Jülich Supercomputing Centre contributes to visionary "Human Brain Project"

HBP
The Human Brain Project

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

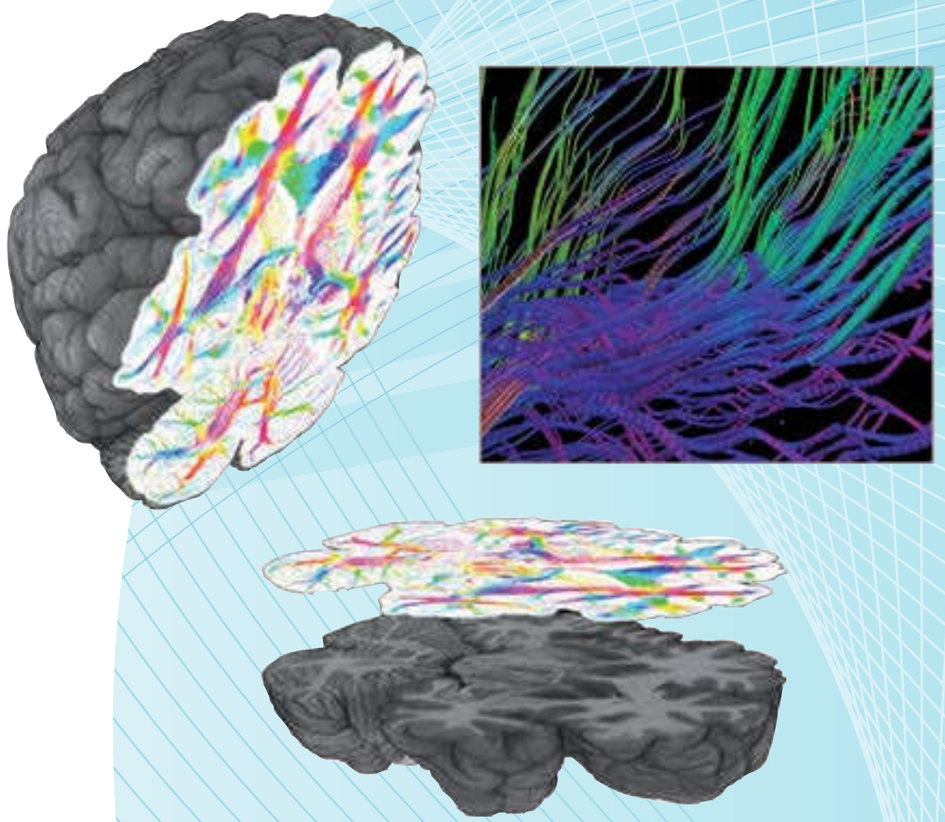

Figure 1: 3D reconstructed block face volume of a human brain with corresponding 3D fiber orientation maps obtained with Polarized Light Imaging (PLI). The pane in the upper right shows tracked fiber tubes. (Source: Axer, Amunts et al., INM, Forschungszentrum Jülich)

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

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

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

References

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

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

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

• Thomas Lippert
• Boris Orth

Jülich Supercomputing Centre (JSC)

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

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

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

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

Reference

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

• Maik Boltes
• Armin Seyfried

Jülich Supercomputing Centre (JSC)

UNICORE Summit 2013

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

Figure 1: Participants of the UNICORE Summit 2013.

Figure 2: Keynote by Michel Drescher.

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

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

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

References

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

• Valentina Huber
• Daniel Mallmann

Jülich Supercomputing Centre (JSC)

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

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

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

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

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

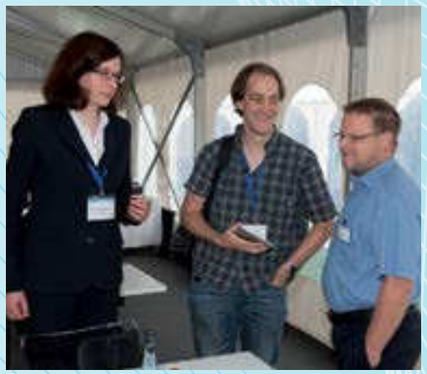

Figure 2: The project team.

Figure 3: Virtual Reality presentation on the factory floor.

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

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

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

Links

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

- Bärbel Große-Wöhrmann
- Uwe Wössner

University of Stuttgart, HLRS

The 17th HLRS-NEC Workshop on Sustained Simulation Performance

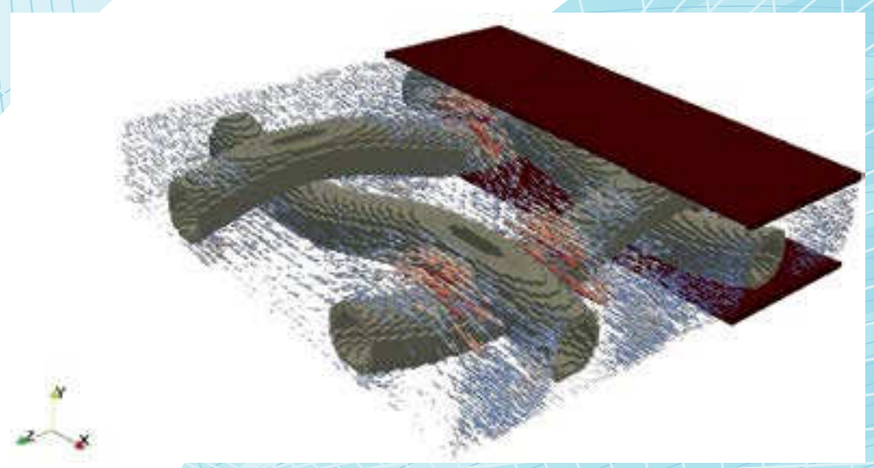

Figure 1: Simulation of a spacer.

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

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

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

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

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

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

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

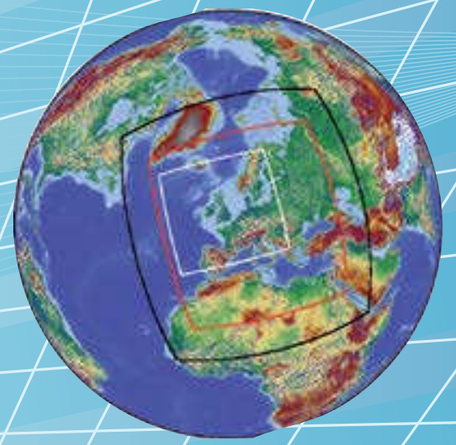

Figure 2: Illustration of the computational domain of various regional climate models with varying resolution. (Courtesy Dr. Warrach-Sagi, University Hohenheim)

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

• José Gracia
• Ralf Schneider

University of Stuttgart, HLRS

ls1 mardyn - a Massively Parallel Molecular Simulation Code

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

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

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

Neighbour Search

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

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

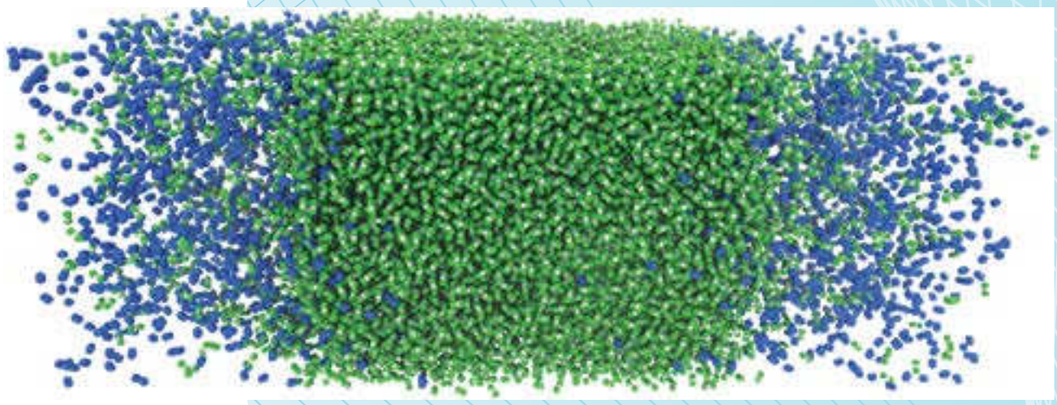

Molecular simulation of vapour-liquid phase boundaries using ls1 mardyn. Green: CO2, Blue: Oxygen, Temperature: -20° C.

Dynamic Load Balancing

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

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

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

Scalability

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

Partners

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

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

For more Informations:

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

University of Stuttgart, HLRS

Activities

GCS at ISC'13 – Review

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

The Gauss Centre for Supercomputing participated in the 2013 edition of ISC, held in Leipzig/Germany, yet again with a 64 m² booth on which its three member centres HLRS, LRZ and JSC showcased their latest research projects and presented sophisticated HPC tools. Countless like-minded HPC users, researchers, technology leaders, scientists, IT-decision makers as well as high tech media representatives stopped by to meet and talk with the directors of the three GCS centres Prof. Bode (LRZ), Prof. Lippert (JSC), the newly elected GCS Chairman of the Board, Prof. Resch of HLRS, and GCS Managing Director Dr. Müller, as well as with many other scientists of the three GCS centres.

Two GCS HPC systems amongst Top Ten of TOP500

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

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

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

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

GCS Booth Highlights

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

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

ISC'13 Gauss Award Winner

Each year at ISC, GCS presents the Gauss Award to recognize the most outstanding paper in the field of scalable supercomputing from all papers accepted for the ISC'13 Research Paper Sessions. This year, the award honored the paper titled „TUE, a new energy-efficiency metric applied at ORNL's Jaguar". Dr. Michael K. Patterson,

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

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

• Regina Weigand

Gauss Centre for Supercomputing

Extreme Scaling Workshop at LRZ

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

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

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

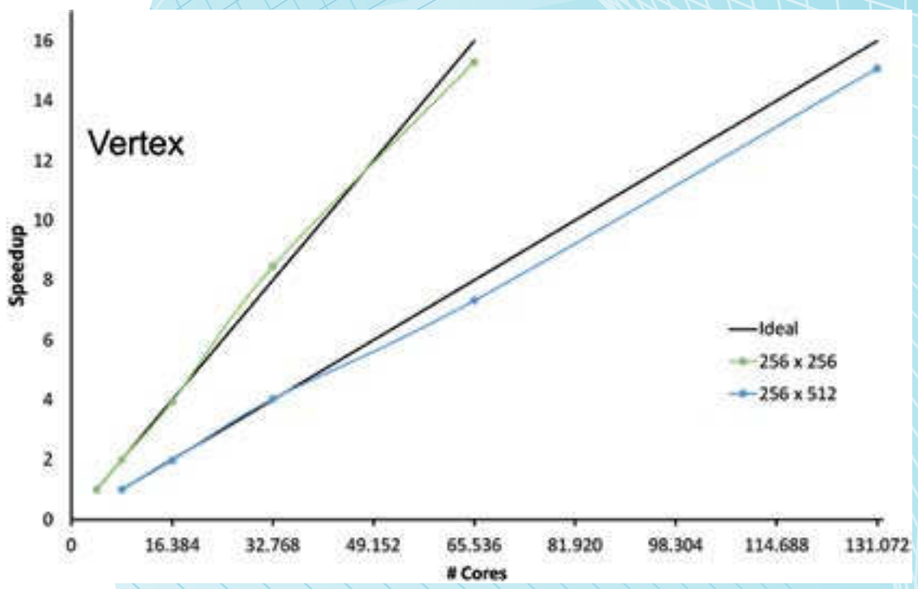

Figure 1: Scaling Plot for Vertex.

• BQCD (16 islands)

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

• CIAO (8 islands)

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

• P-Gadget3-XXL (16 islands)

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

• GROMACS (8 islands)

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

• LAMMPS (16 islands)

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

• Nyx (16 islands)

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

• Vertex (16 islands)

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

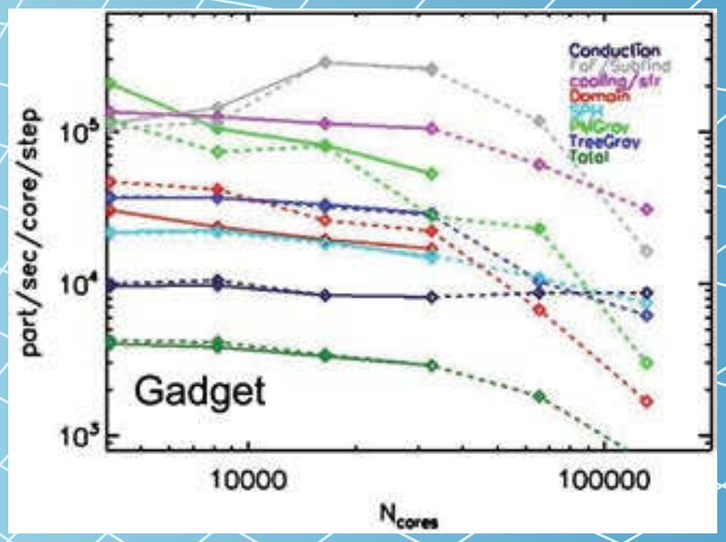

Figure 2: Scaling Plot for P-Gadget3-XXL.

• walBErla (16 islands)

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

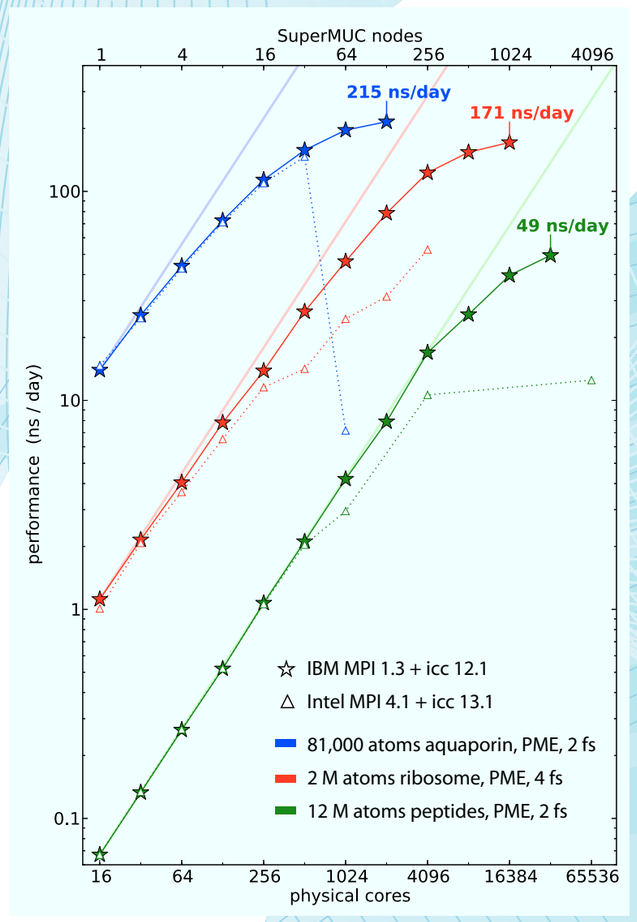

Figure 3: Scaling Plot for GROMACS.

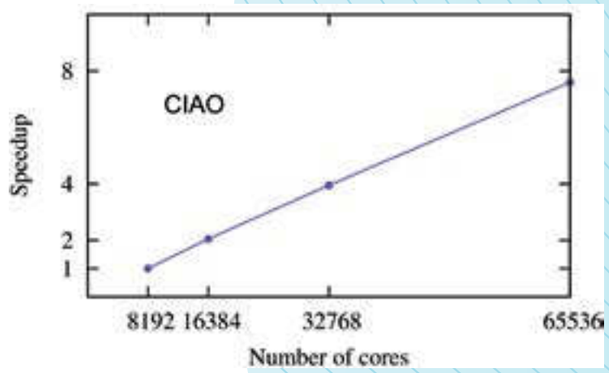

Figure 4: Scaling Plot for CIAO.

• APES (8 islands)

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

• SeisSol (8 islands)

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

• ExaML (4 islands)

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

• ICON (4 islands)

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

Performance Results

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

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

LRZ Extreme Scale Benchmark and Optimization Suite

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

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

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

[1] Leibniz Supercomputing Center

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

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

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

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

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

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

HLRS Scientific Tutorials and Workshop Report and Outlook

• Rolf Rabenseifner

University of Stuttgart, HLRS

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

ISC and SC Tutorials

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

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

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

Programming of Cray XK7 clusters with GPUs is taught in **OpenACC Programming for Parallel Accelerated Supercomputers – an alternative to CUDA from Cray perspective** on April 10 - 11, 2014.

These Cray XE6/XC30 and XK7 courses are also presented to the European community in the framework of the **PRACE Advanced Training Centre** (PATC). GCS, i.e., HLRS, LRZ and the Jülich Supercomputer Centre together, serve as one of the first six PATCs in Europe.

One of the flagships of our courses is the week on **Iterative Solvers and Parallelization**. Prof. A. Meister teaches basics and details on Krylov Subspace Methods. Lecturers from HLRS give lessons on distributed memory parallelization with the Message Passing Interface (MPI) and shared memory multi-threading with OpenMP. This course will be presented twice, in March 2014 at HLRS in Stuttgart and September 2014 at LRZ.

Another highlight is the **Introduction to Computational Fluid Dynamics**. This course was initiated at HLRS by Dr.-Ing. Sabine Roller. She is now a professor at the University of Siegen. It is again scheduled in spring 2014 in Stuttgart and in September/October in Siegen. The emphasis is placed on explicit finite volume methods for the compressible Euler equations. Moreover outlooks on implicit methods, the extension to the Navier-Stokes equations and turbulence modeling are given. Additional topics are classical numerical methods for the solution of the incompressible Navier-Stokes equations, aeroacoustics and high order numerical methods for the solution of systems of partial differential equations.

Our general course on parallelization, the **Parallel Programming Workshop**, October 6 - 10, 2014 at HLRS, will have three parts: The first two days of this course are dedicated to parallelization with the Message passing interface (MPI). Shared memory multi-threading is taught on the third day, and in the last two days, advanced topics are discussed. This includes MPI-2 functionality, e.g., parallel file I/O and hybrid MPI+OpenMP, as well as the upcoming MPI-3.0. As in all courses, hands-on sessions (in C and Fortran) will allow users to immediately test and understand the parallelization methods. The course language is English.

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

We also continue our series of **Fortran for Scientific Computing** in December 2013 and in March 2014, mainly visited by PhD students from Stuttgart and other universities to learn not only the basics of programming, but also to get an insight on the principles of developing high-performance applications with Fortran.

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

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

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

2014 – Workshop Announcements
Scientific Conferences and Workshops at HLRS
12th HLRS/hww Workshop on Scalable Global Parallel File Systems (March/April 2014)
8th ZIH+HLRS Parallel Tools Workshop (date and location not yet fixed)
High Performance Computing in Science and Engineering - The 17th Results and Review Workshop of the HPC Center Stuttgart (October 2014)
IDC International HPC User Forum (October 2014)
Parallel Programming Workshops: Training in Parallel Programming and CFD
Parallel Programming and Parallel Tools (TU Dresden, ZIH, February 24 - 27)
Introduction to Computational Fluid Dynamics (HLRS, March 31 - April 4)
Iterative Linear Solvers and Parallelization (HLRS, March 24-28)
Cray XE6/XC30 Optimization Workshops (HLRS, March 17 - 20) (PATC)
GPU Programming using CUDA (HLRS, April 7 - 9)
Open ACC Programming for Parallel Accelerated Supercomputers (HLRS, April 10 - 11) (PATC)
Unified Parallel C (UPC) and Co-Array Fortran (CAF) (HLRS, April 14 - 15) (PATC)
Scientific Visualisation (HLRS, April 16 - 17)
Parallel Programming with MPI & OpenMP (TU Hamburg-Harburg, July 28 - 30)
Iterative Linear Solvers and Parallelization (LRZ, Garching, September 15 - 19)
Introduction to Computational Fluid Dynamics (ZIMT Siegen, September/October)
Message Passing Interface (MPI) for Beginners (HLRS, October 6 - 7) (PATC)
Shared Memory Parallelization with OpenMP (HLRS, October 8) (PATC)
Advanced Topics in Parallel Programming (HLRS, October 9 - 10) (PATC)
Parallel Programming with MPI & OpenMP (FZ Jülich, JSC, December 1 - 3)
Training in Programming Languages at HLRS
Fortran for Scientific Computing (Dec 2 - 6, 2013 and Mar 10 - 14, 2014) (PATC)
URLs:
http://www.hlrs.de/events/
http://www.hlrs.de/training/course-list/
(PATC): This is a PRACE PATC course

GCS – High Performance Computing Courses and Tutorials

Parallel Programming with MPI, OpenMP and PETSc

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

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

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

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

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

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

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

Node-Level Performance Engineering

(PATC course)

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

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

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

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

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

Fortran for Scientific Computing

(PATC course)

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

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

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

Second JUQUEEN Porting and Tuning Workshop

(PATC course)

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

Contents
The Blue Gene/Q petaflop supercomputer JUQUEEN marks another quantum leap in supercomputer performance at JSC. In order to use this tool efficiently, special efforts by the users are necessary, though. The aim of this hands-on workshop is to support current users of JUQUEEN in porting their software, in analyzing its performance, and in improving its efficiency. This course is a PATC course (PRACE Advanced Training Centres).

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

Programming with Fortran

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

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

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

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

Parallel Programming with MPI, OpenMP, and Tools

Date and Location
February 24 - 27, 2014
Dresden, ZIH

Contents
The focus is on programming models MPI, OpenMP, and PETSc. Hands-on sessions (in C and Fortran) will allow users to immediately test and understand the basic constructs of the Message Passing Interface (MPI) and the shared memory directives of OpenMP. The last day is dedicated to tools for debugging and performance analysis of parallel applications. This course is organized by ZIH in collaboration with HLRS.

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

Parallel Programming of High Performance Systems

Dates and Location
March 10 - 14, 2014
RRZE building, University campus Erlangen, Martensstr. 1: Via video conference at LRZ if there is sufficient interest.

Contents
This course, a collaboration of Erlangen Regional Computing Centre (RRZE) and LRZ, is targeted at students and scientists with interest in programming modern HPC hardware,

specifically the large scale parallel computing systems available in Munich, Jülich and Stuttgart.

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

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

Cray XE6/XC 30 Optimization Workshop

(PATC course)

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

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

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

Iterative Linear Solvers and Parallelization

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

September 15 - 19, 2014
Garching, LRZ

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

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

Eclipse: C/C++/Fortran programming

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

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

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

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

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

Introduction to Computational Fluids Dynamics

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

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

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

Advanced Topics in High Performance Computing

(PATC course)

Date and Location
March 31 - April 03, 2014
LRZ Building, University campus Garching near Munich.

Contents
In this add-on course to the parallel programming course special topics are treated in more depth, in particular performance analysis, I/O and PGAS concepts. It is provided in collaboration of Erlangen Regional Computing Centre (RRZE) and LRZ within KONWIHR.
Each day is comprised of approximately 5 hours of lectures and 2 hours of hands-on sessions.
Day 1
Intel tools: MPI tracing and Checking
Intel tools: OpenMP performance and correctness.
Day 2
Parallel I/O with MPI IO
Performance analysis with Scalasca.
Day 3
Tuning I/O on LRZ's HPC systems.
Portability of I/O: Binary files NetCDF HDF5.
Day 4
PGAS programming with coarray Fortran and Unified Parallel C.
PGAS hands on session.

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

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

GPU Programming using CUDA

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

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

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

GPU Programming

(PATC course)

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

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

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

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

Open ACC Programming Cray XK

(PATC course)

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

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

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

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

(PATC course)

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

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

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

Scientific Visualization

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

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

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

Intel MIC&GPU Programming Workshop

(PATC course)

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

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

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

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

Advanced GPU Programming

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

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

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

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

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

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

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

Parallel I/O and Portable Data Formats

(PATC course)

Date and Location
May 21 - 23, 2014
JSC, Forschungszentrum Jülich

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

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

Authors

Momme Allalen
momme.allalen@lrz.de
Fakher F. Assaad
assaad@physik.uni-wuerzburg.de
Norbert Attig
n.attig@fz-juelich.de
Christian Baczynski
baczynski@uni-hd.de
Gurvan Bazin
gurvan.bazin@gmail.com
Florian Berberich
f.berberich@fz-juelich.de
Christoph Bernau
christoph.bernau@lrz.de
Arndt Bode
arndt.bode@lrz.de
Maik Boltes
m.boltes@fz-juelich.de
David Brayford
david.brayford@lrz.de
Matthias Brehm
matthias.brehm@lrz.de
Dirk Brömmel
d.broemmel@fz-juelich.de
Alexandru Calotoiu
a.calotoiu@grs-sim.de
Alexey Cheptsov
cheptsov@hlrs.de
Paul Clark
p.clark@uni-heidelberg.de
Tanja Cleese
tanja.clees@scai.fraunhofer.de
Holger Dachsel
h.dachsel@fz-juelich.de
Ulrich Detert
u.detert@fz-juelich.de
Klaus Dolag
kdolag@mpa-garching.mpg.de
Jan Frederik Engels
mail@jfengels.de
Dietmar Erwin
d.erwin@fz-juelich.de
Jochen Fröhlich
jochen.froehlich@tu-dresden.de
Andrea Gatto
andreag@mpa-garching.mpg.de
Paul Gibbon
p.gibbon@fz-juelich.de
Philipp Girichidis
philipp@girichidis.com
Colin Glass
glass@hlrs.de
Simon Glover
glover@uni-heidelberg.de
Stefan Gottlöber
sgottloeber@aip.de
Jose Gracia
gracia@hlrs.de
Bärbel Große-Wöhrmann
woehrmann@hlrs.de
Carla Guillen Carias
carla.guillen@lrz.de
Nicolay Hammer
nicolay.hammer@lrz.de
Tim Hender
tim.hender@ccfe.ac.uk
Wolfram Hesse
wolfram.hesse@lrz.de
Steffen Heß
shess@aip.de
Torsten Hoefler
htor@inf.ethz.ch
Martin Hohenadler
martin.hohenadler@
physik.uni-wuerzburg.de
Stefan Holl
st.holl@fz-juelich.de
Frank Holzäpfel
frank.Holzaepfel@dlr.de
Valentina Huber
v.huber@fz-juelich.de
Ferdinand Jamitzky
ferdinand.jamitzky@lrz.de
Ivo Kabadshow
i.kabadshow@fz-juelich.de
Anupam Karmakar
anupam.karmakar@lrz.de
Tobias Kempe
tobias.kempe@tu-dresden.de
Francisco Kitaura
kitaura@aip.de
Ralf Klessen
klessen@uni-heidelberg.de
Bastian Koller
koller@hlrs.de
Carsten Kutzner
ckutzne@gwdg.de
Thomas Lippert
th.lippert@fz-juelich.de
Daniel Mallmann
d.mallmann@fz-juelich.de
Andreas Marek
amarek@rzg.mpg.de
Gerald Mathias
gerald.mathias@
physik.uni-muenchen.de
Bernd Mohr
b.mohr@fz-juelich.de
Thorsten Naab
naab@mpa-garching.mpg.de
Carmen Navarrete
carmen.navarrete@lrz.de
Christoph Niethammer
niethammer@hlrs.de
Boris Orth
b.orth@fz-juelich.de
Ludger Palm
ludger.palm@lrz.de
Thomas Peters
tpeters@physik.uzh.ch
Rolf Rabenseifner
rabenseifner@hlrs.de
Michael M. Resch
resch@hlrs.de
Helmut Satzger
helmut.satzger@lrz.de
Armin Seyfried
a.seyfried@fz-juelich.de
Wolfram Schmidt
schmidt@astro.physik.
uni-goettingen.de
Magnus Schwörer
magnus.schwoerer@
googlemail.com
Anton Stephan
anton.stephan@dlr.de
Sven Strohmer
s strohmer@fz-juelich.de
Godehard Sutmann
g.sutmann@fz-juelich.de
Jan Treibig
jan.treibig@rrze.fau.de
Philipp Trisjono
p.trisjono@itv.rwth-aachen.de
Thomas Ullmann
thomas.ullmann@mpibpc.de
Bernhard Vowinckel
bernhard.vowinckel@tu-dresden.de
Stefanie Walch
walch@mpa-garching.mpg.de
Regina Weigand
weigand@hlrs.de
Volker Weinberg
volker.weinberg@lrz.de
Torsten Wilde
torsten.wilde@lrz.de
Felix Wolf
f.wolf@grs-sim.de
Klaus Wolkersdorfer
k.wolkersdorfer@fz-juelich.de
Richard Wünsch
richard@wunsch.cz

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