# InSiDE

Innovatives Supercomputing in Deutschland

No 26-2 · Autumn 2026



## Thirty Years, One Vision

HLRS welcomed friends and partners to reflect on the center's thirtieth anniversary.

## First year of JUPITER

In the first year since its inauguration, Europe's fastest supercomputer is already a big boon for research.

## Earth Science Support

LRZ and DLR extend their collaboration on terrabyte project to benefit environmental research.

GCS
Gauss Centre for Supercomputing

As high-performance computing centers continue to integrate new technologies related to artificial intelligence and quantum computing, our centers are focused on building a diverse, healthy computing ecosystem that can support researchers across a wide variety of scientific domains in both the public and private sectors. This issue highlights how we not only support that ecosystem with hardware investments, but also through developing next-generation tools and support structures. Taken together, we are creating computing environments that will both support long-time users and bolster emergent fields looking for ways computing can improve competitiveness and expedite scientific progress.



# 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 highlighting our organization's increasingly diverse computing ecosystem and how we are supporting both classical modeling and simulation while also investing in new computing paradigms.

As artificial intelligence applications develop at a rapid pace, we are meeting the moment and developing world-class systems that support training larger AI models and running AI-accelerated research applications. The Jülich Supercomputing Centre's (JSC's) JUPITER system just completed the first year since its inauguration. As the system in Europe with the most accumulated GPU performance, JUPITER has already supported advanced AI training tasks during its early access program JUREAP (page 11). Researchers at the Friedrich Schiller University Jena have also been using high-performance computing (HPC) resources at the High-Performance Computing Center Stuttgart (HLRS) and Leibniz Supercomputing Centre (LRZ) to train new machine learning models to support advanced astrophysics simulations (page 19).

We are also more heavily invested in quantum computing resources at our centers. JSC recently inaugurated the new quantum system JION that uses trapped ion qubits and a new technology called MAGIC for controlling them that could pave the way for building larger quantum systems more efficiently than current superconducting systems (page 30). LRZ is hosting one of the large European quantum systems, Euro-Q-Exa, and the center just began welcoming users. In addition, the center is working with technology provider planqc to deploy a new quantum system that uses neutral atoms (both articles on page 38). We are also profiling an early-career researcher at JSC who has developed a passion for working at the cutting-edge of quantum computing technology (page 39).

Developing and deploying the technology is part of the ecosystem, but building the support, tools, and infrastructure necessary to take advantage of these new systems is also part of developing the next generation of computing. HLRS recently hosted a symposium that focused on how simulation and the development of digital twins can be used to support climate change preparation and mitigation at the local level (page 29). LRZ is an active participant in the EXA4MIND project, which is focused on developing AI-powered data analytics tools for the exascale era (page 24).

As our industry continues to evolve and expand, we will continue to foster a diverse, healthy, thriving computing ecosystem at our centers.

*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


## Editorial 3

## News Features

LRZ Extends Seven-Year Collaboration with German Aerospace Agency with terrabyte Project 6

Thirty Years, One Vision 8

First Year of JUPITER: Research Highlights on Jülich's Exascale System 11

## Science Features

Solving the Mystery of Muons' Magnetic Moment 14

20 Years of SeisSol: Calculating Earthquake Waves and Advancing Earth Science 16

Sharpening Our View Into Neutron Star Mergers 19

## Projects

Performance Optimization Supports Urgent Computing for Earth Systems Simulations 22

EXA4MIND: AI-Powered Data Analytics Tools for Supercomputers 24

New TOAR-II Theme Issue Provides Comprehensive Overview of Global Ozone 26

## Events

Digital Twins Offer Tools for Managing the Environment and Climate Change 28

Official Launch of JION Quantum Computer at JSC 30

TPC Hackathon Fosters Research Collaboration Aimed at Improving AI for Science 32

## News Briefs 34

Staff Spotlight: Far From Home, JSC Post-Doctoral Researcher Stands at the Intersection of Quantum Computing and High-Performance Computing 39

Education and Training 41

## Center Descriptions

High-Performance Computing Center Stuttgart 42

Jülich Supercomputing Centre 44

Leibniz Supercomputing Centre 46

Imprint 48

# LRZ Extends Seven-Year Collaboration with German Aerospace Agency with terrabyte Project

**Long-time partners renew initiative to provide a massive trove of Earth observation data, supporting research across multiple disciplines.**




A real success story: the terrabyte high-performance platform for earth observation and environmental data.

Since 2019, the terrabyte project has supported science and research projects ranging from providing data on energy infrastructure in the world's oceans to creating new monitoring systems for forests. During this period, researchers involved in the effort have generated a wealth of practical knowledge of how to use cutting-edge computing and modelling tools alongside Earth observation data.

The joint operators of this specialized high-performance platform for data analytics, the German Aerospace Centre (DLR) and the Leibniz Supercomputing Centre (LRZ), agreed to extend terrabyte this summer, and the project is now being gradually upgraded and expanded.

## Bringing technological might to environmental research

The organizations agreed that the first step to extending terrabyte needed to start with data management: the data storage capacity will be expanded in 2027. Currently, terrabyte data repositories offer a capacity of around 50 petabytes, which is consistently 95 percent full. In a second step, LRZ and DLR plan to increase the computing power from the current level of around five petaFLOPS. The dedicated supercomputing system for the project currently comprises 636 CPUs and 188 GPUs. In addition to its high processing capacity, terrabyte provides researchers with more than 43 petabytes of Earth observation and geodata for immediate access and also offers useful open-source software and tools for their analysis, calculation and modelling.

Since the collaboration began in 2019, terrabyte has become a real success story: Researchers from DLR and Bavarian universities have come to rely on the platform, which connects the DLR archive in Oberpfaffenhofen with LRZ's Garching-based data storage resources through a high-speed data link. This allows additional satellite data stored there to be retrieved at a rate of 10 gigabits per second and processed on LRZ's supercomputers.

Fifteen long-term projects are currently benefiting from terrabyte. The most extensive use case to date is the DLR project, "Normalise Radar Backscattering." Its aim is to convert data from radar satellites into analysis-ready data. To this end, DLR has developed methods and recently generated a dataset of around eight petabytes. For the "Real-time Earth Observation of Forest Dynamics and Biodiversity" project ROOT, funded by the Bavarian Research Institute for Digital Transformation and the Bavarian Academy of Sciences and Humanities, a team at Julius-Maximilians-Universität Würzburg developed methods for analyzing disturbances and damage in forests. Ecologists, biologists, computer scientists, and Earth observation specialists all actively develop the ROOT software and app, which now integrates real-time satellite and environmental data to monitor forests. This helps researchers and forestry management organizations control pest infestations in trees or mitigate fire risks such as accumulation of damaged trees or fallen timber.




Signed and sealed: Dr. Anke Pagels-Kerp, Exacutive Board Member for Space, and Dr. Doris Klein, Deputy Director of the German Remote Sensing Data Center (both of DLR), signed the collaboration agreement on terrabyte with Prof. Dieter Kranzlmüller, Head of LRZ.

| Terrabyte: Key Figures |
|---|
| • 7,000 jobs per day |
| • around about 600 users |
| • 45 centrally provided data sets |
| • 88 million metadata |

## Digital sovereignty for science

With its computing power, tools, and flexible storage capacity, terrabyte offers an alternative to commercial cloud infrastructure providers: on the platform, researchers can develop and use their own tools, as well as copy calculation results conveniently to their own storage media. Considering current geopolitical trends towards isolationism, terrabyte also promotes digital sovereignty for European and German science. It is no coincidence that the project's success has created another intitive for the project team: it now advises research institutions on the development of comparable platforms and systems. *Susanne Vieser*

# Thirty Years, One Vision

**Marking the 30th anniversary of its founding, the High-Performance Computing Center Stuttgart welcomed friends and partners to reflect on the center's role in the international HPC landscape – and on what the future could hold.**






In a celebration held in Stuttgart's Alte Reithalle, Baden-Württemberg Science Minister Petra Olschowski and representatives from across HLRS's community highlighted its support for scientific excellence, industrial innovation, and international collaboration.

It was the year 1996. "Dolly" the sheep was born, the "Macarena" dominated the popular music charts, and Germany won the European football championship. In the world of computing, the first computer mouse with a wheel for scrolling was introduced, the PNG image format was released, and the first Tomb Raider video game went on sale.

1996 also marked an important milestone in the history of high-performance computing (HPC) in Germany, when Baden-Württemberg Prime Minister Erwin Teufel announced the founding of the country's first national supercomputing center, the High-Performance Computing Center Stuttgart (HLRS). Since then, HLRS has grown into an internationally recognized HPC center, playing important roles in Germany, Europe, and the global HPC community.

On October 9, 2026, colleagues and friends from across HLRS's network gathered under the motto "Thirty Years, One Vision" to celebrate the center's history and contributions to high-performance computing. The gathering offered an opportunity to reflect on HLRS's history, its contributions to the European and international HPC communities, and – with the rise of new computing technologies like artificial intelligence and quantum computing – what the future could hold.

## State-of-the-art technologies for cutting-edge research

HLRS's vision for high-performance computing reaches back to 1986, a decade before its founding in its current form, when University of Stuttgart Professor Roland Rühle traveled to the United States and negotiated the purchase a Cray 2 supercomputer. While its peak performance of 1.9 gigaflops is modest by today's standards, the deal was a milestone for scientific computing in Europe, bringing the world's fastest system to Stuttgart.

Ten years later, Prof. Rühle became the founding director of the High-Performance Computing Stuttgart. At HLRS's inauguration, Minister President Teufel announced the procurement of another Cray system, which ranked seventh among the world's fastest supercomputers.

Since then, innovations in hardware have produced staggering increases in performance. (The theoretical peak performance of HLRS's current Hunter supercomputer is more than 25 million times that of its original Cray 2.) Having the world's fastest system, however, has never been HLRS's primary goal. More importantly, its focus has been to provide its system users with optimized computing capabilities for their research and technology development.

In recent years, this has meant that HLRS has supported its users in transitioning from CPU-based systems to the GPU-accelerated architecture of Hunter. HLRS has also expanded beyond its traditional focus on HPC simulation to embrace data science and artificial intelligence, most notably as coordinator of the HammerHAI AI Factory. In parallel, HLRS continually expands its training and user support programs, helping experts across Europe to make the most of the latest computing technologies.

"HLRS has always been there not just to support us as users but to challenge us as scientists and engineers to keep developing our algorithms and integrate new methods and ideas," says Prof. Andrea Beck, Managing Director of the University of Stuttgart Institute of Aerodynamics and Gas Dynamics. "One recent example is the integration of machine learning into classical simulation science, where HLRS has been at the forefront of providing hardware that helps us to fully leverage both the machine learning and classical simulation science aspects."

HLRS is also preparing for what comes next. Recently, it launched the Future Computing Group to test emerging technologies and evaluate their suitability for the center's user community. As a potential purchaser, HLRS gives valuable feedback to technology startups about their products while also gathering important insights that will help in planning its next-generation computing infrastructure.

## Best practices in HPC for industry

HLRS's history has also been shaped by its location in one of Germany's leading industrial regions. Stuttgart is home to global brands like Mercedes-Benz, Porsche, and Bosch, as well as numerous small and mid-size "hidden champions" that engineer and manufacture technically sophisticated products. For many of these companies, simulation using high-performance computing is essential, but maintaining their own supercomputers would be impractical. HLRS helps to bridge this gap.

In 1995, a consortium of public and private partners founded HWW (High-Performance Computing for Science and Industry), establishing a sustainable framework through which companies like Mercedes-Benz and Porsche access world-class HPC capabilities. HWW was integrated with HLRS when the center was founded a year later, and remains an important framework for supporting German industry.

HLRS's work with the automotive industry evolved further in 2008 with the establishment of the Automotive Solution Center for Simulation (ASCS), a pre-competitive forum for identifying and addressing the automotive industry's simulation needs. The ASCS later became a model for solution centers for other industries, including the CASE4Med medical solution center and the Media Solution Center Baden-Württemberg, both co-founded by HLRS.

To further improve access to HPC capabilities in smaller companies, HLRS and the Karlsruhe Institute of Technology established SICOS-BW in 2010. Now a nonprofit organization based at HLRS, SICOS-BW provides

> "We look back on 30 successful years with pride and gratitude. But now, our task is to do everything we can to ensure that our successors can say the same in 30 years' time." Prof. Michael Resch

consulting services for small and medium-sized enterprises, connects them to technical expertise, and facilitates access to HPC resources like those at HLRS.

When HLRS was founded, this emphasis on industry was unusual for an academic supercomputing center. As policymakers across the world have recognized the importance of HPC for industrial innovation, however, HLRS's approach has become a defining part of its activities.

Earl Joseph of global HPC industry monitor Hyperion Research says, "HLRS is, in our opinion, the world's leader in working with industry. It has developed all kinds of partnerships … with companies and organizations to help make that linkage between an academic large supercomputing center and major industry, and in particular the automotive industry. We use HLRS as an example of how to do it right."

Today, more than 75 companies use high-performance computing capabilities at HLRS that would be impractical for them to maintain independently.

## Global collaboration for global challenges

Stuttgart has a strong culture of engineering and practical problem-solving. For HLRS, this means collaborating in ways that bring together complementary expertise.

HLRS currently maintains 13 formal memoranda of understanding with other leading HPC centers around the world. It also collaborates with HPC centers and other partners across Europe on funded research projects. These typically address technical challenges related to HPC and AI, while at the same time using HPC to solve concrete problems facing society. These include global systems challenges related to climate change, human health, natural catastrophes, or the transition to sustainable energy, for example.

At the European level, HLRS has played a central role in promoting collaboration and coordination across a diverse HPC landscape. As coordinator of the EuroCC and CASTIEL projects, it has facilitated the establishment of European National Competence Centers (NCCs) for HPC and cross-center networking to ensure adoption of best practices.

"Having access to excellent HPC infrastructure, expertise, and support is important for science and industry regardless of which EU member state you are located in," said HLRS Managing Director Bastian Koller, who has spearheaded the EuroCC and CASTIEL projects. "Particularly in light of Europe's need to become more digitally independent, HLRS knows that we are all stronger when we work together." This emphasis on promoting strength through collaboration extends to its activities at the German federal level as a member of the Gauss Centre for Supercomputing.

HLRS also works with local and regional municipalities to explore how HPC could support public administration. One approach uses digital twins, which HLRS has pioneered for applications in urban and regional planning and management. By integrating and visualizing complex datasets in virtual reality, HLRS scientists have demonstrated how digital twins can enable more informed communication among city planners, city officials, and the public. Recent applications of this approach have focused on energy transformation and the conservation of heritage agricultural landscapes. (see page 28)

In these and other ways, HLRS continues to be driven by its early mission: demonstrating how high-performance computing can enable sceintific discovery, support commercial innovation, and help decision makers to make better-informed choices.

By highlighting the center's numerous contributions, HLRS's 30th anniversary celebration was a sign of how valued the center is – for science, for industry, and for society at large. And according to HLRS Director Prof. Michael Resch, the mission continues: "We look back on 30 successful years with pride and gratitude. But now, our task is to do everything we can to ensure that our successors can say the same in 30 years' time."
*Christopher Williams*



NEWS
FEATURES

# First Year of JUPITER: Research Highlights on Jülich's Exascale System

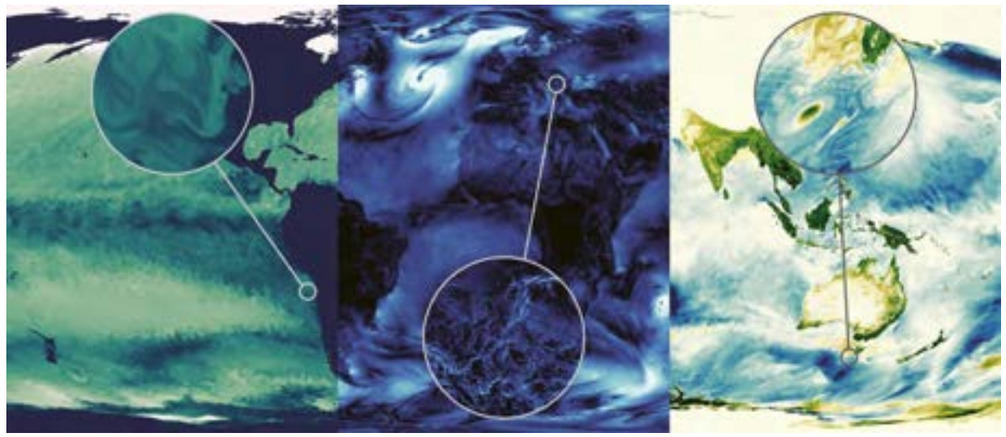
**Advancements in particle physics, climate modeling, AI training, and quantum simulations highlight a productive first year for Jülich's flagship supercomputer.**

On Sept. 5, 2025, high-ranking guests from politics, science, and industry gathered in Jülich, Germany to celebrate a major milestone: the Jülich Supercomputing Centre (JSC) at Forschungszentrum Jülich officially inaugurated Europe's first exascale supercomputer. The Joint Undertaking Pioneer for Innovative and Transformative Exascale Research, or JUPITER, is Europe's most powerful supercomputer. One year after its inauguration, JUPITER is supporting advanced European research across a wide range of scientific domains.

The system offers researchers more accumulated GPU performance than any other European supercomputer, with each of JUPITER's 5,884 nodes containing four NVIDIA GH200 superchips. These chips not only provide a major performance boost for traditional modeling and simulation tasks, such as for computational fluid dynamics or particle physics research, but also offer AI researchers up to 80 ExaFLOP/s in AI-specific 8-bit performance for training large AI models. As of September 2026, JUPITER has so far run jobs for around 400 projects across more than 25 scientific domains. Researchers affiliated with German and European research institutions have multiple ways

Racks of the JUPITER booster





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

to apply for time on the system, including the Gauss Centre for Supercomputing's Large-Scale Calls and the calls through the John von Neumann Institute for Computing for users located in Germany, and calls through the EuroHPC Joint Undertaking for European users.

In the JUPITER Research and Early Access Program, JUREAP, researchers from a wide range of scientific domains have had time to work with JSC staff to port and optimize their codes to the JUPITER architecture, attend GCS training events at JSC (a link to our events calendar can be found on page 41), and generally gain more experience with the system. To date, JUREAP has facilitated the following research highlights, among others:

## JUPITER enabled record-setting quantum simulation

In late 2025, JSC researchers hit a new milestone in quantum computing development– for the first time, the team was able to leverage JUPITER's capabilities, especially the transparent usage of GPU and CPU memory of the GH200 superchip, to fully simulate a 50-qubit universal quantum computer. The work was conducted with a purpose-built version of the Jülich Universal Quantum Computer Simulator (JUQCS) and broke the same group's prior record that ran on Japan's K supercomputer.

"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 team used 4,096 JUPITER nodes and more than a petabyte of memory in its work. Having reached this milestone, the team feels confident that the simulations will help quantum computing researchers get a clearer look at hardware that, due to the complex nature of quantum interaction as well as the sensitivity of the hardware, is difficult for researchers to observe in real-time.

## Climate researchers used JUPITER for Gordon-Bell-Prize-winning modeling advancements

A team that included researchers from JSC used JUPITER to greatly improve the resolution of a simulation with the flagship ICON climate model. The model is used internationally by computational climate scientists and incorporates accurate atmosphere, ocean, land surface, and sea ice modeling. The team optimized the model to run on JUPITER at a resolution of 1.25 kilometers covering the entire globe, when comparable ICON simulations had previously run at 5–10-kilometer resolution.

Being part of the JUREAP program, the team had early access to the hardware to begin optimizing the code before the system was fully online. Building on its early access success, the team scaled the application to run on 5,120 nodes on JUPITER, nearly using the full machine. The work was awarded the Gordon Bell Prize – one of supercomputing's most prestigious prizes – in the climate modeling category during the SC25 conference in St. Louis, USA last fall.

## The Jülich Brain Atlas project trained a neuroscience foundation model in five days

While AI models promise major research advancements across a range of domains, the research will only be as efficient as the models available to scientists. Researchers at Forschungszentrum Jülich's Institute of Neuroscience and Medicine (INM) worked closely with JSC to train a next-generation foundation model that can serve as the basis for many computational neuroscience research tasks.

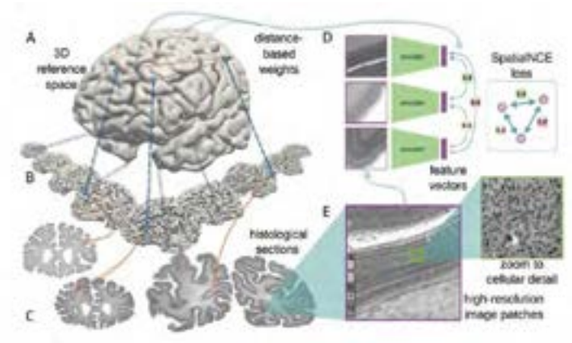
The resulting foundation model, CytoNet, was trained on more than six petabytes of data coming from 21 real post-mortem brains. Using cellular-level imaging for each brain, researchers trained the model on JUPITER in under five days. Prof. Dr. Katrin Amunts, director of the institute division INM-1 "Structural and Functional Organisation of the Brain", indicated that this was a groundbreaking moment in neuroscience research, taking a leap from using AI-accelerated programs to do analyses to developing a model capable of thinking through experiments itself.

## Researchers solved the mystery of muons' magnetic moment

A long-running research collaboration has used JSC computing resources to better study subatomic particle interactions at unprecedented scale and detail. For years, physicists noted a small but important discrepancy between the magnetic moment of muons, an elementary particle, when performing experiments compared to the theory prediction. Such a small difference has large implications – it led physicists to hypothesize that they might have uncovered signs that there were new physics beyond the Standard Model that has long governed our understanding of the physical world.

The team first threw doubt on the idea of new physics with simulations on JSC's JURECA and JUWELS systems in 2021. They then refined the code on these systems to eventually run on JUPITER and take full advantage of its GPU-heavy architecture. The simulations on JUPITER brought theory and experiment into even closer agreement, reducing the likelihood that there were new physics in muons' behaviors. "We needed a system like JUPITER to improve some of the largest systemic errors in our simulated system," said Prof. Kálmán Szabó who coordinated the Jülich contribution, "It was very helpful to be able to test our codes on other systems before we ran it on JUPITER. This way, we could get them working close to optimally with little effort." (For the full story, please visit page 14).

> "We needed a system like JUPITER to improve some of the largest systemic errors in our simulated system." Prof. Kálmán Szabó

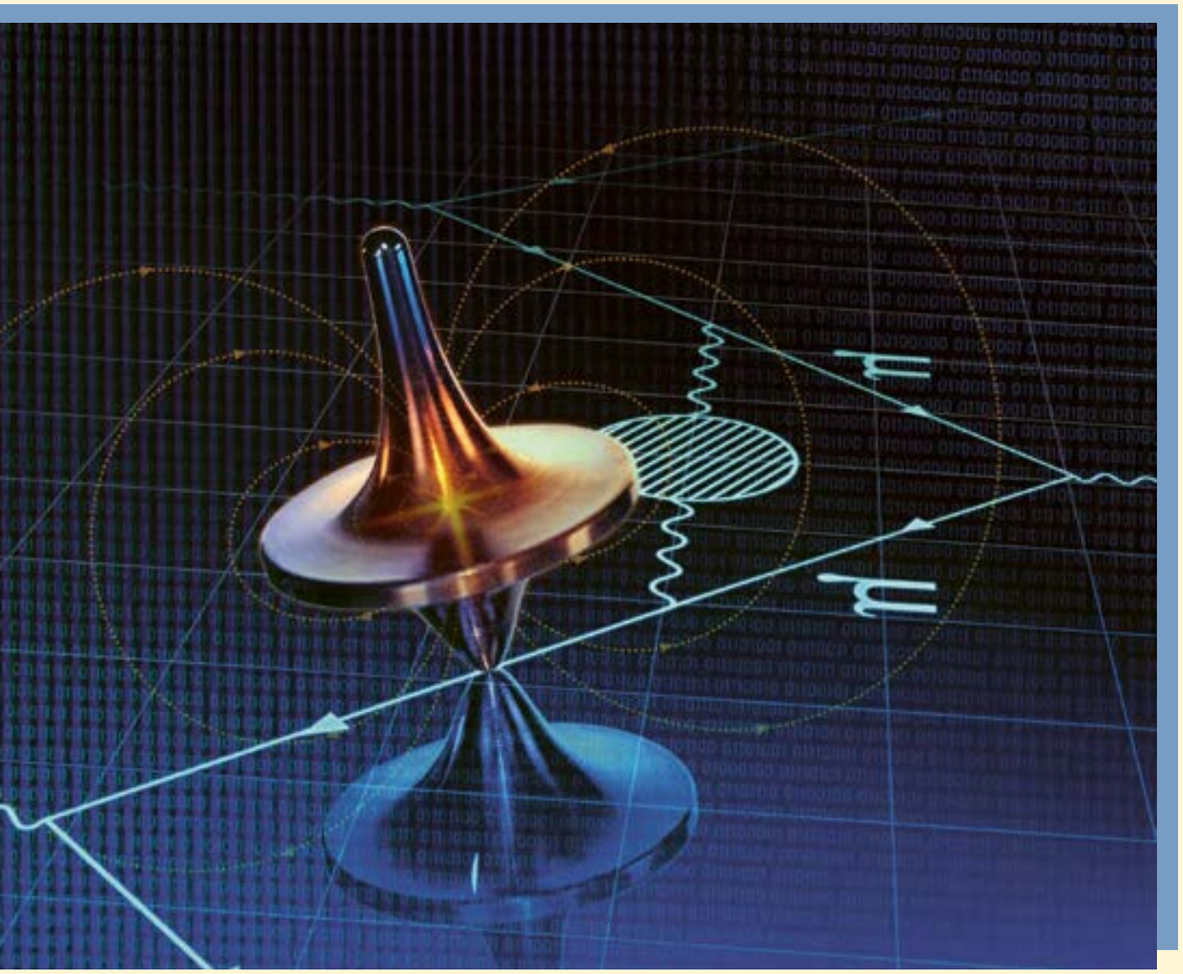



Self-supervised pretraining workflow used to train CytoNet

The JUREAP early-access program is now entering its final phase, with research teams preparing and scaling a further set of scientific use cases. Some research projects have already begun or are in the starting blocks to use the power of the JUPITER Booster module. In the fourth quarter of 2026, the JUPITER Booster will be joined by the JUPITER Cluster module, which will enable the execution of simulations that require pure CPU-focused workloads in a more power-efficient way than the Booster module. Additionally, projects like the JUPITER AI Factory (JAIF) are gaining momentum, empowering start-ups, small and medium-sized enterprises, industrial partners, researchers, and public institutions to harness the full potential of artificial intelligence on JUPITER.
*Eric Gedenk*

# Solving the Mystery of Muons' Magnetic Moment

**Researchers at Forschungszentrum Jülich and the University of Wuppertal collaborated with an international team, using the power of several supercomputers, including the exascale machine JUPITER, to help resolve a long-standing discrepancy between physics experiments and theory.**



A muon (mu) – represented here by a spinning top – moves along its path (solid line), while it interacts with photons (wiggly lines) and other particles (shaded region).





The Jülich Supercomputing Centre (JSC) offers world-class computational power for researchers across a range of research domains. Moreover, the center's directors and long-time power users have played a leading role in combining theory and experiment to advance our understanding of elementary particle physics.

In recent years, JSC researchers in the Nuclear and Particle Physics group, led by Prof. Dr. Kálmán Szabó, participated in an international collaboration that focused on studying muons, an elementary particle. For many years, scientists noticed a tiny but nonetheless significant discrepancy between experimental measurements of the muons' magnetic behaviour and theoretical predictions. Such a small difference has large implications – it led physicists to hypothesize that they might have uncovered signs that there were new physics beyond the Standard Model of particle physics that has long governed our understanding of the physical world.

In 2021, the team used JSC's computing resources to bring this discrepancy much closer together, calling into question the prior research. Recently, the same team continued to refine these calculations using several supercomputing systems at Gauss Centre for Supercomputing (GCS) centers, including JURECA, JUWELS, and JUPITER at JSC; SuperMUC-NG at the Leibniz Supercomputing Centre (LRZ); and Hawk and Hunter at the High-Performance Computing Centre Stuttgart (HLRS). Using these systems, the team further reduced uncertainty in its simulations – by a factor of 1.6 – ultimately bringing theory and experimental work back into nearly perfect agreement. "In 2021, we were the only group with this result, meaning we had to convince ourselves and others in the community we did not make a mistake in our research," Szabó said. "We performed several cross checks and used additional computational resources to improve our theoretical calculations. At the same time, our experimental colleagues also improved their experiment by four-fold, and taken together with other groups' results, we were able to confirm our findings." The team recently published its newest findings in *Nature*.

## Improved theory and experiment reconfirm physics' Standard Model

The team improved the accuracy of their check on the Standard Model of particle physics by refining its methods to take full advantage of new, more powerful computational resources. Szabó indicated that these new, accurate calculations were approaching the level of accuracy now found in cutting-edge experimental work done by an international research collaboration. "A typical bathroom scale has an uncertainty of around one kilogram," he said. "Imagine that your bathroom scale was so precise that it could tell your weight within ten micrograms – about the weight of one eyelash. That is the precision that our experimental colleagues have achieved, and while theory is not in perfect alignment with experiment, we are getting close."

To close that gap between experiment and theory, the team focused its most recent simulations on more accurately representing the strong force – the most powerful force in nature, which binds together subatomic particles to form protons and neutrons. To that end, the researchers started improving their state-of-the-art lattice quantum chromodynamics (LQCD) code on JURECA and JUWELS. By refining their computational method on these systems, Szabó and his collaborators were able to easily port their application to run at scale on JUPITER, taking advantage of the system's GPU-heavy architecture. "We needed a system like JUPITER to help improve some of the largest systemic errors in our simulated system," he said. "It was nice to test our codes on the other systems before running them on JUPITER. This way, we could get them working close to optimally with little effort."

The team's simulations helped bring muons' magnetic behavior in closer agreement with the Standard Model, reducing the likelihood that this discrepancy was a sign of new fundamental physics.

Szabó indicated that the team generated a large amount of data in this current round of simulations, and that it would spend time parsing out the data for other observations, such as more accurate mass and decay rates for other short-lived particles that appear in the simulations. He added that work on elementary particle physics takes the deliberate, iterative approach to prioritize accuracy. This is important as physicists seek to not only better understand the fundamental laws governing our world, but also continue to search for behaviors that call those laws into question.

"We will keep doing what we have always done – making large-scale computations and comparing them to increasingly precise measurements," he said. "Maybe someday we will find something that actually does not fit into the Standard Model."

The researchers will continue to analyze the data generated by their simulations and apply their methods to other questions in particle physics. By comparing increasingly precise theoretical predictions with experimental measurements, they hope to deepen our understanding of the fundamental laws of nature – and remain prepared for any future results that challenge current theories. *Eric Gedenk*

# 20 Years of SeisSol: Calculating Earthquake Waves and Advancing Earth Science

**Over two decades, LRZ researchers have supported a multi-institutional effort to better understand the risks connected to earthquakes and subsequent tsunamis. With researchers gaining access to increasingly powerful supercomputers, the SeisSol code shows no signs of slowing down.**

Of all natural hazards, earthquakes and tsunamis still claim the most lives. Accordingly, researchers continue to work intensively to uncover their causes and better understand how certain types of earthquakes can trigger more devastating tsunamis, structural damage, or both. During the last two decades, researchers in this domain have continued to develop the SeisSol software to incorporate new numerical methods to gain deeper insights.

This work has attracted attention and awards – successes that Michael Bader, computer science professor at the Technical University of Munich (TUM), attributes to years of interdisciplinary and trust-based collaboration.

Created at the Ludwig-Maximilians-Universität (LMU), SeisSol has been developed in cooperation with TUM as well as the Leibniz Supercomputing Centre (LRZ). "It is by no means a given that a research group so readily brings a code it has developed to computer science or to a computing center for optimization," Bader said. "Such code often contains the crown jewels of a scientific community – fundamental insights embodied in modelling and algorithms. A great deal of trust is required for such collaboration."



Devastation that threatens lives: Earthquakes like the one that struck eastern Turkey in 2023 claim the lives of thousands of people and destroy landscapes.

## Precise and smooth modelling

Earthquakes generally occur where tectonic plates are moving beneath the Earth's surface. If these plates are unable to slide past one another, stress builds up along their boundaries. Layers of rock eventually give way, abruptly releasing the stored energy. Geophysicists refer to these events as transform earthquakes. Subduction faults are even more destructive: following a collision, one plate is forced beneath another. These earthquakes often create more intense tsunamis and destructive shaking of buildings and infrastructure.

Anyone investigating the causes of these events examines stress and fracture energy, as well as rock deformation, temperature, subsurface conditions and many other factors. There are equations for all of these phenomena: "SeisSol uses methods that have been applied in computational fluid dynamics, but not yet in geophysical applications," explains geophysicist Alice-Agnes Gabriel, now a professor at the Scripps Institution of Oceanography at the University of San Diego.

Gabriel joined the SeisSol team at LMU in 2012, developed a range of simulations, and refined fundamental equations used in the software. The code itself, however, originally goes back to LMU mathematician and geophysicist Martin Käser, who, together with LMU numerical analyst Michael Dumbser, computationally analyzed the propagation of seismic waves. Alongside differential equations, they employed the discontinuous Galerkin method (DG) and combined it with a new time-stepping approach, the Arbitrary High-Order Derivatives (ADER) method.

The ADER-DG method is based on equations for representing seismic waves. These equations are broken down into many small spatial and temporal segments to allow for more precise calculations of wave behavior, enabling each computational step to be efficiently distributed across thousands of nodes in high-performance computing (HPC) systems. Non-specialists can imagine the difference between DG and ADER-DG results as similar to that between a photograph and a film: while DG calculations produce individual snapshots, ADER-DG provides smooth transitions that put the simulation in motion. Waves are represented more realistically, making it possible to see what accelerates ground motion. "In SeisSol, the aim is to simulate the fundamental physics of earthquakes and tsunamis on the basis of sensor data and measurements, and to develop an ever deeper understanding of the processes involved," Bader said.

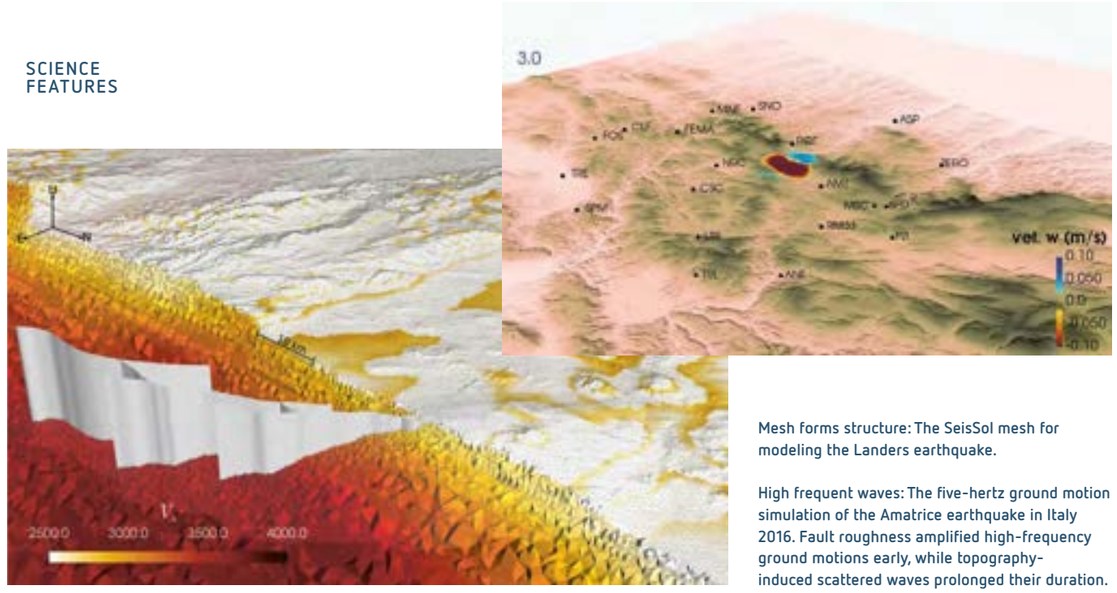


Trustful team: At Prof. Heiner Igel's (center) department at LMU, Michael Dumbser (2nd from left) and Martin Käser (right) create the innovative formulas for SeisSol. Alice Gabriel (2nd from right) and Michael Bader (left) further refined the code. All those involved have since been awarded their habilitation.

## From closed to open code

In its early years, SeisSol made a name for itself as a numerical method. With the addition of ever more parameters, it became increasingly complex, even representing wave propagation in three dimensions: "In 2012, SeisSol was the most accurate, but also the most computationally expensive, computer-based method for solving the wave equation," Gabriel said. The code delivered precision but required substantial computing time. When experts from LMU, TUM, and LRZ came together, they revised the code with efficiency in mind. At the time, the geophysicists at LMU were considering how to model fault rupture processes and integrate them into SeisSol. The computer science teams at TUM and LRZ also planned to release the code as an open-source program so that more researchers could use SeisSol. To achieve this, the geophysics team expanded the set of equations and algorithms to include rupture processes. Meanwhile, computer scientists at TUM and specialists at LRZ improved algorithms and replaced performance-critical computations by new programming approaches. SeisSol's input and output routines were modernized, and the Fortran code had finally evolved into a C++ package with elements of Python, setting it up for use in next-generation computing architectures.

The tremors of the earthquake around Landers, California in 1992 were the first time that SeisSol ran a petascale simulation: SeisSol proved highly scalable and efficient, achieving several quadrillion floating-point operations per second on LRZ's SuperMUC. With this work, the team reached the final of the Gordon Bell Prize in 2014. Three years later, it won the Best Paper Award at the SC conference series (SC17) for the longest and most comprehensive simulation to date of the Sumatra–Andaman earthquake, the most powerful earthquake in the modern times, which claimed more than 200,000 lives.




Mesh forms structure: The SeisSol mesh for modeling the Landers earthquake.

High frequent waves: The five-hertz ground motion simulation of the Amatrice earthquake in Italy 2016. Fault roughness amplified high-frequency ground motions early, while topography-induced scattered waves prolonged their duration.

"What really matters is its usability in science."
Michael Bader

### Out of geophysics, into multi-physics

During this period of growth for the software, the SeisSol team created the website SeisSol.org where it could link source code, tutorials, and publications. SeisSol became firmly established within the international geophysics community with the EU-funded Centre of Excellence for Solid Earth (ChEESE).

Between 2019 and 2023, this project produced a software package of open-source tools for modelling earthquakes and tsunamis. Today, not only academic researchers use these tools: local, state, and federal authorities, as well as private companies also use these tools. "SeisSol is currently evolving into a community code," says Gabriel. "More and more papers are being published where SeisSol provides the foundation. SeisSol has come of age."

Nevertheless, Michael Bader still regards SeisSol as a "young" code. "There is a difference between developing a new numerical method and having individual HPC achievements such as performance and scalability recognized," he said. "What really matters is its usability in science." Recent years have been shaped by technological developments in the HPC industry: GPUs have become widespread in HPC systems, enabling the team to leverage them to extend using statistical methods from artificial intelligence (AI).

Today, SeisSol is no longer used solely in geophysics; engineers and acousticians also employ it to compute wave motion. Geophysics has shifted its focus towards multi-physics applications: alongside seismic waves, combined tsunami and acoustic wave propagation, as well as complex rupture processes, have come to the fore. Recently, researchers at the University of Texas used SeisSol to model an earthquake scenario for the Cascadia subduction zone in the north-west of the United States and Canada, ultimately developing a digital twin for tsunami early warning. The team ultimately won the Gordon Bell Prize in 2025 for this work. "Following devastating earthquakes and tsunamis, we can now set up SeisSol simulations very quickly, so we have become part of the rapid scientific response to disasters," Gabriel explains. "SeisSol now brings supercomputing and very complex calculations into pre-processing and post-processing, which we have in some sense revolutionized."

With support from the United States Geological Survey, the software is now also being used to further improve earthquake early-warning systems. SeisSol additionally helps to produce so-called "shake maps" for particularly at-risk regions by combining real earthquake data with simulation results. Surrogate models, which combine physical modelling with statistical approaches, also help to explain earthquakes more effectively. And the next advance is already on the horizon: "Quantum computing is an exciting development for many geophysical applications," says Gabriel. SeisSol will play a role in this as well. "We are not yet able to model details such as the subsurface structure of earthquakes with sufficient accuracy. To achieve this, we need to bring the partial differential equations into quantum computing." *Susanne Vieser*





# Sharpening Our View Into Neutron Star Mergers

**A team at Friedrich Schiller University Jena used supercomputing resources at the High-Performance Computing Center Stuttgart and Leibniz Supercomputing Centre to advance the state-of-the-art for simulations of binary neutron star mergers.**

Since the first detection of gravitational waves from a binary neutron star (BNS) merger in 2017, astrophysics has entered a new renaissance. As their name implies, binary neutron star mergers occur when two extremely dense stars – think the mass of our sun packed into a sphere the size of a large city – approach one another and then collide, creating a larger, denser neutron star or black hole in the process. This process releases an outflow of matter that is rich in neutrons. Scientists believe that as these outflows spread across the interstellar medium, they could be one of the main sites for the creation of rare metals in our universe.

"In my own career, I started off interested in gravitational wave physics and numerical aspects surrounding the Einstein equations, but this observation of a BNS merger provided data that opened up new worlds in my research," said Prof. Sebastiano Bernuzzi, professor at Friedrich Schiller University's Theoretical Physics Institute.

Bernuzzi and his team, together with international collaborators specializing in computational relativity, use high-performance computing (HPC) resources at the High-Performance Computing Center Stuttgart (HLRS) and Leibniz Supercomputing Centre (LRZ) to interpret the observations, running highly accurate simulations of BNS mergers. Using Hawk at HLRS and SuperMUC-NG at LRZ, as well as resources at other HPC centers, the team recently completed the development of GR-Athena++, one of the first numerical-relativity codes with pre-exascale capabilities. Bernuzzi also worked with postdoctoral researcher William Cook to develop a neural network to improve modeling of magnetic flows and turbulence in BNS simulations.

"The original code was already one of the leading codes for computational astrophysics, having advanced methods for solving magnetohydrodynamics equations" Prof. Sebastiano Bernuzzi

### GR-Athena++ improves understanding of neutrino transport

For years, computational physicists have used the Athena++ code, developed at Princeton University, to model astrophysical magnetohydrodynamics (MHD). Bernuzzi and his collaborators expanded on this public code by adding Einstein's equations of general relativity. GR-Athena ++ enables researchers to study complex astrophysical phenomena under conditions of extreme gravity, such as black holes, binary neutron star mergers, gravitational collapse, and other phenomena characterized by dynamic spacetimes.

Among other advancements the researchers integrated into GR-Athena++, the work included a radiation transport module to model neutrinos – small, neutrally charged subatomic particles that are produced in a BNS merger in large numbers. Neutrino-matter interactions – or weak interactions – change the composition of matter by transforming neutrons into protons and electrons, which,

SCIENCE
FEATURES




With access to HPC resources, Bernuzzi and his collaborators were able to run world-class simulations of binary neutron star mergers that included more detailed physics, like Kelvin-Hemholtz instability and the formation of actinides.

crucially, determine the conditions for nucleosynthesis of heavy elements, like gold.

"The original code was already one of the leading codes for computational astrophysics, having advanced methods for solving magnetohydrodynamics equations," Bernuzzi said. "We included a spacetime solver, which is a significant addition, as well as neutrino transport and state-of-art technology for extracting gravitational waves and dealing with dynamic black holes and neutron star spacetimes."

With these improvements, astrophysicists can now use GR-Athena++ to make precision calculations to study gravitational waves from black hole mergers and magnetic field instabilities in isolated neutron stars. The code also incorporates enhanced microphysics that permit new BNS simulations for studying elements' formation. The team published its results in *Physical Review D*. They have also contributed to another code, called AthenaK. Primarily developed by US colleagues, the code uses modern heterogenous computing architectures with the programming ecosystem Kokkos for exascale computing.

## Machine learning helps scientists better model BNS turbulence

More recently, Bernuzzi and his collaborators have used GCS resources to address one of the other major complications when modeling complex astrophysical phenomena on supercomputers – turbulence.

Neutron stars host some of the most extreme magnetic fields in the Universe, up to 100 million times stronger than the most extreme magnetic fields generated on Earth. The strength of the magnetic fields in plasma can be further amplified in a binary collision, as the pair of stars pull toward one another. As a result, the plasma surrounding these celestial bodies starts to move in a turbulent way, making the accurate simulation of how it behaves difficult, even using a supercomputer.

In fluid dynamics research – whether simulating astrophysical plasmas or fuel being injected in a combustion engine – computational scientists and engineers often use something called large-eddy simulations (LES) to balance accuracy and efficiency. Essentially, LES allows researchers to solve equations that accurately represent large-scale turbulent motions, or eddies. LES also include so-called "sub-grid" models that describe the impact of small-scale dynamics not resolved in their simulations. This approach allows scientists to simulate more realistic scenarios but requires them to carefully calibrate sub-grid models on reference data from extraordinary high-resolution direct numerical simulations or local simulations.

Bernuzzi and postdoctoral researcher William Cook want to apply this approach to BNS merger simulations, hoping that LES could improve the accuracy of plasma dynamics. While other teams have previously used LES for this purpose, Cook and Bernuzzi took a novel approach by using machine learning methods to train neural networks.

**"Neutron stars host some of the most extreme magnetic fields in the Universe, up to 100 million times stronger than the most extreme magnetic fields generated on Earth."**

As their name implies, neural networks are machine learning models that work in a similar way to a human brain as it makes decisions. Researchers feed in large amounts of training data from other turbulence experiments and simulations to the neural network and give feedback to the model concerning how well it represents known behavior. Over time, neural networks start to find patterns in large, complex amounts of data, including connections that are not accessible in global 3D simulations.

The researchers trained the model and ran preliminary tests to ensure that their neural network faithfully reproduces the Kelvin-Helmholtz instability, one of the main mechanisms for magnetic field amplification in BNS. The researchers used the neural network to run a relatively low-resolution simulation and demonstrated that the model performed well in driving turbulence over time. The team published its results in a second article appearing in *Physical Review D* and is applying for more computing time through the Gauss Centre for Supercomputing and other international high-performance computing (HPC) facilities. Its goal is to further refine sub-grid models that can comprehensively and accurately represent the complex physics happening at the smallest turbulent scales.

Bernuzzi indicated that this work gives the team a powerful new tool that will enable it to continue improving its models of plasma dynamics. At the same time, it speaks to how the availability of more powerful HPC resources has advanced what is possible in the field. "I'm very impressed with the developments in our numerical relativity community in the last 10 years," he said. "Many of these methods have become very mature, and in the medium to long-term future, I can see us applying these methods to other astrophysical systems that involve strong gravity."
*Eric Gedenk*

Related Publications:

Cook, W. and S. Bernuzzi (2026). *Subgrid modeling for relativistic magnetohydrodynamics with machine learning*, Physical Review D. DOI: https://doi.org/10.1103/t81t-vx23

Daszuta, B. et al. (2026). *Binary neutron star merger simulations with neutrino transport using gr-athena++*, Physical Review D. DOI: https://doi.org/10.1103/9zlj-fngj

# Performance Optimization Supports Urgent Computing for Earth Systems Simulations

**One of 11 ChEESE flagship codes, OpenPDAC makes it possible to simulate phreatic eruptions of volcanoes.**

In many regions of the world, volcanoes, earthquakes, and tsunamis pose grave threats for infrastructure, property, and human life. In recent decades, high-performance computing (HPC) has become an essential tool for earth scientists, making it possible to simulate complex geophysical systems with increasing accuracy. Using mathematical models of dynamics in the Earth's crust and mantle, scientists are developing more effective methods for predicting such natural catastrophes, while also exploring how simulation could be deployed quickly to provide insights when natural disasters arise. Urgent computing, for example, could quickly help to guide airflight diversion, evacuation planning, or resource allocation for the management of disruptive or destructive geological events.

By the late 2010s, scientists across Europe had developed and validated new software for simulating systems inside the Earth. For these models to achieve the scale and resolution necessary to answer important scientific questions, however, they realized that they needed more computing power. In parallel, high-performance computers like those operated by the Gauss Centre for Supercomputing were gaining unprecedented performance capabilities. Geoscientists recognized the opportunities such supercomputers could offer, but they faced a problem: the codes they used did not run well on large, highly parallelized systems. To achieve the precision at a speed that would make it feasible to respond to natural disasters in real time, a focus on performance optimization became essential.

To address this challenge, an international community of scientists came together in 2018 to establish ChEESE, the Centre of Excellence for Exascale in Solid Earth. The effort was renewed in a second project phase called ChEESE-2P, which was funded by the EuroHPC Joint Undertaking and concludes at the end of 2026. Coordinated by the Barcelona Supercomputing Centre (ChEESE) and the Consejo Superior de Investigaciones Científicas (ChEESE-2P), the projects have brought together Earth scientists and experts in high-performance computing to improve scientific codes for solid Earth simulations in ways that would enable them to take maximal advantage of the opportunities that European HPC systems could offer.

To date, the ChEESE projects have focused on the scaling and optimization of 11 flagship scientific codes for solid earth simulations. This has included identifying performance bottlenecks, monitoring code performance to track progress in scaling and efficiency, and using this information to achieve higher performance. ChEESE and ChEESE-2P have supported the publication of 32 scientific papers and – even more importantly – have helped in the development of computing workflows that could be deployed to more effectively predict and manage natural disasters.

## ChEESE goes POP

At the same time that HLRS has been a partner in ChEESE, it has also been involved in a European consortium called POP, the EuroHPC Performance Optimization and Productivity Center of Excellence. Working together with other leading high-performance computing centers, HPC experts at HLRS have helped to develop analysis tools for evaluating software performance. They also regularly use them to improve code performance on HPC systems.

POP is a free service for academic and industrial HPC users. Once inserted into scientific codes, its performance analysis tools monitor key performance indicators including parallel efficiency, load balance, communication efficiency, scaling efficiency, serialisation efficiency, and data transfer efficiency. HPC experts within POP use the results to identify bottlenecks and make recommendations that would improve the code's efficiency and scalability to larger computing systems.

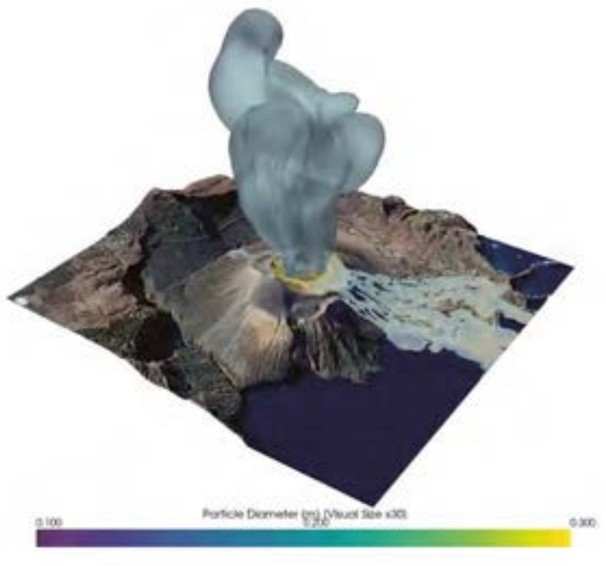
As the the EuroHPC Joint Undertaking has rolled out multiple large-scale high-performance computing systems across Europe, it now recommends that all users of these systems submit their codes for POP analysis. This step not only ensures that users gain more scientific benefit from their computing grants, but also that EuroHPC systems are as productive as possible.

## Collaboration with civil protection agencies

Within the ChEESE projects, performance was not merely a technical matter. One unique feature of the projects has been that they have included an advisory board made up of public officials from communities under threat from potential natural disasters. Representatives of civil protection agencies in Norway, Iceland, and Italy, for example, made specific recommendations concerning the kinds of prediction tools that would be most useful to them.

"Civil Protection Italy, for example, might come to ChEESE to describe what kind of maps they need and the resolution they require," says Dr. Jose Gracia, head of the HLRS Department of Scalable Programming Models and Tools. "The cooperation baked into ChEESE means that there is real interest from real users of our results."

In its most recent project phase, ChEESE-2P extended its focus beyond performance optimization of individual codes to include the workflows into which HPC simulations must be integrated in crisis management situations. To respond to a potential tsunami, for example, seismic sensors might trigger an alert and take initial measurements. The data they collect must be automatically fed into a model, which then quickly analyzes the simulation to predict the extent of danger. In severe cases, reliable, easily interpretable results must then immediately be transmitted to public agencies responsible for alerting the public and protecting their safety. Executing such a workflow efficiently means designing a computing workflow that rapidly moves data through multiple stages of analysis and processing.




Performance optimization within this Centre of Excellence is supporting the software's development for potential use in urgent computing situations. Taking advantage of higher performance in HPC, OpenPDAC could provide early warning of potential eruptions and identify areas of high danger.

## Improved workflows could enable volcano prediction

One ChEESE simulation case was led by Mattia de' Michieli Vitturi and Esposti Ongaro at the National Institute of Geophysics and Volcanology (INGV) of Italy. Focusing on the Campi Flegrei Volcano in Naples, Italy, their team focuses on methods for simulating multiphase flows of phreatic volcanic eruptions, which occur when hot magma heats ground water to high pressure, leading to an explosion. The goal is to develop an early warning tool that could be used when a danger of a phreatic eruption is identified to predict likely eruption scenarios.

Vitturi and Ongaro use the ChEESE-2P flagship code OpenPDAC for their simulations, and have been focused on optimizing it for the EuroHPC system LUMI for urgent computing applications. Their approach involves running an ensemble of parallel simulations to search the parameter space and find the likeliest scenario for an eruption. The workflow also involves pre-processing steps such as parameter setting and generation, and the processing of mesh grids. During post-processing of the ensemble simulation, the simulation results are visualized into figures and movies. Job scheduling and delegation of computing resources are also integral parts of the workflow. Although pre- and post-processing steps account for a very low percentage of the overall compute time, they are pivotal to the correctness and performance of the simulations, and the interpretation of simulation results.

The performance optimization team at HLRS worked with colleagues from INGV to improve the Python and Bash scripts used in pre- and post-processing. After using tools such as Pylint and Shellcheck to do an initial correctness check, the researchers conducted test runs to identify the most time consuming section. Employing the Numba module to convert functions that process mesh grids to machine codes, they achieved a speed up of approximately 27% in the pre-processing stage. In the post-processing, HLRS also suggested moving video generation from CPU to GPU and sped up this process up to a factor of thousands.

These kinds of performance improvements will help to make it more feasible to use OpenPDAC within urgent computing settings, and demonstrate the important role that HPC expertise can play in translating scientific innovation into applications that protect property and life.
*Christopher Williams*

# EXA4MIND: AI-Powered Data Analytics Tools for Supercomputers

**The European research project EXA4MIND provides tools and preconfigured workflows that make it easier to collect, manage, and analyze extremely large and diverse datasets.**



The European research project, "Extreme Analytics for Mining Data Services," (EXA4MIND) has quite the track record: More than 25 scientific papers and studies over nearly four years, 30 training courses, as well as an online platform featuring links to useful tools. Ten universities, research institutes, and companies leveraged the resources of European high-performance computing (HPC) centers and developed tools and interfaces that use open-source software and libraries to create modular workflows. These solutions enable researchers to access databases, manage research data and results, and analyze big data using artificial intelligence (AI) methods. The project was coordinated by the Czech National Supercomputing Center, IT4Innovations, with support from the Leibniz Supercomputing Centre (LRZ).

"EXA4MIND brings together high-performance and cloud computing, AI, and data analytics, as well as professional data management across Europe to unlock datasets that are more diverse and larger than conventional big-data collections," explains Dr. Stephan Hachinger, Head of Research Data Management at LRZ. "Researchers and small and medium-sized enterprises can find tools on the platform that allow them to automate their own analytics workflows across different computing resources, while professionally indexing and publishing their data according to the FAIR (findable, accessible, interoperable, reusable) principles. This makes research data and analytical results discoverable online for both humans and machines."

## AI workflows for HPC systems

One of the main drivers behind these efforts is the continuous growth in data volumes and the increasing use of AI in science and industry. At the same time, more researchers in companies and universities are working with cloud technologies and require additional expertise to establish their own data analytics workflows while integrating supercomputers into their processes. Today, most high-performance computing systems rely on GPUs, which support classical simulations and AI applications such as computer vision and machine learning. While industry often uses databases, object storage systems, and data-streaming technologies to support AI, HPC centers typically focus on adapting file systems, network connections, and input/output libraries. They generally assume that experienced users can efficiently navigate these environments. However, AI is attracting new user groups to HPC centers, many of whom use pre-existing models to analyze research data. As a result, there is growing demand for interoperable tools and standardized workflows that simplify entry into AI-driven supercomputing.

For work at supercomputing centers, EXA4MIND provides guidance on installing data backends, such as vector databases and object storage systems, as well as recommendations on when these solutions offer advantages over traditional storage approaches. EXA4MIND's Advanced Query and Indexing System (AQIS) supports the setup of high-performance analytics pipelines based on the open-source frameworks Airflow and Dask. AQIS also supports AI inference and querying data using natural language. "EXA4MIND opens opportunities for researchers and small companies to explore the services of supercomputing centers and work more efficiently with massive datasets," says Hachinger. "The modules of the platform can be deployed on systems ranging from laptops to HPC clusters, depending on the computational requirements." Through the interfaces developed by the project partners, researchers can access international data spaces as well as data from the European Open Science Cloud and EUDAT. "We explicitly encourage users to connect their workflows with these environments in order to search for data and share results," Hachinger added

| Workflows & Tools | |
|---|---|
| EXA4MIND | https://exa4mind.eu |
| APPLICATION CASES | https://exa4mind.eu/application-cases |
| WORKSHOPS | https://exa4mind.eu/exaki/training-centre |
| AQIS | https://docs.exa4mind.eu |
| EOSC | https://open-science-cloud.ec.europa.eu |
| EUDAT | https://eudat.eu |
| PROJECT LEXIS | https://lexis-project.eu/web |
| LEXIS TOOLS | https://docs.lexis.tech |

## Data analytics workflows in practice

To demonstrate the effectiveness of its workflows across different domains, EXA4MIND validated them through practical use cases from research and industry. The analytics tasks, tools and technical processes behind them are thoroughly documented and linked on the project website. Users can also access research papers and workshop materials. In addition, the team developed workflows specifically for supercomputing centers, enabling their computing power to be used more effectively for AI applications and scaled as required.

EXA4MIND represents an important contribution to Europe's data-driven research landscape. As with its predecessor project, LEXIS, which created a platform for orchestrating geographically distributed HPC and cloud systems, the practical experiences can feed into strategic IT initiatives. The EuroHPC Federation Platform, currently being developed by the EuroHPC Joint Undertaking, aims to provide future access to supercomputers, quantumcomputers, and AI systems.
*Susanne Vieser*

# New TOAR-II Theme Issue Provides Comprehensive Overview of Global Ozone

**Researchers at the Jülich Supercomputing Centre worked with 200 scientists and experts to evaluate the current state of ozone in the lower atmosphere.**

The Tropospheric Ozone Assessment Report (TOAR) is an international initiative launched by the International Global Atmospheric Chemistry (IGAC) project. Its goal is to provide an up-to-date assessment of the global distribution and trends of tropospheric ozone, from the Earth's surface to the tropopause (the boundary between the lowest layers of the atmosphere, meaning between the lower troposphere and the stratosphere), and to make the underlying data accessible to the research community. A new theme issue in the journal *Philosophical Transactions A* brings together the main scientific outcomes of the Tropospheric Ozone Assessment Report Phase II (TOAR-II), providing an updated picture of how ozone is distributed around the world, how it has changed over time, and the ozone distribution (and its changes) affect people, ecosystems, and the climate. More than 200 experts from over 35 countries contributed to this work. TOAR receives financial and logistical support from IGAC, Forschungszentrum Jülich (FZJ), the US National Oceanic and Atmospheric Administration (NOAA), and the World Meteorological Organization (WMO).

Tropospheric ozone is found in the lowest part of Earth's atmosphere. Unlike the protective ozone layer in the stratosphere, ozone near the Earth's surface is considered an air pollutant and greenhouse gas. This pollutant can harm human health and vegetation, but also plays an important role in the climate system.

The first phase of the assessment, TOAR-I (2014–2019), created the first comprehensive global assessment of tropospheric ozone, which brought together observations and analyses from around the world to improve our understanding of ozone's distribution in the atmosphere and its impact on climate, health, and vegetation.

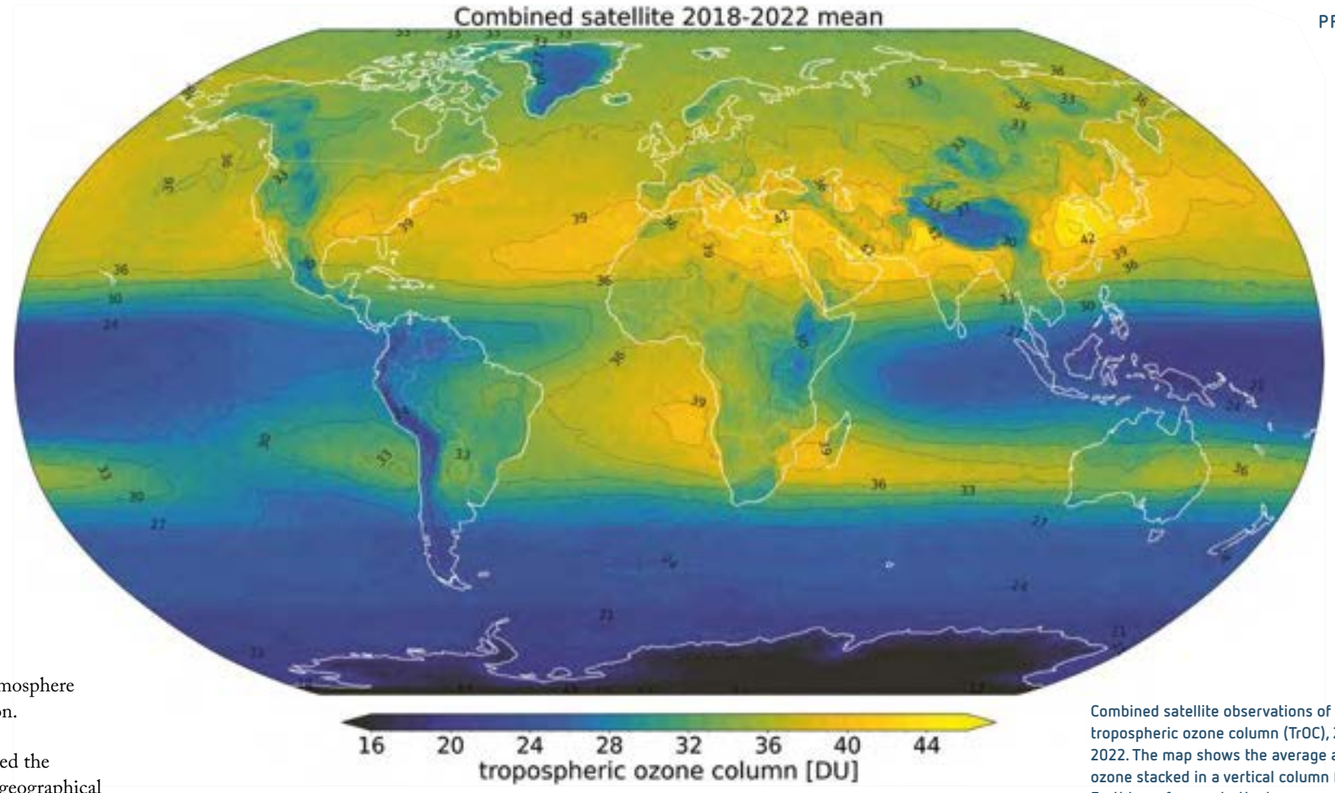
Building on this foundation, TOAR-II expanded the available information while also extending the geographical and scientific scope of the assessment. The project generated a special issue of 58 peer-reviewed publications, bringing together contributions from the international TOAR-II research community on a wide range of topics related to tropospheric ozone.

The Jülich Supercomputing Centre at FZJ, and in particular its Earth System Data Exploration (ESDE) research group, played a central role in developing and maintaining the terabyte-scale data infrastructure supporting TOAR. The TOAR Data Centre provides information and access to global air quality observations used by the International Tropospheric Ozone Assessment Report initiative. It brings together harmonized and verified surface ozone observations and related data, and provides a range of services and tools for exploring and working with the data, including an interactive dashboard, chatbot, data services and APIs, and software for integrating TOAR data into research workflows.

The TOAR Data Portal provides a central access point to a wide range of ozone-related datasets and their descriptions, including surface, satellite, aircraft, ozonesonde instruments, and LiDAR observations, as well as model data. Taken together, these resources provide an important foundation for analyzing changes in ozone. They also demonstrate the role of data infrastructure in enabling large-scale international scientific assessments such as TOAR.





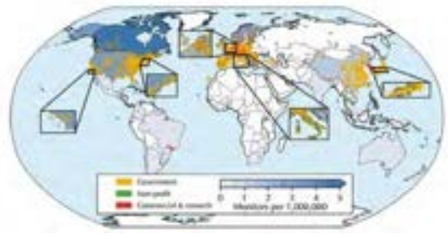



Combined satellite observations of average tropospheric ozone column (TrOC), 2018–2022. The map shows the average amount of ozone stacked in a vertical column from the Earth's surface up to the tropopause, based on 12 satellite records combined at a 1° spatial resolution. The colour scale shows ozone amounts in Dobson Units (DU)(1 DU = 4.46 10-4 mol per m²).

Martin Schultz, the ESDE group leader, served for over ten years as co-chair of TOAR-I and TOAR-II, helping shape the initiative and coordinate the work and publications arising from TOAR.

The new theme issue includes eight peer-reviewed publications. The TOAR-II overview paper provides a description of the TOAR-II process, its history, and main outcomes of the initiative. Seven additional publications take a closer look at specific aspects of the global ozone landscape. Together, these different assessments give a broad overview of tropospheric ozone – from its distribution and atmospheric transport to its impact on human health, ecosystems, and climate change.

The TOAR-II assessment represents the culmination of several years of international collaboration. Its results, together with the data and tools developed through TOAR, offer a valuable resource for scientists investigating air quality and atmospheric change, while also improving our understanding of the role of tropospheric ozone in a changing world. The work will continue with TOAR-III, the next phase of the international assessment.
*Martin Schultz & Sabine Schröder*



Location of in-situ, surface-level $O_3$ monitors by type.
Monitors shown in the global and regional maps are those included in the TOAR-II Database and which meet the inclusion criteria for our assessment (≥ 3 years of valid university-operated monitors). Multiple sites could be overlaid. Shading shows the number of monitors per 1,000,000 people using national-level 2020 population estimates from the UN World Population Prospects, and countries shaded in white have no monitors in the Database that meet inclusion criteria.

Related publication:

M. Schultz, R. J. Seguel & H. M. Worden (editors). Tropospheric Ozone Assessment Report II. In *Philosophical Transactions of the Royal Society A*. 2026. https://go.fzj.de/toar-ii

EVENTS

# Digital Twins Offer Tools for Managing the Environment and Climate Change

**A two-day symposium at HLRS explored how simulation can support city planning and conservation in adapting to a warming planet.**

Hotter temperatures, more intense storms, droughts and wildfires – the effects of climate change are visible everywhere. At the level of individual cities and communities, however, predicting specific local climate change impacts, developing effective adaptation or decarbonization strategies, and mustering the political will to implement change can be difficult. As a recent two-day symposium hosted by the High-Performance Computing Center Stuttgart (HLRS) showed, digital twins offer a toolbox for addressing such challenges.

A digital twin is a digital representation of a physical object, system, or process that is continuously synchronized with its real-world counterpart through data exchange. In this way, digital twins offer powerful tools for monitoring, analysis, simulation, and prediction. Building on this concept, an urban or regional digital twin integrates three-dimensional city models with real-time data from sensors, administrative databases, and other sources. Urban digital twins enable planners, decision-makers, and citizens to monitor urban processes, test future scenarios, and evaluate the impacts of planning decisions before they are implemented in the real world.

Researchers at HLRS have in recent years demonstrated numerous applications of digital twins that are relevant for environmental management, including for city and regional planning. Building on this experience, the HLRS symposium brought scientific experts in digital twins from multiple universities together with representatives of public agencies and small enterprises to discuss current applications of the approach, as well as the challenges and opportunities facing the field. Covering applications of digital twins for climate and weather modeling, landscape management, infrastructure development, and disaster prevention, the event illuminated how digital twins could help make our cities, industry, and landscapes more sustainable.



Thomas Obst of the HLRS Visualization Department has been developing a digital twin of wine growing terraces along the banks of the Neckar River. The approach could support landscape conservation efforts.



At an HLRS symposium, Prof. Kimmo Lylykangas of Tallinn University of Technology described how use of urban digital twins could help cities to reduce greenhouse gas emissions.

## Urban digital twins

Digital twins can be powerful tools for building community understanding, dialogue, and consensus in decision-making. As a partner in the Stuttgart Research Initiative "Discursive Transformation of Energy Systems" (DiTEnS), HLRS has been testing how digital twins can support participatory planning of new, more sustainable energy and heating concepts. HLRS scientist Kilian Türk explained that the DiTEnS team recently completed its first study, focusing on energy and heating systems at the University of Stuttgart's Vaihingen campus. HLRS created a digital twin that integrated multiple architectural datasets and simulation data of the campus' energy system. Presented using virtual reality in the HLRS CAVE visualization facility, it served as a basis for workshops that brought together stakeholders to plan the university's transition toward carbon neutrality. DiTEnS will soon begin a new pilot study focusing on a neighborhood in Stuttgart, supporting local residents and stakeholders in evaluating options for a more sustainable heating infrastructure.

Talks at the symposium also described other applications of digital twins for environmental management in cities. Prof. Kimmo Lylykangas of Tallinn University of Technology, Estonia proposed a concept using digital twins to better characterize the origins of greenhouse gas emissions in a city, knowledge that could empower policy makers to identify more impactful strategies for mitigating emissions. As Andres Maremäe of the City of Tallinn explained, real-time data integration in digital twins can also improve monitoring and timely repair of city infrastructure.

## Digital twins for landscape conservation

Multiple talks at the HLRS symposium also explored how digital twins can help protect the health of important regional and local landscapes. One dedicated session focused on the preservation of heritage vineyards on the terraced hills along the Neckar River. HLRS visualization scientist Thomas Obst described an ongoing project with the district of Ludwigsburg and a winegrowing startup to develop a digital twin of vineyards near Hessigheim. The model combines photogrammetry data captured using an aerial drone to create a three-dimensional model of the steep hillside landscape in virtual reality. Using AI-supported visual recognition software, data from drone flights taken over several months is used to identify specific locations in the vineyards that are undergoing erosion. The researchers can also simulate how water will flow down the hillside during heavy rainfall events, making it possible to identify and reinforce areas at danger of landslide.

Talks at the HLRS symposium considered other applications for landscape preservation as well. Andreas Hänsler of the University of Freiburg described his lab's work to develop a digital twin that could support real-time, high-resolution early warning of flash floods and help forecast potential property damage. Julian Frey, also from the University of Freiburg, introduced an approach that uses drones and laser scanners to collect high-resolution datasets in forests, and converts the data into digital twins to support more efficient forest management planning. Bettina Joa of the Breisgau/Black Forest Highlands has also been working with Freiburg researchers to visualize data sets that make it possible to identify heat islands that will be amplified by rising temperatures. This knowledge could support building strategies that improve airflow across the landscape and mitigate the effects of climate change.

During discussions, symposium participants remarked on current challenges facing the field, including the need for better standardization and coordination in digital twins across Germany's 16 states. Although it was clear that the technical challenge of doing so would be formidable, improved computing and data storage capabilities could improve model integration, enabling richer and more comprehensive simulations. *Christopher Williams*

# Official Launch of JION Quantum Computer at JSC

**During an early September event at Forschungszentrum Jülich, JSC, start-up eleQtron and high-ranking guests from politics, science, and industry inaugurated JION, a trapped-ion quantum computer developed in North Rhine-Westphalia.**

On September 3, 2026, staff from Forschungszentrum Jülich and Siegen-based start-up eleQtron came together to officially inaugurate the quantum computer JION. The ceremony served to underline the close collaboration between politics, research, and industry. North Rhine-Westphalia's (NRW's) Minister-President, Hendrik Wüst, and the NRW Minister for Culture and Science, Ina Brandes, emphasized the importance of JION for NRW as a research and industrial hub. As a symbolic part of the ceremony, a computing operation was performed publicly on JION for the first time.

JION, short for "Jülich trapped-ION quantum computer," is a digital, gate-based trapped-ion quantum computer. The system's qubits consist of ionized ytterbium atoms, which are held in a vacuum by electromagnetic fields and precisely controlled for computations. Unlike quantum computers based on superconducting qubits, JION's quantum processor does not need to be operated at temperatures close to absolute zero. To control the qubits, JION utilizes the magnetic gradient induced coupling (MAGIC) technology developed in Siegen. Microwaves and deliberately generated variations in the magnetic field allow the qubits to be controlled individually and connected directly to one another. This simplifies control and could facilitate the development of larger, more powerful quantum systems in the future.

## From project launch to an integrated system

JION is being realized as part of the EPIQ project – a development partnership for trapped-ion quantum computers in NRW. When the project was launched in Siegen in March 2024, the Jülich Supercomputing Centre (JSC) and eleQtron announced that they would jointly develop a modular supercomputer consisting of a quantum module and several classical digital modules. The Ministry of Culture and Science of NRW is providing around €21 million in funding to the partnership over a four-and-a-half-year period. EPIQ was born out of the quantum computing network,"EIN Quantum NRW."

With the inauguration, the project has developed into a fully operational system at Forschungszentrum Jülich. JION is set to open new prospects for applications in science and industry, such as in optimization tasks in logistics, transport, and process engineering, as well as in physics, chemistry, biology, medicine, materials research, and machine learning.

JION will be integrated into the Jülich UNified Infrastructure for Quantum computing (JUNIQ). JUNIQ provides access to various quantum systems and connects them to JSC's high-performance computing (HPC) infrastructure. With JUNIQ, certain computing operations can be offloaded to quantum systems when using classical JSC supercomputers.

## Next step: scaling and industrial application

During the inauguration ceremony, two funding approvals for the projects SQALING and Q-STAR.NRW were announced. SQALING aims to further advance the trapped-ion technology developed in Siegen into scalable, chip-based quantum computing platforms that could potentially also be used in industry. The objective of Q-STAR.NRW is to procure and integrate a scalable semiconductor quantum computer into JUNIQ to strengthen research and industrial collaboration in the Rhenish mining area. SQALING is being funded by NRW with up to around €25 million from European Union funds through the ERDF/JTF program in NRW while Q-STAR.NRW receives funding of up to around €25 million from structural funding for the Rhenish mining area under the Investment Act for Coal Regions (InvKG) and by the state of NRW. *Forschungszentrum Jülich Corporate Communications*







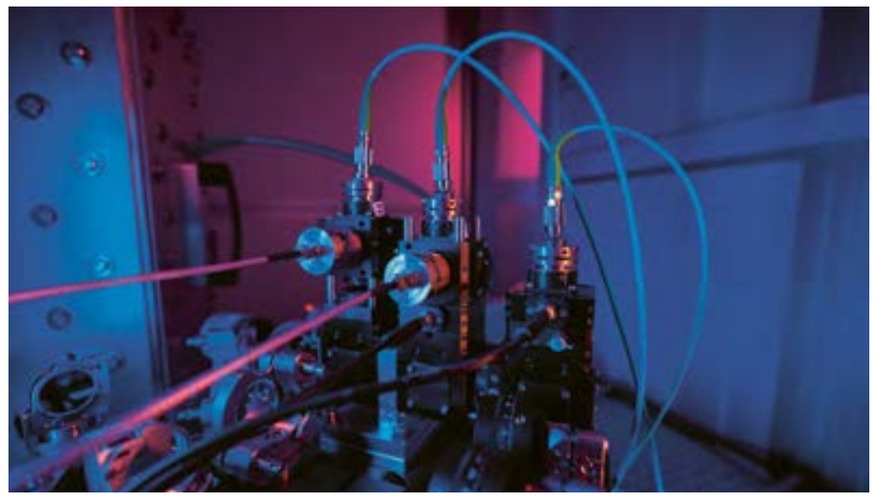

Inside the computer: Prof. Dr. Dr. Thomas Lippert, Director of JSC, and Prof. Dr. Kristel Michielsen, Director of JSC, Dr. Michael Johanning, CTO of eleQtron, Prof. Dr. Astrid Lambrecht, Chair of the Executive Board of Forschungszentrum Jülich, Minister Ina Brandes, Minister President Hendrik Wüst, Jan Leisse, CEO of eleQtron, are standing in the room with the JION ion-trap quantum computer.

**"With JION, we are expanding this unique user platform to include a trapped-ion system from North Rhine-Westphalia and connecting it to our HPC infrastructure – for research and industrial applications."**

Prof. Dr. Kristel Michielsen, Director of Jülich Supercomputing Centre, Head of the division HPC for Quantum Systems and Head of JUNIQ




Official launch of JION quantum computer at Jülich: Dr. Michael Johanning, CTO of eleQtron, Jan Leisse, CEO of eleQtron, Minister-President Hendrik Wüst, Minister Ina Brandes, Prof. Dr. Kristel Michielsen, Director of JSC, Prof. Dr. Astrid Lambrecht, Chair of the Board of Directors of Forschungszentrum Jülich, and Prof. Dr. Dr. Thomas Lippert, Director of JSC



Details of the quantum computer JION

# TPC Hackathon Fosters Research Collaboration Aimed at Improving AI for Science

**Agentic AI is currently a hot topic in the scientific community. During the Trillion Parameter Consortium's autumn hackathon at LRZ, teams presented practical solutions designed to make supercomputers more accessible and research projects more efficient.**



At the TPC AI Hackathon, interdisciplinary teams worked across fields and perspectives to develop solutions for a range of AI-related challenges.

In the Leibniz Supercomputing Centre's (LRZ's) seminar rooms, small groups stand in front of flipcharts discussing diagrams. Or they sit together, coding on laptops, only occasionally breaking the concentrated silence. This was the scene in early September, as 30 researchers from European and US supercomputing centers worked intensively on artificial intelligence (AI) systems or models for scientific research. The Trillion Parameter Consortium (TPC) Europe had invited participants to LRZ for its autumn hackathon. "At the TPC hackathons, plans and discussions turn into action," said Prof. Charles Catlett, Senior Computer Scientist at Argonne National Laboratory (ANL) and a member of the TPC strategy team. "During the event, participants can get to know potential collaboration partners and their working methods, plan projects together, work on ideas in a practical way, or share experiences and ensure that these spread more quickly throughout the community."

## Targeted research and promotion of AI

In light of the rapid global spread of generative AI tools and the risks they pose, the TPC was launched during the SC23 supercomputing conference by ANL, the Japanese RIKEN center, and Europe's Barcelona Supercomputing Centre (BSC). "TPC aims to become a sustainable movement for the use of AI in science and to promote research into it, as well as developing trustworthy programs," explained Catlett. In addition to regular hackathons, TPC organises conferences and workshops worldwide and publishes tutorials on AI-supported research.

Agentic AI featured prominently as a topic during the event. Special tasks were already formulated for the TPC hackathon in the context of agentic use of AI workflows, in which large language models coordinate jobs together with other tools in data processing and interact with them to carry out certain actions. For example, one of the teams at LRZ worked on agentic workflows to develop hypothesis generation for proteins.

Another working group from the Gauss Centre for Supercomputing (GCS) developed initial workflows to enable an AI agent to adapt code to various supercomputers, thereby simplifying access to a wide range of HPC systems. "For me, the most important thing during the TPC hackathons is the exchange of ideas; it also allows me to learn about different approaches or how technology is used," says Mitja Sainio, a Machine Learning (ML) engineer at the Finnish CSC – IT for Science center, who collaborated with the GCS team. In addition, the hands-on work helps building networks of contacts that AI specialists can turn to for practical advice in their day-to-day work.

> "For me, the most important thing during the TPC hackathons is the exchange of ideas; it also allows me to learn about different approaches or how technology is used"
>
> Mitja Sainio, Machine Learning (ML) engineer

## Creating networks to address practical issues

Energy efficiency was another major topic for hackathon participants. "The more Agentic AI is used, the more energy is consumed due to inference," noted LRZ researcher Ajay Navilarekal. "We need to limit that." In light of initial incidents and risky system attacks by AI agents, the risks and boundary violations associated with agent-based systems were also on the agenda: to protect data and prevent these digital assistants from going rogue, another team from GCS worked on developing secure strategies on how authentication and authorization workflows would look like for AI agents.

"The seven working groups pursued different objectives at the TPC Hackathon," observed Navilarekal. "Some wanted to develop an initial, simple prototype that they could continue working on separately, while others arrived with code that they wanted to optimise and expand upon. In any case, working together has led to progress." It also provided inspiration – Catlett from ANL prompted the discussions to consider how AI agents could safeguard and enrich urban life by analysing environmental data from cameras and sensors. Dr. Miguel Vazquez from BSC and Dr. Arvind Ramanathan from ANL, who develop AI systems for bioinformatics and genomics research, gained new impetus, potential collaboration partners, and ideas for projects. "The TPC community brings together a unique pool of expertise," notes Vazquez. "This will enable us to successfully establish AI and agentic AI across all areas of research." *Susanne Vieser*

# News Briefs

JSC



## 19th ADAC Symposium and Workshop Takes Place in Cologne and Jülich

The virtual Accelerated Data Analytics and Computing Institute (ADAC) serves as a forum for exchange related to accelerated HPC for large HPC centers from the US, Europe, Asia and Australia. In the first week of September, the 19th ADAC Symposium and Workshop took place in Cologne. The symposium featured a full schedule of interesting talks, starting with the keynote, "Supercomputing in Jülich: Quo Vadis?" by JSC director Thomas Lippert. The talk was followed by experience reports, future prospects and outlooks from various international HPC representatives. Within the workshop, the ADAC working groups discussed a variety of current issues, with agentic AI emerging as a key topic across all groups. Event participants toured the Forschungszentrum Jülich campus to visit Europe's fastest supercomputer, JUPITER, and the quantum infrastructure, JUNIQ. JSC contributed to the successful event by acting as the host and main organizer, as well as by delivering talks and inviting experts.

Learn more: https://www.fz-juelich.de/en/jsc/news/news-items/news-flashes/2026/19th-adac-symposium-and-workshop




## Jülich-Aachen Start-up Paves the Way for Scalable Quantum Computers

ARQUE Systems, a spin-off from Forschungszentrum Jülich and RWTH Aachen University, is developing a quantum computer based on conventional semiconductor technology. The aim is to use established materials and manufacturing processes to overcome a key challenge in quantum computing: scaling up to a larger number of qubits (quantum bits). The system is currently being installed in Jülich and, once the testing phase is complete, is set to become the first quantum computer of its kind to go into operation at JSC. The integration of the ARQUE system named JAQUI – Jülich ArQUe semIconductor quantum computer – into the JUNIQ quantum computing infrastructure will ultimately give researchers and businesses access to another innovative technology platform. This approach is an example of how existing semiconductor technologies can be utilized for the further development of scalable quantum computers.

Learn more: https://www.fz-juelich.de/en/news/archive/press-release/2026/julich-aachen-start-up-arque-systems

## GCS Industry Day 2026 – Supercomputing and AI Driving Innovation in Industry

HPC and AI have the potential to accelerate innovation and strengthen the economic competitiveness. Bringing these transformative digital technologies closer to industry was the main objective of the GCS Industry Day, held on June 17, 2026 at JSC. Participants learned how they can benefit from JSC's world-class computing infrastructure, extensive expertise, and the latest projects in HPC and AI. The program focused on practical industrial applications, with company representatives sharing first-hand experiences of using supercomputing and AI in their businesses. Presentations highlighted AI for autonomous driving, HPC for fusion power plant development, and computational fluid dynamics (CFD) simulations to optimize large-scale industrial facilities. A particular strength of the event was the vibrant exchange between researchers and industry representatives during panel discussion, Q&A sessions and networking breaks. The event concluded with guided tours of JUPITER, Europe's first Exascale supercomputer. The event was organized by JSC in cooperation with the Gauss Centre for Supercomputing (GCS) and the Chambers of Industry and Commerce (IHK) of Aachen, Cologne, and Middle Lower Rhine.

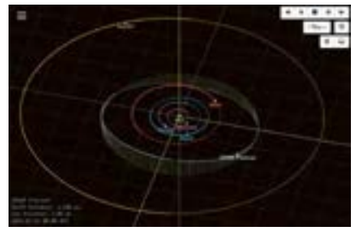
Learn more: https://www.fz-juelich.de/en/jsc/news/news-items/news-flashes/2026/gcs-industry-day-2026-supercomputing-and-ai-driving-innovation-in-industry



## Ericsson and Forschungszentrum Jülich to Develop Advanced AI for 6G

Ericsson, a global leader in telecommunications infrastructure, and Forschungszentrum Jülich are collaborating to develop technologies for continued evolution of 5G, and for the introduction of 6G networks across Ericsson's product portfolio. The project aims to leverage JUPITER, Europe's first exascale supercomputer, to design and test new artificial intelligence solutions for the complex demands of 6G. Key objectives include exploring highly efficient, brain-inspired computing approaches (neuromorphic computing) to handle intense network tasks and strengthen Europe's digital infrastructure. The first commercial 6G services are expected around the year 2030.

Learn more: https://www.fz-juelich.de/en/news/archive/press-release/2026/mou-ericsson-fzj



## Asteroid Named After Susanne Pfalzner

The Working Group for Small Body Nomenclature (WGSBN) of the International Astronomical Union (IAU) regularly names small bodies in space after scientists who made significant contributions to our understanding of the universe. In July 2026, they named asteroid (20596) = 1999 $RX_{188}$ after astrophysicist Susanne Pfalzner in recognition of her contribution to the study of minor bodies of the solar system. Susanne leads the Simulation and Data Lab Astronomy and Astrophysics at the JSC. To give an example of her extraordinary work, in 2024, her team and collaborators solved a cosmic mystery: Their HPC simulations revealed that the flyby of a star billions of years ago explains some dynamics of the solar system beyond Neptune. Apart from her work at JSC, she is an associate professor at the University of Cologne and performs research at the Max-Planck for Radioastronomy, Bonn.

Learn more: https://www.fz-juelich.de/en/jsc/news/news-items/news-flashes/2026/asteroid-pfalzner




## JSC Training Staff Publish Blueprint for HPC Introductory Courses

Training course leaders at JSC recently presented in a new paper the modular concept of the JSC course, "Introduction to Supercomputing at JSC – Theory & Practice". The concept is based on many years of experience in delivering and optimizing the course and can be used by other HPC centres to develop and improve their own introductory training activities. The paper examines the history of the course, its structure and content, as well as the adjustments made in recent years. The authors focused on several key findings: A modular structure helps to successfully integrate participants with varying levels of prior knowledge; Practical, hands-on exercises and how they play an important role in effectively reinforcing knowledge; And the importance of providing materials for independent study to promote asynchronous learning.
The paper also explains how to boost the motivation of participants, how evaluations can be used to improve the course, and which checklists have been effective in the planning and implementation phases.

Learn more: https://www.fz-juelich.de/en/jsc/news/news-items/news-flashes/2026/blueprint-for-creating-an-hpc-introductory-course-at-an-hpc-centre

HLRS



## European Agentic AI Bootcamp

From July 20-26, 2026, 250 participants from across Europe along with 30 trainers, tutors, and technical assistants gathered online for the European Agentic AI Bootcamp. The event was designed for developers, researchers, businesses, public sector organizations, start-ups, and SMEs interested not only in understanding AI agents, but also in building, evaluating, and integrating them into real-world workflows. Jointly organized by AI Factory Austria, HammerHAI, Gaia AI Factory, and PHAROS-CY – as well as partners from the high-performance computing and AI ecosystem including HLRS and LRZ – the bootcamp focused on practical building blocks for modern AI agents. Participants learned how to develop Agent Skills, build HTTP-based MCP clients and servers, and integrate these technologies into LLM workflows. Near the end of the program, participants applied what they had learned in a hands-on challenge using an existing LLM workflow. A key focus of the bootcamp was to facilitate collaboration across multiple AI Factories. As a collaborative program, the partners not only pooled technical expertise but also created a shared European environment for learning and practical collaboration. Participants gained access to knowledge from different national AI ecosystems and prepared for applications that can be further developed through future European hackathons and AI projects.

Learn more: https://www.hlrs.de/news/detail/european-agentic-ai-bootcamp

## RItrain-DITGIT to Improve Skill Development for Research Infrastructure Staff

The rise of artificial intelligence, machine learning, and federated data access promise exciting opportunities for many scientific domains. However, the diversity and complexity of compute offerings as well as the need to manage rapidly expanding data resources can make it difficult for staff in research infrastructures (RIs) to identify, access, and integrate the most relevant digital services into sustainable data analysis operations. A new European project called RItrain-DIGIT recently launched to bridge this skills gap. Funded under the European Union's Horizon Europe programme, the 1.5 million Euro initiative is coordinated by the EGI Foundation and brings together organizations from across the European computing landscape, including HLRS, HUN-REN Sztaki, Cyfronet, and PRACE. The project will implement a comprehensive training programme structured into learning pathways addressing topics related to HPC and cloud systems, AI and data techniques, federated access, and IT service management. A "mobility programme" will also support staff exchanges to promote knowledge transfer and collaboration across institutions, countries, and domains. Through these efforts RItrain-DIGIT aims to ensure that research infrastructure staff gain a holistic understanding of the wider compute and data ecosystem, acquire sustainable competencies, and reduce fragmentation across research computing infrastructures, ultimately strengthening the capacity of RIs to contribute to Europe's digital transformation.

Learn more: https://www.hlrs.de/news/detail/new-project-will-improve-skill-development-for-research-infrastructure-staff



## theGATE International Festival for Art and Technology

Emerging technologies – including artificial intelligence, simulation, and digital networking – are transforming contemporary society. To examine their impacts and to illuminate the opportunities they offer, HLRS hosted a weeklong international festival that brought together more than 150 cultural and innovation managers, artists, and scientists. Named theGATE, the festival was the second in an ongoing series that launched in 2024 in Barcelona. The Media Solution Center Baden-Württemberg – cofounded by HLRS – organized this year's event, which featured three interrelated tracks. The first was the eCulture Convention, which offered lectures and panel discussions to foster dialogue among leaders from across Europe's CCSI community. In parallel, the ReACH Workshop offered artists the opportunity to interact directly with scientists and experts in digital technologies and to receive practical guidance on projects in development. In the final track, theGATE presented ten recent artistic projects that employ advanced digital technologies, including several created during artistic residencies at HLRS supported by the European Union's S+T+ARTS E(C)HO program. In the coming months, the Media Solution Center will be developing a publication the organizers describe as a "manifesto," which will document the knowledge produced during the festival and define steps needed to strengthen Europe's culture and creative industries. The next iteration of theGATE Festival is expected to take place in May 2027.

Learn more: https://www.hlrs.de/news/detail/hlrs-hosts-thegate-international-festival-for-art-and-technology



## Democratic Technology? Technical Democracy?

When new technologies emerge, they can influence democratic structures, values, and processes. The spread of printing in Europe beginning in the 15th century led to a proliferation of political pamphlets, for example, creating new venues for public debate. In the twentieth century, television increased the visibility of public figures while creating communities that cross geographical boundaries. More recently, social media, memes, and deep fakes have challenged Western society's assumptions about how politics should function. Technology can open new possibilities for democracy by improving participation, transparency, and collective decision-making. Paradoxically, however, technology can also make democracy more difficult as the diversity of voices and interests represented in public debate increases. In a three-day conference co-organized by HLRS's program in the Philosophy of Computational Sciences in collaboration with the TU Darmstadt Department of Philosophy and editors from the journal *Jahrbuch Technikphilosophie*, historians, social scientists, and philosophers analyzed how technology affects democratic processes, as well as how political concepts shape the design and use of novel technologies. Looking at a wide range of topics – including AI-powered surveillance and decision-making, living laboratories (*Reallabore*), citizen science, cryptocurrencies, nuclear energy, and radio, among others – participants explored how society negotiates concepts like knowledge, power, legitimacy, and authority in the context of technologically mediated politics. Conference proceedings will be published in an upcoming volume of the *Jahrbuch Technikphilosophie*.



## HLRS Scientists Win Research Poster Award at ISC26

Daniel Renschler and Jonathan Schäfer, two student researchers at HLRS, accepted the third prize in the Best Poster Award competition at this year's ISC High Performance Conference. Their poster, titled "SpMV for the Cerebras Wafer-Scale Engine," reports on the successful implementation of a sparse matrix-vector product (SpMV) method on the Wafer-Scale Engine (WSE), a large-scale computer chip manufactured by California-based startup Cerebras. Developed for artificial intelligence training and inference applications, the WSE's unique architecture offers dramatically faster speeds in comparison to conventional AI processors. Renschler and Schäfer used the Wafer-Scale Engine to accelerate a highly parallelized component of a typical simulation workflow that includes the SpMV method. They also performed weak-scaling and strong-scaling experiments, revealing bottlenecks that affected application performance. This enabled them to suggest optimization strategies that could improve performance of SpMV methods on the WSE in the future. The research was conducted within the context of the HLRS Future Computing Group, a multidisciplinary research team at HLRS that tests and evaluates emerging hardware concepts and their suitability for typical HPC applications. Renschler, Schäfer and Future Computing Group leader Johannes Gebert conducted this experiment together with Mark Parsons, director of EPCC in Edinburgh.

Learn more: https://www.hlrs.de/news/detail/hlrs-scientists-win-research-poster-award-at-isc26

LRZ




## Euro-Q-Exa Welcomes Users

Following its pilot phase from March–July 2026, the Euro-Q-Exa system at LRZ officially enters production mode. Researchers and companies can now apply for computing time on the EuroHPC Joint Undertaking quantum computer. Designed for integration into high-performance computers, the system minimizes latency and maximizes computing power, particularly in hybrid workflows. It is based on the superconducting "Radiance" technology from IQM Quantum Computers, offers 54 qubits of performance, and is currently being connected to HPC systems at LRZ. In the medium term, Euro-Q-Exa will be expanded to up to 150 qubits.

**Learn more: https://www.lrz.de/en/news/detail/euro-q-exa-start-of-operations**

## LRZ and Partners Look Back on One Year of BayernKI

For the past year, BayernKI has been providing a centralized infrastructure for AI applications in science, research, and education. To this end, various high-performance and supercomputing clusters belonging to Gauss Allianz members LRZ and the Erlangen National High Performance Computing Center (NHR@FAU) have been consolidated, including systems specifically designed for AI applications. In total, 320 NVIDIA H100 processors, each with 94 gigabytes of HBM2e memory, are available, and further acquisitions for BayernKI are planned. Since its launch, the BayernKI systems have been used in around 200 projects, with approximately two million computing hours allocated for this purpose. In addition to the development of proprietary AI models and analytical methods, BayernKI also supports inference as well as the deployment and customization of pre-trained models.

**Learn more: https://www.lrz.de/en/news/detail/one-year-of-bayernki**



## With planqc, LRZ Brings in New Quantum Technology

Technology provider planqc is currently developing a quantum computer based on neutral atoms specifically for LRZ. In preparation for the installation, the Garching-based start-up recently delivered 800 kilograms of innovative technology, including a specialized server cabinet that will eventually house the laser technology and quantum processor. The shipment has now been installed in the cleanroom of LRZ's Quantum Integration Centre. The hardware has been connected to the LRZ's infrastructure and is now being thoroughly tested, pushed to its limits, and further evaluated against various reference values. Through this approach, specialists at planqc and LRZ are ensuring that the universal quantum computer will quickly settle into its new home at LRZ and be ready to perform computations for research.

**Learn more: https://www.lrz.de/en/news/detail/new-quantum-technology-for-the-lrz**



## GCS Director Kranzlmüller Receives Bavarian Science Award

Markus Blume, Bavarian Minister for Science, recently honored Prof. Dr. Dieter Kranzlmüller, Chairman of the GCS Board of the Directors and Director at LRZ, with the Pro Meritis Scientiae et Litterarum award in recognition of his contributions to Bavarian science. "As a leading researcher and director of the Leibniz Supercomputing Centre, he combines scientific excellence with a world-class research infrastructure. He creates the conditions that enable Bavaria's researchers to continually break new scientific ground," Blume said during an award ceremony on July 17, 2026.

**Learn more: https://www.lrz.de/en/news/detail/award-for-prof-dieter-kranzlmueller**





# Staff Spotlight: Far From Home, JSC Post-Doctoral Researcher Stands at the Intersection of Quantum Computing and High-Performance Computing

Dr. Vrinda Mehta, JSC




Dr. Vrinda Mehta's passion for mathematics as a child led her across the globe and into the midst of cutting-edge quantum computing research at the Jülich Supercomputing Centre.

Dr. Vrinda Mehta grew an interest in physics from an early age. As early as middle school, she found herself drawn to mathematics. Specifically, she found great satisfaction in working through complex problems step-by-step and arriving at exact solutions. That interest laid the groundwork for what would ultimately become the passion that would drive her into a doctoral degree far beyond her hometown of Delhi, India – physics.

"Mathematics gave me the pleasure of solving abstract problems, but physics added another layer: it allowed me to use that mathematical language to understand how the world works on a fundamental level," she said. That desire to use math to interpret the laws that describe our physical world led Mehta to major in physics at Miranda House College at the University of Delhi.

When she finished her undergraduate studies, Mehta felt compelled to look further afield for graduate school. She had spent her entire life to that point in Delhi and felt compelled to step outside of her comfort zone, meaning she started looking outside of India. In her investigation, she was drawn to RWTH Aachen University. While she knew she loved physics, she was not yet sure what specific research domain she wanted to focus on in graduate school, and RWTH Aachen offered a place for her to explore a variety of options.

"One of the things I appreciated about RWTH was that I could explore elective courses, such as in neuroscience, while pursuing a degree in physics," she said. "It gave me the freedom to follow my curiosity without needing to make an early decision about where my studies would lead."

As Mehta neared the end of her master's degree, she took a quantum information course that set her on the path to her current focus. She noted that the fusion of physics, math, and computation spurred her interest, and she started working with the Jülich Supercomputing Centre's (JSC's) Quantum Information Processing (QIP) group, led by JSC Director Kristel Michielsen. Mehta indicated that the combination of this group's cutting-edge research and connection to world-class computing resources encouraged her to apply to work in Jülich. "I was already increasingly interested in the computational side of quantum physics, so the possibility of working in an environment with access to some of Europe's most powerful supercomputers was very exciting," she said.

Mehta started working on her master's thesis in Jülich in 2018, continued to work there during her PhD, and has been a post-doctoral researcher in the QIP group since she finished her doctoral program. She has primarily focused on quantum annealing – particularly on the effects of introducing modifications to the standard approach. She has worked directly on JSC's D-Wave system as well as modeled quantum annealing systems on traditional high-performance computing (HPC) resources, such as JSC's JUWELS and JUPITER.

She indicated that in her eight years connected to JSC, she has been excited by how much quantum computing activities have grown and expanded during that time. As JSC has invested in prototype systems and quantum simulators, it has taken Mehta's research from a largely theoretical realm to one more closely connected to real-world applications. In the last several years, JSC has also built out the Jülich UNified Infrastructure for Quantum Computing, or JUNIQ, which closely coordinates and integrates the various quantum systems on the Jülich campus with one another and JSC's HPC infrastructure.

For Mehta, proximity to the infrastructure is only one aspect of her job at JSC that she enjoys. "One of the things I particularly appreciate about JSC is being in an encouraging and intellectually open work environment," she said. "I work with people from diverse backgrounds with very different areas of expertise, and that has contributed enormously to my own growth. Some of the best conversations are those that go beyond our day-to-day work and pose broader questions about science, philosophy, and how we understand the world."

In a fast-moving field like quantum computing, Mehta does not exactly know where her career will take her, but she is excited to focus on near-term challenges, such as probing the current capabilities of JSC's suite of quantum systems. As the hardware continues to get more powerful, Mehta and the QIP team are able to design increasingly complex tests and experiments with the hardware, continuing to push forward state-of-the-art quantum computers and methods for quantum systems. Mehta indicated that one of the challenges is to demonstrate how these systems serve complementary roles with traditional modeling and simulation, and that continuing to refine JSC's hybrid workflows will not only advance quantum computing technology and methods, but also support other scientific and engineering domains in leveraging these technologies for their own research problems.

Wherever the field goes, Mehta feels confident that her time in Jülich has prepared her to be a leader in the field. "JSC has been a very important part of preparing me for my career path," she said. "It has given me the opportunity to work with people with different types of expertise and learn how research develops through collaboration. The combination of technical breadth, scientific independence, and exposure to such a diverse research community is one of the strongest foundations I could have for whatever the next stage of my career brings." *Eric Gedenk*



# 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.

CONTACT

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

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





## Compute servers currently operated by HLRS

| System | Size | Peak Performance (Tflop/s) | Purpose | User Community |
|---|---|---|---|---|
| HPE 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 neighboring 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<br>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.

## 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.

Leibniz Supercomputing Centre
of the Bavarian Academy of Sciences and Humanities

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 (4TH 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;<br>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 (HARDWARE SIMULATOR)**<br>N/A | Suitable for computation of quantum algorithms for systems up to 38 qubits | Quantum simulation | Bavarian Universities |
| | **EVIDEN QAPTIVA 2 (HARDWARE SIMULATOR)**<br>N/A | Suitable for computation of quantum algorithms for systems up to 38 qubits | Quantum computation | Bavarian Universities |
| | **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 and, possibly, Transformer Engines. H100: structural sparsity Tensor Core performance and, possibly, second-generation Transformer Engines. B200: Structural sparsity Tensor Core performance and, possibly, Transformer Engines.

## 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/53170-2610-1003

Contributing authors: Eric Gedenk, GCS, eric.gedenk@lrz.de; Christopher Williams, HLRS, williams@hlrs.de; Susanne Vieser, LRZ, susanne.vieser@lrz.de; Martin Schultz, JSC, m.schultz@fz-juelich.de; Sabine Schröder, FZJ, s.schroeder@fz-juelich.de; Forschungszentrum Jülich Corporate Communications

Photo credits:
Cover p.1: istockphoto.com/ugurhan; inside front cover: istockphoto.com/ilyaska p.6: LRZ/V.Hohenegger; p.7: LRZ/A.Podo; p 9: Christopher Williams; p 11: Forschungszentrum Jülich/Sascha Kreklau; p 12: Klocke et al. (2025), licensed under CC BY 4.0; p 13: Christian Schiffer, INM-1 of FZJ; p 14: Dani Zemba, Penn State; p 16: unsplash.com/C.Oskay; p.17: LRZ/S.Vieser; p.18: Mesh forms structure: Heinecke et al., Gordon Bell Prize Finalist, SC14, Wollherr et al., JGR, 2019, High frequent waves: Taufiqurrahman et al., 2022; p 20/21: Sebastiano Bernuzzi; p.23: National Institute of Geophysics and Volcanology (INGV), Italy; p 25: EXA4MIND; p 27: Hubert, D. (2026). Tropospheric Ozone Assessment Report II: Past and present tropospheric ozone using satellite observations. Phil Trans A; p 28/29: HLRS; p 31: Forschungszentrum Jülich/Kurt Steinhausen, Land NRW/Marius Becker, sichtplan; p 32: LRZ/V.Hohenegger; p 34: ADAC, Jülich Supercomputing Centre; p 35: Forschungszentrum Jülich/Kurt Steinhausen, NASA; p 36: AI Factory Austria, Matthias Stroezel, SSC-Services GmbH,; p 37: HLRS, Philip Loeper; p.38: LRZ/V.Hohenegger, Sabrina Schulte; p 39: Michael Bresser, FZJ; p 41: University of Stuttgart/Regenscheit; p.42: HLRS/Julian Holzwarth; p 44: Forschungszentrum Jülich/Sascha Kreklau; p 46: LRZ

**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.

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 of 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