EXASCALE

Newsletter of Forschungszentrum Jülich on Supercomputing

JÜLICH FORSCHUNGSZENTRUM

Issue: 1 | March 2014

Member of the Helmholtz Association

Getting to Grips with Big Data

"The flood of information is growing. Our research helps us to quickly and efficiently gain useful information from big data."

Progress in supercomputing facilitates increasingly more complex applications. This is important in achieving new insights. But more realistic simulations have their price: swelling volumes of data that must be dealt with, for example to create simulations in climate research, basic physical research, or the neurosciences. With the Human Brain Project, which aims to simulate the human brain, we are treading new ground in terms of the expected volumes of data. The problem associated with big data is that storing and transferring large volumes of data takes time and restricts applications. We need approaches that allow us to use data quickly and efficiently. This means filtering out the most important things, perhaps even combining different sources, and then rapidly evaluating them. This is a challenge not only facing supercomputing but also information and communications technology in general.

We are working closely with partners from science and industry in order to find solutions. Together with IBM experts, our researchers in the Exascale Innovation Center have developed a new storage concept: Blue Gene Active Storage (see p. 2). It enables supercomputers to process large data sets and compress them before they are written to storage media. In this way, data-intensive tasks can be processed much faster. To ensure that industry profits directly from our research results, Forschungszentrum Jülich is involved in the new Smart Data Innovation Lab. This project brings companies and research institutions together so that large data volumes from industrial processes can be used more efficiently in the future (see p. 4). Both projects are important steps forward in gaining useful information from big data.

Prof. Achim Bachem
Chairman of the Board of Directors of Forschungszentrum Jülich

IN THIS ISSUE

PAGE 2:
Using Virtual Reality to Improve Weather Forecasts

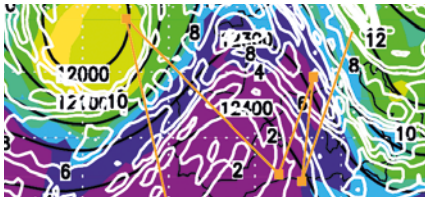

PAGE 3:
Turbulence: Chaos in the Sky

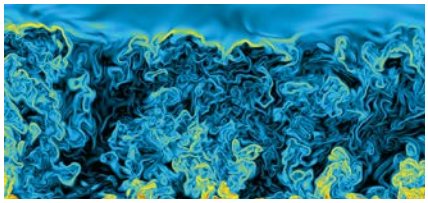

PAGE 4:
News in Brief
Upcoming Events
Publication Details

www.fz-juelich.de/ias/jsc/EN

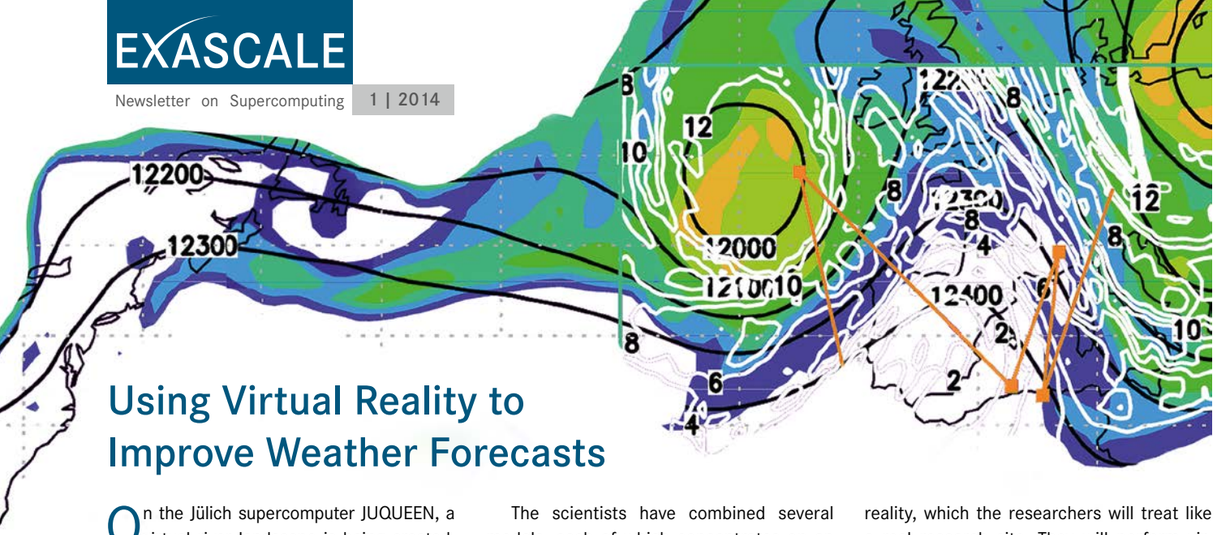

Using Virtual Reality to Improve Weather Forecasts

On the Jülich supercomputer JUQUEEN, a virtual river landscape is being created. A research unit funded by the German Research Foundation (DFG) is simulating the Neckar catchment area measuring some 150 kilometres by 200 kilometres to learn more about water and energy flows. The researchers hope to use the knowledge they gain to improve weather and flood forecasts. The DFG have granted the project entitled Data Assimilation for Improved Characterization of Fluxes across Compartmental Interfaces funding worth around € 2.1 million over the next three years.

The scientists have combined several models, each of which concentrates on an individual aspect like hydrology or the atmosphere, to form one single model. “This allows us to incorporate interactions that have been neglected in the past, for example those between groundwater and the atmosphere. And this helps us to understand more about mass flows,” says co-coordinator Prof. Harrie-Jan Hendricks-Franssen from the Jülich Institute of Bio- and Geosciences. The calculations are based on empirical data from the Neckar catchment area. Models and data are used to create a virtual reality, which the researchers will treat like a real research site. They will perform virtual measurements on parameters such as evaporation and soil moisture. The results of these measurements will then be utilized by the scientists to continuously refine models and forecasts. Data assimilation is what the researchers call this process. At the same time, the project partners hope to infer which measurements are important for improving a certain forecast – for example, precipitation measurements for forecasts on groundwater recharge.

Researchers simulate a river catchment area with the aim of learning more about energy and water flows and optimizing weather forecasts.

www.fz-juelich.de/SharedDocs/Pressemitteilungen/UK/EN/2013/13-12-02-dfg.html

Intelligent Storage Booster for Large Volumes of Data

Thanks to a new storage concept – Blue Gene Active Storage (BGAS) – supercomputers can now process data-intensive tasks even faster. IBM and Forschungszentrum Jülich presented the first-ever installation of the BGAS system in November 2013 at the world’s largest supercomputing conference SC13 in Denver, USA. The system runs on the supercomputer JUQUEEN at the Jülich Supercomputing Centre (JSC). The installation is set to benefit highly complex simulations, particularly those in the field of brain research.

Such applications generate enormous volumes of data. Saving these data on large hard drive systems outside the high-performance cluster has a serious disadvantage: users only have limited access to them. This is due to the slow transfer rate of such systems. “Highly complex simulations such as in brain research generate large volumes of data that can no longer be processed by conventional methods. New I/O concepts are therefore one of the most important steps towards achieving new research results using high-performance computers in these areas,” says director of JSC, Prof. Thomas Lippert.

The concept of “active storage” was developed by Jülich scientists and IBM experts in the jointly run Exascale Innovation Center (EIC). BGAS has its own processors, which enables much faster access. This allows the system to process large data sets and compress them before they are written to the storage media. In addition, extremely fast, non-volatile flash memory is also used.

With the Blue Gene Active Storage System, the Jülich supercomputer JUQUEEN can process data-intensive tasks even faster. Brain research will be the first field to benefit from this.

www.fz-juelich.de/SharedDocs/Pressemitteilungen/UK/EN/2013/13-11-18bgas.html

Turbulence: Chaos in the Sky

Almost every aeroplane passenger has experienced it at some time or other. In the middle of ascent or descent, the plane suddenly starts to vibrate. This is caused by the turbulent movement of air masses. Such turbulence is most common when the aeroplane enters the Earth’s atmosphere from the relatively calm troposphere above it. Dr. Juan Pedro Mellado from the Max Planck Institute for Meteorology in Hamburg hopes to gain new insights into this phenomenon using simulations on Jülich’s supercomputers. The John von Neumann Institute for Computing selected his project on the direct numerical simulation of turbulent mixing processes in the planetary boundary layer as a John von Neumann Excellence Project 2013.

The planetary boundary layer extends to an altitude of around 2,000 metres. In this transition zone between the Earth’s surface and the atmosphere, air masses are continuously in motion. Other factors such as mountains or buildings also affect the air flows. If a flow is strongly disturbed, different masses of air mix, causing turbulence. This is also important for the climate. The boundary layer is where important exchange processes occur, for example heat transfer and water exchange. The problem is that turbulence develops chaotically and evolves rapidly. “Initial progress in supercomputing made it possible in recent years to study mixing processes in more detail and to learn more about their underlying principles,” says Mellado. In his project, he aims to ascertain what effect different surfaces, such as an ice-covered or ice-free ocean, have on turbulence in the atmospheric layers above them. He is also investigating interactions between clouds and turbulence. These complex relationships have not yet been sufficiently accounted for in climate models.

In the atmosphere near the surface of the Earth (lower half of image), air masses mix turbulently.

www.fz-juelich.de/nic/Projekte/mellado.html

The Fastest Kettle in the World

Researchers from Hamburg have discovered a method of bringing water to the boil within a trillionth of a second. Their discovery would be the fastest kettle in the world! The process was simulated on Jülich’s supercomputer JUROPA but has yet to be tested experimentally.

In the simulation, scientists from the Center for Free-Electron Laser Science (CFEL) bombarded water with a focused pulse of terahertz radiation. The frequency of these electromagnetic waves is between that of radio waves and infrared radiation. The pulse causes the water molecules to vibrate. The hydrogen bridge bonds, which hold water molecules close together, are broken. A tiny hot cloud is created that dissipates in less than a millisecond. This is enough time for the researchers to monitor important processes in thermal reactions. Around 200,000 hours of processor time were needed on JUROPA to calculate the interaction of the terahertz pulse with water. A nanolitre (a billionth of a litre) was heated to around 600 °C in just half a picosecond. The calculations would have taken some 20 years if a single processor had been used.

This work opens up new opportunities for experiments with chemically or biologically relevant samples. “Water is the single most important medium in which chemical and biological processes take place. It stabilizes certain chemical compounds and enables specific reactions that would otherwise not occur,” says Dr. Oriol Vendrell from CFEL. The researchers are now investigating the effects of terahertz pulses on various types of molecules dissolved in water.

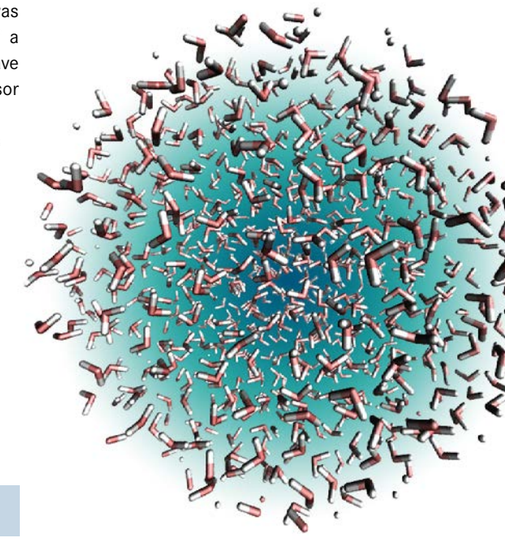

Short-lived hot microclouds: on the supercomputer, researchers made water boil in less than a trillionth of a second.

Applied Chemistry – International Edition: DOI: 10.1002/anie.201305991

NEWS IN BRIEF

Better Performance with Less Electricity

In addition to DEEP Extended Reach (DEEP-ER, see no. 3/2013), a second EU project on exascale computing was launched in October 2013: Mont-Blanc 2. The project builds on and expands the ongoing EU project Mont-Blanc. The Jülich Supercomputing Centre (JSC) is involved in the project, which focuses on energy-efficient supercomputers for the future. With one trillion floating point operations per second, the exascale generation will be 100 to 1,000 times faster than today's petaflop systems. The Mont-Blanc partners are developing a computer architecture that is simultaneously 15 to 30 times more energy-efficient than today's supercomputers. Here, Mont-Blanc is banking on energy-saving processors from the mobile phone sector. JSC researchers are adapting the performance analysis tool Scalasca for hardware prototypes equipped with such processors.

www.deep-er.eu
www.montblanc-project.eu

Making More Efficient Use of Industry Data

In January 2014, companies and research institutions established the Smart Data Innovation Lab (SDIL) as a platform to improve the management and application of large volumes of data from industrial processes. Supercomputing infrastructure is being created for this in Karlsruhe. SDIL will transfer the latest research results to industry as rapidly as possible to generate competitive advantages. The Jülich Supercomputing Centre and the chemicals company Bayer are jointly responsible for the research field of personalized medicine, which aims to tailor drug therapies for diseases and illnesses to meet the specific needs of each individual patient. Other research fields include the transformation of the energy sector, smart cities, and industry 4.0, which links industrial manufacturing with information technology. www.sdil.de

SimLab for Chemistry and Physics

In late 2013, another simulation laboratory (SimLab) was set up at the Jülich Supercomputing Centre: the JARA-HPC Simulation Laboratory for Ab Initio Methods in Chemistry and Physics. It provides support for users of Jülich's supercomputer resources who conduct ab initio simulations in physics, chemistry, nanoscience, and materials science. The calculations for these simulations do not make use of parameters derived in experiments. The SimLab is particularly important for research projects concerned with developing algorithms, optimizations, and increasing performance.

www.jara.org/index.php?id=192&S=0&L=0

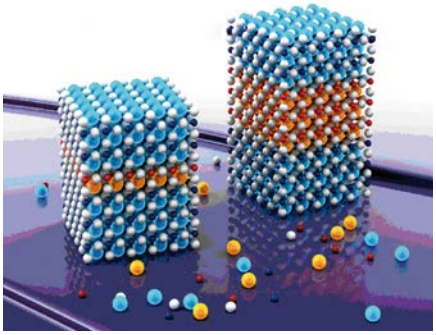

Now available for smartphone und tablet!

Exascale Newsletter
www.exascale-news.de

„RESEARCH in Jülich" (magazine app)
www.fz-juelich.de/app

iOS (iPad) Android

PUBLICATION DETAILS

EXASCALE NEWSLETTER of Forschungszentrum Jülich **Published by:** Forschungszentrum Jülich GmbH | 52425 Jülich, Germany **Conception and editorial work:** Dr. Anne Rother (responsible according to German press law), Tobias Schlößer, Christian Hohlfeld **Graphics and layout:** Graphical Media, Forschungszentrum Jülich **Translation:** Language Services, Forschungszentrum Jülich **Photos:** p. 1, 2, 3 t.: Forschungszentrum Jülich; p. 1: ©fotolia.com/adimas/vladgrin; p. 3 b.: Oriol Vendrell/DESY; p. 4: JARA **Contact:** Corporate Communications Tel: +49 2461 61-4661 | Fax: +49 2461 61-4666 Email: info@fz-juelich.de

UPCOMING EVENTS

GPU Programming

7–9 April 2014
Jülich Supercomputing Centre

This course run by the PRACE Advanced Training Centre (PATC) introduces participants to CUDA, OpenCL, and Multi-GPU programming. The focus is on optimizing and tuning scientific applications.

Instructors: Dr. Jan Meinke, Jochen Kreutz, Peter Philippen, Dr. Andrej Adinets, Anke Zitz, JSC; Jiri Kraus, NVIDIA

www.fz-juelich.de/ias/jsc/events/gpu

Advanced GPU Programming

5–6 May 2014
Jülich Supercomputing Centre

The course concentrates on identifying and eliminating bottlenecks, creating profiles, and advanced programming techniques in GPU programming.

Instructors: Dr. Jan Meinke, Jochen Kreutz, Peter Philippen, Willi Homberg, Dr. Andrej Adinets, JSC; Suraj Prabhakaran, GRS; Jiri Kraus, NVIDIA

www.fz-juelich.de/ias/jsc/events/advgpu

Programming and using supercomputer resources

19–20 May 2014
Jülich Supercomputing Centre

New users of Jülich's supercomputer resources are given an introduction to the existing systems. Amongst other things, they learn how to make optimum use of approved computing resources.

Instructors: Employees from IBM, Intel and ParTec; JSC employees

www.fz-juelich.de/ias/jsc/events/sc-may

Parallel I/O and portable data formats

21–23 May 2014
Jülich Supercomputing Centre

Avoiding performance losses in parallel computers and conducting simulations on different systems: the PATC course provides an introduction to parallel I/O and portable data formats like HDF5 and netCDF.

Instructors: Wolfgang Frings, Dr. Michael Stephan, Dr. Florian Janetzko, JSC

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

You can find an overview of events at the Jülich Supercomputing Centre at:

www.fz-juelich.de/ias/jsc/events