INNOVATIVES SUPERCOMPUTING IN DEUTSCHLAND N° 18-2 · AUTUMN 2020

INSiDE

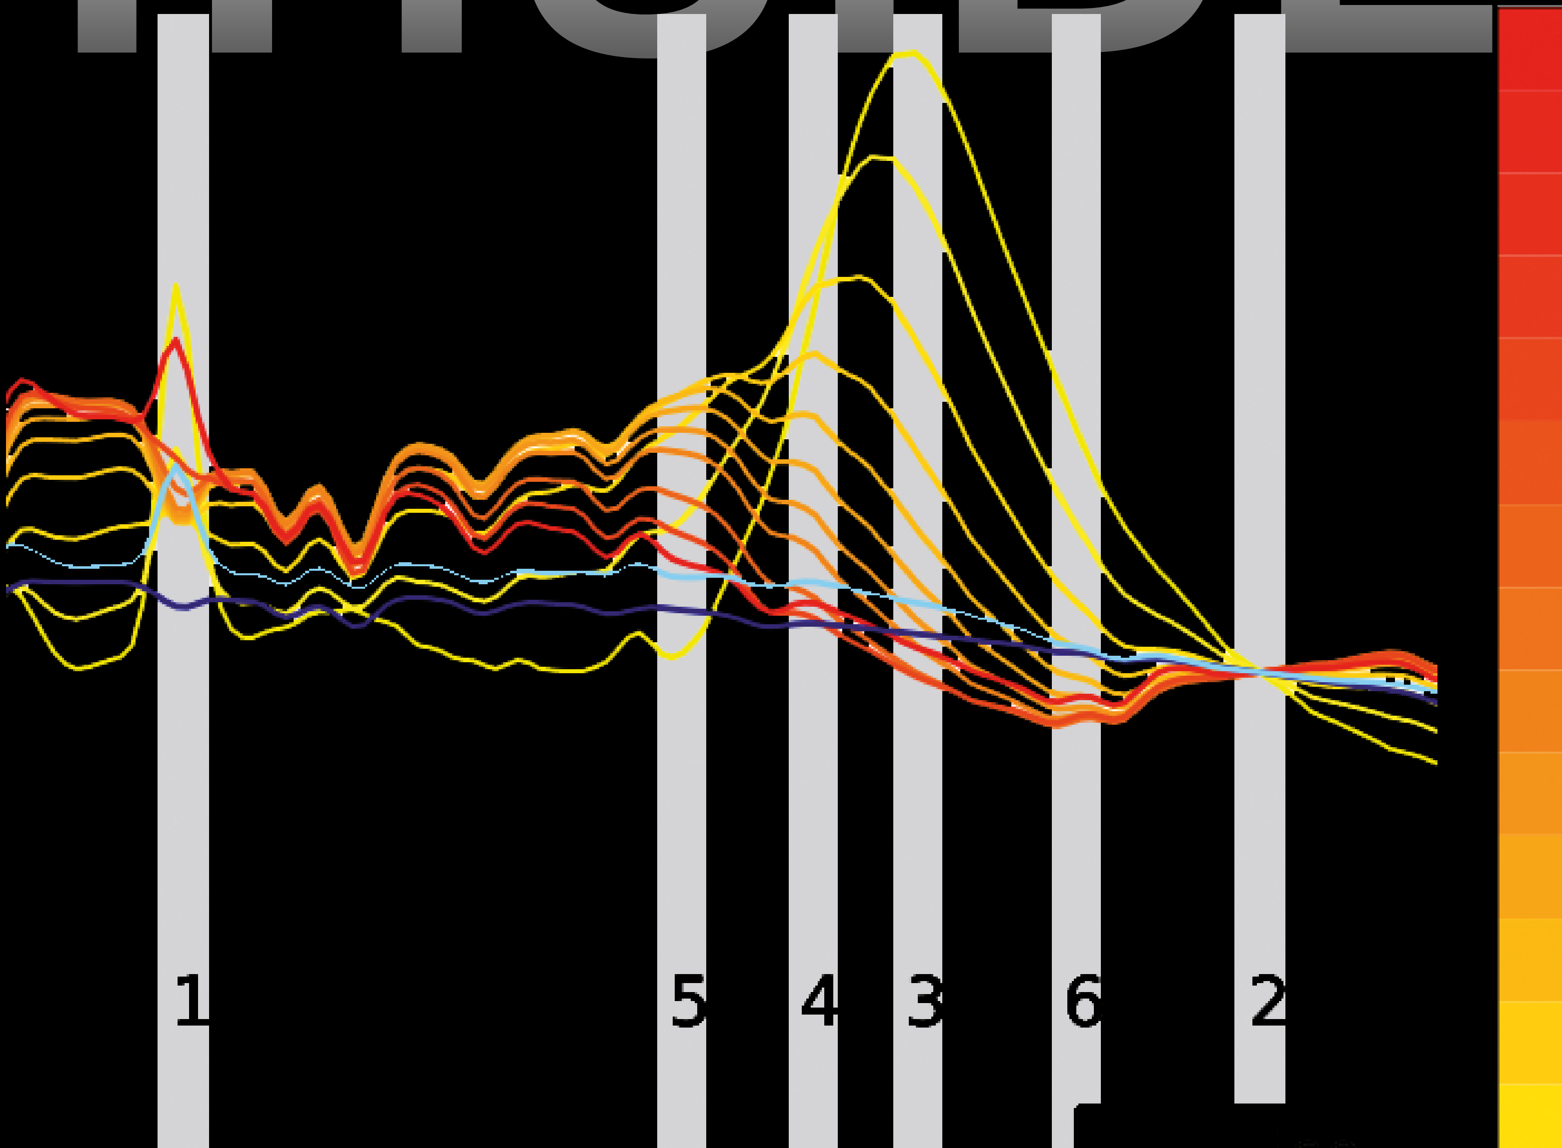

Publishers

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

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

Dr. Alexander Trautmann (*at*), JSC
a.trautmann@fz-juelich.de

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

Susanne Vieser (*sv*), LRZ
susanne.vieser@lrz.de

Visit us at:
www.gauss-centre.eu

Welcome to the latest issue of InSiDE, the bi-annual Gauss Centre for Supercomputing magazine showcasing innovative supercomputing developments in Germany. As the world continues to grapple with the COVID-19 pandemic in the second half of 2020, GCS centres have continued to make high-performance computing (HPC) resources available to researchers in academia, government, and industry to accelerate development of vaccines and other therapeutics, forecast potential emerging hot spots to slow the spread of the disease, and generally ensure that HPC can be quickly accessed by researchers and policy makers in times of emergency.

Despite the complications brought about by COVID-19, our centres remained dedicated to delivering on their key missions—making world-leading HPC resources available for researchers tackling society's most difficult scientific challenges and training current and prospective users how to use these resources as efficiently as possible. While the centres have been unable to host visitors for our wide variety of annual training courses, staff at the centres worked extremely hard to move our robust training program online. In the process, we learned how to expand our training offerings (Page 36). The Jülich Supercomputing Centre and High Performance Computing Center Stuttgart battled through COVID-related complications to stand up their next-generation machines this fall, with JSC welcoming the JUWELS Booster module and HLRS putting its Hawk system online (Page 4).

Our users continue to use our resources to improve core societal problems. Using HPC resources at the Leibniz Supercomputing Centre (LRZ) and HLRS, researchers are collaborating to couple weather and climate models with high-resolution engineering simulations to more efficiently site wind turbines in hilly, complex terrain (Page 12). Researchers at JSC and the University of Wuppertal are partnering with experimentalists to model complex chemical interactions in clouds to better understand ozone depletion (Page 18). University of Stuttgart researchers are using HLRS resources to model laboratory-scale biochemical experiments to understand how they change at the industrial scale.

Whether the end of the pandemic is just around the corner, or the world needs more time to develop and scale effective treatments, GCS centres will continue to ensure our resources are available to help humanity find a way out of this pandemic while continuing to support our staffs and users during these trying times.

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

Imprint	2
Editorial	3
NEWS FEATURES	
As the JUWELS Booster Comes Online, Researchers are Reaching New Computational Heights	4
GCS Leads Efforts to Increase HPC Expertise Across Europe, Including Among SMEs	7
Testing HPC Technologies and Shaping the Future of Computers	9
SCIENCE FEATURES	
GCS Centres Support Multi-Disciplinary Approach to Designing and Siting Wind Farms in Complex Terrain	12
Researchers Use Computation to Scale Up Biochemical Processes	15
Researchers Fly High with JUWELS and Experimental Aircraft to Better Understand Ozone Depletion	18
PROJECTS	
HLRS Will Help Build National Data Research Infrastructure for Catalysis Research	21
Full Steam Ahead to the Exascale SEAs	24
ORIGINS: Answers to Existential Questions	26
NEWS BRIEFS	28
INSIDE GCS	
Staff Spotlight: Long-Time JSC Group Leader Helps Spearhead Next-Generation Networking Technologies	34
During Pandemic-Related Remote Working, GCS Centres Embrace Expanding E-Learning Offerings	36
TRAINING CALENDAR	39
CENTRE DESCRIPTIONS	
Jülich Supercomputing Centre	40
Leibniz Supercomputing Centre	42
High-Performance Computing Center Stuttgart	44

AS THE JUWELS BOOSTER COMES ONLINE, RESEARCHERS ARE REACHING NEW COMPUTATIONAL HEIGHTS

JSC welcomes the newest component to its novel modular supercomputer, a GPU-boosted system to further expand the centre's footprint in supporting emerging technologies.

In 2017, the Jülich Supercomputing Centre (JSC) embarked on an ambitious course to install a "modular" supercomputer as its flagship machine. The plan's first results materialized in Autumn 2018 when the first piece of the centre's newest supercomputer, JUWELS, went online. In essence, JUWELS was conceived as a series of several high-performance computing (HPC) systems, tightly integrated and networked into one larger machine that could enable researchers to use an architecture best suited for their respective research goals.

The traditional supercomputer lifecycle lasts roughly 5 years, then it is "out with the old, in with the new" for hardware. Bucking that trend, JSC leadership analyzed the market options and identified technologies to integrate and evolve the system along the centre's longer-term vision. JUWELS combines the best technologies from the project partners Atos, ParTec, Mellanox and NVIDIA, many of which were just dots on roadmaps toward a long-term integrated system when the project began.

More than 2 years into the project, JUWELS just got a boost—the JUWELS Booster module, the second component in this multi-pronged approach, was installed this fall. The GPU-accelerated module brings JUWELS's performance from 12 petaflops, or 12 thousand trillion calculations per second, to roughly 85 petaflops.

While the performance increase is an obvious advantage, JSC staff see the Booster module as offering far more than just raw computing power—the GPU-centric addition offers new opportunities for JSC to expand further into emerging technologies related to artificial intelligence (AI) and machine learning (ML) applications.

JSC staff members organized a robust early user program that identified and closely worked with long-time allocation awardees across a broad spectrum of scientific disciplines. These staffers knew that working closely with a diverse subset of users in the lead up to the Booster was essential in order to effectively scale up their application codes and take full advantage of the performance increase. Additionally, the GPU-centric architecture offers users the opportunity to expand research into scalable AI and ML applications on the new machine. Early returns have been promising, and the machine is primed to be the JUWEL of European high-performance computing for years to come.

Two parts to a whole

As a member of the Gauss Centre for Supercomputing (GCS), JSC staffers understand that each investment in next-generation HPC hardware needs to serve a broad user base, and with the rise of new artificial intelligence, machine learning, and data analytics technologies, users' needs are only getting more diverse.

When JSC began the first phase of JUWELS, it had a long-term roadmap in mind, but just like any long journey by car, the centre did not have an exact route set in stone. JSC staff knew that the next JUWELS module needed to offer architectural diversity and a large performance boost, but what that looked like exactly only came later.

"When we installed the cluster in 2018, we knew it was just the start of the journey," said Dr. Dorian Krause, Team Lead of JSC's High-Performance Computing Systems group. "From the beginning, the design of the cluster was such that we could extend it later on with something else. Back in 2017 it wasn't 100 percent clear that "Booster" meant GPUs. But at the moment, it's quite clear in the industry that GPUs are the best suited processor technology for the type of highly-scalable systems we are building."

In addition to the raw performance increase, JSC staff knew that having a more GPU-centric addition to JUWELS would enable researchers focused on AI and ML applications to use an architecture suited to their needs. As a GCS centre, though, the procurement had to be focused on balancing the needs of the many scientific stakeholders interested in using Germany's leading compute resources.

To that end, JSC designed the Booster's node architecture to carefully balance the needs of different workloads and provide a consolidated system well suitable for traditional and emerging scientific use cases. "Besides optimizing parameters such as GPU count per node and CPU types, we have paid special attention to the network performance of the system to ensure that the applications will be able to leverage the entire system with high efficiency," Krause said.

On a practical level, JSC users will see the two parts of JUWELS as one machine, submit requests for computing time on either partition or both partitions, and log on to one machine. These characteristics only work due to the innovative approach to high-speed networking the centre has carried out, resulting in more than 200 cables delivering a 5-terabyte-per-second connection between the machine's partitions (for more information on JUWELS network, please visit page 34).

Early user program emphasizes application diversity

Whenever a new supercomputer comes online, HPC centres try to ensure that their users can take full advantage of the machine as soon as possible. Early success requires making sure staff and users are trained on the newest architectural changes, porting users' codes to a new architecture, and helping users scale their codes to maximize performance.

For the Booster module, JSC identified 13 different applications representing a broad swath of JSC's user base to get early access to the first Booster nodes. Most users come from Germany, but several teams come from other European nations. JSC staff identified applications that were representative of JSC's research domain portfolio, with codes from quantum chromodynamics, earth science, neurology, astrophysics, plasma physics, climate modelling, and the JUQCS quantum computing simulator, among others, taking part in the early access program.

Several codes selected came out of close collaborations with JSC's domain-specific "SimLabs," which serve as advanced user support environments where researchers can collaborate with JSC staff who not only can help assist them with basic aspects of porting their applications, but also have scientific backgrounds in the domain in question.

"In quite a few cases, contact was established with help of the SimLabs," said Dr. Andreas Herten, organizer of JSC's Booster early user program. "We hope now that we can learn from each other and find out what strengths different codes have and where there are needs for improvements. We want to ensure we're prioritizing different code bases we deal with in the day to day and ultimately fine-tune the system by looking at how a diverse set of highly scalable codes perform."

The GPU-accelerated JUWELS Booster module, installed this year, brings the modular machine's peak performance to roughly 85 petaflops. **© Forschungszentrum Jülich**

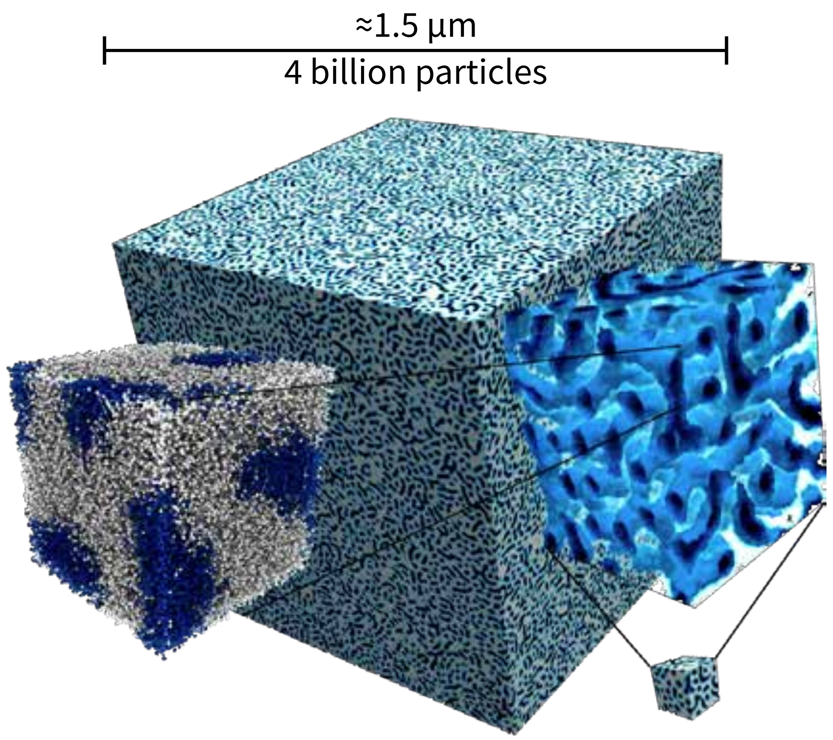

The scale of JUWELS Booster allows SOMA to investigate the fabrication of membranes battery materials with nanometer features up to the size of micrometers—a scale that is also accessible in experiments via common microscopy techniques. With this merging of length scales, JUWELS Booster enables direct comparison to experiment and application, accelerating the development cycle.

© Ludwig Schneider, Georg-August Universität Göttingen

Primed for exascale computing while expanding into new territory

The HPC community draws closer by the year to the once distant exascale horizon. Exascale machines represent roughly 12–20 times the computing power of today's most powerful European supercomputers, and offer the chance for massive advancements in research areas from understanding the human brain to high-precision climate modelling. In 2018, the European Union funded the EuroHPC Joint Undertaking (JU) to chart the path toward European exascale computing. The JU has already funded three European pre-exascale systems, designed to be competitive with the current multi-hundred-petaflop systems among the fastest machines today. The JU plans to follow that up by delivering at least one European exascale system in 2023.

While JUWELS is not a "pre-exascale" system in this sense, JSC staff see the combination of advanced user support, tight interaction with domain experts and via the JSC SimLabs, flexible, diverse architectures, and a robust HPC infrastructure as a model for an eventual European exascale system. GCS centres such as JSC have long played key role in providing users with the necessary data storage and processing infrastructure required to make the most out of the supercomputing infrastructure. Exascale computing will only increase infrastructural and data-storage-related demands. Success with an exascale system will also largely be determined by institutional know-how, and JSC has a long tradition as one of Europe's premier HPC centres, and has regularly focused institutional resources on the front line of HPC research and building machines designed to be among the fastest in the world.

While the next phase of JUWELS offers a massive boost in computing power for traditional modelling and simulation, the Booster module offers 3,700 NVIDIA A100 GPUs, designed with artificial intelligence applications specifically in mind. Herten participates in JSC's Exascale NVIDIA Application Lab, set-up to help users to scale and co-design applications on GPU architectures to push the boundaries of computing performance on the path toward exascale. "We have the ability to really work in detail with applications at varying degrees," he said. "So for some we are just saying, 'hey, you got this module wrong and you forgot this switch, then you are good to go,' but for others, we are really actively contributing to the code and trying to find out what is the bottleneck and what can be done better."

In addition to purely offering hardware and support through the early user program, JSC has continued to regularly host training events focused on helping the broader European HPC community learn more about how GPU programming could benefit different research domains. The centre has one of Europe's longest-running GPU training courses, regularly offers training on the CUDA and OpenACC programming paradigms, and hosted regular seminars on the future of modular supercomputing architectures.

Much like space shuttle using booster rockets to get into orbit, JSC's latest Booster module is helping launch European HPC toward the exascale era.

eg

GCS LEADS EFFORTS TO INCREASE HPC EXPERTISE ACROSS EUROPE, INCLUDING AMONG SMES

Three new interrelated projects will create a Europe-wide network of national HPC competence centres, promote the exchange of knowledge about HPC best practices across borders, and facilitate HPC access for small and medium-sized enterprises.

The EuroHPC Joint Undertaking was organized in 2018 to accelerate the development of a shared European infrastructure and knowledge base in high-performance computing (HPC). On September 1, the Gauss Centre for Supercomputing (GCS) and its partners across Europe celebrated the official kick-off of three closely related, international projects that GCS will be leading to make progress toward EuroHPC's goals.

The High-Performance Computing Center Stuttgart (HLRS), one of the three centres comprising GCS, will be coordinating the three projects. The first, called EuroCC, is leading the development of HPC competence centres in 33 European countries, each of which will become a resource for its national HPC community, including both academia and industry. The second, called CASTIEL, will lead activities to document resources and expertise across Europe, and promote the sharing of knowledge through cross-border collaboration. (For more information about EuroCC and CASTIEL, please see the Spring 2020 issue of InSiDE. https://www.gauss-centre.eu/news/publications/article/inside-spring-2020/

The third pillar of this effort comes with the launch of a new European consortium called FF4EuroHPC. This three-year project, funded by a grant of approximately 10 million EUR from the European Commission's Horizon 2020 program, will undertake targeted outreach to small and medium-sized enterprises (SMEs), facilitating access to HPC technologies and solutions that will make them more competitive.

In addition to spearheading EuroCC and CASTIEL, HLRS will serve as the coordinating centre for FF4EuroHPC, working together with partners at Scapos AG (Germany), CINECA (Italy), Arctur d.o.o. (Slovenia), TERATEC (France), and CESGA (Spain).

FF4EuroHPC builds on the success of Fortissimo

The core team leading FF4EuroHPC includes several HPC centres that were also involved in the EU-funded Fortissimo and Fortissimo 2 projects, which between 2013 and 2020 organized 92 successful business experiments involving SMEs and high-performance computing. FF4EuroHPC will use and extend the lessons learned during Fortissimo to identify and test new HPC applications for industrial research across Europe.

Fortissimo's past successes have enabled SMEs, working with HPC experts, to develop solutions to a wide range of problems. These have included, for example, virtual testing of safety systems for driving, multiphysics simulations for micro speaker design and development, and simulations of ultra-high-temperature furnaces, among many others.

FF4EuroHPC will organize two new open calls for proposals from small and medium-sized enterprises that can demonstrate compelling cases for using HPC technologies and expertise. Whereas in the past Fortissimo and Fortissimo 2 focused primarily on applications of HPC for manufacturing, FF4EuroHPC will be open to a broader range of engineering use cases, as well as innovative cases from other growth sectors. Proposals are welcome not only to use traditional applications of high-performance computing for simulation and modeling, but also for new kinds of applications involving data analytics, machine learning, and artificial intelligence. Outstanding proposals will receive support through FF4EuroHPC to undertake "application experiments" addressing the business problems they propose.

According to HLRS General Manager Dr. Bastian Koller, who will be leading FF4EuroHPC, EuroCC, and CASTIEL, "We see great opportunities for FF4EuroHPC to lower the barriers that typically make it difficult for SMEs to innovate using HPC. We expect that the project will help businesses to improve their product design and development processes, ultimately leading to better products and services. Because FF4EuroHPC will be so closely integrated with EuroCC and CASTIEL, we also expect that this project will raise the global competitiveness of Europe as a whole."

In an application experiment completed during the Fortissimo project, HLRS supported SPICETECH GmbH by providing HPC resources used to develop software for virtual testing of autonomous driving systems. **© SPICETECH**

GCS at centre of European HPC strategy

The fact that coordination of the three new pan-European projects is all based within GCS means that FF4EuroHPC will work closely with the national HPC competence centres that have been created for the new EuroCC network. These competence centres will be integrated into outreach activities surrounding the open calls for proposals from SMEs, as well as in the execution of research experiments.

Extending the prior Fortissimo projects, which involved a smaller number of HPC centres, the new project looks forward to receiving applications from SMEs that will pair with or receive support from national HPC competence centres across the EuroCC network, and potentially other European HPC centres as well. Other kinds of partnerships — involving Digital Innovation Hubs created to support businesses in digital transformation, for example — are also expected to participate. By the end of the project cycle, FF4EuroHPC expects approximately 40 SMEs to participate, all of which will be new users of HPC. Through this approach, FF4EuroHPC will promote industrial HPC usage among SMEs in regions where it has previously been slow to develop.

The FF4EuroHPC organizers anticipate that the success stories resulting from its activities, taking place in combination with the overall EuroHPC strategy for increasing HPC capabilities across Europe, will demonstrate the opportunities that access to HPC offers for small and medium-sized enterprises. This will help not only to attract interest from other SMEs that could benefit from HPC, but also to gain support from governments and other stakeholders for the development and maintenance of robust HPC resources and services. *cw*

TESTING HPC TECHNOLOGIES AND SHAPING THE FUTURE OF COMPUTERS

Researchers at LRZ are building the Bavarian Energy, Architecture, and Software Testbed in order to help proactively use and shape emerging technologies.

The Leibniz Supercomputing Centre (LRZ) is implementing the High Tech Agenda of the Bavarian State Government in its field of expertise and is launching the ambitious "Future Computing" program. Firstly, it is setting up a test environment with the latest computer technologies available on the market. Secondly, it is developing offers to train both staff members as well as the next generation of HPC scientists to exploit and explore new computer technologies and high-performance computing (HPC) systems in collaboration with selected key scientific partners.

"We want to intensively research the latest computer systems and architectures, their energy requirements, and mode of operation, without disturbing the services for our users at the production systems at LRZ," explains the computer scientist Josef Weidendorfer, who heads the Future Computing group g at the LRZ. To his end, login and storage servers are already available in Garching, as well as two AMD Rome systems and two servers with Marvell ThunderX2 processors, each with GPUs as accelerators. The centre recently installed a Cray CS500 system, which uses the same Fujitsu A64FX processors as Fugaku, a Japanese system that is currently the world's fastest supercomputer. Over the next few years, the test environment will be steadily expanded to include more systems and components. LRZ leadership envisions this test environment as becoming a permanent fixture of the research work at the centre and it will also serve to evaluate new computer architectures for Bavaria's largest scientific computing centre.

Super-Tech at LRZ: Preparing for the next generation of high-performance computers

"Bavarian Energy, Architecture, and Software Testbed" or BEAST, is the name the LRZ has chosen for this innovative collection of computer and storage technologies that LRZ's specialists are now putting through their paces.

"BEAST serves as a rich environment to prepare for the next generation of supercomputers," says Prof. Dr. Dieter Kranzlmüller, Director of LRZ. "We use it to investigate which computer architectures are suitable for larger systems and for parallelization. With the experience we gain through BEAST, we will be able to plan the successor to SuperMUC-NG and future services even better and more soundly." In research, the amount of data that supercomputers have to handle is currently growing. Applications such as machine learning and artificial intelligence also require new chip design, or even totally different computer architectures.

The needs posed by these emerging technologies have already begun influencing HPC systems, and in the near future, the need to help organize work or memory performance more efficiently will only grow. As a result, new ideas for computers are needed, and HPC centres need to play an active role in research into the benefits of new technologies and architectures. In the long run, BEAST will therefore contain prototypes and the latest systems, which LRZ wants to further develop and build together with manufacturers and with key research partners. "BEAST is not a conventional LRZ service," Kranzlmüller continues. "However, joint development and co-design of new technology is pushing supercomputing and will ultimately pay off for science and society.

Open for questions from researchers and students

Although BEAST is not a service of LRZ in the traditional sense, selected user groups will be given access to the systems. LRZ specialists will be the first users of these novel systems in order to familiarize themselves with new computer systems and processors. Further, HPC software, such as the monitoring tool DCDB or the smart control system Wintermute, need to be adapted for future systems. "With BEAST, we can prepare ourselves to offer modern, complex architectures and necessary software environments at service quality," says Weidendorfer. "The test environment also enables less goal-oriented and significantly more experimental research projects. We are therefore opening the test environment to selected researchers who are working on next-generation hardware". These researchers will be able to configure operating systems and hardware according to individual needs and modify them for their own applications. The LRZ will actively accompany and support this work and observe how hardware can be built and used more efficiently.

Gather contacts and get to know innovative technology

BEAST is interesting for young scientists and junior technical staff: LRZ, in collaboration with the two Munich universities, now offers an internship for computer science students interested in modern computer architectures and their energy-efficient use, allowing the next generation of HPC experts early exposure to emerging technologies that will likely play major roles in their professional careers. These internship opportunities aim to motivate students to make use of the latest computer technologies for bachelor or master theses.

Participants will use BEAST to familiarize themselves with the technologies of the future and also make valuable contacts. They will regularly solve research questions and practical tasks on and with the new systems. In addition, technology companies will describe and contextualize new system designs and discuss them with young scientists. The BEAST internship will be coordinated by Dr. Karl Fürlinger (Ludwigs-Maximilians University Munich), Dr. Weidendorfer, and the computer scientist Bengisu Elis (Technical Uuniversity of Munich).

“I'm researching how to optimize communication in Graphic Processing Units (GPU),” Elis says. “With BEAST, I can compare combinations of GPU and CPU architectures from different vendors, and I can also test and improve the performance and portability of my code on different systems, BEAST offers valuable resources to increase the quality of my work, the idea that these can also contribute to improve the future systems of the LRZ is an additional incentive.” *sv*

The Bavarian Energy, Architecture, and Software Testbed (BEAST) at LRZ serves as an innovative collection of emerging compute and storage technologies to test. **© LRZ**

GCS CENTRES SUPPORT MULTI-DISCIPLINARY APPROACH TO DESIGNING AND SITING WIND FARMS IN COMPLEX TERRAIN

Supercomputing simulations support the design of a research station to improve wind turbine efficiency in hilly and mountainous regions.

In recent years, Germany has gained the ability to generate almost a quarter of its electricity through wind power—a 15 percent improvement in the last decade. As the country continues to push toward generating 100 percent of its energy through renewable sources, wind turbines will play an increasingly important role.

Despite these gains, wind energy is still geographically limited, with the vast majority of wind turbines located in the flat northern half of the country. Germany's southern states, Baden-Württemberg and Bavaria, have hilly or mountainous terrain that make wind turbine placement more complicated.

From modeling weather patterns and the impact of forest on turbine efficiency, to understanding turbulence's effects on wind turbines, to charting wear and tear on wind turbines, there are myriad issues to measure and take into account to ensure the best return on investment for new wind farms.

Research institutions in southern Germany have been working on improving wind turbine efficiency in their back yards for 70 years. Recently, the Baden-Württemberg Ministry of Science, Research, and the Arts (MWK) sponsored the WindforS project, which brings together multidisciplinary experts to improve wind turbine efficiency. One of the WindforS subprojects, WINSENT, which began in 2017, supports researchers using high-performance computing (HPC) in conjunction with experiment to help design a first-of-its-kind experimental wind energy test station in the Swabian Alb.

WINSENT is a collaboration between the University of Stuttgart, Karlsruhe Institute of Technology (KIT), University of Tübingen, Center for Solar Energy and Hydrogen Research Baden-Württemberg, Technical University of Munich, and Esslingen University of Applied Science (Hochschule Esslingen) and is funded by MWK and the German Federal Ministry for Economic Affairs and Energy (BMWi).

Researchers from Uni Stuttgart and KIT have been using Gauss Centre for Supercomputing (GCS) HPC resources at the High-Performance Computing Center Stuttgart (HLRS) and Leibniz Supercomputing Centre (LRZ) in Garching near Munich to support an innovative, three-part modelling approach that integrates local weather and climate data with high-resolution simulations of turbulence around wind turbine rotor blades.

Three-phase simulations set the stage for more efficient wind energy

Finding the most efficient place for wind turbines in complex, hilly terrain requires much more than just being able to predict wind speed. Often, such terrain is also heavily wooded and has steep inclines and declines. The weather can also be more punishing for wind turbines, as wind gusts and heavy storms can be more severe at higher elevations.

Even modern-day supercomputers are unable to do meaningful simulations that take into account the complex collection of environmental and engineering considerations for siting wind turbines in these environments. As a result, the WINSENT collaborators have developed a three-part simulation approach that focuses on different aspects of the simulation at different resolutions. Each step in the process generates datasets that ultimately inform other steps in the process, making the overall simulation suite more accurate, and thus, more predictive.

Researchers need to ensure that wind turbine investments not only are located where wind will keep them operating at the highest efficiency possible, but also ensure that increased maintenance costs and smaller maintenance intervals do not raise costs too high. Further, by better understanding the turbulent dynamics around turbine rotor blades and the influence of the wake—the complex, altered airflow downstream of a turbine that can cause ripple effects on other turbines close by—researchers are able to better understand how to place groups of turbines to prevent some turbines from negatively affecting the efficiency of others.

Despite focusing on different calculations at different resolutions, the WINSENT collaborators all create a computational mesh for their simulations—that is, they divide their simulations into a grid and solve equations that govern the behavior of wind, turbulence, or other factors within each box of the grid over time. The smaller the mesh or the finer the focus on intricate details, the more computationally expensive the simulations become.

In the first step of the three-phase WINSENT simulation process, KIT researcher Dr. Daniel Leukauf uses SuperMUC-NG at LRZ to run simulations primarily focused on the meteorological aspects of turbine siting,

SCIENCE FEATURES

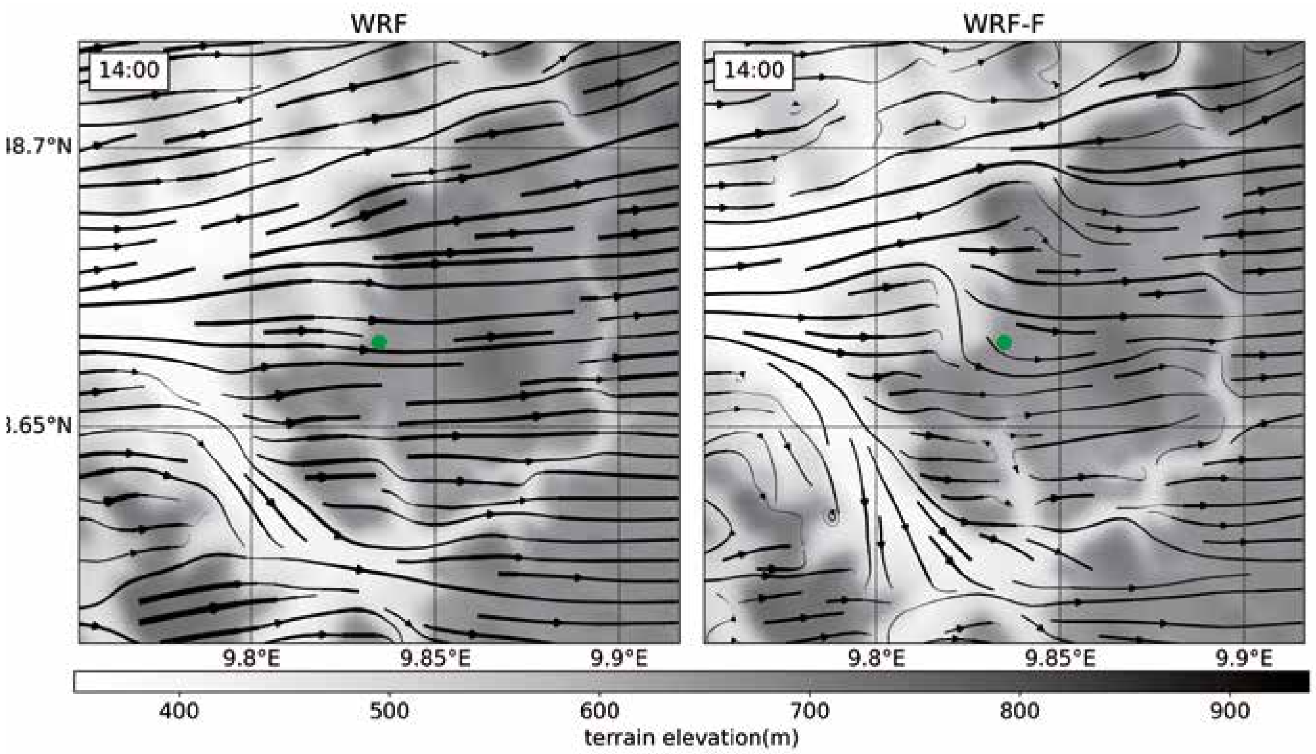

KIT researchers are running meterological models using the Weather Research and Forecasting (WRF) model to support the WINSENT project. In this image, the team shows streamlines of the flow over the proposed test site at 60 metres above ground. The green dot marks the test site and topography is in grey. The slopes are forested. The right picture shows a simulation with forest drag, the left one is the WRF model without modifications. The forest drag has a strong impact on wind direction and wind speed and must be included in site assessment. **© Daniel Leukauf, KIT**

creating meshes at a resolution of roughly 150 metres. These simulations serve as the basis for boundary conditions for increasingly higher-resolution models, and assist the researchers in identifying promising sites for a wind farm.

Leukauf uses the Weather Research and Forecasting (WRF) model, developed by the National Center for Atmospheric Research in the United States and widely used for high-resolution weather forecasting and climate modelling simulations. WRF offers the possibility to model weather conditions at local scales, but its resolution is insufficient for accurately modeling the influence of weather conditions on wind turbines' rotor blades in complex terrain. Despite not being able to include these aspects in his simulation, Leukauf is able to simulate multiple days of detailed meteorological data that can then be passed on to collaborators at Hochschule Esslingen.

The Hochschule Esslingen team uses Leukauf's data to inform the initial conditions for a higher-resolution simulation. The team uses OpenFOAM, a widely used computational fluid dynamics code, to simulate a wind turbine site at roughly 20-metre resolution, further refining complex terrain simulations by resolving air flow and turbulence happening on a level more relevant for the turbines themselves. The team then passes its data to a research group at Uni Stuttgart to inform its ultra-high-resolution simulations.

Patrick Letzgus of the Uni Stuttgart group runs extremely accurate detached eddy simulations (DES). The team's DES

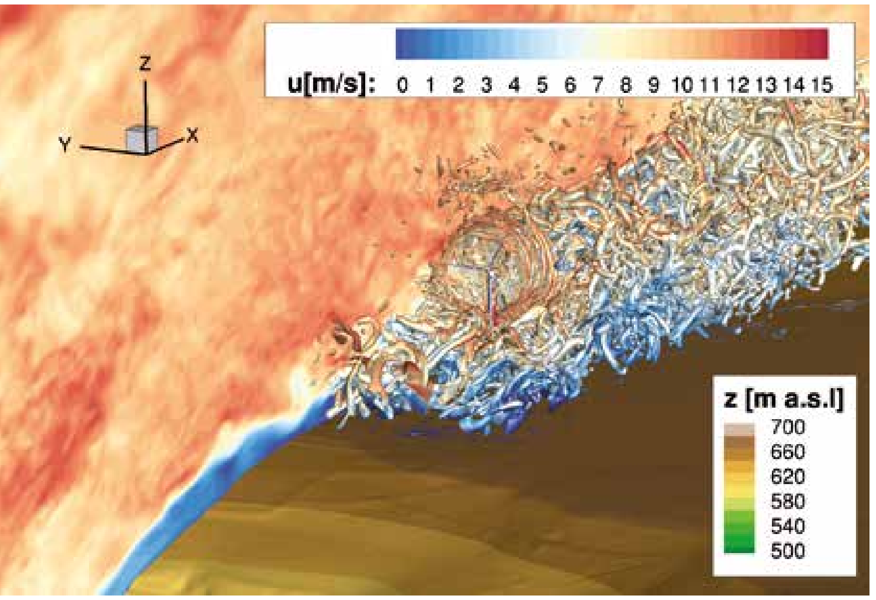

Uni Stuttgart researchers are running ultra-high-resolution detached eddy simulations as part of the WINSENT project. This image shows the topography on the ground, color coded by elevation. Above, a vertical plane depicts the flow fields, colored by the value of the horizontal wind velocity. There are also turbulent structures highlighted in the flat plateau on top of the hills, also colored to represent wind speed. **© Patrick Letzgus, Uni Stuttgart**

simulations can model the subtle but essential turbulence that happens around wind turbine rotors, including the wake, while also taking into account the foliage and local topography near a wind turbine site at much greater detail. At a resolution of one metre, the DES simulations are extremely accurate, but also extremely computationally expensive.

One simulation requires roughly 7,000 compute cores on HLRS's HPC systems and need to run for weeks in order to calculate only a few minutes of simulation time. Early runs were done on the Cray XC40 Hazel Hen machine and more recently on the early access platform of the Hewlett Packard Enterprise Apollo system Hawk as well as SuperMUC-NG at LRZ.

Despite a relatively small amount of real time being modeled in these simulations, these calculations can be extrapolated to understand the turbulent dynamics for a wide variety of weather conditions at wind farm.

Standing up a test site and developing tools for industry

The WINSENT project collaborators have finally gone through the lengthy permitting process for building their experimental facility on the Swabian Alb near Göppingen, Germany, and anticipate that it will be constructed in the coming months. "A wind energy test field of this kind doesn't exist in complex terrain," Letzgus said. "It is great that it is coming so that many groups and disciplines can use it for research."

The WINSENT collaboration has streamlined its simulation chain, and is looking forward to being able to start comparing simulation data directly to experimental data gathered at the test site. Leukauf indicated that doing so will ultimately make it possible to develop software that will help industry engineers without access to HPC to determine where wind turbines will operate most efficiently in complex terrain. Improving wind energy in complex terrain will ultimately help Germany rely more on renewable energy sources.

Despite the end goal being software that can run on a laptop, Letzgus indicated that without access to GCS HPC resources, the team would be unable to have created such an accurate, detailed modelling chain. "Without the access I have to HPC resources, these highly resolved simulations, and my work in general, wouldn't be possible," he said. *eg*

SCIENCE FEATURES

RESEARCHERS USE COMPUTATION TO SCALE UP BIOCHEMICAL PROCESSES

High-performance computing enables bioengineers to predict how laboratory results can be transferred to industrial conditions without loss of performance.

In the multidisciplinary field of bioengineering, scientists use principles from engineering to understand, utilize, and optimize biological processes and systems. This can include, for example, designing biological systems to synthesize new products, developing processes for large-scale applications of biotechnology, or even using biological systems as models for addressing complex human problems. Whether we know it or not, all of us use biotechnology products daily. Methods and technologies based in bioengineering could, for example, help to develop natural alternatives for wasteful or dirty industrial processes, including manufacturing plastics, creating food additives, making technical enzymes for washing powders, or producing pharmaceuticals, among others.

Biotechnology now makes it possible to alter certain parts of a bacteria's DNA—its genetic code—to imbue it with traits that can address a given task. Biotech processes utilize renewable resources such as sugar for carbon supply, and specially engineered microbes often convert that sugar into products of interest. As living cells, they need optimal cultivation conditions to perform as they should.

How microbes, sugars, and other compounds react in a laboratory, though, can be much different than how they behave in industrial-scale bioreactors, which often exceed 500,000 litres.

A research team led by Professor Ralf Takors in the University of Stuttgart's Institute for Biochemical Engineering (IBVT) has been tackling this problem. "As a society we want to develop sustainable processes for the production of biology-based goods, be they commodities, chemicals, or pharmaceuticals," he said. "For this to be practical, however, you have to scale up the processes you develop in the lab to large-scale reactors. This is no easy task."

Takors' multidisciplinary team has partnered with other academic and industrial researchers to investigate how methods from a field called computational fluid dynamics—a discipline often used to model aerodynamics or combustion processes, for example—could be used to simulate how biochemical reactions studied in the lab will occur at industrial scales. By developing detailed computational models, this approach could avoid the need to conduct expensive and time-consuming trial-and-error experiments that would otherwise be impractical.

Recently, the group used high-performance computing (HPC) resources at the High-Performance Computing Center Stuttgart (HLRS), one of the three centres comprising the Gauss Centre for Supercomputing (GCS), to model the behavior of *Corynebacterium glutamicum*, a widely used and well-understood bacterium used in producing amino acids and food additives. The work, which demonstrates close collaboration between computational and experimental research methods, was published in the journal *Biotechnology and Bioengineering*.

Bringing small-scale reactions to the industrial scale

From selective breeding of wheat and corn plants to modern day genetic engineering, humanity has long sought to improve upon naturally occurring biological processes to address its needs. Using bioengineering to increase production of useful materials is another stage in this technological evolution, though it presents some difficult challenges.

"When things get into these large machines, we often observe non-optimal mixing conditions," Takors explained. "Bacteria are living organisms, so in addition to reacting in ways that might be useful in an industrial process, they also initiate processes that are natural but are of no use to the application. This can mean that they underperform at an industrial scale when compared to the ideal conditions in the lab."

In large reactors, bacteria are cultivated in a variety of substrates, or material used to help grow microorganisms. Among other possible challenges, insufficient mixing creates substrate gradients that lead to situations where some reactants are used up too quickly, or reactions occur in certain areas of the reactor but not in others. In other cases, materials don't retain the same mixing characteristics—how soluble a substance is, for instance—when they are brought together in larger quantities. Even the shape of a reactor can influence the efficiency of chemical reactions.

Ultimately, this means that processes can produce less material than intended, make the quality of resulting materials worse than predicted, or in the worst cases, prevent the desired reaction from occurring at all. This loss of efficiency and quality at large scale results in higher expenses and can make it impractical to produce materials commercially, even when they hold the potential to address human needs, or even enable environmental gains by offering cleaner manufacturing processes. Researchers even fear 'the death of innovation' in certain research sectors once lab-scale

SCIENCE FEATURES

expectations cannot be transferred properly to industrial scale after having spent considerable efforts, money, and time.

A feedback loop of simulation and experiment

In order to optimize these industrial processes, investigators in the Institute of Biochemical Engineering at the University of Stuttgart have had to optimize their respective research and development processes. Often, such studies are motivated by national and international industrial partners focused on chemical and biopharmaceutical products.

In their recent paper, Takors and his collaborators use *C. glutamicum* — an industrially important microbe that is used in the manufacture of a variety of common food additives such as the flavor enhancer monosodium glutamate (MSG). Because the bacteria has been so comprehensively studied, it is also a good model system for exploring how computational approaches could help ensure that reactions studied in the lab more reliably function at scale.

The team used data gathered from small-scale laboratory experiments involving *C. glutamicum* as the basis for computational fluid dynamics (CFD) simulations to predict how the bacterium performs in stirred tanks. CFD simulations allow the team to model the mixture of glucose and oxygen in the tank at the molecular level by creating a fine-grained computational mesh, and then solve equations to demonstrate molecule interactions within each box in the grid while advancing in time very slowly.

Based on these simulations, engineers and biologists can then make slight modifications to the bacterium's genetic code, the ratio of ingredients in the mixture, or even the geometry of the bioreactor. By alternating simulations and experiments in an iterative fashion, they gain a better understanding of how the materials will react before testing them in an expensive experiment in a half-million-litre bioreactor.

This process only works if there can actually be efficient iteration, though, and it is here that access to HLRS's HPC systems has played an important role. "One of these simulations takes our team almost a week with our institution's computational resources, but we can do the same simulation in less than 24 hours at HLRS," Takors said. Furthermore, the larger computing power means that the team is able to create more accurate simulations at higher resolution, an approach that minimizes the number of iterations that are necessary to optimize a process. "The way we do the simulation, we create computational meshes, or grids, to solve equations, and the finer the grid, the better the prediction," Takors said.

The team's recent paper demonstrated that its computational approach was capable of simulating large-scale industrial processes at the accuracy necessary to avoid messy trial-and-error work being done inside a large bioreactor. The team was able to simulate the interaction of 120,000 *Corynebacterium glutamicum* cells in realistic industrial mixer conditions. Takors and his collaborators are now exploring other CFD methods and plan to run them on larger amounts of compute cores in parallel to further optimize this approach.

Takors also indicated that his team has been an early adopter of this model of close collaboration and feedback between laboratory scientists, engineers, and computational scientists when it comes to bioeconomy applications. "We are at the forefront of this specific application, and are working to get our approach accepted and adopted by the broader community," he said. *eg*

Uni Stuttgart researchers are using high-performance computing to help model complex biochemical reactions inside of experimental bioreactors in order to make sure these processes will function similarly at the industrial scale. © shutterstock.com

SCIENCE FEATURES

RESEARCHERS FLY HIGH WITH JUWELS AND EXPERIMENTAL AIRCRAFT TO BETTER UNDERSTAND OZONE DEPLETION

Researchers at the University of Wuppertal and Forschungszentrum Jülich combine theory and experiment to study the role of high-altitude clouds on ozone health.

Since being discovered in the 1920s and 1930s, the ozone layer has become synonymous with gauging our planet's health. Protecting the ozone layer has become a common refrain everywhere from elementary school classrooms to industrial board rooms, culminating in the Montreal Protocol and its amendments and adjustments, a sweeping international ozone protection treaty signed in 1987.

This protocol preventing ozone depletion is considered one of the most successful environmental treaties of all time, with broad international adoption and support. However, there are still many unknowns about which natural and artificial processes influence ozone depletion.

Chief among them is the role polar stratospheric clouds (PSCs) play in this complex chemical process. As their name implies, these clouds form in polar regions high up in the Earth's atmosphere, namely the stratosphere, and impact the chemistry of trace gases—small amounts of gases that are naturally part of the Earth's atmosphere apart from the main constituents: oxygen, nitrogen, and argon. While these trace gases account for less than one percent of the gas in the atmosphere, some trace gases play a major role in ozone depletion.

For decades, researchers have refined their experimental toolkits to investigate changes of the ozone layer. Since the 1990s, a research group at the University of Wuppertal has played an active role in designing an instrument to be used by high-altitude planes and satellites. The first, the Cryogenic Infrared Spectrometers and Telescopes for the Atmosphere (CRISTA) was launched on two separate space shuttle flights in 1994 and 1997 and allowed researchers to measure small-scale chemical structures and dynamics within atmospheric trace gases located in the Earth's stratosphere.

After the successful space missions, the Wuppertal researchers partnered with Forschungszentrum Jülich to mount the next-generation CRISTA device, CRISTA-NF, on a high-altitude research aircraft, M-55 Geophysica. While these experiments are capable of measuring trace gas distribution at high altitudes, many questions about the chemical composition of the cloud particles and their role in ozone depletion remain unanswered.

To that end, the team has partnered with researchers in the Climate SimLab at the Jülich Supercomputing Centre (JSC) to use high-performance computing (HPC) to model the impact of the cloud's chemical composition on the CRISTA-NF measurements. The research collective recently published a paper that combines its experimental and computational approaches in *Atmospheric Measurement Techniques.*

Clouds form in simulated space

When studying complex mixtures of gases and small particles in the atmosphere, researchers use multiple instruments and techniques. CRISTA-NF measures electromagnetic radiation emitted by the gases and cloud particles along the "line of sight" and provides the researchers a spectrum in the mid-infrared range with characteristic peaks for each gas and the cloud particles. From the location of the peaks and their shape in the spectra, the scientists can directly tell which major trace gases and particles are present. However, to identify weaker signatures and to extract more properties, such as the trace gas concentration, cloud particle number density, and particle size, advanced simulation and retrieval techniques are required.

"Temperature changes with height, of course, but trace gas concentrations like water vapour also may suddenly change," said Dr. Sabine Grießbach, JSC researcher and meteorologist. "Scientists look at a spectrum and can see, 'this line belongs to CO_2, this line belongs to water vapour, this one belongs to HNO_3,' et cetera. So I may already be able to tell from just looking at the spectrum that there might be this type of cloud or some aerosol, but to know the exact values—the concentration of the gases, the composition of the clouds and their particle sizes—we need to do a retrieval. A retrieval often means inverse modelling and this is what makes this a supercomputing-caliber problem."

The spectra measured by CRISTA-NF in the Arctic at M-55 Geophysica flight altitude and below can only serve as static snapshots of conditions in a certain place at a certain time. These snapshots are not representative for an entire Arctic winter nor the Antarctic winter. Further, even advanced instruments are incapable of accurately measuring particles' sizes, and while the differences may only be micrometres from one another, these minuscule changes can play a large role in how particles impact ozone chemistry.

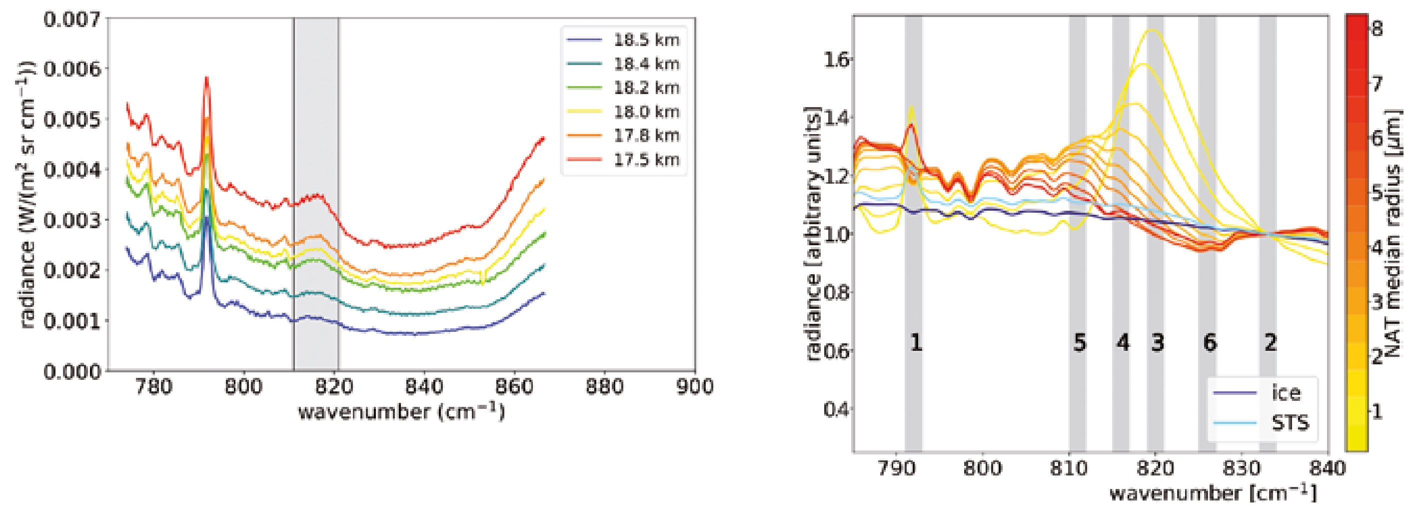

The left shows the spectra of a measured profile, where the NAT-feature is best visible in the measurement at 17.5 km. The right figure shows simulations of the NAT-feature for different particle sizes. For better visualization how the feature changes with particle size, the spectra are scaled to 1 at 833 cm^{-1}. For comparison also spectra for ICE and STS are shown. © Christoph Kalicinsky, University of Wuppertal

When the Wuppertal team noticed an anomalous shape of a certain peak in experimental measurements, it reached out to Grießbach to leverage JSC resources to develop a pipeline between data gathered from CRISTA-NF and simulation.

In this case, inverse modelling means researchers start with an assumption about the state of the atmosphere, then simulate the spectrum and compare it with the measured spectrum. Based on the results of the comparison, the researchers start an iterative process until the difference between simulation and measurement is sufficiently small. However, if clouds are present, scattering on the particles has to be considered in the simulations, which makes the already computationally expensive process even more expensive.

It gets even more complicated, though. There are three "types" of PSC particles: ice, nitric acid trihydrate (NAT), and a chemical phenomenon born of nitric acid and water vapour binding to the sulphuric acid released from natural sources such as volcanoes and certain industrial processes, called supercooled ternary solution (STS). The three types have different scattering properties, which adds a further complexity to the simulations. The researchers decided to perform simulations covering the large variability in terms of particle size and concentration for the three PSC types and mixtures between the three. Based on this comprehensive simulation data set the researchers then developed a computationally cheap method to derive the polar stratospheric cloud composition and the particle size from the many hours of CRISTA-NF measurements. Knowing composition and particle size helps scientists understand how these different PSC types ultimately impact ozone depletion.

In the mid-1970s, scientists identified seemingly inert gases that could be used for everything from aerosol spray cans to Styrofoam production to refrigeration technologies. These gases, made up of only carbon, chlorine, and fluorine, were aptly named chlorofluorocarbons (CFCs). While CFCs were harmless near the surface of the Earth, their interactions with PSC particles in the stratosphere are harmful to the ozone layer. In short, PSC particles provide reaction surfaces for the breakup of CFCs, releasing chlorine gas that actively destroys the ozone layer.

As the measurements exhibited an unexpected shape of the NAT signature, a special focus was put on this particle type. Researchers identified NAT particles as playing a particularly significant role in ozone depletion.

While NAT particles bind nitric acid they can grow to relative large particles sizes. Larger particles sediment, or fall down, more quickly, and remove nitric acid. This impacts ozone chemistry, as nitrate is an important reaction partner for ozone-destroying chlorine by helping convert it to less harmful "reservoir gases".

For the Wuppertal and Jülich researchers, understanding how NAT particles interact with polar chlorine chemistry in the atmosphere is essential to more fully understanding how the ozone layer is depleted and for developing potential remediation strategies.

In fact, in the last 20 years, researchers had identified a peculiar feature in its spectrum data—the data seemed to indicate that the NAT particles had actually "shifted" in the spectrum, something that should not happen when looking at spectroscopy data. The Wuppertal/Jülich team set out to identify why this would have happened, and through its iterative approach, discovered that different sized NAT particles registered differently on the spectrum.

This kind of discovery helps researchers better understand how these complex interactions influence ozone depletion, and ultimately can provide important information in discussions related to geoengineering centred around seeding or injecting the atmosphere with different chemically reactive compounds to cool the planet, among other research areas.

Rising higher

From the experimental side of the research, the next frontier lies in space. The team plans to use its new method to help analyze a decade worth of global experimental data gathered from instruments mounted on satellites, such as Envisat MIPAS.

In order to meet this additional challenge, the team will also be moving into the next computational frontier—JSC's JUWELS Booster module, set to come online in late 2020. JSC is spearheading a novel course for HPC technology, developing a modular supercomputer over the coming decade that will consist of ultra-high-speed connections between modules consisting of different architectures. The initial JUWELS module is capable of roughly 12 petaflops, or 12 quadrillion calculations per second. The second module, the so-called "Booster," will use GPU power to raise the total system performance to 85 petaflops (for more information on the JUWELS Booster, please visit Page 4).

The team used the Jülich-developed JURASSIC code to run its simulations. As part of JSC's Climate Science SimLab and co-developer of JURASSIC, Grießbach will play an active role in ensuring that JURASSIC is ported efficiently to make the best use possible of JUWELS' upcoming enhanced capabilities. "For us, in the simplest sense, we in the SimLabs are supposed to help other users from the community who, of course, use a large variety of different codes," Grießbach said. "JURASSIC is our own code, one that we know very well. While preparing the code and setup for the next scientific study we learn how to make efficient use of the new architecture—which part of a code works well on the GPU partition, and which part of the code should run on the CPU partition?"

eg

HLRS WILL HELP BUILD NATIONAL DATA RESEARCH INFRASTRUCTURE FOR CATALYSIS RESEARCH

By developing a repository for sharing research data and standards for data management, HLRS will create a basis for interdisciplinary research in chemistry and chemical engineering.

A consortium including the High-Performance Computing Center of the University of Stuttgart (HLRS) has been awarded a grant of roughly 10 million EUR from the Deutsche Forschungsgemeinschaft (DFG) to create a National Research Data Infrastructure for Catalysis-Related Sciences (NFDI4Cat). The consortium, led by the nonprofit chemical society DECHEMA (Gesellschaft für Chemische Technik und Biotechnologie e.V.) and involving representatives from 15 additional partner institutions, will develop infrastructure, software, and data management standards to empower the next generation of chemical engineering research. NDFI4Cat is one of nine new consortia that will contribute to the construction of a German National Research Data Infrastructure.

As one of four members of the NFDI4Cat coordination group, HLRS will create and host a data repository for catalysis-related research, including a portal for sharing and accessing data stored at multiple locations. In addition, HLRS will play a major role an effort to establish standardized metadata and ontologies for catalysis research that will ensure compatibility among different data sets, increasing their usability and amplifying their potential impact for scientific progress.

"We are very pleased that HLRS will be participating in the development of a National Research Data Infrastructure," said HLRS Director Prof. Michael Resch. "Working together with partners in the catalysis research community, this project should offer outstanding opportunities to accelerate research in a field that is not only of great economic importance, but that also holds keys to addressing some of our greatest global challenges."

Transforming catalysis into a computational science

Catalysis and chemical engineering are essential disciplines for producing many materials we use in our daily lives. They also offer the potential to address some of humanity's most pressing problems. Developing new technologies for reducing CO_2 emissions, reducing plastic waste, or producing sustainable fertilizers to meet the nutritional needs of a growing global population, for example, are all areas in which catalysis and chemical engineering have important roles to play.

NFDI4Cat

Catalysis research, like many scientific fields, increasingly relies on computational methods that support a continual dialogue between theory, simulation, and experiment. Based on rapidly growing collections of high-throughput experimental data, data science methods can be used to predict relationships between the chemical structures of catalysts and their activities. At the same time, simulation can provide valuable guidance for optimizing the design of chemical reactors and processes. As data accumulate, the ability to integrate knowledge from different disciplines, looking at all levels of catalytic reactions — from the physical chemistry of individual molecules all the way up to process engineering — would also provide fundamental knowledge that would be of great use to researchers across the field.

To exploit the full potential of this trend, however, new kinds of computing infrastructure and methods are needed. Although the usage of data science in catalysis research has been growing, scientists too often work in relative isolation from one another. The result is that data is typically collected in proprietary formats, is not organized using standardized metadata description, is not saved in places where it is accessible to other researchers, and is not linked to related publications and published data sets. The NFDI-4Cat project intends to address these problems by creating a shared, comprehensive framework for pooling and managing catalysis research data.

Establishing catalysis metadata standards

Between 2017 and 2019, HLRS contributed significantly to a research project called DIPL-ING, which developed a system for research data management in computational engineering. The result was a metadata model called EngMeta,

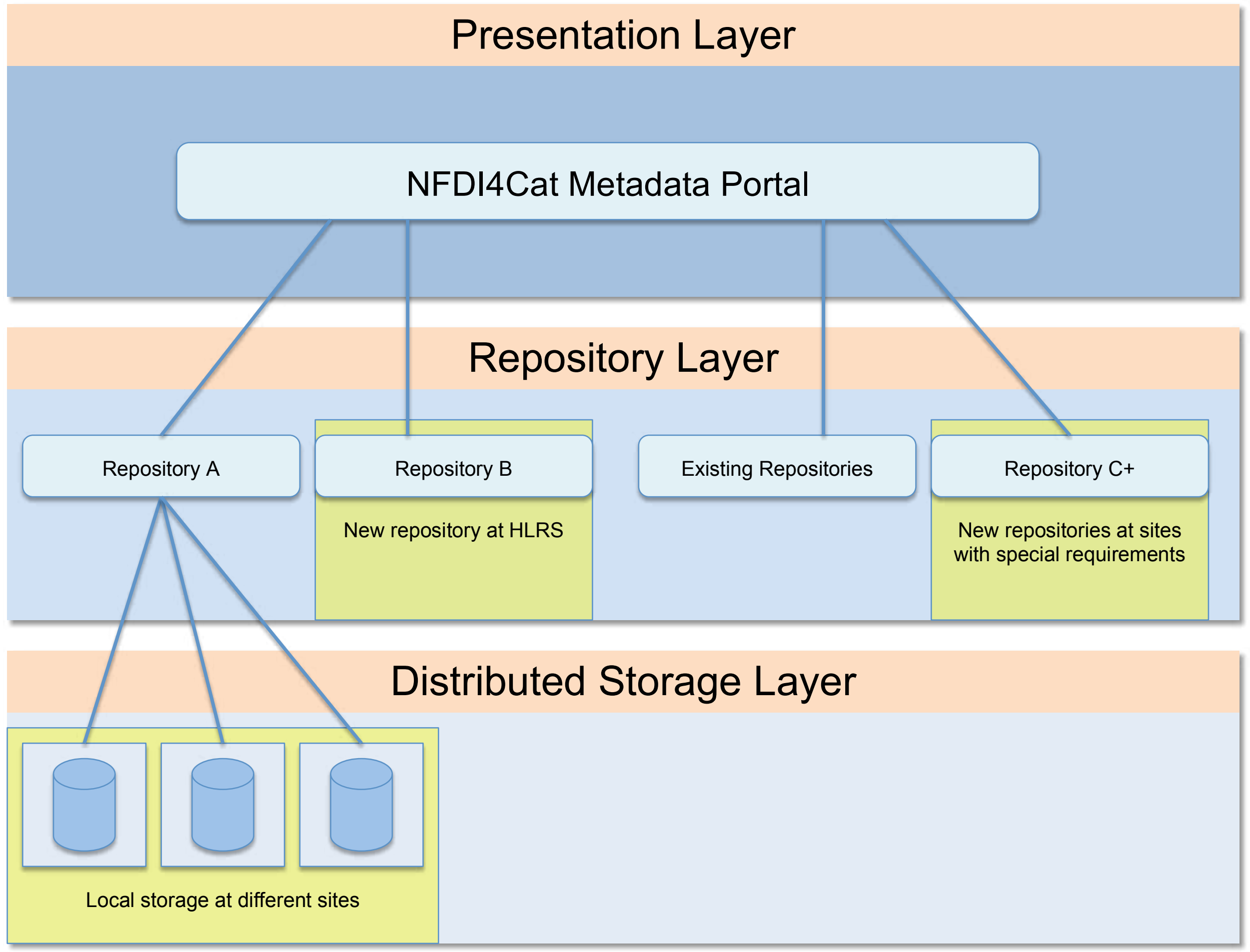

NFDI4Cat is planning a distributed data infrastructure using a shared metadata concept that makes it possible to access data repositories from a central portal. **© NFDI4CAT**

which HLRS and the University of Stuttgart library now use. In addition, researchers developed a method called ExtractIng that could automatically extract metadata from research data sets and transform it to EngMeta. This approach relieves investigators of often tedious and time-consuming work, a goal that will be important to the success of the NFDI4Cat project.

In this new project, HLRS will build on EngMeta to create an ontology — a set of categories for organizing the knowledge contained in all relevant datasets — for the management of data in catalysis research. Such an ontology would include metadata that generally describe the data, technical metadata about the data objects contained in the dataset, process metadata describing the methods and experimental or computational hardware used to generate the data, and domain-specific data related to the specific field of research in which the data were generated.

Researchers in the NFDI4Cat consortium envision integrating the ontology in two other complementary software platforms: Piveau, an open source data management ecosystem developed at FOKUS (Fraunhofer Institute for Open Communication Systems), and CaRMen, software

developed at the Karlsruhe Institute of Technology for analyzing physical and chemical models against experimental data.

Ultimately, by standardizing metadata frameworks, this project aims to bring catalyst research data in line with the so-called FAIR principles for scientific data management (findability, accessibility, interoperability, and reusability). NFDI4Cat will also implement stringent quality assurance methods, including prompts to follow best practices during data ingestion, to ensure that all data that enter the repository are of high quality and labeled correctly.

HLRS to develop and host NFDI4Cat data repository

In addition to developing the conceptual framework for organizing catalysis research data, HLRS will also play a central role in the development of the infrastructure for hosting and sharing data.

NFDI4Cat will be based on a distributed repository infrastructure. Data will be stored across a variety of servers, making it possible for institutions that want to share data with the community to participate, even if policies prevent them from storing it on an external server. A middleware layer at the core of the service will connect the various repositories, and a graphical user interface will provide a portal for users to access data across the network.

To ensure that the catalysis community adopts the repository and its user interface, HLRS, together with FOKUS, will hold meetings with representatives from academia and industry to gather information about their needs. Discussions will explore questions related to the state of the art in data management hardware and software technologies, how to make ingestion of data and metadata into the system as easy as possible, and legal requirements for protecting intellectual property and data access. Once the project requirements are clarified, HLRS will work with FOKUS to develop the hardware and software backbone of the catalysis data repository.

In addition, HLRS will provide two server systems for the repository, as well as hard disk storage of approximately 100 TB and up to 1 PB of background storage on tape.

Networking to leverage the power of data

In addition to maintaining close dialogue with the catalysis research community, the NFDI4Cat will also engage with other NFDI centre on issues of shared interest to ensure that data is useful in other fields. This includes centres focused on developing national research data infrastructures for chemistry (NFDI4Chem), engineering (NFDI4Ing), and photon and neutron experiments (DAPHNE), which share overlapping research concerns. NFDI4Cat will also coordinate with other relevant research institutes, programs, and data resources outside the NFDI network to identify potential synergies.

Developing a more robust infrastructure for sharing and organizing data offers the near-term possibility of transforming catalysis research. By moving toward a more open data structure, scientists could gain new computational power for understanding reaction mechanisms and kinetics, pursue more efficient, rational approaches to catalyst design, and gain new kinds of insights that arise through interdisciplinary research. In this way, NFDI4Cat should itself serve as a catalyst for new kinds of discoveries. *cw*

NFDI4Cat

Runtime:
October 2020 to September 2025

Total Project Grant Amount:
10,150,000 Euro

Funding Source:
German federal government and state governments, based on "Bund-Länder-Vereinbarung zu Aufbau und Förderung einer Nationalen Forschungsdateninfrastruktur"

Project Partners:
A consortium led by DECHEMA (Gesellschaft für Chemische Technik und Biotechnologie e.V.) of 16 institutions and 3 associated partners

FULL STEAM AHEAD TO THE EXASCALE SEAS

DEEP-SEA, IO-SEA and RED-SEA selected for funding by the EuroHPC Joint Undertaking

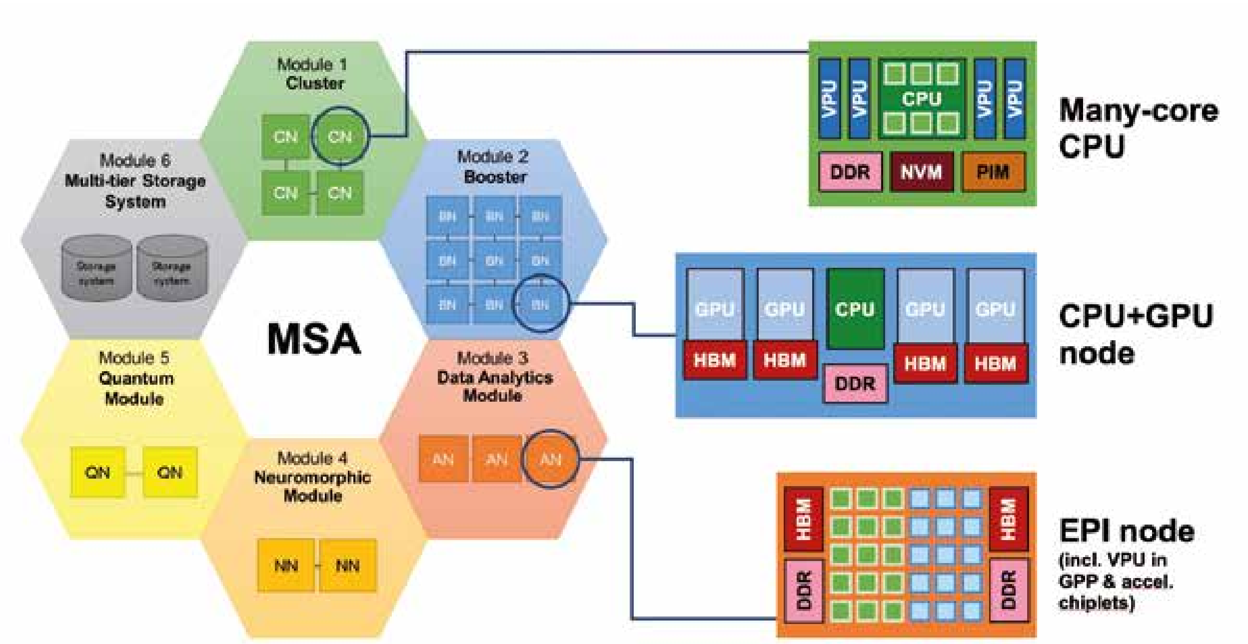

The Modular Supercomputing Architecture is the basis of the DEEP-SEA, IO-SEA, and RED-SEA. It orchestrates heterogeneous compute and memory resources.

After a very positive evaluation, the three members of the new SEA projects family (DEEP-SEA, IO-SEA, and RED-SEA) have been selected for funding in the first call for proposals issued in 2019 by the EuroHPC Joint Undertaking (JU). The SEA projects complement each other and address together some of the most critical challenges for the use and operation of an exascale supercomputer: namely, the programming environment, data management, and interconnects.

DEEP-SEA (DEEP – Software for Exascale Architectures) builds upon the results of the DEEP projects family (www.deep-projects.eu) to develop software for the upcoming European exascale supercomputers, which are expected to combine very diverse compute and memory resources both at node and system level. The goal of DEEP-SEA is to adapt all levels of the software stack – including low-level drivers, computation and communication libraries, resource management, and programming abstractions with associated runtime systems and tools – to efficiently exploit highly heterogeneous compute and memory configurations, and to allow code optimisation across existing and future architectures and systems. The DEEP-SEA software stack will enable dynamic resource allocation, application portability and malleability, programming model composability, and include tools to map applications to Modular Supercomputing Architecture (MSA) like on the JUWELS system at the Jülich Supercomputing Centre (JSC). The result is a software environment enabling applications to run on the best suited hardware in a scalable and energy efficient manner.

IO-SEA (Input/Output Software for Exascale Architectures) aims to provide a novel data management and storage platform for exascale computing based on hierarchical storage management and on-demand provisioning of storage services. The goal of IO-SEA is the efficient management of storage tiers spanning non-volatile memory technologies at the top all the way down to least active data stored with tape-based technologies, removing unnecessary data movements, and considering critical aspects of intelligent data placement for extreme volumes of data. The concept of ephemeral data nodes and data accessors is introduced, which allows using well-known data access paradigms (e.g. POSIX, MPI-IO, etc.) to treat storage resources as dynamic, changing system components, as will be required in exascale-class, highly heterogeneous computing environments.

RED-SEA (Network Solutions for Exascale Architectures) will pave the way to the next generation of European exascale interconnects. In particular, it will develop the next generation BXI (Bull eXascale Interconnect), scalable to hundreds of thousands of nodes. New features will be implemented, including BXI-Ethernet gateways to enable seamless communication within and between compute and storage clusters, and efficient network resource management covering congestion resiliency, virtualization, adaptive routing, and collective operations, among other networking considerations. With these extensions, BXI shall aid high-performance computing (HPC), high-performance data analytics, and artificial intelligence applications benefiting from new hardware and software trends to reach exascale performance.

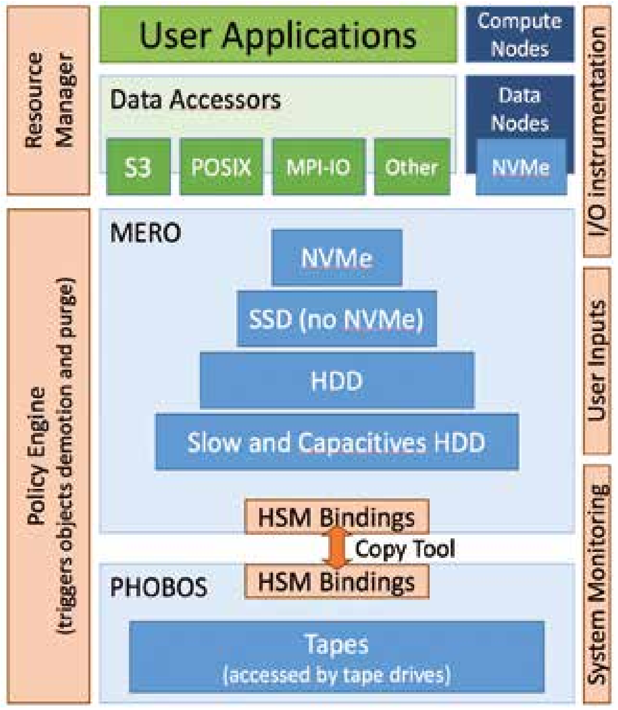

IO-SEA structure, components and relationship.

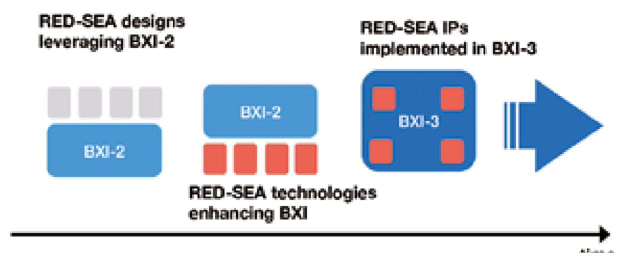

Interconnect development roadmap in RED-SEA.

High technology maturity and co-design

All the three SEA projects target a high technology readiness level by building upon developments from previous European Union-funded projects. With regard to software, these European developments are complemented by international open source packages widely used in the HPC community, extending and adapting both to better deal with challenges related to compute, memory, and storage heterogeneity.

The SEA projects are driven by a collaborative co-design approach with a selected set of high-impact applications. These representative applications help formulate requirements that are addressed by the software and hardware developers, and are also employed to evaluate the project's results and demonstrate their benefits for users of European computing centres. Participation of the wider European user communities into the evaluation process is enabled by early-access programs, thorough which close collaborations within the HPC community and Centres of Excellence (CoEs) is foreseen.

GCS contributions

GCS plays a leading role in the DEEP-SEA project, with JSC as its coordinator and the Leibniz Supercomputing Centre as the leader of dissemination and communication activities while also bringing its expertise in energy efficiency, and power monitoring and management to the project. JSC participates in all three SEA projects. In addition to leading the DEEP-SEA project, JSC contributes the performance analysis tools Score-P and Scalasca, the monitoring tool LLview, the I/O library SIONlib, the benchmarking tools Linktest and JUBE. The centre also provides application use-cases in the areas of Earth system modelling and quantum computing. Additionally, JSC leads the co-design activities in DEEP-SEA and IO-SEA and the overall definition of the software stack in DEEP-SEA.

es

DEEP-SEA
IO-SEA
RED-SEA

Runtime (Funding Period):
2021-2023

Funding Source:
EuroHPC JU Program

Funding:
15 million Euro / 8 million Euro / 8 million Euro

Partner institutions:
13 partners / 10 partners / 11 partners

ORIGINS – ANSWERS TO EXISTENTIAL QUESTIONS

LRZ is providing access to supercomputing resources, storage, and visualization capabilities for the ORIGINS project.

The ORIGINS cluster of excellence serves as a multidisciplinary research initiative to investigate how life develops. **© ORIGINS**

Over 182 days, the eROSITA X-ray telescope orbited the Earth on board the Russian-German "Spectrum-Roentgen-Gamma" (SRG) mission at a distance of 1.5 million kilometers. It delivered around 165 gigabytes of data and millions of images, mainly of black holes in galaxies, but also of hot gas in galaxy clusters in and around our galaxy.

Galaxies act as islands of the cosmos, serving as places where life originates. "Galaxies have a central black hole that permanently attracts matter and thus grows," explains Dr. Klaus Dolag from the University Observatory Munich at the Ludwig-Maximilians-Universität (LMU). "In this process, enormous amounts of energy are released, which in turn influence the development of a galaxy and thus the conditions for the formation of life."

Multidisciplinary collaboration accelerated by supercomputing

Data from eROSITA and the Max Planck Institute for Extraterrestrial Physics (MPE) are highly welcome in the ORIGINS cluster of excellence: Since 2018, the cluster of excellence project has brought together 120 research groups from the fields of astrophysics, particle physics, and biophysics to investigate how life develops. The two Munich universities—the Technical University of Munich and LMU— as well as several Max Planck Institutes and the European Southern Observatory (ESO), together with international research partners, are involved in ORIGINS. The Leibniz Supercomputing Centre (LRZ) provides researchers with supercomputing resources and storage capacity, LRZ's Centre for Virtual Reality and Visualisation (V2C) supports the cluster's Virtual Reality Lab and the ORIGINS Data Science Lab with professional experience in handling data and algorithms. "We need computing time to understand the complex relationships in the universe from the Big Bang to the origins of life. Numerical simulations are the modern workhorses of astro-, particle- and biophysics," explains Dr. Andreas Burkert, Professor at LMU and head of the Chair of Theoretical and Numerical Astrophysics. Burkert serves as one of the coordinators of ORIGINS. "LRZ guarantees access to one of the world's most powerful supercomputers. In addition, we can use the wide range of IT services offered by the LRZ."

A part of LRZ's high-performance computing resources are reserved for the ORIGINS cluster so that models and algorithms can be developed from measurement data. The V2C helps researchers visualize results and calculations by helping to create images, videos, or spatial animations. This collaboration began in 2006, when UNIVERSE, the first Cluster of Excellence project in this area and precursor to ORIGINS, was launched to trace the origins of the universe. By 2018, around 250 researchers and 70 working groups from all over the world had published 4,000 papers and 308 doctoral theses. At the end of the project, they were able to explain supernova explosions, report on new galaxies, explore stellar nucleosyntheses and hypothesize the inner workings of black holes, and discovered the tiny Higgs particle and high-energy cosmic neutrinos. Part of the UNIVERSE results can be seen in the Deutsches Museum, where the film "Ausgerechnet" (Of all things), which was rendered on SuperMUC and vividly reports what happened after the Big Bang some 14 billion years ago, is also shown.

How is the primordial soup made?

As ORIGINS continues to evolve, the project is involving new researchers for a broader view. Recently, the cluster has incorporated biochemists in the collaboration. "Is life an inevitable physical and chemical process that can take place wherever the conditions are right?" Burkert asks. "To find an answer to this question, we now also need biophysics and chemistry." Bringing additional disciplines into the project's collaboration helps brings new ideas, models, and thoughts.

The first year of the project served as an exciting discovery process, Burkert explains "Biologists have different chains of thought than physicists. Before we could work together, the cluster had to make sure that everyone understood the language of the other scientists." As an astrophysicist searching for dark matter signatures in space, Burkert learned, for example, that particle physics already has methods for breaking down these signatures more precisely.

In the search for the origin of life, ORIGINS therefore attaches great importance to interdisciplinary research, and particle physics and biophysics now complements the knowledge ORIGINS researcher have already gained. Through the project, ORIGINS researchers found that phosphorus is formed in stars and is hurled by their explosions into interstellar dust and gas clouds. These clouds serve as the fertile ground in which planetary systems and comets form. "Where phosphorus originates in the stars is something we have not really thought about in astrophysics," Burkert admits.

Generally speaking, interstellar clouds are interesting for researchers in the ORIGINS project, as scientists suspect they contain many more building blocks of life. In the first year of ORIGINS, more than 100 research contributions have already been made not only about stars, black holes or galaxies—in addition to the exploration of the universe, ORIGINS, which will be funded by the German Research Foundation with a total of 45 million Euros until 2026, wants to discover the recipe of the primordial soup, that unknown mixture of inorganic substances which made the emergence of life possible. If its composition can be reproduced in the laboratory, it will be easier to explain how unicellular organisms and the first living beings developed.

What does data report about life?

From what large-scale processes happen across vast expanses of space, to the microscopically small compounds of life, to the tiny atomic building blocks of matter, ORIGINS integrates a broad spectrum of research, and every scientific question produces data via measurements, photographs from telescopes and microscopes or model calculations. Astrophysics and particle physics have already advanced into 3D modeling and into spatial representation with the help of the enormous computing power of supercomputers. "Many of our research areas work with big data and have to evaluate and understand large, multi-dimensional data," says Burkert. In the cluster, the researchers are also working with LRZ on a Data Science Lab for new techniques and methods to collect, analyse and understand data: The 165 gigabytes of information that eROSITA recently sent to Earth are not only interesting in terms of content, but also as a means of testing processing and validation tools. *sv*

ORIGINS

Runtime:
2018-2026

Funding Source:
Deutsche Forschungsgemeinschaft
(EXC 2095-390783311)
plus diverse seed money projects from partners

Funding:
45 million Euro

Partner institutions:
Ludwig-Maximilians-Universität (GER), Technische Universität München (GER), Max-Planck-Institute (GER), Leibniz Supercomputing Centre (GER), Institute des Nanosciences (FR), Laboratoire Astroparticulaire et Cosmologie (FR), Centre for Study Origins and Evolution of Life, Planets and the Universe (NL), European Southern Observatory (EUR), Harvard University (USA), Center for Computational Astrophysics (USA)

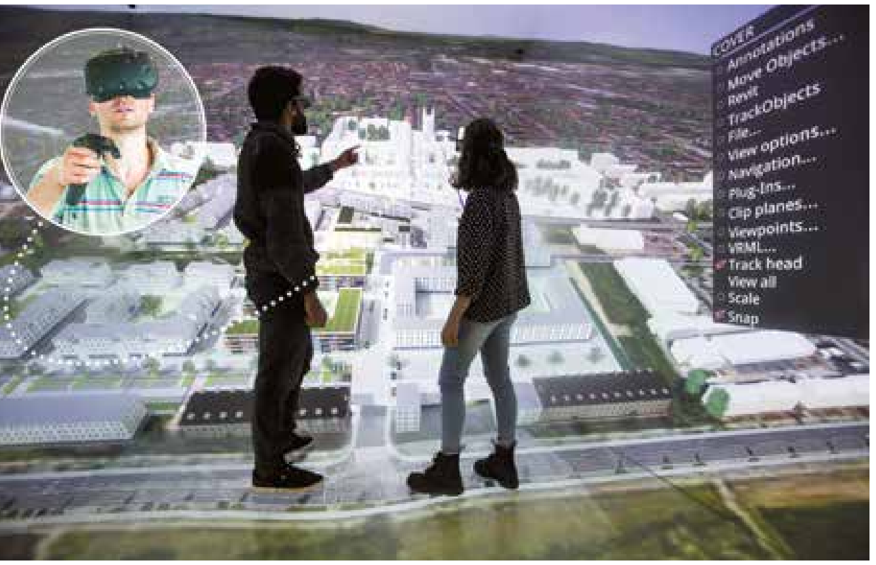

COLLABORATIVE VIRTUAL REALITY ENVIRONMENTS FOR THE HOME OFFICE

Visualization experts from the KoLab project ("Virtual Collaboration Laboratories Baden-Württemberg"), including members of HLRS, have developed a new software platform that enables persons located in different places to meet and collaborate in virtual reality. When combined with inexpensive, commercially available head-mounted displays and controllers, scientists, researchers, developers, and users can connect online and gather in a virtual meeting room, where they are represented by avatars. In virtual reality they can then together observe, analyze, and interact with 3D visualizations—all from their workplaces or even home offices. When spontaneous in-person meetings are difficult—for example, when colleagues work at different locations or cannot meet due to the COVID-19 pandemic—the technology can facilitate cooperation, simplify complicated work processes, and support education and training.

BIG DATA GOES ON A JOURNEY

The EU-funded project LEXIS aims at fostering efficient execution of multipart simulations ("workflows") on distributed cloud and high-performance computing systems. Researchers in both academia and industry are developing programs and infrastructure for the project, and the Research Data Management team at LRZ focuses on the data exchange that makes distributed simulations possible. Through LEXIS, users will find a range of innovative computing and data systems. Lexis integrates offers from the European Open Science Cloud (EOSC) as well as open-source-based solutions such as B2HANDLE, B2SAFE and B2STAGE from EUDAT ("European Data") for data management via specially developed interfaces. Initial tests show that big data can go on the road with these solutions and that location-independent access is possible. A forest fire forecast simulation based on current weather data served as the first test object. The simulations were done in Liguria by the Italy-based CIMA Foundation and are used by the local Environmental Protection Agency, ARPAL. The workflow started with a virtual machine in the LRZ compute cloud, on which weather data was collected and prepared for transfer and simulation. Afterward, the IT4I systems accessed it, calculated small-scale weather phenomena and forest fire hazards, and stored this data back on the LRZ system. There the results can now be downloaded and processed or visualized. Tools for data management ensure that the use of the research data is documented, their integrity is checked and the current work results are safely stored on both sides. For more information, visit: https://lexis-project.eu/web/

STRONG TEAMS AT THE HELMHOLTZ GPU HACKATHON

The computing power provided through GPU-accelerated supercomputers throughout Europe is increasing. This autumn, the JUWELS Booster at JSC added another 70 petaflops of GPU-heavy computing power. In order to develop and port applications for such systems, HZDR, JSC, and NVIDIA have jointly organised GPU hackathons for the last several years. During these events, the accepted teams work together for several days under the guidance of experienced mentors on porting and optimising their codes. Nine teams attended the most recent Helmholtz GPU Hackathon in September. They covered a broad range of science, including simulations in chemistry, hydraulic modelling, theoretical particle physics, and battery research. Other teams applied machine learning methods in the area of experimental particle physics and atomistic modelling. Due to the COVID-19 pandemic, the Helmholtz GPU Hackathon had to be a virtual event. Thanks to the experience collected during earlier hackathons around the world, this did not result in teams being less productive. The series of GPU Hackathons will be continued and application developers are encouraged to watch https://www.gpuhackathons.org/, where future events will be announced.

BAVARIAN MINISTER FOR SCIENCE AND THE ARTS SIBLER AT THE LRZ

In July, the Bavarian Minister for Science and the Arts, Bernd Sibler, visited the Leibniz Supercomputing Centre (LRZ) together with representatives of his ministry. The group discussed future technologies and quantum computing with scientists from Ludwig-Maximilians-Universität (LMU), the Technical University of Munich (TUM) and the Walther-Meißner-Institute. "Visit to the Leibniz Supercomputing Centre, our IT cathedral," Sibler noted in his online diary on Facebook. "Great debate on the quantum strategy of the Free State of Bavaria. Great perspectives for Bavaria." Among those participating in the discussion were (front row, left to right) Dieter Kranzlmüller, Director of the LRZ and Professor of Computer Science at the LMU, Bernd Sibler, Bavarian Minister of Science and Arts, Professor Immanuel Bloch, Quantum Optics (LMU) and Max Planck Institute for Quantum Physics (MPQ). Also present were (top row, left to right) Professor Ignacio Cirac, Theoretical Quantum Optics TUM and MPQ , Ministerialdirigent Michael Greiner, Professor Rudolf Gross, TUM and Walther-Meißner-Institute, Ministerialdirigent Dr. Johannes Eberle.

NIC EXCELLENCE PROJECT 2020

Nuclear lattice effective field theory is a new computational tool to perform *ab initio* calculations in nuclear physics. Because of recent algorithmic developments, the phase diagram of hot and dense nuclear matter can now be explored in this framework. In fact, a research team from the University of Bonn and other partners in both Germany and the United States was named as the 2020 John von Neumann Institute's Excellence Project for using JSC HPC resources to study nuclear systems with a fixed number of nucleons at non-zero temperature, which allows researchers to map out the phase diagram and the location of the critical point. The team's results serve as a major step forward in the understanding of the equation of state of nucleonic matter, which plays an important role in the generation of gravitational waves in neutron star mergers, among other research focuses.

SUPERCOMPUTING-AKADEMIE AWARDS FIRST CERTIFICATES

This summer, the Supercomputing-Akademie for the first time issued its highest certificate—"HPC Expert"—to two course participants. The recipients successfully completed five training modules focused on essential skills for HPC developers. Launched in 2018, the Supercomputing-Akademie has been building a curriculum that enables high-performance computing users, programmers, and IT administrators in industry to gain essential competencies in HPC. Using a blended learning format, the Supercomputing-Akademie is designed to make it easy for specialists in industry to develop expertise in HPC in parallel with their normal work. In recent months, the Supercomputing-Akademie also launched two new course modules—"Visualization" and "Data: Management and Analysis"—adding to its portfolio of modules focusing on simulation, performance optimization, parallel programming, and other topics.

NEW AI PROJECT ALPHANUMERICS ZERO

In the first annual project call for the Helmholtz Artificial Intelligence Cooperation Unit (Helmholtz AI), the proposal "AlphaNumerics Zero" (aN0), led by JSC together with researchers from the Steinbuch Centre for Computing at the Karlsruhe Institute of Technology, has been selected for funding. The objective of this project is to rethink numerical methods on high-performance computers. Until now, a lot of effort has gone into the design, implementation and performance optimization of solvers for differential equations and numerical libraries. New ideas are urgently needed, as developing methods for upcoming extreme-scale supercomputers will become even more challenging. The goal of AlphaNumerics Zero is to use reinforcement learning techniques so that a computer learns the on-average optimal numerical solution method by itself for a given simulation. This is a moonshot project that—if successful—will be the first step to change the paradigm of how numerical simulations are designed and performed on extreme-scale computers.

MEDIA SOLUTION CENTER BADEN-WÜRTTEMBERG WELCOMES TWO NEW MEMBER ORGANIZATIONS

The Media Solution Center Baden-Württemberg (MSC) has welcomed Barcelona's Fundación Épica La Fura dels Baus and the Stuttgart Chamber Orchestra as new members. They join a growing network of arts organizations and media companies that the MSC is building to promote innovation at the intersection of the arts, science, and technology, particularly involving supercomputing, simulation, artificial intelligence, and virtual reality. A non-profit membership organization cofounded with the ZKM (Karlsruhe) and the Hochschule der Medien (Stuttgart) and based at HLRS, the Media Solution Center supports the creative exchange of ideas, expertise, and resources among artists, cultural organizations, media companies, scientists, and technology developers. Through its close relationship with HLRS, the MSC also facilitates access to advanced computing and visualization technologies for artists. The new memberships promise to lead to new kinds of digitally enabled theatrical and musical experiences.

INTERNATIONAL PROJECT WILL CREATE PAN-EUROPEAN DATA INFRASTRUCTURE FOR SARS-COV-2 RESEARCH

HLRS, working as a member of an intercontinental consortium of 27 centres for public health and high-performance computing, will help create a large-scale data repository that will be used for research on the COVID-19 pandemic. The project, called ORCHESTRA, will develop a framework for integrating data currently segregated among many medical centres to create a pan-European cohort of more than 200,000 individuals. With its partners HLRS will help develop standards for data collection and harmonization, as well as a cloud-based portal for researchers to access, share, and integrate large data sets that will be useful for computational research in genetics, epigenetics, immunology, epidemiology, and other disciplines. The project partners intend to provide a resource that will help governments and the medical community to perform more robust studies, better track the development of the current pandemic, and test the effectiveness of vaccines against the SARS-CoV-2 virus. The project will also serve as a model for responding to potential future outbreaks of dangerous infectious diseases.

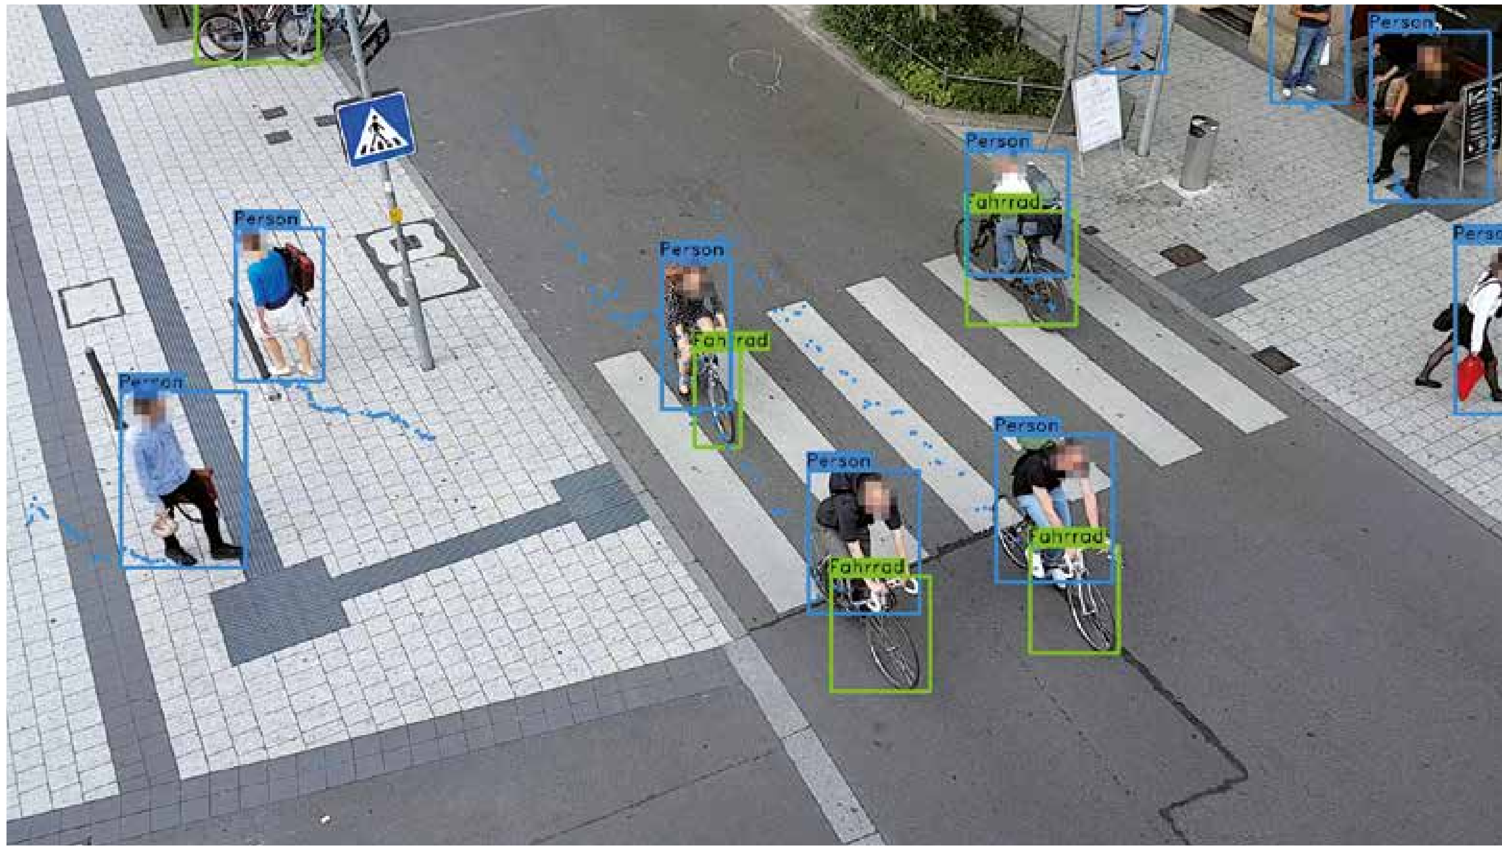

LRZ ISO/IEC CERTIFICATION RENEWED

The auditing company DEKRA recently reviewed the IT security and service management processes in place at LRZ for a second time. The organization was re-certified according to the ISO/IEC standards 20000 and 27001. The DEKRA audits evaluate the effectiveness, development and continuous improvement of internal processes. The auditors were extremely satisfied with the implementation of their recommendations and were surprised that the LRZ had already expanded and standardized considerably more processes than required over the last year.

SAFER CITIES FOR CYCLISTS AND PEDESTRIANS

The Bundesministerium für Verkehr und digitale Infrastruktur (German Federal Ministry of Transport and Digital Infrastructure) has awarded a grant to a team led by HLRS to investigate how computer simulation could help minimize an increasingly common problem in many cities — conflicts that arise between cyclists and pedestrians. The project, called Cape Reviso (Cyclist and Pedestrians on Real and Virtual Shared Roads), will use machine learning, virtual and augmented reality, digital twins, and urban emotion tracking to develop tools for urban planning. The project will conduct a case study in Stuttgart, focusing on a complex location where multiple forms of traffic intersect. The researchers will collect data to create a "digital twin" of the location in the HLRS virtual reality facility, where planners can both study behavior at the location and simulate the effects of interventions to improve traffic flow. These scenarios could then be tested in the "Living Lab" at the physical location, providing an opportunity to compare real-world behaviors and the simulation.

EURONEWS VIDEO ON COMPBIOMED

EuroNews visited University College London and the Leibniz Supercomputing Centre (LRZ) in July for a feature on the EU-funded CompBioMed project. Presenter Katharina Rabillon of Futuris explained the meaning and benefits of digital twins in medicine. The algorithms and software developed at CompBioMed process data from computer and magnetic resonance tomographs (CT, MRT) and X-ray scans as well as data from organs and genomes and more. These models can be used to simulate virtual humans in order to individually adjust treatments and medications, but also to predict health risks like cancer. This major advance in personalized medicine is powered by consistent access to HPC resources. The video is available in the EuroNews media library in both German and English:
https://www.euronews.com/2020/08/24/digital-doppelgangers-to-map-the-future-of-major-surgery

A SCALABLE APPROACH FOR THE ANALYSIS OF HUGE TIME SERIES

Amir Raoofy, PhD candidate at the Chair for Computer Architecture and Parallel Systems at the Technical University of Munich (TUM) headed by Professor Martin Schulz, works with data supplied by thousands of sensors from supercomputers or from monitoring systems installed at power plants. The data has been generated over weeks or even years. Raoofy is interested in how SuperMUC-NG and CoolMUC-3 handle these huge amounts of data. “Using matrix profile algorithms, time series can be searched for patterns and similarities,” says Raoofy, outlining the problem. “But they are difficult to scale and are not suitable for HPC systems.” However, the evaluation of large time series requires supercomputing power: Anyone who wants to know under which conditions a gas turbine will run reliably and when the first components will be in need of repair should be able to check vast amounts of data. The compute power and capabilities provided by supercomputers make such analyses possible only in combination with appropriate and scalable algorithms.
Raoofy and his colleagues have developed the now award-winning scalable approach called (MP)^N. The team’s model can currently be run efficiently on up to 256,000 computer cores - that is around 86 percent of the total computing power of LRZ’s SuperMUC-NG system. The fact that it delivers exact calculations was tested with performance data from SuperMUC-NG. The algorithm is currently being used to analyze data supplied by two gas turbines belonging to Stadtwerke München within the context of the TurbO research project. “In our experiments, we performed the fastest and largest multidimensional matrix profile ever calculated,” reports Raoofy. “We achieved 1.3 petaflops” This means that supercomputers like SuperMUC-NG can quickly and efficiently evaluate data from long-term measurements - science and technology will know how to use this. The publication „Time Series Mining at Petascale Performance“ is open access and can be downloaded from https://link.springer.com/chapter/10.1007/978-3-030-50743-5_6 .

SMARTPHONES MIRROR THE PERSONALITY

A team led by Prof. Markus Bühler at Ludwig-Maximilians-Universität München analysed the smartphone use of 624 students using machine learning on high-performance computers. Their publication in *Proceedings of the National Academy of Sciences* shows that it is possible to draw conclusions about the personality of the users, at least for most of the important personality dimensions. Using HPC resources at LRZ, the team studied six major factors of cell phone usage that can be used to evaluate and predicate personality characteristics—communication and social behavior, music consumption, daytime vs. nighttime use, app usage, mobility, and overall phone activity. For more information, visit: https://doi.org/10.1073/pnas.1920484117

GLACIERS IN HPC

Carlo Licciulli, glaciologist and research associate of the Bavarian Academy of Sciences and Humanities, has just been awarded the Graham Cogley Award of the International Association for Cryospheric Sciences and the International Glaciological Society for the best publication by a young scientist. The publication in the *Journal of Glaciology* is based on his dissertation and was selected from 70 publications. Since 2018, Carlo Licciulli has been using measured data to calculate and model the behavior of the Vernagtferner glacier in the Ötztal Alps on the basis of measured data in order to better explain the retreat of the ice. While still using a PC, he adapted a numerical full-Stokes ice flow model, used it to calculate the ice transport in the glacier and model the three-dimensional distribution of the ice age and used HPC resources available at the Leibniz Supercomputing Centre (LRZ). His research can be found here:
https://doi.org/10.1017/jog.2019.82

NEXT PHASE OF HUMAN BRAIN PROJECT LAUNCHED

The Human Brain Project (HBP) has entered into its next and final phase, which will last another three years until March 2023. Launched in 2013 as one of the first two EU-funded “FET Flagship” projects, the central goal of the HBP in this last phase will be to consolidate and further integrate the developed platforms into a coherent and sustainable research infrastructure for neuroscience and brain-inspired technology development. This digital research infrastructure, called EBRAINS, will provide researchers with tools and services for storing, processing, analyzing and sharing data; for navigating the brain in 3D brain atlases; for running closed-loop AI and robotics workflows; and for creating and simulating models of the brain. JSC leads the EBRAINS Computing Services work package, which operates and integrates the HPC, cloud computing, and storage services of the Fenix Infrastructure (including JSC’s new JUSUF system) with novel neuromorphic computing services. This joint infrastructure layer forms the basis of EBRAINS and enables platforms such as the Neurorobotics Platform and individual solutions to integrate different EBRAINS services within complex workflows.

HLRS RECEIVES “BLUE ANGEL” CERTIFICATION, PUBLISHES PRACTICAL GUIDE FOR SUSTAINABILITY IN COMPUTING CENTRES

HLRS has received certification for Energy-Efficient Data Center Operation under the Blue Angel, the ecolabel of the federal government of Germany. This certification reflects the comprehensive policies that HLRS has established to maximize its energy efficiency and minimize its impact on the environment. The Blue Angel is the latest of several certifications for sustainability and environmental responsibility that HLRS has achieved in the past year, including ISO 14001 (environmental management), ISO 50001 (energy management), and the Eco-Management and Audit Scheme (EMAS). HLRS has now also published a practical handbook that other computing centres can use to improve their sustainability across all levels of the organization. The guide provides tips for developing an environmental management system and an energy management system, advice for navigating the EMAS certification process, and suggestions for increasing staff participation in identifying and realizing sustainability goals.

STAFF SPOTLIGHT: LONG-TIME JSC GROUP LEADER HELPS SPEARHEAD NEXT-GENERATION NETWORKING TECHNOLOGIES

Olaf Mextorf, JSC

A In 1989, after finishing his studies and spending two years in application programming, one of Jülich Supercomputing Centre's (JSC's) longest-tenured employees, Olaf Mextorf, joined the Central Institute for Applied Mathematics (Zentralinstitut für Angewandte Mathematik), the predecessor of JSC. Mextorf was heavily interested in all aspects of data communications, and arrived to work for the Communication and Security Division at a time when centralized mainframes in combination with distributed "terminals" on the campus as well as a few experimental computers dominated the computing landscape.

In the 1990s, the main challenge for a network engineer was to migrate from one network-technology to the next. A lot of time and resources were used for evaluating, testing, deploying, operating and interconnecting radically different network technologies. "This period was a really good education for all my later activities in the area of network design, operation, and management" Mextorf said.

In the early 2000s, JSC began assembling a technical team focused on the increasingly complex aspects of high-speed networks across increasingly large amounts of compute nodes, and Mextorf became head of the Networking Group, responsible for the wired and wireless campus networks at Forschungszentrum Jülich. He is also responsible for the research institute's data centre and HPC networking as well as being a main partner for international project networks through the years.

During the past 30 years, high-performance computing (HPC) resources have only gotten larger and faster, with hundreds of thousands of compute cores strung together to deliver quadrillions of calculations per second. Consequently, demands have put on networking technologies and architectures has also grown almost exponentially. Today, Mextorf draws on over 30 years of experience to help guide networking innovations for what might be JSC's most ambitious HPC project yet—stringing together multiple clusters to form a novel modular supercomputing architecture for its flagship HPC system, JUWELS.

Mextorf has not been a passive observer to changes in networking technologies, though—through his position at JSC, Mextorf has played an active, collaborative role in helping shape technology companies' offerings based on HPC centres' needs. "One of the aspects of my work that has been most appealing for me is that I am able to be in contact all over the globe people developing research on the topic," Mextorf said. "We are not building hardware, but we are providing a lot of information and experience that typically the vendors do not have."

For Mextorf, working to spearhead technology development has always been a major motivation in his job. "We do have a long history of close collaborations, beta and early field tests for upcoming new network products and solutions," he said. "Our expertise and experience as well as testing opportunities when it comes to solutions for high bandwidth and scaling have always been very much appreciated by our partners. We're presently in an early-field test of a new Infiniband-Ethernet-Gateway solution from Mellanox. That product, called Skyway-Gateway, is not yet publically available but has been a long time coming for us, as one of our main challenges in the area of supercomputer networking is the gatewaying between the low-latency/high-speed interconnects like Infiniband and the Ethernet-based I/O-fabric."

The product will support 200 GBit/s links and will aggregate them in a widely scalable way under a common management. JSC is already running the new solution to access storage on the new JUWELS Booster system in its pre-production environment, allowing Mextorf and collaborators to get feedback and gather others' experiences with it and sharing this input with Mellanox.

"Close collaboration and the exchange of experiences with Gauss Centre for Supercomputing (GCS) partners, other national research institutes and universities, international supercomputing sites, and even big global companies has always been a very important and appreciated part of my job" said Mextorf, who is a long-time member of the Cisco Systems US-Data Center Advisory Board. "Sharing experiences with our peers, especially on our high-speed and large-scale deployment of network solutions and being able to influence and contribute to the development of new products, software and protocols is a very appealing aspect of my job."

HPC centres have long served as the testing grounds for new and novel HPC technologies, perhaps none more so than hardware that functions to move data as quickly and smoothly as possible. In 2020, JSC staff brought the next computing module for JUWELS online—the GPU-accelerated Booster module (for more information on the Booster, visit page 4). On the surface, it would seem as if forging tight networking bonds between two "small-scale" supercomputers might be his biggest challenge yet, but Mextorf explained that his years of experience and knowledge of the field more than adequately prepared him for this moment.

Forecasting how network technologies will change over 1 or 2 years can be a challenge, meaning that predicting a decade or more into the future is practically impossible. "With this kind of modular supercomputer, the integrating and the planning aspects from the JSC have been even more important than they normally would be," Mextorf said.

To that end, the centre identified partners to share in the long-term vision on a modular supercomputer. Partnering with Atos, ParTec, and Mellanox (now a subsidiary of NVIDIA), the centre forged bonds with technology partners dedicated to helping the centre realize its plans for JUWELS.

Working with Mellanox, JSC uses hundreds of Infiniband interconnects to connect the JUWELS nodes within and between the Cluster and Booster modules. Despite using slightly newer interconnects inside the Booster, Mextorf indicted that because the centre worked closely with Mellanox and uses Infiniband throughout the system, the connection is seamless. "From the users' perspective when it comes to connectivity and, to certain extent, even latency, they will see one Infiniband interconnect." JUWELS will deliver 200 gigabit-per-second network speed between nodes, ensuring that even users leveraging both the Cluster and Booster modules will be able to efficiently run large-scale, data-intensive simulations efficiently.

While forging close bonds between slightly different eras of technology is a challenge, Mextorf indicated more challenging networking tasks through the years. When JSC staff installed the Booster module for the centre's JURECA system 2 years ago, one of the challenging tasks was to tightly connect Infiniband and Intel's Omnipath and Ethernet—technologies that function mutually exclusive of one another. In order for the JURECA Booster to speak with the original half of the system as well as with the storage-cluster, new gateway technologies were developed that could facilitate data transfer between the three networks.

Challenges such as those connected to JURECA and JUWELS have kept Mextorf interested and engaged in his work for the past 30 years, and over the last 12, tighter integration and collaboration between Germany's leading HPC centres through GCS has further improved the ability for Mextorf and his collaborators to experience, test, troubleshoot, and ultimately help guide the development process for network technologies with hardware partners.

"GCS is really beneficial for all of us, as you get to interact with all these people who are more or less doing similar work, and we have the ability to share our approaches and experiences," Mextorf said. "I have also had a lot of looser, less direct collaborations with other international HPC sites, and these different collaborations have really developed through the years and are really one of the most appealing aspects of my job."

© JSC

Mextorf noted that this knowledge exchange often helps inform different HPC centres' staffs about common issues being encountered with different technology, and can lead to coordinated efforts to guide companies' approaches to further developing technologies.

Mextorf indicated that while it was difficult to predict the path of networking technologies, he envisioned substantial speed increases continuing to double every few years, but that networking technologies would continue to struggle to keep pace with massive increases in processing speeds.

He was more certain that no matter what shape the technological landscape took in the coming years, the enjoyment he gets from collaborating with his co-workers at JSC and partners across GCS and other international HPC centres would remain a highlight of his job. "I really enjoy collaborating with my fellow JSC staffers, and find this work fascinating, as there is always a new challenge for us all to collaborate on and solve," he said.

"Hearing people talk about the network in my area of responsibility as a resource that's always available, like the air around us, I always take it as a compliment for me and my team while knowing about all the complexity in the background and work behind the scenes".

eg

DURING PANDEMIC-RELATED REMOTE WORKING, GCS CENTRES EMBRACE EXPANDING E-LEARNING OFFERINGS

Despite having had only modest plans for online training courses in 2020, COVID-19 demanded that GCS centres' training staffs evolve to ensure the organization delivered on one of its core missions—training scientists to make the best use of HPC resources.

In early March 2020, Dr. Volker Weinberg, Head of Education and Training at the Leibniz Supercomputing Centre (LRZ), was in the midst of conducting a multi-day training course at the Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU). On Wednesday, March 11, he began the course just like he would on any other day, but at noon he was informed that everyone had to leave the lecture room immediately. The novel coronavirus pandemic that began sweeping the world in late 2019 had arrived in Bavaria, and state officials had decided to shutter all in-person meetings and courses at universities.

After an emergency call with LRZ leadership, Weinberg remained in his Erlangen hotel room and feverishly worked to postpone and cancel upcoming courses, inform participants of the situation, and ensure that the LRZ website reflected course cancellations. Despite knowing that courses planned for the coming days or weeks would not happen, Weinberg refused to accept that one of his organization's core missions—training users—would be cancelled indefinitely. "In this sad situation, I had already started thinking about moving everything online, because education in HPC must go on," he said.

All three centres comprising the Gauss Centre for Supercomputing (GCS)—the High-Performance Computing Center Stuttgart (HLRS), Jülich Supercomputing Centre (JSC), and LRZ—are charged with educating their users to leverage HPC resources effectively, and staff at all three centres take great pride in their individual training programs as well as robust collaborations with national and international partners.

In order to continue to deliver on that mission during a pandemic, GCS centres needed to collaborate and innovate their way to rethinking how training courses would look in 2020.

From on-site to online events

Once it became clear that the novel coronavirus disease (COVID-19) was going to cause indefinite interruptions to in-person meetings, GCS training staff began evaluating web conferencing tools in order to identify the most engaging and user-friendly software for online training courses.

"We had to choose the right conference tool for our training courses' needs, especially when it came to our ability to engage with and support participants' work during hands-on exercises," said Dr. Rolf Rabenseifner, Head of HLRS's Parallel Computing Training and Applications group. "For this decision, our national and international cooperations were extremely helpful." Rabenseifner credited his collaboration with the Technical University of Vienna's (TU Wien's) Dr. Claudia Blaas-Schenner for helping HLRS identify its conferencing tools and co-teaching courses in hybrid and parallel programming.

The HLRS trainers ultimately settled on Zoom. Weinberg also spent time evaluating several platforms, and prioritized the ability for training course participants to preserve the interactive nature of in-person courses as much as possible. With the ability to create real-time polls, organize smaller "break-out" rooms for work groups or specialized topics, and communicate through a variety of virtual "signals," from raising one's hand to asking for a coffee break, LRZ trainers also decided on using Zoom.

GCS centre staff had planned to participate in developing and running massively open online courses (MOOCs), sponsored by the pan-European HPC framework PRACE. For several years, HLRS offered working professionals, mainly from industry, the opportunity to do online training through the Supercomputing-Akademie. The online academy primarily focuses on promoting core competencies among industry researchers and other working professionals to enable them to use HPC in their research and development efforts. For the advanced courses, in which trainees use the GCS centres' HPC resources, training always happened on site. Before the COVID-19 crisis began, the core GCS HPC training program had not been planning to offer online courses this year.

The transition to taking training courses online was abrupt, and centres' staffs had to do more than just identify a web conferencing tool—they had to reimagine how courses would run efficiently online. Both Rabenseifner and Weinberg said that ensuring courses remained interactive and engaging was a top priority.

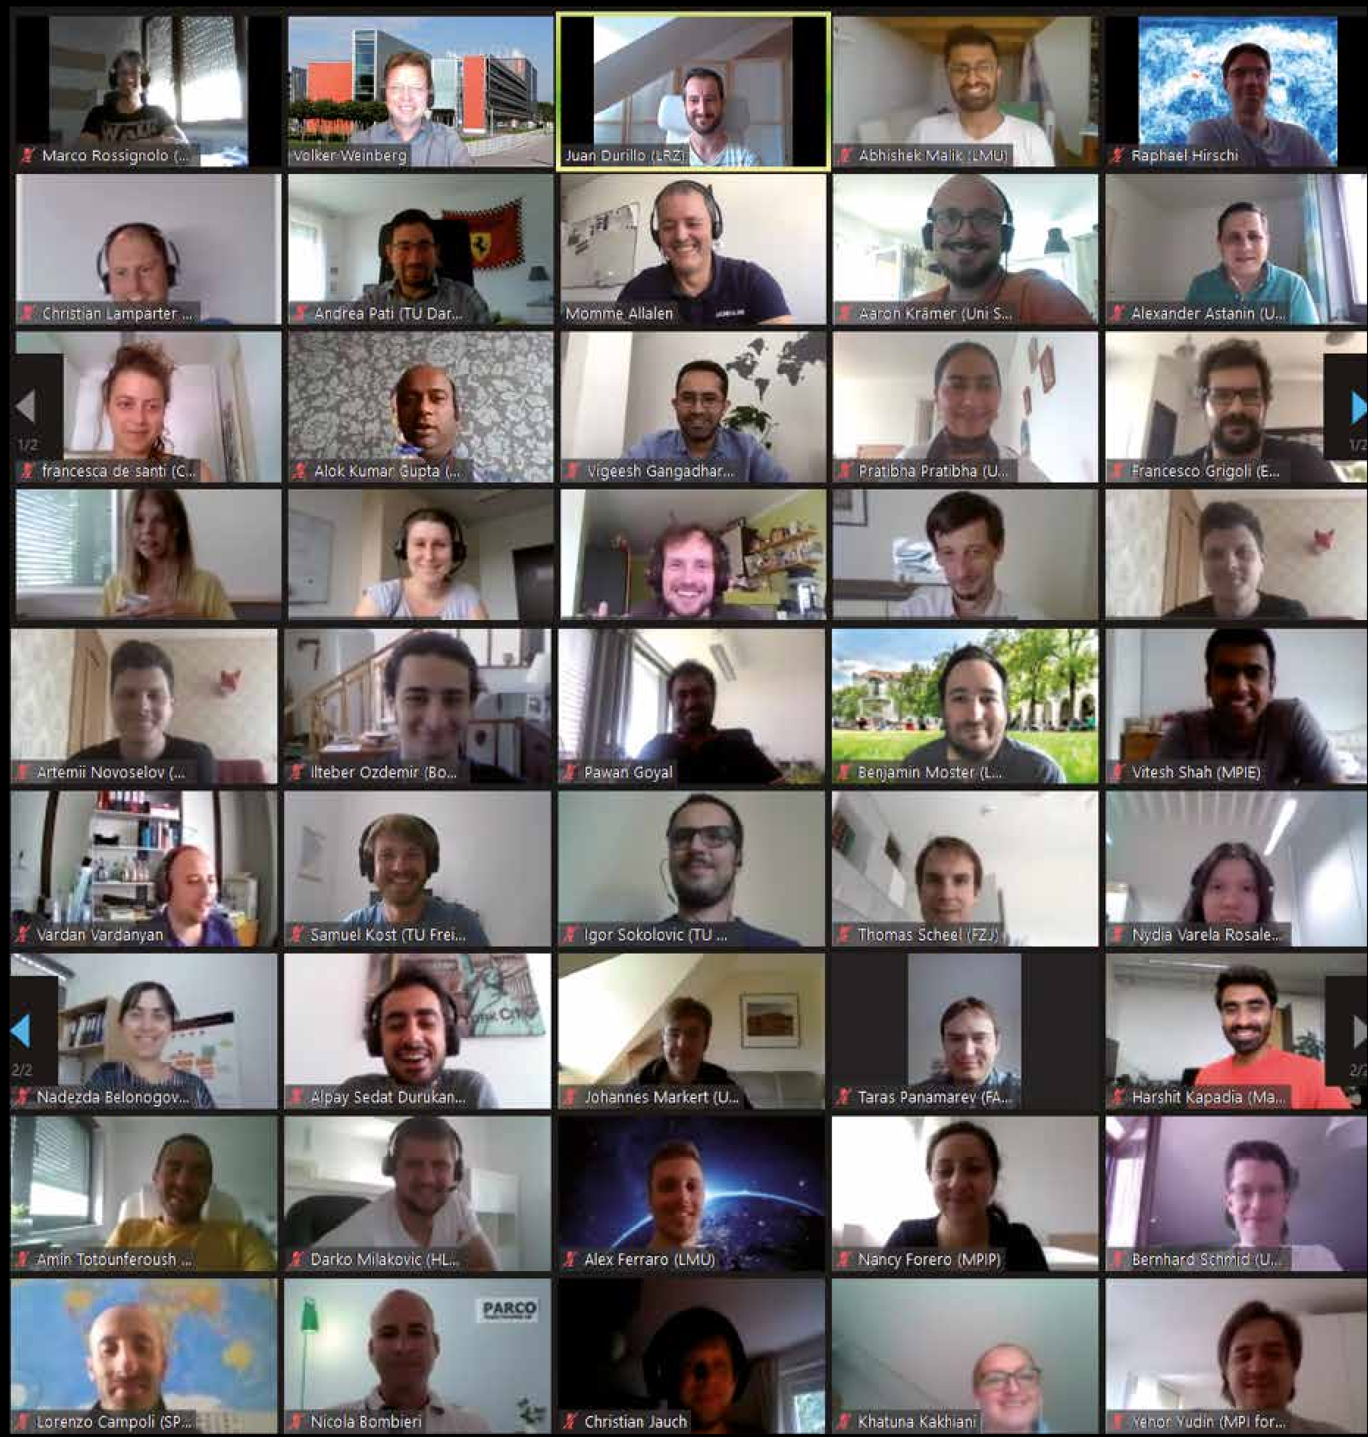

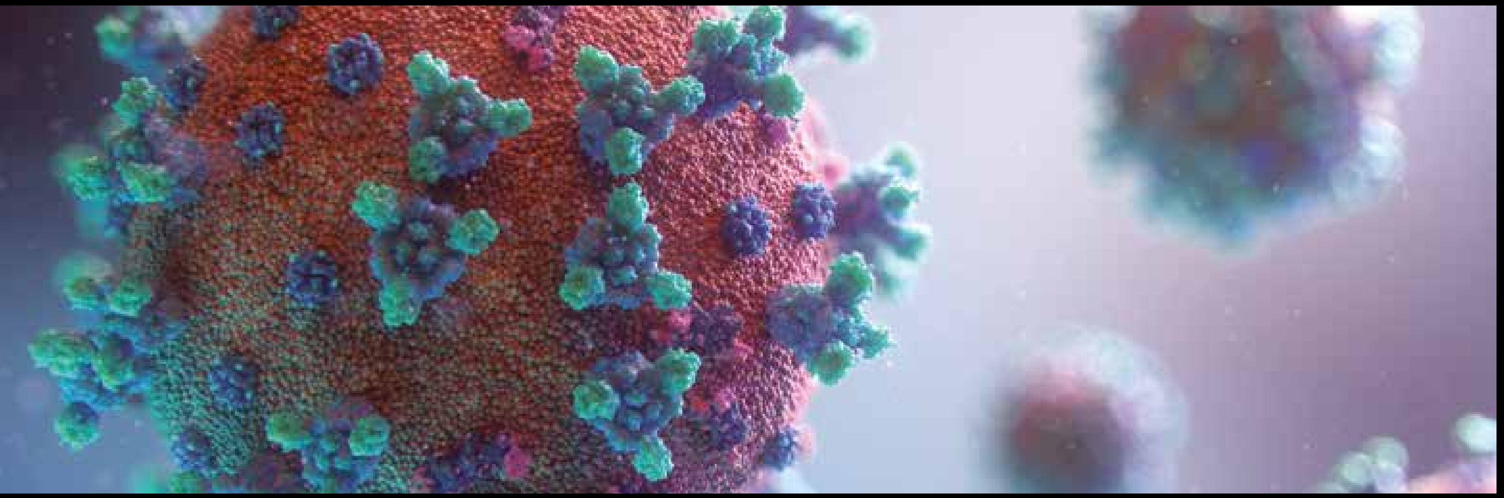

In some cases, that interactivity between trainers and students caused the centres' costs to rise. Rabenseifner pointed out that for traditional in-person training courses, HLRS uses an exercise concept of pairing course participants, which often allowed one trainer to teach a room of 40 to 50 people, offering hands-on assistance to these groups as he or she walked around the room. In online courses, though, HLRS found that to ensure the same quality of assistance, courses were best served with one trainer for every 7-10 participants. While the centre ultimately had to pay enough trainers to provide this level of support, centre staff knew that continuing to provide high-quality training courses was a cornerstone of its annual mission, and delivering on that goal was money well spent.

Centre staff had to reimagine other aspects of training courses as well, such as including social and networking aspects of multi-day events. For Weinberg and LRZ, for instance, that meant starting each course day with virtual polls to help guide the day's activities and supplanting evening dinner trips to Munich-area beer gardens with virtual beer garden events. Rabenseifner also noted that social contacts during the exercises and in the coffee and lunch breaks are lost, even when using breakout rooms for the exercises. He plans to test whether it would help to set up groups of 4-5 participants and build in longer coffee breaks, inviting them to start discussions on non-course-related topics like where they come from or why they are participating in the course.

Broader reach and impact for online training events

According to an old proverb, "necessity is the mother of innovation." With remote work and no travel becoming the norm during the pandemic, GCS had to innovate to keep its training program active. By identifying the right tools and bringing in a variety of motivated and flexible trainers, GCS centres not only ensured training courses for 2020 would continue with only modest interruption, but also wound up laying the framework for a more diverse and inclusive training program for the future.

The vast majority of courses continued as planned, and their attendance and engagement continue to grow. Since the pandemic hit Germany in March, GCS centres virtually hosted and co-hosted 726 participants across 16 different training courses, with 287 participants from outside Germany. The centres have also collaborated with partner organizations at the TU Wien, IT4Innovations in Ostrava, the University of Siegen, and FAU, among others.

Both Rabenseifner and Weinberg said many participants enjoyed participating in training events without having to travel, and despite costs associated with hiring additional training staff, the centres recouped some of the additional trainer costs by not paying for trainers' and teachers' accommodation.

Furthermore, online training has allowed GCS as the largest European PRACE Training Centre to expand its training footprint to HPC professionals in various countries, exposing both participants and trainers to a wider range of applications and problems, enriching the overall experience. Weinberg pointed out that users from 34 different countries have participated in GCS and PRACE training events since the start of the pandemic.

When the world comes out of the COVID-19 pandemic, GCS centres plan to reinstitute in-person training courses, as some courses truly benefit from in-person, hands-on events where trainer and student can sit side-by-side. Based on the positive feedback and success of online courses, though, both Rabenseifner and Weinberg envision virtual training courses remaining a core part of the GCS training portfolio, perhaps comprising as much as 50 percent of the courses offered annually.

"I am looking forward very much to meeting many readers of this article online in one of our upcoming online events," Weinberg said.

eg

TRAINING CALENDAR
HPC COURSES AND TUTORIALS

Editor's Note
Due to the COVID-19 pandemic, the GCS centres made significant changes to their 2020 training plans. Many courses were moved to a digital format and others were cancelled. Due to lingering complications and further unknowns related to the pandemic in the coming months, InSiDE editorial staff ultimately decided to run this feature highlighting the centres' efforts to continue delivering on GCS's core mission of training users to make the best use possible of world-leading HPC resources. We have also decided not to publish the training calendar as usual, as dates, locations, and plans may continue to change. For the most up-to-date information about GCS training courses, please visit:
https://www.gauss-centre.eu/trainingsworkshops

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

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

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

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.

Core tasks of JSC are:

- 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.
- Implementation of strategic support infrastructures including community-oriented simulation and data laboratories and cross-sectional teams, e.g. on mathematical methods and algorithms and parallel performance tools, enabling the effective usage of the supercomputer resources.
- Higher education for master and doctoral students in close coperation with neighbouring universities.

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

Compute servers currently operated by JSC

System	Size	Peak Performance (TFlop/s)	Purpose	User Community
Modular Supercomputer "JUWELS"	Cluster (Atos): 10 cells, 2,567 nodes 122,768 cores Intel Skylake 224 NVIDIA V100 GPUs 275 Tbyte memory	12,266	Capability Computing	European (through PRACE) and German Universities and Research Institutes
	Booster (Atos) 39 racks, 936 nodes 44,928 cores AMD EPYC Rome 3,744 NVIDIA A100 GPUs 629 TByte memory	73,008		
Modular Supercomputer "JURECA"	Data-Centric Cluster (Atos): 768 nodes, 98,304 cores AMD EPYC Rome 768 NVIDIA A100 GPUs 443 TByte memory	18,515	Capacity and Capability Computing	European (only on the Data-Centric Cluster) and German Universities, Research Institutes and Industry
	Booster (Intel/Dell): 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
Atos Cluster "JUSUF"	205 nodes, 26,240 cores AMD EPYC Rome 61 NVIDIA V100 GPUs 52 TByte memory	1,372	Capacity Computing	European and German Universities and Research Institutes through PRACE and the Human Brain Project
Modular Supercomputer "DEEP-EST" (Prototype)	Cluster: 50 nodes, 1,200 cores Intel Xeon Gold 6146 9.6 TByte memory + 25.6 TByte NVM	45	Capacity Computing (low/medium-scalable code parts)	Partners of the EU-project DEEP-EST and interested users through Early Access Programme
	Booster: 75 nodes, 600 cores Intel Xeon Silver 4215 75 NVIDIA V100 GPUs 6 TByte memory	549	Capacity and Capability Computing (high-scalable code parts)	
	Data Analytics Module: 16 nodes, 768 cores Intel Xeon Platinum 8260 16 NVIDIA V100 GPUs 16 Intel Stratix10 FPGAs 7.1 TByte memory + 32 TByte NVM	170	Capacity and Capability Computing (data analytics codes)	

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

For nearly six decades, the Leibniz Supercomputing Centre (Leibniz-Rechenzentrum, LRZ) has been at the forefront of its field as a world-class high performance computing centre dedicated to providing an optimal IT infrastructure to its clients throughout the scientific community–from students to postdocs to renowned scientists—and in a broad spectrum of disciplines–from astrophysics and engineering to life sciences and digital humanities.

Leadership in HPC and HPDA
Located on the research campus in Garching near Munich, the LRZ is a leadership-class HPC and HPDA facility delivering top-tier supercomputing resources and services on the national and European level. Top-notch specialists for HPC code portability and scalability support the LRZ' broad user base and ensure to run operations in the most energy-efficient way.

Future Computing at LRZ
The LRZ is leading the way forward in the field of Future Computing focusing on emerging technologies like quantum computing and integrating AI on large-scale HPC systems. A robust education program for HPC, machine learning, artificial intelligence and big data is complementing the LRZ offer.

IT backbone for Bavarian science
In addition to its role as national supercomputing centre, the LRZ is the IT service provider for all Munich universities as well as research organizations throughout Bavaria.

SuperMUC NG 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 LRZ's web pages: https://doku.lrz.de/display/PUBLIC/Access+and+Overview+of+HPC+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.

Hawk at the High-Performance Computing Center Stuttgart. © Ben Derzian for HLRS

CENTRE DESCRIPTIONS

Contact

High-Performance Computing Center Stuttgart (HLRS), University of 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
HPE Apollo 9000 "Hawk"	5,632 nodes 720,896 cores 1,44 PB memory	26,000	Capability Computing	European (PRACE) and German Research Organizations and Industry
NEC Cluster (Vulcan, Vulcan 2)	622 nodes 18,736 cores 119 TB memory	1,012	Capacity Computing	German Universities, Research Institutes and Industry
NEC SX-Aurora TSUBASA	64 nodes 512 cores 3072 GB memory	137.6	Vector Computing	German Universities, Research Institutes and Industry
Cray CS-Storm	8 nodes 64 GPUs 2,048 GB memory	499.2	Machine Learning Deep Learning	German Universities, Research Institutes and Industry

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

INNOVATIVES SUPERCOMPUTING IN DEUTSCHLAND   Nº 18-2 · AUTUMN 2020

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: Researchers from the University of Wuppertal and Forschungszentrum Jülich are using HPC resources at JSC to better understand the role of high-altitude clouds on ozone health. For more information, visit page 18. © Christoph Kalicinsky