Next-Generation Supercomputing

HLRS

High-Performance Computing Center Stuttgart

2022 Annual Report

2022
Annual Report

The High-Performance Computing Center Stuttgart (HLRS) was established in 1996 as the first German national high-performance computing (HPC) center. As a research institution affiliated with the University of Stuttgart and a founding member of the Gauss Centre for Supercomputing, HLRS provides comprehensive HPC services to academic users and industry. HLRS operates one of Europe's most powerful supercomputers, provides advanced training in HPC programming and simulation, and conducts research to address key problems facing the future of supercomputing. Among HLRS's areas of expertise are parallel programming, numerical methods for HPC, visualization, grid and cloud computing concepts, data analytics, and artificial intelligence. Users of HLRS computing systems are active across a wide range of disciplines, with an emphasis on computational engineering and applied science.

Director's Welcome

Grußwort

Prof. Dr.-Ing. Michael M. Resch, Director, HLRS

With this annual report we present the results of what was in many ways a transitional year for the High-Performance Computing Center Stuttgart (HLRS). At its start, 2022 continued to be strongly influenced by the COVID-19 pandemic, whose effects were noticeable in many areas.

Mit diesem Jahresbericht legen wir die Bilanz eines Übergangsjahres vor. Zunächst war das Jahr 2022 weiterhin stark geprägt von der Pandemie. Sichtbar wurden diese Auswirkungen in einer Reihe von Bereichen. Die Einnahmen aus Drittmittelprojekten sind 2022 zum ersten Mal seit 2013 wieder gesunken. Ein Teil dieses

For one, HLRS's income from our funded research projects fell for the first time since 2013. In part, this was a result of a pandemic that, despite frequent meetings by videoconference, made communication with our national and international network of scientific and industrial partners more difficult than usual. We anticipate that this will only be a short-term effect, however, as we are already seeing a recovery of our external funding in 2023 and could once again reach, or even exceed, our peak funding level from 2021.

We saw some very positive effects in 2022 in our continuing and professional education program for high-performance computing. Digital learning has made it possible for HLRS to reach new user groups. The transition from purely face-to-face learning to a blended concept that combines face-to-face and online learning meant that in 2022 a record number of more than 1,400 course participants took advantage of our diverse course offerings. The Supercomputing Academy, our training program for industry, also played a role in this accomplishment, as it has now transitioned from its project phase to become an established and important component of our curriculum.

From an organizational perspective, HLRS underwent a partial structural reorganization in 2022, welcoming a new steering committee into service. This panel, which is responsible for managing grants of computing time, now consists almost entirely of new members. We are delighted that it includes more women, is younger on average, and brings a wider range of scientific expertise. Together with the new committee chair, Prof. Thomas Ludwig of the German Climate Computing Center, we look forward to a new phase in the scientific usage of our resources.

We extend our thanks to the members of our previous steering committee, particularly to Prof. Wolfgang Nagel, who led the body for almost 20 years. His engagement and good judgment contributed enormously to HLRS's success during this period.

Effekts geht auf die Pandemie zurück, die trotz intensiver Nutzung von Videokonferenzen die Kommunikation mit unserem nationalen und internationalen Netzwerk aus wissenschaftlichen und wirtschaftlichen Partnern erschwert hat. Dieser Effekt ist aber nur kurzfristig, sodass wir für das Jahr 2023 bereits eine Erholung im Drittmittelbreich sehen und den Höchststand von 2021 entweder erreichen oder sogar übertreffen werden können.

Sehr positive Auswirkungen haben wir im Jahr 2022 im Bereich der Weiterbildung sehen können. Digitales Lernen hat dem HLRS neue Nutzergruppen erschlossen. Auch der Übergang vom reinen Vor-Ort-Lernen zu einem gemischten Konzept aus online und Vor-Ort-Lernen hat dazu geführt, dass wir im Jahr 2022 einen neuen Teilnehmerrekord von 1.400 erreicht haben. Eine Rolle spielt dabei auch, dass die Supercomputing-Akademie des HLRS nach einer Projektphase weitergeführt und verstetigt wurde.

Organisatorisch hat das HLRS sich im Jahr 2022 in seinen Strukturen teilweise neu aufgestellt. Zur Mitte des Jahres trat ein neuer Lenkungsausschuss sein Amt an. Das für die Vergabe von Rechenzeit verantwortliche Gremium wurde beinahe vollständig ausgetauscht. Der neue Lenkungsausschuss ist im Durchschnitt weiblicher, jünger und wissenschaftlich breiter ausgerichtet. Gemeinsam mit dem neuen Vorsitzenden des Lenkungsausschusses Prof. Thomas Ludwig vom Deutschen Klimarechenzentrum in Hamburg freuen wir uns auf eine neue Phase im Bereich der wissenschaftlichen Nutzung unserer Ressourcen.

Unser besonderer Dank gilt den Mitgliedern des bisherigen Lenkungsausschusses. Vor allem gebührt Prof. Wolfgang Nagel Dank, der den Lenkungsausschuss des HLRS seit fast 20 Jahren geleitet hat, und der mit seinem Engagement und seiner Umsicht erheblich zum Erfolg des HLRS in dieser Zeit beigetragen hat.

Die Ergebnisse der Nutzung des Supercomputers des HLRS wurden im Oktober beim 25. Results & Review Workshop am HLRS in Stuttgart vorgestellt. Wir heben

Results from the usage of HLRS's supercomputer were presented in October at the 25th Results & Review Workshop, and this annual report highlights a few of our users' applications. All contributions to the last workshop will also soon appear in the Transactions of the High-Performance Computing Center Stuttgart, published by Springer Verlag.

This annual report also presents new developments in our own research within the scope of our funded projects, covering a wide spectrum of activities. One approach that is proving to be particularly versatile is the use of digital twins, which make it possible to simulate and visualize all sorts of highly complex systems. The Stuttgart Research Inititative DiTEnS (Discursive Transformation of Energy Systems), funded by the Carl Zeiss Foundation, will develop methods in which digital twins support decision making processes to make cities more energy efficient. The CIRCE project (Computational Immediate Response Center for Emergencies) is investigating scenarios for the use of digital twins in crisis management. HLRS and its partners are also exploring the potential of digital twins within CASE4Med, a project that is laying the groundwork for a new Solution Center for medical applications.

What all of these projects have in common is that they extend beyond the world of classical simulation, using classical approaches in a targeted way to contribute to advances in new fields. In addition, all of these projects involve the implementation of a concept of digital convergence that has been developed at HLRS. By bringing together technologies from the Internet of things (data gathering), artificial intelligence (data analysis), and high-performance computing (simulation of complex systems), we are demonstrating the potential that exists in combining today's advanced digital tools to solve new kinds of problems.

In two other projects, HLRS continues to coordinate a Europe-wide effort that is establishing a network of national HPC competence centers within the scope of the EuroHPC Joint Undertaking. The projects EuroCC

in diesem Jahresbericht einige Highlights der Benutzeranwendungen heraus. Alle Beiträge des vergangenen Workshops erscheinen demnächst in den Transactions of the High-Performance Computing Center Stuttgart im Springer Verlag.

In diesem Jahresbericht werden auch Forschungsentwicklungen innerhalb unserer geförderten Projekte vorgestellt, die ein breites Spektrum an Aktivitäten abdecken. Als besonders vielseitig erweist sich der Einsatz von digitalen Zwillingen, mit denen sich hochkomplexe Systeme aller Art simulieren und visualisieren lassen. Die von der Carl-Zeiss-Stiftung geförderte Stuttgarter Forschungsinitiative DiTEnS (Discursive Transformation of Energy Systems) wird Methoden entwickeln, mit denen digitale Zwillinge Entscheidungsprozesse unterstützen, um Städte energieeffizienter zu machen. Das Projekt CIRCE (Computational Immediate Response Center for Emergencies) untersucht Szenarien für den Einsatz von digitalen Zwillingen im Krisenmanagement. Das HLRS erkundet auch gemeinsam mit Partnern das Potenzial digitaler Zwillinge für die Medizin im CASE4Med Projekt, das den Grundstein für ein neues Solution Center für medizinische Anwendungen legt.

Allen diesen Projekten ist gemeinsam, dass sie über die klassischen Simulationsbereiche hinausgehen, indem sie klassische Ansätze gezielt in neue Bereiche weiterentwickeln. Allen diesen Projekten ist aber auch gemeinsam, dass sie das am HLRS entwickelte Konzept der digitalen Konvergenz umsetzen. Dabei wird im Zusammenspiel von Techniken des Internet of Things (Datensammlung), der künstlichen Intelligenz (Analyse von Daten) und des Höchstleistungsrechnens (Simulation komplexer Systeme) das Potential aller verfügbaren digitalen Technologien zur Lösung neuartiger Probleme eingesetzt.

Weiterhin koordiniert das HLRS auch die Etablierung von HPC Kompetenzzentren der EuroHPC JU Initiative. Ende des Jahres sind die ersten Phasen von EuroCC und CASTIEL ausgelaufen und EuroCC 2 und

and CASTIEL concluded their first project phases at the end of 2022, but have transitioned seamlessly into their second project phases as EuroCC 2 and CASTIEL 2. These activities are supporting the development of a German national competence center, and continue to strengthen expertise across Europe.

In the midst of our many scientific activities, HLRS remains at its core a national operations and service center for high-performance computing, and moved forward on two key infrastructure projects in 2022. The first has involved planning a new building, which is taking place in collaboration with the University Construction Office for Stuttgart and Hohenheim, and the science and finance ministries of the State of Baden-Württemberg. Our goal is to complete construction of a new structure that will be needed to accommodate a next-generation supercomputer by the end of 2026. Working together with the HLRS steering committee in 2022, we also completed preparations for the procurement process for this new system, which started at the beginning of 2023.

This means that in 2023 HLRS will face new challenges. With the help of its sponsors at the state, federal, and European levels, with its national and international partners, and thanks to the commitment of its staff, however, it is clear that HLRS is well prepared to meet them.

With best regards,

Prof. Dr.-Ing. Dr. h.c. Prof. E.h. Michael M. Resch
Director, HLRS

CASTIEL 2 schließen jetzt hier nahtlos an und entwickeln sowohl das deutsche Kompetenzzentrum als auch die europäische Stärkung der Kompetenzen weiter.

Bei all den wissenschaftlichen Themen bleibt das HLRS aber auch ein nationales Betriebs- und Servicezentrum für das Höchstleistungsrechnen. Daher haben wir im Jahr 2022 zwei zentrale Projekte des HLRS vorangetrieben. Die Planungen für ein neues Rechnergebäude sind gemeinsam mit dem Universitätsbauamt Stuttgart und Hohenheim, dem Wissenschaftsministerium und dem Finanzministerium des Landes Baden-Württemberg im letzten Jahr intensiv vorangetrieben worden. Ziel ist die Bereitstellung des Gebäudes Ende 2026 für eine neue Rechnergeneration. Gleichzeitig wurde im Jahr 2022 in Zusammenarbeit mit dem Lenkungsausschuss des HLRS die Ausschreibung für diese neue Rechnergeneration vorbereitet, die Anfang 2023 gestartet werden konnte.

Das HLRS steht also im Jahr 2023 vor neuen Herausforderungen, auf die es mit Hilfe seiner Fördergeber aus Land, Bund und Europa, mit Hilfe seiner nationalen und internationalen Kooperationspartner und dank des Engagements der Mitarbeiterinnen und Mitarbeiter des HLRS sehr gut vorbereitet ist.

Mit freundlichen Grüßen,

Prof. Dr.-Ing. Dr. h.c. Prof. E.h. Michael M. Resch
Direktor des HLRS

Supercomputing at the Limits

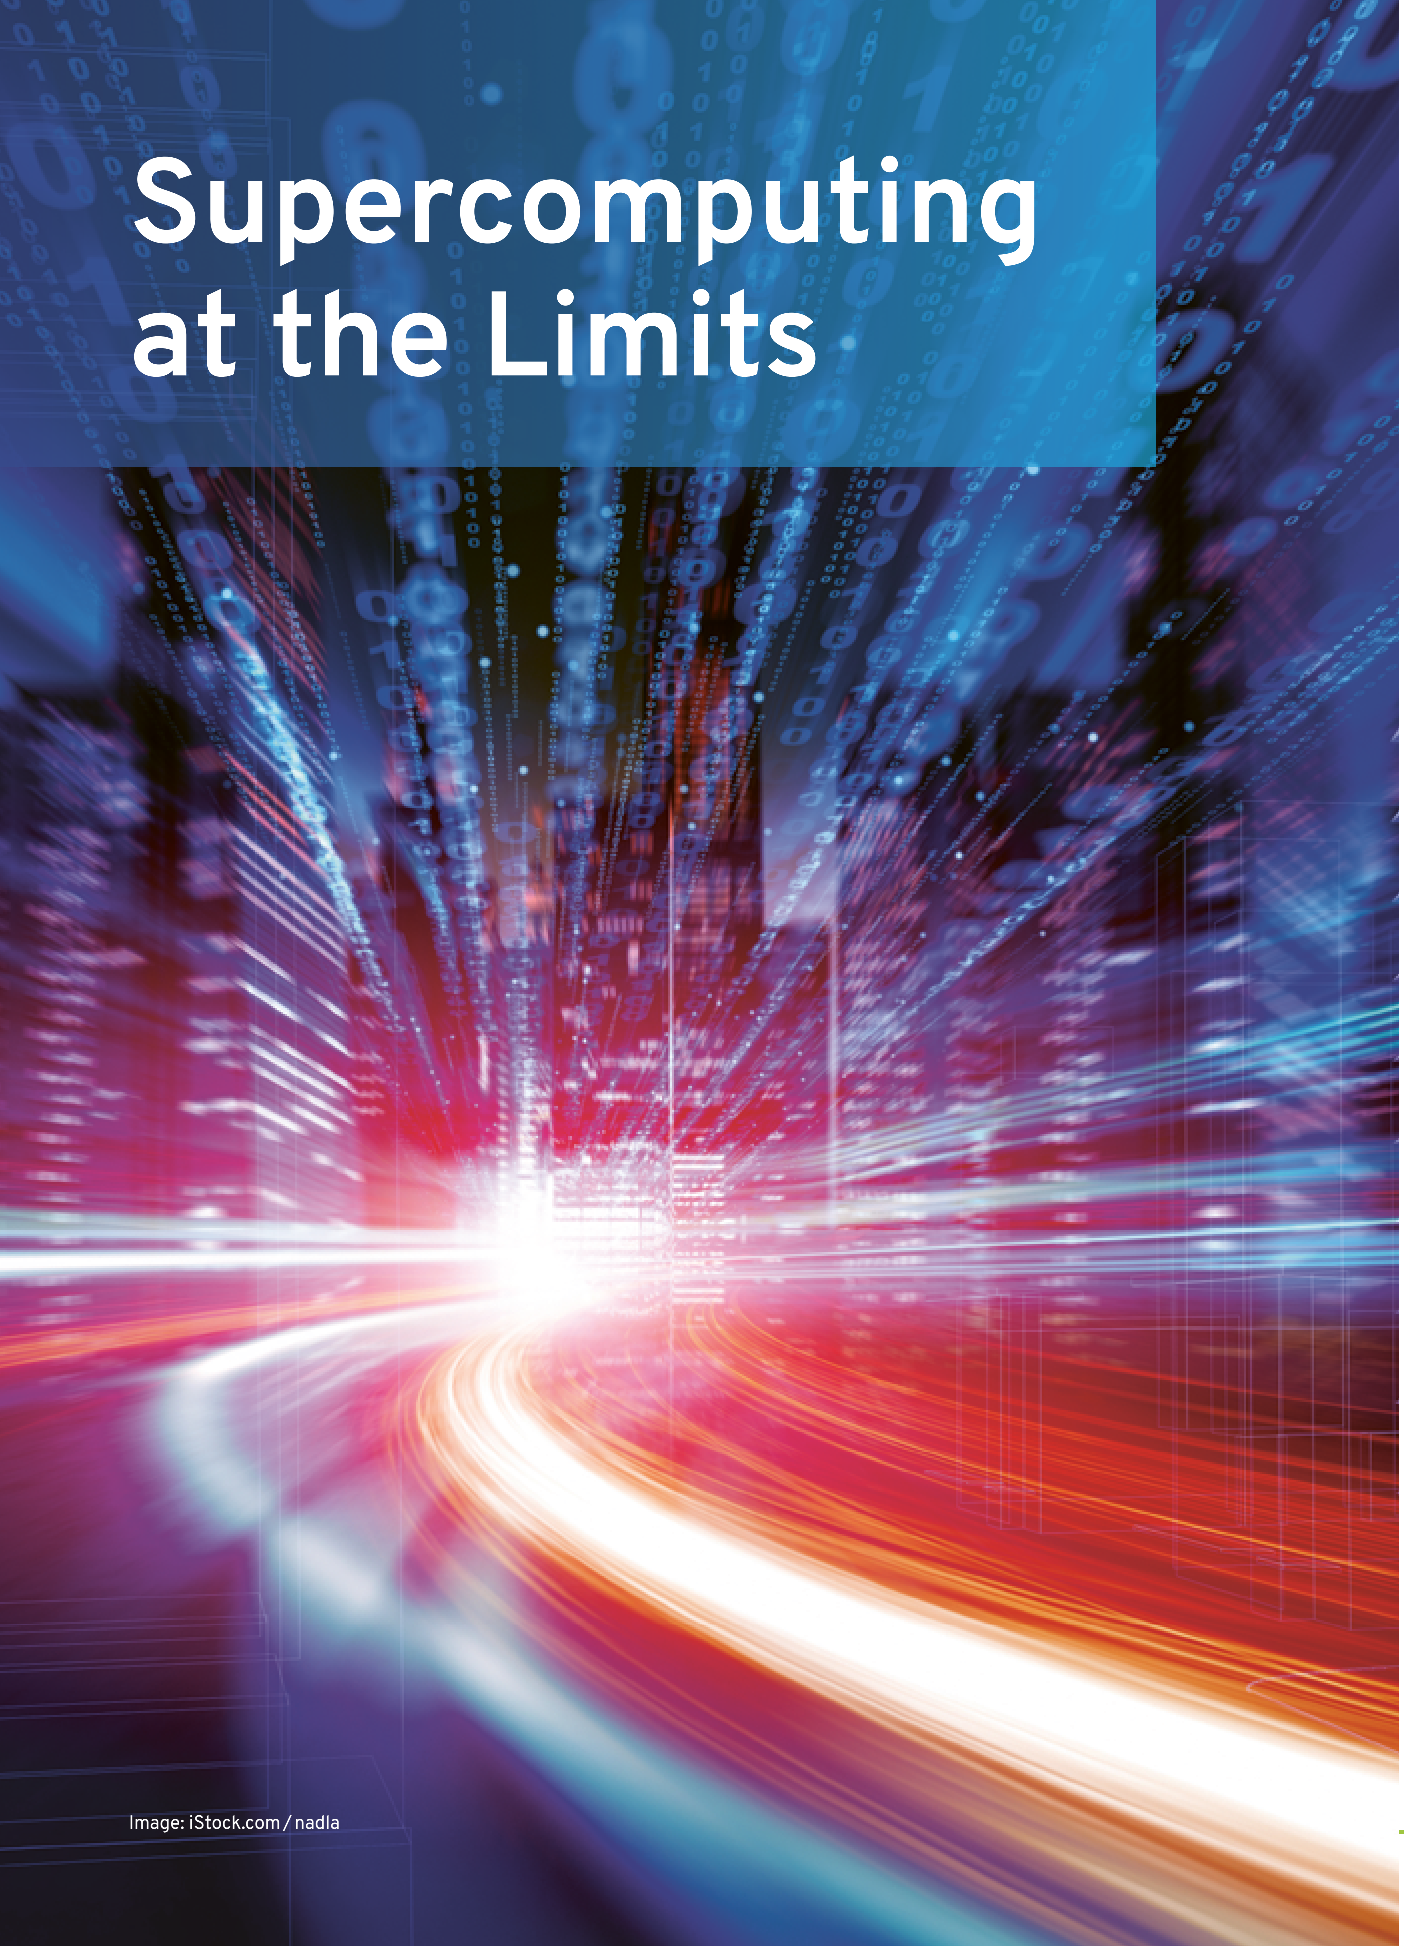

Image: iStock.com / nadla

On the Way to the Next Generation of High-Performance Computing

A survey of new research projects at HLRS offers a glimpse into key technical challenges that HPC faces, and how the center is working to find solutions that will shape the field's future.

In the world of high-performance computing (HPC), 2022 will be remembered as the start of the exascale era. The June 2022 edition of the Top500, a list of the world's fastest supercomputers, heralded the Frontier system at Oak Ridge National Laboratory in the United States as the first exaflop machine – a leap to the next order of magnitude in computing speed that is equivalent to running more than one quintillion (10^{18}) floating point operations per second. Europe is well on its way to breaking this barrier as well. In August 2022, the EuroHPC Joint Undertaking announced that the Jülich Supercomputing Centre (JSC) will soon be home to Europe's first exascale system, while pre-exascale systems are in the process of being rolled out across the continent. JSC's partners in the Gauss Centre for Supercomputing – the Leibniz Supercomputing Centre (LRZ) and the High-Performance Computing Center Stuttgart (HLRS) – are also preparing for their own jumps to exascale. At HLRS planning is underway to install a system at this level by 2027.

Achieving this lightning speed is not simply a matter of building larger machines, however. This is because the pre-exascale and exascale systems being developed today are fundamentally different from supercomputers of previous generations. Crucially, major increases in power are needed to run and cool systems of this size, meaning that maximizing energy efficiency in their operation and usage is more important than ever for ensuring their financial and environmental sustainability. Simultaneously, hardware manufacturers are reaching the physical limits of how much performance can be achieved on each individual computer chip, meaning that the progression predicted by Moore's Law has basically run its course. At the system architecture level, this combination of factors means that whereas supercomputers of an earlier generation grew by simply adding larger numbers of central processing units (CPU), the new generation increasingly combines CPUs with purpose-built accelerators, which, in most cases, are based on graphic processing unit (GPU) technology. Such heterogeneous systems are faster and more energy efficient, but they also require new programming models and software to fully leverage their performance. This is not only because of the larger numbers of processors but also because, currently, many widely used software packages and scientific applications developed for CPU-only systems are hardly supported on accelerated systems, let alone optimized.

Compounding these challenges is the convergence of supercomputers with other digital technologies, leading to the development of increasingly hybrid computing systems and workflows. Sensors, edge and cloud computing, artificial intelligence, and quantum computing offer new opportunities for research, technology development, and public administration, but combining them effectively in ways that utilize their full power requires new programming workflows and systems operations. HPC is no longer just a matter of running a large simulation on a single supercomputer, but is increasingly becoming a complex, distributed process that must be coordinated among computers with different capabilities and programming requirements, often

Driving Factors	Technological Challenges	Current Research at HLRS
Growing demand for HPC resources	Operation of larger, massively parallel systems	**Programming & application development** Code scaling and performance optimization Algorithm migration to accelerated systems Hybrid workflow development User training and support
Energy efficiency and sustainability requirements	Increasing energy demands	
End of Moore's Law	Production of massive datasets	**System management** Adaptive system parameterization Load balancing optimization Improving input / output efficiency Data management
Opportunities offered by new digital technologies	Convergence of HPC with cloud, edge, AI, and quantum computing technologies	
	Hybrid, distributed, accelerated systems and networks	**Sustainability** System performance monitoring and optimization Intelligent regulation of operating temperature Recycling of waste heat Supply chain management

An overview of some key issues and research areas that will shape the future of high-performance computing.

situated at different locations. Orchestrating the processes needed to get these technologies to talk to one another and to move data in a fast and secure way is demanding new approaches for managing tasks across such networks.

In addition, not all hardware components that are essential for high-performance computing are evolving at the same rate. Historically, for example, memory hardware has struggled to keep pace with accelerated systems, meaning that the writing and reading of data can still slow down large-scale simulations. Artificial intelligence (AI) is also changing key characteristics of scientific data. Whereas classical simulations used relatively few input data, AI requires the management of massive datasets composed of millions of small files, meaning that input and output in distributed file systems must be optimized for data processing. And although more processing speed permits more complex ensemble simulations or multiphysics models, for example, it also means generating ever-larger mountains of data. Archiving it for reuse in future studies, for training AI algorithms, or to double-check the results of a scientific paper threatens to overwhelm HPC centers, both because of the space and power requirements for data storage and the time required to backup or transfer large datasets.

"As high-performance computing continues to grow, it is changing in ways that present a host of new challenges," says Prof. Dr. Michael Resch, director of HLRS. "As a federal high-performance computing center working within this emerging landscape, it is our job to provide the infrastructure, solutions, support, and training that will help scientists to navigate these changes, and ultimately to get answers to their complex questions efficiently and in a sustainable way."

In several new research projects launched in late 2022 and early 2023, staff scientists at HLRS are creating and testing potential solutions to some of these big challenges facing HPC. Conducted in partnership with other leading HPC centers and industry, these projects will both contribute to the development of the field and ensure that HLRS continues to provide state-of-the-art support to its community of system users. These projects also offer a window for beginning to understand some of the ways in which high-performance computing and related fields will need to evolve in the coming years.

Software at the exascale

One advantage of larger HPC systems is that they make it possible to run simulations in which potentially billions of parallel calculations are executed simultaneously. In many simulations, for example in computational fluid dynamics (CFD) or climate modeling, programmers create a computational mesh that divides a large simulation into smaller units that are calculated individually and then reintegrated to understand the system as a whole. To use parallel computing systems most efficiently, programmers must adjust the distribution of these units based on the number and types of processing units that are available. As larger computers with hybrid architectures go online, the number of processing units rises, making it more difficult to achieve efficient performance using today's algorithms. This means that for many conventional problems that scientists and engineers need to solve, taking full advantage of the speed that new exaflop-capable systems offer will not happen unless codes are scaled to meet them.

As a member of a new EuroHPC Joint Undertaking Centre of Excellence called **CEEC** (Center of Excellence for Exascale CFD), HLRS is working to improve state-of-the-art algorithms and methods used in computational fluid dynamics to ensure that they perform efficiently at exascale. The project aims to develop exascale-ready workflows for extremely large computing systems, implement methods for reducing the amount of the energy used to run these algorithms, and demonstrate these new algorithms' effectiveness in applications that are important in academic and industrial research. The project is focusing on key algorithms for a variety of fields that rely on CFD simulations, including aeronautical engineering, environmental science, the chemical industry, the wind energy industry, and atmospheric sciences.

HLRS is also coordinator of the project **EXCELLERAT**, the European Centre of Excellence for Engineering Applications, which has been pursuing a related strategy to prepare industry for the next generation of high-performance computing. As a service provider, the project is supporting the development of key codes used in industrial sectors such as automotive, aerospace, and energy to run efficiently on larger, hybrid HPC systems. Research in the project has helped to adapt existing codes to run efficiently across dramatically larger numbers of processors, including on systems that include GPUs and other newer types of processors. The results have shown increases in processing speed of up to 90% as well as dramatic increases in simulation resolution, making it possible to show finer detail in simulations of airflow around airplane wings or in combustion reactions, for example. Near the end of 2022, EXCELLERAT was funded for a second project phase, and will continue to support industry in preparing for the next generation of HPC in the coming years.

HLRS continues to be involved in two additional EuroHPC Centres of Excellence that are focused on software for exascale computing and entered their second phases at the beginning of 2023. The first is **ChEESE**, which is developing exascale-ready codes for solid earth research that could support early warning forecasts, hazard assessment, and emergency responses to geohazards such as volcanoes, earthquakes, or tsunamis. The **HiDALGO** project, for which HLRS serves as technical coordinator, has also been extended, focusing on developing methods that could help to address global challenges using new, hybrid HPC systems.

Smarter systems

In the past, a supercomputer's energy usage could be limited using crude approaches such as reducing the clock frequency that controls processor speed or shutting off sections of the system when they were not in use. Modern HPC systems, however, offer a growing number of options that hold high energy savings potential. For example, adjusting parameters and settings in OpenMP and MPI — two important programming paradigms for parallel computing systems — can improve software performance, leading to more efficient energy usage. When multiple user applications are running on the system at the same time, system administrators can also track and optimize how those applications run in a more holistic, system-wide basis using MPI. Determining the optimal settings for such a systemic approach can be difficult, however, particularly when HPC systems simultaneously run many, diverse applications.

In a project called **EE-HPC**, HLRS is helping to develop and test a new approach initiated at the University of Erlangen-Nürnberg that aims to reduce energy consumption while maximizing computational throughput. Using machine learning, software will dynamically set system parameters to optimize energy usage in hardware based on the jobs and job phases that are running at any particular time. Bringing many years of experience as a member of the MPI-Forum, which sets standards for this widely used programming framework, HLRS will enable the integration of monitoring software into the runtime environment of OpenMP and MPI. A graphic user interface will also offer users transparent insights into the decisions the system is making while running their software.

In the project **targetDART**, HLRS is pursuing strategies to improve the scalability and energy efficiency of applications by optimizing load balancing. Here, the focus is on the programming interface OpenMP, which orchestrates the distribution and execution of computing tasks across a parallel computing system, preventing spikes in activity on some parts of the system while other parts sit idle. The challenge is that because computing tasks in parallelized simulations depend on the output of other tasks, data must constantly be physically moved around the computer, and the time it takes for processors to communicate with one another can slow down the system. On today's largest supercomputers, optimizing load balancing is extremely difficult, and it becomes even more challenging in hybrid systems, particularly as the scale of the entire system and thus the number of components to be monitored and optimized increases. By pursuing new strategies for managing task-dependencies and for monitoring and evaluating the performance of applications, targetDART aims to address this problem. As a member of the MPI-Forum, HLRS will also distribute advances made during targetDART among the wider HPC community.

Digital convergence: putting the pieces together

As supercomputers grow toward exascale, other kinds of digital technologies have also been evolving that could extend the usefulness of high-performance computing far beyond the walls of the traditional HPC center. Sensors of all kinds, for example, now collect measurements that serve as the foundation of new models and simulations. With edge computing, computational tasks can be distributed to sites where data is gathered, making it possible to make decisions faster. Even within HPC centers themselves, new workflows are needed to integrate simulation and data analysis, which run best on different computing architectures. Putting all of these pieces together is one major task that high-performance computing is currently facing.

Emblematic of the challenges of this diversifying landscape is a recently launched project called **DECICE**, which focuses on cloud and edge computing. Such architectures are relevant in domains such as smart cities, industrial automation, and data analytics, where new applications often involve specialized hardware that is located close to users. Integrating these devices with high-performance computers like HLRS's Hawk will mean ensuring low latency and high security during data transmission, as well as location awareness across the network.

DECICE is testing new methods for unifying such distributed networks of devices with a central controlling cluster. HLRS scientists will use KubeEdge, a system derived from the open-source framework Kubernetes, which was designed for deploying, scaling, and managing applications in large-scale hybrid computing systems using so-called containers. DECICE will further develop KubeEdge, which brings Kubernetes' containerized approach to edge computing, using an AI-based approach to assign jobs to the most suitable resources across a distributed system made up of different kinds of devices and processors. HLRS is providing HPC infrastructure for DECICE, as well as its expertise in cloud computing, HPC programming, and HPC system operation. It will lead a work package focused on developing an integrated framework for managing tasks in the cloud, edge, and HPC.

Training users to program new HPC architectures

As the landscape of larger, hybrid HPC system architectures becomes more diverse, HLRS's training program has also been changing to ensure that the center's computing resources are used most effectively. In addition to its traditional menu of courses focusing on programming languages for scientific computing and parallel programming frameworks like MPI and OpenMP, the center expanded its offerings to include new courses on GPU programming, deep learning, and artificial intelligence in 2022. This included a training collaboration with hardware manufacturer NVIDIA, which involved "bootcamp" workshops that introduced the use of artificial intelligence in science and offered a deeper dive into scientific machine learning using physics-informed neural networks. Another course held in partnership with AMD provided specialized instruction in machine learning using the company's Instinct GPUs. Additional new courses at HLRS focused on programming models for adapting existing codes to accelerated architectures, including a collaboration with INTEL focusing on oneAPI, SYCL2020, and OpenMP offloading.

Here as in many disciplines that must work together to prepare for the future of high-performance computing, adaptation is the key word, particularly as the field approaches the limits of what came before. Whether it be limits on energy supplies and natural resources, the physical limits of a traditional CPU chip, limits in the ability to manage the data that HPC systems now produce, or limits in the flexibility of codes written for older architectures, a wide spectrum of challenges is converging in ways that are forcing HPC to evolve into something new and potentially even more powerful. Through its research and training initiatives, HLRS aims to be a protagonist that helps to push this transformation forward. *CW*

In onoing projects such as ENRICH, DEGREE, and SRI DiTEnS, HLRS is investigating new methods for improving its energy efficiency and environmental performance.

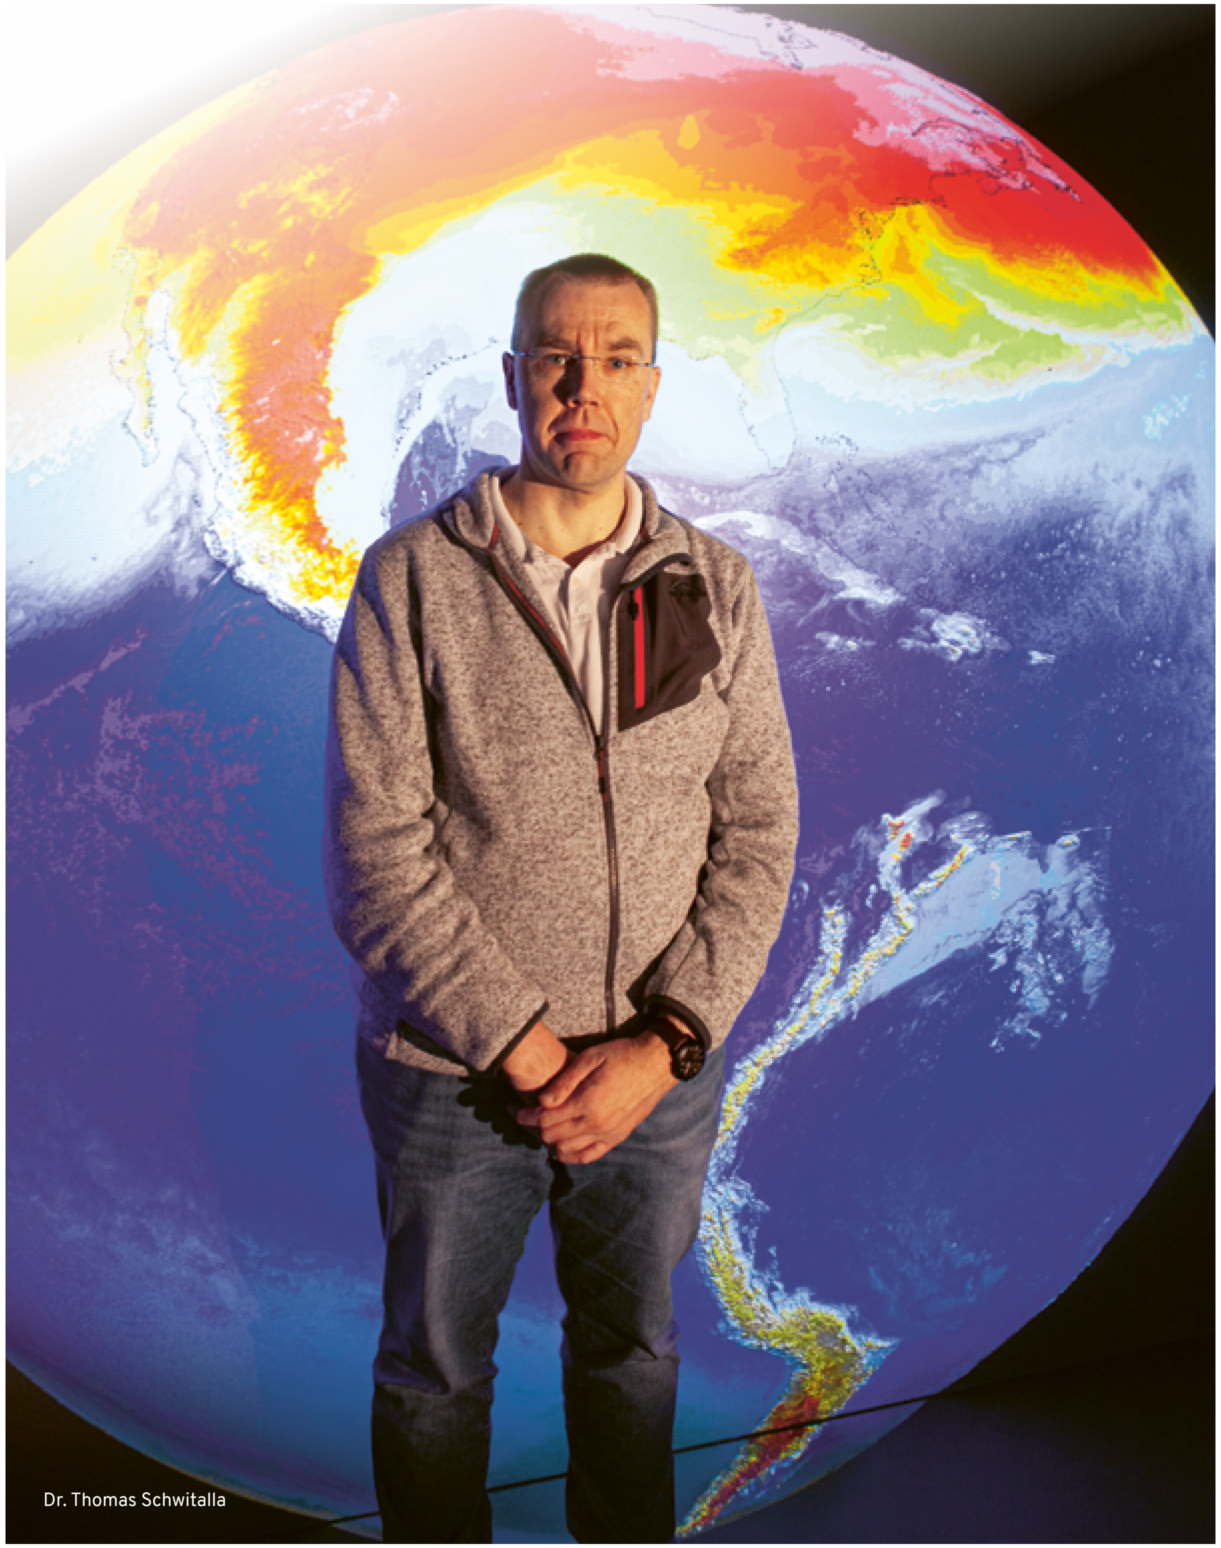

Dr. Thomas Schwitalla

The Future of Atmospheric Simulation: An Interview with Thomas Schwitalla

Dr. Thomas Schwitalla, a physicist at the University of Hohenheim, is a long-time user of HLRS's high-performance computing systems. Using the Weather Research & Forecasting Model (WRF) and the Model for Prediction Across Scales (MPAS), Schwitalla works on weather simulations that have become increasingly precise in the past few years. Recently, he used HLRS's supercomputer, Hawk, to simulate the weather of the entire Earth at an impressively high resolution of just 1.5 km.

As is the case in other research fields that work with large simulations, global weather and climate models pose a challenge for high-performance computing. Data systems in highly parallelized computing systems must constantly move an enormous number of data files between compute nodes and storage. Even as the compute nodes in an HPC system get faster, this so-called "Input/Output (I/O)" process can dramatically affect the efficiency of a simulation.

In a new project called TOPIO, Schwitalla is working together with HLRS scientists to test an approach that could improve the I/O of large-scale simulations. In the following interview, he describes the opportunities and challenges that the march toward exascale poses for HPC.

Why is it important to achieve such a high resolution in global weather and climate modeling?

Currently there are global climate simulations that run at a resolution of 25 to 50 km. This is good, but they aren't sufficient if you want to look at the local level. In the Black Forest, for example, there are a lot of mountains and valleys, and a resolution of 25 km is not capable of capturing this complex topography. You also see this when looking at archipelagoes like in the Philippines, Indonesia, or the Canary Islands, where a mesh size with a resolution of 25 km simply can't represent smaller islands.

The hope is that higher resolution models will make it possible to make more accurate predictions. You can represent atmospheric processes in a much more realistic way, which can help, for example, in predicting tropical storms or sudden heavy rain events. Being able to know what is going to happen in the next 30 to 60 days is also of great interest for agriculture, because farmers want to know the best time to seed or harvest. In one EU-funded project, researchers have conducted the first 40-day simulations and have shown that a higher resolution has a significantly positive effect on the prediction of precipitation.

How much data is produced in such high-resolution, global simulations?

If you were to divide the entire surface of the Earth into 1.5 × 1.5 km boxes, using MPAS you would have 262 million cells. In the configuration I use, 75 vertical layers are also needed to simulate the atmosphere. This means that I have 262 million × 75 cells; that is, approximately 20 billion cells. To model weather realistically, in every one of those cells I need to know the temperature, the air pressure, the wind, and the relative humidity. In addition, you need to account for variables like cloud water, rain, ice, or snow. Every one of these variables has 20 billion data points that need to be saved. Even when computing by using single precision (a format in which floating points are saved using 32 bits instead of 64 bits) the resulting file is approximately a terabyte in size per time step in the output. When simulating 24 hours of time with one file export for every simulated hour, you would have 24 terabytes of data. For a simulation of 60 days that means 1.3 petabytes.

Such huge amounts of data pose a challenge for data systems. During simulations using 10,000 or 100,000 compute cores you run into the situation that on the one hand, the communication between MPI processors isn't fast enough to keep up, and on the other that the data system cannot process the data fast enough because of the limited memory bandwidth.

In the TOPIO project you and scientists at HLRS are addressing the I/O problem. What are you planning to do?

In TOPIO we want to investigate whether it is possible to accelerate the I/O of the data system by using an autotuning approach for data compression. Algorithms for data compression already exist, but they have not yet been parallelized for HPC architectures. We want to determine if such approaches could be optimized using MPI. Scientists from the EXCELLERAT project, which HLRS is leading, have developed a library for compression called BigWhoop, and in TOPIO we will apply the library on MPAS. The goal is to reduce data size without losing information by up to 70 %.

This last point is very important for atmospheric modeling. What is particularly difficult is that specific variables such as water vapor, cloud water, or cloud ice are measured in very small values. If I lose a place in a number during data compression, it can lead to differences in the information that comes out of a longer simulation, as well as the resulting statistical analysis. The problem is also relevant when you compare simulations with satellite data. When information is lost because of compression it can mean, under certain conditions, that specific features simply go missing.

Reducing the amount of data is also extremely important for users of the information, for example in agriculture or flood prediction centers. If the datasets are large and difficult to transfer, no one will use them. If the data are compressed, chances are higher that they will be used as a part of decision-making processes. Another positive side effect is that this can reduce the energy necessary to save and archive the data.

What are some of the next big questions in atmospheric modeling research, and how will HPC need to develop so that it is in the best position to support it?

In the future, scientists would like to be able to better estimate uncertainties in their models. We see this need, for example, in cloud modeling. How clouds develop is very important for the climate because clouds have a big influence on the Earth's radiation balance. If the resolution of a model is coarse, it is not capable of representing small clouds. In some cases this can even mean that the radiation balance is not simulated correctly. If you use this information as part of a climate scenario to project what will happen in the next 100 years, the danger is that the mistake can become large.

Ensemble simulations can help to estimate such uncertainty. In multiphysics ensembles, for example, one might select a variety of different parameter settings for representing clouds or atmospheric processes at

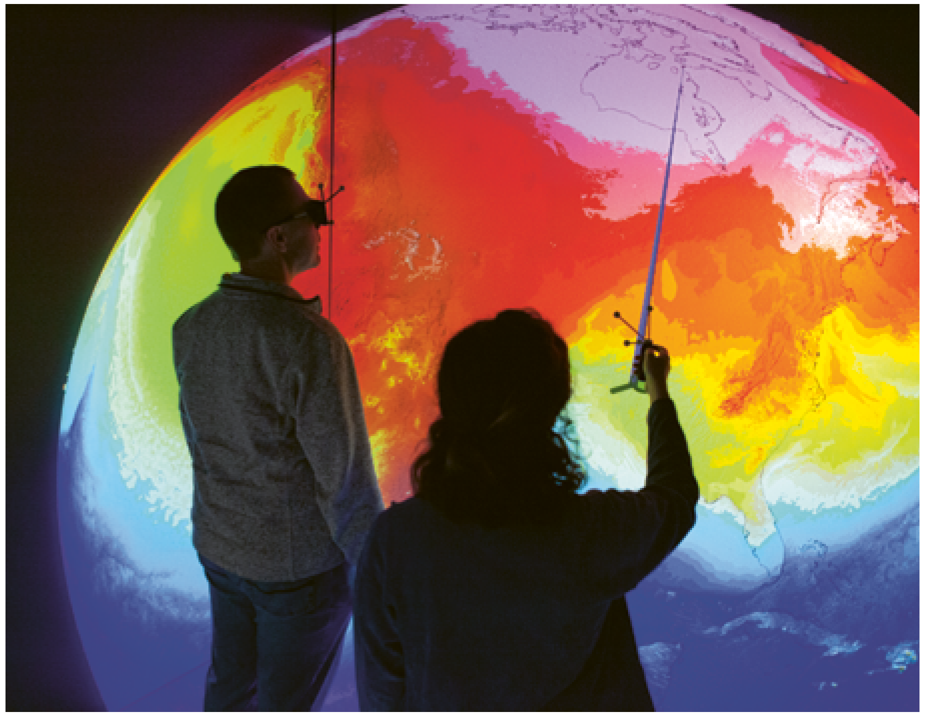

Dr. Schwitalla and HLRS visualization expert Leyla Kern demonstrate a visualization of his high-resolution global model in the HLRS CAVE.

the boundary layer. Another alternative is the time-lagged ensemble, where you could start a second simulation 24 hours later (for example) and compare the results with those of the first. If the results converge, one can be more confident that the model is trustworthy.

Running ensemble models, however, is very computationally intensive. If I wanted to compute an ensemble of 10 individual models, I would need 10 times the compute time and the resulting data set would be 10 times as large as running a single simulation. Computing power is therefore still a limiting factor. Because of this we are going to see a greater need for GPUs, because they are well suited for floating point operations and much faster than GPUs. Many global models now use a combination of GPUs and GPUs, and in the coming years we can expect to see more mixed systems that contain both GPU and GPU components.

This transformation also has consequences for how simulation software is programmed. One can't just simply port the same existing code directly onto GPUs. It must be adapted to be able to run on GPUs. Currently there are a number of approaches for doing this, including OpenACC, CUDA, and OpenMP offloading. This means that even scientists who are experienced in using HPC will need to adopt these new methods. It is great to see that HLRS has already begun addressing this need in its HPC training program. *CW*

News Briefs

In collaboration with the German Federal Institute of Population Research, HLRS's Dr. Ralf Schneider led efforts to implement an ICU occupancy prediction tool.

HLRS and BiB Win HPC Innovation Award for Pandemic Monitoring Tool

In November, HLRS and the German Federal Institute of Population Research (Bundesinstitut für Bevölkerungsforschung, BiB) were named winners in the 18th HPC Innovation Awards. The award recognizes a simulation conceived and led by HLRS and the BiB that uses an automated, agent-based approach to track the progression of infection and predict demand for intensive care units (ICU) across Germany. The predictions are delivered weekly to the Federal Ministry of Health and Federal Ministry of the Interior, where they support political decision making. "HLRS's collaboration with the BiB was the first time that a federal high-performance computing center in Germany provided such a service to the government. It demonstrates how HPC for global systems science can support decision makers in crisis situations," said Dr.-Ing. Ralf Schneider, who contributed to the model development. Among other systems, the simulation was developed on a system donated to HLRS as part of the AMD COVID-19 High-Performance Computing Fund and was implemented on HLRS's flagship supercomputer, Hawk. Coordinated by HPC market analysts Hyperion Research, the HPC Innovation Awards recognize noteworthy achievements by users of high-performance computing across the world. *CW*

HLRS Recertified for Environmental Management

Following a successful, comprehensive review by an external environmental auditor in early November, HLRS was recertified under the Eco-Management and Audit Scheme (EMAS). Developed by the European Union, EMAS is the most demanding environmental management framework worldwide, specifying strict standards to improve environmental performance. In 2020 HLRS became the first supercomputing center to be certified under EMAS. This distinction reflects its implementation of a comprehensive energy and environmental management plan that is now used to guide its activities — from the operation of its supercomputer and cooling systems, to supply chain management and waste management, to preservation of species diversity on its campus, and more. One requirement of EMAS certification is a commitment not just to fulfill legal requirements, but also to strive continually for improvement in environmental performance. To verify this, HLRS undergoes annual environmental audits and completes a more extensive recertification process every three years. With projects including ENRICH, DEGREE, SRI DiTEnS, and EE-HPC, the center also continues to investigate new strategies for improving energy efficiency. These efforts both benefit the operation of its own systems and are helping to develop strategies for reducing environmental impacts in other data centers. *CW*

HLRS Director Prof. Michael Resch (right) welcomed Prof. Wolfram Ressel (Rector, University of Stuttgart) and Petra Olschowski (Minister for Science, Research, and Art of the State of Baden-Württemberg) in the HLRS high-performance computing facility.

New Science Minister Visits the University of Stuttgart

During her first tour of the University since beginning her term as Baden-Württemberg Minister for Science, Research, and the Arts, Petra Olschowski received a tour of the HLRS supercomputer and CAVE visualization facility. During the visit, she learned about HLRS's contributions to fighting the COVID-19 pandemic, efforts to improve its capabilities in crisis computing, advances in climate simulation using HLRS's computing resources, and the use of digital twins in cultural fields. In a press release following the visit, she said, "What I saw today at my alma mater was quite impressive and a fine example of what the University of Stuttgart stands for: Cutting-edge research, supercomputing with international visibility, and industry collaborations that begin at the undergraduate level. I was impressed by the demonstration in the High-Performance Computing Center, which sets Baden-Württemberg apart as a leading digital region. The climate simulation there showed that digitalization can – and will – help us better anticipate and respond to the consequences of human-induced climate change." *CW*

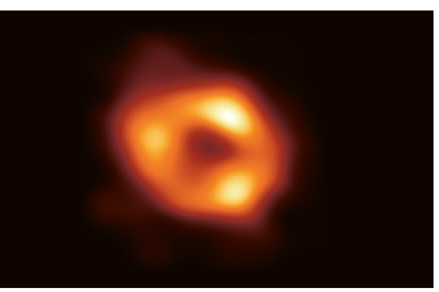

Image: Event Horizon Telescope Collaboration.

Hawk Helps Create First Image of Black Hole at Center of Milky Way

In a blockbuster announcement in May, an international, collaborative team of scientists called the Event Horizon Telescope (EHT) Collaboration revealed an image that for the first time proved the existence of a black hole at the center of our galaxy. Among the contributors to the EHT Consortium was Prof. Dr. Luciano Rezzolla of the Goethe University Frankfurt, whose team used HPC resources at HLRS and the Leibniz Supercomputing Centre to create high-resolution simulations that solved the equations of general-relativistic magnetohydrodynamics (MHD) and of radiative transfer. Their results were then compared with observational data produced by some of the world's leading astronomical observation facilities to help create the image. Rezzolla credited Hawk and the growing power of the Gauss Centre for Supercomputing's infrastructure for enabling his team to run more models than was possible in the past, ultimately resulting in faster analysis. The new announcement followed previous work also done, in part, at HLRS to produce the first image of a black hole ever created in 2019. *EG*

Could Artificial Intelligence Replace Mozart?

In a cooperative research effort facilitated by the Media Solution Center Baden-Württemberg, scientists at HLRS have been working alongside the Stuttgart Chamber Orchestra, the Hochschule der Medien, and the Hertz-Lab at the Center for Art and Media (ZKM) in Karlsruhe to explore a provocative question: Could an algorithm generate a unique composition that is indistinguishable in style from something Mozart might have produced? A panel discussion and concert at the Stuttgart Public Library on October 18 presented preliminary results of this project. SKO General and Artistic Director Markus Korselt, Hertz Laboratory leader Ludger Brümmer, and HLRS AI expert Dennis Hoppe considered questions concerning how artificial intelligence is affecting music now, and how its role might evolve in the future. Following the discussion, the audience experienced the first results of this machine-human collaboration, when members of the SKO performed several compositions generated by the team's experiments. The event took place as part of a series titled "Questions for Our Colleague, AI," which is sponsored by the German Federal Ministry of Education and Research (BMBF). HLRS is participating in the research effort under the auspices of the CATALYST project. *CW*

A concert and panel discussion at the Stuttgart Public Library considered the potential of AI for music composition. Panelists (l–r): Felix Heidenreich (IZKT), Markus Korselt (Stuttgart Chamber Orchestra), Dennis Hoppe (HLRS), Ludger Brümmer (Hertz-Lab, ZKM). Photos: Kai Loges, © die arge lola.

Dr. Rolf Rabenseifner led HLRS's HPC training program until his recent retirement.

Colloquium Honors Dr. Rolf Rabenseifner

For the thousands of scientists who have turned to HLRS for help in developing their high-performance computing skills, Dr. Rolf Rabenseifner is no stranger. Beginning in 1998, Rabenseifner has cultivated the growth of HLRS's HPC training activities into Europe's largest and most experienced program of its kind. His accomplishments have included overseeing the evolution of HLRS's course curriculum, initiating "train the trainer" programs that have amplified HLRS's expertise across Europe, and serving on the steering committee of the MPI Forum, the standardization body for this widely used parallel programming framework. Rabenseifner recently retired from organizational work, but on May 13, 2022 he returned to HLRS as colleagues from as far away as Hawaii gathered for an honorary colloquium celebrating his numerous contributions to the HPC community. Rabenseifner continues to give training courses on parallel programming at many locations. *CW*

OpenBikeSensor Wins German Cycling Award

OpenBikeSensor (OBS), a technology initiated by HLRS research scientist Thomas Obst, was named a co-winner of the 2022 German Cycling Award in the Service & Communication category. OBS is a bicycle-mounted device that uses GPS to continuously track a cyclist's location and his or her distance from nearby hazards such as moving cars. Data from these measurements can be uploaded to a community portal, where software integrates feedback from many contributors to produce a data-based map of cycling routes through a city. This citizen science approach reveals locations where the typical distance between cyclists and cars is dangerously small, providing information that cycling advocates and city planners can use to identify spots where additional barriers, signage, or road markings are needed to make cycling safe. In October, HLRS partnered with the nearby city of Herrenberg in such an experiment, and OBS has also been used in an HLRS project called Cape Reviso, which is demonstrating how urban digital twins could help to improve highly trafficked urban locations. OpenBikeSensor is conceived as an open source project, and interest has been growing across Germany. Plans, including the complete source code, are available at www.openbikesensor.org, making it possible for anyone with some basic programming and electronics knowledge to build his or her own kit. *CW*

Image: OpenBikeSensor

HLRS Hosts HPC User Forum

On October 6–7, the High-Performance Computing Center Stuttgart once again hosted the HPC User Forum. Organized by Hyperion Research and HLRS, the HPC User Forum brings together senior representatives of key HPC initiatives, internationally prominent high-performance computing centers, leading technology manufacturers, and other experts on new HPC, AI, and quantum computing technologies and applications. The event offered a survey of global HPC market dynamics, as well as talks and discussions on topics with broad impact, including perspectives on European and German HPC strategy, the latest developments at HLRS, the Jülich Supercomputing Centre, and the Leibniz Centre for Supercomputing; perspectives on HPC usage in industry; and HPC applications in the arts and culture industries. Additional topics of discussion included lessons being learned in the building of a balanced exascale system at Argonne National Laboratory and an update on the new system at the King Abdullah University of Science and Technology (KAUST); a survey of the state of the art in quantum computing around the world; and new HPC products and technologies. *CW*

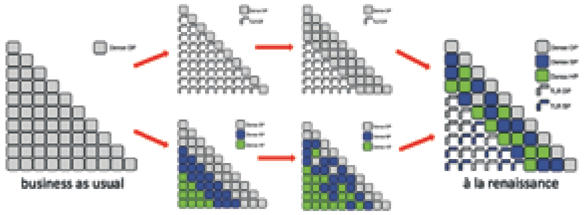

Journey from a global dense double-precision representation of a covariance matrix (left) to a hybrid compressed representation. Image: KAUST

Gordon Bell Award Finalists Develop Method for More Efficient Computing

A team of users of HLRS's Hawk supercomputer was among the finalists for the Association for Computing Machinery's 2022 Gordon Bell Prize. Led by Dr. David Keyes, director of the KAUST Extreme Computing Research Center at King Abdullah University of Science and Technology (KAUST) in Saudi Arabia, the researchers have been developing methods to increase computing performance by distinguishing between components within large simulations that require high precision, and those that can be calculated in a less precise manner. The scientists turned to HLRS's Hawk supercomputer to prove their method's effectiveness and scalability before completing the run on the world's second largest supercomputer – Fugaku at Japan's RIKEN Center for Computational Science – cited in the list of award finalists. Through its work, the team found that it could get a 12-fold performance improvement against traditional state-of-the-art dense matrix calculations, and considering the large energy footprints of high-performance computers, the KAUST team's algorithmic innovations could one day complement efforts to improve HPC hardware efficiency. Although their Gordon Bell-nominated work focused on large-scale climate simulations, the team has extended its approach to signal-decoding in wireless telecommunications, adaptive optics in terrestrial telescopes, subsurface imaging, and genotype-to-phenotype associations, and plans to extend it to materials science in the future. The Gordon Bell Prize recognizes the year's most innovative work in computer science. *EG*

Media Solution Center (MSC) Joins EU Initiative to Promote Innovation in the Cultural and Creative Sectors

A consortium called Innovation by Creative Economy (ICE) is one of 50 partners selected to join EIT Culture & Creativity, an EU-wide Knowledge and Innovation Community (KIC) announced by the European Institute of Innovation & Technology (EIT) in late June. EIT Culture & Creativity is the newest of nine European KIC's, which work strategically to promote the development of new products and services, new companies, and training opportunities in domains that are important for European economic development. ICE was founded through the collaboration of arts, culture, and business organizations across Germany, including the Media Solution Center Baden-Württemberg (MSC). The acceptance of ICE into the EIT Culture & Creativity general assembly means that the Media Solution Center, launched in 2018 by HLRS together with the Hochschule der Medien and Center for Art and Media (ZKM), will contribute to the development of this exciting new initiative at the national and EU levels. Matthias Hauser, General Manager of the MSC, is a co-chair within ICE, serving as chairman of a European platform called The Next Renaissance. As a founding partner of the Media Solution Center, the High-Performance Computing Center Stuttgart is the only European supercomputing center participating in EIT Culture & Creativity. *CW*

A digital twin of the City of Stuttgart visualizes air pollution near City Hall.

Stuttgart Mayor Dr. Frank Nopper Tours HLRS

The visit on September 5 highlighted Stuttgart as a center for science and technology, and how HLRS's activities and resources can support city planning and management. In a press release issued by the city following the visit, Nopper called HLRS an "exemplary research center for the future," adding, "Internationally recognized research takes place here in Stuttgart at the High-Performance Computing Center. Its high-performance computers support research in the natural sciences and engineering in our region and far beyond its borders." During the visit, scientists at HLRS introduced Nopper to projects addressing challenges facing the city and region. These included Cape Reviso, which is using virtual reality and other methods to address dangerous traffic situations, and Open Forecast, which developed methods for visualizing regional climate and air pollution models. Accompanying Dr. Nopper were senior representatives of the City of Stuttgart department responsible for coordinating Stuttgart 21, the new Rosenstein Quarter, and future city planning efforts. *CW*

Winners of the 2022 Golden Spike Awards (l–r): Anna Neuweiler (University of Potsdam), Martin P. Lautenschläger (German Aerospace Center), Johanna Potyka (University of Stuttgart).

Golden Spike Awards Presented at 25th Annual HLRS Results and Review Workshop

On October 4–5, 2022, scientific users of HLRS's high-performance computing systems gathered at the center and online to present their recent research at the annual Results and Review Workshop. This meeting offers an opportunity to learn about current applications of high-performance computing at HLRS and to discuss approaches for using the center's supercomputer most effectively. Twenty-four talks and a lively poster session covered a wide range of topics in fields such as computational fluid dynamics, climate research, computer science, chemistry and materials science, bioinformatics, structural mechanics, and physics, among others. At the conclusion of the meeting, the winners were named of this year's HLRS Golden Spike Awards, which recognize excellence in research and the usage of high-performance computing systems. Representing their respective projects, this year's Golden Spike Award winners were Anna Neuweiler (University of Potsdam) for "Simulating Binary Neutron Star Mergers," Martin P. Lautenschläger (Institute of Engineering Thermodynamics, Computational Electrochemistry, German Aerospace Center, Ulm) for "Lattice Boltzmann Simulation of Flow, Transport, and Reactions in Battery Components," and Johanna Potyka (Institute of Aerospace Thermodynamics, University of Stuttgart) for "Towards DNS of Droplet-Jet Collisions of Immiscible Liquids with FS3D." *CW*

Collaboration Agreements Renewed with Partner Universities

HLRS is globally networked, maintaining formal collaboration agreements with leading academic centers for high-performance computing and simulation across Europe, Asia, and the Americas. In 2022, the center renewed its collaboration agreements with four longstanding partners: the Ukrainian administration of the Donetsk National Technical University (DNTU) in Ukraine; Tohoku University in Sendai, Japan; the University of Science and Technology of China in Hefei; and the Shanghai Supercomputing Center, China. By extending these partnerships HLRS will continue working with the institutions on scientific issues of common interest, staff exchanges to promote the sharing of knowledge and expertise, and the organization of collaborative meetings and publications. The renewal of the collaboration with the DNTU was concluded just a few days after the invasion of Ukraine by Russian forces, and was announced as part of the University of Stuttgart's pledge of solidarity with scientists and students at academic institutions in the two countries. With approximately 25,000 students, Donetsk National Technical University is one of the largest technical universities in Ukraine. *CW*

News Highlights

EuroCC Project Coordinator
Dr. Bastian Koller welcomed attendees
to the conference, providing an overview
of the project's goals and accomplishments.
Photo: Slaven Vilus

EuroCC and CASTIEL Advance Toward Second Funding Phase

Building on successful efforts to establish and facilitate collaboration among national competence centers for HPC, HLRS organized the projects' first all-hands conference and prepared to start their second funding phase.

Since 2020 HLRS has coordinated EuroCC and CASTIEL, a pair of projects funded by the EuroHPC Joint Undertaking (JU) to maximize the impact of high-performance computing, high-performance data analytics (HPDA), and artificial intelligence (AI) across Europe. EuroCC helped to establish national competence centers (NCCs) for HPC, HPDA, and AI in 33 European nations, each of which completed a national HPC competency audit and has become a central contact point for the HPC communities in their home countries. Simultaneously, CASTIEL has enhanced the impact of EuroCC by implementing common standards and facilitating international collaboration and knowledge exchange across the NCC network. These efforts aim to raise the quality of HPC services continent-wide, including regions where usage of high-performance computing has been slower to develop. CASTIEL has supported numerous training, mentoring, twinning, and workshop activities, and through a web portal called EuroCC Access has created a gateway for representatives of industry, public administration, and academia to find the HPC expertise and resources they seek.

In September, HLRS organized the first EuroCC all-hands conference in Bečići, Montenegro under the motto "Uniting Competencies for a Stronger Europe." Gathering representatives of 32 of the 33 EuroCC nations, the meeting offered an opportunity to reflect on the successes of the projects thus far, share insights and best practices, start new collaborations, and plan for the second phase of the projects, beginning in early 2023.

Throughout the conference, discussion returned repeatedly to the question of what the NCCs can do to promote usage of HPC in new communities. As several speakers explained, this will require that the NCCs conduct outreach, communicate effectively with industry, present relevant success stories that demonstrate the potential value of HPC, and provide the training and user support they will need to use HPC effectively.

In a keynote address, JU Executive Director Anders Dam Jensen discussed the role of EuroCC and CASTIEL in the context of Europe's HPC strategy. "The NCCs are

In a panel discussion, participants reflected on the importance of EuroCC for European HPC strategy. L–R: Natalie Lewandowski (HLRS), Tomas Karasek (IT4I), Daniel Opalka (EuroHPC Joint Undertaking), Espen Flage-Larsen (University of Oslo), Milena Miljonić (Ministry of Science and Technological Development, Montenegro), Guy Lonsdale (scapos AG). Photo: Slaven Vilus

one of our most important projects," he remarked, suggesting that increasing understanding of high-performance computing will increase European technological independence and lead to new applications with EU-wide benefits.

Daniel Opalka, who leads the EuroHPC Joint Undertaking's Research and Innovation Sector, also commented on the unique role of EuroCC in within the overall JU strategy. "The network of national competence centers is … not only one of the largest projects, it's also one of the most comprehensive and inclusive European projects that the JU currently manages," he said. "We reach with the national competence centers an unprecedented number of communities to support the adoption of HPC, in particular in SMEs that are notoriously difficult to approach."

In the next project phase CASTIEL 2 will move beyond working only with the NCCs to coordinate strategic collaboration among the EU-funded Centres of Excellence (CoEs). In contrast to the NCCs, which bundle competences on a regional level, the CoEs gather expertise by sector (such as engineering or biology) and conduct research to adapt important computing codes for future exascale systems.

Integrating the CoEs and CASTIEL 2 will help to ensure that collaboration between CoEs and NCCs happens in a more direct way, and that awareness of the technical achievements of the CoEs is disseminated Europe-wide, increasing their application within research and industry. CASTIEL 2 will also support the release of the CoEs' enhanced HPC, AI, and data analytics applications on the next generation of EuroHPC petascale, pre-exascale, and exascale supercomputers, leveraging the reach of the EuroCC network to ensure that HPC users across Europe can access and profit from them. *CW*

EuroHPC Joint Undertaking Executive Director Anders Dam Jensen provided an overview of European HPC strategy. Photo: Slaven Vilus

A detailed 3D model of the ElbX tunnel in virtual reality makes it easier for engineers, architects, and other specialists to plan all aspects of the structure's design and construction.

Virtual Reality Supports Transition to Renewable Energy

HLRS created a digital twin of the ElbX tunnel, a critical component of SuedLink that will transport green energy to southern Germany.

HLRS has been a partner in a major engineering project that is supporting Germany's energy transformation. Working with energy transmission system operator TenneT and architectural firm Kieferle & Benk, a team led by HLRS Visualization Department Head Dr.-Ing. Uwe Wössner created a digital twin of the so-called ElbX tunnel. The structure is a critical component of SuedLink, which according to project leaders TenneT and TransnetBW will transport 4 Gigawatts of power underground over a distance of nearly 700 kilometers from northern Germany to the south. The ElbX tunnel, the largest special structure in the SuedLink project, will cross under the Elbe with a length of more than 5 kilometers.

The digital twin of ElbX developed at HLRS integrates computer aided design (CAD), BIM, Geographic Information Systems (GIS), and simulation results into an immersive environment that can be displayed in HLRS's CAVE visualization facility. Over the course of several months, experts involved in the tunnel's planning visited HLRS. Wearing 3D glasses they can navigate through a detailed, realistic representation of the structure to observe, for example, plans for power transmission line routing, safety infrastructure, lighting, or ventilation. Throughout this process, the realism of the model and the interdisciplinary discussions it enabled improved the project team's ability to understand the project, adhere to the budget and project timeline, and achieve the goal of designing a high-quality structure. The project partners anticipate that even after building begins, the digital twin will continue to be a valuable resource in construction planning and management.

"HLRS's contributions to ElbX demonstrate how useful virtual reality and digital twins can be in planning complex engineering projects," said Wössner. "HLRS has been excited to contribute to this undertaking, and ultimately to support efforts that will help Baden-Württemberg and Germany to meet climate and sustainability goals." *CW*

Community Outreach Demonstrates Uses of HPC for Sustainability and City Planning

In several public events, HLRS met with interested citizens, public officials, city planners, and community leaders to discuss supercomputing and its applications.

The value of high-performance computing tools lies not in the technologies themselves, but in how they are used to help solve society's challenges. In several events in 2022, representatives of HLRS organized and participated in public events to raise awareness of the supercomputing center and understanding of some of these applications.

In March, HLRS's Dr. Karin Blessing led an online seminar in which more than 60 participants from across the region and across Germany learned how digital twins can support planning in cities and communities. "Digital planning is an area of great concern in many areas," Blessing said, "and whether it be developing large-scale city plans, redesigning a public square, or addressing traffic congestion, modern approaches using simulation and virtual reality offer an excellent way to help understand and discuss the potential effects of planning decisions."

Blessing also represented HLRS in the Baden-Württemberg Society's "Society and Nature" roadshow, a traveling event organized in the summer to promote public dialogue concerning nature, species conservation, climate, and sustainability. At a demonstration booth in central locations in five cities across the state, HLRS met local residents and public officials to discuss the supercomputing center's sustainability efforts, and to show how research using HPC is helping the development of more sustainable technologies and communities.

As a participant in the Stuttgart Festival of Science in late June, HLRS presented "Digital Twins in Action" at the Stuttgart City Hall. The demonstration covered a range of topics, including BIM (Building Information Management), privacy-compliant traffic recording, climate simulation, and the use of mobile sensors in traffic simulation.

In September, HLRS also took part in Stuttgart's second "Mobility Week," a city-wide event focused on sustainable transportation. In a booth at the city's vibrant Marienplatz, the center's visualization team presented a digital twin of the location that it has been developing within the project Cape Reviso. Stuttgart Mayor Frank Nopper kicked off the festival with a visit to the HLRS booth, and local residents learned how research methods developed by HLRS scientists could help reduce conflicts between cyclists and pedestrians. *CW*

CASE4Med Will Bring Supercomputing to the Medical Technology Industry

A new Medical Solution Center will improve access to simulation, data analytics, and artificial intelligence resources for developers and manufacturers of medical products.

Simulation and artificial intelligence could revolutionize how Baden-Württemberg's medical technology community develops and tests devices and pharmaceuticals. Image: Marcel Scholte, Unsplash.

The German state of Baden-Württemberg is home to a large and economically robust medical technology community, but has typically made limited use of simulation, machine learning, or artificial intelligence on high-performance computing (HPC) systems. To explore what opportunities they could offer, HLRS this year launched a new Medical Solution Center called CASE4Med. With the support of a five-year grant from the Baden-Württemberg Ministry for Science, Research and the Arts, HLRS – in cooperation with the Innovation and Research Center Tuttlingen of Furtwangen University and SICOS BW GmbH – will build on a model that HLRS has implemented in the past to provide industry-specific HPC solutions.

CASE4Med is creating a network that brings together stakeholders from Baden-Württemberg's medical technology community and experts in the development of IT-based and data-driven approaches for medical applications. Participants will determine what kinds of HPC-based solutions could most benefit the community and what resources and skills are needed to make them a reality. By cultivating contacts and completing pilot projects that demonstrate relevant applications, CASE4Med aims to become a self-sustaining membership organization by the end of the grant period.

As in other engineering fields, simulation could offer medical device companies tools that make the development and testing of new products faster and less expensive. Simulation could be used, for example, to assess the suitability of materials or components for medical instruments and implants, to optimize the design of electronic and software components, to extend medical devices' lifetime performance, or to address safety regulations. New applications of data collection and analysis could also provide insights into production or quality control processes, or support the data-driven development of new products or business models.

Prof. Dr.-Ing. Martin Haimerl, Scientific Director of the Innovation and Research Center Tuttlingen, foresees great potential in CASE4Med. "In the medical technology sector, use of simulation and high-performance computing is uncommon and will need to be built up in a systematic way," he explained. "The collaborative network that the Medical Solution Center plans to build could take medical technology across the state to a new level." *CW*

Air flow in the simulation of a ventilation device with two radial fans. Image: ebm-papst

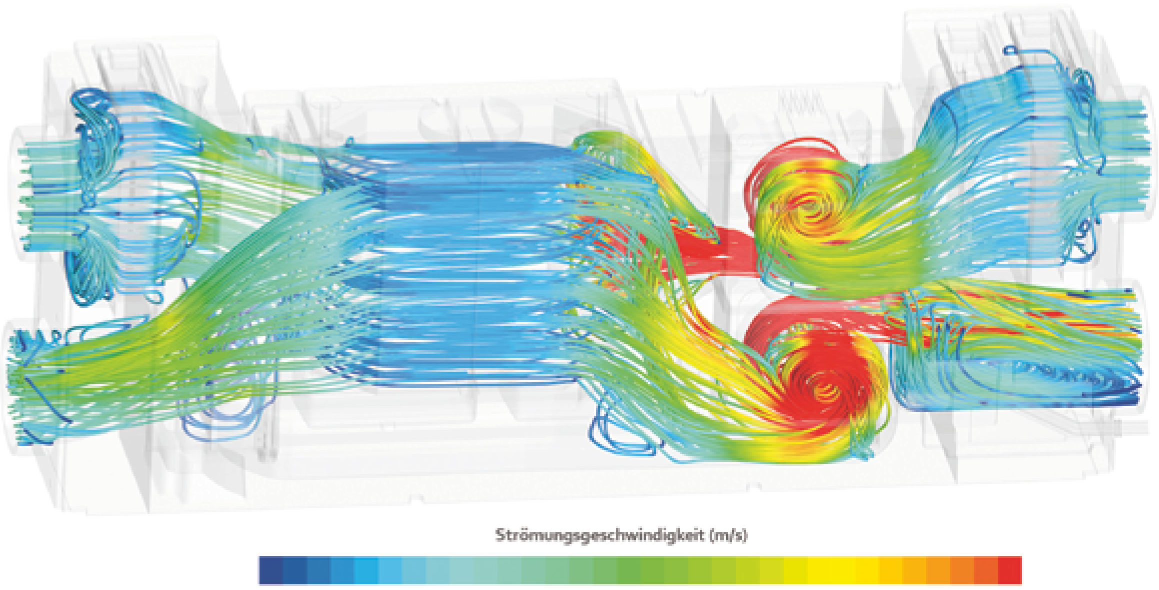

HPC Helps in Design of Quieter Fans and Motors

Engineers at manufacturer ebm-papst run aeroacoustic simulations on HLRS's Hawk supercomputer to accelerate product design.

Nestled in a small town in the scenic Jagst river valley in northern Baden-Württemberg, ebm-papst is one of the world's leading manufacturers of high-quality ventilation fans and electric motors. Its industrial clients include manufacturers of heating, air conditioning, and refrigeration systems, and the company's products can be found worldwide.

Because so many of ebm-papst's products are used in locations that are in close proximity to people, one focus at the company is the development of low-noise fans. As part of this R&D effort, the company has turned to the High-Performance Computing Center Stuttgart (HLRS) for access to high-performance computing (HPC) resources. Using the center's flagship supercomputer, Hawk, ebm-papst engineers run high-resolution aeroacoustic simulations that help them to gain a deeper understanding of the complex mechanisms that give rise to noise.

Supervising this work is Dr. Andreas Lucius, an engineer in ebm-papst's pre-development department who evaluates and develops new computational methods for aeroacoustic simulation. Although the company has used simulation to model aerodynamics on its in-house compute cluster for many years, aeroacoustic simulation has other requirements that can be best addressed using a supercomputer like Hawk.

"We have a small computing cluster that is sufficient for many traditional aerodynamics simulations, but aeroacoustic simulations require more performance," Lucius explained. "We need to use a very fine computational mesh and other kinds of modelling approaches in order to resolve small turbulent structures. This can quickly use up more than 10,000 CPU hours, requires highly paralleled hardware, and takes too long on our own systems. In situations like these having access to HLRS's supercomputer helps us to get our results much faster."

Simulation complements measurement in localizing sources of noise

Before being sold on the market, ebm-papst's fans must undergo rigorous testing. This includes experimental studies conducted in a wind tunnel at the company's Mulfingen headquarters. In addition to performing standardized measurements of noise performance the engineers use an array of highly sensitive microphones, a so-called "acoustic camera", to determine locations on the rotating fan blade where noise is produced. The system is very helpful during product development, but requires that many sophisticated, expensive microphones be placed around the laboratory. In an aeroacoustic simulation, however, little additional work is required to place as many virtual microphones in space as needed. This approach improves the quality of the team's ability to localize sources of noise.

When modeling a spinning fan, the engineers use design data about the machine and computationally subdivide the area surrounding it into a fine-grained mesh. Using commercially available software, they then calculate how air moves in each of those cells over time. Based on physical principles, the simulations help the engineers to identify locations where unwanted turbulence or fluctuations in air pressure might be occurring. In both situations, pressure waves are produced in air that the human ear perceives as sound.

At ebm-papst, aeroacoustic simulation is used in an iterative way as a complement to experiments. When a new fan is in development, simulations can provide early information that make it easier for engineers to arrive at quieter designs more quickly. At the same time, data generated by experiments can be used to validate and improve the precision and trustworthiness of computational models.

This approach has also helped the company in optimizing its standards testing site. During laboratory tests tones are measured together with a fan's blade repetition frequency and multiples thereof. In theory, the generation of tones should not be significant for undisturbed wind tunnel flow. In one case, however, the measurement identified an unexpected sound. When Lucius ran a simulation on Hawk that accounted for the fan and the surrounding laboratory, it clearly identified the source of the noise as a turbulence structure located above the fan that was being produced by the interaction between moving air and the wind tunnel's sound insulation panels. When compared with experimental data captured by the microphone array, the technicians confirmed the results of the simulation, and changed the shape of the insulation panels to eliminate the tones. In the standardized noise performance measurement the level of the tones could be reduced by up to 10 dB and the level of acoustic power by up to 1 dB(A).

For the company, the lower cost of using HLRS's system in comparison to commercial cloud providers is showing economic benefits. "Whenever we conduct an experiment we need to build a prototype, operate the wind tunnel, and make measurements. All of those things cost time and money," Lucius explains. "As the precision of aeroacoustic simulations gets better, this approach has enormous potential to save money and speed up the development process."

Future directions for simulation in engineering

Lucius anticipates that new methods and access to larger and more diverse kinds of HPC systems could open new opportunities for R&D at ebm-papst.

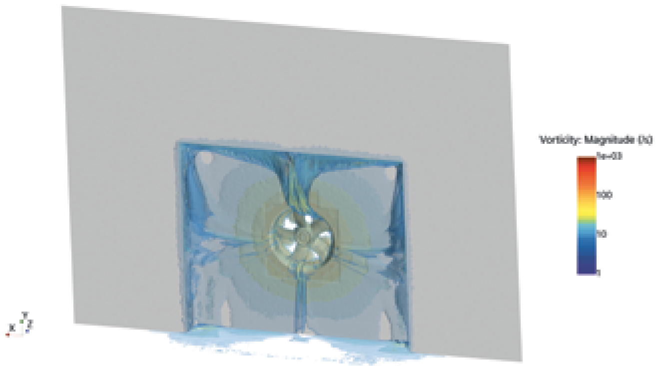

Visualization of flow structures in the wind tunnel using vorticity.
Image: ebm-papst

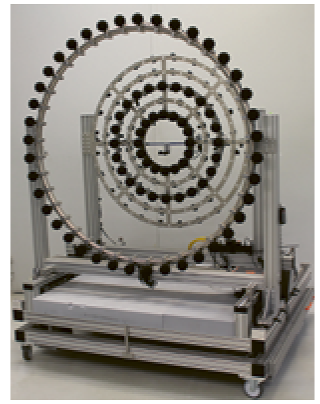

A microphone array used to experimentally localize sources of noise in the wind tunnel at ebm-papst. Photo: ebm-papst

Providing customized solutions for manufacturers is already a hallmark of the company, and when a customer is interested in building one of its fans into a specific air conditioning unit, for example, it is possible to test a prototype experimentally in the wind tunnel. In such complex systems, however, simulation is currently very computationally expensive. “Here we’re no longer talking about 10,000 CPU-hours, but more like 100,000 CPU-hours,” Lucius says. “These are situations that are very interesting and we are currently searching for methods that could reduce computing time even for such complex cases. Our goal is not only to be able to analyze the aeroacoustics of such complex installation cases, but also to optimize them mathematically using many simulation runs.”

Lucius computes primarily on CPUs, the traditional workhorse processors for HPC simulation. With the growing availability of GPU-enabled flow solvers, it will be possible to perform calculations faster in the future. One relevant approach is the lattice Boltzmann method, which is well suited to GPUs because of the processors’ architecture. As such methods become more mature, they could open new opportunities for modeling complex aeroacoustic systems. GPUs also offer the added benefit of using significantly less energy in comparison to traditional CPUs, a factor that is also important for ebm-papst.

In addition, Lucius sees possibilities in the new field of artificial intelligence. “Through our experiments, ebm-papst has accumulated measurements of all kinds, including fluid dynamics, aeroacoustics, electric, and thermal data,” he said. “The question we are considering now is how this data could be leveraged for the future.” Here, too, access to high-performance computing could be helpful, as algorithms using neural networks could perform data analyses that enable faster evaluation of designs during product development. AI could also make it possible to run large numbers of optimization simulations, an approach that is currently computationally expensive using CPU-based systems.

Although it will take time for such methods to develop, these perspectives suggest that high-performance computing will continue to provide critical tools for pioneering companies working at the frontiers of engineering. *CW*

New HLRS Steering Committee Begins Term

A multidisciplinary panel of advisors will help to shape the center’s continuing evolution.

Comprised of members of the center’s user community and other experts in the field of high-performance computing and its applications, the HLRS steering committee provides guidance and advice regarding the center’s activities. This includes overseeing system usage, setting policies concerning computing time allocations, and participating in decision-making in the selection of hardware and software, among other duties.

At its first meeting in October, the steering committee elected Prof. Dr. Thomas Ludwig, Director of the German Climate Computing Center, as the committee chair. Prof. Dr. Peter Bastian, Group Leader for Scientific Computing at the University of Heidelberg, was elected vice chair. They will serve for a three-year term. *CW*

Members of the new HLRS steering committee (l–r):
Andrea Beck, Roland Potthast, Kira Rehfeld, Miriam Schulte, Anita Schöbel, Thomas Ludwig, Andreas Frommer, Peter Bastian, and Volker Wulfmeyer.
(Not pictured: Lars Pastewka, Jörg Schröder, Birgit Strodel).

The following are the members of the new HLRS Steering Committee.

Chair

Prof. Dr. Thomas Ludwig
German Climate Computing Center

Vice Chair

Prof. Dr. Peter Bastian
University of Heidelberg

Steering Committee Members

Prof. Dr. Andrea Beck
University of Stuttgart

Prof. Dr. Andreas Frommer
University of Wuppertal

Prof. Dr. Lars Pastewka
University of Freiburg

Prof. Dr. Roland Potthast
German Weather Service

Prof. Dr. Kira Rehfeld
University of Tübingen

Prof. Dr. Anita Schöbel
Fraunhofer Institute for Industrial Mathematics, ITWM

Prof. Dr.-Ing. habil. Jörg Schröder
University of Duisburg-Essen

Prof. Dr. rer. Nat. Miriam Schulte
University of Stuttgart

Prof. Dr. Birgit Strodel
Heinrich Heine University Düsseldorf

Prof. Dr. Volker Wulfmeyer
University of Hohenheim

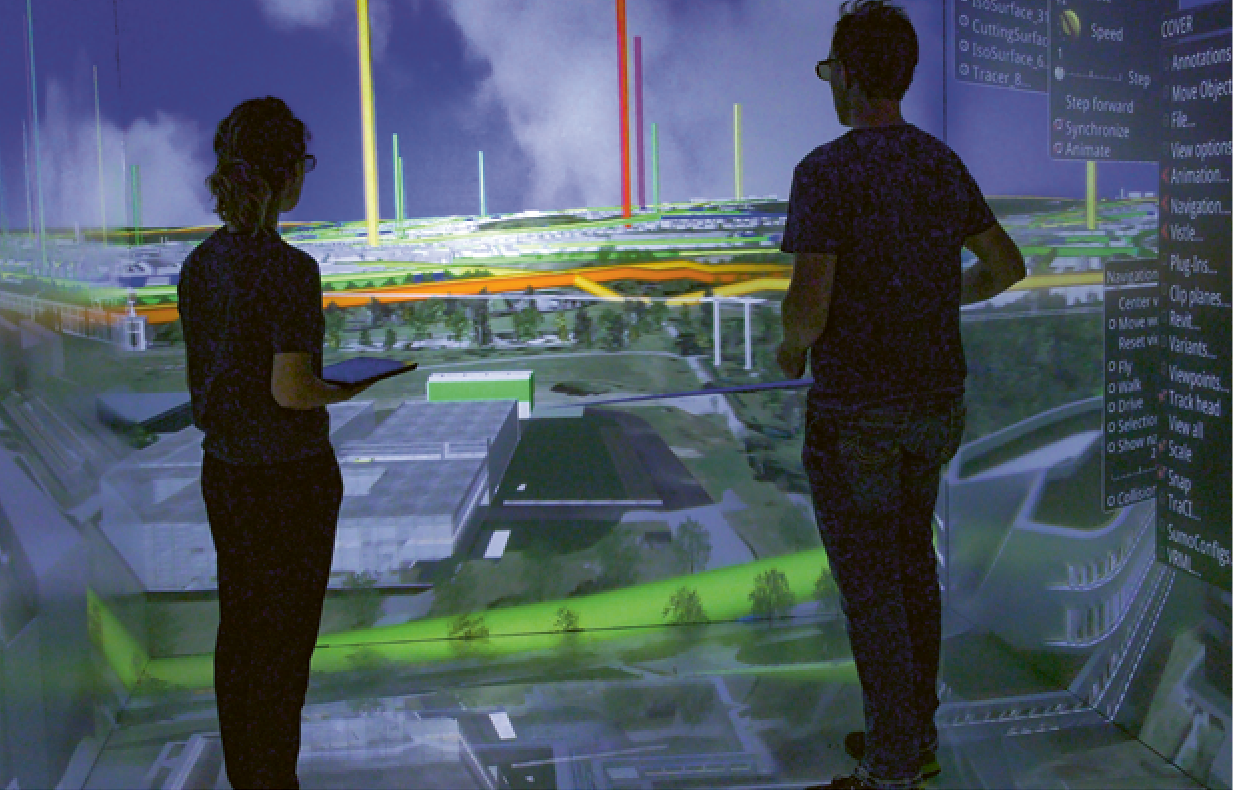

Within the DiTEnS project, the HLRS visualization team will enhance its digital twin of the University of Stuttgart to help identify strategies for making the campus carbon neutral.

Stuttgart Research Initiative Will Support Energy Transformation in Local Communities

As a participant in the Stuttgart Research Initiative DiTEnS, HLRS will use digital twins to improve the processes needed to make urban energy systems climate neutral.

The European Union's Climate Target Plan and Germany's Climate Change Act have set ambitious goals for reducing greenhouse gas emissions in Europe. For these efforts to be successful, one focus needs to be on reducing emissions that result from the heating of buildings, which is responsible for approximately one-fifth of greenhouse gas emissions and consumes about one-third of electricity in Germany.

Increasing efficiency in buildings and shifting to renewable energy sources for modern heating infrastructures are not just technological challenges, but also community-based activities. Many different kinds of stakeholders — including energy providers, homeowners, renters, the building trades, heating technology manufacturers, and others — each have perspectives and interests that need to be considered. For many municipalities and stakeholders, the technical, economic, and social dimensions of this effort pose major challenges.

A new, multidisciplinary research initiative launched by the University of Stuttgart aims to support communities and relevant stakeholders in making this transition. The Stuttgart Research Initiative (SRI) DiTEnS (Discursive Transformation of Energy Systems) will unite HLRS researchers with scientists at the University of Stuttgart to develop methodologies and technologies for planning energy transformation. Using modern algorithms based on artificial intelligence and scientific visualization, it will support discursive processes that bring together stakeholders needed to implement energy transformation at the local level.

SRI DiTEnS is funded in part by a six-year grant from the Carl Zeiss Foundation (CZS) as part of its CZS Breakthroughs Program in Resource Efficiency and Future Energy Systems. The University of Stuttgart will contribute additional funding, and will establish a permanent research initiative and a graduate education program in the field. Prof. Kai Hufendiek of the University of Stuttgart's Institute of Energy Economics and Rational Energy Use (IER) will serve as spokesperson for the project.

In the first steps undertaken in SRI DiTEnS, University of Stuttgart researchers will develop methods for efficiently investigating complex conditions in streetscapes and neighborhoods, and will help to identify customized, climate neutral options for energy transformation in these locations. From the regional level down to individual buildings, the approach will help to determine what opportunities community leaders, property owners, and other stakeholders might have for utilizing renewable energy; improving energy efficiency; making use of intelligent networks, available waste heat, electromobility, and energy storage technologies; and implementing a more flexible approach to managing energy demand.

Scientists in HLRS's Visualization Department led by Dr. Uwe Wössner will use the results of these models to develop digital twins of building and energy systems within their urban contexts. When simulated in virtual reality in a 3D visualization facility like the HLRS CAVE, digital twins make it easier to demonstrate complex systems such as a town's energy and building infrastructure. This can help stakeholders to understand, for example, how changes implemented at a specific location or within a complete system will affect buildings, neighborhoods, or city districts, including their networked operation within the complete system.

With expertise of Center for Interdisciplinary Risk and Innovation Studies (ZIRIUS) Director Prof. Cordula Kropp, DiTEnS will also use perspectives from the social sciences to develop recommendations for a successful energy and heat transition together with the various stakeholder groups. These efforts will be crucial for the success of energy-saving measures. The project will promote interdisciplinary dialogue to reduce conflict among participants with differing requirements and promote shared perspectives.

The University of Stuttgart has set itself the goal of becoming climate neutral by the year 2030. To support this transformation, DiTEnS will begin with a study to comprehensively understand the university's energy and building infrastructure, and to recommend potential steps that could be taken. Among the questions the investigators will consider will be how excess heat generated by HLRS's supercomputer could be reused to heat other buildings.

Using the knowledge gained, DiTEnS will also conduct additional case studies, working with communities surrounding Stuttgart and across the state of Baden-Württemberg to investigate other strategies for sustainable transformation of urban energy systems, and for using digital twins within participatory planning processes. Ultimately, the researchers look forward to helping to accelerate Germany's energy transformation and enhance its ability to realize its climate and sustainability goals. *CW*

Combining HPC with Computers in Wind Farms

A new research project called WindHPC will test strategies for making high-performance computing systems, software, and workflows more energy efficient.

In partnership with several German academic research centers and WestfalenWIND IT GmbH & Co KG | windCORES, HLRS is coordinating a new research project that, for the first time, will combine computing infrastructure located at wind energy generation sites with a high-performance computing center. The project aims both to use surplus energy produced at wind farms efficiently and to increase the amount of green energy used in computationally demanding research.

WindHPC will pursue a holistic strategy focusing both on hardware and other elements of the problem-solving process that affect energy efficiency in simulations. This means looking closely at how compute tasks are assigned within a distributed computing architecture, how data resulting from a simulation are managed, and how simulation algorithms are chosen. At all of these levels, WindHPC will monitor power consumption and performance metrics as a basis for cost-benefit analyses that could improve sustainability in HPC.

Computer scientists at HLRS will optimize the workflows necessary to move and store data when distributing a simulation across networks, such as one combining an HPC center with compute clusters at wind parks. Working with project partners, they will also develop new methods for using performance data to guide the auto-tuning of algorithms and intelligent scheduling of a simulation, in order to reduce power consumption. At the cluster level, WindHPC will study system life cycle management and the effect of fluctuations in power production capacities resulting from, for example, changing wind conditions.

As use cases, the project will focus on HPC applications for process engineering in the chemical industry. These applications will be developed into digital twins, and the WindHPC team will investigate the power consumed in the visualization of scientific results. By setting performance benchmarks and conducting cost-benefit analyses of simulation methods, researchers also aim to better understand when to use an approach called “approximate computing,” which balances the need for precise scientific results with the demand of using as little energy as necessary. The results could help to begin answering important questions facing the future of HPC: Is the knowledge acquired from certain kinds of simulations commensurate with the energy they consume? In what situations could smaller-scale, less precise, and less energy-intensive algorithms and simulations still deliver scientists the information they need?

WindHPC is funded by the German Federal Ministry of Education and Research (BMBF) as part of its Green-HPC initiative. *CW*

WindHPC will for the first time connect computers located at a wind park with an HPC center.
Image: iStock.com / ezypix

The spread of misinformation, disinformation, and malinformation has created a widespread sense of what scientists call epistemic uncertainty.

Building Trust in the Face of Disinformation

A conference organized by HLRS explored the origins and nature of disinformation, its effects on public opinion, and potential strategies for fighting against it.

The tendency of individuals to believe and act on bad information is nothing new. The rise of social media and its effects on the polarization of public opinion, however, have made the need to understand why this happens more urgent than ever. With support from the Baden-Württemberg Ministry of Science, Research and the Arts, the HLRS Department of Philosophy of Computational Sciences has been researching how to assess and improve the trustworthiness of information. In a three-day, international conference titled “Trust & Disinformation” HLRS invited researchers to Stuttgart to explore how disinformation – together with other technological, sociological, institutional, and political factors – can both lead to mistaken trust in false information and damage the shared sense of trust necessary for society to function.

As several conference speakers pointed out, an increasingly complex information landscape has resulted not only in a proliferation of false information, but also in a crisis of trust. Confusion and disagreement concerning fundamental facts make public conversation and decision making difficult. Moreover, speakers explained, the tendency of an individual to believe disinformation is not purely a result of social media algorithms, but results from a complex set of factors including individual beliefs and cognitive processes, social relationships, education, and even pattern-seeking functions in the brain. Countering disinformation successfully will require a more refined understanding of how these and other factors interact.

Several speakers at the conference also discussed the potential advantages and feasibility of measures for fighting the spread of disinformation. Some proposals include media literacy campaigns, web plugins that rate trustworthiness or facilitate fact checking, algorithms for deleting disinformation from the web, or even crowdsourcing alerts of disinformation. Although such measures sound attractive and some might even be technically feasible, conference participants repeatedly returned to a critical problem that calls their desirability into question: Who should have the authority to distinguish “good” information from “bad” information? This political question raises a host of issues concerning regulation of individual rights and free will, and presents new challenges with respect to legitimacy and trust. *CW*

HLRS by Numbers

138 Staff

2 Visiting Researchers

99 Scientists

34 Nonscientists

2 Student Assistants

3 Research Assistants

57 Talks by HLRS Staff

System Usage

Correction, system usage (November 2023): In 2022, 2.887 billion core-hours were approved for large-scale projects on HLRS systems. A total of 4.364 billion core-hours were utilized by GCS and PRACE projects. The total usage, including academic and industry users was 4.639 billion core-hours.

2.887 billion
Core Hours Produced

130
User Projects

66
Industrial Customers

176
User Publications

Staff Publications

26 Papers in Journals, Books, and Conference Proceedings

4 Books

Education and Training

54 Continuing Education Courses	8 Scientific Workshops	17 University Lectures
1,416 Participants 140 Course-Days	391 Participants 13 Days	963 Participants 37 SWS

718
Vistors at HLRS

Third-Party Funds

6,534,118 €

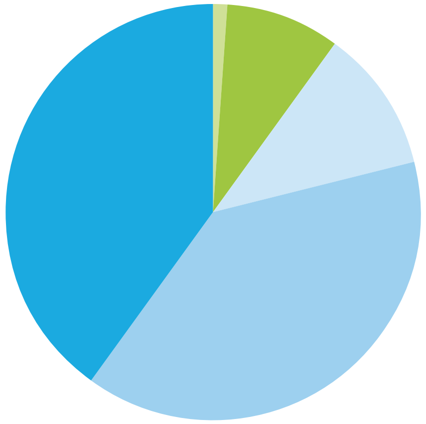

- EU 40 %
- State Gov. 39 %
- Federal Gov. 11 %
- DFG 9 %
- Industry 1 %

Simulation of a neutron-star–black-hole coalescence in which the neutron star is tidally disrupted during the merger. Image: T. Dietrich (Potsdam University and Max Planck Institute for Gravitational Physics), N. Fischer, S. Ossokine, H. Pfeiffer (Max Planck Institute for Gravitational Physics), S.V. Chaurasia (Stockholm University), T. Vu

User Research

When Stars Collide

With the help of new observational data of gravitational waves and electromagnetic signatures, University of Potsdam researchers are using supercomputers to understand binary neutron star mergers.

In 2015, astrophysicists, astronomers, and astrophiles celebrated an exciting development. For roughly 100 years, researchers had hypothesized the existence of gravitational waves — waves of gravity in space-time caused by large, violent events in the cosmos such as supernovas or neutron star mergers — but had never seen direct evidence confirming their existence. When the Laser Interferometer Gravitational-Wave Observatory (LIGO) in the United States definitively detected a gravitational wave event, researchers set out to build on this finding to advance astrophysics research.

"Since 2015, we've seen about 100 gravitational wave events. This kind of research is extremely new, and it has a huge potential for additional inputs that we can't currently get through other observations," said Prof. Tim Dietrich, researcher at the University of Potsdam. Since that time, Dietrich and his research group have been using high-performance computing (HPC) resources to simulate cosmic phenomena that produce gravitational waves. Specifically, the team has been using the High-Performance Computing Center Stuttgart's Hawk supercomputer to simulate what happens when binary neutron stars collide. Over the last three years, the team has made significant strides modelling these complex celestial events in unprecedented detail, and simulations on Hawk have contributed to more than a dozen scientific journal articles in the process, including publications in the leading journals *Nature* and *Science*.

Come together

Neutron stars are essentially the fossil relics of massive stars that have reached their ends. When stars run out of fuel, they start to collapse before their outer layers expand explosively outward in a supernova. These massive events not only produce gravitational waves, but also project heavy elements and other materials across the universe. The remaining material cools (relatively speaking) to a brisk 1,000 degrees Celsius and further consolidates, becoming an ultra-dense neutron star. (The name comes from the fact that after a supernova, heavy neutrons comprise the bulk of the remaining materials.) If two of these objects drift too close to one another, strong gravitational pull causes them to merge, forming a much larger neutron star or creating a black hole in the process.

Astrophysicists can detect neutron star mergers through their observational signatures, such as gravitational waves. To complement these methods, scientists also simulate these events using supercomputers. This makes it possible to understand at a fundamental level how these events produce gravitational waves and electromagnetic signals, and eject materials across the universe.

To do that, though, researchers need to have a reliable model that can accurately represent the complex physics interactions taking place at a wide variety of scales within these massive systems. This requires world-leading HPC resources, and even today's most powerful machines cannot completely simulate these events from first principles. For Dietrich and his collaborators, this has meant finding ways to improve computational efficiency without sacrificing realistic physics in their simulations.

For the team, simulating neutron star mergers realistically means including so-called multi-messenger

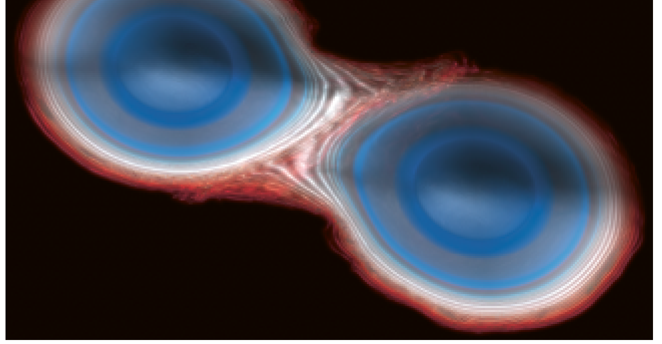

Simulation of two merging neutron stars, each with a mass of 1.35 solar masses. From red to blue, increasing densities are shown. Image credit: T. Dietrich (Potsdam University and Max Planck Institute for Gravitational Physics).

physics information. As the name implies, multi-messenger physics collates information describing multiple physical phenomena to get a more comprehensive picture of materials' behaviors at a fundamental level. Measurements of features including photons (light), a mysterious class of elementary particles called neutrinos, high-energy cosmic rays, and gravitational waves provide valuable, detailed information for researchers at both small and large scales, but are very difficult to integrate into a single simulation that accurately represents the entire system. "We need to perform 5,000 operations for the evolution of a single point in our computational grid," said Anna Neuweiler, a PhD candidate in Dietrich's group and collaborator on the project. "Of course, our grid is comprised of many points, so for even just one time evolution, we need a lot of capacity to compute and solve our equations."

Using Hawk, the team has been able to get closer to an accurate simulation of neutron star mergers based on first principles by selectively lowering the resolution of portions of its simulation that are less pertinent to the research. In addition, the researchers compare multi-messenger physics data in their simulations with complementary heavy-ion collision experiments being run at specialized experimental facilities on Earth. This combined approach has enabled the team both to advance the state of the art in researching binary neutron star mergers and to create a reliable application that they hope to enrich with even more first-principles physics calculations in the years to come. While the team has run larger simulations on other systems, access to Hawk laid the foundations for their successful simulation approach in the recent past.

"I don't point to one big accomplishment in our work, because all these developments aim at getting a better understanding of the physics. This means that even simulations that move incrementally are still necessary, and might become even more important down the line. It is more like a marathon than a sprint," Dietrich said.

Turbulence ahead

Having successfully improved its code's computational efficiency, the team is now focused on ways to include even more detail in its simulations. As part of her PhD research, Neuweiler has begun including first-principles magnetic field calculations in the team's code, leading to a significant increase in computational demands. "Understanding the role that magnetic fields play is mostly important for simulating what happens after neutron stars merge so that we have a more accurate description of how matter flows," she said. "We would like to gain a more accurate description, and in principle we have additional equations and variables that we can use, but it will be more computationally expensive than what we are currently doing."

Dietrich also indicated that in the future the team would like to include accurate descriptions of turbulence at the smallest scales of their simulations, as well as details about neutrino physics that are becoming available as astrophysicists learn more about these mysterious particles. Researchers are also looking forward to the next predicted binary neutron star merger observation run in May 2023, and another one in 2026. With each new event the team will gain access to valuable observational data, enabling them to refine their simulations further. "There will be a lot of instances where we can use our simulations to interpret things better, and we need the resources to do the analysis of the computational data," Dietrich said. "So, one thing is for sure – we will definitely not be asking for less computational time moving forward." *EG*

Simulation for Better Batteries

Computational methods originally developed for hydrology research and oil and gas extraction find new applications in understanding electrolyte flow physics in battery cells.

Since the Industrial Revolution, scientists and engineers have spent significant effort trying to understand how fluids behave in complex, porous spaces. To identify risks for groundwater pollution or to develop more efficient extraction methods for oil and natural gas, for example, researchers have had to determine how fluids behave beneath the Earth's surface. Because of the difficulty in directly observing subterranean interactions, high-performance computing (HPC) simulations play an important role in understanding these complex environments, and the computational methods have continued to improve through the decades.

Recently, researchers at the German Aerospace Center (DLR) in Ulm have begun to apply such approaches in an entirely new context: Specifically, they are using computational methods well-known from hydrology research to better understand how highly porous battery materials and cells can be filled by an electrolyte – a solution containing ions that carry an electric charge. By refining this method, researchers are gaining a better understanding of real-world conditions inside battery cells that will enable them to refine current manufacturing methods and develop new battery designs.

To that end, the team has been using HPC resources at the High-Performance Computing Center Stuttgart (HLRS). With help from the center's flagship supercomputer, Hawk, it runs computationally demanding lattice Boltzmann simulations to study the complex fluid-solid interactions taking place inside a battery cell.

"Battery materials are quite porous, and a lot of the physics interactions happen at the surfaces of the pores at the interface between electrolyte and surrounding solid material," said Dr.-Ing. Martin Lautenschläger, a scientist at the DLR Institute of Engineering Thermodynamics and principal investigator on the project. "Because of the complex physics, geometries, and the variety of scales we have to account for here, there aren't many methods that can tackle this problem. However, the lattice Boltzmann method can and we're among the first to apply this method in this context for real battery materials."

Using Hawk, the team has been able to modify the traditional lattice Boltzmann method, improving computational efficiency in the process. The team has accurately and efficiently simulated multiphase fluid flow simultaneously in pores that range from nanometers to micrometers, which is relevant for batteries. It published its results in the journal *Advances in Water Resources* and currently has its sights set on further developing the method to also model complex chemical and electrochemical reactions taking place inside of a battery cell.

Selective simplification

Researchers face two significant challenges when simulating fluids in motion. The first is scale, as investigators must simulate both a volume that is large enough to represent a real-world system and small-scale interactions that can influence how the fluid as a whole behaves. The second is complexity. When simulating oil extraction or the filling of a battery with electrolyte, researchers do not just simulate a liquid, but rather the interaction between the liquid and the air or other gas that occupies the pores before filling. This multiphase

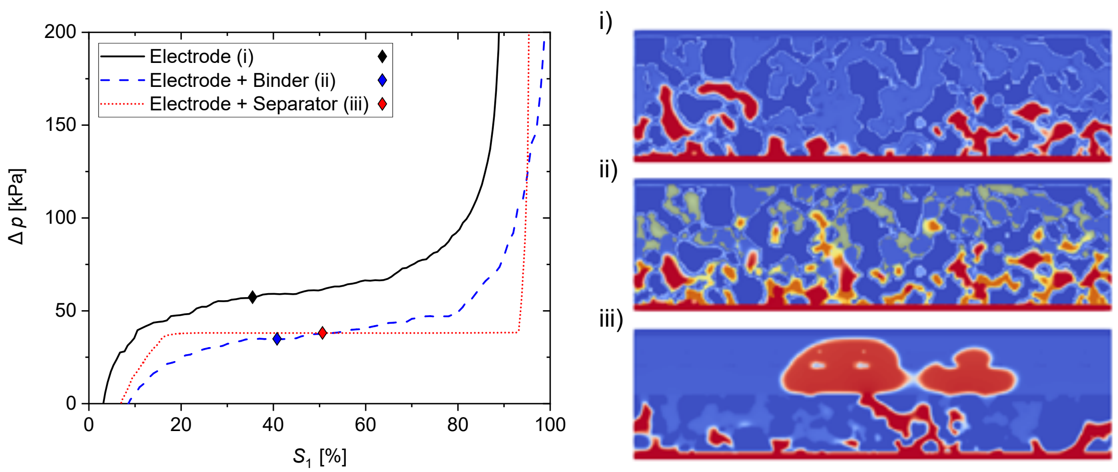

Electrolyte filling of different electrode structures. The pressure-saturation behavior (left) is shown together with snapshots of cross sections (right) in which the electrode is depicted dark blue, gas is depicted blue, the electrolyte is depicted red, and the binder is depicted yellow. Image: Martin Lautenschläger

simulation is further complicated by the need to account for other structurally complex battery components, which often consist of a mixture of materials.

Because of the physics and geometries involved, the lattice Boltzmann method shows advantages over standard computational fluid dynamics approaches. It is based on the so-called Boltzmann equation, and treats fluids as a large collection of particles on a computational grid, or lattice. In contrast to conventional methods that are based on the equations of fluid mechanics, the lattice Boltzmann method solves numerically much simpler equations, offering a more computationally practical approach for addressing these challenges in a way that is especially favorable for HPC.

In a battery cell, the negative and positive poles called electrodes — the anode and cathode — are separated by an electrically insulating porous interlayer that prevents their direct contact. However, this separator and other materials used in batteries contain tiny, nanoscale pores that are ideally filled with electrolyte and, in turn, partially enable ion transport. When running lattice Boltzmann simulations, researchers must pay special attention to how the electrolyte behaves when passing through different pores that vary in shape and size.

The team adapted a so-called "homogenization" approach from conventional fluid dynamics and transferred it to the lattice Boltzmann method for applications of multiphase flow. Here, in a computationally efficient manner, they realistically simulate how electrolytes behave when passing through both the larger mesoscale pores and the tiny nanoscale pores, but simplify their calculations. To reduce the computational efforts, only the electrolyte flow through the larger pores is fully resolved and strictly calculated. The tiny pores are not resolved. Instead, assumptions based on observation are made about how they affect the electrolyte flow. While the resulting model does not fully account for all of the smallest-scale interactions in the system, the team sees excellent agreement with analytical solutions. The comparison with results obtained from their experimental collaborators, also at DLR, is currently ongoing.

"We rely on experiments for informing our models, because many material properties cannot be guessed," Lautenschläger said. "Then, we can set up a system that can be reproduced by computers and use experiments to validate it. Once a model is validated, we can use the simulations to predict things that can't be resolved in experiments. In experiment, you can measure a result, but you don't always get a reason for that result; using simulations this can give you a better understanding. Therefore, experiment and simulation should always go hand-in-hand."

Charging forward

The team's initial successful simulations of electrolyte filling in a battery (published in *Batteries & Supercaps*) fold neatly into the DEFACTO Project, a European Union-funded initiative focused on optimizing material development and manufacturing processes for lithium-ion battery cells. The work also earned Lautenschläger praise from HLRS leadership: He was named one of three winners of the 2022 Golden Spike Awards during HLRS's annual Results and Review Workshop.

With this proof-of-concept work complete, the team is now focused on adding additional complexity to its simulations. "The motivation goes in two directions: make common battery technologies better by optimizing manufacturing processes and start improving the early stages of development of next-generation battery technologies," Lautenschläger said. The researchers have begun exploring how to improve current-generation batteries using more realistic simulations for manufacturing and of electrochemical processes. They are also investigating promising next-generation battery technologies such as lithium-sulfur batteries. Their goal is to optimize their design and to prevent degradation processes and unwanted side reactions to improve their long-term performance.

Regardless of whether the team focuses on further optimizing today's battery technologies or continues to design battery cells of the future, Lautenschläger made one thing clear: "What is sure is that our need for computational resources will only increase in the years to come." *EG*

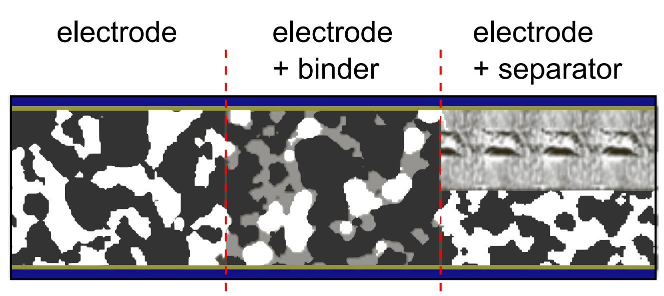

Microstructures (from left to the right) of a pure electrode (black), an electrode with binder, and an electrode attached to a separator. Mesoscale pores are depicted white and materials with nanoscale pores are depicted grey. Image: Martin Lautenschläger

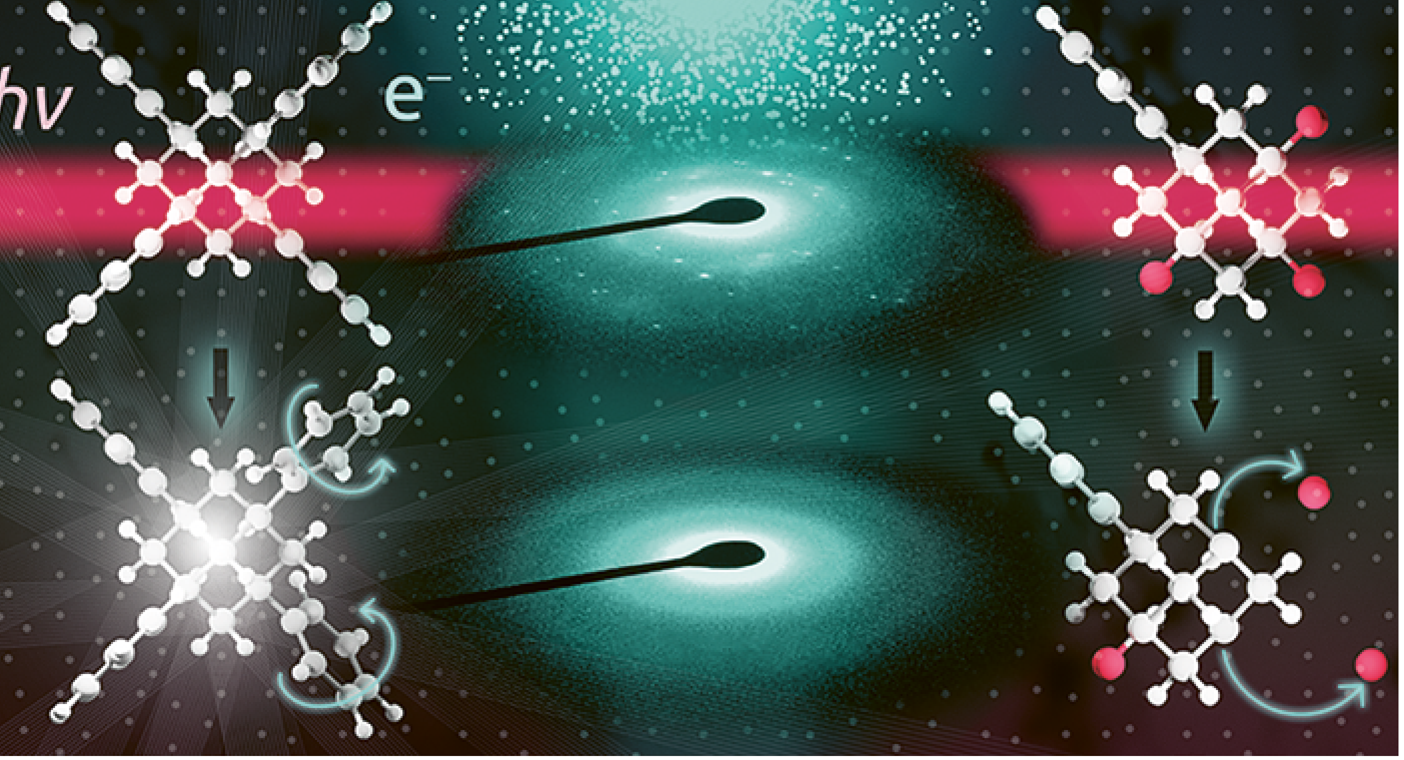

Structural modifications of the the molecular clusters that lead to the formation of amorphous compounds can be induced by electron or laser irradiation.
Image: Elisa Monte, Justus-Liebig-Universität Gießen

HPC Helps Identify New, Cleaner Source for White Light

Researchers at Justus Liebig University Giessen used HLRS supercomputing resources in the discovery of cluster glass, a new class of materials.

When early humans discovered how to harness fire, they were able to push back against the nightly darkness that enveloped them. With the invention and widespread adoption of electricity, it became easier to separate heat from light, work through the night, and illuminate everything from train cars to highways. In recent years, old forms of electric light generation such as halogen lightbulbs have given way to more energy efficient alternatives, further cheapening the costs to brighten our homes, workplaces, and lives generally.

Unfortunately, however, white light generation by newer technologies such as light-emitting diodes (LEDs) is not straightforward and often relies on a category of materials called "rare-earth metals," which are increasingly scarce. This has recently led scientists to look for ways to produce white light more sustainably. Researchers at Justus Liebig University Giessen, the University of Marburg, and Karlsruhe Institute of Technology have recently uncovered a new class of material called a "cluster glass" that shows great potential for replacing LEDs in many applications.

"We are witnessing the birth of white-light generation technology that can replace current light sources. It brings all the requirements that our society asks for: availability of resources, sustainability, biocompatibility," said Prof. Dr. Simone Sanna, Giessen University Professor and lead computational researcher on the project. "My colleagues from the experimental sciences, who observed this unexpected white light generation, asked for theoretical support. Cluster glass has an incredible optical response, but we don't understand why. Computational methods can help to understand those mechanisms. This is exactly the challenge that theoreticians want to face."

Sanna and his collaborators have turned to the power of high-performance computing (HPC), using the Hawk supercomputer at the High-Performance Computing Center Stuttgart (HLRS) to better understand cluster glass and how it might serve as a next-generation light source. They published their findings in *Advanced Materials*.

Clear-eyed view on cluster glass formation

If you are not a materials scientist or chemist, the word glass might just mean the clear, solid material in your windows or on your dinner table. Glass is actually a class of materials that are considered "amorphous solids;" that is, they lack an ordered crystalline lattice, often due to a rapid cooling process. At the atomic level, their constituent particles are in a suspended, disordered state. Unlike crystal materials, where particles are orderly and symmetrical across a long molecular distance, glasses' disorder at the molecular level make them great for bending, fragmenting, or reflecting light.

Experimentalists from the University of Marburg recently synthesized a particular of glass called a "cluster glass." Unlike a traditional glass that almost behaves as a liquid frozen in place, cluster glass, as the name implies, is a collection of separate clusters of molecules that behave as a powder at room temperature. They generate bright, clear, white light upon irradiation by infrared radiation. While powders cannot easily be used to manufacture small, sensitive electronic components, the researchers found a way to re-cast them in glass form: "When we melt the powder, we obtain a material that has all the characteristics of a glass and can be put in any form needed for a specific application," Sanna said.

While experimentalists were able to synthesize the material and observe its luminous properties, the group turned to Sanna and HPC to better understand how cluster glass behaves the way it does. Sanna pointed out that white light generation isn't a property of a single molecule in a system, but the collective behaviors of a group of molecules. Charting these molecules' interactions with one another and with their environment in a simulation therefore means that researchers must both capture the large-scale behaviors of light generation and also observe how small-scale atomic interactions influence the system. Any of these factors would be computationally challenging. Modeling these processes at multiple scales, however, is only possible using leading HPC resources like Hawk.

Collaboration between experimentalists and theoreticians has become increasingly important in materials science, as synthesizing many iterations of a similar material can be slow and expensive. High-performance computing, Sanna indicated, makes it much faster to identify and test materials with novel optical properties. "The relationship between theory and experiment is a continuous loop. We can predict the optical properties of a material that was synthesized by our chemist colleagues, and use these calculations to verify and better understand the material's properties," Sanna said. "We can also design new materials on a computer, providing information that chemists can use to focus on

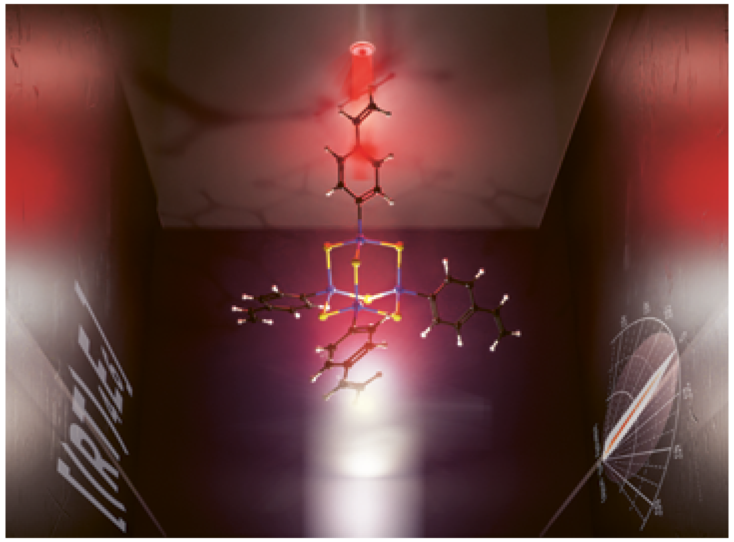

Upon irradiation by infrared light, adamantane-based molecular clusters with the general composition [(RT)4E5] (with R = organic group; T = C, Si, Ge, Sn; E = O, S, Se, Te, NH, CH2, ON•) emit highly directional white light.
Image credit: Elisa Monte, Justus-Liebig-Universität Gießen

synthesizing compounds that have the highest likelihood of being useful. In this way, our models inspire the synthetization of new compounds with tailored optical properties."

In the case of cluster glass, this approach resulted in an experiment that was verified by simulation, with modelling helping to show the researchers the link between the observed optical properties and the molecular structure of their cluster glass material, which can now move forward as a candidate to replace light sources heavily reliant on rare-earth metals.

HPC expedites R&D timelines

HPC plays a major role in helping researchers accelerate the timeline between new discovery and new product or technology. Sanna explained that HPC drastically cut down on the time to get a better understanding of cluster glass. "We spend a lot of time doing simulation, but it is much less than characterizing these materials in reality," he said. "The clusters we model have a diamond-shaped core with four ligands (molecular chains) attached to it. Those ligands can be made of any number of things, so doing this in an experiment is time consuming."

Sanna pointed out that the team is still limited by how long they can perform individual runs for their simulations. Many research projects on supercomputers can divide a complex system into many small parts and run calculations for each part in parallel. Sanna's team needs to pay special attention to long-distance particle interactions across large systems, so they are limited by how much they can divide their simulation across computer nodes. He indicated that having regular access to longer run times — more than a day straight on a supercomputer — would allow the team to work more quickly.

In ongoing studies of cluster glass Sanna's team hopes to thoroughly understand the origin of its light generating properties. This could help to identify additional new materials and to determine how best to apply cluster glass in light generation.

Sanna explained that HPC resources at HLRS were essential for his team's basic science research, which he hopes will lead to new products that can benefit society. "The main computational achievement in this journal article was only possible through our access to the machine in Stuttgart," he said. *EG*

Selected Publications by Our Users in 2022

Ala-Lahti M, Pulkkinen TI, Pfau-Kempf Y, et al. 2022. **Energy flux through the magnetopause during flux transfer events in hybrid-Vlasov 2D simulations.** Geophys Res Lett. 49(19): e2022GL100079.

Alho M, Battarbee M, Pfau-Kempf Y, et al. 2022. **Electron signatures of reconnection in a global eVlasiator simulation.** Geophys Res Lett. 49(14): e2022GL098329.

ALPHA Collaboration. 2022. **Determination of $\alpha_s(m_Z)$ by the non-perturbative decoupling method.** Eur Phys J C. 12: 1092.

Antolovic I, Vrabec J. 2022. **Vapor–liquid–liquid equilibria of nitrogen + ethane by molecular simulation.** Ind Eng Chem Res. 61(8): 3104-3112.

Bangga G, Parkinson S, Lutz T. 2022. **Utilizing high fidelity data into engineering model calculations for accurate wind turbine performance and load assessments under design load cases.** IET Renew Power Gen. ePub Dec 3.

Bangga G, Seel F, Lutz T, Kühn T. 2022. **Aerodynamic and acoustic simulations of thick flatback airfoils employing high order DES methods.** Adv Theor Simul. 5(8): 2200129.

Bell IH, Fingerhut R, Vrabec J, Costigliola L. 2022. **Connecting entropy scaling and density scaling.** J Chem Phys. 157: 074501.

Beyer J, Pfeiffer M, Fasoulas S. 2022. **Non-equilibrium radiation modeling in a gas kinetic simulation code.** J Quant Spectrosc Radiat Transfer. 280: 108083.

Bhowmik A, Alon OE. 2022. **Longitudinal and transversal resonant tunneling of interacting bosons in a two-dimensional Josephson junction.** Sci Rep. 12: 627.

Bi MN, Rost S, Auge M, et al. 2022. **Low-energy Se ion implantation in MoS_2 monolayers.** NPJ 2D Mater Appl. 6: 42.

Bocchini A, Gerstmann U, Bartley T, et al. 2022. **Electrochemical performance of $KTiOAsO_4$ (KTA) in potassium-ion batteries from density-functional theory.** Phys Rev Materials. 6: 105401.

Bocchini A, Gerstmann U, Schmidt WG. 2022. **Oxygen vacancies in $KTiOPO_4$: optical absorption from hybrid DFT.** Phys Rev B. 105: 205118.

Borsanyi S, Kara R, Fodor Z, et al. 2022. **Precision study of the continuum SU(3) Yang-Mills theory: how to use parallel tempering to improve on supercritical slowing down for first order phase transitions.** Phys Rev D. 105: 074513.

Borsanyi S, Guenther JN, Kara R, et al. 2022. **Resummed lattice QCD equation of state at finite baryon density: strangeness neutrality and beyond.** Phys Rev D. 105: 114504.

Chandola S, Sanna S, Hogan, et al. 2022. **Adsorbate-induced modifications in the optical response of the Si(553)–Au surface.** Phys Status Solidi RRL. 16(6): 2200002.

Chepkasov IV, Smet JH, Krasheninnikov AV. 2022. **Single- and multilayers of alkali metal atoms inside graphene/MoS_2 heterostructures: a systematic first-principles study.** J Phys Chem C. 126(37): 15558-15564.

Chiocchetti S, Dumbser M. 2022. **An exactly curl-free staggered semi-implicit finite volume scheme for a first order hyperbolic model of viscous two-phase flows with surface tension.** J Sci Comput. 94: 24.

Djukanovic D, von Hippel G, Koponen J, et al. 2022. **Isovector axial form factor of the nucleon from lattice QCD.** Phys Rev D. 106: 074503.

Doulis G, Atteneder F, Bernuzzi S, Brügmann B. 2022. **Entropy-limited higher-order central scheme for neutron star merger simulations.** Phys Rev D. 106: 024001.

Duan J, Chava P, Ghorbani-Asl M, et al. 2022. **Self-driven broadband photodetectors based on $MoSe_2$/$FePS_3$ van der Waals c–p type-II heterostructures.** ACS Appl Mater Interfaces. 14(9): 11927-11936.

Dudi R, Dietrich T, Rashti A, et al. 2022. **High-accuracy simulations of highly spinning binary neutron star systems.** Phys Rev D. 105: 064050.

Dues C, Müller MJ, Shatterjee S, et al. 2022. **Nonlinear optical response of ferroelectric oxides: first-principles calculations within the time and frequency domains.** Phys Rev Materials. 6: 065202.

Dunleavy N, Ballance CP, Ramsbottom, et al. 2022. **A Dirac R-matrix calculation for the electron-impact excitation of W+.** J Phys B. 55: 175002.

Eberheim K, Dues C, Attaccalite C, et al. 2022. **Tetraphenyl tetrel molecules and molecular crystals: from structural properties to nonlinear optics.** J Phys Chem C. 126(7): 3713-3726.

Eckart S, Pio G, Zirwes T, et al. 2022. **Impact of carbon dioxide and nitrogen addition on the global structure of hydrogen flames.** Fuel. 335. ePub Dec 13.

Frankel N, Pillepich A, Rix HW, et al. 2022. **Simulated bars may be shorter but are not slower than those observed: TNG50 versus MaNGA.** Astrophys J. 940: 61.

Gan Z, Paradisanos I, Estrada-Real A, et al. 2022. **Chemical vapor deposition of high-optical-quality large-area monolayer Janus transition metal dichalcogenides.** Adv Mater. 34(38): 2205226.

Gao M, Kuhn T, Munz CD. 2022. **On the investigation of oblique shock-wave/turbulent boundary-layer interactions with a high-order discontinuous Galerkin method.** Int J Numer Methods Fluids. 94(8): 1331-1357.

Gieg H, Schianchi F, Dietrich T, Ujevic M. 2022. **Incorporating a radiative hydrodynamics scheme in the numerical-relativity code BAM.** Universe 8(7): 370.

Glahn LJ, Ruiz Alvarado IA, Neufeld S, et al. 2022. **Clean and hydrogen-adsorbed AlInP(001) surfaces: structures and electronic properties.** Phys Status Solidi B. 2200308.

Hafner R, Guevara-Carrion G, Vrabec J, Klein P. 2022. **Sampling the bulk viscosity of water with molecular dynamics simulation in the canonical ensemble.** J Phys Chem B. 126(48): 10172-10184.

Han P, Bester G. 2022. **Determination of the phonon sidebands in the photoluminescence spectrum of semiconductor nanoclusters from ab initio calculations.** Phys Rev B. 106: 245404.

Heinen M, Hoffmann M, Diewald, et al. 2022. **Droplet coalescence by molecular dynamics and phase-field modeling.** Phys Fluids. 34: 042006.

Hoedl MF, Ertural C, Merkle R, et al. 2022. **The orbital nature of electron holes in $BaFeO_3$ and implications for defect chemistry.** J Phys Chem C. ePub Jul 21.

Huth S, Pang PTH, Tews I, et al. 2022. **Constraining neutron-star matter with microscopic and macroscopic collisions.** Nature. 606: 276-280.

Ibach M, Vaikuntanathan V, Arad A, et al. 2022. **Investigation of droplet grouping in monodisperse streams by direct numerical simulations.** Phys Fluids. 34: 083314.

Isotta E, Mukherjee B, Bette S, et al. 2022. **Static and dynamic components of Debye–Waller coefficients in the novel cubic polymorph of low-temperature disordered Cu_2ZnSnS_4.** IUCrJ. 9: 272-285.

Jain M, Gerstmann U, Schmidt TW, Aldahhak H. 2022. **Adatom mediated adsorption of N-heterocyclic carbenes on Cu(111) and Au(111).** J Comput Chem. 43: 413-420.

Jain M, Kretschmer S, Höflich K, et al. 2022. **Atomistic simulations of defects production under ion irradiation in epitaxial graphene on SiC.** Phys Status Solidi RRL. 202200292.

Janssen M, Falcke H, Kadler M, et al. 2022. **Event Horizon Telescope observations of the jet launching and collimation in Centaurus A.** Nat Astron. 5: 1017-1028.

Jöns S, Munz CD. 2022. **Riemann solvers for phase transition in a compressible sharp-interface method.** Appl Math Comput. 440: 127624.

Keim J, Munz CD, Rohde C. 2022. **A relaxation model for the non-isothermal Navier-Stokes-Korteweg equations in confined domains.** J Comput Phys. 474: 111830.

Kempf D, Munz CD. 2022. **Zonal direct-hybrid aeroacoustic simulation of trailing edge noise using a high-order discontinuous Galerkin spectral element method.** Acta Acust. 6: 39.

Kiwitt T, Fröhlich K, Meinke M, Schröder W. 2022. **Nusselt correlation for ellipsoidal particles.** Int J Multiphase Flow. 149: 103941.

Kopper P, Copplestone SM, Pfeiffer M, et al. 2022. **Hybrid parallelization of Euler-Lagrange simulations based on MPI-3 shared memory.** Adv Eng Softw. 174: 103291.

Köster J, Storm A, Ghorbani-Asl M, et al. 2022. **Structural and chemical modifications of few-layer transition metal phosphorous trisulfides by electron irradiation.** J Phys Chem C. 126(36): 15446-15455.

Kozub AL, Gerstmann U, Schmidt WG. 2022. **Third-order susceptibility of lithium niobate: influence of polarons and bipolarons.** Phys Status Solidi B. 2200453.

Krenz M, Gerstmann U, Schmidt WG. 2022. **Bound polaron formation in lithium niobate from ab initio molecular dynamics.** Appl Phys A. 128: 480.

Kretschmer S, Ghaderzadeh S, Facsko S, Krasheninnikov AV. 2022. **Threshold ion energies for creating defects in 2D materials from first-principles calculations: chemical interactions are important.** J Phys Chem Lett. 13(2): 514-519.

Kumar S, Gosselet P, Huang D, et al. 2022. **Parallel multiphysics simulation for the stabilized optimal transportation meshfree (OTM) method.** J Comput Sci. 62: 101739.

Kurz M, Offenhäuser P, Viola D, Resch M, Beck A. 2022. **Relexi — a scalable open source reinforcement learning framework for high-performance computing.** Software Impacts. 14: 100422.

Lagemann C, Lagemann K, Mukherjee S, Schröder W. 2022. **Generalization of deep recurrent optical flow estimation for particle–image velocimetry data.** Meas Sci Technol. 33(9): 094003.

Larsson HR, Schröder M, Beckmann R, et al. 2022. **State-resolved infrared spectrum of the protonated water dimer: revisiting the characteristic proton transfer doublet peak.** Chem Sci. 37. ePub Aug 30.

Lasek K, Ghorbani-Asl M, Pathirage, et al. 2022. **Controlling stoichiometry in ultrathin van der Waals films. $PtTe_2$, Pt_2Te_3, Pt_3Te_4, and Pt_2Te_2.** ACS Nano. 16(6): 9908-9919.

Lasek K, Li J, Ghorbani-Asl M, et al. 2022. **Formation of in-plane semiconductor–metal contacts in 2D platinum telluride by converting $PtTe_2$ to Pt_2Te_2.** Nano Lett. 22(23): 9571-9577.

Lautenschlaeger MP, Prifling B, Kellers B, et al. 2022. **Understanding electrolyte filling of lithium-ion battery electrodes on the pore scale using the lattice Boltzmann method.** Elec Soc S. 5(7): e202200090.

Lee J, Moon JS. 2022. **Merger effects on the spin and shape alignments of galaxy stellar, cold gas, hot gas, and dark matter components.** Astrophys J. 936(2): 119.

Lesnicki D, Wank V, Cryan JD, et al. 2022. **Lower degree of dissociation of pyruvic acid at water surfaces than in bulk.** Phys Chem Chem Phys. 24: 13510.

Letzgus P, Guma G, Lutz T. 2022. **Computational fluid dynamics studies on wind turbine interactions with the turbulent local flow field influenced by complex topography and thermal stratification.** Wind Energ Sci. 7: 1551-1573.

Li J, Joseph T, Ghorbani-Asl M, et al. 2022. **Edge and point-defect induced electronic and magnetic properties in monolayer $PtSe_2$.** Adv Funct Mater. 32(18): 2110428.

Lolicato F, Saleppico R, Griffo A, et al. 2022. **Cholesterol promotes clustering of $PI(4,5)P_2$ driving unconventional secretion of FGF2.** J Cell Biol. 221(11): e202106123.

Luu TD, Shamooni A, Stein OT, et al. 2022. **Flame characterisation of gas-assisted pulverised coal combustion using FPV-LES.** P Combust Inst. ePub Sep 6.

Luy JN, Henkel P, Grigjanis D, et al. 2022. **Bonding character of intermediates in on-surface Ullmann reactions revealed with energy decomposition analysis.** J Comput Chem. 44(3): 179-189.

Marx J, Kohns M, Langenbach K. 2022. **Systematic study of vapour-liquid equilibria in binary mixtures of fluids with different polarity from molecular simulations.** Mol Phys. e2141150.

Mastrikov YA, Gryaznov D, Sokolov MN, et al. 2022. **Oxygen vacancy formation and migration within the antiphase boundaries in lanthanum scandate-based oxides: computational study.** Materials. 15(7): 2695.

Mastrikov YA, Gryaznov D, Zvejnieks G, et al. 2022. **Sr doping and oxygen vacancy formation in $La_{1-x}Sr_xSC_{3-\delta}$ solid solutions: computational modeling.** Crystals 12(9): 1300.

Mausbach P, Fingerhut R, Vrabec J. 2022. **Thermodynamic metric geometry and the Fisher-Widom line of simple fluids.** Phys Rev E. 106: 034136.

Molignini P, Lévêque C, Keßler H, et al. 2022. **Crystallization via cavity-assisted infinite-range interactions.** Phys Rev A. 106: L011701.

Moritz DC, Ruiz Alvarado IA, Zare Pour MA, et al. 2022. **P-terminated InP (001) surfaces: surface band bending and reactivity to water.** ACS Appl Mater Interfaces. 14(41): 47255-47261.

Müller C, Mossier P, Munz CD. 2022. **A sharp interface framework based on the inviscid Godunov-Peshkov-Romenski equations: simulation of evaporating fluids.** J Comp Phys. 473: 111737.

Müller MJ, Ziese F, Belz J, et al. 2022. **Octave-spanning emission across the visible spectrum from single crystalline 1,3,5,7-tetrakis-(p-methoxyphenyl) adamantane.** Opt Mater Express. 12(9): 3517-3529.

Nitzke I, Pohl S, Thol M, et al. 2022. **How well does the Tang-Toennies potential represent the thermodynamic properties of argon?** Mol Phys. e2078240.

Norouzi MJ, Andric J, Vernet A, Pallares J. 2022. **Shape evolution of long flexible fibers in viscous flows.** Acta Mech. 233: 2077-2091.

Papadakis K, Pfau-Kempf Y, Ganse U, et al. 2022. **Spatial filtering in a 6D hybrid-Vlasov scheme to alleviate adaptive mesh refinement artifacts: a case study with Vlasiator (versions 5.0, 5.1, and 5.2.1).** Geosci Model Dev. 15: 7903-7912.

Papenfort LJ, Most ER, Tootle S, Rezzolla L. 2022. **Impact of extreme spins and mass ratios on the post-merger observables of high-mass binary neutron stars.** Mon Not R Astron Soc. 513(3): 3646-3662.

Pèrez-Montaño LE, Rodriguez-Gomez V, Bervantes Sodi B, et al. 2022. **The formation of low surface brightness galaxies in the IllustrisTNG simulation.** Mon Not R Astron Soc, ePub Jun 21.

Peter JMF, Kloker MJ. 2022. **Direct numerical simulation of supersonic turbulent flow with film cooling by wall-parallel blowing.** Phys Fluids. 34: 025125.

Qian W, Kronenburg A, Hui X, et al. 2022. **Effects of agglomerate characteristics on their collision kernels in the free molecular regime.** J Aerosol Sci. 159: 105868.

Ren J, Kloker M. 2022. **Instabilities in three-dimensional boundary-layer flows with a highly non-ideal fluid.** J Fluid Mech. 951: A9.

Rojas-León I, Christmann J, Schwan S, et al. 2022. **Cluster-glass for low-cost white-light emission.** Adv Mater. ePub Jun 24.

Rößler J, Antolovic I, Stephan S, Vrabec J. 2022. **Assessment of thermodynamic models via Joule-Thomson inversion.** Fluid Phase Equilib. 556: 113401.

Ruiz Alvarado IA, Schmidt WG. 2022. **Water/InP(001) from density functional theory.** ACS Omega. 7(23): 19355-19364.

Rusevich LL, Kotomin EA, Zyejnieks G, et al. 2022. **Effects of Al doping on hydrogen production efficiency upon photostimulated water splitting on $SrTiO_3$ nanoparticles.** J Phys Chem. 126(50): 21223-21233.

Sabater E, Sola M, Salvador P, Andrada DM. 2022. **Cage-size effects on the encapsulation of P_2 by fullerenes.** J Comp Chem. 44(3): 268-277.

Santini P, Castellano M, Fontana A, et al. 2022. **The stellar mass function in CANDELS and frontier fields: the buildup of low-mass passive galaxies since z ~ 3.** ApJ. 940(2): 135.

Satcunanathan S, Meinke M, Schröder W. 2022. **Impact of porous media on boundary layer turbulence.** Fluids. 7(4): 139.

Schmidt F, Kozub AL, Gerstmann U, et al. 2022. **A density-functional theory study of hole and defect-bound exciton polarons in lithium niobite.** Crystals 12(11): 1586.

Schmitt S, Vo T, Lautenschlaeger MP, et al. 2022. **Molecular dynamics simulation study of heat transfer across solid–fluid interfaces in a simple model system.** Mol Phys. 120(10): e2057364.

Schröder M, Gatti F, Lauvergnat D, et al. 2022. **The coupling of the hydrated proton to its first solvation shell.** Nat Commun. 13: 6170.

Sebastian R, Schreyer AM. 2022. **Flow fields around spanwise-inclined elliptical jets in supersonic crossflow.** Eur J Mechanics B Fluids. 94: 299-313.

Sebastian R, Schreyer AM. 2022. **Influence of jet spacing in spanwise-inclined jet injection in supersonic crossflow.** J Fluid Mech. 946: A39.

Secchi F, Gatti D, Frohnapfel B. 2022. **The wall-jet region of a turbulent jet impinging on smooth and rough plates.** Flow Turbul Combust. ePub Nov 23.

Secchi F, Häber T, Gatti D, et al. 2022. **Turbulent impinging jets on rough surfaces.** GAMM Mitteilungen. 45: e202200005.

Shamooni A, Stein OT, Kronenburg A, et al. 2022. **Fully-resolved simulations of volatile combustion and NO_x formation from single coal particles in recycled flue gas environments.** P Combust Inst. ePub Oct 1.

Shen X, Vogelsberger M, Nelson D, et al. 2022. **High-redshift predictions from IllustrisTNG – III. Infrared luminosity functions, obscured star formation, and dust temperature of high-redshift galaxies.** Mon Not R Astron Soc. 510(4): 5560-5578.

Soares PMM, Careto JAM, Cardoso RM, et al. 2022. **The added value of km-scale simulations to describe temperature over complex orography: the CORDEX FPS-Convection multi-model ensemble runs over the Alps.** Climate Dynam. ePub Dec 23.

Sotillo-Ramos D, Pillepich A, Donnari, et al. 2022. **The merger and assembly histories of Milky Way- and M31-like galaxies with TNG50: disc survival through mergers.** Mon Not R Astron Soc. 516(4): 5405-5427.

Steinhausen M, Zirwes T, Ferraro F, et al. 2022. **Turbulent flame-wall interaction of premixed flames using Quadrature-based Moment Methods (QbMM) and tabulated chemistry: an a priori analysis.** Int J Heat Fluid Flow. 93: 108913.

Stiskalek R, Bartlett DJ, Desmond H, Anbajagane D. 2022. **The scatter in the galaxy-halo connection: a machine learning analysis.** Mon Not R Astron Soc. 514(3): 4026-4045.

Strabac A, Greiwe DH, Hoffmann F, et al. 2022. **Piloted simulation of the rotorcraft wind turbine wake interaction during hover and transit flights.** Energies. 15(5): 1790.

Trcka A, Baes M, Camps P, et al. 2022. **UV to submillimetre luminosity functions of TNG50 galaxies.** Mon Not R Astron Soc. 516(3):3728-3749.

Ujevic M, Rashti A, Gieg H, et al. 2022. **High-accuracy high-mass-ratio simulations for binary neutron stars and their comparison to existing waveform models.** Phys Rev D. 106: 023029.

Walzer U, Hendel R. 2022. **Mantle evolution and continental growth events.** Earth-Sci Rev. 2323: 104130.

Wang W, Lozano-Duran A, Helmig R, Chu X. 2022. **Spatial and spectral characteristics of information flux between turbulent boundary layers and porous media.** J Fluid Mech. 949.

Warrach-Sagi K, Ingwersen J, Schwitalla T, et al. 2022. **Noah-MP with the generic crop growth model Gecros in the WRF Model: effects of dynamic crop growth on land-atmosphere interaction.** J Geophys Res-Atmos. 127(14): e2022JD036518.

Wen X, Shamooni A, Nicolai H, et al. 2022. **Flame structure analysis and flamelet modeling of turbulent pulverized solid fuel combustion with flue gas recirculation.** P Combust Inst. ePub Sep 18.

Wollny P, Menser J, Engelmann L, et al. 2022. **The role of phase transition by nucleation, condensation, and evaporation for the synthesis of silicon nanoparticles in a microwave plasma reactor — simulation and experiment.** Chem Eng J. 453(1): 139695.

Wu Y, Römer T, Axtmann G, Rist U. 2022. **Transition mechanisms in a boundary layer controlled by rotating wall-normal cylindrical roughness elements.** J Fluid Mech. 945: A20.

Yogi P, Joch J, Sanna S, Pfnür H. 2022. **Electronic phase transitions in quasi-one-dimensional atomic chains: Au wires on Si(553).** Phys Rev B. 105: 235407.

Zamponi R, Satcunanathan S, Moreau S, et al. 2022. **Effect of porosity on Curle's dipolar sources on an aerofoil in turbulent flow.** J Sound Vib. 542: 117353.

Zhang F, Kujata M, Sebbar N, et al. 2022. **Numerical study on flame stabilization and NO_x formation in a novel burner system for sulfur combustion.** Energy Fuels. 36(7): 4094-4106.

Zhang F, Zirwes T, Wachter S, et al. 2022. **Numerical simulations of air-assisted primary atomization at different air-to-liquid injection angles.** Int J Multiphase Flow. 158: 104304.

Zhang F, Zirwes T, Wang Y, et al. 2022. **Dynamics of premixed hydrogen/air flames in unsteady flow.** Phys Fluids. 34: 085121.

Zimmermann NER, Guevara-Carrion G, Vrabec J, Hansen N. 2022. **Predicting and rationalizing the Soret coefficient of binary Lennard-Jones mixtures in the liquid state.** Adv Theor Simul. 5(11): 2200311.

Zirwes T, Zhang F, Bockhorn H. 2022. **Memory effects of local flame dynamics in turbulent premixed flames.** Combust Inst. ePub Sep 20.

About Us

Inside Our Computing Room

Hewlett Packard Enterprise Apollo (Hawk)

HLRS's flagship supercomputer, called Hawk, was ranked #16 in its November 2020 debut on the Top500 List of the world's fastest supercomputers. Based on second-generation EPYC processors from AMD, the system is optimized for the sustained application performance and high scalability required for large-scale simulation, particularly for engineering and the applied sciences. In September 2021, HLRS announced the beginning of production of an expansion of Hawk that includes HPE Apollo systems with NVIDIA graphic processing units (GPUs). The upgrade has enhanced the center's capacity for deep learning and artificial intelligence applications, and enables new kinds of hybrid computing workflows that integrate HPC with Big Data methods.

System Type: Hewlett Packard Enterprise Apollo
CPU Type: AMD EPYC Rome 7742, 64 core, 2.25 GHz
Number of compute nodes: 5,632
Number of compute cores: 720,896
System peak performance: 26 petaflops
Total system memory: ~ 1.44 PB
Total disk storage capacity: ~ 25 PB

System Type: Apollo 6500 Gen10 Plus
GPU Type: NVIDIA A100
Number of GPUs: 192
Performance: 120 petaflops AI performance

Funding for Hawk was provided by the Baden-Württemberg Ministry of Science, Research and Arts, and by the German Federal Ministry for Education and Research through the Gauss Centre for Supercomputing (GCS). Hawk is part of the GCS national supercomputing infrastructure.

Baden-Württemberg
MINISTRY OF SCIENCE, RESEARCH AND ARTS

Federal Ministry of Education and Research

Cray CS-Storm

The Cray CS-Storm is optimized for artificial intelligence (AI) workloads, including processing-intensive applications for deep learning. Based on a GPU architecture, the CS-Storm provides a high-performance platform for deep learning frameworks such as TensorFlow and PyTorch, while also supporting use of classical machine learning tools such as Apache Spark and scikit-learn. The system is installed with the Cray Urika-CS AI and analytics suite, enabling HLRS users to address complex problems and process data with higher accuracy.

Deep learning partition: 64 NVIDIA Tesla V100 GPUs
Cray CS500 Spark partition: 8 CPU nodes
Software compiler: Urika-CS AI Suite
Interconnect: HDR100 Infiniband

AMD GPU System

Installed in 2021, this GPU-based system was donated to HLRS by hardware manufacturer AMD as a part of AMD's COVID-19 High-Performance Computing Fund. The system is dedicated to providing computing resources for medical research related to the COVID-19 pandemic and other diseases, and provides data analytics capacity for addressing sudden demands for simulation and data analytics that can occur in crisis situations. This system is integrated into HLRS's Vulcan cluster.

Processors: 10 × AMD EPYC
Accelerators: 80 × AMD Instinct
Performance: 530 TFlops, 64-bit

NEC Cluster (Vulcan)

This standard PC cluster was installed in 2009. Its configuration has been continually adapted to meet increasing demands and provide requirement-optimized solutions, including CPU, GPU, and vector computing components. The current configuration is as follows.

Intel Xeon Gold 6248 @2.5GHz (CascadeLake)
Number of nodes: 96
Memory per node: 128 GB

Intel Xeon Gold 6138 @2.0GHz (SkyLake)
Number of nodes: 100
Memory per node: 192 GB

Intel Xeon E5-2660 v3@ 2.6 GHz (Haswell)
Number of nodes: 88
Memory per node: 256 GB

Intel Xeon E5-2680 v3 @ 2.5 GHz (Haswell)
Number of nodes: 168
Memory per node: 384 GB

AMD Radeon
CPU: Intel Xeon Silver 4112 @ 2.6 GHz (Skylake)
Number of nodes: 6
Memory per node: 96 GB
CPU: 1 × AMD Radeon Pro WX8200
CPU memory: 8 GB

Intel Xeon E5-2667 v4 @ 3.2 GHz (Broadwell) mit P100
Number of nodes: 10
Memory per node: 256 GB
CPU: 1 × Nvidia P100
CPU memory: 12 GB

NEC SX-Aurora TSUBASA A300-8 @ 2.6 GHz
Number of nodes: 8
Memory per node: 192 GB
Vector engines: 8 × NEC Type 10B @ 1.4 GHz
Vector engine memory: 48 GB @ 1.2 TB/second

Interconnects
Infiniband EDR/FDR/HDR/QDR

User Profile

In 2022 the Gauss Centre for Supercomputing approved 7 new large-scale projects (each project requiring more than 35 million core hours) for HLRS's flagship supercomputer, Hawk, for a total of 2.887 billion core-hours. The Partnership for Advanced Computing in Europe (PRACE) also approved 3 international simulation projects for HLRS, for a total of 318 million core-hours. In total, 130 projects, including test projects, were active on Hawk in 2022, for a total of 4.364 billion core hours.

System Usage by Scientific Discipline

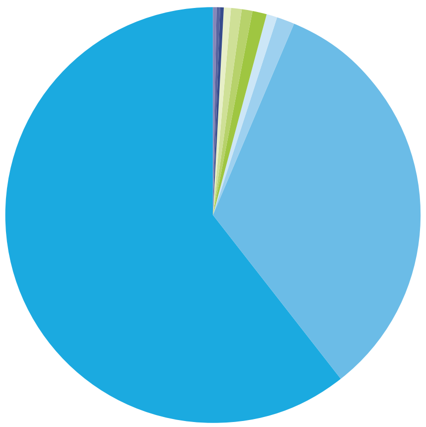

- CFD **60.57%**
- Physics **33.23%**
- Structural Mechanics **1.23%**
- Astrophysics **1.03%**
- Solid State Physics **1.00%**
- Materials Science **0.93%**
- Chemistry **0.74%**
- Other **0.44%**
- Reactive Flows **0.36%**
- Bioinformatics **0.21%**
- Computer Science **0.20%**
- Electrical Engineering **0.03%**
- Transportation and Climate **0.02%**

System Usage by State

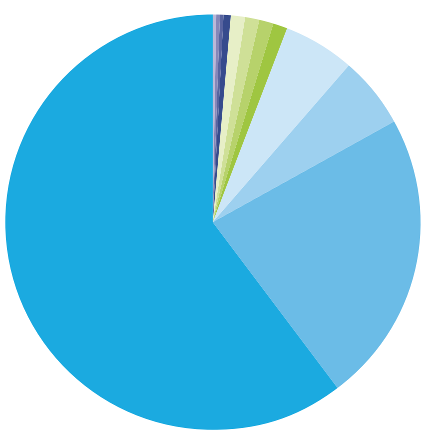

- North Rhine-Westphalia **60.47%**
- Baden-Württemberg **22.82%**
- Rheinland-Palatinate **5.53%**
- Hesse **5.36%**
- Thuringia **1.23%**
- Federal Research Center **1.22%**
- Berlin **1.12%**
- Brandenburg **0.90%**
- Saxony **0.66%**
- Hamburg **0.26%**
- Bavaria **0.23%**
- Saxony-Anhalt **0.21%**

Third-Party Funded Research Projects

In addition to providing supercomputing resources for scientists and engineers in academia and industry, HLRS conducts its own funded research on important topics relevant for high-performance computing (HPC), artificial intelligence, visualization, and high-performance data analytics. These activities, many of which are conducted in collaboration with investigators at other institutes and in industry, address key problems facing supercomputing and are opening up new opportunities for addressing key German, European, and global challenges. The following is a list of funded projects in 2022.

For more information about our current projects, visit *www.hlrs.de/projects.*

3×a ■
November 2022 – October 2025 (BMBF)
Will develop scalable methods for the simulation of three-body interactions in particle systems, applying vectorized kernels, dynamic load balancing approaches, and adaptive resolution schemata.

aqua3S
September 2019 – December 2022 (EU)
Developing a new system for detecting threats in drinking water safety and security, combining data from state-of-the-art sensors and other detection mechanisms.

bwHPC-S5
July 2018 – June 2023 (MWK)
Coordinates support for HPC users in Baden-Württemberg and the implementation of related measures and activities, including data intensive computing and large-scale scientific data management.

Cape Reviso
July 2020 – June 2023 (BMVI)
Developing planning and decision support tools for conflict analysis and reduction between cyclists and pedestrians in cities.

CASTIEL 2 ■
January 2023 – December 2025 (JU)
This coordination and support action will enhance the activities of the EuroCC 2 project and the European HPC Centres of Excellence by promoting collaboration and the exchange of knowledge and skills among national competence centers for HPC across Europe.

CATALYST
October 2016 – December 2022 (MWK)
Researched methods for analyzing large datasets produced by modelling and simulation, with the goal of implementing a framework that combines HPC and data analytics.

■ New in 2022 ■ Grant awarded, starts in 2023

Funder Abbreviations:
BMBF – Federal Ministry of Education and Research | BMVI – Federal Ministry of Transport and Digital Infrastructure | BMWi – Federal Ministry for Economic Affairs and Energy | CZS – Carl Zeiss Foundation | DBU – German Federal Environmental Foundation | DFG – German Research Foundation | ESF – European Social Fund | EU – European Union | ICM – InnovationsCampus Mobilität der Zukunft | JU – European High Performance Computing Joint Undertaking | MWK – Baden-Württemberg Ministry for Science, Research, and Art | UM – Baden-Württemberg Ministry of the Environment, Climate Protection and the Energy Sector | WAT – Baden-Württemberg Ministry of Economic Affairs, Labor and Tourism

CEEC ■
January 2023 – December 2026 (JU, BMBF)
The Center of Excellence for Exascale CFD is developing new and improved algorithms and workflows for computational fluid dynamics on exascale systems and pursuing strategies to improve energy efficiency in highly parallelized HPC architectures.

ChEESE 2 ■
January 2023 – December 2026 (JU, BMBF)
Preparing European flagship codes for upcoming pre-exascale and exascale supercomputing systems focusing on fields such as computational seismology, magnetohydrodynamics, physical volcanology, tsunamis, and the monitoring of earthquake activity.

CIRCE
November 2021 – October 2024 (BMBF, MWK)
A study to assess potential applications of high-performance computing (HPC) in crisis situations, and what organizational procedures are needed to ensure that HPC resources are immediately available.

CYBELE
January 2019 – March 2022 (EU)
Integrated tools from high-performance computing, high-performance data analytics, and cloud computing to support the development of more productive, data driven methods for increasing agricultural productivity and reducing food scarcity.

DECICE ■
December 2022 – November 2025 (EU)
Developing an open and portable cloud management framework that will enable the automatic and adaptive optimization of software applications for heterogeneous computing architectures.

DEGREE
June 2021 – June 2023 (DBU)
Investigating a method for increasing energy efficiency in data centers by dynamically controlling cooling circuit temperatures, and developing guidelines for implementing the resulting concepts.

EE-HPC ■
September 2022 – August 2025 (BMBF)
Testing an approach for improving energy efficiency in HPC systems by automatically regulating system parameters and settings based on current job requirements.

ENRICH
April 2021 – March 2023 (UM)
Analyzing current developments in IT and the operation of high-performance computing (HPC) centers regarding their resource efficiency and sustainability potential.

EuroCC 2 ■
January 2023 – December 2025 (JU, BMBF)
HLRS is the coordinating center of this Europe-wide project to establish national competence centers for HPC and to develop a shared, high level of expertise in high-performance computing, high-performance data analytics, and artificial intelligence.

exaFOAM
April 2021 – March 2024 (EU)
Working to reduce bottlenecks in performance scaling for computational fluid dynamics applications on massively parallel high-performance computing systems.

EXCELLERAT P2 ■
January 2023 – December 2026 (JU, BMBF)
Facilitates the development of important codes for high-tech engineering, including maximizing their scalability to ever larger computing architectures and supporting the technology transfer that will enable their uptake in industry.

FF4EuroHPC
September 2020 – August 2023 (JU)
Conducts outreach and provides support to Europe's small and medium-sized enterprises (SMEs) to enable them to profit from the advantages offered by high-performance computing technologies and services.

FocusCoE
December 2018 – March 2022 (EU)
Coordinated strategic collaboration and outreach among EU-funded Centres of Excellence to more effectively exploit the benefits of extreme scale applications for addressing scientific, industrial, or societal challenges.

Gaia-X4ICM
May 2022 – December 2024 (MWK, ICM)
The goal of Gaia-X4ICM is to implement a scaling production platform based on the Gaia-X ecosystem for the InnovationCampus Mobility of the Future (ICM) to make Gaia-X more usable for production of planning systems, industrial controls, and sensor data, among other applications.

HiDALGO 2
January 2023 – December 2026 (JU, BMBF)
Develops novel methods, algorithms, and software for HPC and high-performance data analytics to accurately model and simulate the complex processes that arise in connection with major global challenges such as air pollution and wildfires.

HPC-Europa 3
May 2017 – April 2022 (EU)
Fosters transnational cooperation among EU scientists (especially junior researchers) who work on HPC-related topics such as applications, tools, and middleware.

IKILeUS
December 2021 – November 2024 (BMBF)
HLRS is the coordinating center for this project to integrate artificial intelligence (AI) topics into curricula at the University of Stuttgart, and to implement AI technologies to improve instruction.

InHPC-DE
November 2017 – December 2023 (BMBF)
Coordinated integration among Germany's three Tier-1 supercomputing centers to create a standardized and distributed HPC ecosystem. It provided funding for 100 Gbit networking and opportunities for high-speed data management and visualization.

KoLab BW
March 2021 – December 2024 (MWK)
Developing tools for meeting and collaborating from remote locations in three-dimensional virtual reality environments.

MERIDIONAL
October 2022 – September 2026 (EU)
Developing a tool for assessing the performance and loads experienced by onshore, offshore, and airborne wind energy systems.

NFDI4Cat
October 2020 – September 2025 (DFG)
As a participant in the German National Research Data Infrastructure initiative, this consortium is creating a national platform for data integration in catalysis and chemical engineering research.

ORCHESTRA
December 2020 – November 2023 (EU)
The development of a networked platform for sharing data is enabling the creation of a new, large-scale, pan-European cohort that will improve research and responses to the SARS-CoV-2 pandemic and provide a model for addressing future public health threats.

PRACE
May 2019 – December 2022 (EU)
Supports high-impact scientific discovery and engineering R&D to enhance European competitiveness for the benefit of society.

SDC4Lit
May 2019 – April 2023 (MWK)
An interdisciplinary research project to sustainably organize the data lifecycle in digital literature. The resulting infrastructure will offer a data repository and research platform for the digital humanities.

SEQUOIA End-to-End
January 2023 – March 2024 (WAT)
Will develop transparent, automated, and controllable end-to-end solutions for the industrial use of hybrid quantum applications and algorithms through holistic quantum software engineering.

SERRANO
January 2021 – December 2023 (EU)
Introducing a novel ecosystem of cloud-based technologies, from specialized hardware resources to software toolsets, to enable application-specific service instantiation and optimal customization.

SimTech
July 2019 – March 2023 (DFG)
An interdisciplinary Excellence Cluster at the University of Stuttgart that is developing simulation technologies to enable integrative systems science. HLRS supports the development of efficient methods for uncertainty quantification and management.

Simulated Worlds
January 2011 – August 2024 (MWK)
Offers students opportunities to develop and execute simulation projects in collaboration with HLRS scientists.

SiVeGCS
January 2017 – December 2025 (BMBF, MWK)
Coordinates and ensures the availability of HPC resources of the Gauss Centre for Supercomputing, addressing issues related to funding, operation, training, and user support across Germany's national HPC infrastructure.

SODALITE
February 2019 – January 2022 (EU)
Aims to provide an optimized, resilient, heterogeneous execution environment that enables operational transparency between cloud and HPC infrastructures.

SRI DiTEnS
April 2023 – March 2029 (CZS)
Developing methods for discursive transformation in local energy systems, using urban digital twins involving virtual reality to support decision making among stakeholders.

TargetDART
October 2022 – September 2025 (BMBF)
Developing a task-based approach for highly scalable simulation software that mitigates load imbalance on heterogenous systems through dynamic, adaptive, and reactive distribution of computational load across compute resources.

TOPIO
November 2022 – October 2025 (BMBF)
Focusing on a large-scale, high-resolution Earth system model, TOPIO is investigating read and write rates for large amounts of data on high-performance file systems, as well as approaches that use compression to reduce the amount of data without causing a significant loss of information.

Trust in Information
August 2020 – August 2023 (MWK)
Multidisciplinary research led by the HLRS Department of Philosophy that is developing perspectives for assessing the trustworthiness of computational science and limiting the spread of misinformation.

WindHPC
October 2022 – September 2025 (BMBF)
In the first ever project to connect computers in wind farms with an HPC center, WindHPC aims to reduce energy consumption by improving efficiency in simulation codes, HPC workflows, and data management.

HPC Training Courses in 2022

HLRS offered 54 courses in 2022, providing continuing professional education on a wide range of topics relevant for high-performance computing. The courses took place over 140 course-days (compact courses), online and in Stuttgart and in cooperation with other institutes in Germany and internationally. A total of 1,416 trainees participated in these activities.

For a current listing of upcoming courses, please visit *www.hlrs.de/training*.

Date	Location	Topic	Host
Jan 10–Feb 21	online	HPC-Cluster – Auslegung, Kosten & Nachhaltigkeit	HLRS (SCA)
Jan 20–21	online	Data analytics for engineering data using machine learning NEW	Fraunhofer-SCAI/HLRS
Feb 1–2	online	AI for Science Bootcamp NEW	NVIDIA/HLRS/JSC/LRZ
Feb 7–11	online	Parallel Programming with MPI & OpenMP and Tools	ZIH/HLRS
Feb 14–18	online	Introduction to Computational Fluid Dynamics	HLRS/DLR
Feb 21–25	online	Fortran for Scientific Computing *	HLRS
Feb 21–Mar 28	online	Performance Optimierung - Kommunikation	HLRS (SCA)
Mar 8–11	online	Modern C++ Software Design (Intermediate)	HLRS
Mar 14–15	online	N-Ways to GPU Programming Bootcamp NEW	NVIDIA/HLRS/LRZ
Mar 14–Apr 11	online	Datenmanagement	HLRS (SCA)
Mar 22–25	Mainz	Parallelization with MPI and OpenMP	ZDV/HLRS
Mar 28–Apr 1	Stuttgart	Iterative Solvers and Parallelization	HLRS
Apr 5–7	online	Hybrid programming in HPC - MPI+X *	VSC Vienna/HLRS
Apr 7–8	online	AI for Science Bootcamp	NVIDIA/HLRS/JSC/LRZ
Apr 25–May 9	online	Paralleles Programmieren mit OpenMP	HLRS (SCA)
Apr 26–29	online	Optimization of Node-level Performance and Scaling on Hawk	HLRS
Apr 28–29	online	Shared memory parallelization with OpenMP @ VSC Vienna * (TtT)	VSC Vienna/HLRS
May 2–30	online	Visualisierung: Grundlagen & Anwendung	HLRS (SCA)
May 3–6	Stuttgart	Modern C++ Software Design (Advanced)	HLRS
May 9–11	online	Fortran for Scientific Computing	HLRS
May 16–Jul 4	online	HPC-Cluster – Aufbau & Betrieb	HLRS (SCA)
May 17–20	online	Parallelization with MPI @ VSC Vienna * (TtT)	VSC Vienna/HLRS
May 23–25	online	Data analytics for engineering data using machine learning	Fraunhofer-SCAI/HLRS
Jun 13–15	online	Data analytics for engineering data using machine learning	Fraunhofer-SCAI/HLRS
Jun 13–Jul 25	online	Simulation: Grundlagen & CFD	HLRS (SCA)
Jun 15–17	online	MPI and OpenMP in Scientific Software Development @ SURF * (TtT)	SURF/HLRS
Jun 20–21	online	Efficient Parallel Programming with GASPI *	HLRS
Jun 21–24	Stuttgart	Modern C++ Software Design (Intermediate)	HLRS
Jun 22–24	online	Hybrid Programming in HPC - MPI+X	LRZ/HLRS/NHR@FAU/VSC
Jun 28–Jul 1	online	Node-Level Performance Engineering *	HLRS
Jul 12–14	Stuttgart	Deep Learning and GPU programming using OpenACC	HLRS/LRZ
Aug 17–25	online	Six-day course in parallel programming with MPI/OpenMP	ETH/HLRS
Sep 12–16	Stuttgart	Introduction to Computational Fluid Dynamics	HLRS/DLR
Sep 12–Oct 24	online	Datenanalyse mit HPC	HLRS (SCA)
Sep 20–23	Stuttgart	Julia for High-Performance Computing * NEW	HLRS
Sep 26–27	online	Machine Learning with AMD Instinct GPUs and ROCm Software NEW	HLRS/AMD
Sep 29–30	Stuttgart	Scientific Visualization	HLRS
Sep 29–30	online	AMD Instinct GPU Training	HLRS/AMD
Oct 10–14	Stuttgart	Parallel Programming Workshop (with TtT)	HLRS
Oct 17–21	Stuttgart	CFD with OpenFOAM	HLRS
Oct 24–25	online	NVIDIA/HLRS SciML GPU Bootcamp NEW	HLRS/NVIDIA/LRZ
Oct 26–28	online	Introduction to oneAPI, SYCL2020 and OpenMP offloading NEW	HLRS/INTEL
Nov 7–11	Stuttgart/online	Optimization of Scaling, I/O and Node-level Performance on Hawk	HLRS
Nov 7–30	online	Paralleles Programmieren mit MPI	HLRS (SCA)
Nov 10–11	online	Shared memory parallelization with OpenMP @ VSC Vienna * (TtT)	VSC Vienna/HLRS
Nov 14–17	online	Introduction to GPU Programming using CUDA NEW	HLRS
Nov 21–23	online	Data analytics for engineering data using machine learning	Fraunhofer-SCAI/HLRS
Nov 22–25	online	Modern C++ Software Design (Advanced)	HLRS
Nov 22–25	Vienna/online	Parallelization with MPI @ VSC Vienna TtT	VSC Vienna/HLRS
Nov 28–30	online	Advanced Parallel Programming with MPI and OpenMP	JSC/HLRS
Dec 1–2	online	Introduction to NEC SX-Aurora TSUBASA vector platform	HLRS/NEC
Dec 5–9	online	Fortran for Scientific Computing *	HLRS
Dec 12–14	online	Hybrid Programming in HPC – MPI+X @ VSC Vienna * (TtT)	VSC Vienna/HLRS/NHR@FAU
Dec 12–16	online	From Machine Learning to Deep Learning: a concise introduction	HLRS

- Parallel Programming
- Computational Fluid Dynamics (CFD)
- Performance Optimization and Debugging
- Data in HPC
- Programming Languages for Scientific Computing
- Scientific Visualization
- Compute Cluster – Usage and Administration
- Training for special communities

* PRACE courses: HLRS is a member of the Gauss Centre for Supercomputing (GCS). GCS is one of ten PRACE Training Centres in the EU. The marked courses are in part sponsored by PRACE and are part of the PRACE course program.

TtT: Train the Trainer Courses

AMD – AMD | DLR – German Aerospace Center | ETH – Scientific IT Services, ETH Zurich | Fraunhofer-SCAI – Fraunhofer Institute for Algorithms and Scientific Computing | HLRS – High-Performance Computing Center Stuttgart | HLRS (SCA) – Supercomputing Academy | INTEL – INTEL | JSC – Jülich Supercomputing Centre | LRZ – Leibniz Supercomputing Centre | NEC – NEC | NHR@FAU – Erlangen National High Performance Computing Center | NVIDIA – NVIDIA | SURF – SURFSara (Dutch National Supercomputing Center) | VSC Vienna – Vienna Scientific Cluster | ZDV – Data Center, University of Mainz | ZIH – Center for Information Services and High Performance Computing (TU Dresden)

Workshops and Conferences in 2022

February 24
CIRCE Kickoff Workshop: How Urgent Computing Can Support Crisis Management
Organized for representatives of public administrative agencies, this workshop focused on how large-scale computer simulations can enable better forecasting and decision making in critical situations.

March 24
What Do Municipal Planning and Digital Models Have in Common?
This online workshop demonstrated opportunities that high-performance computing, virtual reality, digital twins, and related technologies offer for urban planning, city management, and sustainable development.

May 23 – 24
33rd Workshop on Sustained Simulation Performance
Organized in cooperation with NEC, this annual meeting brings scientists, application developers, and hardware designers from different continents together to discuss hardware architectures, programming styles, and strategies for achieving the highest possible sustained application performance.

August 31 – September 2
SAS Conference 2022: Trust in Information
This interdisciplinary conference explored the origins and nature of disinformation, its effects on public opinion, and potential strategies for fighting against it.

October 4 – 5
25th Results and Review Workshop
Scientists and engineers, including users of HLRS's computing infrastructure, presented and discussed research results as well as challenges and best practices in using HPC systems.

October 6 – 7
HPC User Forum
Organized by Hyperion Research and HLRS, the HPC User Forum brought together senior representatives of key HPC initiatives, internationally prominent high-performance computing centers, and leading technology manufacturers, as well as other experts on new HPC, AI, and quantum computing technologies and applications.

December 1
6th Industrial HPC User Round Table (iHURT)
The annual iHURT meeting facilitates dialogue between HLRS and its industrial user community, focusing on innovative applications of HPC for research and development as well as challenges that industry faces in using HPC.

December 12
Sustainable Procurement and Sustainable Computing Center Operations
Based on experience gained in the ENRICH project, this workshop focused on strategies for making procurement and disposal of large-scale computer systems more sustainable, as well as approaches for improving energy efficiency related to system operations, efficient programming, the use of intelligent networks, and other infrastructure.

Organization Chart

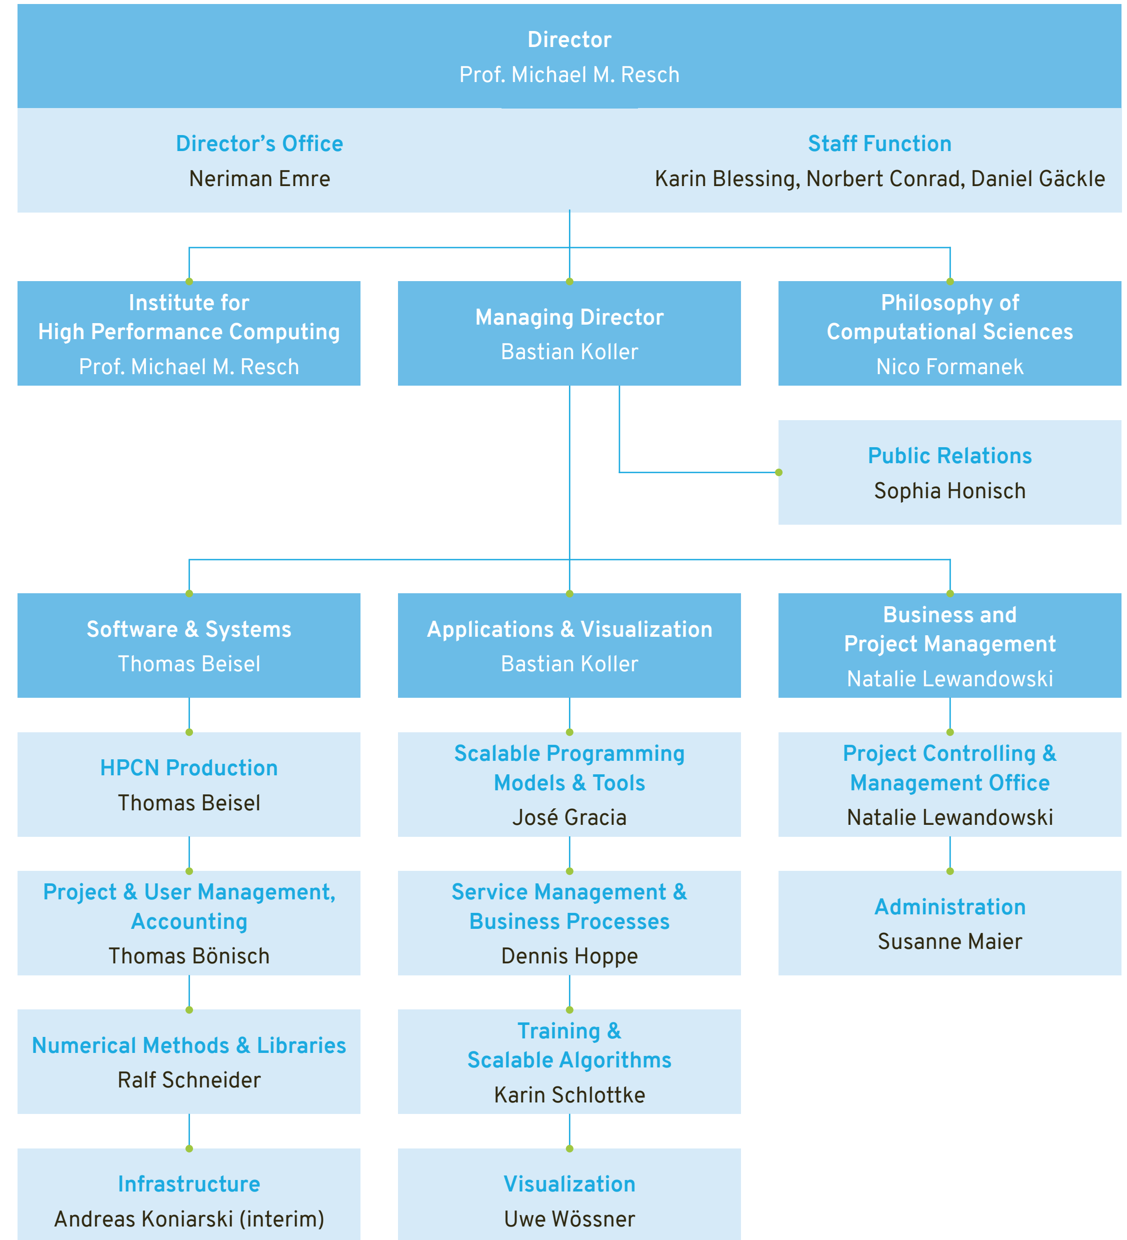

Departments

Administration

Leader: Susanne Maier

Manages issues related to the day-to-day operation of HLRS. Areas of responsibility include financial planning, controlling and bookkeeping, financial project management and project controlling, legal issues, human resources development, personnel administration, procurement and inventory, and event support.

High-Performance Computing Network – Production (HPCN Production)

Leader: Thomas Beisel

Responsible for the operation of all platforms in the compute server infrastructure. This department also operates the network infrastructure necessary for HPC system function and is responsible for security on networks and provided platforms.

Infrastructure

Interim Leader: Andreas Koniarski

Responsible for planning and operating facilities and infrastructure at HLRS. This division ensures reliable and efficient operation of the HLRS high-performance computing systems, provides a comfortable working environment for HLRS staff, and fosters all aspects of energy efficient HPC operation. It is also responsible for HLRS's sustainability program, which encourages and supports the entire HLRS staff in acting according to principles of sustainability.

Numerical Methods and Libraries

Leader: Dr.-Ing. Ralf Schneider

Provides numerical libraries and compilers for HLRS computing platforms. The department has expertise in implementing algorithms on different processors and HPC environments, including vectorization based on the architecture of modern computers. Department members also conduct research related to the simulation of blood flow and bone fracture in the human body, and are responsible for training courses focused on programming languages and numerical methods that are important for HPC.

Philosophy of Computational Sciences

Leader: Nico Formanek

Examines both how computer simulation and machine learning are changing science and technology development, and how society and politics react to these changes: Does simulation and machine learning change our understanding of knowledge and how we justify scientific results? How can computer-based methods help to overcome uncertainties about the future? And how do we deal with the uncertainties of simulation and machine learning itself?

Project Controlling and Management Office

Leader: Dr. Natalie Lewandowski

The Project Controlling and Management Office (PCMO) is responsible for the controlling and quality assurance of current research projects at HLRS or with HLRS as a beneficiary, and the management of large-scale third-party funded projects, including coordination and business development tasks. The PCMO also assists coordination at the proposal planning and writing stage and acts as a supporting and coordinating entity between the HLRS management, department heads, and HLRS administration in project-related matters.

Project and User Management, Accounting

Leader: Dr. Thomas Bönisch

Responsible for user management and accounting, including creating and maintaining web interfaces necessary for (federal) project management and data availability for users. The department also conducts activities related to the European supercomputing infrastructure (PRACE) and data management. This involves operating and continually developing the high-performance storage system as well as conceiving new strategies for data management for users and projects working in the field of research data management.

Public Relations

Leader: Sophia Honisch

Responsible for all areas of HLRS's external communications, from media relations to the management of HLRS's website and social media accounts: It is the main contact point for press and the broader public. The PR department communicates about HLRS's wide range of scientific and engineering disciplines, its research (projects) as well as its services, and disseminates results, new findings, and insights gained.

Scalable Programming Models and Tools

Leader: Dr. José Gracia

Conducts research into parallel programming models and into tools to assist development of parallel applications in HPC. Currently the focus is on transparent global address spaces with background data transfers, task-parallelism based on distributed data-dependencies, collective off-loading of I/O operations, and parallel debugging. As a service to HLRS users, the group also maintains part of the software stack related to programming models, debugging, and performance analysis tools.

Service Management and Business Processes

Leader: Dennis Hoppe

Evaluates novel technologies that will have a strong impact on the future use of high-performance computing. These technologies include artificial intelligence, cloud and edge computing, as well as quantum computing. For example, the group promotes the convergence of high-performance computing and artificial intelligence with the aim of incorporating AI methodologies into classical simulations to create hybrid HPC/AI workflows. This includes the development of AI solutions, especially in a business context, using cutting-edge technologies for Big Data, machine learning, and deep learning. The group also researches related virtualization technologies such as containers, orchestration, and job scheduling. By exploiting synergies between virtualization and HPC, the group has gained expertise in developing and operating dynamic and scalable federated cloud computing services.

Training and Scalable Algorithms

Leader: Karin Schlottke

Organizes and implements HLRS's training activities focusing on a variety of topics in high-performance computing, artificial intelligence, and modeling and simulation. These include compact, high-intensity courses, blended learning modules, and public outreach activities. In each area, our goal is to provide an outstanding learning experience by offering training on relevant topics, with up-to-date and audience-focused content, and given by highly-qualified instructors. Besides our teaching and outreach activities, we conduct research on the development of efficient algorithms for scientific computing applications.

Visualization

Leader: Dr.-Ing. Uwe Wössner

Supports engineers and scientists in the visual analysis of data produced by simulations on high-performance computers. By providing technologies capable of immersing users in visual representations of their data, the department enables users to interact directly with it, reducing analysis time and enabling new kinds of insights. The department is developing tools for visualization in virtual reality, augmented reality, and has designed a software system for integrating processing steps spread across multiple hardware platforms into a seamless distributed simulation and visualization environment.

High-Performance Computing Center Stuttgart (HLRS)
University of Stuttgart
Nobelstraße 19
70569 Stuttgart, Germany

Tel: +49 711 685-87269
Fax: +49 711 685-87209
Email: info@hlrs.de
Web: www.hlrs.de

Director, HLRS
Prof. Dr.-Ing. Dr. h.c. Dr. h.c. Prof. E.h. Michael M. Resch

Head, Department of Public Relations
Sophia Honisch

Editor
Christopher M. Williams

Contributing writers
Christopher M. Williams (CW)
Eric Gedenk (EG)

Translation support
Louise Acupanda

Photography and images
Unless otherwise indicated, all images property of HLRS.

Printing
gutenberg beuys, Feindruckerei GmbH, Langenhagen

Design
GROOTHUIS. Gesellschaft der Ideen und Passionen mbH
für Kommunikation und Medien, Marketing und Gestaltung;
groothuis.de

Institutional affiliations

University of Stuttgart Germany

GCS Gauss Centre for Supercomputing

Funding for Hawk provided by:

Baden-Württemberg Ministry of Science, Research and Arts

Federal Ministry of Education and Research

This magazine has been printed climate-neutral on paper that has been certified by FSC®.

RECYCLED Paper made from recycled material FSC® C009051

Carbon neutral Print product ClimatePartner.com/10951-2304-1007

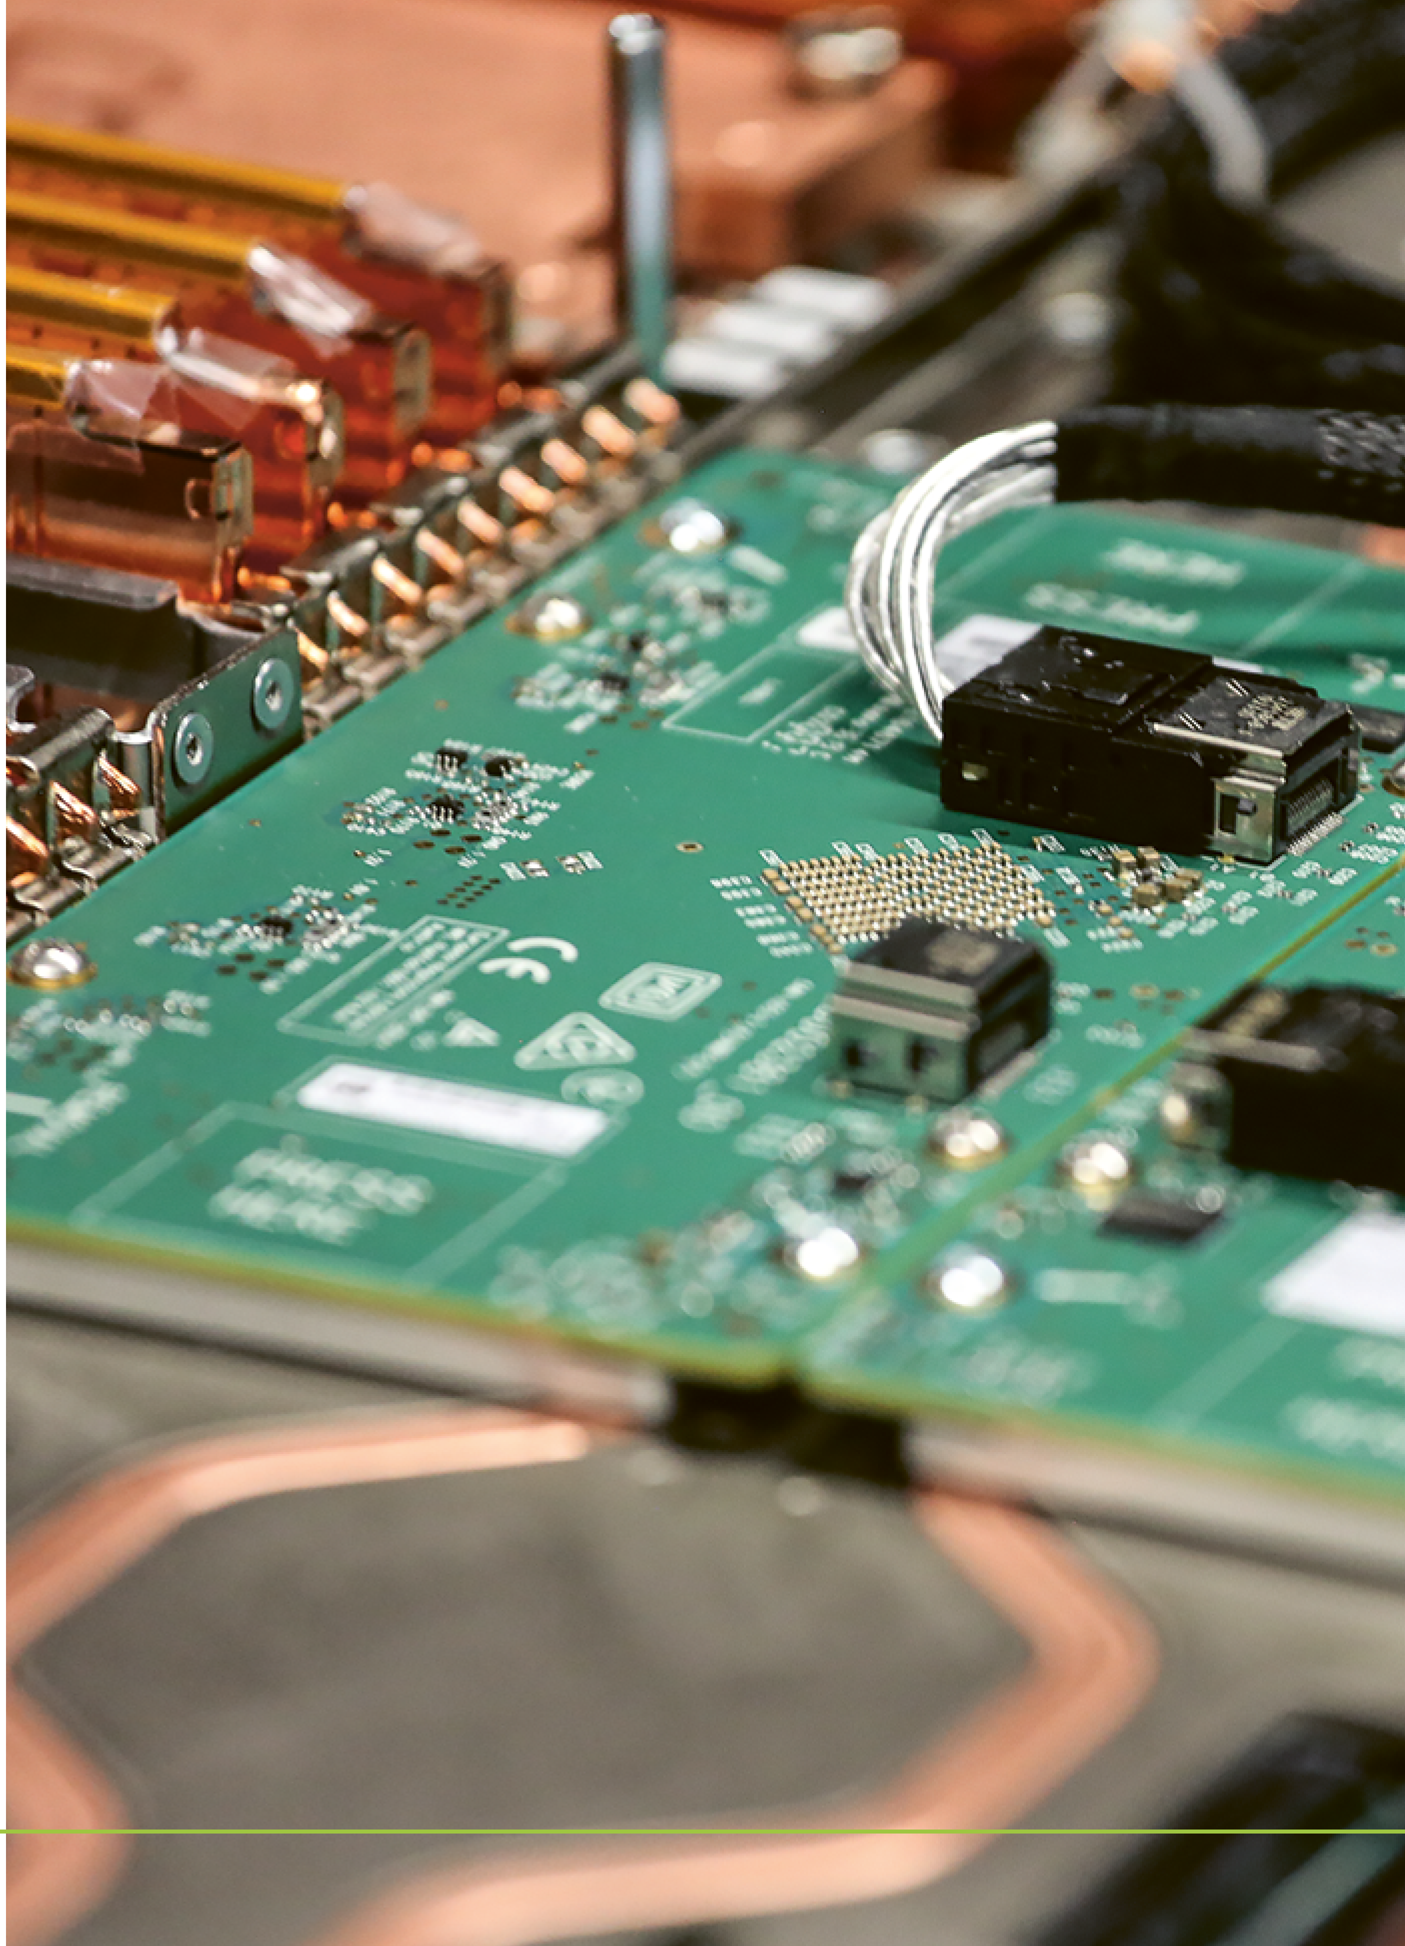

Follow us on:

LinkedIn Twitter

High-Performance Computing Center Stuttgart
www.hlrs.de

Cover:
Scientists in the HLRS Visualization Department worked with engineers and architects to develop a digital twin of the ElbX tunnel. The largest special structure of the SuedLink project, the tunnel will transport offshore wind energy from northern Germany under the Elbe River to the south.

HLRS is certified for environmental management under the Eco-Management Audit Scheme (EMAS) and Blue Angel Ecolabel. This magazine has been printed climate-neutral on paper that has been certified by FSC®.