held in Japanes style and brought together culture and science in a relaxed mode.

German-Russian Conference Supercomputing in scientific and industrial Problems

From March 9th till March 11th Russian and German scientists met for a conference on Supercomputing in scientific and industrial Problems. The Conference was organized together by the Keldysh Institute of Applied Mathematics of the Russian Academy of Sciences and the High Performance Computing Center Stuttgart (HLRS) of the University of Stuttgart. Prof. Chetverushkin and Prof. Resch were co-chairing the conference which was organized in Moscow. Prof. Chetverushkin giving an impressive talk about algorithmic methods and applied mathematics. The conference showed that there is a wealth of knowledge, expertise and leading edge research both on the theoretical side and the applied side of High Performance Computing which can be fruitfully brought together to make supercomputing applications a toolfor both, scientific and industrial applications.



Eight German scientists joined the conference which was held for the first time and met with about 30 Russian colleagues. In these talks a variety of fields was covered. Mathematical algorithms were one focus of the conference with The conference dinner was one of the highlights of the conference and brought together participants in a friendly atmosphere where future collaborations were discussed. Another conference is planned to be held in Stuttgart in 2017. Activities

• Michael M. Resch

University of Stuttgart, HLRS, Germany

DEEP Project at CeBIT16

The Jülich Supercomputing Centre (JSC) presented the innovative technology developed within the EU research project DEEP [see page 64] at this year's CeBIT. The advantages the system brings to scientific and industrial applications were at display at the shared booth "Innovationsland NRW" in hall 6.

With 3,300 exhibitors from 70 different countries and more than 200,000 attendees in 2016, CeBIT is still the world's largest and most international computer expo. It is considered a barometer of current trends and a measure of the state of the art in information technology. DEEP therefore increased the visibility for its impressive achievements in the fields of hardware, software and application development and greatly benefited from the large and diverse audience at CeBIT.

One important field of research in DEEP, and one of the greatest challenges facing future supercomputers, is the need for reducing the overall energy consumption in HPC systems. The costs of cooling the individual computer components could become completely prohibitive. Because of that DEEP tested two cooling systems: direct water cooling and liquid immersion cooling.

The large DEEP Prototype - a machine with 500 PFLOP/s peak performance operates with direct water cooling and is up and running at JSC. At Hannover, additionally to an electronic board from the DEEP Prototype, the researchers showcased the GreenICE Booster: a smaller prototype that explores an innovative and very efficient immersive cooling system.

In the GreenICE the electronic assemblies are immersed in a special hightech liquid which evaporates even at moderate temperatures. The phase transition from liquid to gaseous maximizes the cooling effect. This means that no waste heat is given off into space and the energy requirements for cooling are cut to about one percent of the overall system consumption.

The GreenICE prototype, in which electronic components seem to "boil", was a crowd puller for the CeBIT audience, who had the chance to learn about this and other technological achievements of the DEEP project.

contact: Estela Suarez, e.suarez@fz-juelich.de

CECAM Tutorial - Atomistic Monte Carlo Simulation

An efficient Tool for studying large Scale conformational **Changes of Biomolecules**

Molecular simulation has become an indispensable tool to study the molecular mechanisms that underlie biological function. For fast processes, such as ligand binding, molecular dynamics (MD) is a popular and effective tool. In contrast, processes that involve large conformational changes like protein folding and peptide aggregation are typically much slower, acting on time scales between milliseconds and seconds. Such time scales are difficult to sample in MD simulations, and alternative techniques circumventing these difficulties should be integrated into scientific workflows to derive the necessary biophysical Atomistic Markov chain insights. Monte Carlo (MC) has been recently demonstrated to be a computationally efficient alternative to MD simulations [1] for such problems. While excellent training courses for atomistic molecular dynamics simulations already exist, students rarely get a useful exposure to Monte Carlo techniques. To fill this gap, a five-day tutorial "Atomistic Monte Carlo Simulations of Biomolecular Systems" funded by CECAM (Centre Européen de Calcul Atomique et Moléculaire) will take place at the Jülich Supercomputing Centre (JSC), September 19-23, 2016.

The goal of the tutorial, organized for the second time by the SimLab Biology of the JSC, is to introduce atomistic Monte Carlo simulations in sufficient detail, that researchers can apply it productively to their own research tasks. As a demonstration tool for the highly transfer-



able MC techniques, the open source protein folding and aggregation package ProFASi, developed at the SimLab Biology, will by used. For realistic tests of advanced parallel simulation techniques like replica exchange MC or Wang-Landau, participants will have access to JURECA, Jülich's recently installed 2 PetaFlop Supercomputer.

Interested researchers can apply and obtain detailed information on the tutorial contents on the web page for the CECAM school:

http://www.cecam.org/ workshop-1339.html

References

1] Mohanty, S., Meinke, J.H., Folding of Top7 in unbiased all-atom Monte Carlo simulations. Proteins 81:1446-1456, 2013

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

Estela Suarez



Figure 1: Estela Suarez (JSC, left) explaining the DEEP project to Minister Svenja Schulze (right) from the Ministery for Innovation, Science and Research of NRW. (Copyright: MIWF/Schweizer)

Sandipan Mohanty

Leaping forward with HPC Education

To receive an in-depth training in HPC architectures and large-scale numerical simulations, fellows of the European Marie Curie training network HPC-LEAP (High Performance Computing in Life sciences, Engineering and Physics) [1] and other students attended a 3-week school at Europe. Jülich Supercomputing Centre (JSC). The goal of the school was not only to improve practical skills of young numerical scientists, but to provide them with good understanding of computer architectures as well as knowledge about today's and future supercomputers.

This European Joint Doctorates (EJD) program started in September 2015 and brings together 16 partners from all over Europe and is coordinated by Cyprus Institute.

Fellows of an EJD are supervised by multiple academic institutions and receive a joint degree. The European Commission has established this new form of training networks to promote international and interdisciplinary collaboration in doctoral training in

HPC-LEAP follows this spirit and brings together applications experts from different research areas, experts for mathematical methods and algorithms as well as experts for HPC architectures and technologies. As disruptive evolution in computer technologies is required for attaining exascale performances, it requires interdisciplinary approaches to enable future computational scientists to use these technologies for their research. The collaborative network



https://hpc-leap.eu/ also includes commercial operators

like HPC solution providers as IBM, NVIDIA and Eurotech as well as numerical service providers as OakLab. The application areas featured in HPC-LEAP include turbulent flows, lattice Quantum Chromodynamics and computational biology.

At the school [2] experts from JSC, other Jülich institutes and from outside held lectures on computer architectures, parallel algorithms, performance analysis, modeling and optimization, MPI and OpenMP programming, GPU programming, visualization, parallel I/O and mathematical libraries. To promote thinking outside the box, the program was augmented by lectures on novel HPC architectures, brain simulation and scalable material research. During various exercises students could practice their new knowledge. The students were challenged by projects, where small teams worked on the parallelization and optimization of different numerical tasks. All students left JSC enthusiastic about the opportunities of using supercomputers for their research.

contact: Marcus Richter, m.richter@fz-juelich.de

References

HPC-LEAP

http://indico-jsc.fz-juelich.de/e/

• Marcus Richter Dirk Pleiter

97

NIC Symposium 2016 at the Forschungszentrum Jülich

The John von Neumann Institute for Computing (NIC) supports research projects from a broad scientific spectrum including topics from Astrophysics, Biology and Biophysics, Chemistry, Elementary Particle Physics, Material Sciences, Condensed Matter, Soft Matter Sciences, Earth and Environment, Computer Sciences and Numerical Mathematics, Fluid Mechanics and Plasma Physics. The NIC symposium is held biennially in February and gives an overview of the activities and results of those projects which received computing time on the supercomputers at the Forschungszentrum Jülich through the NIC. This year the NIC symposium reached a new record of 200 scientists who attended the talks and poster session.

The participants were welcomed by the Forschungszentrum Jülich's Chairman of the Board of Directors Prof. Wolfgang Marquardt and the Director of the Jülich Supercomputing Centre (JSC) Prof. Thomas Lippert. Prof. Marquardt focussed on Big Data and its significance in scientific simulation. He shortly discussed a selection of emerging key challenges. Prof. Lippert expanded on the challenges of Scientific Big Data Analytics (SBDA) in High Performance Computing (HPC). Additionally, he gave an overview of the JURECA system, the successor to the general purpose system JUROPA, which has been greatly accepted by the users of the JSC computer facilities.

Recent results and outcomes were presented in 14 insightful talks and in an overwhelming number of 120 posters. The symposium produced fruitful discussions after the talks and during the poster session and provided plenty of space for the exchange of ideas and experiences in an interdisciplinary scientific environment.

All accompanying materials such as the programme, talks, posters, proceedings, and photographs are available at http://www.john-vonneumann-institut.de/nic/nic-symposium-2016

contact: Alexander Schnurpfeil, a.schnurpfeil@fz-juelich.de



HPC in Science and Engineering - The 18th Results and Review Workshop of the HLRS

The 18th Results and Review Workshop of HLRS had taken place from October 5-6, 2015 at the HLRS in Stuttgart and brought together 64 attendees from German research institutions, the steering committee and scientific support staff of HLRS. During the two day event, the participants were given the opportunity to present and discuss the research projects they had been working on over the last 12 months.

The presentations covered all fields of computational science and engineering ranging from computational fluid dynamics and reacting flows to bioinformatics, chemistry, solid state physics, nanotechnology, physics, and astrophysics to climate research, structural mechanics and earth sciences with a special emphasis on industrially relevant applications. State-of-the-art scientific simulations on leading edge supercomputer technology again emphasized the world-class research done at HLRS obtaining outstanding results in achieving highest performance for production codes which are of particular interest for both scientists and engineers.

These outstanding results of research problems in science and engineering are milestones of modern highest level scientific computing also by using exceptional complex models and methods, and therefore provide an excellent overview of recent developments in high performance computing and simulation techniques.



Out of all, 44 preselected projects have been presented at the HLRS Results and Review Workshop. The steering committee of HLRS, a panel of twelve top-class scientists in their research fields and also responsible for project proposal reviews, appreciated the high quality of the work carried out as well as the spectacular scientific results and the efficient usage of supercomputer resources by awarding the three most outstanding contributions of the workshop, combining the quality of the paper and the lecture. These three papers were awarded the traditional Golden Spike Award of HLRS.

Alexander

Prof. Dr. Dietmar Heinrich Kröner of University of Freiburg, Co-Chairman of the HLRS steering committee, announced the following three research projects as winners of the 2015 The third awardee is Dr. Ana-Catalina Golden Spike Awards:

 Ab initio calculations of the vibrational properties and dynamical processes in semiconductor nanostructures, Prof. Dr. Gabriel Bester, Institut für Physikalische Chemie, Universität Hamburg

• Advances in Parallelization and High-fidelity Simulation of Helicopter Phenomena, Patrick P. Kranzinger, Institut für Aero- und Gasdynamik, Universität Stuttgart

 Large Scale Numerical Simulations of Planetary Interiors, Dr. Ana-Catalina Plesa, Institute of Planetary Research, German Aerospace Center, Berlin.

Two of the three awardees of the 2015 Golden Spike Award, Prof. Dr. Gabriel Bester of University of Hamburg (left), Patrick P. Kranzinger of University of Stuttgart (middle) and the Co-Chair-

man of the HLRS steering committee Prof. Dr. Dietmar H. Kröner of University of Freiburg (right).

Plesa of German Aerospace Center, Berlin.



The Golden Spike Award winners have published a detailed project report in this issue of the inSiDE magazine.



All scientific projects presented at the HLRS Review Workshop are published in the conference proceedings:



High Performance Computing in Science and Engineering '15, Transactions of the High Performance Computing Center, Stuttgart (HLRS) 2015, Wolfgang. E. Nagel, Dietmar H. Kröner, Michael M. Resch (Editors), Springer Cham Heidelberg New York Dordrecht London, © Springer International Publishing Switzerland 2016

ISBN 978-3-319-24631-4 ISBN 978-3-319-24633-8 (eBook) • Heinz Pöhlmann

SPXXL - IBM/Lenovo User **Group Meeting @ LRZ**

SPXXL was founded as a worldwide user group for large-scale, High Performance Computing centers using IBM hardware. With the divestiture of IBM's x86-based server business to Lenovo in 2014, SPXXL dropped the vendor requirement from its bylaws and opened up its membership to any High Performance Computing (HPC) organization. SPXXL now provides advanced content and collaboration among multiple vendors - including IBM, Lenovo, Intel, Mellanox, and others - along with many of the world's largest scientific and technical computing centers.

Unlike other vendor-organized user groups, SPXXL is a self-organized, self-sufficient entity, registered in the United States and the State of California as a 501(c)(6) non-profit corporation. Members and affiliates participate actively in SPXXL meetings and cover their own costs for participating. At the moment SPXXL has 45 member sites. The goal of the organization is to work together with vendors to increase the capabilities and advance the technology of largescale, parallel technical computing hardware and software and to provide quidance to vendors on essential development and support issues for HPC at scale. Some of the areas covered are: Applications, Code Development Tools, Communication, Networking, Parallel I/O, Resource Management, System Administration, Security, and Training. The group addresses topics across a wide range of issues that are important to maximizing performance

and energy efficiency of scientific/ technical computing on highly scalable parallel systems.

The last SPXXL Winter Meeting took place from February 15 - 19 at the Leibniz Supercomputing Centre in Garching near Munich and was focussed around the Lenovo HPC roadmap. It included presentations from Mellanox about the latest Infiniband developments and Intel about their upcoming CPUs and OmniPath interconnect. Member sites also gave presentations on research projects, best practices, and the latest developments in system administration and user support to round out the week-long programme.

With the rotational elections at the end of the meeting Dr. Michael Stephan (JSC) was elected as new president of SPXXL.

References

[1] www.spxxl.org

contact: Michael Stephan, m.stephan@fz-juelich.de

New PATC Course: Introduction to hybrid **Programming in HPC**

Most HPC systems are clusters of shared memory nodes. These SMP nodes can be small multi-core CPUs up to large many-core CPUs. Parallel programming may combine the distributed memory parallelization on the node interconnect (e.g., using MPI) with the shared memory parallelization inside of each node (e.g., using OpenMP or MPI-3.0 shared memory).

As such hybrid programming techniques are getting more and more important in HPC, GCS as one of the 6 European PRACE Advanced Training Centres (PATC) has extended its curriculum by a new PATC course on hybrid programming techniques, which took place at LRZ on January 14, 2016 for the first time. The new course attracted around 40 international participants.

Similar tutorials about hybrid programming have been very successfully presented by the two lecturers Dr. habil. Georg Hager (RRZE, winner of the "Informatics Europe Curriculum Best Practices Award: Parallelism and Concurrency") and Dr. Rolf Rabenseifner (HLRS, member of the steering committee of the MPI-3 Forum) during various supercomputing conferences in the past, but have never been presented as a course with hands-on sessions in a GCS centre before.

The course analysed the strengths and weaknesses of several parallel programming models on clusters of SMP nodes. These models were compared



12, 2017.

References

programming_hpc/

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

Michael Stephan

Figure 1: Participants of the new PATC Course: "Introduction to hybrid programming in HPC" @ LRZ, January 14, 2016

with various hybrid MPI+OpenMP approaches and pure MPI. Numerous case studies and micro-benchmarks were presented to demonstrate the performance-related aspects of hybrid programming. Tools for hybrid programming such as thread/process placement support and performance analysis were presented in a "howto" session. Hands-on exercises gave attendees the opportunity to try the new MPI shared memory interface and explore some pitfalls of hybrid MPI+OpenMP programming.

Due to the great success of the course, it will be repeated at HLRS on June 13, 2016 and at LRZ on January

https://www.lrz.de/services/compute/ courses/x lecturenotes/hybrid

Volker Weinberg

15th HLRS/hww Workshop on Scalable Global Parallel **File Systems**

NTAINE	251	ORA			8801 (SE 410
EFER					NUMERON DE LA COMPANY
h HLRS / hww We lable Global Par	orkshop allel File	on Systems			Êſ
1 111	1		11	1 11	<u> </u>
25-27.2016	4751				#SU3 823
Stuttgart, Germany	4751 B		MR03 827 536 2	*	*
125-27, 2016 , Stuttgart, Germany .hlrs.de			MRU E275362	*	
, Stuttgart, Germany chlrs.de			MRU 827 506 2 4851	*	

From April 25th to April 27th, 2016, representatives from science and industry working in the field of global parallel file systems and high performance storage did meet at HLRS for the fifteenth annual HLRS/hww Workshop on Scalable Global Parallel File Systems "Container Storage". About 70 participants did follow a total of 20 presentations which have been on the workshop agenda.

Prof. Michael Resch. Director HLRS. opened the workshop with an opening address on Monday afternoon.

In the first talk, NEC's Masashi Ikuta, gave an introduction and an overview of the NEC ScaTeFS file system. He showed the file system structure, the tasks of the different components and he provided performance information and an outlook into further developments. The second talk in the first file system session was dedicated to BeeGFS: Franz-Josef Pfreundt, ITWM; showed new features of BeeGFS and BeeOND as well as BeeGFS in the Cloud.

As a representative of the European projects targeting exascale storage and data management, Sai Narasimhamurthy, Seagate, gave an introduction to the SAGE project: Percipient StorAGe for Exascale Data Centric Computing.





In the second Parallel File System session, Willard Davis provided recent updates and an outlook on IBM Spectrum Scale, formerly known as GPFS. In addition, Gabriele Paciucci showed news about the Intel Enterprise Edition for Lustre. Dockerization and Virtualization in the high performance file system and storage field were additional topics covered by Reiner Jung, RedCoolBeans, and Jan Heichler, DDN.

Further presentations from Cray, DDN, EMC, Lenovo and Seagate gave insight into the different hardware and software solutions for high performance data storage including flash tiers and burst buffers. A high performance HSM system named Xtreem-Store was the focus of GrauDatas' Uli Lechner and IBM provided information about the future of tape.

In the traditional session on networks, Matthias Wessendorf, Cisco, talked about issues and solutions concerning real time and lossless requirements in Ethernet fabrics. Energy efficiency in ICT and especially in networking was the topic of Klaus Grobe, Adva.

The last session has been more research oriented. Anastasiia Novikova, Hamburg University, showed Compression of Scientific Data in AIMES. Nico Schlitter, KIT, introduced the Smart Data Innovation Lab and Xuan Wang, HLRS, talked about Optimizing I/O by machine learning. Finally, Thomas Bönisch, also HLRS, provided his view about Challenges and Prospects of NVRAM.

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

contact: Thomas Bönisch, boenisch@hlrs.de

• Thomas Bönisch

Date(s)	Course Title	Location
2016, Jun 27-28	High-performance scientific computing in C++	Jülich
2016, Jun 27-28	Node-Level Performance Engineering (PATC course)	Stuttgart
2016, Jun 27-29	Intel MIC Programming Workshop (PATC course)	Garching
2016, Jun 29-Jul 01	Scientific Workshop: High Performance Computing for Water Related Hazards	Garching
2016, Aug 09-12	Introduction to parallel programming with MPI and OpenMP	Jülich
2016, Sep 05-09	Iterative Linear Solvers and Parallelization	Garching
2016, Sep 12-16	Advanced Fortran Topics (PATC course, tbc.)	Garching
2016, Sep 12-16	Introduction to Computational Fluid Dynamics in HPC	Stuttgart
2016, Sep 26-30	CFD with OpenFOAM®	Stuttgart
2016, Oct 17-21	Parallel Programming Workshop (MPI, OpenMP and advanced topics) (PATC course, tbc.)	Stuttgart
2016, Oct 24-25	Introduction to GPU programming using OpenACC	Jülich
2016, Oct 25-28	Cray XC4O-Workshop on Optimization at Scale	Stuttgart
2016, Nov 03-04	Scientific Visualization	Stuttgart
2016, Nov 24-25	Introduction to the Programming and Usage of the Supercomputer Resources at Jülich	Jülich
2016, Nov 28-30	Advanced Parallel Programming with MPI and OpenMP	Jülich
2016, Dec 01-02	Node-Level Performance Engineering (PATC course, tbc.)	Garching
2016, Dec 05-09	Fortran for Scientific Computing	Stuttgart
2017, Jan 12	Introduction to hybrid programming in HPC (PATC course, tbc.)	Garching

Courses and Tutorials

Introduction to hybrid **Programming in HPC**

Date & Location

June 13, 2016, HLRS, Stuttgart and January 12, 2017, LRZ, Garching near Munich (PATC course, tbc.)

Contents

Most HPC systems are clusters of shared memory nodes. Such SMP nodes can be small multi-core CPUs up to large many-core CPUs. Parallel programming may combine the distributed memory parallelisation on the node interconnect (e.g., with MPI) with the shared memory parallelization inside of each node (e.g., with OpenMP or MPI-3.0 shared memory). This course analyses the strengths and weaknesses of several parallel programming models on clusters of SMP nodes. Tools for hybrid programming such as thread/process placement support and performance analysis are presented in a "how-to" section. Hands-on exercises give attendees the opportunity to try the new MPI shared memory interface and explore some pitfalls of hybrid MPI+OpenMP programming.

Prerequisites

Basic MPI and OpenMP knowledge as presented in the LRZ course "Parallel programming of High Performance Systems".

Webpages

http://www.hlrs.de/ training/2016/HY-S and

http://www.lrz.de/services/com pute/courses/

High Performance scientific Computing in C++

Date & Location

June 27-28, 2016 JSC, Forschungszentrum Jülich

Contents

Modern C++ with its support for procedural, objected oriented, generic and functional programming styles, offers many powerful abstraction mechanisms to express complexity at a high level while remaining very efficient. It is therefore the language of choice for many scientific projects. However, achieving high performance on contemporary computer hardware, with many levels of parallelizm, requires understanding C++ code from a more performance centric viewpoint.

In this course the participants will learn how to write programs in C++ that better utilize the modern processors. The participants will learn how to identify and avoid performance degrading characteristics, such as unnecessary memory operations, branch mis-predictions, and unintentionally strong ordering assumptions. Two powerful open source libraries to help write structured multi-threaded applications will also be introduced: Intel® Threading Building Blocks and NVIDIA Thrust.

Webpage

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

Node-Level Performance Engineering (PATC course)

Date & Location

June 27-28, 2016, HLRS, Stuttgart and December 1-2, 2016, 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.

- Introduction to computer architecture
 - x86 architectures
 - Pipelining, out of order
 - processing
 - Cache hierarchy
 - SIMD, SMT, ccNUMA
- LIKWID tools: topology, pin, perfctr
- Microbenchmarking for architectural exploration - Scalable and bottle

necked resources - OpenMP overhead

- Chip-level performance modeling - Roofline model - Example: vector summation
- Case study: The Jacobi solver - Modeling with the Roofline model
 - Model-guided optimizations
- Case study: sparse matrix-vector multiply
 - Phenomenology - CRS storage format
 - Roofline modeling
- Making best use of parallelizm - Simultaneous multi-threading (SMT) - ccNUMA
 - Single Instruction multiple Data (SIMD) - Examples
- Further case studies (optional)
- Energy & Parallel Scalability - Energy consumption of modern processors - The energy-to-solution
 - metric
 - Performance engineering == power engineering
- Practical performance engineering with patterns

Prerequisites

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

Webpages

http://www.hlrs.de/training /2016/NLP and

http://www.lrz.de/services/com pute/courses/

Intel MIC Programming Workshop

(PATC course)

Date & Location June 27-29, 2016

LRZ, Garching near Munich

Contents

The course discusses Intel's new Many Integrated Core (MIC) architecture. It covers various programming and optimization techniques for Intel Xeon Phi coprocessors. The first 2 days provide an introduction about the Intel MIC architecture and various Intel Xeon Phi programming models, interleaved with many hands-on sessions. The hands-on sessions are done on the SuperMIC system at LRZ. The last day presents advanced topics and talks about experiences and best practices using Intel Xeon Phi based systems like e.g. Salomon @ IT4Innovations (Czech Republic), the largest Intel Xeon Phi based system in Europe.

The course material is developed within PRACE and the joint German-Czech project CzeBaCCA. A two-day scientific workshop on "High Performance Computing for Water related Hazards" of this project is taking place at LRZ directly after this course, on June 29 - July 1, 2016.

Prerequisites

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

Webpage

http://www.lrz.de/services/com pute/courses/

Scientific Workshop: High Performance Computing for Water Related Hazards

Date & Location

June 29-July 1, 2016 LRZ, Garching near Munich

Contents

The Czech-Bavarian Competence Team for Supercomputing Applications invites all interested researchers to a workshop on large-scale simulation of water-related hazard scenarios. Contributed presentations are welcome (for Jun 30) on all related topics, in particular on demands and desired features of (future) simulation software, parallelisation for current and novel HPC platforms, as well as establishing scalable simulation workflows on supercomputing environments. Simulation scenarios may include rain-induced floods, storm surges, tsunamis, and similar events Wednesday: Plenum session with invited talks on MIC experience and best practice recommendations. Thursday: Contributed talks on large-scale simulation of water-related hazard scenarios. Friday: Special session on "Floods in Bavaria".

Prerequisites

This workshop is targeted at PhD students, postdocs and researchers interested in water related hazards.

Webpage

http://www.lrz.de/services/com pute/courses/

Courses and Tutorials

Introduction to Parallel **Programming with MPI** and OpenMP

Date & Location August 9-12, 2016 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

Compact Course: Iterative Linear Solvers and Parallelization

Date & Location

September 5-9, 2016 LRZ, Garching near Munich

Contents

The focus of this compact course is on iterative and parallel solvers, the parallel programming models MPI and OpenMP, and the parallel middleware PETSc.

Different modern Krylov Subspace Methods (CG, GMRES, BiCGSTAB ...) as well as highly efficient preconditioning techniques are presented in the context of real life applications. Hands-on sessions (in C and Fortran) will allow users to immediately test and understand the

- basic constructs of iterative solvers
- the Message Passing Interface (MPI),
- the shared memory directives of OpenMP.

The course is organized by LRZ, the University of Kassel, the High Performance Computing Centre of Stuttgart (HLRS) and IAG.

Prerequisites

Participants must have basic knowledge in linear algebra and programming under Unix with C, C++ or Fortran.

Webpage

http://www.lrz.de/services/com pute/courses/

Advanced Fortran Topics

(PATC course, tbc.)

Date & Location

September 12-16, 2016 LRZ, Garching near Munich

Contents

This course, partly a PRACE Advanced Training Center (PATC) course (to be confirmed), is targeted at scientists who wish to extend their knowledge of Fortran to cover advanced features of the language. Days 1-3

- Best Practices
- Object-Oriented Programming
- IEEE features and floating point exceptions
- Interoperability with C

 Fortran 2003 I/O extensions. Days 4-5 (PATC course, support by PRACE still has to be approved)

- OO Design Patterns: application of object-oriented programming
- PGAS concepts and Coarrays.

Prerequisites

Course participants should have good knowledge of the Fortran 95 standard (such as covered in the February course at LRZ), and also basic UNIX/Linux knowledge.

Webpage

http://www.lrz.de/services/com pute/courses/

Introduction to Computional Fluid Dynamics in **High Performance** Computing

Date & Location

September 12-16, 2016 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/training /2016/CFD-S

CFD with OpenFOAM®

Date & Location September 26-30, 2016 HLRS, Stuttgart

Contents

open-source code and a powerful framework for solving a variety of problems mainly in the field of CFD. The five-day workshop gives an introduction to OpenFOAM® applied on CFD phenomena and is intended for beginners and people with CFD background knowledge. The user will learn about case setup, meshing tools and available OpenFOAM® utilities that can be used for pre- and postprocessing tasks as well as job monitoring during runtime.

OpenFOAM® is a widely-used

Additionally, basic solvers and major aspects of code structure are highlighted. Lectures and hands-on session with examples will guide through this course

Webpage

http://www.hlrs.de/training /2016/0F2

Parallel Programming Workshop

(PATC course, tbc.)

Date & Location October 17-21, 2016 HLRS. Stuttgart

Contents

Distributed memory parallelization with the Message Passing Interface MPI (Mon, for beginners): On clusters and distributed memory

architectures, parallel programming with the Message Passing Interface (MPI) is the dominating programming model. This course-day gives an introduction into MPI.

Shared memory parallelization with OpenMP (Tue, for beginners): The focus is on shared memory parallelization with OpenMP, the key concept on hyper-threading, dualcore, multi-core, shared memory, and ccNUMA platforms. Racecondition debugging tools are also nresented

Advanced topics in parallel programming (Wed-Fri): The course will teach newest methods in MPI-3.0/3.1 and OpenMP-4.0/4.5, which were developed for the efficient use of current HPC hardware. Topics with MPI are the group and communicator concept, process topologies, derived data types, the new MPI-3.0 Fortran language binding, one-sided communication and the new MPI-3.0 shared memory programming model within MPI. Topics with OpenMP are the OpenMP-4.0 extensions, as the vectorization directives, thread affinity and OpenMP places. (The GPU programming with OpenMP-4.0 di-

Courses and Tutorials

rectives is not part of this course.) Further aspects are domain decomposition, load balancing, and debugging. The course also contains performance and best practice considerations, e.g., with hybrid MPI+OpenMP parallelization.

Hands-on sessions (in C and

Fortran) will allow users to immedi-

ately test and understand the taught

constructs of the Message Pass-

ing Interface (MPI) and the shared

memory directives of OpenMP. This

course provides scientific training in

Computational Science, and in addi-

tion, the scientific exchange of the

participants among themselves.

http://www.hlrs.de/training

Webpage

/2016/PAR

Contents

Date & Location

October 24-25, 2016

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

Introduction to GPU Programming using OpenACC

JSC, Forschungszentrum Jülich

GPU-accelerated computing drives

Cray XC40-Workshop on **Optimization at Scale**

Date & Location

October 25-28, 2016 HLRS, Stuttgart

Contents

In order to help users running applications efficiently on this new large system, HLRS and Cray offer an optimization workshop for applications running at scale. The intent of this workshop is to tune the performance of such codes by detecting, locating and solving bottlenecks.

Together with specialists from Cray, the users will enhance the efficiency of their code. In order to do so, users will be required to bring their codes to the workshop ("bring your own code", "BYOC"). The specialists from Cray will assist in profiling application runs, analyzing the profile and optimizing the code based on the analysis. All categories of potential bottlenecks (CPU, memory subsystem, communication and I/O) will be addressed.

In addition to working on the code, there will be lectures on:

- profiling tools
- node-level performance tuning
- I/O (MPI-IO, NetCDF, HDF5)

Webpage

http://www.hlrs.de/training /2016/XC40-2

Scientific Visualization

Date & Location

November 3-4, 2016 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.

Prerequisites

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

Webpage

http://www.hlrs.de/training /2016/VIS2

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

Date & Location

November 24-25, 2016 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 28-30, 2016

JSC, Forschungszentrum Jülich

Contents

The focus is on advanced programming with MPI and OpenMP. The course addresses participants who have already some experience with C/C++ or Fortran and MPI and OpenMP, the most popular programming models in High Performance Computing (HPC). It will teach newest methods in MPI-3.0/3.1 and OpenMP-4.5, which were developed for the efficient use of current HPC hardware. Topics with MPI are the group and communicator concept, process topologies, derived data types, the new MPI-3.0 Fortran language binding, one-sided communication and the new MPI-3.0 shared memory programming model within MPI. Topics with OpenMP are the OpenMP-4.0 extensions, as the vectorization directives, thread affinity and OpenMP places. The course also contains performance and best practice considerations, e.g., with hybrid MPI+OpenMP parallelization. The course ends with a section presenting tools for parallel programming. The course is given in 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

Courses and Tutorials

Fortran for Scientific Computing

Date & Location

December 5-9, 2016 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/training /2016/FTN2

