INNOVATIVES SUPERCOMPUTING IN DEUTSCHLAND N° 16-2 · AUTUMN 2018

INSiDE

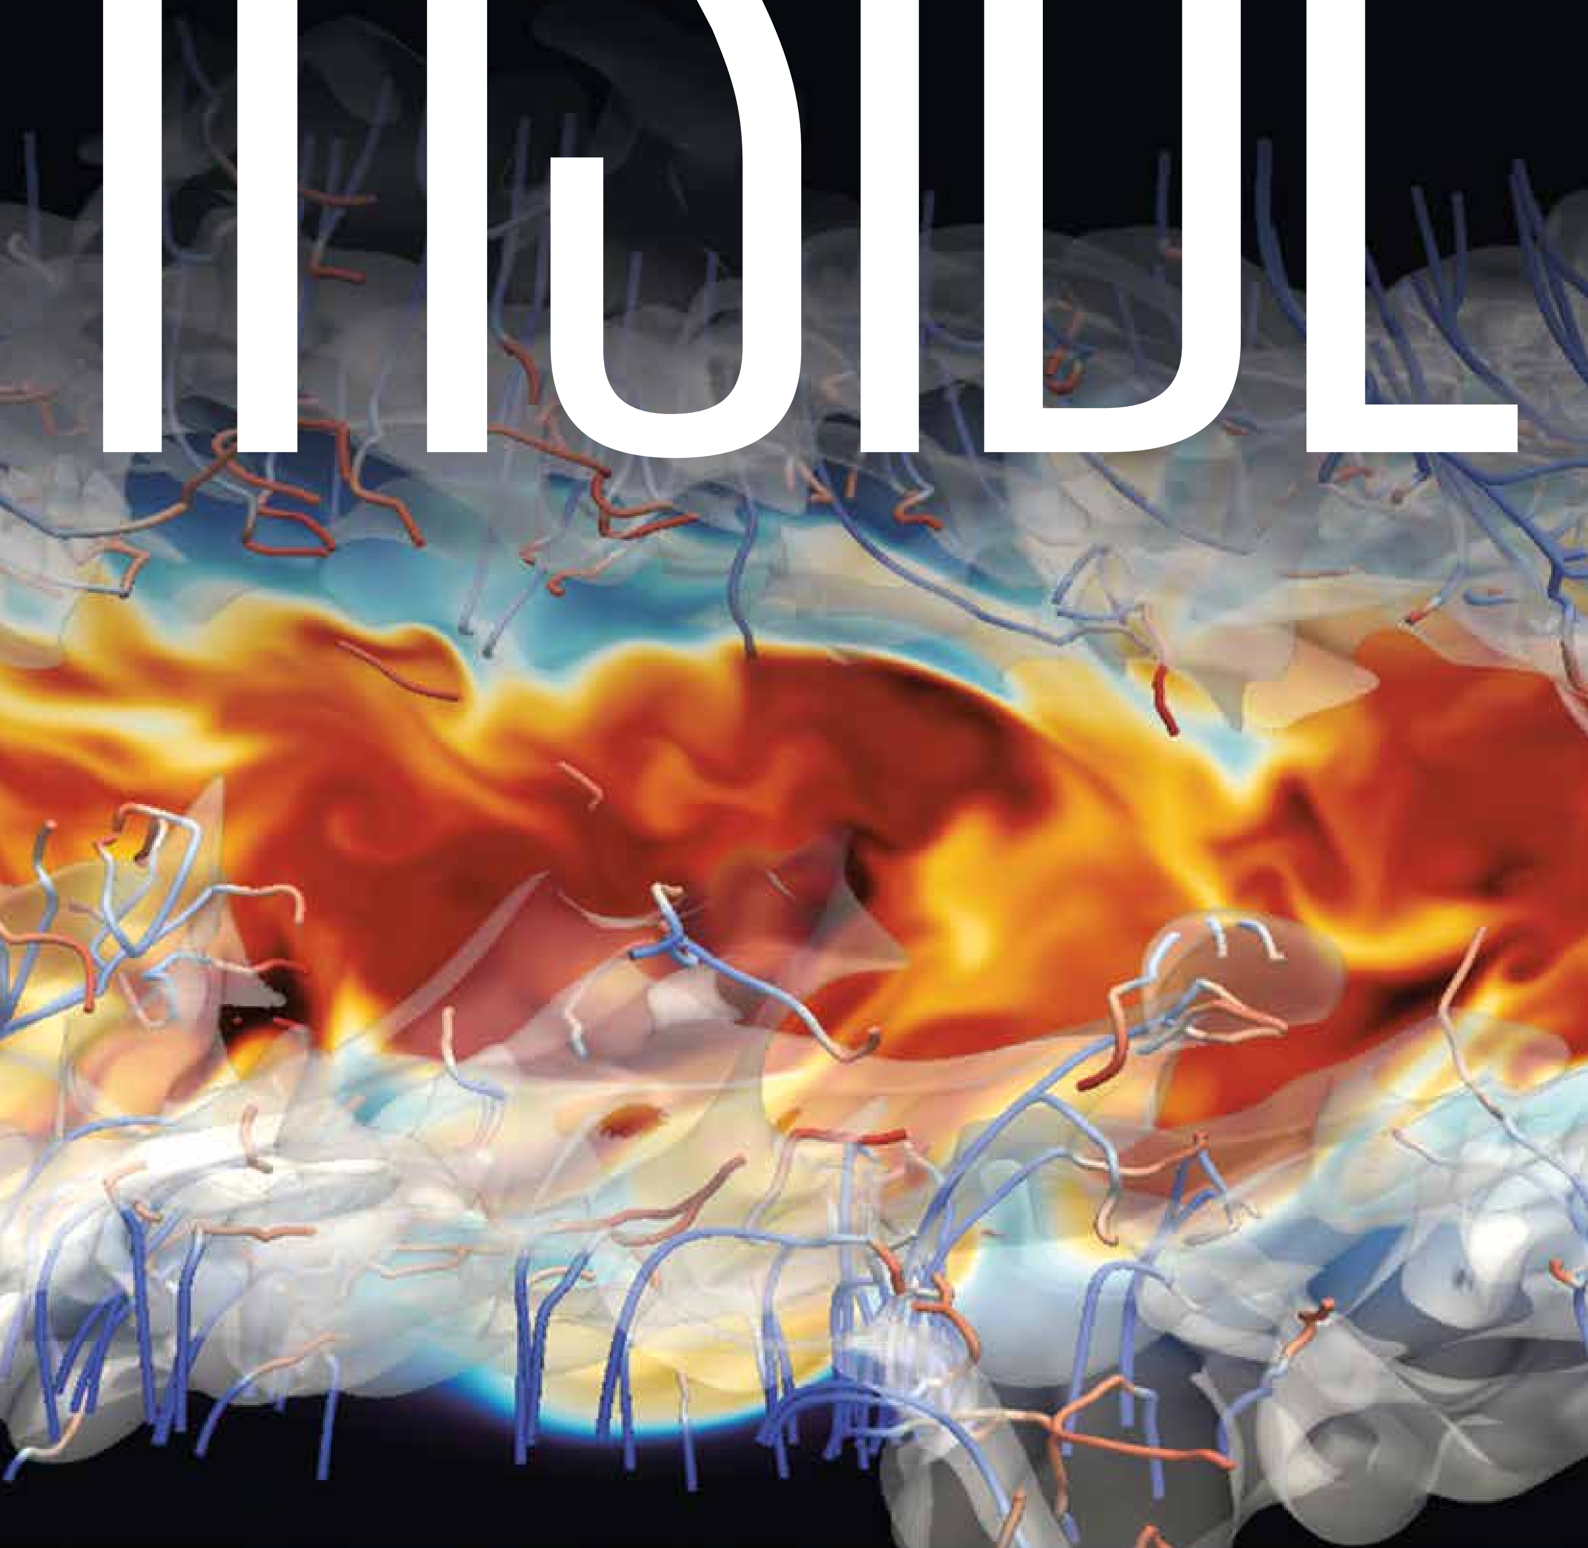

Publishers

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

Science Writers

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

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

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

Stefan Holl (*sh*), JSC
st.holl@fz-juelich.de

Lukas Arnold (*la*), JSC
l.arnold@fz-juelich.de

Sabrina Eisenreich (*se*), LRZ
sabrina.eisenreich@lrz.de

Visit us at:
www.gauss-centre.eu

GCS
Gauss Centre for Supercomputing

Welcome to the newest issues of InSiDE, the bi-annual Gauss Centre for Supercomputing (GCS) magazine highlighting innovative supercomputing in Germany. In 2018, GCS further advanced its computing capabilities, and deepened its commitment to helping academic and industrial users make full use of next-generation supercomputing architectures.

Recently, the Jülich Supercomputing Centre inaugurated the 12-petaflop Jülich Wizard for European Leadership Science (JUWELS), a novel modular supercomputer that went into operation in September (page 4). The Leibniz Supercomputing Centre is installing its next-generation SuperMUC machine, set to be operational in 2019. Users are looking forward to nearly 30 petaflops all based on a homogenous architecture (page 5). At the High-Performance Computing Center Stuttgart (HLRS), we are finishing up the procurement process for our next-generation system, and have been hard at work to ensure our users get up to speed as quickly as possible. To that end, we applied and received funding for participating in 4 European high-performance computing (HPC) Centres of Excellence (page 7), including leading the EXCELLERAT CoE, dedicated to facilitating the knowledge transfer between the HPC community and industrial users in engineering.

In this issue, we highlight how our users are already leveraging supercomputing technology to make advancements in studying embryotic development (page 11) in addition to studying how to make power plants safer, cleaner, and more efficient (pages 8 and 13).

As all three centres migrate their flagship supercomputers, GCS will continue to focus on providing our users with a variety of architectures designed to support the widest possible range of scientific applications. We will also maintain our commitment to training and advancing our user support structures so that our staff members are not only helping users efficiently leverage our machines, but also ensuring that they are spending most of their time focused on scientific challenges rather than clearing computational hurdles.

2018 marked the beginning of the second decade of GCS, and we look forward to continually building on the success of the last decade

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

CONTENTS

GERMANY RINGS IN ITS NEXT GENERATION OF SUPERCOMPUTERS AT JUWELS INAUGURATION

JSC finishes installation of modular HPC system, the first since GCS received its second decade of funding.

North Rhine-Westphalia Minister President Armin Laschet, Federal Minister for Education and Research Anja Karliczek, and FZ Jülich Chairman Wolfgang Marquardt press the lauch button for the new JUWELS supercomputer at JSC. © FZ Jülich

During a ceremony led by state and federal officials at the Jülich Supercomputing Centre (JSC) on September 18, 2018, the fastest supercomputer in Germany, the Jülich Wizard for European Leadership Science (JUWELS), joined the Gauss Centre for Supercomputing (GCS) family. The machine, capable of 12 petaflops, is the latest member of the high-performance computing (HPC) systems offered by GCS and is one of the first modular architectures in the world to go into operation.

JUWELS, built in partnership and co-design between French hardware provider Atos, German software company ParTec, Israel-based Mellanox, and JSC, is a milestone on the road to a new generation of ultra-flexible modular supercomputers, targeting a broader range of scientific applications. The first module is now in operation, and the additional modules will be subsequently added. Although the machine is the first double-digit petaflop system in Germany, JSC Director Thomas Lippert emphasized JUWELS' multiple modules are focused more on delivering a flexible portfolio of architectures to enable researchers across a wide swath of scientific disciplines to achieve new heights in their research.

"For us, modular supercomputing is the key to a forward looking, affordable, and energy efficient technology which will facilitate the realization of forthcoming exascale systems," he said. "As one of the largest German research centres, we are in a position to work together with our partners Atos and ParTec to develop the next generation of supercomputing ourselves."

After a greeting from Forschungszentrum Jülich Director Wolfgang Marquardt, Anja Karliczek, Germany's Minster for Education and Research, spoke about the importance of investing in high-performance computing (HPC) and called on researchers to continue to use these increasingly complex devices to benefit society at large. Afterwards, the Minister President for the State of North Rhine-Westphalia, Armin Laschet, spoke about how JSC was continuing its leadership in supercomputing both in terms of performance and in terms of environmental stewardship.

The system, which replaces JSC's IBM BlueGene/Q system JUQUEEN, consists of 2,511 compute nodes with two Intel Xeon 24 core Skylake CPUs each and 48 accelerated compute nodes with two Intel Xeon 20-core Skylake CPUs each equipped with 4 Nvidia V100 GPUs. A second booster module, optimized for massively parallel workloads, is currently scheduled for the beginning of 2020. The entire system will consist of multiple, architecturally diverse but fully integrated modules designed for specific simulations and data science tasks.

JUWELS embodies the German Smart Scale strategy by delivering a stable, diverse architecture for researchers to spend more time focused on science and less time focused on optimizing their codes.

The projects supported come from science fields as diverse as atomic and nuclear physics, condensed matter, elementary particle physics, and scientific engineering. "We are excited to further develop and implement the modular architecture concept with JUWELS. Our next challenge is to work with our users on their applications to use the new system in the most efficient way," says Professor Thomas Lippert, director of the Jülich Supercomputing Centre. *eg*

USERS LOOK TO ADVANCE SIMULATION SIZE AND SCOPE AS SUPERMUC-NG COMES ONLINE

The newest iteration of the SuperMUC high-performance computing system at LRZ will approach 30 petaflops peak performance.

On September 24, 2018, government officials from the state of Bavaria and the Bavarian Academy of Sciences and Humanities came to the Leibniz Supercomputing Centre (LRZ) to celebrate the start-up phase of LRZ's newest supercomputer, SuperMUC-NG. The combination of the last two generations of SuperMUC achieve a peak performance of 6.8 petaflops, or 6.8 quadrillion calculations per second. When the next-generation SuperMUC is fully online in 2019, the machine will be capable of 26.9 petaflops.

The machine, a joint development between Intel and Lenovo, also continues to build on the trend of energy efficiency in supercomputing—SuperMUC-NG will be cooled using warm water, and the system is designed to capture the machines waste heat and reuse it to heat the building and generate cool water to cool the storage and networking parts of the system.

In addition, SuperMUC-NG will not be using accelerators to get its performance boost. While there will be a "gateway" for pre- and post-processing that uses GPUs, the machine itself is purely CPU-based. Researchers whose simulations have difficulty being ported and scaled effectively to hybrid architectures are excited at the prospect of such a large system with a homogenous architecture. One long-time LRZ user, Ludwig-Maximilians-University Munich (LMU) researcher Dr. Klaus Dolag, has long used supercomputing to study galaxy formation and other large-scale cosmological phenomena. He sat down with the Gauss Centre for Supercomputing (GCS) to discuss his thoughts on the next-generation machine.

GCS: Dr. Dolag, how long have you been computing at LRZ, and can you explain a bit about your research?

KD: I have been involved in using high-performance computing to study how cosmological structures form for 20 years, and have been doing simulation at LRZ for the last 10. The last decade has been interesting if you are interested in large-scale cosmological structures and galaxy formation. We realized that supermassive black holes are really important in shaping galaxies or the central parts of galaxy clusters, and we've learned a lot about the numerical techniques used to study these objects.

GCS: What makes this type of cosmological modeling a supercomputing-caliber problem?

SuperMUC-NG is in the startup phase, and will be fully operational in 2019. © LRZ

KD: The systems that we're interested in are highly non-linear, which means we have to use numerical methods to model them in order to look at how they evolve and how many different processes interact. What makes it even more challenging is that if you want to get a clear sample of galaxies, you need simulate a large volume of the cosmos, but if you want to understand how they actually form and evolve, you need to be able to resolve the simulation in high detail. The processes governing galaxy formation are happening at the parsec (slightly more than 3 light years) or sub-parsec scale, but influence whole galaxies by driving gases out of galaxies on the hundreds of kiloparsec scale. In addition, individual galaxies are smaller parts of larger cosmological structures reaching hundreds of megaparsec scales (a million times farther than a single parsec). So as I said, these huge ranges need to be attacked using numerical methods and also require so much resolution that we need supercomputers in order to get useful models.

GCS: How has supercomputing impacted your work over the last decade?

KD: The interesting thing about galaxy formation is that many of the biggest breakthroughs have happened over the last 5 to 10 years. With the last generation of supercomputers, you had machinery that for the first time allowed you to do statistically useful simulations, meaning that we could finally get galaxy properties and large sample sizes in a reasonable amount of time. These types of simulations were only enabled by having access to machines like SuperMUC. These successes have really just opened the door for us. We have just now realized what we are capable of studying through modelling, and we have to work on improving our models. Now that we have made this technological leap forward, we can really ask interesting questions about how galaxies are born and evolve, the role of the surrounding environment and the role of additional physical processes related to plasma physics (such as determining the detailed properties of astrophysical fluids), radiation (such as the interaction of the astrophysical fluids with the radiation field created by various astrophysical objects) or even related to dust (such as the build-up of dust grains within astrophysical fluids) and things like that.

GCS: SuperMUC-NG will have roughly 4 times the performance of the last generation SuperMUC, but gets its speedup without using accelerators. What do you think about the architecture?

KD: The problem for us is, if we go to accelerators, we have to substantially change our programming. If you look at the programs used in our simulations, there are dozens of PhD theses that went into our models. It is not practical for us to just rewrite everything for new hardware. SuperMUC-NG's current configuration is perfect for us, as we can continue our development as-is. We will still have to make improvements and optimizations with (application programming interface) OpenMP, but this architecture allows us to continue on the path that we have already started with SuperMUC.

GALAXY CLUSTER: Researcher Klaus Dolag has been modelling cosmological structures, such as galaxy clusters, using LRZ resources for 10 years. With the next-generation SuperMUC set to come online in 2019, he looks forward to improving the size and resolution of his simulations.

GCS: What major computational hurdles have limited your simulations to-date? With access to the computing power of SuperMUC-NG, where do you see your research going?

KD: This is a more difficult question. There are two primary things we need in order to improve our simulations: size and resolution. So far, we have performed simulations at a reasonable resolution but where the volume isn't big enough to build up really representative volumes of the universe. That means we have to increase the size of the simulation, which will ultimately further increase the need for more resources. This problem is primarily memory bound, meaning the more memory we have, the better our simulation will run. Additionally, we would like to be able to numerically simulate more physical and chemical processes than we are currently able to include in the simulation, and that requires running a lot of additional or improved sub-scale models, which essentially account for the physical processes happening at scales too small for us to resolve in the main simulation. Ultimately we would like to even bridge these scales directly. By having access to greater CPU power, we would be able to keep our simulations the same size, but drastically increase resolution in order to more fully resolve some of these fine-scale phenomena. One way to explore our ultimate goal are applications where we improve the physics and resolution by taking smaller parts of the whole system and do much higher resolution simulations with this smaller section. With the node-size of SuperMUC-NG being so much larger (meaning a lot more memory-per-node, and therefore per core) we can substantially improve these simulations. *eg*

EUROPEAN CENTRES OF EXCELLENCE TO KICK OFF AT THE TURN OF THE YEAR

HLRS receives 7 million Euro in European grants to support technological progress in simulation.

In its most recent funding round for grants for "Centres of Excellence (CoE) on HPC" the European Union (EU) approved all proposals containing an HLRS contribution, supporting the future development of innovative, new high-performance computing (HPC) applications. With this achievement, HLRS is involved in four out of nine successfully funded EU CoEs in this call, making it one of the leading computing centres in application development in Europe. The Centres of Excellence will address topics related to HPC applications for industry and academia. With three-year funding for HLRS totaling 7 million Euro, the projects will start operating near the end of 2018.

Centre of Excellence for Engineering Applications flagship CoE for HLRS

HLRS will play the biggest part in EXCELLERAT (European Centre of Excellence for Engineering Applications), acting as project coordinator thanks to its long-term experience with industrial engineering partners. The project will support industrial HPC users by improving complex simulation technologies for engineering applications, making the development and testing of high-tech products in many engineering fields safer and more efficient.

The ultimate goal of EXCELLERAT is to enable industrial applications to scale on (pre-)exascale HPC systems. As HLRS director Michael Resch pointed out, simply giving industry access to supercomputers will not serve this purpose. "It is not enough to have a powerful machine and make it available," said Resch. "Users from industry need software and overall solutions that cover their individual needs by running efficiently on these complex computer architectures. To achieve this, we work closely with code developers, computer manufacturers, and users from the private sector and academia."

This proactive approach has met the approval of industry leaders in Stuttgart and beyond. Automotive and aeronautic engineering companies including Airbus, Rolls-Royce PLC, Daimler and Porsche, and independent software vendors such as the simulation software engineering company ANSYS welcomed the funding for EXCELLERAT in letters of support to the European Commission.

CoEs aim to support industry and society with advanced simulation technology

POP 2, which follows up on the first POP project, will support scientific users and software developers by helping them more efficiently parallelize and, in turn, improve the performance of their codes. Researchers who optimize their codes with HPC experts ultimately make better use of European supercomputers and improve their individual productivity.

Further, HLRS is participating in two projects addressing broader societal challenges. The CheeSE project will use HPC to predict natural disasters quicker and more reliably in order to reduce response time. HiDALGO will combine data analysis, machine learning, and HPC to address global challenges in finance, health care, migration, and energy. In this unique consortium, the computational and social sciences will join forces, with HLRS acting as an HPC expert and providing its infrastructure for high-performance data analysis.

In addition, HLRS is part of FocusCoE, a collaboration support action that will bring together all nine CoEs to facilitate knowledge exchange. HLRS will contribute mainly in training activities that will identify problems common to the CoEs and organize events to help solve them.

Centres of Excellence play important role in the European HPC ecosystem

Centres of Excellence play an important role in the European HPC ecosystem, providing essential applications for basic research, such as testing the functionality of codes and algorithms on complex computer architectures. They are also structured as part of the pan-European HPC infrastructure, ensuring that partners from academia and industry have access to the best European computing centres.

This ecosystem drives the development of new HPC technologies, making it possible to perform increasingly complex simulations and ultimately laying the foundations for scientific breakthroughs. *lb*

SUPERCOMPUTING SIMULATIONS AND MACHINE LEARNING HELP IMPROVE POWER PLANTS

High-performance computing resources and data-driven machine learning help University of Stuttgart researchers model how coal, nuclear, and geothermal power plants could be retrofitted for cleaner, safer, and more efficient and flexible operation.

In conventional steam power plants, residual water must be separated from power-generating steam. This process limits efficiency, and in early generation power plants, could be volatile, leading to explosions.

In the 1920s, Mark Benson realized that the risk could be reduced and power plants could be more efficient if water and steam could cohabitate. This cohabitation could be achieved by bringing water to a supercritical state, or when a fluid exists as both a liquid and gas at the same time.

While the costs associated with generating the temperature and pressure conditions necessary to achieve supercriticality prevented Benson's patented Benson Boiler from being widely adopted at power plants, his concepts offered the world its first glimpse at supercritical power generation.

Almost a century later, researchers at the University of Stuttgart's Institute of Nuclear Technology and Energy Systems (IKE) and Institute of Aerospace Thermodynamics (ITLR) are revisiting Benson's concepts to explore how it can improve safety and efficiency in modern power plants. Using high-performance computing (HPC), the researchers are developing tools that can make supercritical heat transfer more viable.

"Compared with subcritical power plants, supercritical power plants result in higher thermal efficiency, elimination of several types of equipment, such as any sort of steam dryer, and a more compact layout," said team member Sandeep Pandey, a PhD candidate at IKE.

Mr. Pandey and Dr. Xu Chu of ITLR are leading the computational aspects of this research, and in conjunction with computer science researchers at the Singapore Institute of Technology (SIT), are employing machine learning techniques informed by high-fidelity simulations on a supercomputer, while also developing a tool that can be easily employed using commercial computers.

In order to make an accurate tool to use commercially, the team needed to run computationally intensive direct numerical simulations (DNS), which is only possible using HPC resources. The High-Performance Computing Center Stuttgart's (HLRS's) Hazel Hen supercomputer enabled the high-resolution fluid dynamics simulations they required.

The heat of the moment

While power generation and other industrial procedures use a variety of materials to generate steam or transfer heat, using water is a tried and true method—water is easily accessible, well-understood on a chemical level, and predictable under a wide range of temperature and pressure conditions.

That said, water predictably enters its critical point at 374 degrees Celsius, making supercritical steam generation a sizzling process. Water also needs to be under high pressure—22.4 megapascals, or more than 200 times the pressure coming out of a kitchen sink, in fact. Further, when a material enters its critical state, it exhibits unique properties, and even slight changes to temperature or pressure can have a large impact. For instance, supercritical water does not transfer heat as efficiently as it does in a purely liquid state, and the extreme heat needed to reach supercritical levels can lead to degradation of piping, and, in turn, potentially catastrophic accidents.

Considering some of the difficulties of using water, Pandey and his colleagues are investigating using carbon dioxide (CO_2). The common molecule offers a number of advantages, chief among them being that it reaches supercriticality at just over 31 degrees Celsius, making it far more efficient than water. Using carbon dioxide to make power plants cleaner may sound like an oxymoron, but Pandey explained that supercritical CO_2 (sCO_2) is a far cleaner alternative.

"sCO_2 actually has zero ozone depletion potential, and little global warming potential or impact when compared to other common working fluids, such as chlorofluorocarbon-based refrigerants, ammonia, and others," Pandey said. In addition, sCO_2 needs far less space and can be compressed with far less effort than subcritical water. This, in turn, means that it requires a smaller power plant—an sCO_2 plant requires ten-fold less hardware for its power cycle than traditional subcritical power cycles.

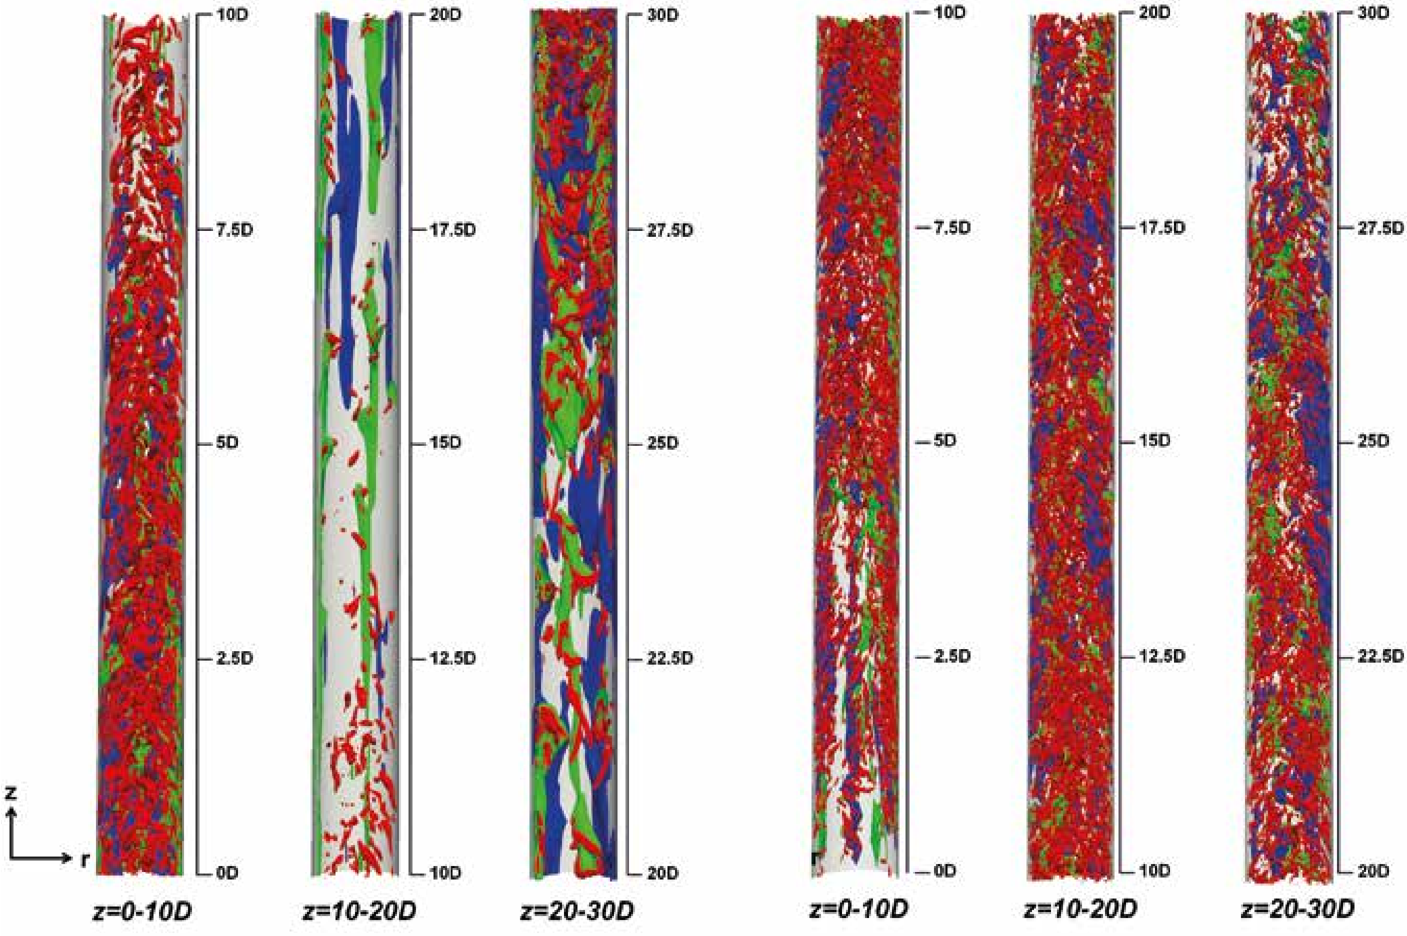

Researchers at IKE and ITLR at the University of Stuttgart are studying supercritical carbon dioxide that could replace supercritical water as a working fluid at power plants. This simulation shows the structure and the (red) high and (blue) low speed streaks of the fluid during a cooling process. The researchers observed a major difference in turbulence between downward flowing (left) and upward flowing (right) supercritical carbon dioxide. **© IKE, University of Stuttgart**

In order to replace water with carbon dioxide, though, engineers need to thoroughly understand its properties on a fundamental level, including how the fluid's turbulence—or uneven, unsteady flow—transfers heat, and in turn, interacts with machinery.

When doing computational fluid dynamics simulations related to turbulence, computational scientists largely rely on three methods: Reynolds-Averaged Navier-Stokes (RANS) simulations, large eddy simulations (LES), and direct numerical simulations (DNS). While RANS and LES methods both require researchers to include some assumptions using data coming from experiments or other simulations, DNS methods start with no preconceived notions or input data, allowing them to be far more accurate, but much more computationally expensive.

"LES and RANS models are usually used for simpler fluids," Pandey said. "We needed a high-fidelity approach for a complex fluid, so we decided to use DNS, hence our need for HPC resources."

Neural networks for commercial computers

Using the stress and heat transfer data coming from its high-fidelity DNS simulations, the team worked with SIT's Dr. Wanli Chang to train a deep neural network (DNN), a machine learning algorithm modeled roughly after the biological neural networks, or the network of neurons that recognize and respond to external stimuli.

Traditionally, researchers train machine learning algorithms using experimental data so they can predict heat transfer between fluid and pipe under a variety of conditions. When doing so, however, researchers must be careful not to "overfit" the model; that is, not make the algorithm so accurate with a specific dataset that it does not offer accurate results with other datasets.

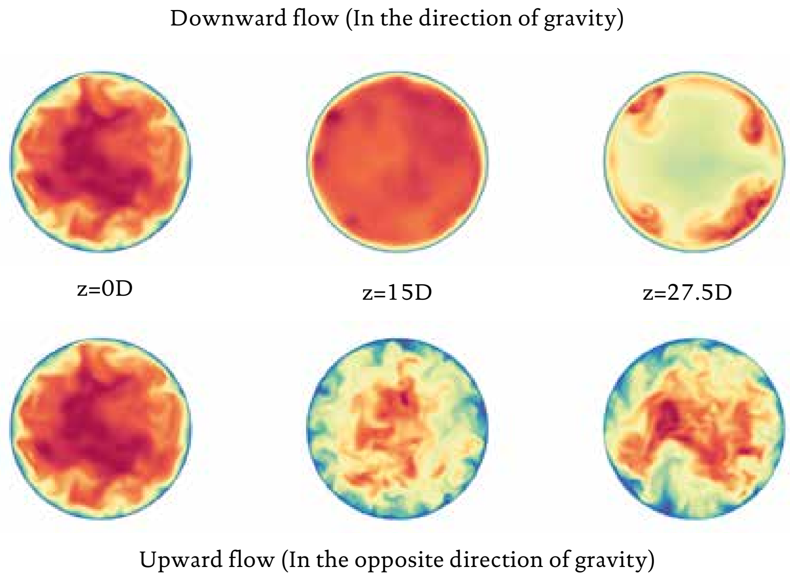

Figure showing the instantaneous velocity field in cooling process. A tremendous difference in the turbulence for upward and downward flow of sCO2 can be seen.
© IKE, University of Stuttgart

Using Hazel Hen, the team ran 35 DNS simulations, each focused on one specific operational condition, and then used the generated dataset to train the DNN. The team uses inlet temperature and pressure, heat flux, pipe diameter, and heat energy of the fluid as inputs, and generates the pipe's wall temperature and wall sheer stress as output. Eighty percent of the data generated in the DNS simulations is randomly selected to train the DNN, while researchers use the other 20 percent of data for simultaneous, but separate, validation.

This "in situ" validation work is important to avoid overfitting the algorithm, as it restarts the simulation if the algorithm begins showing a divergence between the training and datasets. "Our blind test results show that our DNN is successful in counter-overfitting, and has achieved general acceptability under the operational conditions that we covered in the database," Pandey said.

After the team felt confident with the agreement, they used the data to start creating a tool for more commercial use. Using the outputs from the team's recent work as a guide, the team was able to use its DNN to simulate the operational condition's heat energy with new data in 5.4 milliseconds on a standard laptop computer.

Critical next steps

To date, the team has been using OpenFOAM, a community code, for its DNS simulations. While OpenFOAM is a well-established code for a variety of fluid dynamics simulations, Pandey indicated that the team wanted to use a higher-fidelity code for its simulations. The researchers are working with a team from University of Stuttgart's Institute of Aerodynamics and Gas Dynamics (IAG) to use its FLEXI code, which offers higher accuracy and can accommodate a wider range of conditions.

Pandey also mentioned he is using a method called implicit LES in addition to the DNS simulations. While implicit LES simulations do not have quite the same high resolution present in the team's DNS simulations, it does allow the team to run simulations with higher Reynold's numbers, meaning it can account for a wider range of turbulence conditions.

The team wants to continue to enhance its database in order to further improve its DNN tool. Further, it is collaborating with IKE experimentalists to conduct preliminary experiments and to build a model supercritical power plant in order to test the agreement between experiment and theory. The ultimate prize will be if the team is able to provide an accurate, easy-to-use, and computationally efficient tool that helps engineers and power plant administrators generate power safer and more efficiently.

"Researchers at IKE are working with both experiments and numerical simulations," he said. "As part of the numerical team, we are seeking answers for poor heat transfer. We study the complex physics behind fluid flow and turbulence, but the end goal is to develop a simpler model. Conventional power plants help facilitate the use of renewable energy sources by offsetting their intermittent energy generation, but currently aren't designed to be as flexible as their renewable energy counterparts. If we can implement sCO2 based working fluids, we can improve their flexibility through more compact designs as well as faster start up and shut down times." *eg*

RESEARCHERS USE SUPERCOMPUTING TO DIVE DEEPER INTO STUDYING CELLULAR COMMUNICATION

Team combines simulation and experiment to study how proteins guide cells' development.

Long before a bird takes its first flight, a fish embarks on its life swimming through a lake or ocean, or a human takes its first steps, cells have to communicate to one another, ensuring that when an embryo reaches maturation, an organism has the traits and abilities needed to be successful.

Organs and their constituent parts, tissue and cells, govern a living being's ability to move, react to external stimuli, or perform any other of a being's many voluntary and involuntary motor processes. But cells have to talk to one another in order to divide up the labor necessary to grow into tissues and ultimately govern more complex processes. While biologists have long understood the basic principles guiding this process, understanding specific signaling processes is an important ongoing effort across many domains of biological research.

While experiments are still the primary driver for understanding biological processes at the cellular level, simulations have become increasingly used to complement experiments and accelerate research. To begin creating accurate representations of cellular processes in three dimensions, though, researchers have turned to supercomputers.

"Supercomputers are really pushing research on many levels, as they provide access to huge systems on long time scales across many disciplines," said Dr. Alexander Schug, head of the Computational Structural Biology group at the Jülich Supercomputing Centre (JSC). "Famous examples range from huge systems such as the universe or climate modeling to simulating tiny yet computationally complex biomolecular systems. In all these examples, supercomputers are indispensable."

Recently, Schug and a team of researchers have been using computing resources at JSC to model cellular signaling in the developing zebrafish brain.

Pass it on

Information passing through living tissues has been studied for a long time. Researchers know that certain cells release signaling molecules, and other cells' receptors receive these signals in order to guide an embryo's development. Theory long assumed that once these signaling molecules were released, they simply diffused across larger expanses of tissues, assigning specific cell fates or functions as they passed. Schug pointed out, though, that this alone cannot explain the great variability of signaling found in living organisms.

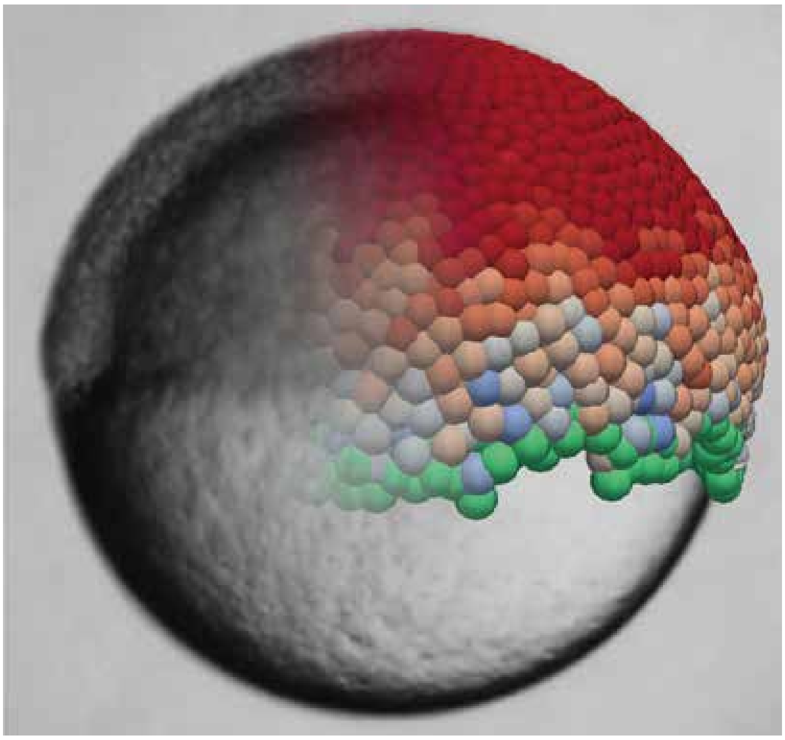

Overlay 3D simulation of early embryotic growth (coloured by density of signaling molecules) with experimental measurements (gray) © A. Schug, J. Rosenbauer.

"Look at people. Despite the differences in size, the proportions between head and feet are always about the same," he said. "This is tricky to achieve with simple diffusion. If molecules were to just simply diffuse through tissue, you get something called exponential depletion, meaning extremely quickly there is almost no signal the further away they get from their source."

The researchers modelled a fairly new concept in developmental biology research—direct transport. In essence, the cells that send out instructions develop tentacle-like protrusions called filopodia that pass information from the signaling cell to the receiving cell's receptors. While most prior similar simulations have primarily modelled cells in only one dimension—meaning that researchers cannot fully visualize many of the processes happening in tissue—Schug and his collaborators were able to develop a 2D model of this process occurring in a zebrafish embryo and are now working on a full 3D model. These simulations can closely mirror the measurements obtained from the team's experimental collaborators.

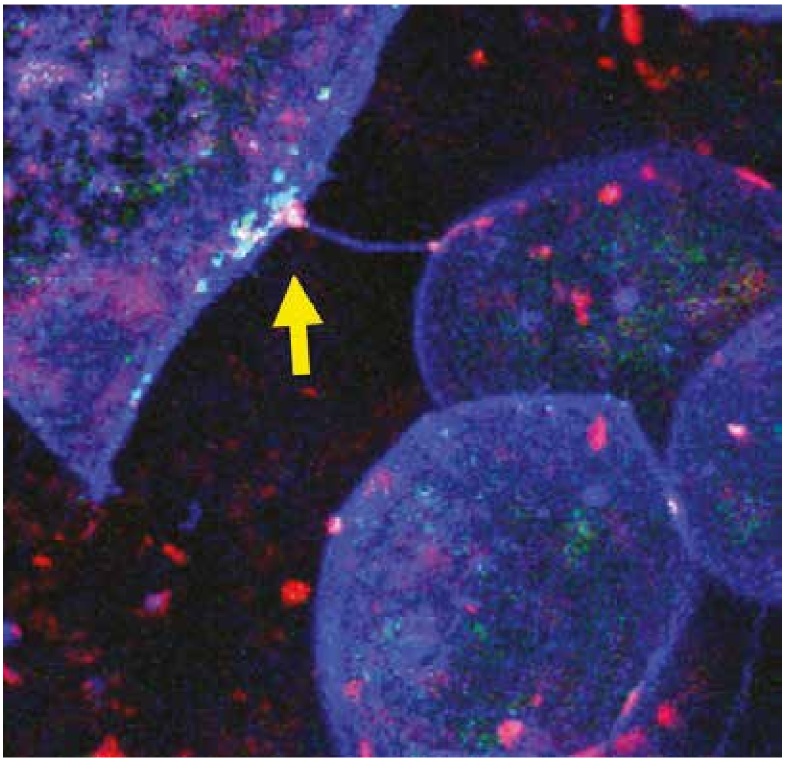

Cell extrusion signal (red) to a neighbouring cell, causing clustering of the receptor (green) at the membrane of the target cell. © JSC

Zebrafish are considered a model organism for scientific research due to their fully sequenced genome and transparent, relatively large embryo.

"Experimentalists can tag zebrafish cells with fluorescent molecules, and since the zebrafish embryo is transparent, you can nicely see where the cells are, which makes experiments easier," Schug said. "It is one of the standard models for this reason. It is a simple vertebrate that you can investigate. It is one of the easiest examples you can have, because you can just look at everything—literally."

Many rivers to cross

With access to supercomputing, the researchers were able to advance their simulation capabilities significantly by designing more complex models that accurately simulated the interactions of the filopodia passing along protein signals to receptors during a zebrafish embryo's brain development. Despite these advances, though, there is still a lot of room for growth in simulation size and detail and experiments are still an essential observational tool.

Schug indicated that the collaboration between experimentalists and computational scientists is key to advancing the field, and by having access to leading-edge supercomputing resources, simulation will become an increasingly important tool in developmental biology. "My original background is in molecular dynamics, where people have been developing computational models for simulating molecules for many years now," Schug said. "They have now become so realistic that they are driving experiments and make new hypotheses, ultimately complementing each other. Right now, in the field of tissue development we have fantastic data from experiments. I want to get the simulations to a comparable level of sophistication as well, which would enable us to test hypotheses to see how the entire biochemistry works during early stage development."

Understanding organisms' early development opens doors for researchers, improving understanding of the processes that govern healthy organism development and, in turn, what environmental factors can increase risk of genetic mutations or stunted development. Research in developmental biology is also important for studying human health, as knowledge of how cells communicate is pertinent for drug design and cancer research.

Considering the researchers have just begun optimizing their code for supercomputing architectures, it will take some time to extrapolate this approach to simulating larger biological systems, but the early returns for the team's 3D modelling approach are promising. As next-generation supercomputers, such as JUWELS at JSC—inaugurated this fall—continue to come online, research groups such that of as Schug and his collaborators will have even more capabilities to push their simulations forward in terms of size and detail.

"To grow larger in system size, we need bigger computers," Schug said. "But for the first time, we have taken our mathematical models into 3D space. We now have to continue to improve our models, so they don't become unrealistic, as they've never been used for a system of this size." *eg*

RESEARCHERS USE SUPERCOMPUTING FOR SUSTAINABLE COMBUSTION

New model simultaneously considers both chemistry and turbulence to test alternative fuels for power generation.

In the last decade, scientists have improved the energy efficiency and storage technology for green energy sources such as wind and solar power. Nevertheless, the world still generates roughly 80 percent of its power through the tried and true method of combustion.

In combustion, a fuel and an oxidizer (most often oxygen) meet under the right conditions to start the energy conversion process of the fuel, converting chemical energy into heat and some unwanted byproducts. For power generation, society has long relied on coal as the fuel for steam-based power generation, though it has the unfortunate side effect of generating large amounts of pollutants and carbon dioxide that contribute to air pollution, acid rain, and global warming.

For these reasons, researchers are interested in making combustion a cleaner, more efficient process; one approach involves gaining a more precise understanding of how combustion works, including during combustion of alternative fuels.

Prof. Dr. Christian Hasse, professor at TU Darmstadt, is principal investigator on a computational project aimed at this goal. Through a German Research Foundation (DFG) grant, he and his collaborators have been using the SuperMUC supercomputer at the Leibniz Supercomputing Centre (LRZ) in Garching near Munich to develop models that accurately capture the complex interaction between chemical and turbulent processes that happen during combustion.

Laying the foundation

Combustion is a multiscale process, meaning that when researchers simulate chemical reactions in a computer, they must simultaneously model how a system is behaving at the molecular level while also accounting for larger-scale interactions that can influence the reaction. Studying either one of these phenomena through computer modeling is a challenge that requires access to supercomputing resources.

Hasse and his collaborators are investigating the interaction of turbulence and chemical reactions. Understanding turbulence at the smallest, most fundamental scale is one of the great, unsolved scientific challenges that pushes even the most advanced supercomputers to their limits. Studying turbulent flows in the context of combustion is even more complicated.

"All the demands you have for a turbulence simulation are amplified when it comes to combustion," Hasse said. "The interactions related to combustion chemistry and turbulent motion are all happening on the finest scales."

Gaining a realistic understanding of how such a system behaves means that researchers must model a sufficiently large system. At the same time, they must study both chemical mixing and turbulent motions at the smallest scales.

In computational research, scientists sometimes use experimental data to inform computer simulations, and at other times run simulations to verify experimental results. For Hasse's group, simulation is leading the way in scientific discovery—the researchers' simulations start with no specific modelling assumptions for the turbulence-chemistry interactions, and calculate the fundamental physics underpinning these interactions. "Experiments in our case are impossible," Hasse said. "You need a 3D resolution as well as time. We wouldn't be able to experimentally generate a 4D dataset with the level of detail we can in simulation."

Most researchers doing HPC-caliber modelling of turbulent fluids use one of two approaches. The first, large-eddy simulations (LES), use input data from experiment or other simulations to make assumptions about the fine-scale interactions; its goal it to accurately resolve the chaotic, rotational motions of a turbulent fluid—the large eddies.

A more computationally intensive approach, direct numerical simulation (DNS), starts with no modelling assumptions, and solves equations for all of the most fundamental interactions in a turbulent system. The added computational demands result in a more accurate simulation.

Using SuperMUC, Hasse and his collaborators developed a general DNS-based model capable of simulating the chemistry and fluid dynamics of a variety of gaseous fluids used in power generation, allowing the team and other researchers to study this complicated relationship in a variety of scenarios.

Even the world's most powerful supercomputers are currently incapable of modelling full engine-sized reactions at DNS-level resolution, so the team is primarily focused on understanding the physics on a fundamental level. They are using that knowledge to inform and improve LES simulations with the aim of simulating this process more accurately on larger scales that can approach real conditions in power plants or engines.

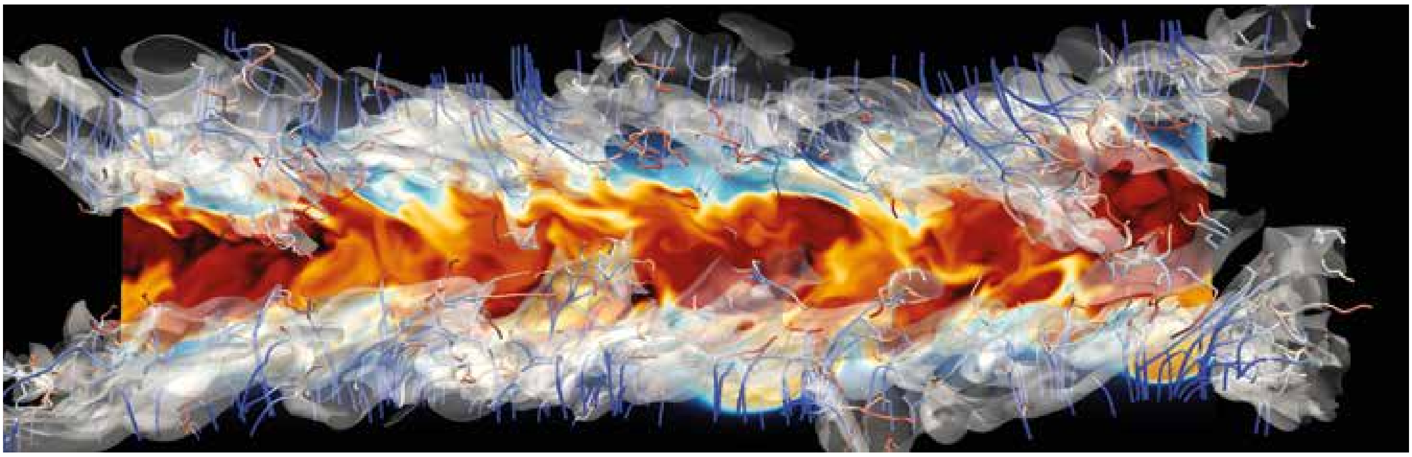

Illustration of the fine-scale mixing process (the lines) passing through the flame (gray surface). © Felix Dietzsch

"Our ultimate goal is to have physically consistent LES models that accurately account for the fine-scale features of turbulent chemistry interactions and that can scale to size we need," Hasse said.

High hopes for hydrogen

In addition to building its DNS model, the team uses the data collected from simulations to compare to combustion flames that can be studied in laboratory conditions, allowing it to validate its simulation data and give it more confidence in moving to larger, more realistic scenarios.

Hasse pointed to the collaboration with GCS centre staff as key to their productivity. "One key aspect to the team's success was efficient integration of our DNS code," he said. "We had valuable support from HPC specialists, in this case from Jülich Supercomputing Centre (JSC), and it really helped us to use the available hardware at LRZ as efficiently as possible."

The team is also using the model for more than traditional fossil fuel power generation. Germany has invested heavily in expanding its wind and solar power production. While technology has improved the efficiency with which wind turbines and solar panels can produce electricity, effectively storing it for situations in which sunlight or wind are not sufficient to meet energy needs is still a challenge.

Recently, researchers have explored a novel concept for storing power that offers opportunities not available using today's battery technologies. "One intelligent way for us to use this power is to store it chemically, this process is referred to as power to gas (P2G)," Hasse said. "We can use the electricity generated from renewable sources to produce hydrogen through electrolysis (a process that breaks water into hydrogen and oxygen). We could then use the resulting hydrogen directly for combustion or further process it for natural gas or other synthesis products. Germany has an especially wide distribution network for gas, though wind and solar are not practical everywhere, so producing hydrogen is a very interesting and promising way to store the energy generated from these other sources—not as electricity, but as a chemical."

The long-term health of the planet demands that humanity power daily life through other means than burning fossil fuels. A wide variety of renewable energy sources are making gains, but for the foreseeable future, combustion will still do the heavy lifting when it comes to power generation. From Hasse's perspective, that means researchers have an obligation to make combustion a clean and sustainable process for the here and now.

"For the foreseeable future, humanity will rely on a broad range of energy conversion technologies," Hasse said. "Combustion will be a part of this, so we have to make it as clean and sustainable as possible."

"Maybe in 30 or 40 years, we won't need combustion anymore," Hasse said. "But right now there is so much to be done so that we don't ruin the planet in that time. By making combustion more sustainable and less hazardous, that will help us quite a bit."

eg

EXA-DUNE – FLEXIBLE PDE SOLVERS, NUMERICAL METHODS, AND APPLICATIONS

Researchers across a broad range of computational science and engineering disciplines use simulations that solve partial differential equations (PDEs). The equations' ever-growing complexity and hardware requirements pose implementation challenges for domain scientists. Software frameworks offer a solution, as they provide a broad range of readily available numerical methods and state-of-the-art techniques, such as adaptive mesh refinement, higher-order discretisation schemes, or efficient parallel iterative solvers beyond "Matlab." They are designed with flexibility in mind so that new models and new numerical methods can be more easily implemented while also alleviating the burden of long-term code maintenance. One such software framework for numerical simulations is the Distributed and Unified Numerics Environment (DUNE) framework.

While software frameworks offer flexibility for the user, researchers also raise the concern that these frameworks fail to offer decent performance over a range of platforms, meaning that for broad use, developers have to focus more on dedicated hardware and software co-design. The EXA-DUNE project is a joint effort of seven workgroups, comprising scientific and high-performance computing groups that aims at extending DUNE by hardware-oriented numerical methods and hardware-aware implementation techniques developed in the FEAT3 project to provide an exascale-ready software framework solving a large variety of PDE systems.

In recent years, hardware architectures have undergone a dramatic change. As the clock rates grow much more slowly, increased computation power is mainly achieved by increasing parallelism. As the performance of most classical numerical methods for PDE simulations is memory bandwidth bound, these methods will hardly benefit from modern hardware and will not be able to sufficiently utilize upcoming exascale HPC systems, which represent a thousand-fold increase in current petascale performance. The paradigm of hardware-oriented numerics states that fundamental numerical and algorithmic research must accompany (long-term) technology trends--prospective, disruptive hardware trends enforce research into novel numerical techniques that are in turn better suited for the hardware. Hardware-oriented numerics thus involves much more than just hardware-efficient implementations--the explicit goal is to simultaneously maximise numerical efficiency (accuracy, asymptotically optimal convergence, flexibility), parallel efficiency (weak and strong scalability, independence of size and number of patches), and hardware efficiency (flop rates, saturation of memory bandwidth).

In the EXA-DUNE project, we have identified important challenges and possible solutions to improve the overall performance and reliability of the DUNE framework specifically and numerical codes in general, without sacrificing the current flexibility and maintainability.

Accelerators such as GPUs and modern CPUs feature wider and wider single-instruction-multiple-data (SIMD) units, meaning a single operation is performed on multiple input data. Current Intel Xeon processors, for example, feature 512 bit units, where either 8 double values or 16 float values can be computed in parallel, and NVIDIA GPUs even execute on warps of 32 data items simultaneously. To reach exascale, numerical algortihms must take full advantage of these SIMD units. A common misconception is that this is only possible by writing low-level machine code, explicitly using so-called intrinsic instructions. The EXA-DUNE project develops an high-level interface to these low-level instructions, such the flexibility of the DUNE framework is not compromised.

Funding:
2.3 Million Euro

Run Time:
2013 – 2019

Funding Source:
SPPEXA, the German Science Foundation Priority Programme 1648 – Software for Exascale Computing

Collaborators:
WWU Münster, University of Stuttgart, TU Dortmund, TU Clausthal, Heidelberg University, TU Kaiserslautern

All development of the EXA-DUNE project is available as open-source software. Several improvements, are already merged back to the official DUNE repositories (www.gitlab.dune-project.org/core/).

For more information, please visit: www.dune-project.org

Authors: Peter Bastian, Dominik Göddeke, Christian Engwer, Oleg Iliev, Olaf Ippisch, Mario Ohlberger, Stefan Turek

EXCELLERAT TO BRING HPC APPLICATIONS TO ENGINEERING INDUSTRY

A new European Union Centre of Excellence coordinated by HLRS aims to support industrial HPC users by improving simulation technologies for engineering applications.

High-performance computing (HPC) specialists are looking forward to the technological improvements that should arrive in the coming years as supercomputers approach the exascale. New approaches in hardware design (including new processors and high-bandwidth memory) and in application development (for example, code parallelization and data processing) will expand the power of supercomputing and therefore make it possible to solve new kinds of complex problems.

Industrial research and development in engineering is one application area of high-performance computing that is likely to benefit from these advances. Located in a region famed for its strong automotive sector, the High-Performance Computing Center Stuttgart (HLRS) collaborates closely with industrial engineering companies, making its research staff keenly aware of both the potential and the challenges that face the practical implementation of computational improvements in this arena.

Now, in an effort to expand its transfer of HPC expertise to the engineering industry, HLRS has launched a new European Union Centre of Excellence (CoE) for engineering applications called EXCELLERAT. With its roughly 8 million Euro in EU funding, the centre will accelerate technology transfer of leading-edge HPC developments to the engineering sector.

EXCELLERAT's goal is to facilitate the development of important codes for high-tech engineering, including maximizing their scalability to ever-larger computing architectures and supporting the technology transfer that will enable their uptake within the industrial environment. These activities will support engineers through the entire HPC engineering application lifecycle, including data pre-processing, code optimization, application execution, and post-processing. In addition, EXCELLERAT will provide training that prepares engineers in industry to take advantage of the opportunities that the latest HPC technologies offer.

Keys to the future success of EXCELLERAT are six codes – namely Nek5000, Alya, AVBP, Fluidity, FEniCS, and FLUCS – that are being brought into the project by consortium partners. The codes have been developed for academic applications in engineering-related fields such as aerospace, automotive, combustion, and fluid dynamics. To support their integration into real-life industrial applications, all consortium partners will work closely with end users outside the consortium. This will ultimately lead to fast feedback-cycles in all areas of the HPC engineering application lifecycle, from consultation on methodology and code implementation to data transfer and code optimization. End users will benefit from their commitment by gaining first-hand access to the project results. This concept of a one-stop-shop for all services is unique in the area of industrial HPC.

Past collaborations between HLRS and several consortium partners on large-scale engineering simulations ultimately led to the realization that EXCELLERAT is necessary for the future success of the European high-tech engineering sector. Project partner RWTH Aachen, for example, computed airflow simulation in an aircraft turbine on HLRS's last-generation supercomputer Hornet, using 95,000 cores to run simulations aimed at reducing noise. This simulation is not even close to reaching the limits of its scaling potential, though improved HPC technologies could make this possible.

The consortium of 13 partners in 7 European countries cover the following areas of expertise:

Project coordinator **HLRS** will provide consulting services for industrial HPC users on how to ensure their supercomputers are used most effectively. To support post-processing, HLRS will offer its extensive expertise in the implementation and application of parallel visualization tools for analyzing the large datasets produced by simulations.

Teratec and **SICOS-BW** are HPC advisors for industry and active at the intersection between computer science and the private sector. They will identify potential industrial users of EXCELLERAT services in their networks.

For machine learning-applications, the **Fraunhofer Institute for Algorithms and Scientific Computing SCAI** will provide expertise on computational algorithms and data analytics.

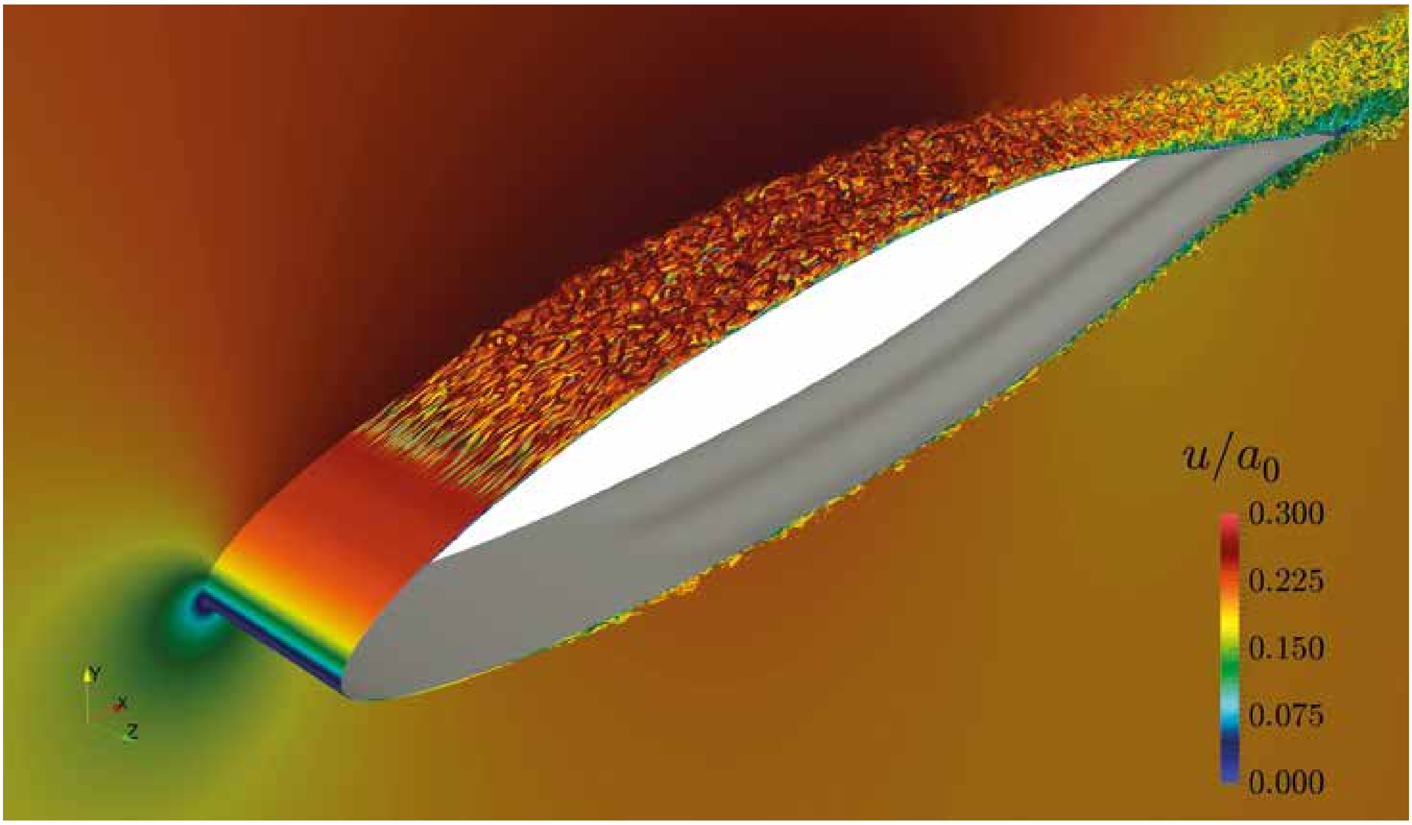

Codes being optimised through the EXCELLERAT CoE will help researchers to, among other things, model more effective wing designs. **© Institute of Aerodynamics, RWTH Aachen University**

The **Barcelona Supercomputing Center (BSC), Royal Institute of Technology (KTH), German Aerospace Center (DLR), RWTH Aachen, and French research centre CERFACS**, will supply the codes.

The participating HPC centres **HLRS, ARCTUR, CINECA, EPCC, BSC,** and **KTH** will provide access to their respective supercomputing systems in Germany, Italy, Slovenia, Scotland, Spain and Sweden, and will establish links to other systems, such as the forthcoming European pre-exascale machines.

IT service provider **SSC** will ensure secure data transfer using their SWAN platform for proprietary use cases. *lb*

Funding:
8 Million Euro

Run Time:
12/2018 – 11/2021

Funding Source:
European Comission, H2020 Program

Collaborators:
Arctur, BSC, CINECA, Cerfacs, DLR, EPCC, Fraunhofer SCAI, HLRS, KTH, RWTH Aachen, SICOS, SSC, Teratec

HOW TO OPTIMIZE TRAIN STATIONS FOR LARGE-SCALE EVACUATION

The project "KapaKrit - Optimization of the Traffic Capacity of Railway Stations in Crisis and Disaster Scenarios" started in August 2018, as an effort to leverage computational capabilities to help plan for mitigating potential disasters. The project, which runs for three years, is funded by the Federal Ministry of Education and Research (BMBF) in the program, "Civil Security – Transport Infrastructures." The project will be coordinated by the newly established institute 'Civil Safety Research' (IAS-7) at Forschungszentrum Jülich. Bochum University of Applied Sciences, and the Büro für Forschung, Entwicklung und Evaluation Wuppertal serve as partners in the project. Other institutions, including Deutsche Bahn AG and the non-state National Express Rail GmbH (NEX), are also involved as associated partners.

In crisis and disaster situations, such as chemical or nuclear accidents, railway stations become safety-critical traffic junctions for large-scale evacuation. However, the day-to-day capacity of the stations is not sufficient in the event of a large-scale evacuation. Railway preparedness includes both the maximization of transportation capacity and guaranteed, effective operations transfer to a crisis mode. Researchers involved in KapaKrit will investigate how people can be brought to safe areas via the railway network in such crisis scenarios. This modelling work covers the investigation of the macroscopic foot traffic flows toward stations as well as the microscopic dynamics inside the stations and on the platforms. All studies are accompanied by associated public safety authorities and organizations to ensure realistic scenarios and practical approaches. The project partners will evaluate the developed methods on the main train station in Dortmund as a case study, with respect to its central role in the Ruhr area. *sh, la*

Funding:
1.1 million Euro

Run time: 07/2018 – 06/2021

Funding source: BMBF

Collaborators:
Forschungszentrum Jülich, Büro für Forschung, Entwicklung und Evaluation Wuppertal, Hochschule Bochum, DB Station&Service AG, Dortmund, DB Sicherheit GmbH, Dortmund
Dortmunder Stadtwerke AG, Dortmund
Kompetenzcenter Sicherheit NRW, Gelsenkirchen, National Express Rail GmbH, Köln

Working with partners from academia and industry, JSC will be using the Dortmund main train station to better understand and plan for evacuations in the event of disasters. **© FZ Jülich**

NEXT-GEN HPC: THE PATH TO EXASCALE—ARTIFICIAL INTELLIGENCE AND PERSONALIZED MEDICINE

Experts from HPC and personalized medicine meet to talk about the future of the two fields.

Next-Gen HPC speakers (from left to right): Prof. Dr. Peter Coveney, Prof. Dr. Dieter Kranzlmüller, Dr. Frederick Streitz, Dr. Thomas Friese, Prof. Dr. Rick Stevens, Dr. Barbara Schormaier, Prof. Dr. Martin Schulz, Dr. Al Gara, Dr. Satoshi Matsuoka. © LRZ

Internationally renowned researchers from the HPC and personalized medicine communities converged at the Leibniz Supercomputing Centre (LRZ) in Garching, Germany on June 29, 2018 to offer their expertise and insight in a symposium titled "Next-Gen HPC: The Path to Exascale – Artificial Intelligence and Personalized Medicine."

Keynote speakers were Dr. Alan Gara, Chief Architect at Intel, with his presentation, "Technology Challenges and Trends—a Look through a 2030 Crystal Ball," and Dr. Fred Streitz, Director of the High-Performance Computing Innovation Center at the Lawrence Livermore National Lab speaking on "Machine Learning and Predictive Simulations–HPC and the US Cancer Moonshot."

Distinguished speakers included: Prof. Satoshi Matsuoka, Director of the Riken Center for Computational Science (Japan); Prof. Rick Stevens, Deputy Director of the Argonne National Laboratory and Professor at the University of Chicago; Prof. Peter Coveney, Director of the Centre for Computational Science at University College London (UK) and Dr. Barbara Schormair, Deputy Director of the Institute of Neurogenomics at the Helmholtz Zentrum München, German Research Centre for Environmental Health. The panel was moderated by Dr. Thomas Friese, Siemens Healthineers and Prof. Martin Schulz, Technical University of Munich.

The speakers provided detailed insights into trends and challenges, showing how closely the three topics of HPC, artificial intelligence and personalized medicine are interlinked. All speakers agreed that future supercomputers and their greatly increased computing capacity will enable new interactions of computer simulations, machine learning and the analysis of large amounts of data. However, as pointed out by Dr. Gara, the HPC community has to solve some challenges along the way, such as fundamental improvements in data transport and storage, changes in system architectures to scale memory, compute and communications or power consumption. Prof. Matsuoka added that the big data and machine learning (ML) communities must bring some of the HPC rigor in architectural, algorithmic, system software performance and modeling into their fields to make the convergence of HPC and ML fully realized.

In his keynote, Streitz, one of the lead researchers in the US Department of Energy's (DOE) Cancer Moonshot Project, assured that increasing the compute capacity of supercomputers from the present petaflop to the future exaflop horizon fundamentally changes the scientific questions that can be asked and, above all, answered. Pointing to his team's work on machine-learning-guided multiscale simulations for investigating Ras cancer proteins, such increased computing amplifies the abilities of machine learning in predictive simulations. Prof. Coveney (UCL) agreed, noting the fledging use of predictive simulation in the field of medical and pharmaceutical research, and highlighting the impact possible through presentation of his work using predictive simulations to determine which breast cancer treatment is best suited to treat patients based on their individual genetic profile.

Dr. Barbara Schormair, an expert in neurogenetics, asserted the necessity of an interdisciplinary approach bringing together technicians, lab scientists, engineers, computational scientists, biologists, neurologists and doctors to increase scientific knowledge at the crossroads of personalized medicine, HPC and machine learning. She noted the USA, Great Britain, and the Netherlands have already made considerable progress with major projects in the field of human genomics and presented concrete examples of how patients benefit directly from current research at the Helmholtz Zentrum München and the University Hospital rechts der Isar at the Technical University of Munich (TUM). For example, advances in genetics such as exom sequencing allow more precise diagnoses, such as in dystonia, and in certain cases, even the use of tailor-made therapies.

In the subsequent panel discussion, the interactions between the experts and the audience focused on questions around ethical handling of patient data and on the possibilities now available through the new technologies. The experts also discussed how early career scientists can become educated in cross disciplinary fields and how interdisciplinary cooperation between experts in supercomputing and machine learning, as well as physicians and geneticists, can be promoted and successfully shaped.

The Next-Gen HPC event attracted representatives from industry, high-ranking officials of the Bavarian Ministry of Science and the Arts, as well as professors and PhD students from Germany's leading universities of excellence, the Ludwig-Maximilians-Universität, Munich (LMU) and the Technical University of Munich (TUM).

"We are very pleased to bring internationally renowned speakers to LRZ for this discussion," said Prof. Dr. Dieter Kranzlmüller, Director of LRZ. "In the presentations and the panel, it became clear how well positioned Bavaria is for next-generation supercomputing and its application to artificial intelligence and personalized medicine. To continue pushing forward and to achieve the breakthroughs we know are possible, building and fostering rigorous interdisciplinary exchange and international cooperation is paramount to our common success and future computing." *se*

JSC DIRECTOR NAMED CHAIR OF PRACE COUNCIL

Prof. Dr. Dr. Thomas Lippert begins two-year term at the helm of trans-European supercomputing organization.

The Gauss Centre for Supercomputing (GCS) has been a hosting member in the Partnership for Advanced Computing in Europe (PRACE)—a trans-European organization dedicated to providing high-performance computing (HPC) resources across Europe—since its 2010 inauguration. As a key member of PRACE, GCS and the centre directors have long served as advocates and advisors for the organization, and one centre director was recently elected to help lead PRACE.

Prof. Dr. Ir. Anwar Osseyran of SURF/SURFsara in the Netherlands hands over the title of PRACE Chair to Prof. Dr. Dr. Thomas Lippert of JSC. © PRACE

Prof. Dr. Dr. Thomas Lippert, member of the GCS board and Director of the Jülich Supercomputing Centre (JSC), was elected as Chair of the PRACE Council during the group's 30th meeting, held in Amsterdam on June 18. Lippert follows Prof. Dr. Ir. Anwar Osseyran of the Netherland's SURF/SURFsara. Dr. Janne Ignatius, Programme Director at CSC—the IT Center for Science in Finland, will serve as vice-chair, succeeding Dr. Sergi Girona from the Barcelona Supercomputing Centre in Spain.

"PRACE is deeply grateful to Anwar and Sergi for their great commitment to shaping the second, very successful phase of the partnership," Lippert said. "The challenge ahead is to provide true exascale computing capability for science and industry in Europe. As PRACE Council Chair, and together with my colleagues and our scientific and industrial advisors, I am keen to bring the top-level HPC provision and support of PRACE together with the exascale systems of the coming EuroHPC joint undertaking."

Lippert will serve a two-year term in the post. The PRACE Council serves as the highest decision-making body in PRACE, overseeing the research infrastructure and the users' access to resources across the PRACE member organizations. He was selected as the German delegate in January, 2018, but has long been involved in the creation and guidance of PRACE since before it was officially founded. He served as an advisor to Forschungzentrum Jülich Director and GCS founding member Prof. Dr. Achim Bachem, one of the founding Council Chairs of PRACE. *eg*

HPC USER FORUM CONSIDERS STATE OF THE ART IN EUROPEAN SUPERCOMPUTING

The meeting brought together representatives of Germany's national supercomputing centres, HPC hardware manufacturers, and HPC users in industry to discuss current trends and challenges in the field.

Organized by HPC industry consultants Hyperion Research, the 2018 HPC User Forum took place at HLRS on October 1-2, 2018. It offered an intimate venue for promoting interactions among HPC users and technology suppliers, and provided insights into how supercomputers and their uses are evolving.

Following an HPC industry overview by Hyperion's Earl Joseph and Steve Conway, and an introduction to the Gauss Centre for Supercomputing's smart scaling initiative by GCS director Michael Resch, the three GCS centres described their latest improvements. JSC Deputy Head Norbert Attig introduced JUWELS, one of the world's first supercomputers based on a modular architecture. He described plans to boost its capabilities in the coming years, investments that GCS is making in facilitating better data transfer between supercomputing centres, and GCS initiatives to provide high-level user support. LRZ Director Dieter Kranzlmüller highlighted the Garching centre's new system, SuperMUC-NG. He focused on its unique energy efficiency features, LRZ's deployment of cloud infrastructure for data pre- and post-processing, and some exciting application areas in which the new supercomputer is being used. HLRS Managing Director Bastian Koller described HLRS's growth in personnel and funding, highlighting its recent successes in securing EU centre of excellence grants and expanding professional training offerings through the Supercomputing-Akademie.

Leonardo Flores of the European Commission provided an overview of the EuroHPC Joint Undertaking. © *cw*

Earl Joseph and Steve Conway of Hyperion Research opened the HPC User Forum. ©*cw*

Leonardo Flores of the European Commission provided an overview of the EuroHPC Joint Undertaking, an initiative to implement an EU strategy for exascale computing. He reported on the current roadmap, which intends to support development of both infrastructure and HPC expertise. The program plans to contribute to the creation of two to three more petascale computers and two larger, pre-exascale machines by 2020, two exascale machines by 2023, and the first post-exascale system by 2027. Flores reported that EuroHPC also intends to focus on the development of expertise in related areas such as cybersecurity and artificial intelligence.

Representatives of several leading HPC hardware manufacturers also gave overviews of their current product development strategies. Represented were Lenovo, Verne Global, Cray, IBM, and Inspur. Among the current challenges the manufacturers discussed were improving energy efficiency, managing increased power and cooling requirements, addressing exploding needs for data storage, and accelerating input and output among computing nodes in HPC systems.

Industry users of HPC were also represented, including Daimler, ETAS, and Bosch. Several talks focused on current methods for developing autonomously driving vehicles. In addition, presentations by representatives of Sicos-BW, the Fortissimo Project, and a newly funded project based at HLRS called EXCELLERAT looked at challenges facing HPC users in industry.

An additional highlight was a provocative presentation by GCS Director Michael Resch that urged the supercomputing community and its users to begin to rethink how HPC technologies are developed. Considering the signs that Moore's Law may have run its course, Resch suggested that the time is ripe to begin reconceiving HPC hardware and software so that they more effectively and efficiently manage today's most important simulation challenges. *cw*

ASC(S HOLDS 10-YEAR GALA

HLRS director Michael Resch joined the Automotive Simulation Center Stuttgart in celebrating its successes in networking car manufacturers, engineering service providers, and research institutes to support virtual automobile development. HLRS's resources have enabled the asc(s to maintain its focus on leveraging HPC tools.

EU COMMISSIONER GÜNTHER OETTINGER VISITS HLRS

The European Commissioner for Budget and Human Resources met with HLRS Director Michael Resch to discuss the importance of a European strategy for high-performance computing, including how European science and industry could be best supported through the EuroHPC Joint Undertaking.

STUDENT DIVES INTO DATA TO PREDICT TRAIN DELAYS

As a participant in the HLRS program Simulated Worlds, Niklas Knöll used machine learning to study this common source of complaints in the Stuttgart area. Simulated Worlds, an initiative funded by the Baden-Württemberg Ministry of Science, Research, and the Arts, aims to prepare high school students for careers in the digital realm.

HBP COLLOQUIUM AT FORSCHUNGSZENTRUM JÜLICH

The first Human Brain Project (HBP) Colloquium at Forschungszentrum Jülich took place on October 4. More than 180 participants exchanged views with the speakers on the successes and challenges of the Europe-based HBP. The keynote speech of Prof. Rainer Goebel from University of Maastricht, the panel discussion on Artificial Intelligence and Deep Learning as well as the initial talk of Prof. Andreas Pinkwart, North Rhine-Westphalia Minister of Economics were among the highlights.

EARTH EXHIBIT COMES TO NEW YORK CITY THANKS TO SUPERMUC SIMULATIONS

Results of computationally intensive simulations, aimed at studying processes in the Earth's mantle, can now be admired by the visitors of the American Museum of Natural History (AMNH) in New York City. The new permanent exhibit is based on the results of a research project led by Professor Hans-Peter Bunge of Ludwig-Maximilians-Universität in Munich (LMU), who leveraged the computing power of the at LRZ in order to study inner-Earth phenomena. He and his team have developed a computer model for simulating processes that take place in Earth's deep interior and illustrating their impact on the planet's surface. The new exhibit, which depicts convection in the Earth's mantle, is a 3-D printed physical model based on a snapshot from such a computer simulation.

SUPERCOMPUTING-AKADEMIE COMPLETES FIRST COURSE

In July, this newly launched continuing education program for HPC users in industry awarded certificates for its inaugural course in parallel programming. HLRS and partners continue developing additional modules that will be offered in a blended learning format, including a new course in simulation that is scheduled to launch early 2019.

EXTREME DATA WORKSHOP 2018 AT JSC

On September 18 and 19, 2018, 25 experts in weather, climate, environmental sciences, materials, and neurobiology gathered at the JSC for the first workshop on “Extreme Data: Demands, Technologies, and Services.” Exascale computing will require exascale data handling, meaning scientific progress and insights will hinge on the ability of the next-generation compute and data infrastructures to cope with these data streams. The workshop identified some common demands but also different workflows across various research fields and highlighted the need to strengthen IT training domain scientists’ educations so that extreme data workflows can be built and used by the community. A proceedings volume will be published.

PROF. ARNDT BODE RECEIVES HIGHEST AWARD FOR PUBLIC SERVICE IN GERMANY

Prof. Dr. Arndt Bode, former Chairman of the Board at LRZ and current President of the Bavarian Research Council, was awarded in August the Cross for Distinguished Service of the Federal Republic of Germany (first class) for his major contribution to the field of high-performance computing. The award is the highest national award for public service in Germany. The award was presented by Prof. Dr. med. Marion Kiechle, Bavarian State Minister for Science and Arts, who was acting on behalf of German President Frank Walter Steinmeier.

STUDENT RACING TEAM WINS INTERNATIONAL COMPETITION

For the third consecutive year, the Formula Student Racing Team of the University of Stuttgart won the 2018 Formula SAE Michigan, an international engineering design competition for students. HLRS has supported the team by sponsoring computing hours on its NEC cluster since 2016.

SUPERMUC STATUS AND RESULTS WORKSHOP 2018

In July, users of HPC system SuperMUC met at the bi-annual SuperMUC Status and Results Workshop at LRZ to share their experience, present the scientific results of their work and use the opportunity to discuss the usage and challenges of HPC systems. Almost 40 simulation projects from a broad range of scientific disciplines, such as astrophysics, elementary particle physics, computational and scientific engineering, life sciences, and geophysics were introduced by oral presentations and complemented by a research poster exhibition. During the user forum, the scientists discussed their current and future HPC needs with system administrators and application experts from LRZ.

TRAINING CALENDAR
HPC COURSES AND TUTORIALS

Course / Workshop Title	Location	Date
Cray XC40 Workshop on Scaling and Node-Level Performance	Stuttgart	Nov 5-9, 2018
C++ Language for Beginners	Garching	Nov 6-9, 2018
Software Development in Science	Jülich	Nov 19-20, 2018
Advanced C++ with Focus on Software Engineering	Stuttgart	Nov 19-22, 2018
Intro. to Progr. and Usage of Supercomputer Resources at Jülich	Jülich	Nov 22-23, 2018
Advanced Parallel Programming with MPI and OpenMP	Jülich	Nov 26-28, 2018
Fortran for Scientific Computing	Stuttgart	Dec 3-7, 2018
Introduction to C	Jülich	Dec 3-12, 2018
Deep Learning Workshops	Garching	2019 (tba)
Parallel and Scalable Machine Learning (PRACE course)	Jülich	Jan 29-31, 2019
Introduction to hybrid programming in HPC (PRACE course)	Garching	Jan 2019 (tba)
Tuning and Scaling Workshop	Jülich	Feb 2019 (tba)
Node-Level Performance Engineering (PRACE course)	Garching	Feb 20-21, 2019
Introduction to Python	Jülich	Mar 2019 (tba)
OpenMP GPU Directives for Parallel Accel. Supercomp. (PRACE couse)	Stuttgart	Mar 2019 (tba)
Programming with Fortran	Garching	Mar 6-8, 2019
CFD with OpenFOAM®	Stuttgart	Mar 11-15, 2019
Introduction to ParaView for the visualization of scientific data	Jülich	Mar 14, 2019
Introduction to parallel programing with MPI and OpenMP	Jülich	Mar 18-22, 2019
Iterative Linear Solvers and Parallelization	Stuttgart	Mar 25-29, 2019
Advanced Topics in High Performance Computing (PRACE course)	Garching	Mar 2019 (tbc)
Introduction to ANSYS Fluid Dynamics	Garching	Apr 2019 (tba)
GPU Programming with CUDA (PRACE course)	Jülich	Apr 1-3, 2019
Fortran for Scientific Computing (PRACE course)	Stuttgart	Apr 1-5, 2019
Interactive High-Performance Computing (PRACE course)	Jülich	Apr 9-10, 2019
Scientific Visualization	Stuttgart	May 2019 (tba)
Programming in C++	Jülich	May 6-9, 2019
Cray XC40 Workshop on Scaling and Node-Level Performance	Stuttgart	May 6-10, 2019
From zero to hero, Part I: Understanding and fixing on-core performance bottlenecks	Jülich	May 14-15, 2019
Advanced C++, Focus on Software Engineering	Stuttgart	May 14-17, 2019
Usage of VTK for scientific-technical visualization	Jülich	May 20, 2019
Introduction to Programming and Usage of Supercomputer Resources at Jülich	Jülich	May 20-21, 2019
High-performance scientific computing in C++ (PRACE course)	Jülich	May 27-29, 2019
Fortran Modernization Workshop	Stuttgart	Jun 2019 (tba)
Introduction to Deep Learning Models	Jülich	Jun 2019 (tba)
Introduction to hybrid programming in HPC	Stuttgart	Jun 2019 (tba)
HPC code optimization workshop (PRACE course)	Garching	Jun 2019 (tbc)
Intel Manycore Programming Workshop (PRACE course)	Garching	Jun 2019 (tbc)
Cluster Workshop	Stuttgart	Jun 25-26, 2019
Node-Level Performance Engineering (PRACE course)	Stuttgart	Jun 27-28, 2019
Concepts of GASPI and Interoperability with other communication APIs (PRACE course)	Stuttgart	Jul 1-2, 2019
Advanced C++, Focus on Software Engineering	Stuttgart	Jul 9-12, 2019

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

lrz 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 Phase 1" IBM System x	IBM iDataPlex 9216 nodes, 147456 cores, 288 TByte, FDR10	3,185	Capability Computing	
	205 nodes, 8,200 cores Westmere EX 52 TByte, QDR	78	Capability Computing	
	32 accelerated nodes Knights Corner 76 GByte, FDR14	100	Prototype System	
"SuperMUC Phase 2" Lenovo Nextscale	3072 nodes, 86016 cores, Haswell EP 197 TByte, FDR 14 IB	3,580	Capability computing	German universities and research institutes, PRACE (Tier-0 System)
"CooLMUC2" Lenovo Nextscale	252 nodes, 7,056 cores Haswell EP 16.1 TByte, FDR 14 IB	270	Capability computing	Bavarian Universities (Tier-2)
"CoolMUC3" Megware Slide SX	148 nodes, 9472 cores, Knights Landing, 14.2 TByte, Omnipath	383	Capability Computing	Bavarian Universities (Tier-2)
Compute Cloud Linux-Cluster	200 nodes, 2700 cores	18	Capability Computing	Bavarian Universities, LCG Grid

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	Laki: 652 TFlops Laki2: 74 TFlops	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 bi-annual 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: Illustration shows fine-scale mixing process (the lines) passing through the flame (gray surface) during combustion. More information on page 13. © Felix Dietzsch

GCS
Gauss Centre for Supercomputing