# InSiDE

Innovatives Supercomputing in Deutschland

No 26-1 · Spring 2026

## Building for the Future

Construction advances on HLRS's next-generation data center in preparation for two new supercomputers, Herder and HammerHAI.

## New Quantum Computer at LRZ

The Leibniz Supercomputing Centre inaugurates Euro-Q-Exa, the first European-funded quantum computer in Germany.

## Record-Breaking Quantum Simulation

JSC and NVIDIA simulate 50 qubits on JUPITER as exascale supercomputing supports development of quantum systems.

GCS
Gauss Centre for Supercomputing

Just like small tributaries that converge to feed larger rivers, an increasingly diverse mix of technologies and computing paradigms are converging to feed the larger world of high-performance computing. In this issue, we celebrate this convergence by highlighting how our centers' staffs and users are evaluating, deploying, and leveraging new and novel computing methods and technologies. Whether it is researchers performing world-record quantum computing simulations, staff evaluating new photonic computing systems, or scientists using the power of exascale for new levels of detail in climate modeling, our centers are embracing what is possible as these powerful new tools combine and converge in the interest of advancing science.



# Welcome!

Welcome to the latest issue of InSiDE, the biannual Gauss Centre for Supercomputing magazine showcasing innovative supercomputing developments in Germany. In this issue, we are focusing on how high-performance computing (HPC) centers like ours are helping shape the convergence of new, novel computing technologies with traditional modeling and simulation.

In addition to their flagship systems, all our centers are investing in new quantum systems, systems designed specifically for artificial intelligence applications, and other new technologies while also evaluating how new paradigms like photonic and neuromorphic computing can be used to support researchers' needs.

The Leibniz Supercomputing Centre (LRZ) started operation of the Euro-Q-Exa system, the first European quantum computing system deployed in Germany by the EuroHPC Joint Undertaking (page 9). The Jülich Supercomputing Centre (JSC) has long played a pioneering role in quantum computing systems and recently broke the record for the largest quantum computer simulation on a classical HPC system, simulating 50 qubits on the JUPITER system (page 12). The High-Performance Computing Center Stuttgart (HLRS) just celebrated breaking ground on its new facility, HLRS III, which will house the center's next flagship system, Herder, and make room for a new suite of next-generation computing technologies (page 7). LRZ and JSC are also both in the midst of evaluating novel photonic computing chips, partnering with Stuttgart-based company Q.ANT to build prototype systems and evaluate their performance in HPC environments (pages 23 and 34 respectively).

Seeing the need to focus on evaluating these distinct computing technologies, HLRS started a new Future Computing Group in the last year. In this issue, we highlight the group's first event (page 32) and speak with its group leader, Johannes Gebert, about how he came into the role and how he views the future roles of HPC centers in this evolving computing paradigm (page 39).

While these new technologies offer exciting new potential, our centers ensure that providing world-class research tools to scientists remains at the forefront of our efforts. Researchers at JSC are using JUPITER to build a large foundation model for AI-accelerated protein research (page 17). Scientists at the Max Planck Institute for Plasma Physics have been using LRZ computing resources for modeling experimental fusion power plants at unprecedented detail (page 19).

While new computing technologies continue to multiply, the Gauss Centre for Supercomputing and its centers are focused on making sure these exciting technologies converge in the service of scientific excellence and societal benefit.

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

**Prof. Dr. D. Kranzlmüller**
**Director, Leibniz Supercomputing Centre**
**Chairman of the GCS Board of Directors**

**Prof. Dr. Dr. Th. Lippert**
**Director, Jülich Supercomputing Centre**
**Co-Chair of the GCS Board of Directors**

**Prof. Dr.-Ing. Dr. h.c. Dr. h.c. Hon.-Prof. M. M. Resch**
**Director, High-Performance Computing Center Stuttgart**
**Co-Chair of the GCS Board of Directors**

# Contents

"We are constructing a state-of-the-art building, promoting innovation and advanced research in the state, and investing in our future. This will create space for a new generation of computers. HLRS III in Stuttgart will be a top-class facility in Europe and a major European cluster for high-performance computing and artificial intelligence." Danyal Bayaz, Baden-Württemberg Finance Minister


Construction has begun on HLRS III, a new data center in Stuttgart. The new facility will include a waste heat processing center, seen at the left of this rendering, which will capture all waste heat generated by the center's supercomputers and distribute it to the district heating network of the University of Stuttgart campus.





# Building for the Future in Stuttgart

**HLRS announced details of its next-generation supercomputer, Herder, and celebrated the laying of the cornerstone of its new data center, HLRS III.**

The future of the High-Performance Computing Center Stuttgart has come into sharper focus. In December, 2025 HLRS signed an agreement with HPE to build and install its next supercomputer, called Herder. Based on the HPE Cray Supercomputing GX5000 platform, Herder will not only support the kinds of large-scale numerical simulations that are essential within HLRS's traditional user communities, but also data-science approaches for AI model training and generative AI. Bringing these capabilities together within a single system will enable the development of new kinds of converged computing workflows that seamlessly integrate the two approaches.

Continuing with an approach introduced with HLRS's current Hunter system, Herder will contain next-generation processors from AMD, including the AMD Instinct™ MI430X GPU and AMD EPYC™ "Venice" CPU. Each MI430X supports 432GB of HBM4 memory and 19.6TB/s of memory bandwidth, offering powerful capabilities for data-intensive operations in both HPC and AI. The Venice CPU is the world's first processor to use TSMC's (Taiwan Semiconductor Manufacturing Company's) 2-nanometer fabrication methods, which make it possible to pack more transistors onto each computer chip than was possible in the past. By bringing the transistors closer together, information can be moved through the system more quickly. With the increased memory and communications density that these processors offer, HLRS anticipates that Herder will achieve a peak performance of more than seven times the capabilities of Hunter.

## Cornerstone laid for HLRS III

This jump in performance that Herder will offer means that it will use substantially more power than Hunter, exceeding the limits of HLRS's current facilities. To prepare for Herder's arrival, construction has been underway to build a new data center on the HLRS campus, called HLRS III.

Friends and supporters of HLRS gathered at the center on December 15, 2025 to celebrate the laying of the HLRS III cornerstone. In a press release, Baden-Württemberg Finance Minister Danyal Bayaz said, "We are constructing a state-of-the-art building, promoting innovation and advanced research in the state, and investing in our future. This will create space for a new generation of computers. HLRS III in Stuttgart will be a top-class facility in Europe and a major European cluster for high-performance computing and artificial intelligence."

HLRS III will be home to the upcoming Herder and HammerHAI supercomputers.

On December 15, 2025, HLRS celebrated the laying of the cornerstone for HLRS III.

HLRS III will expand on existing HLRS facilities to provide an additional 7,000 square meters of new space. A hybrid construction using reinforced recycled concrete, steel, and wood, HLRS III will contain a state-of-the-art computer room with a redundant power supply, a highly energy-efficient cooling infrastructure, and a building structure capable of supporting the higher weight of future supercomputers. In addition to Herder, HLRS III will house the supercomputer for the AI Factory HammerHAI once construction has been completed.

### Energy efficiency and waste heat reuse

Maximizing energy efficiency is an essential element of the HLRS III building concept. To take advantage of the heat that its future systems will generate during operations, a new waste heat processing facility is also being constructed next to HLRS III. The facility will capture all waste heat generated by the center's supercomputers, increase the temperature with large-scale heat pumps, and distribute it to the district heating network of the University of Stuttgart's Vaihingen campus. This approach will cover up to one-third of the university's heating needs during winter months and all of its process heat requirements in the summer. When used to its full capacity, it is estimated that the HLRS III waste heat processing facility could reduce $CO_2$ emissions on campus by nearly 50%.

University of Stuttgart Rector Prof. Peter Middendorf highlighted the significance of HLRS III for the university community saying, "The visionary new construction of HLRS III is a milestone for the digital future of the University of Stuttgart, and is an excellent example of how we bring cutting-edge technology and sustainability together. The supercomputer Herder will not only make high-performance science possible, but will also actively contribute to the decarbonization of our campus."

Sustainable technologies on the facades of the two buildings will complement the high-performance computers contained within. The south and east sides of HLRS III will be covered with photovoltaic panels to generate electricity. The roof of HLRS's existing building was also recently outfitted with solar panels.

HLRS III has been identified as a component within the State of Baden-Württemberg's Green IT strategy. HLRS received the German Environmental Management Award 2025 in recognition of the particularly energy-efficient and sustainable operation of its computers. HLRS also won a Datacenter Strategy Award for "Transformation," highlighting the sustainable measures being implemented in HLRS III. *Christopher Williams*



NEWS
FEATURES

# Euro-Q-Exa starts Operation at LRZ, Becoming the First European-Funded Quantum System in Germany

**Euro-Q-Exa, the first EuroHPC Joint Undertaking-funded European quantum computer to be installed in Germany, recently went into operation at the Leibniz Supercomputing Centre. The system features 54 qubits and will be expanded by the end of 2026 with an additional system offering more than 150 qubits.**

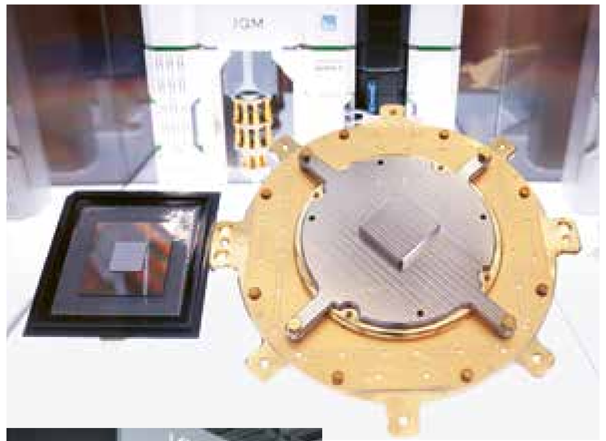


Euro-Q-Exa: the real deal inside LRZ's data centre (below) and LEGO replica and exhibits shown during the inauguration ceremony.

The Leibniz Supercomputing Centre has set its sights for a journey into unknown dimensions: Early in 2026, LRZ inaugurated the Euro-Q-Exa system in the presence of politicians from across Germany and Europe. The European quantum computer, or Euro-Q-Exa, is based on the Radiance system by IQM Quantum Computers and offers 54 quantum bits (qubits) made from superconducting circuits. The system is expected to be expanded by the end of 2026 with a second, more powerful quantum computer featuring more than 150 qubits. Dieter Kranzlmüller, Chairman of the Board of Directors of the LRZ, is already looking forward to exciting research projects that will use the system. "With Euro-Q-Exa, we combine the strengths of quantum computing and traditional supercomputing," Kranzlmüller said. "Researchers gain the opportunity to test new approaches, carry out groundbreaking computations, and thus open up new scientific dimensions – using technology made in Europe."

Euro-Q-Exa is one of six quantum systems being integrated into European high performance computers and procured by the EuroHPC Joint Undertaking (EuroHPC JU) in order to achieve technological independence in quantum computing. The total cost of €25 million – covering system costs, operations, and the planned expansion – is shared by the European Union, the German Federal Ministry for Research, Technology and Space, and the Bavarian State Ministry of Science and the Arts.

LRZ hosts and operates Euro-Q-Exa. Over the past years, the computing center has gained extensive experience integrating different quantum technologies into supercomputers,

NEWS
FEATURES

**"The inauguration of Euro-Q-Exa represents another milestone in our journey towards a world-class European quantum computing infrastructure. This new EuroHPC quantum system reinforces our commitment to providing researchers, industry, and the public sector with cutting-edge computational resources, fostering innovation and technological sovereignty across Europe."**

Executive Director of the EuroHPC JU, Anders Jensen


Joint kick-off for Euro-Q-Exa: Silke Launert (BMFTR), Dieter Kranzlmüller (LRZ), Maximilian Böltl (Bayr. Landtag), Markus Blume (StMWK), Henna Virkkunen (EU) und Sylwia Barthel de Weydenthal (IQM) from left to right.

as well as operating and utilizing them. Euro-Q-Exa is the second hybrid quantum computer installed by the LRZ together with IQM Quantum Computers in a production environment and is available to researchers via the Munich Quantum Portal (MQP). Coupled with the LRZ supercomputer, Euro-Q-Exa also enables hybrid workflows or the combination of classical supercomputing with cutting edge quantum computing.

### High computing power for large scale algorithms

Euro-Q-Exa uses IQM's Radiance technology, which is specifically designed for integration into supercomputing environments by minimizing latency and maximizing computing performance, particularly in hybrid workflows. The computer offers tunable couplers and high-fidelity gates that allow for a lattice topology. It is optimized for running large-scale algorithms and is cooled to below negative 273 degrees Celsius by a cryostat to stabilize its delicate computing units.

In principle, quantum computers operate differently from classical systems. They make use of the concepts of quantum superposition (essentially a principle that subatomic particles exist in multiple states simultaneously) and entanglement (the concept that two quantum particles are so interconnected they cannot be described individually) between multiple qubits. This allows them to process data simultaneously and perform certain computations in parallel. Theoretically, each additional qubit doubles computing power, enabling these innovative machines to solve mathematical problems that currently overwhelm classic high-performance computing (HPC) systems.

Researchers also hope to model electron interactions within atoms, molecular behavior, and other quantum mechanical states more precisely and efficiently with quantum computers. "Systems of the size and quality of Euro-Q-Exa enable interesting proof of concept simulations," says Luigi Iapichino, researcher at LRZ. "This is true especially when HPC and quantum resources are combined in hybrid workflows." Although coherence times, noise, and susceptibility to interference still limit the performance of quantum computers, larger experiments and meaningful results are already possible when they interact with classical supercomputers.

### Scalable tools for hybrid workflows

As part of the European initiative, Euro-Q-Exa is open to researchers from across the European Union. "Euro-Q-Exa will be available to a wide variety of European researchers, both through the Munich Quantum Portal and as an accelerator for HPC workflows that incorporate quantum resources," explains Iapichino. Additionally, the EuroHPC JU is also developing its own access pathways for Euro-Q-Exa. The quantum computer can be used on its own, in combination with SuperMUC NG, and ultimately in the future with Blue Lion, the next LRZ flagship supercomputer.

For programming, the system supports not only widely used quantum software packages such as Qiskit and PennyLane, but also various new programming languages. These are provided through the Munich Quantum Software Stack (MQSS), developed within the Munich Quantum Valley (MQV) by universities, research institutes, and companies, and now available for Euro-Q-Exa. It supports hybrid algorithms and workloads, the development of quantum programs, and provides interfaces to useful software tools.

The first research groups from across Europe and within MQV have already expressed interest in the machine. Early researchers aim to use Euro-Q-Exa to explore new dimensions – such as decoding the causes of neurodegenerative diseases, expanding methods in computational pharmacology, refining climate models, or improving power grids. Other fundamental research will focus on determining how to harness the power of quantum systems to determine the so-called "quantum advantage," where certain computational tasks can be executed more efficiently on a quantum system than traditional HPC resources. Others hope to finally determine how the long anticipated quantum advantage comes about, as recent findings suggest that achieving quantum advantage requires much more than simply a high number of qubits and efficient entanglement. *Susanne Vieser*

NEWS
FEATURES

# JSC Partners with NIVIDA for Record-Breaking Quantum Simulation on Classical Supercomputer

**Researchers used JUPITER to fully simulate 50 qubits, a new record for simulating a quantum computer. The work not only gives researchers insights into how to better design quantum systems, but also shows how exascale computers will play a significant role in the development of quantum computing.**

In late 2025, researchers from the Jülich Supercomputing Centre (JSC) together with NVIDIA achieved a quantum computing milestone – the team was able to fully simulate a universal quantum computing system with 50 qubits, breaking the previous record of 48 qubits simulated by Jülich researchers and their Japanese counterparts on Japan's K supercomputer in 2019. The team used JSC's JUPITER supercomputer, Europe's fastest high-performance computing (HPC) system and the first in Europe to cross the exascale threshold, for its achievement.

"Building a quantum computer is incredibly expensive and the hardware is still very noisy and prone to errors," said Dr. Hans De Raedt, researcher in JSC's Quantum Information Processing (QIP) group and collaborator on the project. "HPC-based simulators can act as a perfect flight simulator. They allow us to test their 'flight plans,' meaning the quantum algorithms, in a controlled environment where we know exactly what the answer should be. This helps us separate errors in the algorithm from errors in the hardware."

The work was born out of JSC researchers' interest in preparing a new version of the Jülich Universal Quantum Computer Simulator (JUQCS) that could take full advantage of JUPITER's new heterogeneous computing architecture that uses NVIDIA's Grace CPUs and Hopper GPUs for its GH200 superchip architecture. While JUPITER offered a massive performance increase, it also presented a new porting and scaling challenge for the research team.

## Memory and data topography challenge simulating quantum systems

Simulating a quantum system gets complicated quickly. In a quantum system – be it qubits in a quantum computer or subatomic particles interacting in space – researchers must account for the superposition of each qubit or particle. For particles, that usually means that researchers must consider that each particle represents a positive and negative charge simultaneously. Accordingly, researchers must represent qubits behaving as the computational binary of one and zero at the same time. De Raedt explained that each qubit possibility takes computational memory, meaning that tracking one qubit might only require a computer to keep up with two possibilities, but 10 qubits requires tracking 1,024 distinct possibilities, and 50 qubits requires tracking more than a quadrillion possibilities at once.

To simulate this large, complex system, researchers are essentially creating a massive spreadsheet that tracks each possible mix of ones and zeroes and updates the likelihood of each outcome based on how gates – the elementary operations of a quantum computer – change them during a simulation. The team used 4,096 nodes of JUPITER and created more than a petabyte of data that needs to be readily accessed by the system's memory. In fact, memory is one of the primary challenges in simulating larger volumes of qubits. "Since so much data cannot fit on a single chip, we must chop up the data and spread it across thousands of JUPITER's nodes, and we orchestrate data movement so that when gates need to interact with different parts of the data, they can find them without creating a 'traffic jam' on the system's network," De Raedt said.

The exascale supercomputer JUPITER at JSC.

To solve the data management problem, JSC and NVIDIA researchers improved JUQCS to operate more efficiently on either GPUs or CPUs and developed a novel memory compression technique that reduced memory requirements eight-fold. The new JUQCS version, JUQCS-50, will serve as a major foundational piece in JSC's quantum user facility, the Jülich UNified Infrastructure for Quantum computing (JUNIQ).

## JUQCS-50 as research tool and benchmark

Unlike a prototype quantum system, JUQCS offers researchers the ability to peer inside the "black box" of a quantum computer. De Raedt indicated that due to the hardware's sensitivity, trying to measure qubits inside a physical quantum computer causes the system state to collapse. This "noise" can come from programming, but also from vibrations to the system or other physical disturbances. Having the ability to accurately simulate an entire 50-qubit system allows researchers to analyze patterns and find correlations in the system in a way that is impossible on real hardware.

Further, because the team added artificial noise into the simulation, scientists can study how these disturbances impact a quantum system's ability to deliver clear, coherent results. Researchers can then compare simulations of noisy systems to one operating perfectly to better identify weak points in hardware design.

JUQCS-50 also serves as a powerful benchmarking tool for future quantum systems. De Raedt indicated that proving quantum supremacy means that a quantum system needs to clearly demonstrate it can do something that a classical computer cannot. This requires not only moving beyond relatively modest qubit counts, but also being able to benchmark gate fidelity and throughput – essentially the processing speed of a quantum computer. "If JUQCS can calculate a specific circuit in 10 minutes, a real quantum system needs to be able to do this with high accuracy in seconds to be a viable alternative to the current state-of-the-art in classical computing," he said.

For the moment, JUQCS-50 provides users at the JUNIQ facility access to a powerful tool to serve as a digital twin of a 50-qubit quantum system. "The future isn't just a simulator; it is a supercomputer with a quantum chip attached to it," he said. "We are exploring the integration of JUQCS into hybrid workflows at the center, and we already have several applications where our HPC system and quantum processors can work in a loop. We want to keep optimizing how our simulator handles these interactive tasks on JUPITER, ensuring the classical system's data processing bottleneck is minimized when we eventually swap the simulator for a physical quantum processing unit."
*Eric Gedenk*

To learn more about JUNIQ, please visit: https://www.fz-juelich.de/en/jsc/systems/quantum-computing/juniq-facility

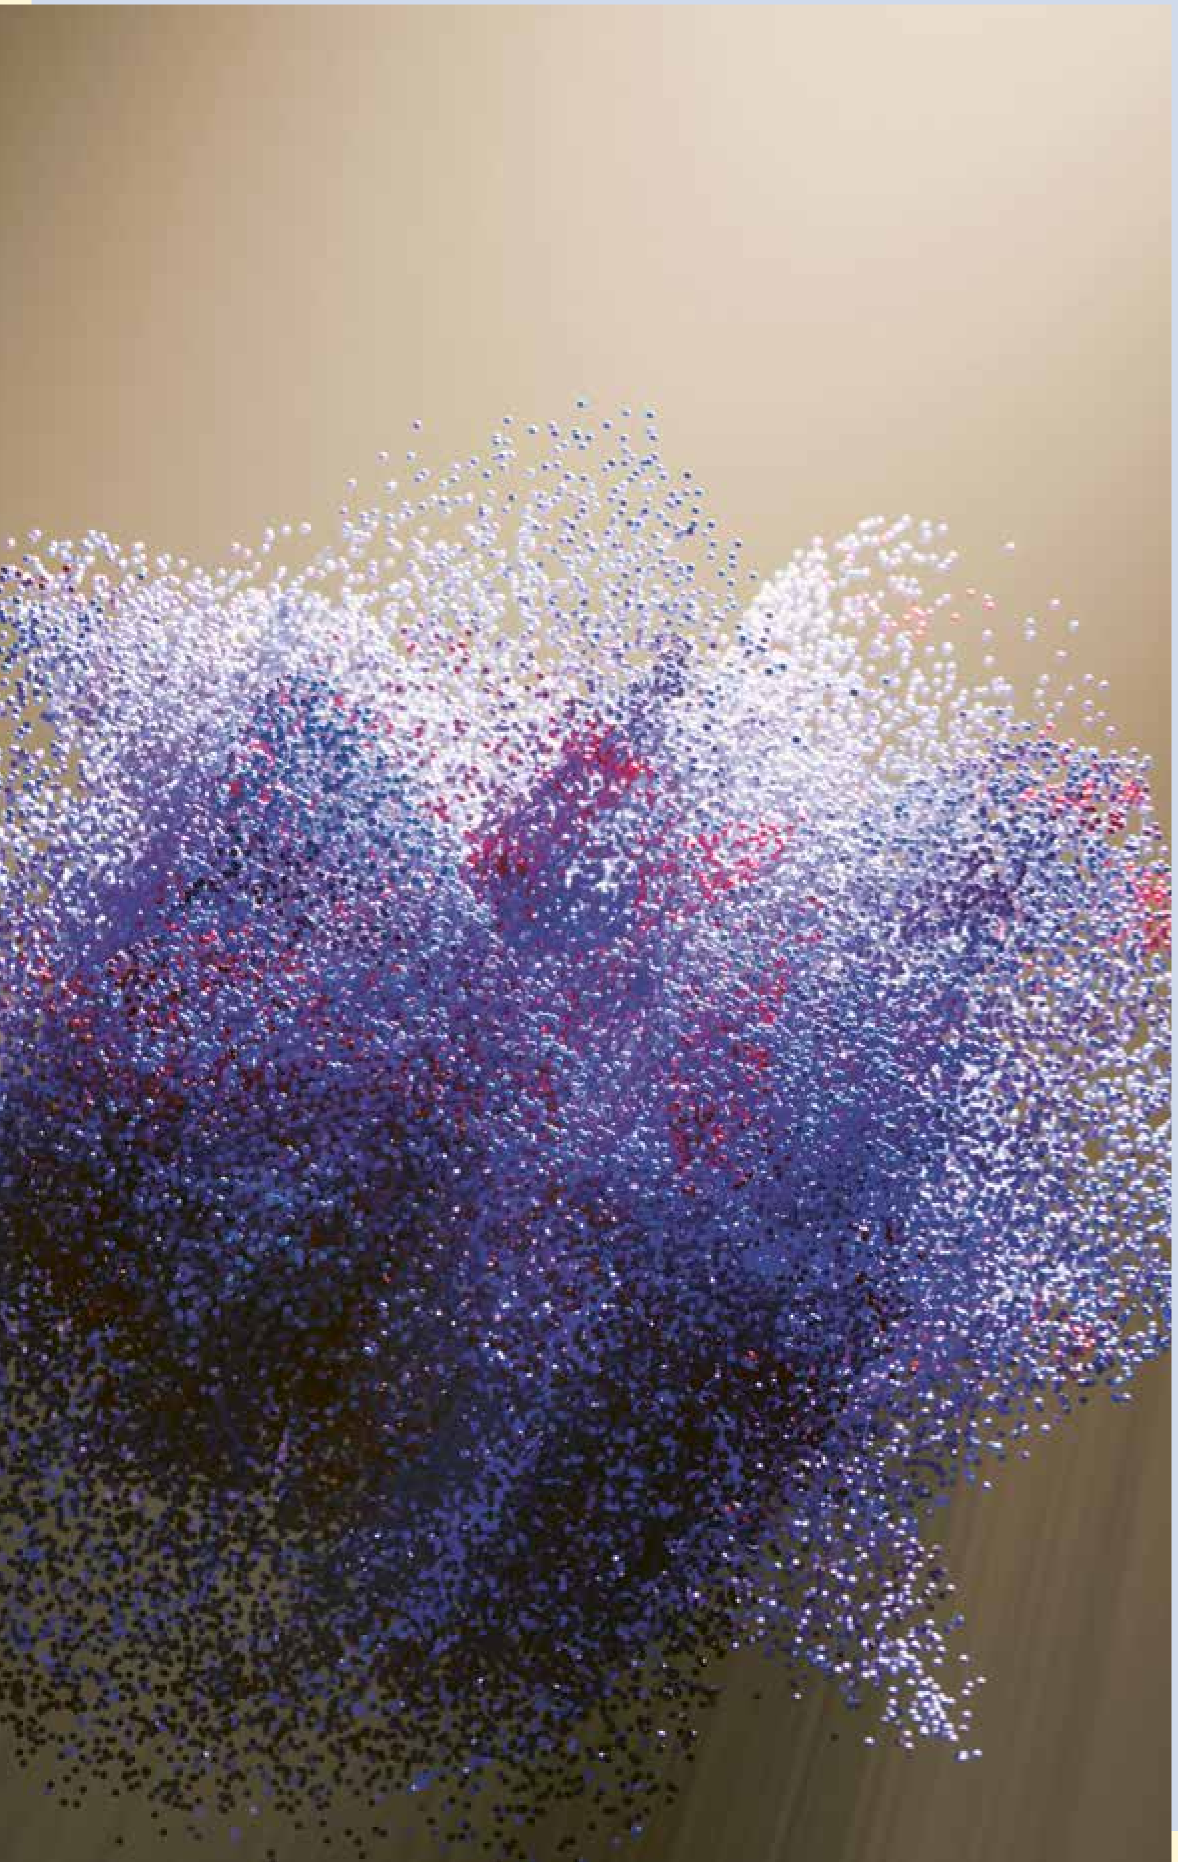

Researchers use powerful models on high-performance computing systems to study how particles and fluids interact in exotic environments, such as fuel injectors or inside power plants. As computing power has grown, researchers have been able to create more realistic simulations that better account for odd-shaped particles.

SCIENCE
FEATURES

# Stretching Our Understanding of Spheroidal Particles

**Researchers at RWTH Aachen University's Institute of Aerodynamics and Chair of Fluid Mechanics use high-performance computing to develop and improve models for simulating turbulent fluid-particle interactions. In recent years, graduate and postdoctoral researchers have spearheaded projects using HLRS supercomputing resources to model particles more common in nature. The research is helping to improve understanding of biomass combustion, which will support the design of cleaner next-generation biomass power plants.**

Understanding how fluids and particles behave during turbulent flows remains one of physics' greatest challenges. RWTH Aachen University's Institute of Aerodynamics and Chair of Fluid Mechanics (AIA) is among the world's leading institutions for advancing understanding of turbulence in engineering flows. For many years, the institute has leveraged high-performance computing (HPC) resources at the High-Performance Computing Center Stuttgart (HLRS) to model turbulent flows computationally. Typically, such simulations have assumed that particles are spherical in shape. In recent papers, however, the AIA team has been investigating the effects of other particle shapes, such as ellipsoids, on simulations of particle interactions. These shapes are more often found in the context of biomass combustion, a scalable, cleaner alternative to coal-fired power plants.

"The institute has been working on this problem for several years before I arrived as a graduate student, and it is clearly motivated by what we see in nature," said Laurent André, graduate researcher at AIA. "Most particles we see in nature are not perfectly spherical, but most modeling approaches assume spherical particle shapes. We are investigating the impact of this assumption – how valid is it? When does this assumption start to impact the accuracy of a model?"

André and recent graduate Dr. Thede Kiwitt worked closely with Professor Wolfgang Schröder, Head of AIA, and Dr. Matthias Meinke, Head of Numerics at AIA, to evaluate various modeling approaches for ellipsoidal particles and run high-fidelity simulations to improve their accuracy.

## Accurate modeling leads to better simulation

When researchers use computers to simulate turbulent flows, they must either include models that inform how certain particles behave or solve equations for each particle interaction during each time step. The direct simulation of particles without model assumptions is called direct particle-fluid simulation (DPFS). Although DPFS is prohibitively expensive for even modest system sizes, researchers have been able to increase the scale of direct particle-fluid simulations as computers have gotten more powerful.

In fact, AIA researchers in January, 2026 published a paper in the *International Journal of Multiphase Flow,* where they focused on doing DPFS of both spherical and ellipsoidal particles through pipes and jets. DPFS solves the particles' flow field directly, meaning that the boundaries and behaviors of each individual particle are numerically resolved without any additional model. Although the flow problem was relatively simple compared to real-world conditions, the team still needed 1.3 billion cells to simulate approximately 80,000 particles. The team demonstrated that although most particles were represented as spheres in such simulations, non-spherical particles behaved significantly differently. Specifically, it discovered that ellipsoidal particles with larger aspect ratios – the ratio of length to width in a particle – move further towards the center of a pipe compared to spherical particles. In the jet flow, the concentration peaks of particles with larger aspect ratio move farther off the centerline. The team used both HLRS's Hawk and Hunter supercomputers for its direct particle-fluid simulations.

Most scientists and engineers do not have access to the computational power needed for DPFS simulations, however. For this reason, André has been evaluating and improving models of particle flow produced using high-performance computing. These models can be integrated into less computationally demanding types of fluid-particle flow simulations called large-eddy simulations (LES). LES numerically calculates particle behavior at the larger turbulence scales but represents the impact of smaller eddies on the fluid-particle mixture in models. This makes LES more accessible for researchers who do not have access to large-scale supercomputers, including in industry.

André and his predecessors have worked with the team on evaluating point-particle models that can improve the accuracy of fluid-particle simulations. In point-particle models, researchers assume that particles have a consistent mass, even if they have different shapes. This simplifies the problem, making it possible to concentrate computational resources on simulating particle-fluid interactions.

Using HLRS resources, André and his collaborators evaluated three different point-particle models: a model with spherical particles, a model with ellipsoidal particles, and a model that includes ellipsoidal particles in a more complex system. This third category better accounts for their different behaviors regarding lift, drag, and torque. The team found not only that the point-particle model with additional physics outperformed the others, but also that it showed good agreement with more expensive direct particle-fluid simulations in terms of the preferential particle orientation in the free-jet. The team published its results in *Fuel.*

### Moving toward simulating realistic combustion conditions

The team's efforts to improve fluid-particle models lays a foundation for larger simulations that could realistically represent particle-laden combustion processes. While researchers at AIA have a long track record of simulating single and multi-phase flows, next-generation HPC systems offer the chance to improve these simulations on a fundamental level. "Our simulations so far are really reference simulations, as they are not yet at the scale to use in an industrial context," André said. "With access to leading HPC sources, we can increase the level of detail. This will be necessary to produce the highly resolved data we need to develop more accurate models, which will offer industry more reliable predictions." In this way, the team's simulations will continue to improve LES done in academic and industrial settings.

With continued access to leading HPC systems like those at HLRS, the team also aims to run DPFS at an increasingly large scale. Because of the size and complexity of the system, DPFS for an entire biomass reactor is still computationally impossible. With its current allocation of HLRS resources, however, the team is now focused on how to further improve point-particle models to better account for particle orientation and rotation during a simulation.
*Eric Gedenk*

*Funding for HLRS's Hawk and Hunter supercomputers was provided by the Baden-Württemberg Ministry for Science, Research, and the Arts and the German Federal Ministry of Research, Technology and Space through the Gauss Centre for Supercomputing (GCS).*

Related Publications
Kiwitt, T. et al. (2026). "Direct particle-fluid simulation of spherical and ellipsoidal particles in turbulent pipe-free-jet flow," *International Journal of Multiphase Flow.*
DOI: https://doi.org/10.1016/j.ijmultiphaseflow.2025.105443

Andre, L. et al (2026). "Comparison of Point-Particle Models and Direct Particle Fluid Simulations for Non-Spheroidal Particles," *Fuel.*
DOI: https://doi.org/10.1016/j.fuel.2026.139560

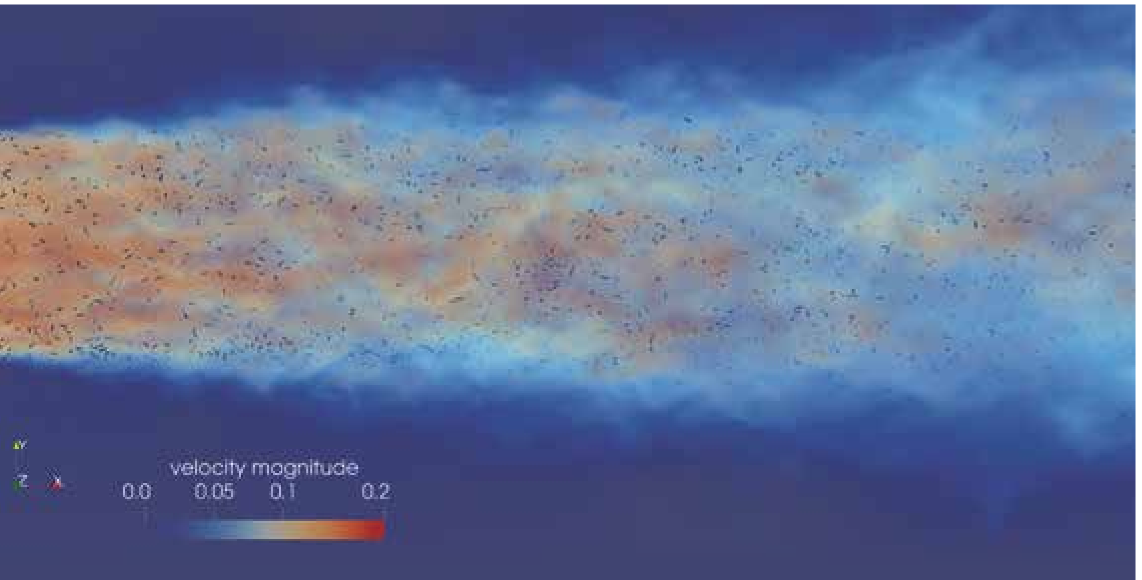


Using Hawk, the RWTH Aachen team was able to accurately model spheroidal particles, which more accurately represent biomass particles burning in a power plant.



# OneProt Initiative Pioneers a Multi-Modal Foundation Model for Protein Research

**A multi-institution effort is bringing protein research into the artificial intelligence era. A research team has built the foundation model OneProt for protein structure research, and, as part of the Gauss AI Compute Competition, is now also working on combining it with a large language model.**

Computational biologists and public health researchers rely on modeling and simulation to design new medications and to better understand how to deliver those drugs where they belong in the body. From speeding up vaccine development to gaining a deeper understanding of how certain viruses invade our body, simulations using high-performance computing (HPC) have become essential tools for advancing healthcare.

While traditional modeling and simulation methods remain important pillars for research into protein structure and function, rapidly advancing artificial intelligence (AI) methods offer powerful new tools to speed up research. "With AI, we are now capable of processing even larger amounts of information and biological detail on HPC systems, which has had a great impact on this kind of research," said Dr. Alina Bazarova, AI consultant at the Jülich Supercomputing Centre (JSC). "Now, as more diverse AI architectures are being developed, we are able to take more protein properties into account, and we can train more complex models for longer amounts of time."

Bazarova is JSC's institutional contact for the OneProt initiative. OneProt, which includes collaborators from Helmholtz AI, the Technical University of Munich, and other research groups at Forschungszentrum Jülich, focused on producing a multi-modal foundation model for protein research. This essentially means that the model is trained to process disparate datasets related to proteins and can be adapted by various research teams to focus on their specific research questions. The initial model was revealed in the team's recent paper in *PLOS Computational Biology* and sets the stage for the next phase of the team's work – OneProtGPT, which integrates large language models (LLMs) into this powerful AI tool.

### Multi-modal foundation models support larger research communities

The original motivation for OneProt began with researchers planning to focus on creating embeddings – essentially small numerical representations of proteins that allow protein variants to be studied by machine learning algorithms – for different protein mutations. During the development of the system, the researchers realized that training an AI system for this task could be readily extrapolated into other relevant research tasks. Therefore, the development team devised an easy-to-use fine-tuning process that only requires one to two additional layers in the model's neural network and can be adapted to any task without retraining. Additionally, to make the OneProt framework successful as a foundation model, the team designed it so researchers could easily incorporate new kinds of data during pre-training.

Accordingly, the model is not only capable of supporting research into protein mutations, but also can be adapted for studying the molecular structure of proteins and gene ontology – a classification system for gene function – among other research questions. OneProt was trained using massive amounts of data related to proteins' 3D structure, amino acid sequences, and binding site details,

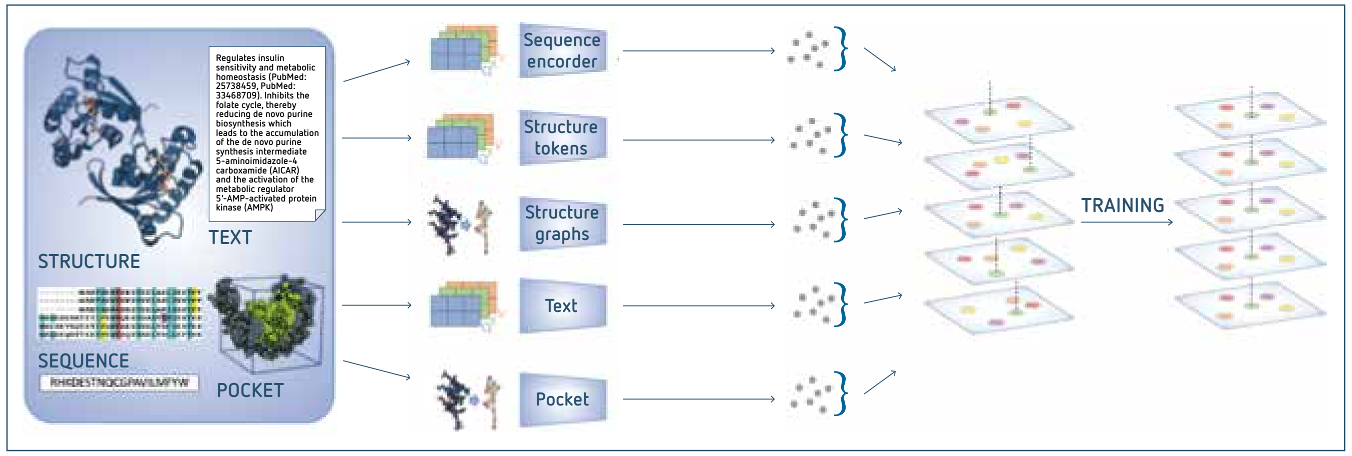


The OneProt model aligns multiple modalities, including primary protein sequence, 3D protein structure, protein binding sites, and text annotations. Each modality is processed by its respective encoder, generating embeddings aligned in a shared latent space, facilitating cross-modal learning and integration.

and the model has performed especially well at tasks predicting enzyme function and analyzing protein binding sites. The team attributes some of that success – as well as the model's relatively lean training footprint – to training using multiple modalities, meaning using various different types of data related to protein function and behavior. "Considering the improvements to AI models, we designed OneProt so that it could function as somewhat more of a blackbox technology," Bazarova said. "Rather than using an HPC system to take into account every possible interaction between protein molecules, we created some functions that help us get the desired outputs by allowing the AI model itself to infer some of these comparisons, as it is better at disentangling this data than we are on our own."

### OneProtGPT promises to further streamline foundation model via large-language models

Having demonstrated success with OneProt, the team applied for time on JSC's newest flagship supercomputer – JUPITER, Europe's first exascale system – as part of the Gauss AI Compute Competition. The program, organized by the Gauss Centre for Supercomputing, focuses on advancing next-generation generative AI models that can benefit diverse research communities. The work is funded by the German Federal Ministry of Research, Technology, and Space, the North Rhine-Westphalia Ministry of Culture and Science, the Baden-Württemberg Ministry of Science, Research and Arts, and the Bavarian State Ministry of Science and the Arts.

"OneProt was trained on the JUWELS Booster system, which uses NVIDIA A100 GPUs, and we are now moving to JUPITER, which offers GH200 Grace Hopper superchips," said Javad Kasravi, HPC AI research engineer at JSC and collaborator on the initiative. "The GH200 configuration combines larger memory, higher compute capacity, and faster data transfer. With roughly three times the GPU memory compared with the A100, an NVIDIA Grace CPU, and a highspeed link between the CPU and GPU, we can scale training more efficiently – helping us iterate faster and improve model accuracy in less time."

With access to JUPITER, the OneProt team is actively improving the model by integrating protein data with large language models. By integrating LLMs into OneProt, researchers could generate more detailed protein descriptions when using various data sources. "For now, we have finetuned OneProtGPT using protein sequence data, but we are working on the ability to replace this sequence data by other data modalities, which would greatly expand the model's flexibility for a variety of applications," Bazarova said. *Eric Gedenk*

*Funding for JUWELS was provided by the Ministry of Culture and Research of the State of North Rhine-Westphalia and the German Federal Ministry of Research, Technology and Space through the Gauss Centre for Supercomputing (GCS).*

Related publication
Flöge K, Udayakumar S, Sommer J, Piraud M, Kesselheim S, et al. (2025). "OneProt: Towards multi-modal protein foundation models via latent space alignment of sequence, structure, binding sites and text encoders," *PLOS Computational Biology 21(11).*
DOI: https://doi.org/10.1371/journal.pcbi.1013679





# Modeling Fusion Physics and Technology with HPC

**With the GENE code, researchers can simulate physics in various fusion devices on supercomputers. To enable larger and more realistic simulations through computing power, the code family gets a new sprout: GENE-X is currently being prepared for systems accelerated by graphics processing units.**

Expectations for nuclear fusion are high. Although scientists know that this energy source will not be as cheap to build as solar or wind energy, the energy generated from the controlled fusion of hydrogen isotopes into helium could, help close gaps in the supply of electricity from renewable sources – likely without the long-term risks associated with nuclear power, such as highly radioactive waste that is difficult to manage and assess.

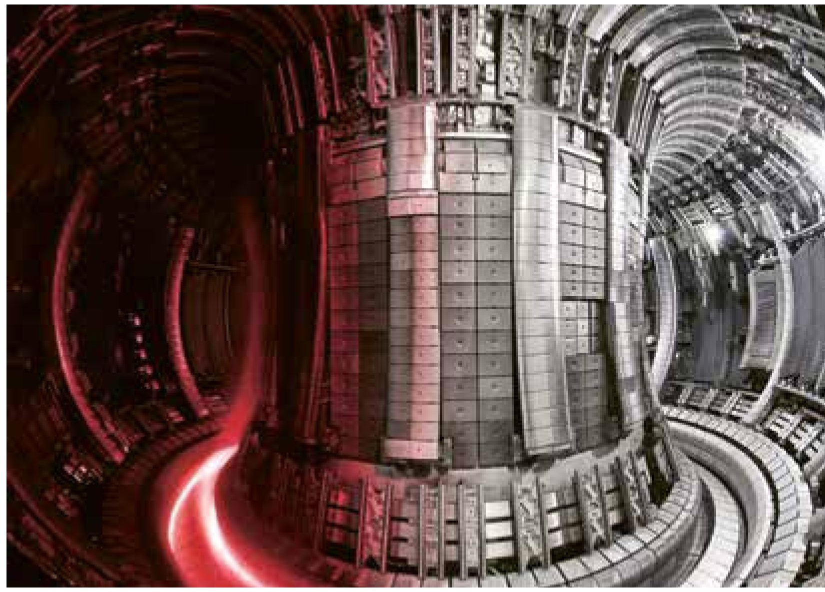

Inside the Joint European Torus tokamak (JET) at Culham, UK, with a superimposed image of the hot plasma

Researchers and startups worldwide are working on concepts for fusion power plants. They model plasma dynamics under the influence of strong magnetic fields which are used to stabilize particle mixtures under extreme conditions. At the Max Planck Institute for Plasma Physics (IPP) in Garching near Munich, such simulations are an everyday activity. "With GENE-X, we want to simulate plasma turbulence to study the causes of heat and particle losses which help us understand the performance of a fusion device," said Dr. Philipp Ulbl, a computational physicist working at IPP. "We also want to target modeling designs of future fusion reactors so that we can predict, among other things, the temperature and density of plasma in steady state reactor operation."

### GENE-X for Accelerated HPC

The GENE code family and especially its latest offspring, GENE-X, has helped with simulations of fusion physics. While many plasma physics codes model particle mixtures found in space, the sun, or stars, GENE was developed specifically for simulating conductive plasmas in facilities, which are based on magnetic confinement. The first GENE code was developed in 1999 at IPP as an open-source application and was gradually optimized for more detailed, three-dimensional simulations as computing power increased. By now, researchers use the code also to simulate the two most common magnetic confinement reactor concepts, the tokamak and the stellarator.

The newest version, GENE X, is now being prepared for deployment on supercomputers acceler-ated by graphics processing units (GPUs). The team is currently working

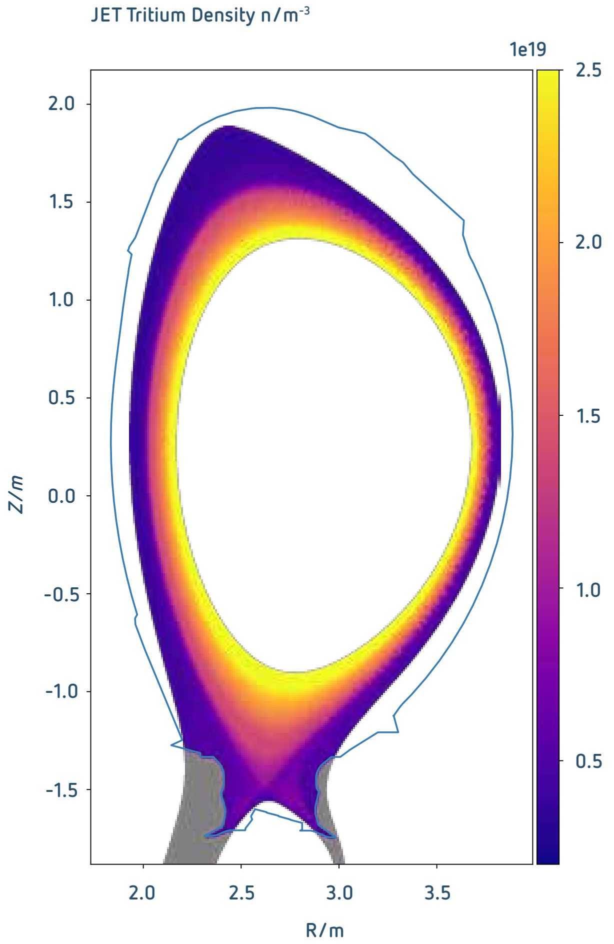


This figure shows the plasma density in a two dimensional cross section of the Joint European Torus (JET) and takes into account tritium, a heavy isotope of hydrogen. The blue line marks the wall, and the grey regions represent the divertor plates used in the experiment for heat removal.

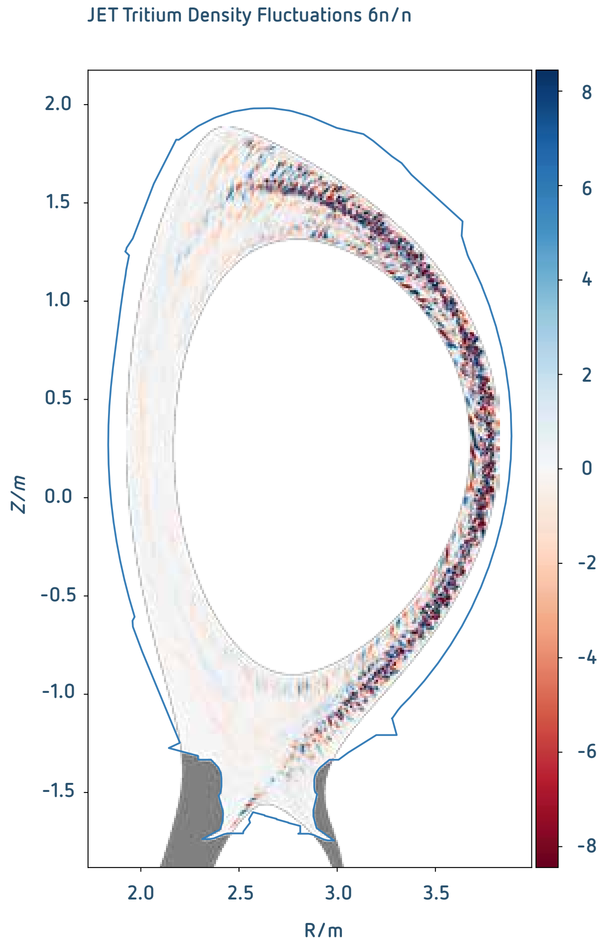


This figure shows the fluctuating component of the plasma in a two dimensional cross section of JET (in percentages relative to the background), which is caused by turbulence. It also considers tritium, a heavy isotope of hydrogen. The blue line marks the wall, and the grey regions represent the divertor plates used in the experiment for heat removal.

on implementing GENE X on the SuperMUC-NG Phase 2 (SNG 2) at Leibniz Supercomputing Centre (LRZ), which uses Intel chips. "Intel is a comparatively new player in the GPU market, where manufacturers such as NVIDIA and AMD have been dominating the market a while", says Jordy Trilaksono from IPP, who is responsible for adapting the code. "If we want to make our application widely available, GENE X should run as many of the major processor types as possible. We also want to become independent of specific vendors."

Thanks in part to the GENE codes, computer-supported fusion simulations have become increasingly precise over recent decades, and their results can now be compared with observations from experimental fusion facilities.

This iteration helps reduce costs for expensive reactors and experiments. For the modeling activities planned, the IPP team aims to process measurement data from the Joint European Torus (JET), a tokamak. This experimental fusion facility was the largest in Eu-rope to date. It operated until 2023 in Culham, UK, and at the end of its operational life managed to generate 69 megajoules of energy from just 0.2 milligrams of plasma or fuel – enough to power a TV, a stove, and a kettle for several hours. Beyond this record, JET delivered massive amounts of data, most of which remains unanalyzed. "JET contains unique data on hydrogen isotope mix-tures, which is particularly valuable for studying plasma turbulence", says Dr. Baptiste Frei, a researcher at IPP, who's conducting simulations for the project.





Of the more than 100,000 discharges of JET, the IPP team plans to analyze only three, from which GENE-X generates approximately 50 terabytes of data. "If the GENE X code runs faster, you can integrate more parameters into the model and process the resulting larger volumes of data much more quickly," explained Dr. Sajjad Azizi. The astrophysicist is part of the Computational X Support Team at LRZ and advises the IPP group in the transformation and implementation of GENE X. "By adapting the code to GPUs, researchers can leverage the full computing power of HPC systems like SNG 2. Scientific simulations are computationally demanding because they must account for many parameters and complex interactions between computing components. Pure CPU execution can easily become a bottleneck." The researchers have big plans: With GENE-X, they aim to model not just parts of the fusion facility, as they have done so far, but the entire plant – from the walls to the core.

## Combining Classical Modeling and Simulation with AI Tools

Given the growing use of artificial intelligence (AI) methods in plasma physics as well, the IPP team is taking a more fundamental approach to reworking GENE X. The code, originally written in Fortran, is first being expanded with an auxiliary C++ layer. This not only increases the range of programming models and tools with which the code can be adapted for future simulation tasks, but also makes it easier to accelerate simulations with GPUs and to combine with AI models or AI supported tools built in C++ languages. These could complement high-performance computing (HPC) methods. This porting work brings computational research another step closer to creating digital twins of fusion plasma and technologies. "The Fortran/C++ hybrid mode of GENE X can run on CPUs using OpenMP or on GPUs using OpenACC or OpenMP offload," Trilaksono added, highlighting further advantages of the code transformation. As a result, GENE X gains flexibility, making it easier to adapt to different hardware and operating systems.

Occasionally AI also helps with modifying the code. "It still makes mistakes, of course, but it is often faster at finding errors," said LRZ specialist Azizi. He recounted a build issue where code lines were repeatedly rewritten debugging messages and tested on the SuperMUC-NG 2 to narrow down the cause step by step. The team worked on this for some time, but AI found the simple solution very quickly – an outdated programming tool the specialists had started out with.

## Supporting Commercial Concepts

All researchers involved agree that the computational paradigm shifts make this porting work more complicated. "Previously, the transition to a new computing generation took days or weeks. The switch to GPUs is much more demanding," Ulbl said. "We already have a CUDA version of GENE X that we can certainly rewrite into SYCL, but that will take time." The first version of GENE X was adapted on supercomputers at the Max Planck Institute as well as on the Spanish supercomputer Mare Nostrum 5 in Barcelona. Early versions confirm researchers' expectations that GENE X would show substantial performance gains. The researchers saw an increase of up to a factor of 10, meaning it runs ten to thirteen times faster than its predecessor on only CPU systems.

The new code is currently being adapted to Intel GPUs. For this purpose, the LRZ support team regularly implements and tests parts of the program on SuperMUC-NG and its companion system SNG 2. Operating metrics from these test runs, along with comparison figures from similar applica-tions, provide clues about where algorithms can be modified and functions optimized. By autumn this year, the team hopes to run the first GENE-X-simulations on Intel GPUs. The resulting data is intended to help train AI models in the medium term and accelerate the optimization of fusion concepts.

"With GENE X, we want to enable the simulation of power plants and reactors that achieve at least a performance factor of 10 times larger than JET," said Ulbl. He noted growing interest in the fusion code. In many countries, politics are promoting fusion research and the development of new technologies, and the first startups and companies are entering the field. "This is another reason why we want to adapt and improve GENE X," Ulbl explained, "to support people and companies working on nuclear fusion and building the first commercial reactors." *Susanne Vieser*

**"Scientific simulations are computationally demanding because they must account for many parameters and complex interactions between computing components. Pure CPU execution can easily become a bottleneck."** Dr. Sajjad Azizi

Ehab Saleh checking on the Q.ANT processors installed at LRZ.

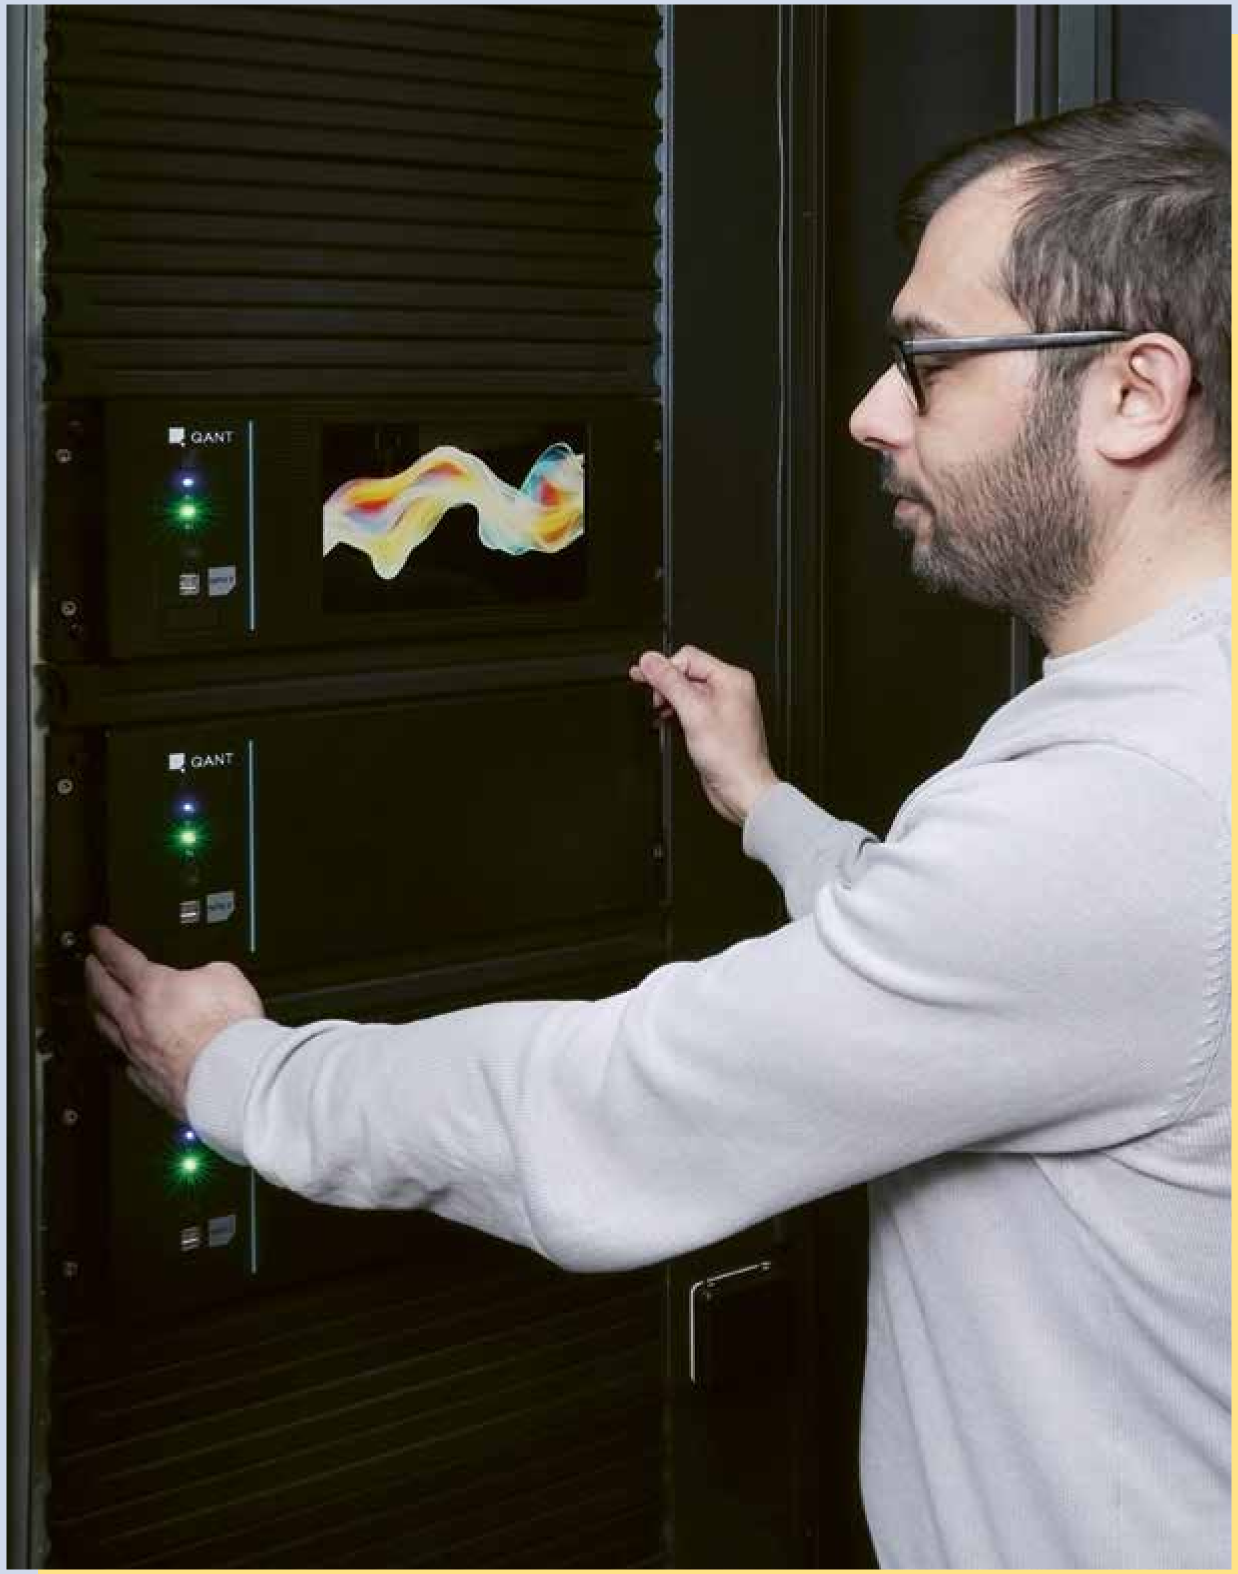



# LRZ Puts Photonic Computers to the Field Test

**LRZ experimented with photonic processors from Q.ANT in everyday operations. After working with the novel technology, LRZ staffers see the technology as having strong development potential.**

Colorful waves move from left to right across the screen, revealing cutting-edge technology working behind the scenes: Q.ANT's photonic specifically, Q.ANT's photonic analog accelerators, which the start-up calls it Native Processing Units (NPUs), use laser light, which photodetectors convert into digital electrical impulses when results need to be processed further by other components of a computer system. Recently, the Leibniz Supercomputing Centre (LRZ) has been experimenting with Q.ANT's technology, and according to the center, the results are promising.

"Within LRZ's Future Computing team, we evaluated the first two generations of these Q.ANT chips and compared their performance and potential uses in data centers, methods of artificial intelligence (AI), and high performance computing (HPC)," said computer scientist Dr. Josef Weidendofer. "In principle, they represent a highly promising technology that can significantly increase energy efficiency in data centers, especially during the training and inference of AI applications."

Photonic accelerators promise solutions for both the rapidly rising energy demand of AI clusters and for performance gains in computations such as matrix vector multiplications or nonlinear equations, which are widely used in HPC. Q.ANT discovered the computing power of light during its search for quantum computing technologies. Founded in 2018, the Stuttgart-based startup now has production capacity for stringing together a small series of its innovative co processors – a prerequisite for allowing potential users to test them. To this end, the Federal Ministry for Research, Technology and Space (BMFTR) financed a research procurement so that LRZ could also explore the potential of photonic computing for science and industry.

Second-generation Q.ANT photonic processor at the LRZ – an avant-garde computing architecture harnessing light to push the frontiers of high-performance and energy-efficient computation.

## Test Tasks from HPC and AI

Q.ANT supplies its analog photonic chips as plug-in cards for computer mainboards and currently also in its own servers. "The co-processors cannot perform all operations using light," explains LRZ researcher Dr. Ehab Saleh. "Therefore, a continuous data exchange between photonic and conventional processing units is required. This is why so-called complementary metal-oxide semiconductor (CMOS) technology is also included on the cards." The package additionally contains a function and software library: "You don't need to understand exactly how the accelerator works," Saleh added. "Users can retrieve commands for switching between technologies, functions for their own code in programming languages like C, Python, or Rust, as well as initial standard AI models."

While current graphics processing units (GPUs) draw up to 1,000 watts of power, Q.ANT's systems – consisting of server and card – require only 350 to 420 watts, the cards only 25 watts. For the tests at LRZ, three servers were procured to evaluate performance and analyze parallelization of photonic nodes. While the first-generation model contained a single card each, their successors were equipped with three – an approach to increase performance.
In addition to various computations, the processors were tested with typical AI tasks such as pattern and image recognition as well as the training and inference of smaller neural networks (MNIST, ReSet, KAN). LRZ staff measured metrics such as execution time, performance, energy consumption, error rates, prediction reliability, and training loss.

## Greater Performance and Higher Efficiency

Ultimately, LRZ researchers found the results are encouraging: while the optical element of the first generation performed around 300 million operations per second, the next generation achieved two billion operations per second. Q.ANT also integrated eight lasers into these newer cards. "In the best case, the second-generation chips operated 50 times faster," Weidendorfer said. A convolutional neural network accelerated image recognition by a factor of 25 thanks to this improvement.

The evaluation team indicated that future development stages for the optical elements would likely focus on frequency and workload management. The current-generation chip uses a two-gigahertz frequency, which could rise to 50 or 100 gigahertz if a given system's electronics can support the increase.

They also suspect that future cards could accommodate thousands or even tens of thousands of optical processing units. In addition, using different wavelengths of light could further increase computing power, and users could expand performance by adding more cards.

Further, performance gains do not come at the expense of energy efficiency. For typical HPC workloads, power consumption dropped by 50 to 84 percent from generation 1 to generation 2, depending on computational complexity; in more demanding tasks, the newer Q.ANT chip required roughly the same amount of power as a central processing unit (CPU). "Energy efficiency is highest when computations remain on the optical unit as long as possible," Weidendorfer noted. "Otherwise, components on the chip should be placed as close together as possible to minimize energy used for data transfer."

Like GPUs, analog photonic co-processors do not calculate exactly: they convert digital and analog data with 16 bit precision, but the noise generated by photodetectors must be taken into account. "AI applications operate well with lower precision, which makes analog photonic computing particularly suitable for them," Saleh observed.

At LRZ, researchers believe that Q.ANT chips could be ready for deployment in AI clusters within one or two further development cycles – boosting both energy efficiency and performance. Weidendorfer and Saleh also indicated that the processor can also emulate higher precision and accelerate supercomputing. For more precise results, complex tasks – similar to hand calculations – must be broken into multiple operations, which requires further and more sophisticated software tools. However, to achieve more precise results, complex tasks – such as mental arithmetic – would need to be broken down into multiple calculations, which requires further and more complex enhancements to the included software tools. A key step in the further development of photonic computing would be the emergence of a more diverse programming environment for this technology, featuring extensive libraries, languages, and additional software: "For evaluating next-generation systems," the researchers add, "it would be interesting to work with real application programs and optimize them for photonic processors. This would allow us to identify where the function library needs to be adapted and expanded." *Susanne Vieser*





HammerHAI is a collaborative project involving HLRS as well as the Gesellschaft für wissenschaftliche Datenverarbeitung mbH Göttingen, Leibniz Supercomputing Centre, Karlsruhe Institute of Technology, and SICOS BW. It has already begun offering AI computing and services to European startups and SMEs.

# HammerHAI Supercomputer Announced

**Designed to handle medium to large-scale AI workloads, the system will be deployed at the High-Performance Computing Center Stuttgart as a central component of the AI Factory HammerHAI.**

Launched in April 2025, the EuroHPC Joint Undertaking (EuroHPC JU) AI Factory HammerHAI has had a busy year establishing its service portfolio and building alliances across Baden-Württemberg's thriving artificial intelligence (AI) community. In an important step forward for the project, the EuroHPC JU in March 2026 signed a contract to locate an AI-optimized supercomputer at the HighPerformance Computing Center Stuttgart (HLRS), coordinator of HammerHAI. The system will provide powerful new capabilities for artificial intelligence, machine learning, and data science, strengthening European science, industry, small and medium-sized enterprises, and startups.

Hewlett Packard Enterprise (HPE) will manufacture the new HammerHAI supercomputer, which is based on the liquid-cooled NVIDIA GB200 NVL4 architecture. Combining NVIDIA Grace CPUs with NVIDIA Blackwell GPUs and scaled with NVIDIA Quantum-X800 InfiniBand networking, the NVIDIA GB200 NVL4 by HPE will offer more than 15 exaflops of peak AI inference performance. It will integrate the VAST Data DASE storage architecture as well as a partition based on AI-optimized inference engines and hardware accelerators from Netherlands-based Axelera AI. The HPE Morpheus Enterprise software will be used as a unified AI control plane, enabling automated provisioning, governance, and workload lifecycle management.

This configuration will make the HammerHAI supercomputer a powerful tool for handling medium to large-scale AI workloads for machine learning and artificial intelligence. It will incorporate a cloud-native software stack familiar to the AI community, making it straightforward to migrate or scale applications from local systems or commercial cloud environments. The system will be configured to support research and technology development in disciplines prioritized with the HammerHAI consortium, with an emphasis on engineering, manufacturing, automotive, and mobility.

Delivery of the HammerHAI supercomputer is scheduled for the second quarter of 2026 and it is expected to go into operation in the second half of 2026. The system will support the types of applications typically handled by commercial cloud AI services, while being operated in Germany and in accordance with EU data security regulations. Access to this publicly funded resource will be free of charge to eligible European system users.

HLRS is coordinating HammerHAI in collaboration with the Gesellschaft für wissenschaftliche Datenverarbeitung mbH Göttingen, Leibniz Supercomputing Centre, Karlsruhe Institute of Technology, and SICOS BW. The AI Factory is co-funded by the European Commission, the German Federal Ministry of Research, Technology and Space (BMFTR), the Baden-Württemberg Ministry of Science, Research and the Arts, the Bavarian State Ministry of Science and the Arts and the Lower Saxony Ministry of Science and Culture.

### HammerHAI partners with EPCC for UK AI Antenna

In other related news, the EuroHPC Joint Undertaking in late 2025 announced the creation of 13 "AI Factory Antennas," a new category of competence centers for artificial intelligence that will complement and extend the activities of AI Factories across Europe. Among the winning proposals was the UK AI Factory Antenna (UKAIFA), led by EPCC at the University of Edinburgh. This action includes a partnership with HammerHAI, which will serve as a bridge between the United Kingdom and the European AI Factory network. The partners will also collaborate on a variety of efforts to accelerate the adoption of artificial intelligence in SMEs, startups, industry, and the public sector.

Under the AI Antenna agreement, HammerHAI will provide access for the UKAIFA user community to its AI-optimized supercomputer, and will work with UKAIFA to enable secure, federated data sharing and cross-platform compatibility. This collaboration will also enable UKAIFA to integrate with the wider EuroHPC AI Factory network, leveraging HammerHAI's connections to participate in European AI infrastructure development efforts.

The UK AI Factory Antenna will build on the expertise gained over more than 20 years of collaboration between EPCC and HLRS. Working together, the two organizations have focused on bringing high-performance computing and artificial intelligence to European industry. EPCC launched the highly successful Fortissimo project series, which since 2012 has worked with more than 3,000 companies to support their use of supercomputing to develop innovative products, improve efficiency, and compete more effectively in the global marketplace. HLRS continues to coordinate the latest iteration of this project, called FFplus. *Christopher Williams*

**"With the HammerHAI AI factory, HLRS and its partners are pooling their resources and expertise along the entire AI value chain in order to drive scientific and economic progress in Baden-Württemberg and beyond. HammerHAI is thus actively shaping our AI landscape while strengthening Europe's innovative strength and technological sovereignty."**
Petra Olschowsky, Baden-Württemberg Minister of Science, Research and Art





# ICON Team Wins Gordon Bell Prize for Climate Modelling

**Multi-institutional team uses JUPITER to run well-known ICON climate model at improved 1.25-kilometer resolution.**

Representatives of the research team accepted the Gordon Bell Prize for Climate Modelling at SC25 in St. Louis.

A research team involving the Jülich Supercomputing Centre (JSC) has won the prestigious Gordon Bell Prize in the category of climate modelling. The Gordon Bell Prizes are among the most important awards in the field of high-performance computing (HPC). They are granted annually by the Association for Computing Machinery (ACM) and acknowledge outstanding achievements in distinct research categories. ACM in 2023 introduced "Climate Modelling" as a new category, and the award recognizes innovations that improve our understanding of the Earth's climate system and climate change.

The jury was convinced by the team's novel approach, which enables the simulation of the Earth system with an unprecedented level of spatial detail. To achieve this, the team used Europe's first exascale supercomputer, JUPITER at Forschungszentrum Jülich, as well as the Swiss supercomputer Alps – both of which rank among the fastest systems in the world.

### Full-Earth Simulation at 1.25-Kilometer Resolution

At the heart of the work lies an enhanced version of the ICON climate model that incorporates the atmosphere, oceans, surfaces, and sea ice, and operates at a horizontal resolution of approximately 1.25 kilometers. For comparison, climate simulations with the ICON model have typically operated at a resolution of around 5 to 10 kilometers until now. This new level of detail enables the model to accurately simulate key processes – such as convection, extreme precipitation, and ocean eddies – to be represented far more realistically. Leveraging 85% of JUPITER's computing resources, the team achieved a record-breaking rate of 145.7 simulated days per real-time day.

### Algorithmic Innovations and Powerful Hardware

The record-breaking run was made possible by algorithmic innovations and the targeted use of JUPITER's powerful NVIDIA GH200 hardware, which the team used as part of the JUPITER Research and Early Access Program (JUREAP). The program gave researchers access to the system during its early testing phase.

The team, led by the Max Planck Institute for Meteorology (MPI-M) and the German Climate Computing Centre (DKRZ), succeeded in fully exploiting the capabilities of NVIDIA's GH200 superchips. These combine a Grace CPU and a Hopper GPU into a single, large superchip. They were supported in this effort by researchers from the University of Hamburg, JSC, the Swiss National Supercomputing Centre (CSCS), ETH Zurich, and NVIDIA.

This simulation is one of the first model runs to use the Grace CPU and Hopper GPU simultaneously for different, yet closely linked, model components. This marks a crucial

"This achievement reflects not only powerful hardware, but also the close collaboration needed to use JUPITER efficiently for Earth system modeling at scale." Dr. Lars Hoffmann, JSC

step toward making efficient use of JUPITER's heterogeneous exascale architecture and enabling ever more complex and accurate climate simulations.

For the calculations, ICON was scaled to 5,120 of JUPITER's compute nodes. These runs are among the largest and most technically demanding simulations conducted on the system to date. From JUPITER's first development module, JEDI, to the full system deployment, JSC staff configured the supercomputer for the ICON code and prepared JUPITER for the simulations. JSC also supported the optimization of ICON for the hybrid Grace-Hopper architecture by conducting tests and evaluating performance and played a key role in scaling the simulations to the exascale level.

The foundation of this success lies in the persistent, long-term development efforts of the ICON community and its partner institutions. The award highlights ICON's ability to utilize emerging computing technologies on a large scale, while also demonstrating the benefits of prolonged collaboration between climate scientists, software engineers, and HPC experts.

The award ceremony took place at the Supercomputing Conference 2025 (SC25) in St. Louis. The team also received an HPCwire Award in recognition of their work, demonstrating how the next generation of supercomputers can improve our understanding of global climate dynamics.
*Eric Gedenk*

Related publication
Klocke, D., Frauen, C., Engels, J. F., Alexeev, D., Redler, R., Schnur, R., Haak, H., Kornblueh, L., Chegini, F., Römmer, M., Hoffmann, L., Griessbach, S., Bode, M., Coles, J., Gila, M., Sawyer, W., Calotoiu, A., Budanaz, Y., Mazumder, P., Copik, M., Weber, B., Herten, A., Bockelmann, H., Hoefler, T., Hohenegger, C., Stevens, B. (2025). Computing the Full Earth System at 1 km Resolution. SC '25: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis, St. Louis, MO, USA, 2025, 125-136. https://doi.org/10.1145/3712285.3771789

The further development of ICON for the efficient use of exascale HPC systems was provided by the German Federal Ministry of Research, Technology and Space (BMFTR) funded project WarmWorld, under grant numbers 01LK2202, 01LK2203 and 01LK2204 and the BMFTR funded project IFCES2-Scalexa under grant 16ME0692.

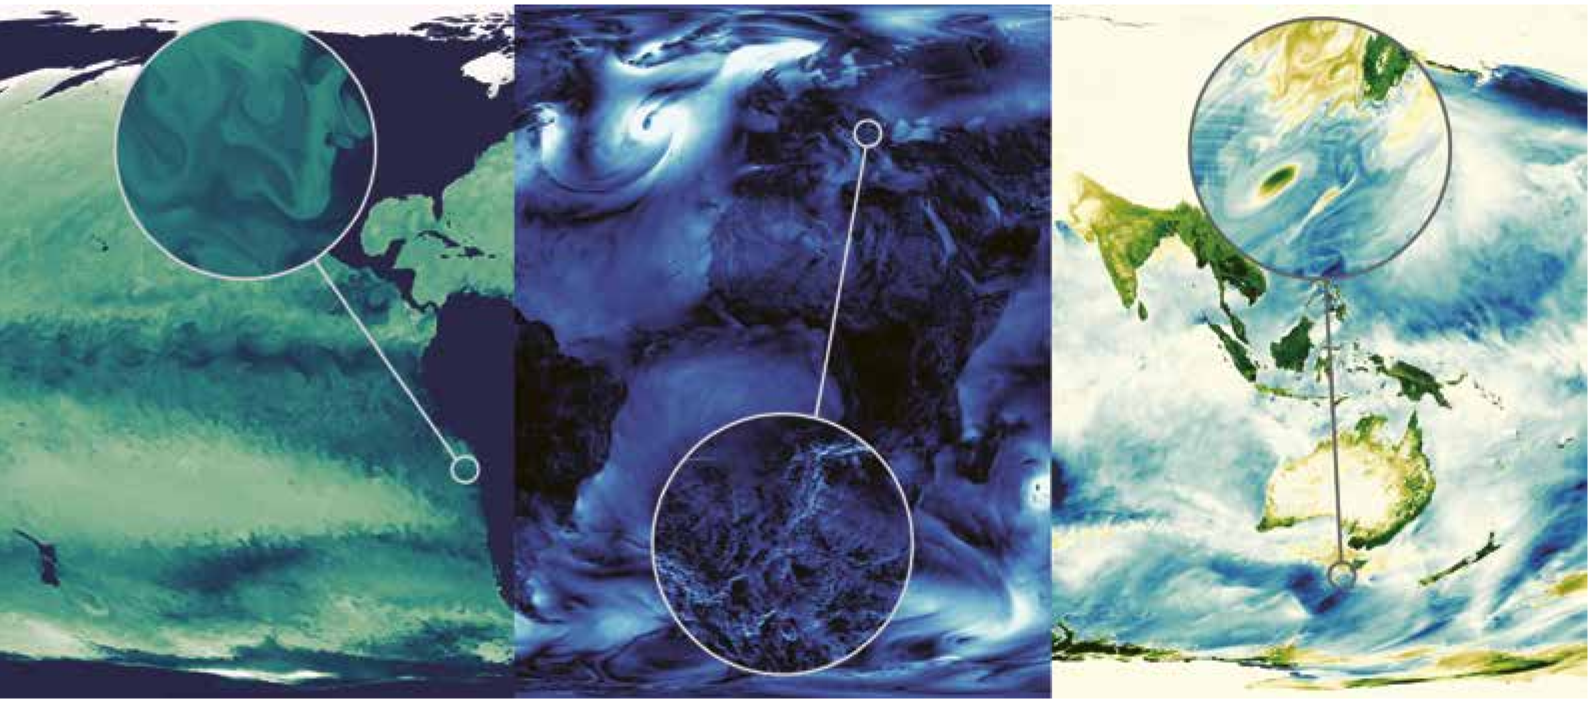

Highlights of the ICON simulation presented in the study, showing phytoplankton, surface wind, and air-sea $CO_2$ flux (from left to right).





# Towards Europe's Hybrid HPC-Quantum Infrastructure: JADE and Ruby Quantum Processors Inaugurated

**Dignitaries from government and industry gathered at the Jülich Supercomputing Centre to welcome two new quantum processors into the European Union's increasingly diverse mix of quantum systems.**

The project, "High-Performance Computing and Quantum Simulator hybrid" (HPCQS) inaugurated the two quantum processors JADE and Ruby in a festive event on November 13, 2025. JADE (Julich neutral Atom DEvice) is hosted at Forschungszentrum Jülich (FZJ) in Germany and Ruby at *Commissariat à l'énergie atomique et aux énergies alternatives* (CEA) in France. The inauguration spanned three locations, with coordinated celebrations at FZJ, CEA and the European Commission's premises. A live stream was also available for remote participants. The event started with a presentation of the HPCQS project, with two use cases demonstrated to showcase the integration of the two quantum processors.

During the joint part of the event, Oscar Diez (Head of Quantum Computing, European Commission), Daniel Opalka (Head of Unit Research & Innovation, EuroHPC JU), Bruno Bonnell (representing the French government as Lead of the France 2030 investment plan) and Roland Krüppel (Deputy Head of Division for Microelectronics, Supercomputing and Robotics, German Federal Ministry of Research, Technology and Space; BMFTR) congratulated the project for this great achievement and highlighted the importance of this milestone for European leadership and technological sovereignty in the field of quantum computing. The project was represented by Wilhelmine Bach, Communication Manager of the HPCQS project, EURICE – European Research and Project Office GmbH; Francesca Vaccaro, Research & Impact Manager at EURICE GmbH and responsible for Communication, Dissemination and Exploitation in the HPCQS project; and Kristel Michielsen, JSC Director and HPCQS project coordinator. A panel with the HPCQS work package leaders concluded the joint part of the event.

"We are entering a new phase in which quantum and classical computing begin to evolve together. This is more than a technical achievement – it signals a new paradigm for computation, one that will shape the future of scientific discovery."
Prof. Dr. Kristel Michielsen, JSC Director

EVENTS

To round off the event locally, participants in Jülich had the opportunity to attend a guided tour of JADE and the other quantum computers integrated into JUNIQ, the Jülich UNIfied Infrastructure for Quantum Computing, located in the JUNIQ building, as well as in the JSC's Cryolab.

The quantum processing units (QPUs) built in JADE and Ruby were prefabricated by Pasqal. The technology is based on arrays of neutral atoms that are trapped and manipulated by lasers in programmable geometries for quantum operations. These systems operate at room temperature and with low power consumption, making them particularly robust and scalable.

As demonstrated during the inauguration, the quantum processors have been successfully integrated into existing high-performance computing (HPC) environments, which marks a major milestone of the HPCQS project. A specific HPC-quantum computing software stack was developed by the consortium, relying on industrial and open-source software components. It supports applications in optimization, simulation, and machine learning.

This integration provides a key step towards a federated quantum–HPC infrastructure, enabling users from research and industry to develop hybrid algorithms and solve complex problems such as industrial battery design, drug discovery, and optimization challenges in finance and traffic management.

The HPCQS project is funded by the EuroHPC JU, which receives support from the EU's Horizon 2020 research and innovation program and Germany, France, Italy, Ireland, Austria, and Spain in equal parts. *Anna Lührs*

More information: https://www.fz-juelich.de/en/news/archive/press-release/2025/jade

Quantum processor JADE at JSC.

# LRZ and RIKEN R-CSS Leadership Meet to Sign Cooperation Agreement Toward Greater Computing Energy Efficiency

**On the sidelines of the SC Asia event, management from the two centers met to sign an agreement aimed at improving sustainability in high-performance computing.**

The RIKEN Center for Computational Science (R-CCS) and the Leibniz Supercomputing Centre (LRZ) have long traditions of improving energy efficiency in high-performance computing (HPC). Recently, during SC Asia 2026 in Osaka, Japan, leadership from the two centers met and signed a memorandum of understanding (MoU) that formalizes a collaboration to exchange experiences as well as launch their own research projects.

At the event, Prof. Satoshi Matsuoka, Director of R-CCS, and Prof. Dieter Kranzlmüller, Director of the LRZ, spoke about how the MoU will combine their expertise in the energy-efficient and resource-efficient operation of supercomputers. Both computing centers are currently planning the successors to their current flagship HPC systems. While the concepts for LRZ's Blue Lion and R-CCS's Fugaku-Next differ, both systems will integrate accelerators that enable new artificial intelligence (AI) capabilities, but will also significantly increase power consumption.

Through their MoU, both institutes agreed to establish the exchange of data and metrics related to the operation of directly water-cooled high-performance systems under heavy load. The agreement includes sharing technical resources and bringing together support and HPC staff through workshops and mutual visits.

"High-performance computing is an essential tool for science and research, from astrophysics and life sciences to environmental sciences," said Prof. Dieter Kranzlmüller, Chairman of the Board of Directors at LRZ. "However, these highly complex systems consume a lot of power, but the potential for optimization is just as great. I am very much looking forward to working closely with RIKEN – it will benefit science and help us further increase the efficiency of our supercomputers."

The two academic data centers will jointly analyze the thermal behavior of hardware components and explore the thermal sweet spot to determine the optimal operating temperature that balances high energy efficiency, computing performance, and system availability. The MoU also encourages the centers to assess the risks of water cooling and the development of benchmarks for optimal system temperature. The focus of these efforts is connected to everyday operations – HPC specialists will evaluate how scientific compute and simulation workloads can be implemented, optimized, and managed with energy awareness, and will develop tools together for this purpose. The centers have also planned practical scenarios with guidelines for users on energy-conscious programming and computing.

In light of the increasing energy demand of accelerated HPC systems, the centers are also focused on heat recovery and the development of district heating systems in the vicinity of data centers. The collaboration between R CCS and LRZ is planned to run until 2030, with considerations to further expand the cooperation and strengthen the partnership between the two centers. *Susanne Vieser*

EVENTS

# HLRS Future Computing Group Holds First Annual Workshop

**The event brought vendors of emerging computing hardware together with researchers in the HPC community, highlighting opportunities and challenges in a diversifying technology landscape.**

In the foreseeable future, traditional technologies for high-performance computing (HPC) based on CPU and GPU processors could reach their limits in terms of performance, economic viability, and environmental sustainability. This has launched the hardware community on a search for alternative paradigms and architectures that could open the door to improvements in speed and energy efficiency. Quantum computing and neuromorphic computing are among the best known approaches, although numerous other avenues are also being explored, including concepts that build on and extend traditional technologies found in x86, GPU-accelerated, and ARM-based architectures.

In parallel, progress is being made on software development for new computing frameworks. These include programming models and libraries that would make it easier to use established codes with new hardware, and to implement heterogeneous workflows across multiple hardware types. Many anticipate that the develolpment of hybrid approaches holds more potential for advancing HPC capabilites than expecting a single hardware type to replace all others, meaning that one goal for future HPC systems would be to integrate different hardware types and programming methods seamlessly. In this way, various elements within complex algorithms could then be distributed to specialized processors that run them most efficiently.

Recently, the High-Performance Computing Center Stuttgart (HLRS) launched a new initiative called the Future Computing Group to explore which emerging technologies for high-performance computing hold the greatest potential to address the needs of HLRS's scientific and industrial user communities (for more information on the Future Computing Group, visit page 38). Working closely with technology companies, the group will also provide HPC expertise that will support industrial research on next-generation computing hardware. In its first collaboration, for example, it is working with Barcelona-based firm Openchip to investigate its RISC-V architecture-based, system-on-a-chip solutions for high-performance computing workflows.

On March 16-17, 2026, HLRS also hosted its first Future Computing Workshop, an event designed to promote discussion, networking, and collaboration surrounding new computing paradigms. Hardware vendors, researchers, and computing facility operators offered in-depth looks at emerging technologies and methods, including the potential advantages they offer and the challenges they face in practice. By including perspectives from the academic research community, the event also enabled hardware vendors to gain insights into user requirements they will need to consider to ensure that their products are widely adopted.

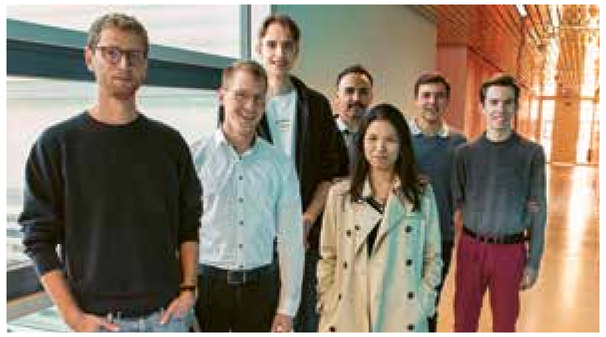

Led by Johannes Gebert (second from left), the Future Computing Group brings together a dozen specialists in HPC hardware and software from across multiple HLRS departments.

The first Future Computing Workshop attracted more than 65 participants, including hardware vendors, computer science researchers, and operators of HPC centers.

Dr. Johannes Gebert leads HLRS's Future Computing Group and organized the event. "Computer scientists, domain-specific researchers, computing centers, and hardware vendors deal with widely different challenges and incentives," he explained. "We established the workshop as a platform for people to understand one another, accelerating the deployment of high-end computing paradigms."

In addition to spotlighting innovations in hardware, the Future Computing Workshop illuminated challenges facing the HPC community resulting from a diversifying technology landscape. Existing research codes have often taken years to develop, and cannot easily be ported to new computing hardware, if at all. Because of the length of procurement cycles, HPC centers must also be able to plan for the future, requiring a clear understanding of what the user community will need in 5-7 years. As technologies evolve, it will be important that new hardware is well-suited to the scientific problems that need to be solved, that the answers it delivers are reliable, and that it is easily programmable. Moreover, regardless of how fast new processing technologies become, the trend toward ever larger simulations and data-driven methods means that improving memory bandwidth capabilities will be at least as important. Otherwise, the ability to move and manage the resulting data efficiently could remain a major rate-limiting step.

These kinds of observations suggest that while new computing architectures hold great potential, they will be most successful if developed in partnership with potential users. The Future Computing Workshop aims to support this essential dialogue among stakeholders in order to facilitate the success of next-generation technologies.

More than 65 participants attended the inaugural workshop. Considering the very positive response from across the hardware and research communities, a second annual event will take place on March 15-17, 2027.
*Christopher Williams*

> **"We established the workshop as a platform for people to understand one another, accelerating the deployment of high-end computing paradigms."**
> Dr. Johannes Gebert

# News Briefs

## JSC

### Computing With Light: Forschungszentrum Jülich and Q.ANT Launch Collaboration on Photonic Computing

JSC and Stuttgart-based technology company Q.ANT are launching a four-year development partnership to systematically explore the potential of photonic computing for high-performance computing, and to identify new application opportunities. Together, they aim to integrate Q.ANT's photonic processor into JSC's modular computing infrastructure, and to develop new methods and algorithms that make optimal use of the unique properties of computing with light.

Photonic computing uses light waves instead of electronic currents to process data. This enables complex mathematical operations to be carried out with virtually no heat generation, and with significantly lower energy consumption. Given the growing energy demands of data centres, light-based computing is considered a promising approach for creating more sustainable computer architectures for the future.

As part of the partnership, JSC will procure a Photonic Native Processing Server from Q.ANT to test hardware and software under real HPC conditions.

To read the full article, please visit: https://www.fz-juelich.de/en/news/archive/press-release/2025/computing-with-light-forschungszentrum-julich-and-q-ant-launch-collaboration-on-photonic-computing

### JUPITER is Europe's First Supercomputer to Reach 1 ExaFLOP/s

JUPITER at JSC is the first supercomputer in Europe to reach the milestone of 1 ExaFLOP/s in 64-bit precision. For 8-bit calculations, which are commonly used to train large AI models, the theoretical performance exceeds 40 ExaFLOP/s. JUPITER ranks as the fourth-fastest supercomputer in the world and the most energy-efficient system in the exascale class. Not only does the system represent a technological milestone, but it also strengthens Europe's digital and scientific sovereignty.

JUPITER opens new possibilities for training large AI models and performing scientific simulations with unprecedented complexity and detail in fields such as climate, energy, medicine, and materials research. Its enormous computing power is expected to enable the prediction of extreme weather events, such as heavy rainfall or heatwaves, at previously unattainable spatial resolutions; accelerate the development of sustainable energy systems; and provide deeper insights into complex biological processes in proteins, cells and the brain, forming the basis for new therapies.

To read the full article, please visit: https://www.fz-juelich.de/en/news/archive/press-release/2025/europes-first-supercomputer-reaches-1-exaflop-s



### Project "European HPC Professional Traineeship Programme" (HPCTRAIN) Launched

In January 2026, the HPCTRAIN project was officially launched. Spanning four years, it will connect emerging talent with Europe's leading supercomputing centers and industry to develop advanced digital skills and drive innovation. The project consortium, coordinated by Forschungszentrum Jülich, comprises 12 universities, supercomputing centres and companies. Forschungszentrum Jülich is one of the hosting sites for the trainees during the project.

The HPCTRAIN project aims to strengthen Europe's high-performance computing (HPC) skills ecosystem by offering hands-on training and mentorship opportunities to students and early-career professionals across Europe. HPCTRAIN will deliver 200 structured traineeships, matching talented students and early-career professionals with leading EuroHPC Hosting Entities, research institutions and companies using HPC to innovate and grow. Trainees gain hands-on experience with advanced HPC technologies, collaborate with industry on real-world challenges, and build the digital skills employers demand across Europe. Industry, meanwhile, gains access to HPC to help drive its innovation.

To read the full article, please visit: https://www.fz-juelich.de/en/jsc/news/news-items/news-flashes/2026/hpctrain-launched

### Quantum Pioneer Kristel Michielsen is Director of the JSC

Kristel Michielsen is Director of the Jülich Supercomputing Centre (JSC) since August 2025, in addition to Thomas Lippert who continues to perform in this role. Her new role is aligned with a professorship for Quantum Computing & Modular Supercomputing at the University of Cologne.

The computational physicist has led the Quantum Information Processing (QIP) research group at JSC since 2009 and held a professorship for Quantum Information Processing at RWTH Aachen University from 2009 until 2025. Her QIP group has ample experience in performing large-scale simulations of quantum computers and annealers, and in benchmarking and studying prototype applications for this new computing technology in optimization, simulation, machine learning, and quantum AI. She has also led the Jülich UNified Infrastructure for Quantum computing (JUNIQ) since 2019. In 2023 she became the co-leader of the division HPC for Quantum Systems at the JSC.

To read the full article, please visit: https://www.fz-juelich.de/en/jsc/news/news-items/news-flashes/2025/michielsen-director-jsc

## HLRS

### HLRS Wins Environmental Management Prize

HLRS was awarded the 2025 Environmental Management Prize in the category "Best Measure for Improving Energy Performance." Organized by the German Federal Ministry for the Environment, Climate Action, Nature Conservation, and Nuclear Safety (BMUV) together with the Austrian Federal Ministry of Agriculture and Forestry, Climate and Environmental Protection, Regions, and Water Management (BMLUK), the award was presented in Berlin during the conference "30 Years of EMAS – a Sustainable Success". The Environmental Management Prize recognizes HLRS's use PowerSched, an innovative energy management solution cdeveloped in collaboration with Hewlett Packard Enterprise (HPE). Implemented in 2024 on HLRS's Hawk supercomputer, PowerSched distributes available power across a supercomputer's compute nodes in a targeted manner, continuously monitoring and regulating the distribution of power in order to maintain an optimal balance and adapt to changes in HLRS's energy and load budget. Using PowerSched, HLRS reduced the power consumption of Hawk by approximately 20 percent without negatively affecting computational performance. In other sustainability news, HLRS successfully renewed its certification under the Eco-Management and Audit Scheme, which it has held continuously since 2019.

A new Environmental Statement is also now available on the HLRS website at hlrs.de/sustainability.

To read the full article, please visit: https://www.hlrs.de/news/detail/hlrs-receives-2025-environmental-management-prize

## Hawk Gets a Second Life Through Refurbishing

Following the decommissioning of its supercomputer, Hawk, HLRS and its technology partner HPE took a sustainable approach to its disposal. In an effort coordinated by HPE Financial Services, Hawk's hardware – including 4,096 computer nodes, 8,192 processors, and 65,536 memory modules – was carefully dismantled and shipped to the company's refurbishing center in Erskine, Scotland. There, the data were securely erased and the hardware was tested for functionality. Following this audit, HPE was able to refurbish and sell more than 90% of Hawk's infrastructure, giving it a second life with several customers, including in the aerospace technology industry. According to HPE this refurbishing avoided approximately 2,800 metric tons of $CO_2$ emissions – similar to emissions generated by 25,000 hours of individual air travel. Around 13.7 metric tons of electronic waste was diverted from landfills, equal to the annual e-waste disposal of 1,700 individuals in Germany. This effort resulted in energy savings of almost 64,000 kWh, a similar amount to that required to charge an electric car fully 1,000 times.

**To read the full article, please visit: https://www.hlrs.de/news/detail/hawk-gets-a-second-life-through-refurbishing**

## A New Initiative to Strengthen the European Auto Industry

The European Connected and Autonomous Vehicle Alliance (ECAVA) is a European Commission-led initiative to strengthen Europe's competitiveness and digital sovereignty in the automotive sector. Bringing together experts and stakeholders from across Europe, it aims to promote collaboration and accelerate innovation in this important industry, including in the integration of improved software, chips, AI applications, and autonomous driving technologies. The project officially launched on February 5-6, 2026 at a workshop held in Brussels. Following a selective application process, HLRS was chosen to join ECAVA and participated in the Brussels conference as a member of the AI Model Development and Data Working Group. HLRS's participation builds on its long history of providing high-performance computing infrastructure and expertise for companies in the automotive sector. As coordinator of the AI Factory HammerHAI, HLRS is also in the process of developing new capabilities that will be highly relevant for the development of next-generation automobile technologies. As an ECAVA member, HLRS and HammerHAI will gain an improved understanding of what manufacturers in the automotive industry need from HPC centers and AI Factories in order to use their resources efficiently. This includes end user requirements with respect to data sharing, data pooling, and foundational AI model development for autonomous vehicles.

**To read the full article, please visit: https://www.hlrs.de/news/detail/hlrs-joins-eu-initiative-for-digital-innovation-in-the-auto-industry**

## Bastian Koller Named an HPCwire "Person to Watch"

Dr. Bastian Koller, HLRS Managing Director, was named one of twelve "People to Watch 2026" by HPCwire. The selection focuses on Dr. Koller's leadership roles in international collaborative projects like EuroCC, CASTIEL, and FFplus, EuroHPC Joint Undertaking Centers of Excellence like EXCELLERAT, and the AI Factory HammerHAI. These efforts have contributed to expertise development in high-performance computing and artificial intelligence across Europe, HPC and AI uptake in industry and SMEs (small and mid-size enterprises), and initiatives to improve Europe's digital sovereignty. HPCwire is a leading publication covering the international high-performance computing industry. For 24 years, its "People to Watch" program has recognized HPC industry professionals who are driving innovation and increasing the benefits of high-performance computing for society.

**To read the full article , please visit: https://www.hlrs.de/news/detail/bastian-koller-named-an-hpcwire-person-to-watch-2026**

## HLRS at the Hannover Messe

For the first time, HLRS participated in the Hannover Messe, an international trade fair that gathers leading players from industry, research, and technology to present innovations in automation, energy, digitalization, industrial IT, and advanced manufacturing. As a partner in the Baden-Württemberg International booth, HLRS highlighted the projects EXCELLERAT P2 and HammerHAI, which are focused on improving capabilities for simulation and artificial intelligence in industry. HLRS research scientist Rishabh Saxena presented a talk in the convention conference program explaining lessons being learned in the HammerHAI AI Factory. He described how modern AI software stacks, orchestration, and optimization techniques can interact with high-performance compute environments to enhance industrial AI scalability and support data sovereignty. The session highlighted design principles, architectural patterns, and future perspectives for building robust and adaptable industrial AI platforms.





LRZ

## LRZ Officially EMAS certified

In late 2025, LRZ was certified under the Eco-Management and Audit Scheme (EMAS) after implementing an updated environmental management system. The certification serves as international recognition that the centre operates at the highest sustainably standards. To this end, staff from all departments collaborated with the two LRZ environmental management officers for over a year to compile information and key performance indicators on the centre's environmental impact. Staff also installed measuring devices, reviewed legal requirements, and adapted processes. The efficiency data showed excellent results, with LRZ achieving a power usage effectiveness (PUE) rating of 1.23, a cooling efficiency ratio (CER) of 8.78, and a water efficiency ratio (WER) of 2.30. Since 2012, LRZ has used electricity exclusively from renewable sources. In 2024, consumption stood at around 40,500 megawatt hours, 98% of which was used to power computing resources.

**To read the full article , please visit: https://tiny.badw.de/2deo9C**

## LRZ Collaborates On A New Era of Biomolecular Simulation

Researchers have developed a hybrid quantum classical simulation pipeline to model complex biomolecules, combining GPU accelerated supercomputing with quantum processors. The collaboration includes international universities, public research centers, and technology providers, with the work presented around NVIDIA GCT summit in March 2026. Using CUDA Q, quantum models simulate chemically active regions, while classical methods handle the broader environment, enabling detailed, scalable simulations. An initial study focused on a G protein coupled receptor (GPCR), which is important for drug design. Quantum computations ran on Euro-Q-Exa, a 54 qubit system at LRZ, with GPUs managing electronic sampling and postprocessing on the center's BayernKI and NVIDIA's EOS system. This hybrid approach enhances computational power and precision, paving the way for realistic, large-scale biomolecular simulations and deeper insight into molecular mechanisms.

**https://tiny.badw.de/IUD2I**

## LRZ Supports TUfast Racing Team with Computational Muscle

The Technical University of Munich's student club, TUfast Racing, is using LRZ's CoolMUC supercomputer to improve the design of its electric Formula Student race car, including aerodynamics, battery performance, and motor performance for both driver piloted and autonomous runs. About 100 students work on modeling all vehicle components and behaviors, such as designing the rear wing and chassis, simulating airflow and thermals, analyzing large datasets, and optimizing parts before production. CoolMUC enables full vehicle simulations that small local computers cannot accomplish. TUfast's current vehicle model, the xb026, builds on a lightweight carbon monocoque, four 35-kilowatt motors and a 600-volt battery. Simulations help place cooling inlets and reduce drag, improving vehicle performance. The project also teaches students HPC skills and collaborative engineering. Beyond race speed, Formula Student competitions judge teams on their vehicle's concept, design, cost management, and business planning. TUfast plans to compete in four European races in summer 2026, aiming for wins and strong overall results at events including Formula Student Germany.

**https://tiny.badw.de/rQiZ1J**

Johannes Gebert leads HLRS's new Future Computing Group.



# Staff Spotlight: From Automotive Engineering to the Future of Computing

## Johannes Gebert, HLRS

Johannes Gebert grew up close to and with a strong interest in computers and engineering. During his bachelor's studies in automotive engineering at the University of Applied Sciences Esslingen, he worked at his parents' audio and video technology company. He then transitioned to the University of Stuttgart, earning a master's degree in automotive and engine engineering and a PhD in biomechanics on high-performance computing (HPC) systems.

During his studies, Gebert interned at both Audi and Mercedes-Benz. Those experiences helped him realize his interest in cars was closely coupled with his exposure to the broader information technology industry growing up. "Through my parents' company, I had a connection to IT from early on," he said. "In the end, I realized that automotive engineering and computer science are both dealing with complex systems. The merging of these two things is an interesting place for me academically and professionally." His internship experiences exposed Gebert to automotive engineers who use computational power to model LeMans Prototype and F1 crash simulations, design vehicle components, and simulate aerodynamics. He realized specializing in computer science while working on his PhD would increase the chances he could work on the research side of such challenges.

To pursue his doctoral studies, he found an appointment at the High-Performance Computing Center Stuttgart (HLRS). In his first role at HLRS, Gebert worked on a biomechanics project, again covering structural mechanics. He saw how two different disciplines used high-performance computing (HPC) in exciting and similar ways. He also learned that success came from bringing together computer scientists, domain experts, mathematicians, and engineers to advance the state-of-the-art.

As he finished his studies, Gebert's vision of how disparate research domains overlap dovetailed nicely with rapid changes happening in HPC. For decades, computers got faster largely due to computer chips shrinking and becoming more efficient, resulting in large performance gains from one chip generation to the next, but that rapid growth has slowed in the last decade. The inability to maintain the same pace of performance development is colliding with both the rapid rise and commercialization of artificial intelligence and increased investment in new computing paradigms such as quantum computing and neuromorphic computing. Together, industry changes are both challenging scientists and engineers to learn new computational methods and asking HPC centers to reexamine their roles. In addition to offering the scientific community a single large, flagship system, HPC centers are increasingly building heterogeneous computational ecosystems, offering researchers a suite of hardware and software solutions.

To that end, HLRS in 2025 launched the Future Computing Group and tapped Gebert to lead the team. The group is charged with exploring emerging computing technologies and evaluating how each may fit into the ecosystem being built at the center. As group leader, Gebert works

> "As a taxpayer funded research institute, we serve a huge research community as well as a huge industrial community."
> Johannes Gebert

closely with HLRS's technology partners and his colleagues at the other Gauss Centre for Supercomputing (GCS) centers to ensure that researchers in Germany and the European Union always have access to world-class computational tools. "As a taxpayer-funded research institute, we serve a huge research community as well as a huge industrial community," he said. "They both need access to these systems to be competitive internationally, and it is our job to provide a mix of technologies that are best suited to accomplishing that task."

Gebert pointed out that as HPC, AI, and quantum computing all continue to grow rapidly, groups like his cannot effectively do their jobs alone. He pointed out that as part of GCS, HLRS staff members are well-positioned to share their experiences with staffers at the Jülich Supercomputing Centre and Leibniz Supercomputing Centre, its partners in the Gauss Centre for Supercomputing. "With the current pace and pressure in the industry, we have to join forces to maintain our competitiveness," he said. "We can work together to evaluate the most promising future computing technologies and minimize the duplication of effort across our centers. Good plans get created through discussion, and we are lucky to have so many different people with different experiences, perspectives, and specializations across the three centers."

No matter what technologies rise and fall, Gebert envisions the next several years of his job closely connected to helping HLRS further integrate its traditional HPC capabilities with increasingly powerful AI systems, and that these two computing disciplines will continue to merge. He also takes great interest in the commodification of HPC and AI, as it will continue to shift the role of public HPC centers and inform his role for years to come. "Thirty years ago, not everyone had a mobile phone, but in the meantime, it's become a basic item, almost a necessity," he said. "Similar things are happening with HPC, because 40 years ago, having an HPC system really was a flagship project. But today, many centers and even private companies run relatively large systems. HPC centers may become more common in the coming decades, and we must keep reinventing ourselves to be even more service focused as the technologies continue to rapidly change and evolve." *Eric Gedenk*

"HPC centers may become more common in the coming decades, and we must keep reinventing ourselves to be even more service focused as the technologies continue to rapidly change and evolve."
Johannes Gebert

Gebert works at the intersection of HLRS's current capabilities and its focus on supporting rapidly evolving research communities in the era of AI, quantum computing, and other next-generation computing technologies.

# Education and Training

GCS and its member centers provide first-class training opportunities for the national and European high-performance computing (HPC) communities. More than 50 highly qualified staff members and scientists work as trainers for over 100 courses per year. This training volume positions GCS as one of Europe's leading training organizations for skills related to classical modeling and simulation, artificial intelligence, and quantum computing.

Together, the centers offer trainees flexibility through a mix of in-person training courses, online training events, and hybrid training courses. Most training courses employ a combination of video conferencing software, chat programs, and other collaboration tools, allowing remote participants to engage in the training work in meaningful ways. Many of the in-person training events are spread across multiple days, building in networking opportunities and tours of Germany's leading HPC centers.

As the GCS centers have diversified their respective HPC architectures, they have come together to collaborate on training events to ensure that trainees can learn across the broadest-possible span of programming languages, hardware architectures, artificial intelligence paradigms, and novel new computing technologies such as quantum computing.

**For a full list of training courses, please visit:**

| GCS | HLRS | JSC | LRZ |
|---|---|---|---|
| https://www.gauss-centre.eu/trainingsworkshops | https://www.hlrs.de/training/hpc-training | https://www.fz-juelich.de/en/jsc/education/training-courses | https://app1.edoobox.com/en/LRZ/ |


The Hunter supercomputer at HLRS.

# High-Performance Computing Center Stuttgart

The High-Performance Computing Center Stuttgart (HLRS) was established in 1996 as Germany's first national high-performance computing center. A research institution affiliated with both GCS and the University of Stuttgart, HLRS provides infrastructure and services for HPC, data analytics, visualization, and artificial intelligence to academic users and industry across many scientific disciplines, with an emphasis on computational engineering and applied science.

## Supercomputing for industry

Through a public-private joint venture called hww (Höchstleistungsrechner für Wissenschaft und Wirtschaft), HLRS ensures that industry always has access to state-of-the-art HPC technologies. HLRS also helped to found SICOS BW GmbH, which assists small and medium-sized enterprises in accessing HPC technologies and resources.

Additionally, HLRS cofounded the Supercomputing-Akademie, a training program that addresses the unique needs of industrial HPC users.

## Guiding the future of supercomputing

HLRS scientists participate in dozens of funded research projects, working closely with academic and industrial partners to address key problems facing the future of computing. Projects develop new technologies and address global challenges where supercomputing can provide practical solutions. With the support of the EuroHPC Joint Undertaking, HLRS is also currently coordinating efforts to build and integrate HPC competencies across Europe, and coordinates the EuroHPC JU AI Factory HammerHAI. The center is certified for environmental management under the EU's Eco-Management and Audit Scheme (EMAS) and for information security under the ISO 27001 standard.

High-Performance Computing Center Stuttgart

CONTACT

**High-Performance Computing Center Stuttgart (HLRS), University of Stuttgart**

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



CENTER
DESCRIPTIONS

## Compute servers currently operated by HLRS

| System | Size | Peak Performance (Tflop/s) | Purpose | User Community |
|---|---|---|---|---|
| HPE CRAY EX4000 (HUNTER) | APU: AMD Instinct MI300A, 188 nodes. 512 GB HBM3 (~5.3 TB/s) memory<br>CPU: AMD EPYC 9374F, 256 nodes. 768 GB DDR5-4800 memory<br>Cray ClustorStor E2000, 25 PB on 2,120 disks<br>HPE Slingshot 11 Dragonfly (APU: 4×200 Gbps per node) | 48.1 PFlop | Capability Computing | German and European research organizations and industry |
| HPE APOLLO 6500 GEN10 PLUS (HAWK AI EXPANSION) | 24 nodes, 192 NVIDIA A100 GPUs | 120 Pflops AI performance | Machine learning and artificial intelligence applications | German and European research organizations and industry |
| NEC CLUSTER (VULCAN, VULCAN 2) | CPU<br>Intel 6230 (8 nodes), Intel 6258 (96 nodes), Intel 6238 (72 nodes), Intel 4112 (9 nodes), AMD 9124 (24 nodes), AMD 9334 (154 nodes), AMD 7302 (3 nodes), AMD 7642 (4 nodes) | CPU performance: 802 TF | Capacity Computing | German Universities, Research Institutes, and Industry |
| | GPU<br>NVIDIA A30 (24 nodes), AMD MI50 (4 nodes, 8 GPUs per node), NVIDIA RTX4000 (3 nodes), AMD WX8200 (6 nodes) | GPU performance<br>Single precision: 763 TF<br>Single precision: 763 TF | | |

Modular Supercomputer JUPITER at the Jülich Supercomputing Centre.

# Jülich Supercomputing Centre Forschungszentrum Jülich

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

## Provision of supercomputer resources:

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

## Core tasks of JSC are:

- Supercomputer-oriented research and development in selected fields of physics and other natural sciences by research groups and in technology, e.g. by doing co-design together with leading HPC companies.
- Implementation of strategic support infrastructures including community-oriented simulation and data laboratories and cross-sectional teams, e.g. on mathematical methods and algorithms and parallel performance tools, enabling the effective usage of the supercomputer resources.
- Cutting-edge quantum computing research and access through the Jülich UNified Infrastructure for Quantum computing (JUNIQ).
- Higher education for master and doctoral students in close cooperation with neighbouring universities.

JÜLICH Forschungszentrum | JÜLICH SUPERCOMPUTING CENTRE

CONTACT

**Jülich Supercomputing Centre (JSC)
Forschungszentrum Jülich**

Prof. Dr. Dr. Thomas Lippert
Phone +49 - 24 61- 61- 64 02 / th.lippert@fz-juelich.de

Prof. Dr. Kristel Michielsen
Phone +49 - 24 61 - 61 - 25 24 / k.michielsen@fz-juelich.de

Wilhelm-Johnen-Straße, 52425 Jülich, Germany
www.fz-juelich.de/jsc





## Compute servers currently operated by JSC

| System | Size | Peak Performance (Tflop/s) | Purpose | User Community |
|---|---|---|---|---|
| JUPITER BOOSTER | 125 racks, 5,884 nodes Eviden BullSequana XH3000<br>23,536 NVIDIA GH200 Superchips<br>1,694,592 CPU cores NVIDIA Grace<br>4,964 TB memory (LPDDR5+HBM3) | 1,258,140 | Capability Computing and AI | European and German Universities, Research Institutes, and Industry |
| JSC CLOUD | 72 nodes, 11,760 cores<br>AMD EPYC & Intel<br>NVIDIA V100/A100/RTX A6000/L40s/H100, AMD MI210<br>36 TByte memory<br>145 TByte NVM | 760 | Cloud Computing (IaaS and PaaS) | European and German Research Institutes |
| MODULAR SUPERCOMPUTER "JUWELS" | Cluster (Eviden):<br>10 cells, 2,567 nodes<br>122,768 cores<br>Intel Skylake<br>224 NVIDIA V100 GPUs<br>275 TByte memory | 12,266 | Capability Computing | European and German Universities and Research Institutes |
| | Booster (Eviden):<br>39 racks, 936 nodes<br>44,928 cores<br>AMD EPYC Rome<br>3,744 NVIDIA A100 GPUs<br>629 TByte memory | 75,020 | | |
| SUPERCOMPUTER "JURECA" | Data-Centric Cluster (Eviden):<br>768 nodes, 98,304 cores<br>AMD EPYC Rome<br>768 NVIDIA A100 GPUs<br>443 TByte memory | 18,515 | Capacity and Capability-Computing | European and German Universities, Research Institutes, and Industry |
| | JURECA-HWAI:<br>32 nodes, 2,048 cores<br>Intel Sapphire Rapids 8462Y<br>128 NVIDIA H100 GPUs<br>16 TByte memory | 8,756 | AI | German Research Institutes and Industry<br>HAICORE3/WestAI |
| HPC/ CLOUD CLUSTER "JUSUF" | 205 nodes, 26,240 cores<br>AMD EPYC Rome<br>61 NVIDIA V100 GPUs<br>52 TByte memory | 1,372 | Capacity and Cloud Computing | European and German Universities and Research Institutes through Human Brain Project |
| D-WAVE QUANTUM ANNEALER "JUPSI" | More than 5,000 qubits | No classical performance measure applicable | Quantum Computing | German Universities and Research Institutes (10%) Industry Applications and D-Wave customers (90%) |


The SuperMUC-NG supercomputer at LRZ.

# Leibniz Supercomputing Centre

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

## Leadership in HPC and HPDA

Located on the research campus in Garching near Munich, the LRZ is a leadership-class HPC and HPDA facility delivering top-tier supercomputing resources and services on the national and European levels. Top-notch specialists for HPC code portability and scalability support the broad user base at LRZ and ensure that the systems are running their operations in the most energy efficient way possible.

Leibniz Supercomputing Centre
of the Bavarian Academy of Sciences and Humanities

## Quantum and Future Computing at LRZ

LRZ is leading the way forward in the field of future computing, focusing on emerging technologies like quantum computing and integrating AI on large-scale HPC systems. A robust education program that touches on HPC, machine learning, artificial intelligence, and big data complements LRZ's offerings.

## IT backbone for Bavarian science

In addition to its role as a national supercomputing center, LRZ is also the IT service provider for all Munich universities as well as research organizations throughout Bavaria.

CONTACT

**Leibniz Supercomputing Centre (LRZ)**

Prof. Dr. Dieter Kranzlmüller
Boltzmannstraße 1, 85748 Garching near Munich, Germany
Phone +49 - 89 - 358 - 31 - 80 00
leitung@lrz.de
www.lrz.de





## Compute servers currently operated by LRZ

| System | Sytem & Size | Peak Performance (Tflop/s) | Purpose | User Community |
|---|---|---|---|---|
| SuperMUC-NG Phase 1 Intel/Lenovo ThinkSystem | 6,336 direct hot-water cooled compute nodes, 304,128 cores, Intel Xeon Platinum 8174, 608 TByte of memory, Omni-Path 100G interconnect | 26,300 | Capability Computing | German universities and research institutes (Tier-1) |
| | 144 direct hot-water cooled compute nodes, 6,912 cores, Intel Xeon Platinum 8174, 111 TByte of memory, Omni-Path 100G interconnect | 600 | Capability Computing | |
| SuperMUC-NG Compute Cloud | 64 air-cooled nodes, 5,120 cores, Intel Xeon Gold 6148, 64 Nvidia Tesla V100 | 644 (CPUs + GPUs)<br>7,168 AI Performance* | Cloud Computing | German Universities and Research Institutes (Tier-1) |
| SuperMUC-NG Phase 2 | 240 direct hot-water cooled compute nodes, 26,880 Intel Xeon Platinum 8480+ compute cores (Sapphire Rapids), 159 TByte of memory, 960 accelerators (Intel Ponte Vecchio), NVIDIA HDR Infiniband interconnect | 27,960 | Capability Computing & Machine Learning, AI applications | German Universities and Research Institutes (Tier-1) |
| CoolMUC (4$^{th}$ Generation) | 119 direct hot-water cooled compute nodes, 12,928 cores Intel Xeon Platinum 8480+ (Sapphire Rapids) and others, 76 TByte memory, NVIDIA HDR Infiniband interconnect | 807 | Capability Computing | Bavarian Universities (Tier-2) |
| LRZ AI Systems | 89 nodes (NVIDIA GPU-based); HDR Infinband<br>384 NVIDIA GPUs; 50 TByte HBMemory | 16,696<br>1,196,174 AI Performance* | Machine Learning, AI applications | Bavarian Universities |
| Cerebras CS-2 | 1 node with 850,000 compute cores, 40GB SRAM, 20 PB/s memory bandwidth and 220Pb/s interconnect | 3,570,000 AI Performance* (estimate based on arXiv:2204.03775) | Purpose-built Deep Learning System | select users, not part of the public-facing AI Systems |
| LRZ Quantum Computing Resources | **EVIDEN QAPTIVA 1 + 2 (HARDWARE SIMULATOR)**<br>N/A | Suitable for computation of quantum algorithms for systems up to 38 qubits each | Quantum simulation | Bavarian Universities |
| | **DAQC SYSTEM 1 + 2**<br>5 and 20 qubits respectively based on superconducting technology, manufactured by IQM | N/A | Quantum computation | DAQC project research partners and internal LRZ staff |
| | **Q-EXA**<br>20 qubits based on superconducting technology, manufactured by IQM | N/A | Quantum computation | Quantum computation German, Bavarian and European users as preliminary system for Euro-Q-Exa |
| | **MARMOT**<br>16 qubits based on ion-trap technology, manufactured by AQT | N/A | Quantum computation | MQV research partners |
| | **EURO-Q-EXA SYSTEM 1**<br>54 qubits based on superconducting technology, manufactured by IQM | N/A | Quantum computation | German, Bavarian and European users |

*AI Performance refers to GPU peak performance for FP16 operations. For Nvidia GPUs, it is specific to different architectures. P100 architecture: CUDA core performance. V100 architecture: Mixed precision Tensor Core performance. A100: Structural sparsity Tensor Core performance. H100: structural sparsity Tensor Core performance and, possibly, Transformer Engines. B200: Structural sparsity Tensor Core performance and, possibly, second-generation Transformer Engines.

Would you like to keep up with the newest research and news happening at the GCS centers?

Please send your postal address to pr@hlrs.de and you will receive InSiDE magazine twice a year. You can also request to receive the magazine as a PDF via email.

## IMPRINT

InSiDE is published two times a year by the Gauss Centre for Supercomputing e.V., Alexanderplatz 1, 10178 Berlin

Editor-in-Chief: Michael Resch, HLRS Höchstleistungsrechenzentrum Stuttgart, Nobelstaße 19, 70569 Stuttgart, resch@hlrs.de

Executive Editor: Eric Gedenk, GCS, eric.gedenk@lrz.de

Art Direction and Design:
GROOTHUIS. Gesellschaft der Ideen und Passionen mbH für Kommunikation und Medien, Marketing und Gestaltung: Lucy Lepstück, Sophie Popp; www.groothuis.de

Print:
Buch- und Offsetdruckerei H. Heenemann GmbH & Co. KG

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

RECYCLED Paper made from recycled material FSC® C020935

Druckprodukt mit finanziellem Klimabeitrag ClimatePartner.com/53170-2605-1002

Contributing authors: Eric Gedenk, GCS, eric.gedenk@lrz.de; Christopher Williams, HLRS, williams@hlrs.de; Susanne Vieser, LRZ, susanne.vieser@lrz.de; Anna Lührs, a.luehrs@fz-juelich.de;

Photo credits:
Cover p. 1: stock.adobe.com / Love You Studio; p. 2: stock.adobe.com / Love You Studio; p. 3: Alessandro Podo, Sascha Kreklau, HLRS; p. 6/7: Benthem Crouwel Architects, in collaboration with Birk Heilmeyer und Frenzel Architekten; p. 8: HLRS; p. 9/10 LRZ / V. Hohenegger; p. 13: Forschungszentrum Jülich / Sascha Kreklau; p. 14 unsplash.de / Zoha Gohar; p. 16: Laurent Andre, RWTH Aachen; p. 18: Alina Bazarova, JSC; p. 19: UKAEA, courtesy of EUROfusion; p. 20Max Planck Institute for Plasma Physics; p. 22/23: LRZ / V. Hohenegger; p. 25: HLRS; p. 27: Lillie Elliot, SC Photography; p. 28: Klocke et al. (2025), licensed under CC BY 4.0; p. 30: Forschungszentrum Jülich / Sascha Kreklau; p. 32/33: HLRS, Jenö Gellinek; p. 34/35: Q.ANT, Forschungszentrum Jülich / Sascha Kreklau; p. 36/37: HPE, HPCwire; p. 38/40: HLRS; p. 41: University of Stuttgart / Regenscheit; p. 42: HLRS / Julian Holzwarth; p. 44: Forschungszentrum Jülich / Sascha Kreklau; p. 46: LRZ

InSiDE magazine (German: Innovatives Supercomputing in Deutschland) is the bi-annual publication of the Gauss Centre for Supercomputing, showcasing recent highlights and scientific accomplishments from users at Germany's three national supercomputing centers. GCS was founded in 2007 as a partnership between the High-Performance Computing Center Stuttgart, Jülich Supercomputing Centre, and the Leibniz Supercomputing Centre. It is jointly funded by the German Federal Ministry for Research, Technology, and Space (BMFTR) and the corresponding ministries of the three states of Baden-Württemberg, North Rhine-Westphalia, and Bavaria.

GCS
Gauss Centre for Supercomputing

www.gauss-centre.eu

Funding for GCS HPC resources is provided by:

Federal Ministry of Research, Technology and Space

Bavarian State Ministry of Science and the Arts

Baden-Württemberg Ministry of Science, Research and Arts

Ministry of Culture and Science of the State of North Rhine-Westphalia