INNOVATIVES SUPERCOMPUTING IN DEUTSCHLAND N° 17-1 · SPRING 2019

INSiDE

Publishers

Prof. Dr. D. Kranzlmüller
Director
Leibniz Supercomputing Centre

Prof. Dr. Dr. Th. Lippert
Director
Jülich Supercomputing Centre

Prof. Dr.-Ing. Dr. h.c. Dr. h.c. Hon.-Prof. M. M. Resch
Director
High-Performance Computing Center Stuttgart

InSiDE is published two times a year by the Gauss Centre for Supercomputing (HLRS, JSC, LRZ)

Editor-in-Chief
Michael Resch, HLRS
resch@hlrs.de

Editor
Eric Gedenk, GCS, HLRS
gedenk@hlrs.de
e.gedenk@gauss-centre.eu

Production Manager
F. Rainer Klank, HLRS
klank@hlrs.de

Design
Zimmermann Visuelle Kommunikation
www.zimmermann-online.info

This magazine is printed on paper that has been certified by FSC, the EU Ecolabel, and the Blue Angel Ecolabel.

If you would like to receive InSiDE regularly, please send an e-mail with your postal address to F. Rainer Klank: klank@hlrs.de

Contributing Authors

Eric Gedenk (*eg*), GCS, HLRS
gedenk@hlrs.de
e.gedenk@gauss-centre.eu

Lena Bühler (*lb*), HLRS
buehler@hlrs.de

Christopher Williams (*cw*), HLRS
williams@hlrs.de

Dr. Ludger Palm (*lp*), LRZ
presse@lrz.de

Dr. Estela Suarez (*es*), JSC
e.suarez@fz-juelich.de

Dr. Dirk Pleiter (*dp*), JSC
d.pleiter@fz-juelich.de

Dr. Brian Wylie (*bw*), JSC
b.wylie@fz-juelich.de

Visit us at:
www.gauss-centre.eu

GCS
Gauss Centre for Supercomputing

Welcome to the latest issue of InSiDE, the bi-annual Gauss Centre for Supercomputing (GCS) magazine highlighting innovative supercomputing in Germany. 2019 is already shaping up to be an exciting year for GCS, and we are only at the halfway point. In addition to completing the centres' transitions to next-generation architectures, we are also hard at work on addressing the challenges of the next generation of high-performance computing (HPC).

The Leibniz Supercomputing Centre is finalizing the installation of its newest machine, SuperMUC-NG, and researchers are already seeing encouraging results from its early user program (*Page 8*). Staff at the Jülich Supercomputing Centre are playing leadership roles in the European Processor Initiative, which aims to design and implement a European-made processor within the next decade (*Page 16*). The High-Performance Computing Center Stuttgart has been collaborating with Deutsche Bahn and Stuttgart-based industry partners to support the design of the massive infrastructure project, Stuttgart 21, a remodeling of the Stuttgart main train station (*Page 4*). In addition, HLRS will start installation on its next-generation supercomputer, Hawk, later this year.

Our users have also had major breakthroughs in their respective scientific fields. Scientists at the Technical University of Munich collaborated with investigators from the Canadian research organization Ouranos to simulate high-resolution climate models over 150-year period (*Page 10*). Physicists at the Helmholtz Zentrum Dresden-Rossendorf used HPC to confirm a major breakthrough in 2D materials (*Page 14*).

As we continue our infrastructure changes, GCS remains committed to providing users a robust training program in order to make the best use of these large investments. In addition, we are further expanding our user support structure to provide users easy access to experts who not only help address the challenge of optimizing codes to run on our resources, but also offer guidance on how to best achieve their scientific goals. As always, GCS prides itself in offering our users multiple HPC architectures so they can use resources that are best-suited to their needs.

As HPC technology continues to evolve, the three GCS centres are embracing the challenges and feel well-positioned to continue to lead in HPC innovation as we head toward the exascale horizon.

Prof. Dieter Kranzlmüller
Prof. Thomas Lippert
Prof. Michael Resch

HLRS VISUALIZATION SUPPORTS STUTTGART 21 CONSTRUCTION

3D virtual reality models created at HLRS helped engineers to successfully build an architecturally unique and complex kind of reinforced concrete column that will be a showpiece of the new Stuttgart train station.

Virtual reality visualization of the Stuttgart 21 columns at HLRS. **© HLRS**

In 2018 construction began on the most architectonically distinctive elements of the new Stuttgart central train station. The so-called "Kelchstütze" are a set of massive, chalice-shaped columns that will cradle windows to bring natural light into the subterranean station. At a height of over 12 meters and a diameter of approximately 32 meters, the elegantly curving forms of reinforced concrete promise to become a landmark in the city.

For construction engineers working on the project — called Stuttgart 21 — the complex geometry of the columns created unique challenges. Although CAD drawings of the structures existed, the curving volumes and need to precisely arrange an enormous number of steel bars to reinforce the concrete proved more challenging than is the case with more conventional pillars.

To better plan the structures' fabrication, engineers at the Deutsche Bahn called upon HLRS to develop a 3D visualization. Using CAD and point cloud data, HLRS staff created an interactive digital model of one of the columns for display in the CAVE, a walk-in virtual reality environment.

Meeting at HLRS, representatives of the Deutsche Bahn and construction managers at Ed. Züblin AG could then explore and interact with the visualization to develop a strategy for erecting the column most efficiently and at the highest possible structural integrity. Once the first column was poured onsite, the HLRS team also created 3D scans of the physical structure to compare the virtual model to the actual column and determine whether it had been erected to specifications.

According to Uwe Wössner, who heads the HLRS Visualization Department, "The effort was extremely productive, as the participants had an easier time visualizing the complex geometry, identified problems in the original design, and were able to develop effective solutions." *cw*

JSC HARDWARE AND EXPERTISE SUPPORT INTERNATIONAL LOW-FREQUENCY ARRAY

As researchers comb through treasure troves of data from Europe's largest radio telescope, they leverage high-performance computing and data analytics to accelerate discovery.

For centuries, humanity has tried to get a better view, and understanding, of the distant celestial objects lighting up the sky. From the invention of the earliest telescopes, people have painstakingly noted slight changes in stars' positions, colours, and brightness.
In 2010, European astronomers made a major leap in studying distant space—with the construction of the Low-Frequency Array (LOFAR), researchers could study low-frequency radio waves emanating from the farthest reaches of space. Primarily located in the Netherlands but with stations in 5 other European countries, LOFAR uses patience to study the subtle signals coming from stars and other celestial bodies by taking snapshots, or pointings, of different places in the night sky, focusing its thousands of antennae on a slice of sky usually for 8 hours at a time.
7 years later, roughly 200 researchers are starting to dig into the massive amounts of data the sky survey collects. The first public release of data, detailing the findings of 2 percent of all data collected thus far, came out in 26 articles, filling up an entire special issue of the journal Astronomy and Astrophysics.

LOFAR is collecting massive amounts of data about the northern hemisphere sky, but in order to make sense of these datasets, astronomers work closely with high-performance computing (HPC) centres to help efficiently store, manage, and analyze that data. As one of the German partners of LOFAR, Forschungszentrum Jülich, specifically the Jülich Supercomputing Centre (JSC), houses roughly a third of the data generated by LOFAR, currently managing about 15 petabytes for the project (a petabyte is 1 million gigabytes).

Astronomers like Dr. Matthias Hoeft, researcher at the Thuringian State Observatory Tautenburg and partner in the LOFAR project, have increasingly come to see HPC facilities as indispensable partners in their quest for scientific discovery.

"The amounts of data generated by these telescopes are continuously growing," Hoeft says. "The data is driving the need for us to have access to large amounts of computing power in order to reconstruct the actual sky brightness distribution, which is a key part of our mission, but we also need to analyze the data in other possible ways to help us understand anomaly events in space, though, such as rare short radio outbursts, a phenomenon recently discovered and whose origins we do not yet understand."

To assist in this process, JSC has become one of the three main centres for storing and managing LOFAR data, becoming the project's long-term storage archive. JSC is one of the three centres that make up the Gauss Centre for Supercomputing (GCS).

Jülich's LOFAR station DE605 consists of two antenna fields for measuring high and low frequencies. The container between them houses the electronics for processing the signals from the individual antennae. **© JSC Jülich / Ralf-Uwe Limbach**

Radio galaxies

LOFAR observations show researchers outcomes of very energetic events in the universe, such as when stars explode (supernovae), the outflow of supermassive black holes in the centre of galaxies, or when huge clusters of galaxies collide. All these events can accelerate electrons to a velocity very close to the speed of light. Magnetic fields, a ubiquitous part of space, influence these electrons, causing them to lose their energy and, in turn, emit radio waves.

Particles are not the only objects emanating through space, though—energetic, high-frequency electromagnetic radiation also shoots out from stars and supernovae. As with any wave, the longer these excited waves travel, the slower they move and oscillate. Waves that have been travelling a long time eventually become radio waves that fall on the low end of the electromagnetic spectrum.

LOFAR enables astronomers to "see" far deeper into space than is possible with the human eye. By focusing on extremely low-frequency wavelengths—specifically,

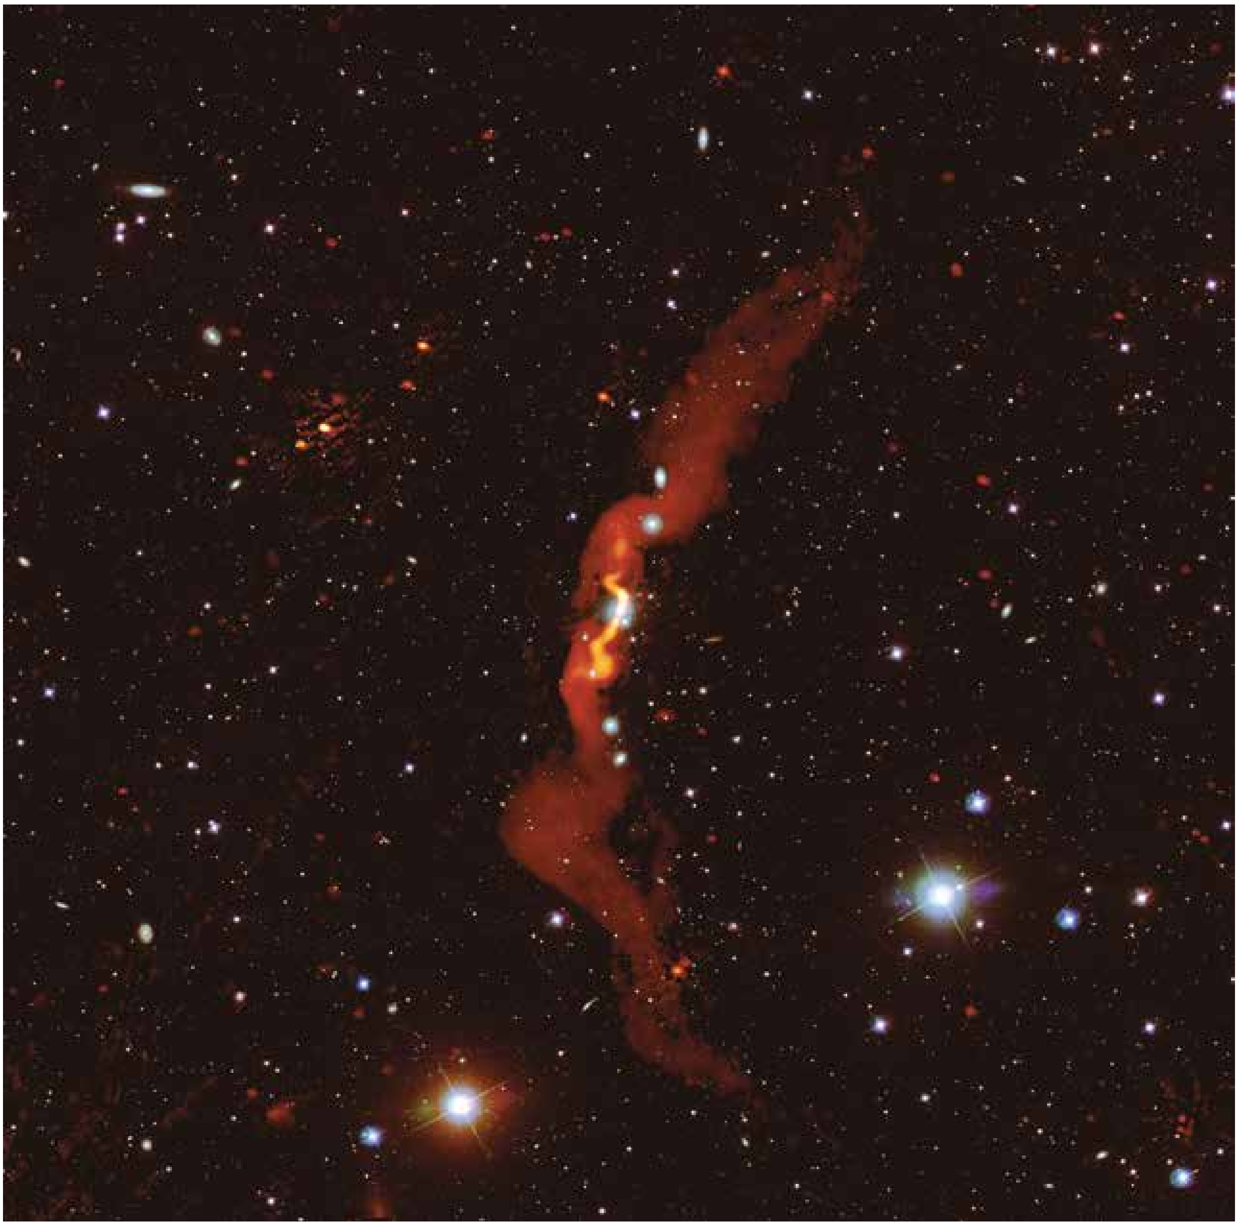

the ranges 10–80 megahertz and 110–240 megahertz—LOFAR can pick up on stars and galaxies that are billions of light-years away from Earth or illuminate cosmic phenomena taking place on previously invisible wavelengths.

Data generated from LOFAR is already helping inform astronomical research in a variety of areas, including understanding the influence of black holes on celestial bodies, trans-galactic magnetic fields, the creation of galaxy clusters (merging galaxies), and the properties of gas clouds that surround these clusters.

Sharper Images

Because LOFAR is collecting information about electromagnetic radiation sources, astronomers need to convert raw data into useable images. Creating detailed images through massive amounts of data, while also accounting for observation distortion caused by the Earth's ionosphere, is an extremely time consuming process, as is moving large amounts of data from one place to another.

To that end, JSC offered to become the LOFAR project's long-term storage archive, helping researchers accelerate scientific discovery. "Processing these gigantic data sets poses a great challenge," says Prof. Dr. Dr. Thomas Lippert, Director of the Jülich Supercomputing Centre. "What would have taken centuries on conventional computers was achieved within one year thanks to the use of innovative algorithms and extremely powerful computers."

The next level of integration and collaboration is focused on making sure that astronomers using the LOFAR data for research can spend as much of their respective time as possible on scientific discovery rather than spending large amounts of time managing datasets.

Researchers connected to LOFAR are in talks to extend the framework European Open Science Cloud to make the LOFAR datasets openly available one year after observation, and also able to be processed at the long-term storage archives so that astronomers interested in doing research with these datasets can get meaningful imagery rather than just large amounts of raw data.

"The intent of the international LOFAR telescope is really to evolve the work to integrate with the European Open Science Cloud so that these resources are really combined and we have a well-defined infrastructure," Hoeft says. *eg*

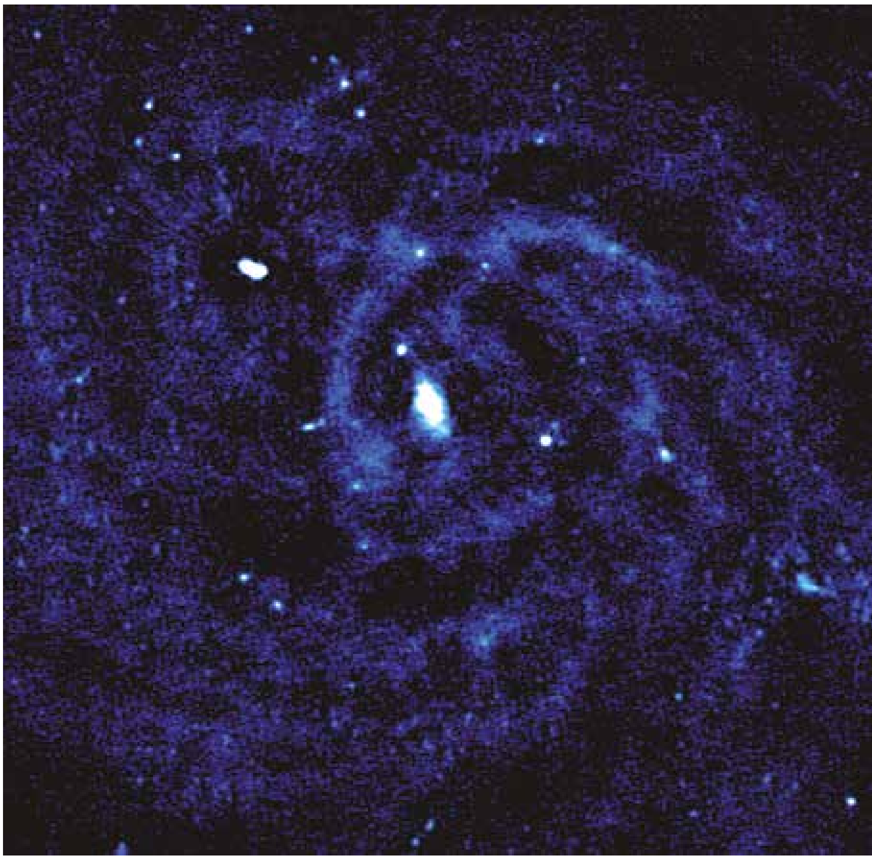

The nearby spiral galaxy IC 342 seen with LOFAR. Radio emission comes from the spiral arms, where supernova explosions accelerate electrons to high energies. Many background active galaxies can be seen behind the spiral galaxy itself.

© Maya Horton and the LOFAR surveys team

Left page: **The radio galaxy 3C31, observed with LOFAR by Heesen et al (2018), is shown in red on top of an optical image. LOFAR reveals the radio galaxy to be more than 3 million light years in size.**

© Volker Heesen and the LOFAR surveys team

SUPERMUC-NG: EARLY USERS' FIRST EXPERIENCES

With every successive generation of high-performance computing (HPC) technology, researchers across the science and engineering spectrum gain more insight into their respective research fields by accessing supercomputers that allow them to run simulations longer, faster, in finer detail, or across larger areas.

Generally speaking, astrophysics and computational fluid dynamics (CFD) have always been among the first scientific domains where users can scale up to take full advantage of the most recent compute resources.

Not surprisingly, researchers from these disciplines were also among the first "Early Users" to put their hands on SuperMUC-NG, the Gauss Centre for Supercomputing's (GCS') new supercomputer at the Leibniz Supercomputing Centre (LRZ) in Garching.

Staff working at LRZ's Astrolab (is supporting Christoph Federrath (Australian National University) in optimizing FLASH, a multiphysics, multiscale simulation code with a wide and international user base.

The code, originally used for modelling astrophysical thermonuclear flashes, has also been repurposed to use in other numerical simulations. FLASH had already been highly optimized and used to great effect on SuperMUC Phase 2, the predecessor to SuperMUC-NG. Federrath and AstroLab staff tested the optimized FLASH code for SuperMUC-NG, and early returns show a perfect linear scaling up to 24,576 cores, filling 512 nodes in memory and processors.

The optimization performed on this version further improved parallel communication, as well as vectorization. As Salvatore Cielo, one of LRZ AstroLab's scientists said, "The performance of SuperMUC-NG is really impressive. At (compute) core level, it visibly surpasses that on SuperMUC Phase 2. The user was so surprised that he thought something was wrong!"

The team ran the largest-ever simulations of astrophysical turbulence in order to understand the effects of hydrodynamic turbulence on the seeding and formation of new stars in interstellar space, verifying predictions of the most recent theoretical models of star formation.

This is a problem that notoriously needs very high resolution, as a large fraction of the energy resides in the smallest simulation scales. Besides the need to compute fine-grained detail over a large area, the influence of magnetic fields seeded by nearby stars in interstellar turbulence qualitatively change the prediction of the models. In order to test these effects, one needs to introduce magnetohydrodynamics (MHD) in the simulation code. These MHD calculations roughly double the effort needed, both in terms of computation as well as memory footprint, as the three spatial components of the magnetic fields have to be stored and evolved throughout the simulation as well. While FLASH can incorporate MHD, the computational costs were prohibitive before this optimization work and the advent of SuperMUC-NG. Besides the impressive peak performance of the machine, the architecture of SuperMUC-NG takes huge advantage of the optimization. FLASH can effectively scale up to a number of cores unlikely achieved elsewhere.

In addition to FLASH, another CFD code showing excellent performance on SuperMUC-NG. This in-house developed

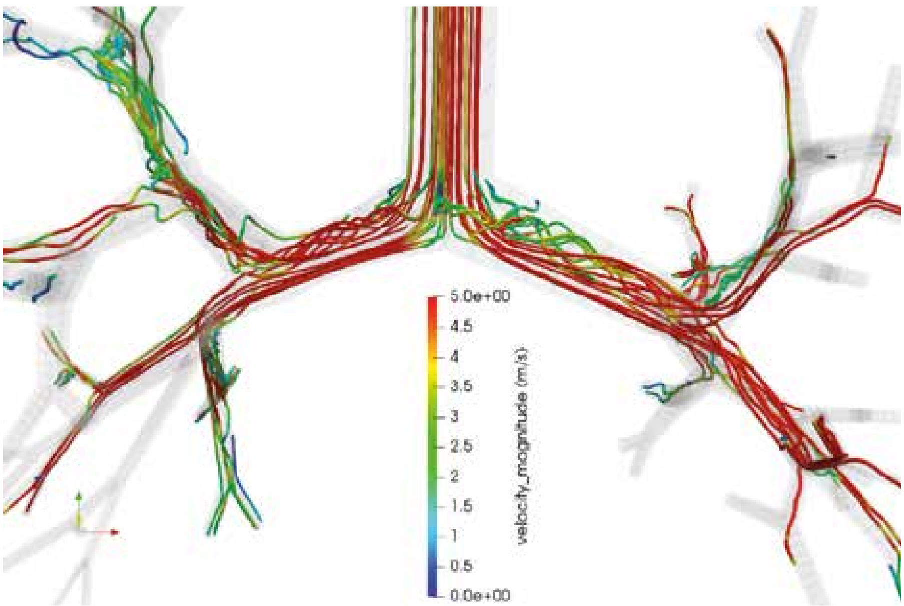

Working together with LRZ's CFDlab staff, researchers from the Technical University of Munich have been modelling fluid dynamics in neonatal lungs of pre-term infants. © TU Munich

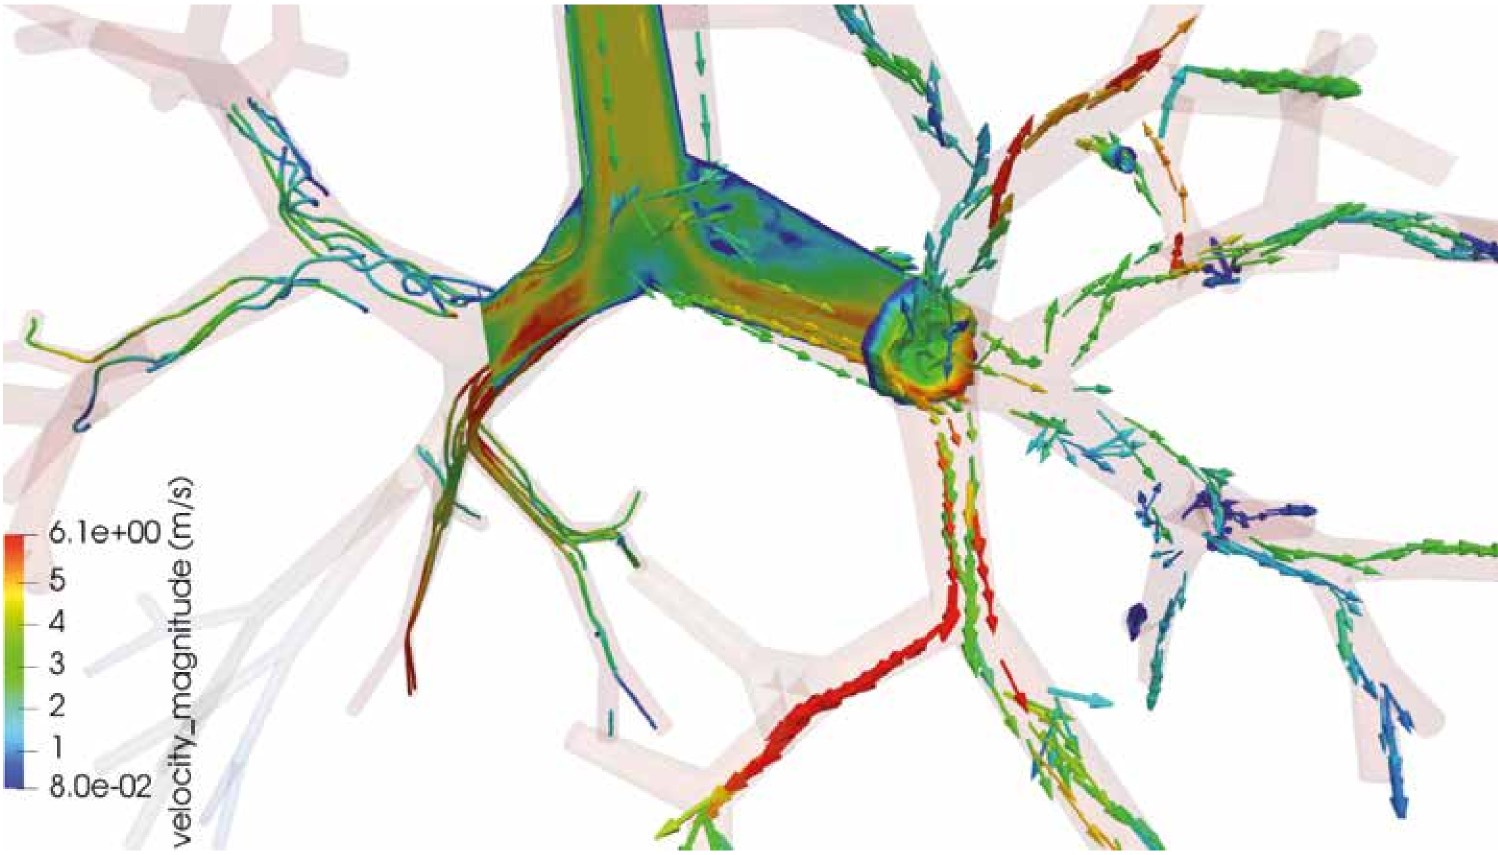

Working together with LRZ's CFDlab staff, researchers from the Technical University of Munich have been modelling fluid dynamics in neonatal lungs of pre-term infants. **© TU Munich**

source code that solves the Navier-Stokes equations continuously developed by Dr. Martin Kronbichler's team with Prof. Wolfgang A. Wall at the Institute for Computational Mechanics of Technical University Munich. The code is intended for high-fidelity biomedical simulations to perform initial simulations of fluid flow in the respiratory system of pre-term infants. These simulations are needed for understanding the gas transport and gas exchange in high-frequency ventilation protocols, ultimately assisting in a deeper understanding of airflow to provide unprecedented insight into the entire neonatal lung.

Together with LRZ's CFDLab team, the code was adapted and optimized for SuperMUC-NG. Due to the early state of the machine, it was only possible to run experiments on up to half of the machine, in this case, 3,168 nodes. (152,064 out of 304,128 cores). Despite these limitations, the code shows perfect linear scaling up to 152,064 cores, and the performance is excellent compared to the previous architectures. The code uses matrix-free methods to balance arithmetic operations, memory access, and cache usage in an optimal way. "We obtained a very high performance value on our iterative solver—an arithmetic throughput up to 780 GFlop/s on a single node and up to 1.53 PFlop/s on 2,048 nodes, more than 10 times higher than typical matrix-based implementation," said Momme Allalen, leader of the CFDLab at LRZ.

"Work is still in progress to get the numbers for the half and the full machine runs," resumes Momme Allalen, leader of CDFLab@LRZ. The code achieves about 18 % of the peak performance of SuperMUC-NG with full machine runs. "Currently, there is no machine with similar size to SuperMUC-NG with more than 6,000 Skylake nodes. SuperMUC-NG is the only one in the world where we can perform such simulations. And we expect the linear scaling to be seen when we will be able to use the whole machine," Allalen added. *lp*

INTERNATIONAL RESEARCH COLLABORATION COMPUTES CLIMATE PAST, PRESENT, AND FUTURE

Canadian-German partnership simulates the climate in Quebec and Bavaria across 150 years.

Many of the world's largest cities are built near coasts, be it along rivers or oceans. Humanity relies on waterways for transportation, trade, and sustenance. However, waterways can also unleash devastating floodwaters that lead to billions in damage, loss of life, and years-long clean-up efforts.

Two geographically distinct areas that are prone to dealing with the fury of flood waters, the Canadian province of Quebec and the German state of Bavaria, have been collaborating for a decade to investigate the impacts of climate change on water resources. The latest endeavor of this partnership, the ClimEx Project, aims to improve researchers' understanding of severe flooding dynamics under changing climate conditions. "This knowledge is of fundamental importance," says Prof. Dr. Ralf Ludwig, Geography Professor at Munich's Ludwig-Maximilians-Universität (LMU) and project leader of ClimEx. "Understanding these phenomena helps us better prepare and improve adaptation to the increasing extreme events we expect to face in our future."

"The goal of ClimEx is to investigate extreme floods associated with long return periods," said Martin Leduc, climate researcher at the non-profit research organization Ouranos and partner in the ClimEx project. "If you look at observations, you only have a relatively short time period to reference—often less than 30 years of accurate, detailed data. For the most extreme floods, these are once-a-century phenomena."

In order to efficiently model long-term climate trends, the ClimEx collaborators are using the SuperMUC supercomputer at the Leibniz Supercomputing Centre (LRZ). The team published its most recent results in the Journal of Applied Meteorology and Climatology, simulating the climate in Quebec and Bavaria from 1950 to 2100.

Zooming in, spreading out

To study climate trends and changes computationally, researchers use a climate model to divide the area of study into a grid for simulating the countless meteorological processes and properties that form an area's climate.

With a finite amount of computing power available for any given simulation, researchers have to simulate a representative area of the globe across a long enough period of time to establish climate trends while also capturing enough detail to verify a model's ability to predict past climate behavior and, in turn, predict future climate events. Therefore, such an experiment involves a balance between the length of the simulations, the level of detail (resolution) of the grid, and the size of the covered area.

To balance these demands, climate scientists use a combination of a global climate model (GCM) and a regional climate model (RCM). While GCMs simulate the climate over the entire globe, they have to sacrifice the level of detail, meaning that the distance between two neighbouring grid cells must be more than 100 kilometres. Using the Canadian Regional Climate Model version 5 (CRCM5) that was developed by the ESCER Centre of l'Université du Québec à Montréal in collaboration with Environment and Climate Change Canada, the researchers are able to study areas of the globe at much higher detail, using grids with a 12-kilometre resolution. This allowed the LMU-Ouranos-LRZ team to run its simulations including relevant climate phenomena in high resolution.

In order to better understand and predict flooding, the team further downscales the ClimEx simulations statistically to provide input data for hyper-accurate, high-resolution hydrological modeling. Not only does this level of detail help better anticipate and plan for large-scale flooding events in Bavaria and Quebec, but it also helps provide higher-quality information for other impact models and decision makers.

Leduc also brought up the "butterfly effect" as it related to climate simulations—even the highest resolution simulations cannot account for all the miniscule changes that can influence climate. Further, researchers have no way of knowing how much humanity will curb its emissions in the coming decades, which could significantly influence climate patterns. In ClimEx, the team ran 50 simulations for Bavaria and 50 for Quebec, with each iteration introducing slight changes in input data, giving them a total of 7,500 years of climate data for each location.

Not only are these simulations computationally expensive, but they also generate an extremely large amount of data—more than 500 terabytes, in fact. To get meaningful results from these simulations and the data analysis that follows, researchers need access to world-leading computing resources.

"Doing these simulations required an incredible amount of computational resources and the calculations last for more than 6 months," Leduc said. LRZ staff helped ensure that the team could run its simulations as efficiently as possible, and helped the team get access to a full computational island on

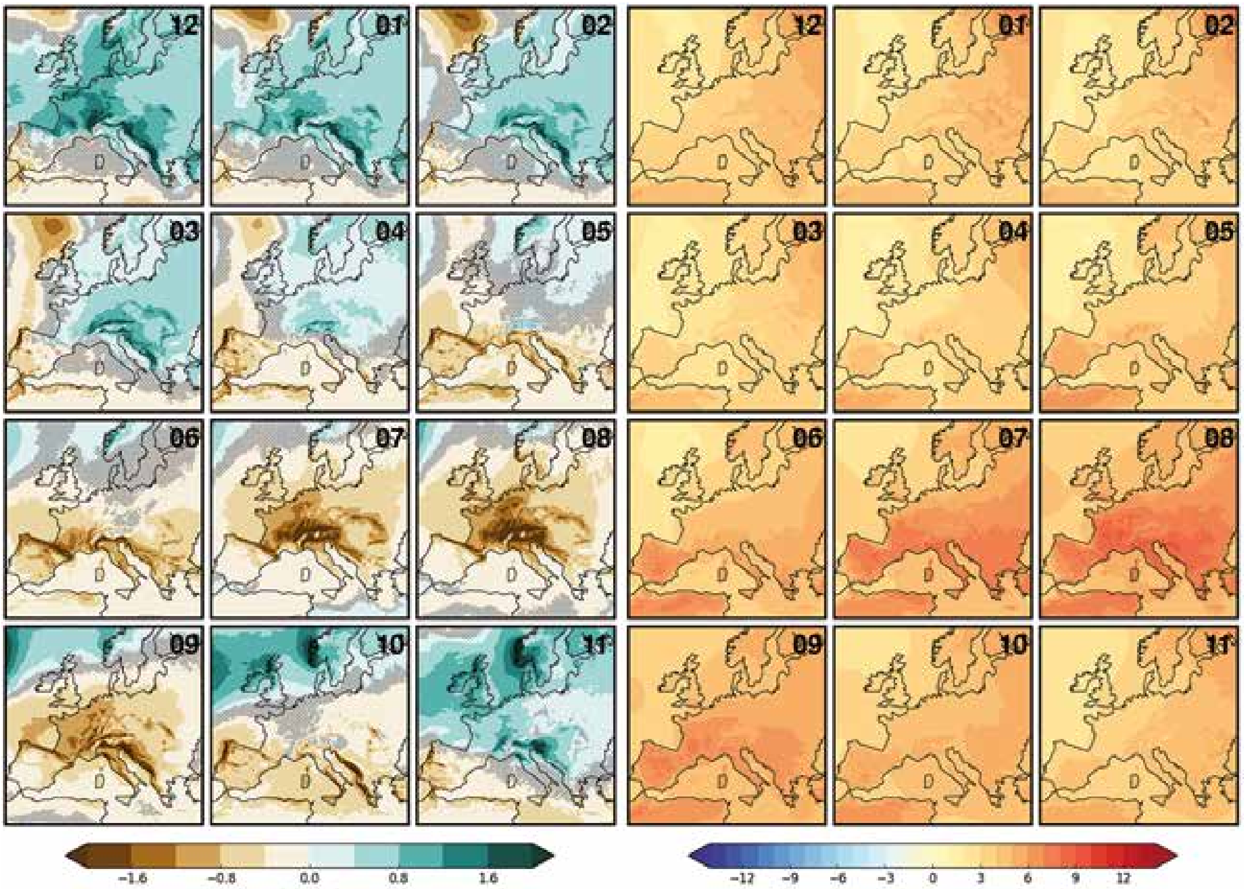

Using the CRCM5 climate model, researchers calculated climate changes in Europe and northeastern North America from 1950-2100. Comparing data from 2000–2019 and 2080–2099, the image on the left shows changes in monthly precipitation over Europe, while the image on the right shows the mean change in monthly surface-air temperature. Credit: Leduc, et al. © IKE, University of Stuttgart

SuperMUC to expedite its simulations, and was able to help the team optimize its code and manage its massive amount of data.

Forecasting the future

The team's simulations showed good agreement with historical climate data, leaving them confident in its predictive power and its ability to help improve impact models and regional adaptation strategies. Ludwig confirms that the team is sharing its data with the research community, and explains that the ClimEx experiment can help researchers study the future probabilities of extreme events such as heat waves, floods, and fires, and linking meteorological patterns with the development of these extreme events. This dataset helps scientists and government officials better evaluate flood risk projections and develop more robust methods to mitigate floods' impacts.

Under the worst-case scenario, in which carbon emissions continue to grow by roughly one percent per year, the model predicts that European summers be hotter by an average of up to 8 degrees Celsius per year starting in 2080, and that Quebecois winters would be up to 12 degrees warmer during the same time period.

"These projections are referencing the summer in Europe, and this is important, because this warming will occur at the same time as a decrease in precipitation, meaning that Europe could have much warmer and drier summers, which raises the possibility for more extreme heat waves and drought," Leduc said. "We should keep things in perspective, though. The model assumes a pathway for future greenhouse gas emissions, and that part is still uncertain. We don't know how much we will limit CO2 emissions in the future."

eg

SIMULATING ATOMIC NUCLEI, PHYSICISTS FOCUS ON LIVING LABORATORIES

University of Bonn researchers use simulation to understand subatomic phase transitions.

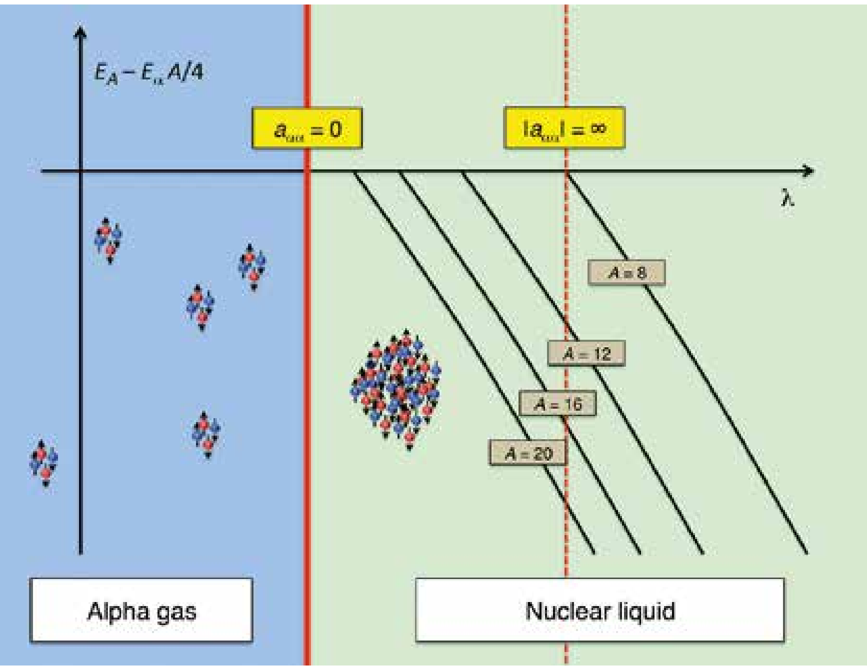

Zero-temperature phase diagram. **© University of Bonn & FZJ**

As technology has advanced, researchers from many different scientific backgrounds continue to take advantage of new tools for studying things at previously intractable sizes—both big and small. With the invention of the microscope, humans were first able to study how cells work, leading to breakthroughs in cellular biology. The same can be said for telescopes enabling new advancements in astronomy.

In the case of understanding the fundamental building blocks of matter—the building blocks of protons, neutrons, and electrons that, in turn, form atoms—extremely large experimental facilities, such as the Large Hadron Collider in Switzerland or the Spallation Neutron Source in the United States, were built so that researchers could study these objects experimentally, but as computing technology has advanced, researchers have come to rely upon high-performance computing (HPC) as a catalyst for new scientific discovery.

Recently, researchers from the University of Bonn, along with collaborators from other parts of Germany and the US, have been using HPC resources at the Jülich Supercomputing Centre (JSC) in order to understand the laws underpinning the subatomic world.

The long-time JSC users, led by Prof. Dr. Dr. Ulf Meißner, professor at the University of Bonn, have been studying how protons and neutrons bind to form nuclei, which is not only the centre of an atom, but also plays a central role in defining and ordering all matter.

Effective estimation and conscientious computing

All matter interacts in 4 basic forms—the weak and strong forces, electromagnetism, and gravitation. As its name implies, the strong force is the strongest fundamental interaction known to humanity. The strong interaction exerts force on quarks, elementary particles that, under the right circumstances, are held together by the strong force to form the neutrons and protons that comprise atomic nuclei.

Due to their extremely small size and extremely fast speed, quarks are very difficult to observe experimentally. In fact, quarks do not exist as free particles. When calculating subatomic particle interactions, researchers have to make certain assumptions about how subatomic particles structure themselves as well as about their positions at any given point in time. To this end, researchers have to use HPC to create a fine-grained grid, or lattice, then solve equations to determine particles' positions and properties.

Researchers have developed many different theories to try and calculate subatomic particles' behavior. As the field continues to evolve, it has become clear that many of these theories are capable of computationally recreating subatomic particle interactions and behavior under certain circumstances, but researchers have to make different assumptions or approximations when atomic nuclei are in higher or lower energy states. Meißner and his collaborators use effective field theory, a method that does not make assumptions about particles' behaviors and focuses instead on doing calculations that describe a system fundamentally under a specific energy range.

"In nuclear physics there are many well-known methods to study this topic," Meißner says. "They all have some uncontrolled approximations, though. We use effective field theory for studying the interactions of nucleons (protons and neutrons) and then use the lattice-based calculations to solve the many-body problem." The many-body problem is a fundamental equation in physics that deals with the feedback particles feel from one another in a system of many particles moving and interacting with one another)."

In its recent research, the team has been using its methods to study how alpha particles form and fade under a variety of circumstances. Alpha particles are particles whose nuclei are made of two protons and two neutrons. Their uniformity makes them relatively difficult to perturb.

The team found that alpha particles formed in lighter elements when at their respective ground energy states, but seemed to disappear as elements got heavier. The researchers realized that this quantum phase transition—a change of state dictated by energy rather than temperature or pressure—was actually a result of quantum entanglement of particles, which means that as particles move and increasingly come into contact with one another across their respective trajectories, they are influencing and changing the character of the system as a whole. This is the first time researchers were able to support this theory with simulation data.

Quantum collaboration

The team's runs on JSC's recently decommissioned JUQUEEN supercomputer regularly used 20,000 cores, and the team did scale its application to take advantage of the full machine. They recently ported their code to also take advantage of GPU accelerators, and are looking forward to further their goals with JSC'S newest machine, JUWELS.

Meißner credits the long-running collaboration with JSC physics experts for helping them more efficiently achieve their results. "In the beginning, this was a completely new code," Meißner says. "We had and have very good people in my group on the programming side, but we also had close contact with JSC staff and got tips to improve our algorithms. Our institute has a strong collaboration with JSC staff." In addition, the group was a big proponent of and offered advice and support for JSC's new modular supercomputing concept.

As supercomputing capabilities grow, the team looks forward to being able to more exhaustively chart many more nuclear reactions. Meißner referenced that chief among the team's goals is charting a computational course to solving the "holy grail" of nuclear astrophysics—accurately simulating oxygen generation occurring through the fusion of an alpha particle and carbon-12, which happens at extremely low energy scales that are difficult to study experimentally. Further, the team wants to exhaustively chart the "drip lines" across the periodic table, which requires a significantly different computational approach.

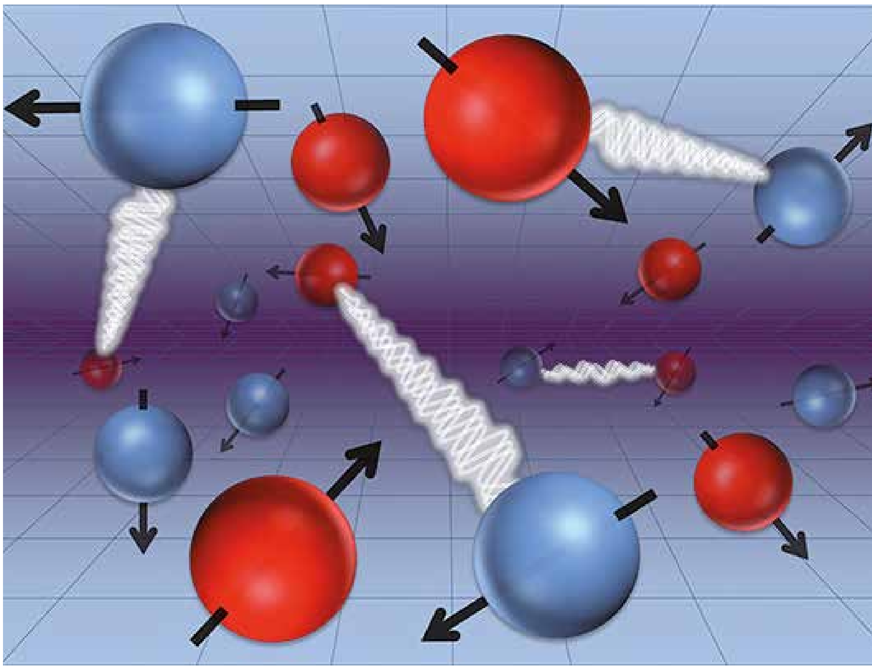

Schematic representation of the entanglement of nucleons that leads to the disappearance of alpha-clusters in heavier nuclei. **© University of Bonn & FZJ**

"In the nuclear chart, you can add protons or nucleons to different elements, but at a certain point these nuclei become instable and neutrons or protons starts to drip off," Meißner says. "Of course, where these boundaries are located is a difficult question, because these structures are weakly bound. To find out where in the nuclear chart these drip lines are located, you are essentially figuring out how many protons and neutrons can you add to an atom while keeping it stable. You need high precision calculations which need a lot of computing time." *eg*

SUPERCOMPUTING HELPS STUDY TWO-DIMENSIONAL MATERIALS

HPC helps researchers understand experiments for observing real-time motion of lithium atoms in bi-layer graphene, paving the way for designing new materials for batteries and other electronics.

Whether it is high-temperature superconductors and improved energy storage or bendable metals and fabrics capable of completely wicking liquids, materials scientists study and understand the physics of interacting atoms in solids to ultimately find ways to improve materials we use in every aspect of daily life.

The frontier of materials science research lies not in alchemical trial and error, though; to better understand and improve materials today, researchers must be able to study material properties at the atomic scale and under extreme conditions. As a result, researchers have increasingly come to rely on simulations to complement or inform experiments into materials' properties and behaviours.

A team of researchers led by Dr. Arkady Krasheninnikov, physicist at the Helmholtz-Zentrum Dresden-Rossendorf, partners with experimentalists to answer fundamental questions about materials' properties, and the team recently had a big breakthrough—experimentalists were able to observe in real time lithium atoms' behaviour when placed between two graphene sheets. A graphene sheet is what researchers consider a 2D material, as it is only one atom thick, which made it possible to observe lithium atom motion in transmission electron microscopy (TEM) experiments.

With access to supercomputing resources through the Gauss Centre of Supercomputing (GCS), Krasheninnikov's team was able to use the High-Performance Computing Center Stuttgart's (HLRS') Hazel Hen supercomputer to simulate, confirm, and expand on the team's experimental findings. The collaborative work was recently published in Nature.

"2D materials exhibit useful and exciting properties, and can be used for many different applications, not only as a support in TEM," Krasheninnikov says. "Essentially, 2D materials are at the cutting edge of materials research. There are likely about a couple thousands of these materials, and roughly 50 have actually been made."

Under the microscope

To better understand 2D materials experimentally, researchers routinely use TEM nowadays. The method allows researchers to suspend small, thin pieces of a material, then run a high-energy electron beam over it, ultimately creating a magnified image of the material that researchers can study, much like a movie projector takes images from a reel and projects them onto a larger screen. With this view into a material, experimentalists can better chart and estimate atoms' positions and arrangements.

The high-energy beam can do more than just help researchers observe materials, though—it is also a tool to study 2D materials' electronic properties. Moreover, researchers can use the high-energy electrons from TEM to knock out individual atoms from a material with high precision to see how the material's behavior changes based on the structural change.

Recently, experimentalists from Max Planck Institute for Solid State Research, Stuttgart and the University of Ulm wanted to better understand how lithium particles interacted between two atom-thin graphene sheets. Better understanding lithium intercalation, or placing lithium between layers of another material (in this case, graphene), helps researchers develop new methods for designing better battery technologies. Experimentalists got data from TEM and asked Krasheninnikov and his collaborators to rationalize the experiment using simulation.

Simulations allow researchers to see a material's atomic structure from a variety of different angles, and they also can help speed up the trial-and-error approach to designing new materials purely through experiment. "Simulations cannot do the full job, but they can really limit the number of possible variants, and show the direction which way to go," Krasheninnikov says. "Simulations save money for people working in fundamental research and industry, and as a result, computer modelling is getting more and more popular."

In this case, Krasheninnikov and his collaborators found that the experimentalists' atomic coordinates, or the positions of particles in the material, would not be stable, meaning that the material would defy the laws of quantum mechanics. Using simulation data, Krasheninnikov and his collaborators suggested a different atomic structure, and when the team re-ran its experiment, it found a perfect match with the simulation.

"Sometimes you don't really need high theory to understand the atomic structure based on experimental results, but other times it really is impossible to understand the structure without accurate computational approaches that go hand-in-hand with the experiment," Krasheninnikov says.

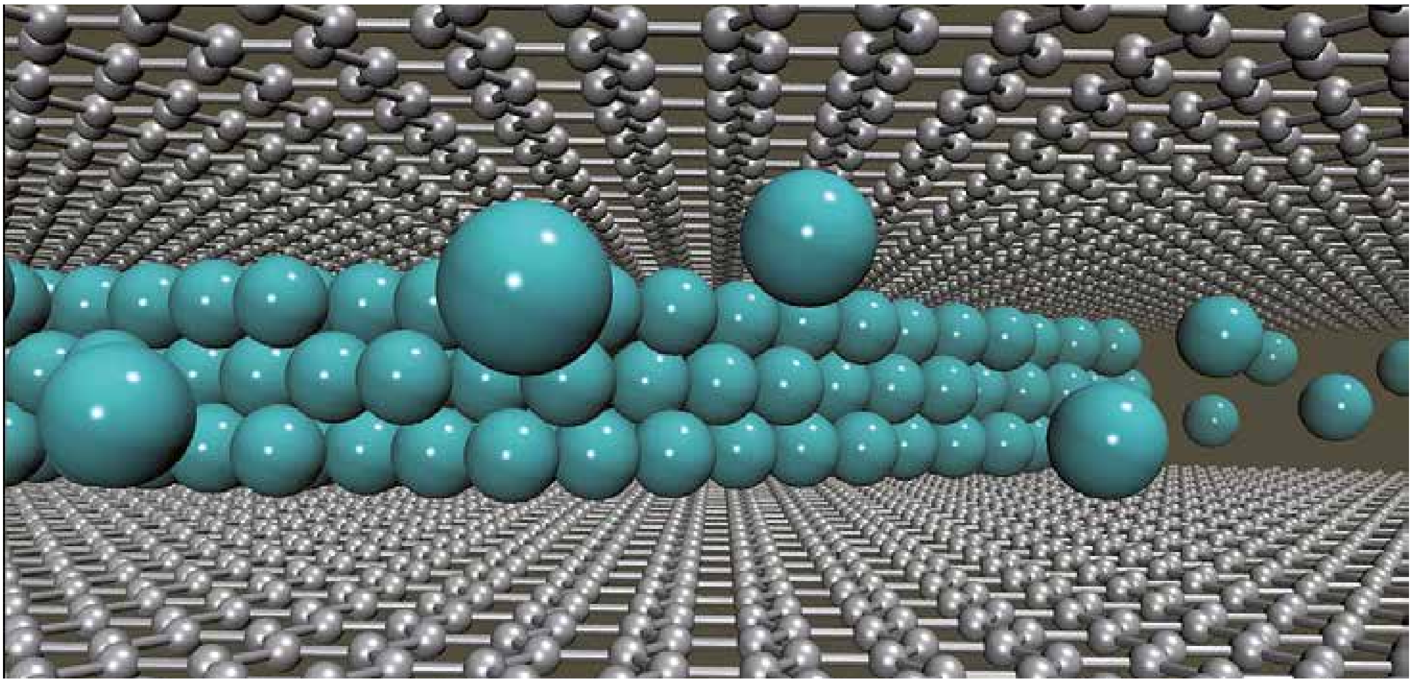

Atomistic model illustrating a multilayer of lithium atoms between two graphene sheets.

© Courtesy of Dr. Mahdi Ghorbani-Asl, HZDR. See M. Kühne, et al., Nature 564 (2018)

The experimentalists were able to, for the first time, watch in real-time how lithium atoms behave when placed between two graphene sheets, and with the help of simulations, get insights into how the atoms were arranged. It was previously assumed that in such an arrangement, the lithium would be structured as a single atomic layer, but the simulation showed that lithium could form bi- or trilayers, at least in bi-layer graphene, leading researchers to look for new ways to improve battery efficiency.

Charging forward

Krasheninnikov noted that, while simulation has made big strides over the last decade, there is still room for improvement. The team can effectively run first-principles simulations of 1,000-atom systems over periods of time to observe short-term (nanosecond time scale) material interactions. Larger core counts on next-generation supercomputers will allow researchers to include more atoms in their simulations, meaning that they can model more realistic and meaningful slices of a material in question.

The greater challenge, according to Krasheninnikov, relates to how long researchers can simulate material interactions. In order to study phenomena that happen over longer periods of time, such as how stress can form and propagate a crack in metal, for example, researchers need to be able to simulate minutes or even hours to see how the material changes. That said, researchers also need to take extremely small time steps in their simulations to accurately model the ultra-fast atomic interactions. Simply using more compute cores allows researchers to do calculations for larger systems faster, but cannot make each time step go faster if a certain “parallelization” threshold is reached.

Breaking this logjam will require researchers to rework algorithms to more efficiently calculate each time step across a large amount of cores. Krasheninnikov also indicated that designing codes based on quantum computing could enable simulations capable of observing material phenomena happening over longer periods of time—quantum computers may be perfect for simulating quantum phenomena. Regardless of what direction researchers go, Krasheninnikov noted that access to supercomputing resources through GCS and PRACE enables him and his team to keep making progress. “Our team cannot do good research without good computing resources,” he said. *eg*

EUROPEAN PROCESSOR INITIATIVE BUILDS IN-EUROPE PROCESSOR TECHNOLOGIES

JSC staff lead co-design of EU-based hardware technologies for HPC and beyond.

Very few countries and regions in the world are currently able to deliver the high-end processor technologies that drive our digital era. As of today, Europe still relies on others for processing power. This is a significant risk for Europe's future competitiveness and innovation, as such technologies are export-controlled, meaning that access to them strongly depends on the policies of governments outside of Europe.

To change this scenario, the European Commission is setting up a strategy to support the development of European processor technologies through the EU-funded project: European Processor Initiative (EPI) [1]. EPI aims at creating a new family of high-performance processor technologies designed and manufactured in Europe. The EPI processing units address the requirements of high-performance computing (HPC), but also keep much larger market sectors in mind, including the automotive, cryptography, artificial intelligence industries, and trusted IT infrastructures, among others. With this approach, EPI will deliver competitive and commercially successful computing technologies, reducing Europe's dependency from international (USA and Asia) providers.

EPI's first phase started on the December 1, 2018, and will run for 3 years with a budget of €80 million. The project is based on a long-term framework partnership agreement (FPA) between the European Commission and 23 partner institutions, under the leadership of Bull/ATOS. As an active member of the consortium, JSC is leading the co-design approach, applied to ensure that the EPI processors meet the requirements of the eventual users. EPI researchers analyse a wide range of applications and develop processor architecture models in order to identify how hardware and software design parameters influence the performance achievable with real codes. This information is fed back to the system designers, who can then establish design trade-offs and take informed decisions.

In the first 3 years, the EPI partners will develop a processor combining general purpose and accelerator units. Both will

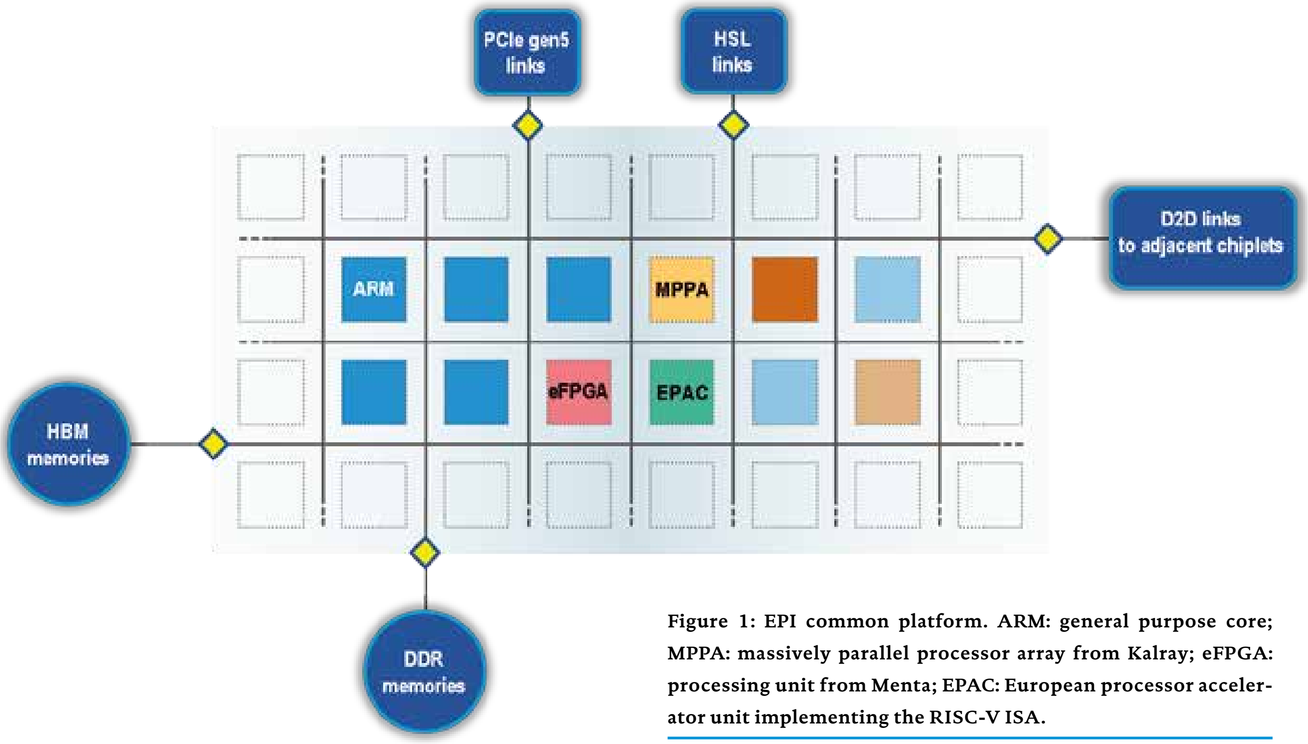

Figure 1: EPI common platform. ARM: general purpose core; MPPA: massively parallel processor array from Kalray; eFPGA: processing unit from Menta; EPAC: European processor accelerator unit implementing the RISC-V ISA.

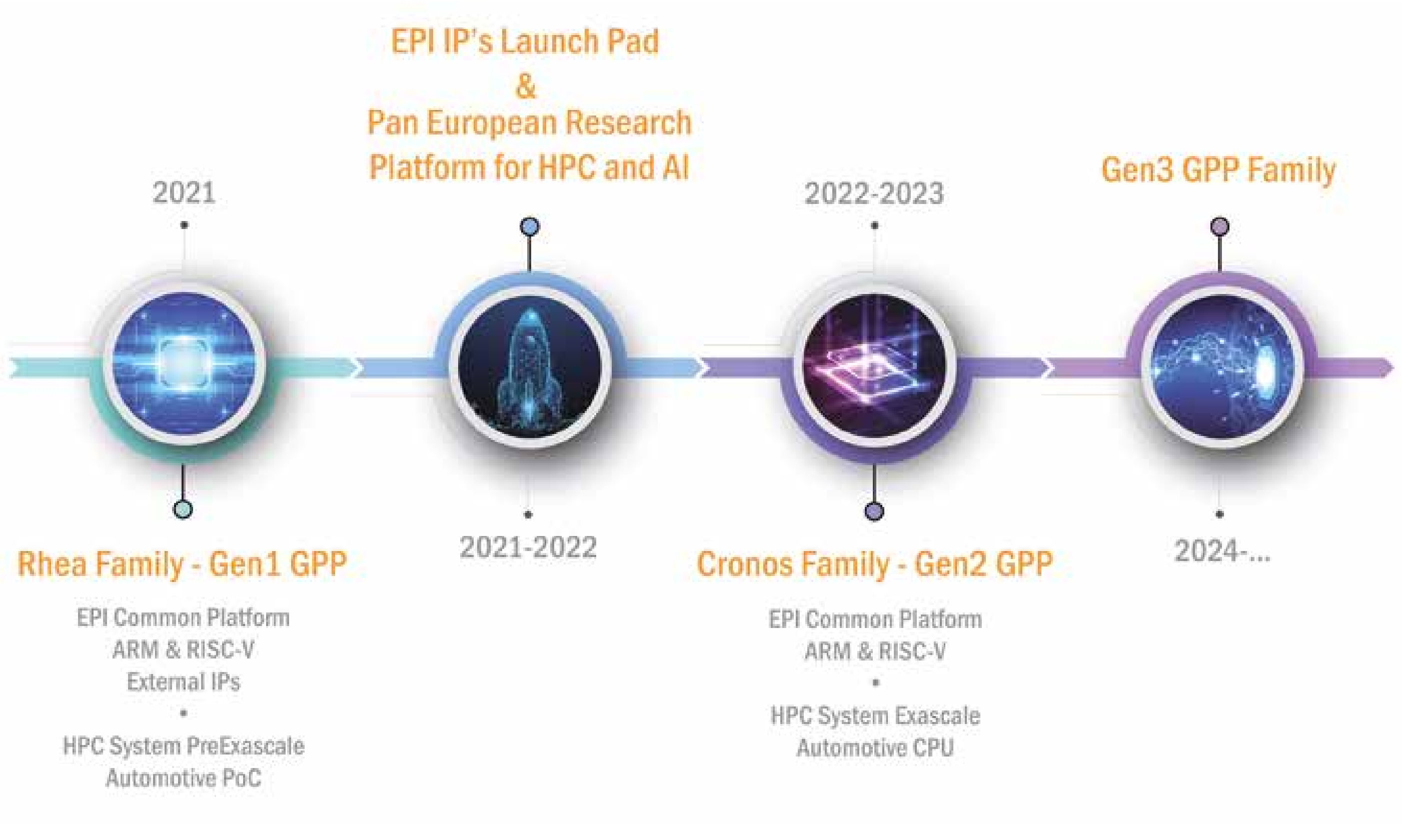

be integrated as chiplets into a common platform: a modular approach to chip construction that allows building different processor "flavours" by selecting the type and number of chiplets according to the targeted goal and market (see Figure 1). The first generation chip, called 'RHEA', is planned to be built and validated in 2020, with production-ready chips to become available in 2021. RHEA will include general purpose cores of the ARM architecture, and acceleration units implementing the RISC-V instruction set architecture, the Kalray MPPA architecture, and the eFPGA architecture from Menta. The RHEA chips will be integrated into test platforms in order to validate the hardware units, develop the necessary software interfaces, and run applications.

The long-term EPI roadmap goes well beyond the RHEA generation and foresees further processing units coming in in 2022-2023 (see Figure 2). The "Cronos" processor family will be the second EPI chip generation and aims to support Exascale performance. In fact, it is one of EPI's target to deliver competitive European processor units for at least one of the Exascale systems that EuroHPC targets for 2023. *es, dp*

Figure 2: EPI processor development roadmap

Reference:
https://www.european-processor-initiative.eu/

Acknowledgements:
This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 826647

Author:
Estela Suarez, Jülich Supercomputing Centre
e.suarez@fz-juelich.de

INVESTMENTS IN INFRASTRUCTURE HELP LRZ ADVANCE AI USERS THROUGH HPC

State-funded program supports Bavarian digitalization efforts, portability of AI applications.


© shutterstock

As high-performance computing (HPC) technologies have continued to advance, so too has the number of scientific communities benefitting from supercomputing to accelerate discovery. In recent years, the rise of artificial intelligence (AI)—along with its associated fields of data analytics and machine learning—have presented HPC centres with both new challenges and new opportunities to innovate.

Data science researchers create and use new methods for modelling complex phenomena from vast amounts of data with many variables. Massive computing is needed to crunch such calculations and to deliver significant results for many socially relevant research topics–among them are city planning, public health, and environmental response.

While HPC plays a significant role in developing and improving AI, requirements for research teams using these applications differ significantly from those primarily using long-established modelling methods. Data researchers are also challenged with figuring out how to use general-purpose AI software efficiently for their specific workflows.

Realizing the need to facilitate the merger of AI with HPC for greater scientific gain, the Bavarian Ministry for Science and the Arts invested to build upon LRZ's existing capabilities and to establish it as a Bavarian Big Data Competence Centre. As part of the funding, LRZ has made significant investments in staffing, training, and infrastructure to help not only further integrate data-intensive technologies into an HPC environment, but to also allow researchers to spend less time on optimizing and porting AI software, thereby gaining them more time to focus on science.

Recently, LRZ staff members have focused on workflows for high-performance artificial intelligence (HPAI) applications. When research teams use commercial machine learning applications on HPC resources, they have to work to port the application to run efficiently on supercomputers. When researchers need multiple AI applications in their workflows, these applications can adversely impact the efficiency or functionality of previously ported applications.

To address the problem, LRZ has integrated CharlieCloud containers into the centre's capabilities. Developed by the US Department of Energy specifically for specialized HPC systems, these containers act as secure environments that can structure sets of software so that AI applications perform their specific calculations without interfering with one another. In addition to making complex workflows far simpler to deploy on LRZ systems, this containerization work also enables researchers to more easily port their workflows onto different computing architectures, such as those at fellow GCS centres the High-Performance Computing Center Stuttgart (HLRS) and the Jülich Supercomputing Centre (JSC).

The first results of this containerized approach show promising results, especially on LRZ's current leadership HPC system SuperMUC-NG. In addition to expanding and streamlining AI workflows on HPC systems, LRZ has added further dedicated hardware solutions to its AI capabilities, such as two NVIDIA DGX-1 deep learning systems for development and refinement of AI applications on a supercomputing platform.

Hardware, software, and workflows all come together in LRZ's education and training program. A broad catalogue of AI-specific courses caters to anyone from beginners to experienced HPC practitioners and includes programming basics with R and Python, machine and deep learning, and HPC cluster usage. The catalogue is completed with offerings from LRZ's partner vendors including Intel€ AI and NVIDIA DLI. For more information on current courses, visit LRZ's website at https://www.lrz.de/services/schulung/.

"LRZ is well equipped to leapfrog AI workflows through smart integration with HPC systems," says Prof. Dieter Kranzlmüller, Chairman of the Board and Managing Director of LRZ. "We're responding to the demand and challenges associated with rapid digitalization and we are clearing the way for the Bavarian scientific ecosystem to capitalize on the future of extreme-scale AI." *eg*

BIG DATA TECHNOLOGIES ON THE RISE IN STUTTGART

The CATALYST projects combines HPC and high-performance data analytics (HPDA) to push innovation in academia and industry.

The High-Performance Computing Center Stuttgart (HLRS) prides itself in having a diverse user base across a variety of scientific disciplines, as well as researchers coming from both academia and industry.

Despite their different needs, the need for executing more and more data-intensive applications seems to be an increasingly common theme. Since high-performance computing (HPC) systems, and, in turn, overall computing power continually increases, several petabytes of data at the end of a simulation is no longer uncommon, making the original approach of manually validating these datasets no longer feasible.

To address this increasingly complex problem, HLRS acquired the Cray Urika-GX in 2016 to enable users to more efficiently process and analyze huge amounts of complex data, placing HLRS among the first HPC centres in the world to operate such hardware, enabling the centre to delve deeper into HPDA-related projects.

Federal funding helps regional organizations exploit Big Data Technologies

However, with HPDA being a relatively fresh topic in the HPC community, hardware has to be fed with real-life use cases in order to assess a new HPC technology's potential. In an effort to bolster the Stuttgart region as an innovation hub and find out how local organizations can benefit from dedicated hardware running HPDA workflows, the Baden-Württemberg Ministry of Science and the Arts (MWK) funded the project CATAYST in 2016.

Stuttgart-based software company nFrames was among the first users of the Urika-GX system. The company develops the popular 3D surface reconstruction software SURE, which supports large-scale mapping and monitoring projects through photogrammetry – a set of measures to determine the surface and position of objects derived from airborne images. With modern camera systems, both the data size of each individual image as well as the sheer volume of images has increased drastically.

HPDA proves to be a powerful tool for fast-response applications

While technological advances have led to a huge leap in the quality of digital models in terms of precision and resolution it also increases the processing time for these large datasets. Emerging application areas for 3D digital modeling, such as disaster response and the insurance sector, however, require up-to-date data with immediate availability and high resolution.

In order to help solve this problem, nFrames teamed up with HLRS to test distributed processing capabilities that were introduced to the SURE software in order to run a project on high-performance hardware with high CPU core counts, fast disk storage, and a high-bandwidth interconnect – such as the Cray Urika-GX system – ultimately leading to a shorter time-to-solution. Showcasing that their rapid-response application is effective on massive datasets when run on supercomputers, the team chose a dataset consisting of roughly 8,000 images covering a 76-square-kilometer area of San Francisco, California, resulting in a total resolution of 796 Gigapixels.

While the processing of such a massive dataset would have taken about three months on the company's most powerful system, the entire 3D model could be reconstructed and made ready for delivery within just three days thanks to the collaboration with CATALYST. "With combined efforts we could prove that the combination of HPC with aerial image processing in 3D reconstruction shows great potential for the development of new fast-response applications," nFrames CTO Thomas Zwölfer states. "In the long term, this approach could allow for quick reaction times required for disaster management."

CATALYST puts Big Data on the school curricula

Part of CATALYST's mission is also to educate the scientific community and public about HPDA's value as a tool to tackle a wide variety of challenges. In September 2018, CATALYST provided summer school participants from the Albstadt-Sigmaringen University with access to the Cray Urika-GX system, including pre-installed software stacks

This illustration shows a textured 3D Mesh of San Francisco. The data was provided courtesy of Geomni © nFrames

for HPDA, such as Apache Hadoop and Apache Spark. The university students were granted access to the system via virtual machines, enabling them for the first time to work on real-life use cases from two partnering companies within professional HPDA environments.

Additionally, researchers working on CATALYST are collaborating with the outreach project "Simulated Worlds," an initiative also funded by the MWK aimed at preparing high school students for careers in the digital realm. In 2018, one student used the HPDA environment to successfully predict train delays in the Stuttgart area with the help of machine learning. Recently, a number of high school students are working with a large movie database in order to apply big data analytics principles to identify, among other things, the main ingredients to make a movie a box office hit, considering a wide range of variables like cast and genre.

"CATALYST successfully showcases the opportunities that data analytics offers academia and, in particular, industry to strengthen their competitiveness in today's data-centric markets," project manager Dennis Hoppe says. "That's why we are currently working on additional funding opportunities in order to keep up with this development in the future." *lb*

Funding:
1.7 Million Euro

Runtime:
10/2016 – 9/2019

Funding Organisation:
Baden-Württemberg Ministry of Science and the Arts

Collaborators:
HLRS, Cray Inc.

JSC HOSTS HPC OPTIMISATION AND SCALING WORKSHOP

To help make the most effective use of its supercomputing systems, JSC staff organized the HPC Optimisation and Scaling Workshop, which took place February 18–22. The event combined training from Intel and JSC in the use of their parallel application performance analysis tools and gave participants the the opportunity to do large-scale runs and measurements to validate correctness, as well as investigate and improve scaling their codes. Over twenty participants from all over Germany (often working in teams of two to four) were supported by a similar number of staff from the JSC simulation laboratories and cross-sectional teams, three instructors from Intel, and additional HPC systems support staff assisting in the background. JUWELS and JURECA (cluster and booster parts) were reserved during the workshop so that participants could run jobs with large numbers of compute nodes.

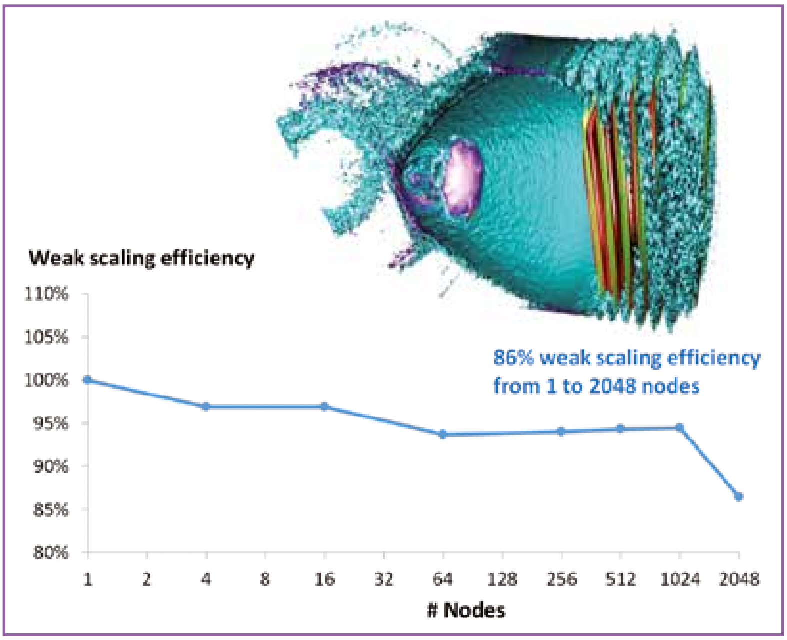

Figure 1: PIConGPU weak scaling efficiency on JUWELS.

Participants were encouraged to provide their codes with test cases and/or execution measurements for auditing by the Performance Optimisation and Productivity (POP) Centre of Excellence [2] in advance of the workshop. Preliminary analyses of several codes could be briefly presented at the start of the workshop and incorporated within plans developed during the event. For instance, a team from RWTH Aachen identified a load imbalance in their CIAO code for turbulent reacting flows that could be reduced with an alternative workload specification. A team from Göttingen University found that communication within and between blocks of processes could be optimised in its SOMA simulation, which focuses on polymer melts. Assessments are continuing in most cases.

The PIConGPU plasma simulation code, developed by HZDR in collaboration with TU Dresden, ran its largest CPU-only execution on 2,048 compute nodes of JUWELS. 95% weak scaling efficiency was maintained to 1,024 compute nodes when running a single MPI process on each processor with OpenMP threads on each core which minimized both inter-node communication and non-local memory accesses within nodes [*Figure 1*].

Two other codes which effectively employed multi-threading combined with MPI could also readily scale to run successfully with up to 2,048 compute nodes of JUWELS, and two codes using only MPI (and no threading) required much more tuning effort to use 1,024 compute nodes with very large numbers of MPI processes. Using the 1D-NEGF code, a team from SimLab Quantum Materials explored large-scale modular supercomputing combining JURECA cluster and booster nodes for inelastic quantum charge transport in semiconductors: the less scalable Poisson equation solver being run on cluster nodes and the highly scalable remainder on the booster nodes. Other code-teams focused on optimising memory accesses and vectorisation. The workshop participants were very grateful for the opportunity and assistance provided.

In the course of this thorough exercising of system hardware and software, participants and staff encountered a variety of issues, not all of which could be resolved during the workshop, but are still being actively being addressed. The workshop's findings will inform improvements to usage of the supercomputing systems, allowing users to solve larger computational problems and use HPC resources more efficiently. *bw*

[1] http://www.fz-juelich.de/ias/jsc/EN/Research/Projects/_projects/sivegcs.html
[2] https://www.pop-coe.eu/

LOOKING INSIDE SIMULATION'S BLACK BOX

At the 2018 Science and Art of Simulation Workshop, philosophers, sociologists, historians, and computer scientists sought to develop a more rigorous framework for evaluating the trustworthiness of computational science.

Dr. Andreas Kaminski, leader of the HLRS Philosophy of Science & Technology of Computer Simulation Program, introduced the concept of epistemic opacity. **© HLRS**

Science and mathematics have long been celebrated for their ability to shed light on previously unknown features of the natural world. Through logic, experiment, proof, and self-critique within the scientific community, science was seen as a march toward increasing clarity.

In recent years, however, the rise of computer simulation and machine learning has complicated this view of science. The ability to create and apply complex algorithms running on computers has undoubtedly opened up new kinds of opportunities for scientific discovery and technological innovation. Nevertheless, as "black boxes" whose inner workings are partly impossible to observe and understand, even the scientists who use them often do not know exactly how their computer programs generate their results. In such a nontransparent situation, how can we be confident that we can trust the simulations our technologies produce?

Philosopher Paul Humphreys coined the term "epistemic opacity" in 2004 to describe this situation. Since then the concept has spurred much discussion among philosophers, sociologists, and simulation scientists, although developing a consistent, comprehensive, and theoretically grounded understanding of epistemic opacity in all of its forms has been difficult. Doing so would not only be intellectually significant but could also have the practical effect of providing a useful framework for evaluating the results of simulation and machine learning. Such knowledge could also potentially help to improve the design of new applications in these fields.

On November 28-30, 2018, the Philosophy of Science & Technology of Computer Simulation Program at the High-Performance Computing Center Stuttgart (HLRS) hosted an interdisciplinary conference to take steps toward this goal. Bringing together an international community of investigators, the event sought to achieve a more refined understanding of how the simulation sciences differ from more traditional disciplines. The event also sought to formulate a deeper understanding of the many sources and flavors of epistemic opacity.

The meeting spotlighted sources of epistemic opacity that are unique to computer-intensive research. In particular, the combination of advanced mathematics and information technology in high-performance computing constitute a methodological opacity. Here, scientists must rely on results produced by algorithms whose mechanistic operation is impossible to observe directly.

In addition, social characteristics of the scientific communities that make up interdisciplinary teams affect almost all research in science and engineering. In large-scale projects, for example, individual scientists might bring expertise in applying one specific type of analysis, but lack in-depth knowledge about the workings of other parts of the project on which their own work relies.

Because science relies on the comprehensibility of the scientific method, there was a broad consensus at the meeting that epistemic opacity touches upon fundamental questions of what science is and how the scientific method should be understood. The rise computer-intensive science, leading to new sources of opacity, suggests that the scientific method may be undergoing further change.

Proceedings from the 2018 Science and Art of Simulation workshop will be published by Springer in 2019. *cw*

NEXT-GENERATION SUPERCOMPUTING FOSTERS NEXT-GENERATION SCIENCE

Multidisciplinary researchers converge on LRZ for the Next-Generation Science Symposium.

When it comes online in 2019, the SuperMUC-NG supercomputer at the Leibniz Supercomputing Centre (LRZ) will enable researchers in Germany and the rest of Europe to climb to previously unattainable heights, resulting in new scientific discovery.

To get a first impression of SuperMUC-NG's capabilities and application areas, researchers converged on LRZ in November, 2018 for the "SuperMUC-NG Next-Gen Science Symposium." In addition to researchers, representatives of the German Federal Ministry for Education and Research (BMBF) and the Bavarian State Ministry for Science and the Arts also attended the event.

Prof. Dieter Kranzlmüller, Chairman of the Board of LRZ, spoke of the importance and relevance of the event in the context of closer collaboration between scientists and HPC centres. "In addition to installing the hardware for SuperMUC-NG, one of our major missions is supporting our users by providing the necessary support and know-how in order to make sure their projects are in a position to succeed."

The symposium covered a wide spectrum of themes, ranging from the general question, "Quo vadis HPC? (where is HPC going?)" in the realm of classical HPC disciplines such as astrophysics or engineering as well as earth science, climate research, health, and the use of gigantic data sets from satellite measurements in order to improve living conditions on Earth. With the performance increase offered by SuperMUC-NG, researchers will be able to develop a deeper understanding of the origins and development of our universe, or dive deeper into the Earth's interior to understand the geological processes that impacted the planet's development.

In the realm of health and medicine, researchers can use SuperMUC-NG for other valuable breakthroughs. Scientists focused on health-related research will be able to create hyper-accurate simulations to study how blood flow changes with aneurysms, develop high-resolution simulations of air flow in the lungs, or better study the effects of various medications on individuals. Throughout the course of the event, it became clear that supercomputing simulations have only deepened their value in many different scientific disciplines, becoming indispensable tools for helping solve not only some of the world's most intensive scientific challenges, but also the smaller-scale challenges in health and coexistence that humanity faces in its collective day-to-day life. *lp*

The Leibniz Supercomputing Centre (LRZ) in Garching near Munich. **© LRZ / Christoph Rehbach**

HLRS, PARTNERS SIGN CONTRACT FOR NEXT-GENERATION SUPERCOMPUTER

On November 30, 2018, representatives from Hewlett Packard Enterprise, the University of Stuttgart, and HLRS signed a contract to build Hawk, HLRS's next-generation supercomputer. The machine, which will have a theoretical peak performance of 24 petaflops, represents a nearly four-fold increase in computing power over its current flagship system, Hazel Hen. When installed this fall, it will be one of the first HPC systems in the world to use AMD's newest EPYC processor. "The real winners will be our user community of computational engineers in academic and industrial research who will benefit from the ability to run more complex simulations," says Prof. Michael Resch, Director of HLRS.

HLRS CELEBRATES FIRST ANNIVERSARY OF THE SOCIOPOLITICAL ADVISORY BOARD

One year ago, HLRS began a sociopolitical advisory board to help guide its mission. On April 1, 2019, the board members—who come from different areas of society and representing a wide range of socially relevant topics, such as nursing and education—met again to discuss current developments at HLRS. During the meeting, the board learned that some of last year's discussion topics would be playing a major role in a new project proposal that aimed to extend the HLRS Philosophy of Simulation Department's work. The group focuses on assessing simulation's ability as an instrument for exploring wide-ranging possibilities in solving problems of particular social importance.

NRW MINISTER PFEIFFER-POENSGEN VISITS JSC

On March 22, Isabel Pfeiffer-Poensgen, Minister for Culture and Science of North Rhine-Westphalia, visited Forschungszentrum Jülich and was also a guest at the Jülich Supercomputing Centre (JSC). During her visit, Prof. Thomas Lippert, head of JSC, informed her about the new supercomputer JUWELS, half of which was financed by the state of North Rhine-Westphalia. He also elucidated to her the centre's long-time expertise and competence in the field of supercomputing and presented plans on how the research centre's expertise in supercomputing, machine learning, and artificial intelligence (AI) can be used for structural change in the Rhenish lignite mining area. In addition, Prof. Kristel Michielsen (JSC) demonstrated how new quantum technologies are being researched in Jülich and what role quantum computing plays at JSC.

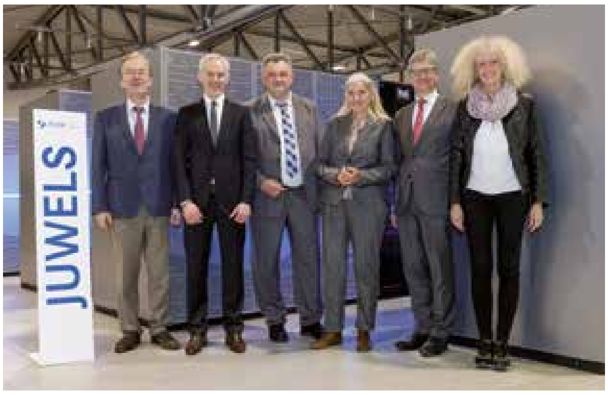

HELMHOLTZ GPU HACKATHON AT JSC

JSC hosted the 2019 Helmholtz GPU Hackathon in April, 2019. Sixty participants from all over the world participated in accelerating their scientific applications on GPU devices. The ten teams came from a variety of scientific domains, including physics (specifically, quantum chromodynamics), neuro-science, climate modeling, and materials science. Paired with mentors from academia and industry, they worked intensely from morning until evening filling up the GPU queues of JUWELS. For some participants, the Hackathon was their first steps toward optimizing their applications for GPUs, while other applications were already apt in GPU usage and used the present expert knowledge to further fine-tune their codes. All teams were able to improve the GPU performance of their code and could formulate a plan for further improvement for the near future.

LRZ WELCOMES NEW FACES TO BOARD OF DIRECTORS

At the beginning of February, 2019, LRZ announced its new directorate. Former board members, both long-running board members and former Chairmen of the Board of LRZ—Prof. Dr. Heinz-Gerd Hegering (1989–2008) and Prof. Dr. Arndt Bode (2008–2017)—were seen off with a celebration of their time and commitment to LRZ. LRZ welcomes Prof. Dr. Martin Schulz, Chair of Computer Architecture and Parallel Systems at the Technical University of Munich (TU Munich), and Prof. Dr. Thomas Seidl, Chair of Database Systems and Data Mining at Ludwig-Maximilians-Universität München (LMU) to the board of directors. In addition to the newcomers, Prof. Dr. Hans-Joachim Bungartz, Chair of Scientific Computing at TU Munich, and Prof. Dr. Dieter Kranzlmüller, Munich Network Management Team (MNM-Team) at LMU, remain part of the LRZ directorate.

From left to right: Dieter Kranzlmüller, Hans-Joachim Bungartz, Martin Schulz, Thomas Seidl, Heinz-Gerd Hegering, and Arndt Bode

SIMTECH NAMED A NATIONAL CENTER OF EXCELLENCE

In 2007 the University of Stuttgart founded the SimTech Center of Excellence to foster a robust, multidisciplinary community focused on simulation. Recognizing the university's role as a national leader in simulation technology and applications, the German Research Foundation (DFG) announced that SimTech was selected as a national center of excellence. The new, seven-year grant will focus on increasing capacities for data-driven simulation. HLRS is a core partner in the SimTech initiative, and HLRS Director Michael Resch serves as the principal investigator for high-performance computing systems within the research consortium.

LONG-TIME GCS USERS' COSMOLOGY SIMULATION HONORED ON POSTAGE STAMP

The record-breaking galaxy formation simulation, Illustris, can now adorn letters across Germany. In 2014, a multi-institution team led by researchers at the Heidelberg Institute of Theoretical Studies (HITS) created the largest cosmological simulation of galaxy ever created. The computational project, dubbed Illustris, allowed researchers better understand how galaxies sprouted up across the universe, and was run on the Gauss Centre for Supercomputing's (GCS') SuperMUC supercomputer at the Leibniz Supercomputing Centre. In 2018, the team broke its own record with the "Illustris: The Next Generation" simulation at another GCS member centre, the High-Performance Computing Center Stuttgart. In honor of the team's contributions to cosmology over the last several years, the German Federal Post Office (Deutsche Post) recently released as postage stamp commemorating the simulation.

MEDIA SOLUTION CENTER WILL PROMOTE INNOVATION IN TECHNOLOGIES FOR FILM AND DIGITAL ART

In October, 2018, HLRS partnered with the center for Art and Media Karlsruhe (ZKM) and Hochschule der Medien Stuttgart (HdM) to found the Media Solution Center Baden-Württemberg (MSC). The MSC will support collaborative research projects aimed at solving problems that media producers face, particularly related to animation, simulation, and visual effects. The MSC will also facilitate knowledge and technology sharing, including continuing education activities and events focused on themes important to media professionals. The MSC is open to membership of media professionals and companies, as well as leaders of relevant scientific and technical research groups.

SUPERMUC-NG ENTERS ITS IMPLEMENTATION PHASE

After the setup of SuperMUC-NG at LRZ had been finished, Bavaria's Prime Minister, Dr. Markus Söder, symbolically started the start-up phase. He was accompanied by Bavaria's State Minister for Science and the Arts, Prof. Dr. med. Marion Kiechle, Prof. Dr. Dieter Kranzlmüller, Head of LRZ, and Prof. Dr. Thomas O. Höllmann, President of the Bavarian Academy of Sciences and Humanities. SuperMUC-NG follows as the "Next Generation" to its precursor, SuperMUC, which went into operation in 2012 and was extended in 2015.
From left to right: Marion Kiechle, Dieter Kranzlmüller, Thomas O. Höllmann, and Markus Söder starting the start-up phase of SuperMUC-NG.

HLRS JOINS NEW PROJECT TO IMPROVE DIGITAL LITERARY RESEARCH

HLRS has partnered with the German Literature Archive and the University of Stuttgart's Institute for Literary Studies and Institute for Natural Language Processing to develop a new platform for the management and analysis of data for literary research. Focusing on "born-digital" literature — such as writers' digital estates and texts published on blogs and other websites — the project aims to provide a sustainable and scalable archive for storing, organizing, and accessing digital data; an integrated pipeline of tools for digital literary and linguistic analysis; and a framework for disseminating research findings. The project is focusing on the needs of scholars in the digital humanities, a quickly evolving field that uses computational methods to ask new kinds of questions in humanistic research. HLRS will lead the implementation and deployment of the archive's hardware and software infrastructure and will provide expertise to support the adaptation of digital text processing methods to HPC resources. The project is one of four new Science Data Centers sponsored by the Baden-Württemberg Ministry of Science, Research, and Art.

SUPERCOMPUTING-AKADEMIE LAUNCHES NEW COURSE IN SIMULATION

On March 11, the Supercomputing-Akademie launched a new training module focusing on numeric simulation on HPC systems. Designed to address the needs of HPC users and IT professionals in industry, the 120-hour continuing education course is structured using a blended learning format. Participants meet at HLRS at the beginning and end of the course, but have the flexibility to complete the majority of their training online at their convenience. The Supercomputing-Akademie opened in 2018 with its first course on parallel programming and is in the process of developing a comprehensive training program aimed at increasing HPC expertise among industry professionals.

LRZ, LENOVO, AND INTEL EARN HPCWIRE EDITOR'S CHOICE AWARD FOR ENERGY EFFICIENT HPC

During the SC18 conference in Dallas, Texas, trade magazine HPCwire announced the winners of its annual Readers' and Editor's Choice Awards, acknowledging outstanding technical and scientific achievements in HPC. LRZ's new supercomputer, SuperMUC-NG, was awarded the Editor's Choice award for energy efficient HPC. The award acknowledged LRZ's long-running commitment to energy efficiency, pointing out that by using free warm water cooling for the system and reusing the system waste heat to heat the building as well as to produce cold water with adsorption chillers, SuperMUC-NG's installation would make LRZ the most effective HPC centre in the world for energy reuse.
From left to right: Dieter Kranzlmüller, Head of LRZ, Erika Krimmer, LRZ, Tom Tabor, CEO of Tabor Communications Inc. at the booth "Bits, Bytes, Brezel & Bier – Supercomputing in Bavaria" at SC18 in Dallas.

GCS ANNOUNCES NEW CHAIRMAN OF THE BOARD

In May, GCS announced Prof. Dr. Dieter Kranzlmüller as the new Chairman of the Board. Prof. Kranzlmüller, who also serves as Director of the Leibniz Supercomputing Centre in Garching, is serving his first term as GCS Chair. He succeeds Prof. Dr. Michael Resch, Director of the High-Performance Computing Center Stuttgart, who will serve as vice chair alongside Prof. Dr. Dr. Thomas Lippert, Director of the Jülich Supercomputing Centre. The three GCS directors rotate the Chair position every two years.

"From the validation of black holes at the farthest reaches of the universe, to the interaction of a medicine's molecules with a cancer cell, to massive, real-time data processing of autonomous vehicles operating in a hive fashion on a highway, HPC has been and will be pivotal in shaping our scientific understanding, our industrial and economic competitiveness, and the quality and prosperity of our society," Kranzlmüller said.

"It is an immense pleasure to serve as GCS' Chair and to continue building our nation's leadership-class supercomputing assets with my fellow directors and our centres' staffs of talented, dedicated HPC experts, scaling and domain specialists, educators, and researchers," he added.

HLRS FEATURED IN PROMINENT ART INSTALLATION

This spring, supercomputing came to the Kunstmuseum Stuttgart, a cultural landmark in the heart of the city. In an installation called Image Capital, internationally exhibited photographer Armin Linke and historian of photography Estelle Blaschke used images and video created at HLRS to raise provocative questions about the aesthetics, economics, and politics of image making, ownership, and distribution.

The exhibit prominently features large-scale scientific visualizations and photos created at HLRS, as well as a video in which HLRS staff members Uwe Wössner and Thomas Bönisch describe the importance of supercomputing, data management, and 3D visualization to research and development today. By highlighting HLRS, Linke and Blaschke invite visitors to consider the historical context of the cutting edge technologies used at the center, particularly with respect to other technologies associated with the history of photography and photographic archives. The exhibit was part of the competition for the 2019 Kubus. Sparda Kunstpreis, an annual award presented to one outstanding artist who lives in or has a close connection to the state of Baden-Württemberg. In May, Linke was named winner of the competition, which was accompanied by a 20,000 euro prize.

STAFF SPOTLIGHT: CONNECTING THE DATA DOTS

Dr. Martin Schultz, Jülich Supercomputing Centre

Since joining the Jülich Supercomputing Centre (JSC) two years ago, Dr. Martin Schultz has aimed at connecting data and people. As group leader of JSC's Earth System Data Exploration (ESDE) group, he coordinates projects and initiatives aimed at integrating disparate datasets related to climate and weather processes, ultimately helping researchers improve insight into complex earth system processes by having more comprehensive data sets to work with. Based on his background in atmospheric science and long-running research interests in both natural science as well as computing, he is perfectly suited for the job.

From a young age, Schultz showed a strong interest in computers. During his university education in physics and subsequent graduate work in physical chemistry he focused on the natural sciences, but always used computers to control experiments and analyze data.

For years, he saw parallel computing as a means to help solve specific science problems, but that as computing technology advanced, he found himself drawn to working more directly with hardware and software solutions to manage increasingly large, complex datasets in order to enable scientific discovery on a range of topics related to earth science and beyond.

Each successive generation of supercomputers has enabled climatologists, earth scientists, and other researchers to make large gains in the level of detail and complexity they could incorporate into simulations, allowing for more accurate, but also more computationally demanding, methods for simulating problems. In the realm of climate simulations, that complexity comes from the manifold physical and biogeochemical interactions between air, seas, ice sheets, and ground, and from the enormous range of timescales involved. Increased resolution also exacerbates the volume of data that is needed to drive earth system models and that is generated from these models.

"For a long time, climate models were very efficient in making use of HPC resources," Schultz says. "However, the general programming paradigm is basically one large, very homogeneous type of code. That whole process is starting to change, because the next-generation supercomputers will inevitably consist of different subsystems with heterogeneous computing, memory, and storage architectures. Therefore, this kind of programming paradigm is breaking down, and we need new modelling concepts, which are more flexible and also more resilient to hardware failures."

To that end, Schultz and his collaborators in ESDE focus on helping the earth system science community address one of the primary concerns of the changing programming paradigm—how to efficiently manage and analyze large, complex datasets coming from multiple sources.

"There are estimates from HPC centres that the data users move winds up being about 10 times the amount that they actually generate. This is a big bottleneck. The typical model for data in HPC centres is focused on dedicated user communities who get certain rights to access certain datasets. In contrast, in the earth science community, there is a relatively open attitude for sharing data. Right now, there is a wide gap between these two worlds."

Schultz explained that to help bridge the gap, the community needs automated workflows on HPC systems, where the huge volumes of model data are immediately augmented with metadata while they are being produced so thatmetadata can be published without delay and without a laborious manual data publication process. Decoupling HPC technology from science will be an important step to increase productivity.

Ultimately, the HPC user, such as an earth system modeler, should not have to worry about understanding the fine points of storage architectures, but rather let the HPC system find the optimum solution for moving input, output, and validation data around.

Currently, Schultz and his group focus on establishing such workflows for machine learning tasks associated with air quality and rainfall predictions. These activities are funded through two different projects that emphasize the collaborative, multidisciplinary nature of the ESDE team and its emphasis on data integration.

The IntelliAQ project (http://intelliaq.eu) uses deep learning techniques to improve the analysis of air quality data with geographic and weather model data. For this, the project is able to draw from one of the world's largest collections of surface-based air quality measurements in the world, currently being housed at FZ Jülich.

The DeepRain project (http://deeprain-project.de) is focused on improving predictions for localized precipitation by using machine learning techniques on large-scale ensemble weather forecast simulations together with radar and topographical data. JSC partnered with the German Weather Service, Jacobs

University in Bremen, and the Universities of Bonn and Osnabrück on the project, and aims to ultimately enable 24-hour predictions of precipitation on a 1-kilometer scale.

For Schultz, JSC proved to be the perfect place to both grow and fuse together his various interests. In addition to the friendly and collegial atmosphere, Schultz values the multidisciplinary, collaborative environment at JSC.

“It is really easy to just walk across a corridor and ask other experts questions about many different topics from system architecture design to software concepts, programming paradigms, and useful software packages,” Schultz says. “I find the broad spectrum of expertise really interesting and appealing. In addition to HPC experts, we have the simulation labs that serve as direct connections to the communities around the various science domains. They are very well positioned to understand user requests arising from the need to answer scientific questions and work with JSC’s IT experts to provide solutions. You need a certain size and breadth to do this type of work, and JSC is well-suited to do it as a multidisciplinary centre.” *eg*

© FZ Jülich

TRAINING CALENDAR
HPC COURSES AND TUTORIALS

Course / Workshop Title	Location	Date
Concepts of GASPI and Interoperability with other communication APIs (PRACE course)	Stuttgart	Jul 1-2, 2019
Parallel Programming with HPX (PRACE course)	Stuttgart	Jul 4-5, 2019
Advanced C++, Focus on Software Engineering	Stuttgart	Jul 9-12, 2019
Deep Learning and GPU programming using OpenACC	Stuttgart	Jul 15-17, 2019
HPE Porting Workshop (tbc)	Stuttgart	Jul 22-26, 2019 (tbc)
Introduction to parallel programming with MPI and OpenMP	Jülich	Aug 12-16, 2019
Parallel Programming with MPI / OpenMP	Zürich	Aug 19-22, 2019
13th International Parallel Tools Workshop	Dresden	Sep 2-3, 2019
Intro. to ANSYS Fluid Dynamics (CFX, Fluent) on LRZ HPC Systems	Garching	Sep 2-6, 2019
CFD with OpenFOAM®	Siegen	Sep 2-6, 2019
Advanced Fortran Topics (PRACE course, tbc)	Garching	Sep 9-13, 2019
Introduction to Computational Fluid Dynamics	Stuttgart	Sep 9-13, 2019
Iterative Linear Solvers and Parallelization	Garching	Sep 16-20, 2019
High Performance Computing in Science and Engineering	Stuttgart	Oct 7-8, 2019
Porting code from Matlab to Python	Jülich	Oct 7-8, 2019
Introduction to Python	Jülich	Oct 7-9, 2019
Parallel Programming Workshop (MPI, OpenMP and advanced topics) (PRACE course, tbc)	Stuttgart	Oct 14-18, 2019
Scientific Visualization	Stuttgart	Oct 24-25, 2019
Introduction to GPU programming using OpenACC	Jülich	Oct 28-29, 2019
Workshop on Scaling and Node-Level Performance	Stuttgart	Nov 4-8, 2019 (tbc)
From zero to hero, Part II: Understanding and fixing intra-node performance bottlenecks	Jülich	Nov 5-6, 2019
Data Science with Python for Computational Biology	Jülich	Nov 6-8, 2019 (tba)
Software Development in Science	Jülich	Nov 19-20, 2019
Advanced C++ with Focus on Software Engineering	Garching	Nov 20-22, 2019
C++ Language for Beginners	Garching	Nov 25-29, 2019
Advanced C++ with Focus on Software Engineering	Stuttgart	Nov 26-29, 2019
Intro. to Progr. and Usage of Supercomputer Resources at Jülich	Jülich	Nov 28-29, 2019
Introduction to C	Jülich	Dec 2019 (tba)
Advanced Parallel Programming with MPI and OpenMP	Jülich	Dec 2-4, 2019
Fortran for Scientific Computing	Stuttgart	Dec 9-13, 2019
Introduction to hybrid programming in HPC	Stuttgart	Dec 2019/Jan 2020 (tbc)

VISIT INSIDE ONLINE FOR DETAILS

For a complete and updated list of all GCS courses, please visit:
http://www.gauss-centre.eu/training
or
http://www.gauss-centre.eu/gauss-centre/EN/Training/Training/_node.html

The German HPC calendar (organized by the Gauss Allianz in cooperation with all German HPC centres) provides an extensive list of training all taking place German HPC centres. More information can be found at:
http://hpc-calendar.gauss-allianz.de/

Further training courses and events can be found on GCS member sites:
http://www.hlrs.de/training/
http://www.lrz.de/services/compute/courses/
http://www.fz-juelich.de/ias/jsc/events

The Rühle Saal at HLRS in Stuttgart **© HLRS**

JÜLICH SUPERCOMPUTING CENTRE FORSCHUNGSZENTRUM JÜLICH

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

Provision of supercomputer resources: JSC provides access to supercomputing resources of the highest performance for research projects coming from academia, research organizations, and industry. Users gain access for projects across the science and engineering spectrum in the fields of modelling and computer science.

- Supercomputer-oriented research and development in selected fields of physics and other natural sciences by research groups and in technology, e.g. by doing co-design together with leading HPC companies.
- Higher education for master and doctoral students in close coperation with neighbouring universities.
- Implementation of strategic support infrastructures including community-oriented simulation laboratories and cross-sectional teams, e.g. on mathematical methods and algorithms and parallel performance tools, enabling the effective usage of the supercomputer resources.

The Cluster module of JSC's Modular Supercomputer "JUWELS". **© Forschungszentrum Jülich**

Contact

Jülich Supercomputing Centre (JSC)
Forschungszentrum Jülich
Prof. Dr. Dr. Thomas Lippert
Wilhelm-Johnen-Straße, 52425 Jülich, Germany
Phone +49 - 24 61 - 61 - 64 02
th.lippert@fz-juelich.de
www.fz-juelich.de/jsc

Compute servers currently operated by JSC

System	Size	Peak Performance (TFlop/s)	Purpose	User Community
Atos BullSequana X1000 Cluster "JUWELS"	10 cells, 2,559 nodes 122,448 cores Intel Skylake 192 graphics processors (NVIDIA V100) 273 Tbyte memory	12,001	Capability Computing	European (PRACE) and German Universities and Research Institutes
T-Platforms Cluster + Intel/Dell Booster "JURECA"	Cluster: 1,884 nodes 45,216 cores Intel Haswell 150 graphics processors (NVIDIA K80) 281 TByte memory	2,245	Capacity and Capability Computing	German Universities, Research Institutes and Industry
	Booster: 1,640 nodes 111,520 cores Intel Xeon Phi (KNL) 157 TByte memory	4,996		
Fujitsu Cluster "QPACE 3"	672 nodes, 43,008 cores Intel Xeon Phi (KNL) 48 TByte memory	1,789	Capability Computing	SFB TR55, Lattice QCD Applications

A detailed description can be found on JSC's web pages: www.fz-juelich.de/ias/jsc/systems

LEIBNIZ SUPERCOMPUTING CENTRE

Leibniz Supercomputing Centre
of the Bavarian Academy of Sciences and Humanities

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

- Giving general IT services to more than 100,000 university customers in Munich and for the Bavarian Academy of Sciences
- Running and managing the powerful communication infrastructure of the MunichScientific Network (MWN)
- Acting as a competence centre for data communication networks
- Being a centre for large-scale archiving and backup
- Providing high-performance computing resources, training and support on the local, regional, national and international level.

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

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

Contact

Leibniz Supercomputing Centre (LRZ)
Prof. Dr. Dieter Kranzlmüller
Boltzmannstraße 1, 85748 Garching near Munich
Germany
Phone +49 - 89 - 358 - 31 - 80 00
kranzlmueller@lrz.de
www.lrz.de

Compute servers currently operated by LRZ

System	Size	Peak Performance (TFlop/s)	Purpose	User Community
"SuperMUC-NG" Intel/Lenovo ThinkSystem	6,336 nodes, 304,128 cores, Skylake 608 TByte, Omni-Path 100G	26,300	Capability Computing	German universities and research institutes, PRACE (Tier-0 System)
	144 nodes, 8,192 cores Skylake 111 TByte, Omni-Path 100G	600	Capability Computing	
"SuperMUC Phase 2" Lenovo Nextscale	3,072 nodes, 86,016 cores, Haswell EP 197 TByte, FDR 14 IB	3,580	Capability computing	German universities and research institutes, PRACE (Tier-0 System)
"CooLMUC-2" Lenovo Nextscale	384 nodes, 10,752 cores Haswell EP 24.6 TByte, FDR 14 IB	447	Capability computing	Bavarian Universities (Tier-2)
"CoolMUC-3" Megware Slide SX	148 nodes, 9,472 cores, Knights Landing, 17.2 TByte, Omnipath	459	Capability Computing	Bavarian Universities (Tier-2)
IvyMUC	Intel Xeon E5-2650 ("Ivy Bridge")	13	Capability Computing	Bavarian Universities (Tier-2)
Teramem	1 node, 96 cores, Intel Xeon E7-8890 v4 ("Broadwell"), 6 TByte RAM	13	Big Data	Bavarian Universities (Tier-2)
DGX-1, DGX-1v Machine Lerning Systems	2 nodes, Nvidia Tesla, 8 x P100, 8 x V100	1,130 (Mixed Precision)	Machine Learning	Bavarian Universities (Tier-2)
Compute Cloud for SuperMUC-NG	64 nodes, 3,072 cores, Intel Xeon ("Skylake"), 64 Nvidia V100	128, 8,000 (Mixed Precision)	Cloud	German Universities and Research Institutes, PRACE

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

HIGH-PERFORMANCE COMPUTING CENTER STUTTGART

High Performance Computing Center | Stuttgart

Based on a long tradition in supercomputing at University of Stuttgart, HLRS (Höchstleistungsrechenzentrum Stuttgart) was founded in 1996 as the first German federal centre for high-performance computing. HLRS serves researchers at universities and research laboratories in Europe and Germany and their external and industrial partners with high-end computing power for engineering and scientific applications.

Service for industry

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

Bundling competencies

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

World class research

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

View of the HLRS Cray XC40 "Hazel Hen". **© HLRS**

Contact

Höchstleistungsrechenzentrum Stuttgart (HLRS), Universität Stuttgart
Prof. Dr.-Ing. Dr. h.c. Dr. h.c. Hon.-Prof. Michael M. Resch
Nobelstraße 19, 70569 Stuttgart, Germany
Phone +49 - 711 - 685 - 8 72 69
resch@hlrs.de
www.hlrs.de

Compute servers currently operated by HLRS

System	Size	Peak Performance (TFlop/s)	Purpose	User Community
Cray XC40 "Hazel Hen"	7,712 nodes 185,088 cores 1 PB memory	7,420	Capability Computing	European (PRACE) and German Research Organizations and Industry
NEC Cluster (Laki, Laki2) heterogenous computing platform of 2 independent clusters	826 nodes 17,420 cores 88 TB memory	726	Capacity Computing	German Universities, Research Institutes and Industry
NEC SX-ACE	64 nodes 256 cores 4 TB memory	16	Vector Computing	German Universities, Research Institutes and Industry

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

InSiDE magazine (German: Innovatives Supercomputing in Deutschland) is the biannual publication of the Gauss Centre for Supercomputing, showcasing recent highlights and scientific accomplishments from users at Germany's three national supercomputing centres. GCS was founded in 2007 as a partnership between the High-Performance Computing Center Stuttgart, Jülich Supercomputing Centre, and the Leibniz Supercomputing Centre. It is jointly funded by the German Ministry of Education and Science (Bundesministerium für Bildung und Forschung – BMBF) and the corresponding ministries of the three states of Baden-Württemberg, North Rhine-Westphalia, and Bavaria.

www.gauss-centre.eu

Cover image: This illustration shows a textured 3D Mesh of San Francisco. The data was provided courtesy of Geomni. © nFrames