EXASCALE
03-2017 | NEWSLETTER ON SUPERCOMPUTING AND BIG DATA

JÜLICH
FORSCHUNGSZENTRUM

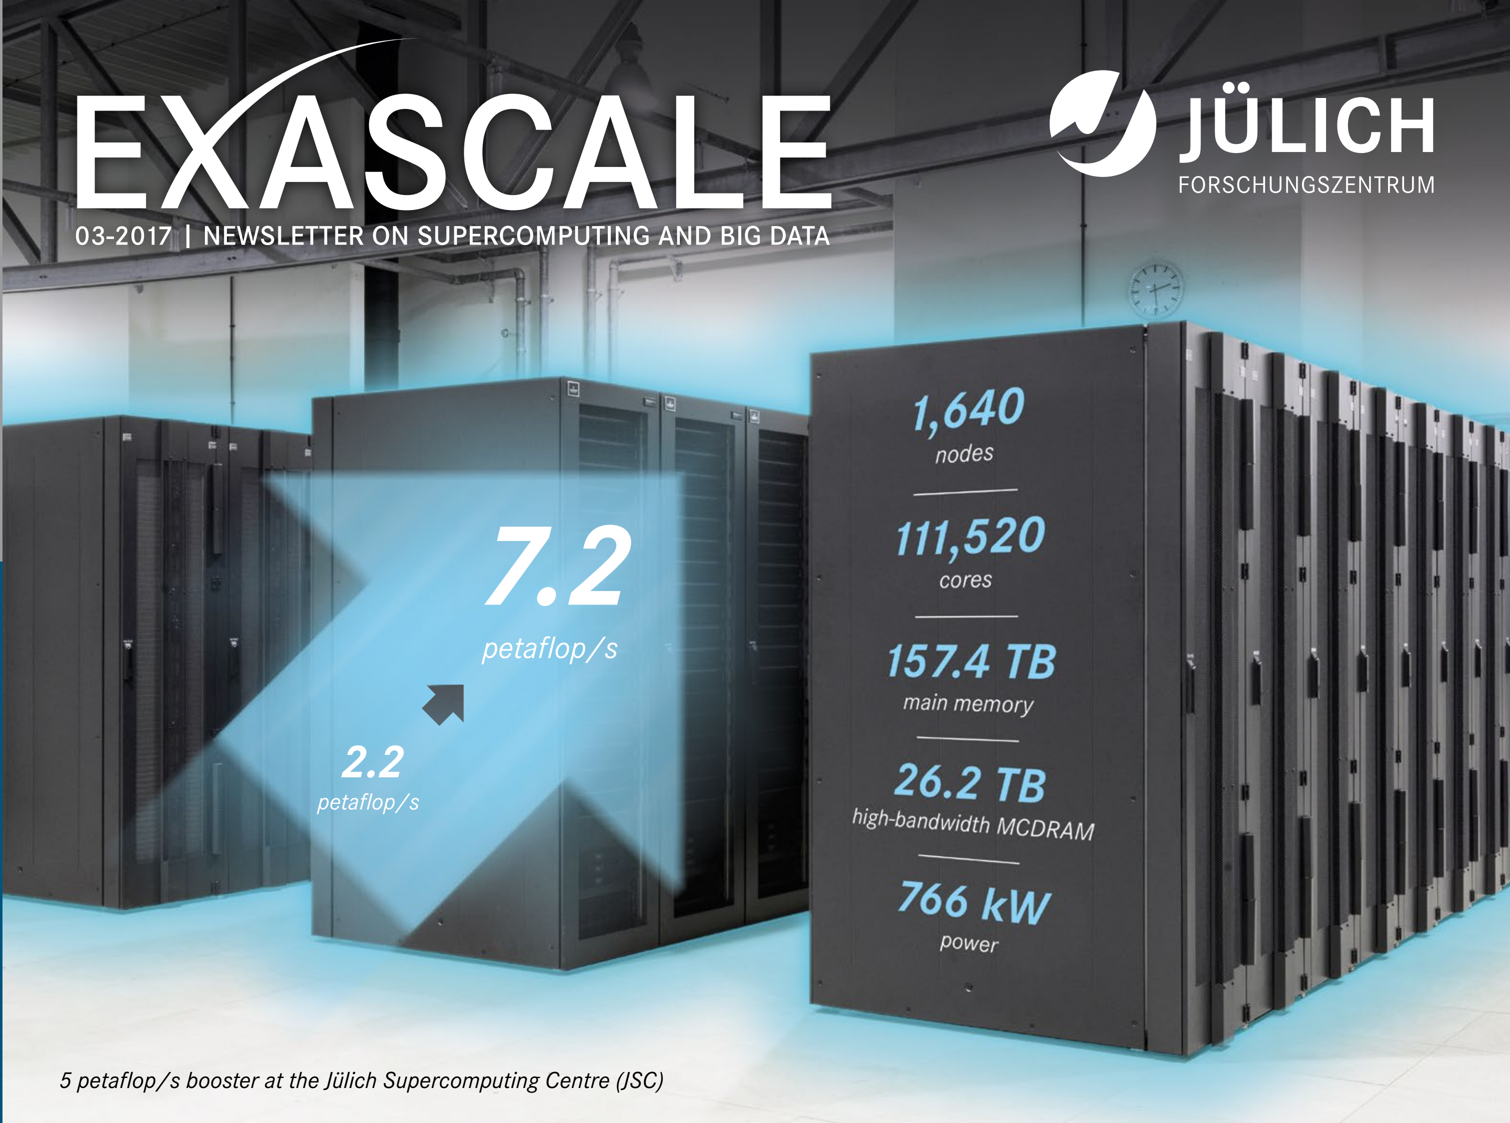

5 petaflop/s booster at the Jülich Supercomputing Centre (JSC)

Member of the Helmholtz Association

Boost for JURECA

Booster module triples computing power

The Jülich Supercomputing Centre's new Booster is as large as an entire supercomputer and accommodated across 33 racks in JSC's computer room. The expansion is integrated in the JURECA supercomputer, with the purpose of significantly increasing the computing power. Thanks to the Booster and its additional 5 petaflop/s – the equivalent of 5 quadrillion mathematical operations per second – the computer, which was developed with Jülich's involvement, is now one of the fastest in Europe. It is also the first supercomputer with an innovative Cluster–Booster architecture to enter productive operation at a computing centre.

The Cluster–Booster concept originated from the DEEP and DEEP-ER projects. Their aim is to develop a modular supercomputer architecture specially tailored to the requirements of advanced simulation codes. Various types of modules can be controlled and dynamically assigned using one uniform software interface, ensuring their individual advantages can take effect.

Any parts of the simulation code that are difficult to process on a multitude of cores simultaneously are instead run on a Cluster, whose general-purpose processors are designed to deal with instructions in quick succession. In contrast, other program parts that can be more easily parallelized are outsourced to the Booster. The latter uses a large number of relatively slow but extremely energy-efficient processor cores, which significantly increases the system's energy efficiency.

JURECA's Booster represents an addition of 1,640 compute nodes using Intel Xeon Phi processors with 68 cores each. The JURECA Cluster module was procured from T-Platforms in 2015. Together, both modules achieve a peak performance of 7.2 petaflop/s. the parts use different network technologies: The JURECA Cluster makes use of InfiniBand EDR while an Intel Omni-Path Architecture is used in the JURECA Booster. They will in future be interlinked by a novel high-speed bridge, which is currently being developed.

STATEMENT

Prof. Paolo Carloni
Head of
➚ IAS-5/INM-9
Forschungszentrum Jülich

HPC-based simulations allow molecular-level aspects of neuronal signalling pathways to be unravelled and new compounds for brain imaging and therapeutic applications to be designed. The JURECA upgrade greatly extends our ability to investigate these fascinating and pharmacologically relevant processes.

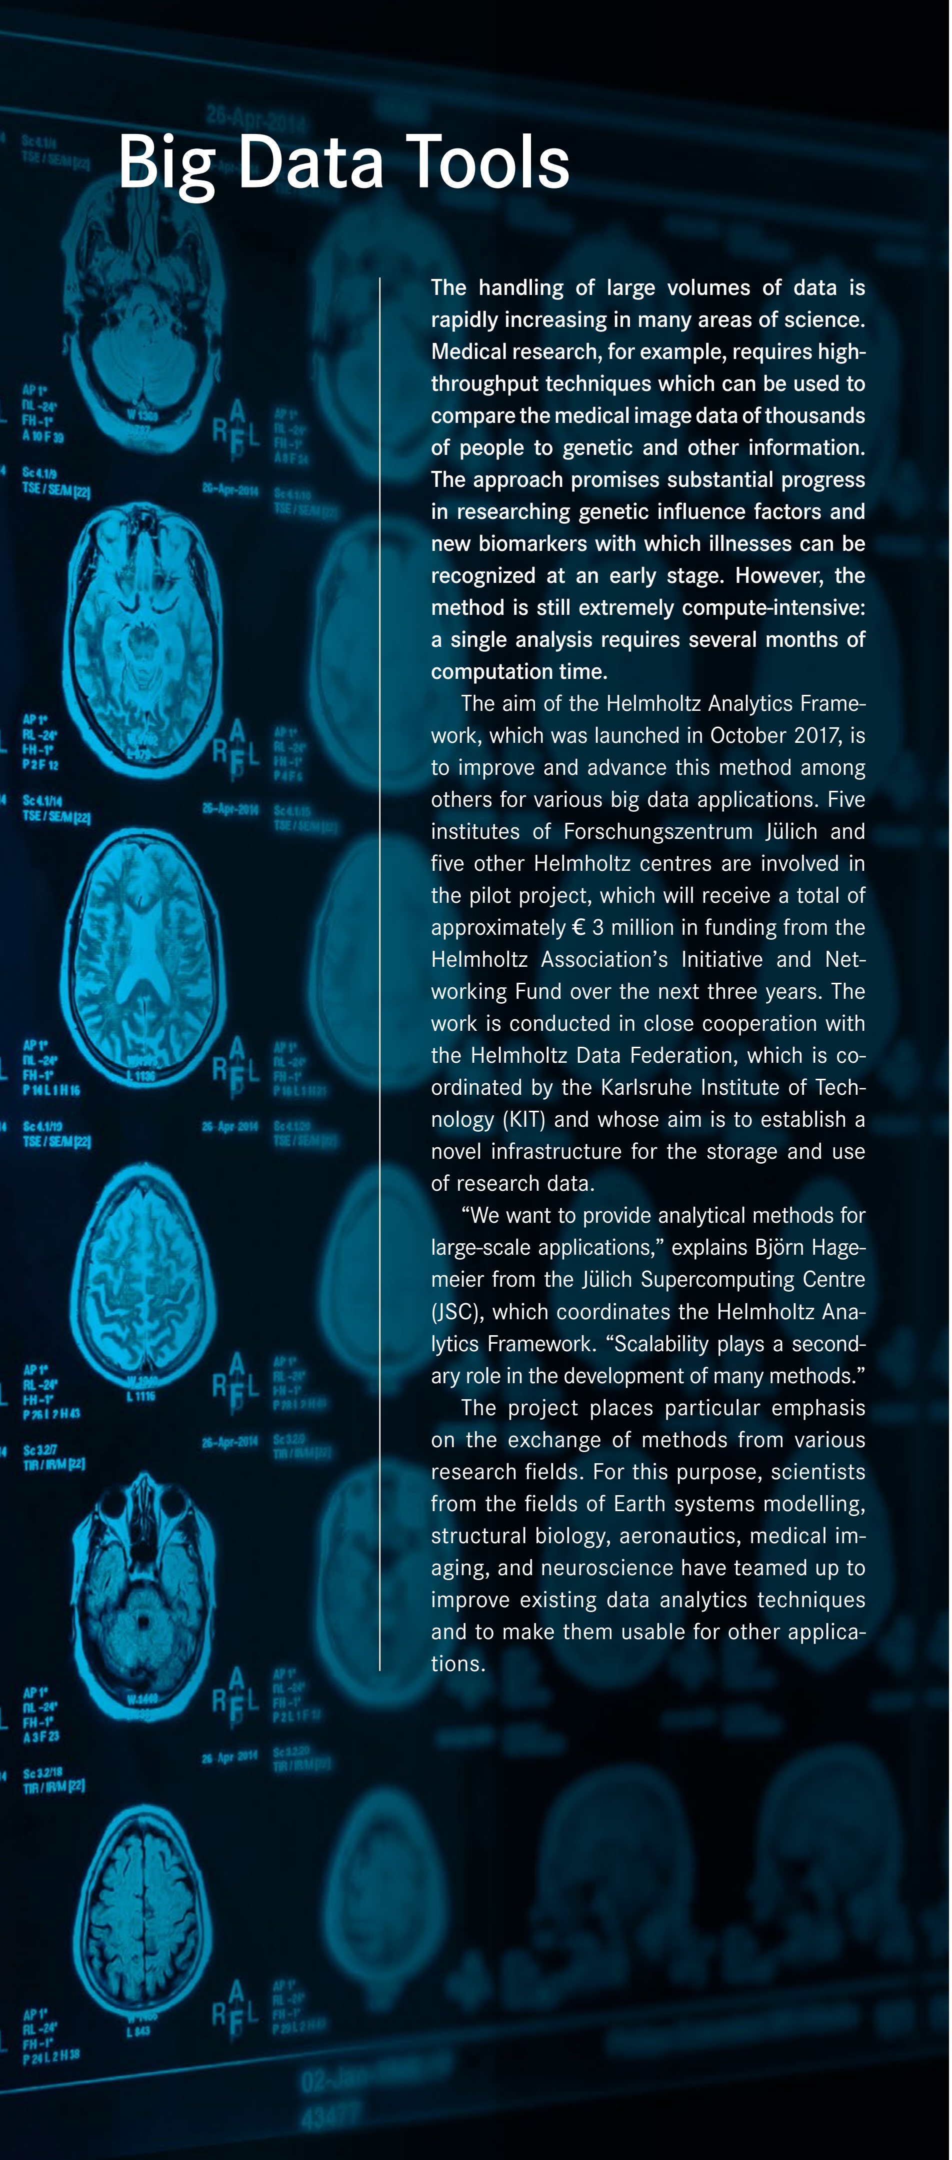

Big Data Tools

The handling of large volumes of data is rapidly increasing in many areas of science. Medical research, for example, requires high-throughput techniques which can be used to compare the medical image data of thousands of people to genetic and other information. The approach promises substantial progress in researching genetic influence factors and new biomarkers with which illnesses can be recognized at an early stage. However, the method is still extremely compute-intensive: a single analysis requires several months of computation time.

The aim of the Helmholtz Analytics Framework, which was launched in October 2017, is to improve and advance this method among others for various big data applications. Five institutes of Forschungszentrum Jülich and five other Helmholtz centres are involved in the pilot project, which will receive a total of approximately € 3 million in funding from the Helmholtz Association's Initiative and Networking Fund over the next three years. The work is conducted in close cooperation with the Helmholtz Data Federation, which is coordinated by the Karlsruhe Institute of Technology (KIT) and whose aim is to establish a novel infrastructure for the storage and use of research data.

"We want to provide analytical methods for large-scale applications," explains Björn Hagemeier from the Jülich Supercomputing Centre (JSC), which coordinates the Helmholtz Analytics Framework. "Scalability plays a secondary role in the development of many methods."

The project places particular emphasis on the exchange of methods from various research fields. For this purpose, scientists from the fields of Earth systems modelling, structural biology, aeronautics, medical imaging, and neuroscience have teamed up to improve existing data analytics techniques and to make them usable for other applications.

SC 17: Flying the Jülich flag

Black, red, and gold – the colours of the German flag have also been chosen for the SC 17 logo, possibly as an allusion to its coordinator. Bernd Mohr is the first ever non-American to organize the most important conference in the field of supercomputing, which will take place in Denver, USA, this year. Preparations have taken three years, with 600 volunteers contributing. Bernd Mohr, a born networker, has already been included in the "People to Watch" list twice – once in 2015 and once in 2017. The list is compiled by the online news service *HPCwire* once a year. The question "Do you know Bernd Mohr?" has become something of a catchphrase in the supercomputing scene.

Dr. Mohr, what is so special about SC?

The world's entire supercomputing community gets together at this conference every year. More than 11,000 experts from 60 nations take part, from universities and research laboratories, but also computer manufacturers and users. It is both a scientific conference and a trade fair: manufacturers present their latest products but research organisations such as Forschungszentrum Jülich also have a booth.

What new trends are developing?

There are two major trends: for a long time, computing was the main concern – how fast can I compute what? But other aspects are now starting to come to the fore. For example, there is a need to process large data volumes stemming from scientific instruments such as satellites or the LIGO experiment, which recently received the Nobel Prize [for the observation of gravitational waves; editor's note]. The keywords here are big data and data analytics.

And the other trend?

The other trend is deep learning. Machine learning processes and artificial intelligence methods are being used to analyse the large volumes of data from scientific experiments. And then of course, you have the race for the exascale computer: There are indications that China, Japan, or the USA could build an exascale computer by 2020 or 2021. Plans for such computers are being developed at the moment – behind closed doors, of course, but everyone hopes to learn more about it at SC.

- *Blog "Do you know Bernd Mohr?" written by Bernd Mohr:* ***blogs.fz-juelich.de/berndmohr***

- *Complete video interview (in German):* ***exascale-news.de/en/sc17***

- ***Supercomputing-Konferenz SC17***

Focus on Biomolecules

Structural biology is the common topic of two new research groups at the John von Neumann Institute for Computing (NIC). The "Computational Structural Biology" group, founded in September 2017, works on new methods which can be used to derive the 3D structure of biomolecules and proteins from various data sources. Proteins are huge chain-like molecules that are involved in almost all processes of life. Their complex folded spatial structure is crucial to their function in the cell.

The exact arrangement of the atoms can often only be determined by laboriously piecing them together. One priority of the group headed by Dr. Alexander Schug is to record "co-evolution data". In doing so, they focus on mutations in particular – changes of individual amino acid components in the structure of complex proteins. The objective is to find mutations that do not occur individually but which are followed by further mutations in their vicinity. Such co-evolved pairs also permit conclusions to be drawn about the spatial proximity of the two partners – and therefore on the spatial structure of the entire protein.

Back in late April 2017, another NIC research group entitled "Computational Biophysical Chemistry" was launched. The group headed by Prof. Holger Gohlke pools the competences of various expert fields and aims to answer complex questions on the structure, dynamics, and function of biomolecules and supermolecules.

The focus of the group is on methods which permit the physiological effects of biomolecules to be predicted and modulated – for example in order to research new medical active substances. Thus, Gohlke has examined a mutation to an essential enzyme in humans, which causes a major disruption of the nitrogen metabolism. Using molecular-dynamics simulations on Jülich's supercomputer JURECA, the functional disturbance was investigated in detail, and thus a candidate drug was tested that counteracts the weakening of molecular binding forces caused by the mutation.

New NIC Research Group „Computational Structural Biology“

New NIC Research Group „Computational Biophysical Chemistry“

PLOS Computational Biology, DOI: 10.1371/journal.pcbi.1004693

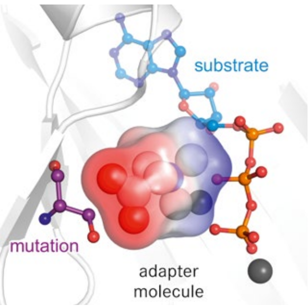

Protein repair: adapter molecules (betaines) enhance ATP binding that is essential for thehuman metabolism

Flow Predictions for Your Nose

No one is immune to the sniffles. But for approximately 11 % of the population, nasal airway obstruction is chronic. As part of the Rhinodiagnost project, supercomputing experts are working together with other specialists to support physicians in making the – often difficult – decision for or against an operation. The aim is to set up a service network that will provide 3D models and flow simulations as additional decision-making aids.

More than 100,000 operations on the nasal cavity and sinuses are performed in Germany every year.

At an estimated 30 %, the error rate of such procedures is relatively high. The current diagnostic methods are too unreliable. Medical imaging methods are usually used for diagnoses, or purely optical methods that provide no information about flow conditions.

"Correctly estimating the chances of success of an operation currently depends heavily on the physician's experience," explains Jens Henrik Göbbert from the Jülich Supercomputing Centre (JSC). "Within the scope of the Rhinodiagnost project, we want to make it possible to run flow simulations for individual patients. These would be based on computed tomography images and would permit the effect of anatomical changes to be depicted."

The healthy, naturally shaped nose is a little marvel of fluid dynamics and is easily damaged during an operation. Over a distance of not even 10 cm, the nose moistens the air that is breathed in and warms it to body temperature – even at outside temperatures below zero. The simulation of air flow on supercomputers is technically well-established and could represent an additional measure, avoiding unnecessary procedures and permitting operations to be adapted individually.

Researchers from the Jülich Supercomputing Centre (JSC) and the Institute of Aerodynamics (AIA) at RWTH Aachen University are developing the necessary software components in Rhinodiagnost to enable physicians to use simulation data interactively and purposefully on modern high-performance computers. In addition, it is planned to demonstrate virtual operations in which relevant flow parameters can be displayed and analysed in real time. The German companies Sutter Medizintechnik and Med Contact as well as the Austrian Applied Information Technology Research Society, which is coordinating the project launched in September 2017, are involved as project partners from industry.

rhinodiagnost.eu

Simulation of nasal air flow using computational fluid dynamics (CFD)

NEWS IN BRIEF

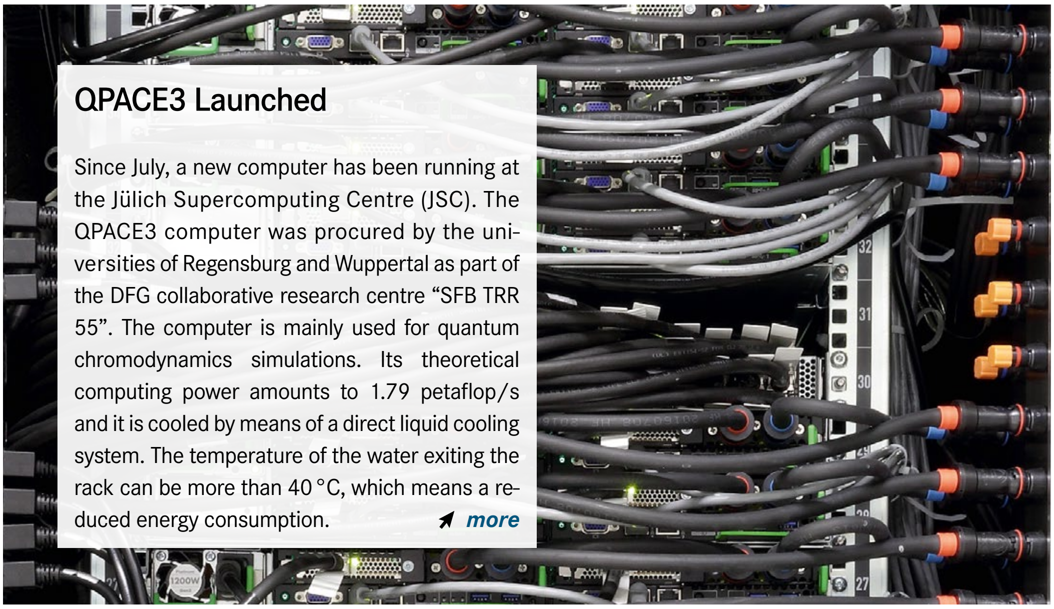

QPACE3 Launched

Since July, a new computer has been running at the Jülich Supercomputing Centre (JSC). The QPACE3 computer was procured by the universities of Regensburg and Wuppertal as part of the DFG collaborative research centre "SFB TRR 55". The computer is mainly used for quantum chromodynamics simulations. Its theoretical computing power amounts to 1.79 petaflop/s and it is cooled by means of a direct liquid cooling system. The temperature of the water exiting the rack can be more than 40 °C, which means a reduced energy consumption. ***more***

Focus on Deep Learning

A new overarching team at JSC is set to drive forward research in the field of deep learning and to optimize user support in this area. In addition, the team will also investigate new computer architectures for unsupervised and reinforcement learning and, together with scientists from other fields, develop applications permitting large volumes of raw data to be analysed and combined in a way that makes it possible to discover valuable, hidden information.

GCS to Receive € 500 Million

The German federal and state governments will provide the Gauss Centre for Supercomputing (GCS) with funding totalling € 500 million until the end of 2025. The budget will serve to expand the computer infrastructure as well as user support and software development as part of the "Smart Exascale" mission. The three large supercomputing centres in Germany – Stuttgart (HLRS), Jülich (JSC), and Munich (LRZ) – have been pooling their activities in GCS since 2007 in order to further expand their role as a global player in the field of high-performance computing. ***more***

European Hardware Procurement

Computing centres from four European countries have joined forces to procure new innovative hardware together for the first time. The process is coordinated by Forschungszentrum Jülich and intended to provide more resources on a European level in future. In addition, the centres involved (BSC, CINECA, JSC, GENCI) aim to increase their influence on hardware development via PPI4HPC in order to better adapt hardware solutions to the requirements of scientists and engineers. ***more***

Read it online!

- *Exascale Newsletter:*
  exascale-news.de
- *effzett – the crossmedia magazine:*
  effzett.fz-juelich.de/en
- *Facts and figures:*
  fakten.fz-juelich.de/en

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 **Graphics und Layout:** SeitenPlan GmbH **Translation:** Language Services, Forschungszentrum Jülich **Photos:** p. 1 top: Forschungszentrum Jülich/Wilhelm Peter Schneider/SeitenPlan; p. 1 bottom: Forschungszentrum Jülich/Sascha Kreklau; p. 2 left: sacura/Shutterstock.com; p. 2 right: Forschungszentrum Jülich/R.-U. Limbach; p. 3 top: Forschungszentrum Jülich/Holger Gohlke; p. 3 bottom: CLIPAREA I Custom media Shutterstock.com/Forschungszentrum Jülich/Jens Henrik Göbbert/SeitenPlan; p. 4: Forschungszentrum Jülich/R.-U. Limbach **Contact:** Corporate Communications | Tel.: +49 2461 61-4661 | Fax: +49 2461 61-4666 | Email: info@fz-juelich.de **Printed by:** Schloemer & Partner GmbH **Print run:** 200

UPCOMING EVENTS

Training course:
"Introduction to the programming and usage of the supercomputer resources at Jülich"
23–24 November 2017
at the Jülich Supercomputing Centre
Instructors: Employees from Intel and ParTec; JSC employees
more

Training course:
"Advanced Parallel Programming with MPI and OpenMP"
27–29 November 2017
at the Jülich Supercomputing Centre
Instructors: Dr. Rolf Rabenseifner, HLRS; JSC employees
more

Training course:
"Parallel and Scalable Machine Learning"
15–17 January 2018
at the Jülich Supercomputing Centre
Instructor: Prof. Morris Riedel, JSC
more

Training course:
"Parallel I/O and Portable Data Formats"
12–14 March 2018
at the Jülich Supercomputing Centre
Instructors: Sebastian Lührs, Dr. Michael Stephan, Benedikt Steinbusch, Dr. Kay Thust, JSC
more

Overview of events at the Jülich Supercomputing Centre:
fz-juelich.de/ias/jsc/events