Seizing the AI Future

HLRS

High-Performance Computing Center Stuttgart

2024/25 Annual Report

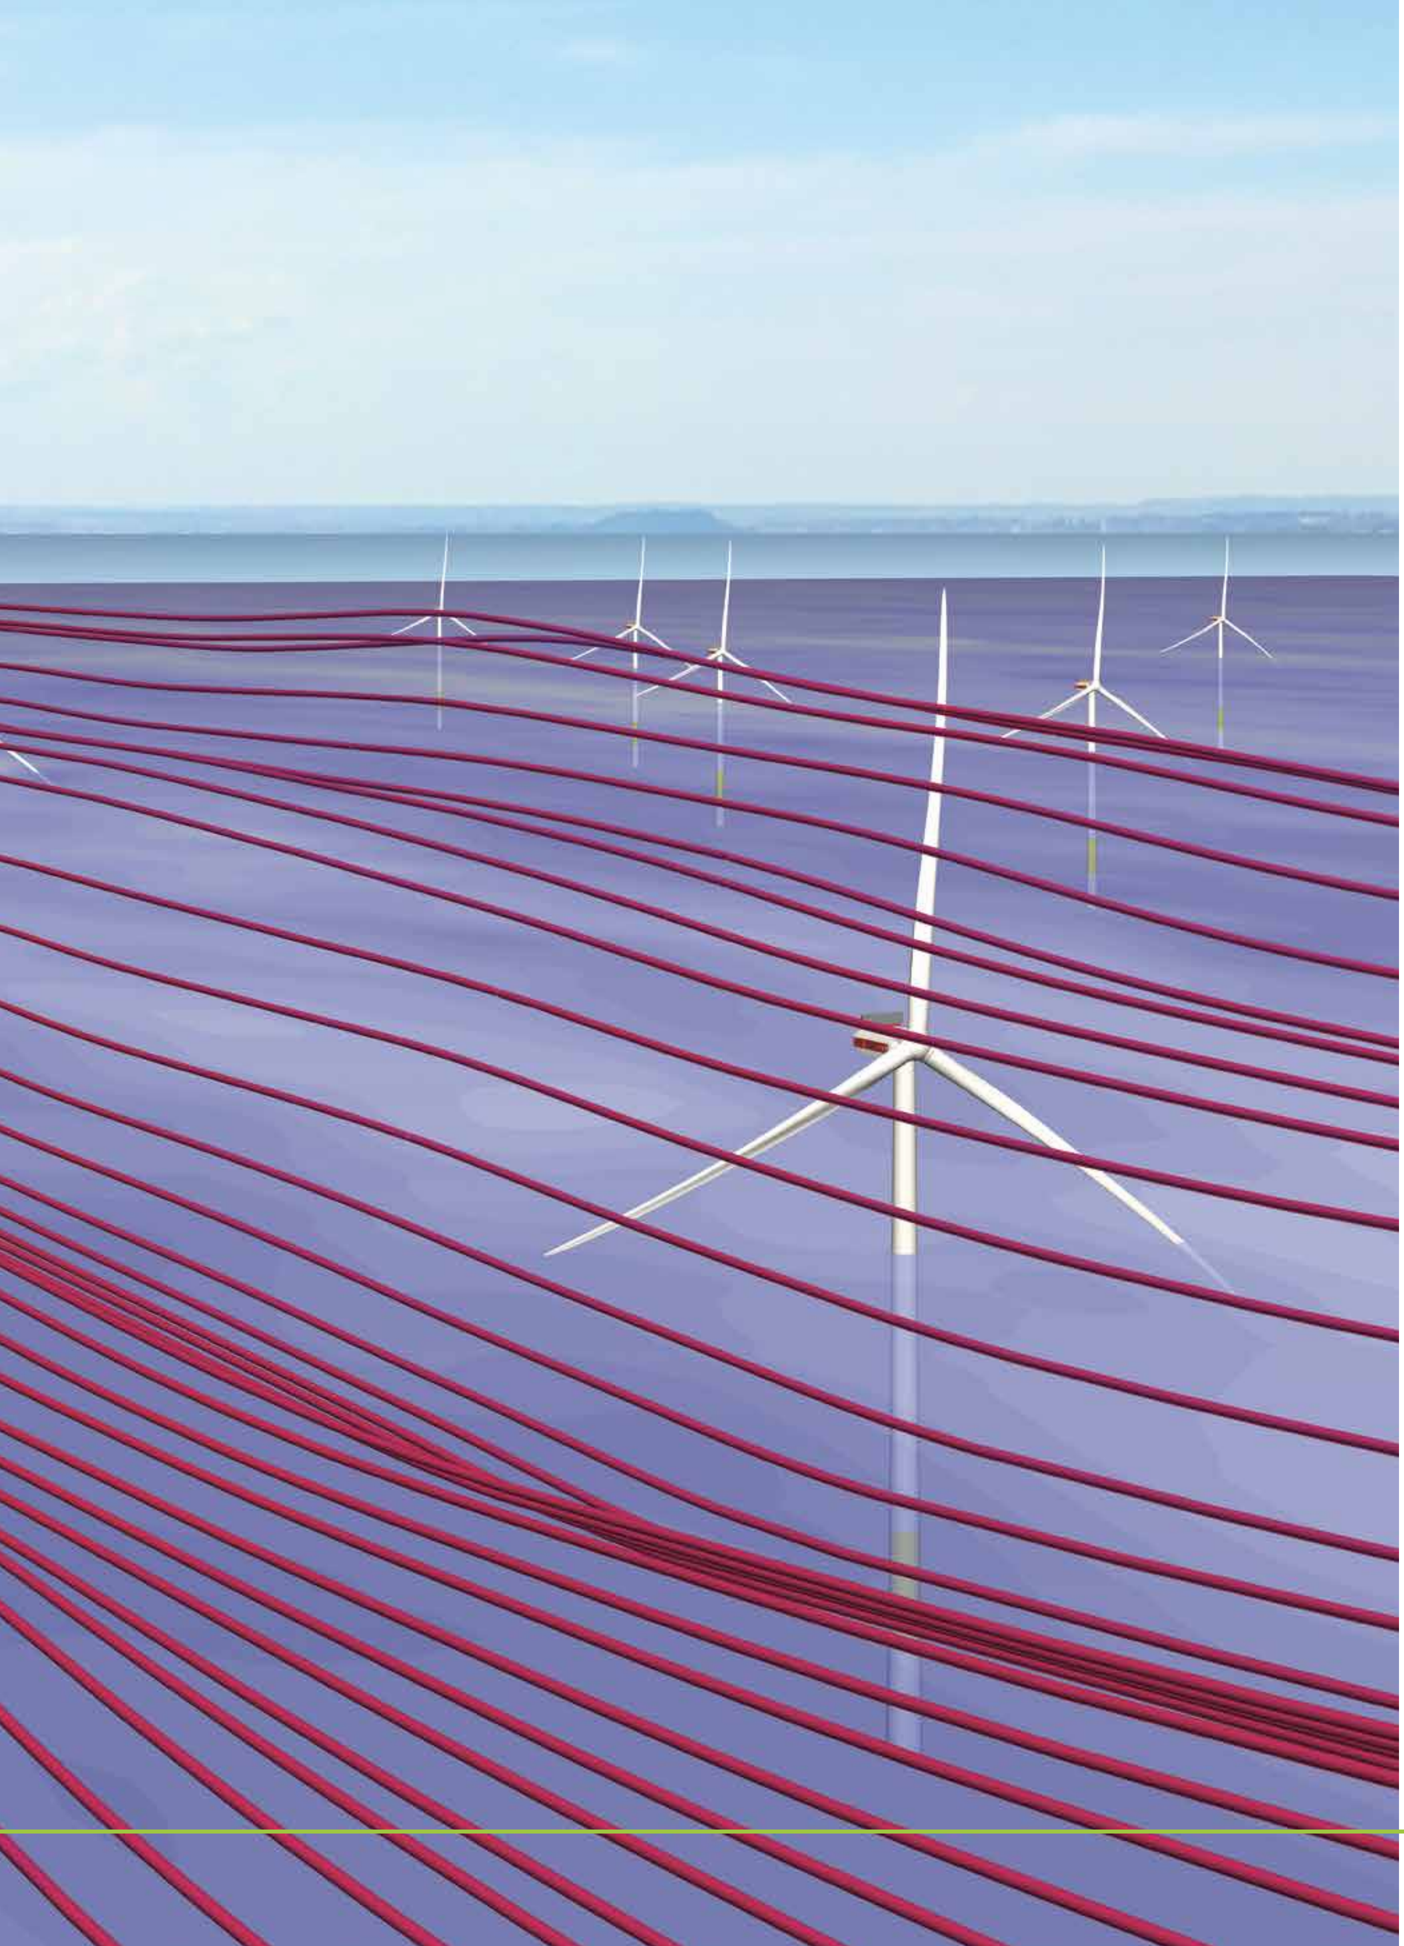

2024/25 Annual Report

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

Director's Welcome

Grußwort

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

Welcome to the 2024/25 Annual Report of the High-Performance Computing Center of the University of Stuttgart (HLRS). It is our pleasure to present the highlights of an exciting and successful year for the center, our user community, and our partners.

Driven by recent advances in high-performance computing (HPC), 2024 was an important transitional year for HLRS. One reason was the installation of Hunter, our newest flagship supercomputer. Manufactured by Hewlett Packard Enterprise, Hunter is a particularly important system for HLRS. It moves away from our longtime operation of CPU-based architectures to one based on AMD's accelerated processing unit (APU), which offers higher performance and greater energy efficiency by closely integrating CPUs and GPU accelerators. A key activity in 2024 involved supporting longtime users of HLRS's systems in porting their codes to GPUs. Helping them to optimize codes for the new architecture will continue to be a high priority in 2025. With Hunter serving as a "stepping-stone" system, this will put our users in an ideal position to fully leverage the possibilities of Herder, which will offer a major leap in performance when it arrives in 2027.

In another significant transition in 2024, HLRS expanded on its traditional emphasis on numerical simulation by moving boldly into the field of artificial intelligence. This included coordinating a successful proposal to the EuroHPC Joint Undertaking to establish HammerHAI, Germany's first "AI Factory." Building on HLRS's expertise in HPC for industry, engineering, and global challenges, HammerHAI will install a large-scale, AI-optimized EuroHPC JU supercomputer at HLRS, and establish a comprehensive suite of solutions, resources, and training opportunities that will lower the barriers that industry and research currently face in using artificial intelligence. We are very pleased to collaborate on HammerHAI with partners at the Leibniz Supercomputing Centre, Karlsruhe Institute of Technology, Gesellschaft für wissenschaftliche Datenverarbeitung Göttingen, and SICOS BW, and that our Gauss Centre for Supercomputing partners at the Jülich Supercomputing Centre will host a second German AI factory, called JAIF. Together, our efforts will establish infrastructure and expertise that Germany and Europe will need to remain globally competitive in a rapidly evolving technological landscape.

In light of these developments, a special feature section in this Annual Report focuses on how high-performance computing centers are evolving with the rise of artificial intelligence. In addition to providing more information about HammerHAI, we report on Stuttgart-based AI consultancy Seedbox.ai, who began training large language models on Hunter even before its inauguration. Our partnership with Seedbox.ai is helping HLRS to develop a comprehensive portfolio of knowledge and services to support the development and scaling of AI applications, and offers a strategically relevant model for promoting AI-based innovation in accordance with the goals and requirements of the European AI Act. These pages also include an interview

Herzlich willkommen zum Jahresbericht 2024/25 des Höchstleistungsrechenzentrums der Universität Stuttgart (HLRS), in dem wir Ihnen die Höhepunkte eines spannenden und erfolgreichen Jahres für das Zentrum, unsere Nutzergemeinschaft und unsere Partner zeigen.

Getrieben von den aktuellen Entwicklungen im Höchstleistungsrechnen (HPC) begann im Jahr 2024 eine Übergangszeit für das HLRS, u.a. mit der Anlieferung und Installation von Hunter, unserem neuen Supercomputer. Hunter wird von Hewlett Packard Enterprise hergestellt und ist sehr wichtig für das HLRS. Es geht weg von CPU-basierten Architekturen und hin zu einem System, das auf der Accelerated Processing Unit (APU) von AMD basiert, die dank des engen Zusammenspiels zwischen CPUs und GPU-Beschleunigern eine höhere Leistung und bessere Energieeffizienz bietet. Eine der wichtigsten Aktivitäten im Jahr 2024 war die Unterstützung langjähriger System-Nutzer:innen bei der Portierung ihrer Codes von CPUs auf GPUs. Hunter als Übergangssystem bereitet unsere Nutzer:innen darauf vor, später das gesamte Potenzial von Herder, unserem kommenden Supercomputer auszuschöpfen. Bei seinem Einzug im Jahr 2027 wird Herder einen großen Leistungssprung bieten.

Seit 2024 erweitert das HLRS seine langjährige Expertise in numerischen Simulationen um Kompetenzen in künstlicher Intelligenz und erschließt damit zusätzliche Anwendungsmöglichkeiten. So koordiniert das HLRS Deutschlands erste AI Factory HammerHAI, die von der EuroHPC Joint Undertaking (EuroHPC JU) gefördert wird. Aufbauend auf unserer Expertise mit HPC für Industrie, Ingenieurwesen und globale Herausforderungen wird HammerHAI einen KI-optimierten EuroHPC JU-Supercomputer am HLRS installieren und Lösungen, Ressourcen, Services sowie Schulungsmöglichkeiten anbieten, die die Hürden für Industrie und Forschung bei der Nutzung von KI abbauen werden. Wir freuen uns sehr, dass wir in HammerHAI mit Partnern des Leibniz-Rechenzentrums, des Karlsruher Instituts für Technologie, der Gesellschaft für wissenschaftliche Datenverarbeitung Göttingen und SICOS BW zusammenarbeiten und dass unsere Partner des Gauss Centre for Supercomputing am Jülich Supercomputing Centre den Zuschlag für die zweite deutsche AI Factory JAIF erhalten haben. Gemeinsam werden wir die Infrastruktur und das Fachwissen schaffen, damit Deutschland und Europa in dieser sich rasch entwickelnden technologischen Landschaft global wettbewerbsfähig bleiben.

Ein Sonderteil dieses Jahresberichts widmet sich der Frage, wie sich Höchstleistungsrechenzentren seit dem Aufkommen der künstlichen Intelligenz weiterentwickeln. Über HammerHAI hinaus geht es u.a. um das Stuttgarter KI-Beratungsunternehmen Seedbox.ai, das bereits vor der offiziellen Einweihung von Hunter begann, große Sprachmodelle auf dem System zu trainieren. Als Teil unserer Partnerschaft mit Seedbox.ai entwickeln wir gemeinsam ein umfassendes Portfolio an Wissen und Services, um die Entwicklung und Skalierung von KI-Anwendungen zu unterstützen. Außerdem bietet die Zusammenarbeit ein strategisches Modell zur with HLRS's longtime partner Guy Lonsdale of scapos AG. An enthusiastic advocate for HPC in small and medium-sized enterprises (SMEs), Guy explains how the latest iteration of the successful Fortissimo projects, called FFplus, is supporting the development of AI applications in European SMEs and startups.

Since its founding, supporting the use of HPC in industry has been one of HLRS's key missions, and with more than 70 industrial clients we continue to provide critical infrastructure for companies large and small. This year we were particularly pleased to sign a formal partnership with TRUMPF, a global leader in high-precision laser technologies for manufacturing. In collaboration with the Media Solution Center Baden-Württemberg, MACK Research (a spinoff of the company that runs the Europa Park amusement park) also turned to HLRS for its expertise and technologies for visualization and simulation. In this Annual Report you can read about Dresden-based SME hydrograv GmbH, which has been using HLRS's supercomputer for computational fluid dynamics simulations of water treatment facilities. Hydrograv's approach is a perfect example of how HPC can improve sustainability — both environmental and economic — in cities and communities.

As a publicly funded HPC center, HLRS intensified its efforts to explore how public sector agencies could benefit from new applications of HPC. In June we organized a conference in Berlin that presented the findings of our project CIRCE, focusing on how simulation could help prevent and manage crisis situations like floods or wildfires. In 2024 we also began a collaboration with the district of Ludwigsburg in which a digital twin of our region's beloved wine-producing landscape along the Neckar River will support planning and conservation efforts. We were also pleased with the enthusiastic turnout for our "Science Goes Society" event in Sersheim, in which community representatives from across the region discussed strategies for addressing climate change at the local level.

With all of these exciting new developments, one also can't overlook HLRS's primary role of providing state-of-the-art computational infrastructure for cutting

Förderung von KI-basierten Innovationen im Einklang mit den Zielen und Anforderungen des Europäischen KI-Gesetzes. Auf diesen Seiten finden Sie ein Interview mit einem langjährigen Partner des HLRS, Dr. Guy Lonsdale von der scapos AG. Guy Lonsdale befürwortet HPC in kleinen und mittleren Unternehmen (KMU) und erklärt, wie FFplus, die aktuelle Förderphase der Fortissimo-Reihe, die Entwicklung von KI-Anwendungen in europäischen KMU und Start-ups unterstützt.

Seit seiner Gründung unterstützt das HLRS die Industrie mit HPC-Leistung. Mit mehr als 70 Industriekunden stellen wir weiterhin Infrastrukturen für Unternehmen bereit. Ein Highlight in diesem Jahr war die Unterzeichnung einer formellen Partnerschaft mit TRUMPF, einem weltweit führenden Unternehmen für hochpräzise Lasertechnologien in der Fertigung. In Zusammenarbeit mit dem Media Solution Center Baden-Württemberg wandte sich auch MACK Research (eine Ausgründung des Betreibers des Europa-Parks) an das HLRS, um seine Expertise und Technologien für Visualisierung und Simulation zu nutzen. In diesem Bericht können Sie auch erfahren, wie das Dresdner KMU hydrograv GmbH mithilfe des HLRS strömungstechnische Simulationen für die Wasserwirtschaft umsetzt. Hydrograv ist ein gutes Beispiel dafür, wie HPC die Nachhaltigkeit in Städten und Gemeinden verbessern kann.

Als öffentlich gefördertes Höchstleistungsrechenzentrum evaluiert das HLRS verstärkt, an welchen Stellen öffentliche Einrichtungen von HPC-Anwendungen profitieren können. Im Juni organisierten wir eine Konferenz in Berlin, auf der wir die Ergebnisse unseres Projekts CIRCE vorstellten. Dieses befasste sich damit, wie Simulationen Naturkatastrophen wie Überschwemmungen oder Waldbrände verhindern oder bewältigen können. Im Jahr 2024 begann außerdem unsere Zusammenarbeit mit dem Landkreis Ludwigsburg, bei der ein digitaler Zwilling der Neckarterrassen die Planungs- und Naturschutzbemühungen unterstützen wird. Erfreulich war auch die rege Teilnahme an unserer Veranstaltung „Science Goes Society" in Sersheim, bei der Vertreter:innen aus Kommunen der Region Strategien für den Umgang mit dem Klimawandel auf lokaler Ebene diskutierten.

edge scientific research. As the publications list in this Annual Report shows, we continue to support impactful research across a wide range of scientific disciplines. We also report in more depth on two very different applications of HPC: engineering a novel helicopter being developed by Airbus, and identifying potential drug targets for personalized cancer therapies.

At a time when provincial nationalism seems to be on the rise globally, HLRS remains committed to international cooperation. We were glad to join the new HANAMI project, an EU initiative to promote scientific collaboration between Europe and Japan, and signed new memoranda of understanding with Tallinn University of Technology (Estonia) and Kumoh National Institute of Technology (South Korea).

As I write this in the spring of 2025, the HLRS building is currently surrounded in scaffolding in preparation for the installation of solar panels on our roof. Excavation has also begun in preparation for the construction of HLRS III, a new building and power plant for our next-generation supercomputer, Herder. Although we are in for a sometimes noisy and dusty year, we are excited that HLRS continues to thrive and grow, enabling numerous valuable activities in science, industry, and in the public sector.

On behalf of all of HLRS's dedicated staff, I thank our funders, partners, and user community for your continued support, creativity, and collaboration. As the field of high-performance computing continues to evolve, we remain grateful for the chance to share in this fascinating journey with you.

With best regards,

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

Bei all diesen Neuigkeiten geht der zentrale Auftrag des HLRS dennoch nicht unter, nämlich die Bereitstellung einer hochmodernen Recheninfrastruktur für die Spitzenforschung. Wie die Liste wissenschaftlicher Veröffentlichungen in diesem Jahresbericht zeigt, unterstützen wir weiterhin Forschung in einem breiten Spektrum von wissenschaftlichen Disziplinen. Außerdem berichten wir ausführlicher über zwei sehr unterschiedliche Anwendungen von HPC: die Entwicklung eines neuartigen Hubschraubers von Airbus und die Identifizierung von Wirkstoffzielen für personalisierte Krebstherapien.

In einer Zeit, in der provinzieller Nationalismus weltweit auf dem Vormarsch zu sein scheint, engagiert sich das HLRS weiterhin für die internationale Zusammenarbeit. Wir freuen uns über die Teilnahme am HANAMI-Projekt, einer EU-Initiative zur Förderung der Zusammenarbeit zwischen Europa und Japan, und haben neue Absichtserklärungen mit der Technischen Universität Tallinn (Estland) sowie dem Kumoh National Institute of Technology (Südkorea) unterzeichnet.

Während ich diese Zeilen im Frühjahr 2025 schreibe, ist das HLRS-Gebäude von einem Gerüst umgeben, um die Installation von Solarzellen auf dem Dach vorzubereiten. Außerdem haben die Aushubarbeiten für den Bau von HLRS III begonnen, einem neuen Gebäude und Zuhause für unseren kommenden Supercomputer Herder. Obwohl uns ein lautes und staubiges Jahr bevorsteht, freuen wir uns, dass das HLRS weiterwächst und wertvolle Aktivitäten in Wissenschaft, Industrie und im öffentlichen Sektor ermöglicht.

Im Namen des HLRS danke ich unseren Geldgebern, Partnern und Nutzern für ihre kontinuierliche Unterstützung. Das Höchstleistungsrechnen entwickelt sich ständig weiter, und wir sind dankbar dafür, diesen Weg gemeinsam zu gehen.

Mit freundlichen Grüßen

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

Seizing the AI Future

HLRS Builds Capabilities for Artificial Intelligence

With new computing systems and partnerships, and as coordinator of Germany's first "AI Factory," HLRS is expnding beyond its traditional focus on simulation to establish high-performance AI capabilities that will benefit industry and scientific research.

For decades, the High-Performance Computing Center Stuttgart (HLRS) has maintained one of Europe's most powerful infrastructures for simulation using high-performance computing (HPC). In fields as diverse as aerodynamics, combustion, molecular dynamics, and climate and weather, HLRS's supercomputers have made it possible to run large-scale, high-resolution simulations using numerical methods that replicate the natural world with high precision. These methods remain essential for both scientific research and industrial engineering, and supporting HPC applications continues to be a core part of HLRS's mission. At the same time, however, HLRS and its partners across the European high-performance computing community now find themselves in a rapidly changing, transitional period that promises to expand upon what a modern HPC center should be.

One major driver of this transition is the rise of artificial intelligence. Since the company OpenAI released the ChatGPT large language model into the world in November 2022, global competition to harness the power of AI has dramatically accelerated. Although it remains to be seen exactly what the ultimate effects of AI will be, it is clear that the field has reached a new level of maturity and across the scientific and business communities, a race is on to take advantage of the new opportunities that are now available. Major investments in infrastructure for training neural networks in the United States, China, and elsewhere are signs of the potential of AI, and many predict that it will have dramatic impacts across all levels of society.

AI grows at HLRS

Developing advanced AI applications requires suitable computing systems, and HLRS has been operating AI-capable infrastructure and gaining expertise in artificial intelligence over the past several years. In 2020, it deployed a Cray CS-Storm system designed to handle machine learning workloads for big data. This was followed in 2021 with the addition of a GPU extension to the center's flagship Hawk supercomputer, which expanded its ability to support artificial intelligence, high-performance data analytics, and converged computing approaches combining simulation and AI. Early adopters began testing applications of AI on the new systems, developing models for predicting physical properties in metal alloys, using AI to accelerate computationally expensive fluid dynamics simulations, and even producing AI-generated artworks.

In 2024, HLRS took several necessary steps towards a future in which artificial intelligence will become a more integral part of its portfolio. At a fundamental level, the arrival of its new supercomputer Hunter signals a shift in computing architecture that will improve the center's ability to support artificial intelligence applications (see page 29). In addition, the center is the coordinator of a new consortium called HammerHAI, which was selected by the EuroHPC Joint Undertaking (EuroHPC JU) to create Germany's first "AI Factory." As part of this effort, HLRS will host a new EuroHPC JU AI-optimized supercomputer beginning in early 2026 (see page 11). In 2024, HLRS also signed partnerships with small and medium-sized companies (SMEs) in

southwestern Germany to use its AI systems. Among its earliest users is the Stuttgart-based AI consultancy Seedbox.ai Ventures, which began training large language models on Hunter even before it officially went into service (see page 14). And finally, HLRS is coordinating the latest version of the successful Fortissimo project series, now called FFplus, which expanded in 2024 to fund Innovation Studies that will enable SMEs to scale up AI applications on supercomputers to achieve their business goals (see page 16). Through these and other developments, AI arrived at HLRS in 2024 in a big way.

Opening the door for new user communities

The recent installation of Hunter marked the beginning of a crucial technical shift for HLRS. By transitioning away from a primarily CPU-based system to a platform that uses AMD MI300A accelerated processing units (APUs) — which combine CPUs and GPU accelerators on a single chip — the center has laid a foundation that will make it possible both to increase the performance of traditional simulation codes and to better support AI-related applications like training large language models and running high-performance data analytics. Doing so also offers a state-of-the-art infrastructure for running new kinds of converged computing workflows that combine simulation and AI. Early research has demonstrated the potential of AI to accelerate traditional simulation workflows by taking over time-consuming components of algorithms. At the same time, simulation could be used to generate massive synthetic datasets needed to train AI models. Eventually, this could lead to surrogate models of traditional simulation methods that achieve comparable accuracy much more quickly.

It is not just traditional users of HPC who are bound to benefit from these capabilities. Importantly, Hunter's architecture potentially opens the door to supporting new kinds of user communities. In the past, HLRS has focused primarily on computational engineering and the applied sciences. As AI becomes more widespread in industry and the public sector, however, the center's infrastructure for artificial intelligence could offer an attractive option for companies and organizations in other fields — from finance, to law, to city planning and management, for instance — interested in scaling up their AI and data science applications.

In the following special section, we highlight recent developments at HLRS that demonstrate how it is evolving to meet the needs of a changing high-performance computing landscape.

HLRS Will Coordinate HammerHAI, Germany's First "AI Factory"

As part of a European Union initiative to increase artificial intelligence capabilities and usage, HLRS will operate a new AI-optimized supercomputer and develop solutions to make AI more accessible and secure for industry and science.

Many scientists and companies across Europe face hurdles that limit their ability to adopt artificial intelligence technologies. These include the prohibitive cost of installing AI computing hardware, the shortage of AI expertise, and data security concerns arising from Europe's current reliance on offshore cloud AI service providers.

In 2024, the EuroHPC Joint Undertaking announced a major initiative to address these challenges. It will fund the creation of new "AI Factories," a robust and interconnected network of AI infrastructure and expertise across Europe.

Among the first seven AI Factories to be announced in December 2024 was HammerHAI (Hybrid and Advanced Machine Learning Platform for Manufacturing, Engineering, and Research), a consortium coordinated by the High-Performance Computing Center Stuttgart (HLRS) in partnership with the Leibniz Supercomputing Center, the Gesellschaft für wissenschaftliche Datenverarbeitung Göttingen, the Karlsruhe Institute of Technology, and high-performance computing (HPC)/AI consultancy SICOS BW.

HammerHAI (also the German word for hammerhead shark) was the sole German AI factory named in the initial EuroHPC JU announcement. An additional six AI Factories were announced in March 2025, including the JUPITER AI Factory (JAIF) led by the Forschungszentrum Jülich.

AI infrastructure for industry and research

The primary objective of HammerHAI is to lower the barriers for start-ups, SMEs, large enterprises, and scientific institutions that need powerful computing capabilities to create and deploy AI-driven solutions. By offering state-of-the-art infrastructure and support, HammerHAI will speed up innovation, help train an AI-ready workforce, and foster a robust and secure AI ecosystem in Germany and across Europe.

During its first year, HammerHAI will procure a new, large-scale AI-optimized supercomputer, which HLRS will begin operating in Stuttgart in 2026 on behalf of the EuroHPC JU and the project funders. Based on accelerated compute nodes, the system will be tailored to the computational requirements of typical AI workloads in industry, including the training and usage of customized large language models, deep learning, and complex data analytics.

Within Germany's HPC landscape, the Gauss Centre for Supercomputing (GCS) and its member centers (HLRS, LRZ, and the Jülich Supercomputing Centre) follow a coordinated approach to support Europe's AI needs. In this context, HammerHAI will complement JAIF by building on HLRS's established focus on industry, engineering, and global challenges. HammerHAI's system and services will be particularly tailored to the needs of manufacturing, automotive and mobility, start-ups and SMEs across all sectors, as well as academia. (JAIF's focus is on healthcare, energy, climate change

Members of the HammerHAI consortium met at HLRS in early 2025 to plan the start of the project.

and the environment, education, media and culture, the public sector, finance, and insurance.)

"The European and German governments have identified the creation of a robust, native European infrastructure for artificial intelligence as a high priority. HammerHAI will quickly help to address this need," said Prof. Dr. Michael Resch, Director of HLRS. "At the same time, the project will play an important role in implementing the development strategy regarding AI at HLRS and the Gauss Centre for Supercomputing as we evolve to support the changing landscape of high-performance computing."

Breaking down barriers to AI uptake

Using "cloud-like" features, the HammerHAI system will make it easy for users to migrate and run existing applications on its large-scale, AI-optimized system on an as-needed basis. In addition, HammerHAI's service portfolio will help users integrate AI solutions at all stages of the application life cycle. It will offer a "concierge" service that provides end-to-end personal guidance through HammerHAI's offerings, from first inquiry to solution deployment. HammerHAI's support staff will be able to help identify the most relevant AI applications, provide technical support, estimate costs, and optimize scaling. Moreover, the AI Factory will provide pre-trained models and container templates that empower users to start quickly, as well as collaboration tools that allow teams to share datasets, code, and prebuilt models while maintaining strict data governance.

To expand the workforce of AI experts, HammerHAI will also offer opportunities for skills development, including workshops and training courses, hands-on labs enabling real-world experimentation, and university and vocational programs.

Notably, the secure, Germany-based platform and data repository created by HammerHAI will follow principles included in the European Union's AI Act of 2023. It will place strict controls on access to the resulting AI models, the data on which they are based, and the query interfaces through which they are utilized. This will make it easier for European researchers, startups, and SMEs to adhere to compliance requirements concerning data security and the ethical use of AI.

With a budget of approximately 85 million Euros, HammerHAI is co-funded by the EuroHPC JU, the German Federal Ministry of Research, Technology and Space, the Baden-Württemberg Ministry of Science, Research, and Art, the Bavarian Ministry of Science and Art, and the Lower Saxony Ministry for Science and Culture.

Applications of AI in Engineering and Manufacturing

Recent research at HLRS has been investigating how artificial intelligence could deliver answers to complex engineering problems more quickly.

Reducing the Computational Cost of CFD Simulations

Researchers at the University of Stuttgart's Institute of Aerodynamics and Gas Dynamics have been exploring how AI could help make high-resolution fluid simulations more accessible. A framework called Relexi combines FLEXI, a computational fluid dynamics solver, and Tensorflow, one of the most popular machine learning packages. Simulation data from FLEXI are fed into a reinforcement learning algorithm, becoming training data for parameter optimization in a turbulence model. The turbulence model then predicts the eddy viscosity as input data for the larger FLEXI simulation, which generates data for another round of training. After multiple iterations, the optimization of the turbulence model converges to a point where the simulation remains stable and accurate.

Simplifying Prediction of Metal Properties for Sheet Metal Forming

Sheet metal forming processes for manufacturing products such as automobile doors require a thorough understanding of the material properties of the metals being used. Typically, this involves taking measurements of test materials placed under various stress conditions and then performing computationally expensive simulations. To make this more feasible for engineers who don't have access to supercomputers, researchers at the University of Stuttgart's Institute for Metal Forming Technology, together with HLRS, ran two billion simulations of material properties on HLRS's Hawk supercomputer. They then used the results to train a neural network capable of quickly and accurately predicting parameters that are essential for metal forming processes, based on experimental measurements of any material of interest. The resulting application could potentially be run on a typical desktop computer.

Designing Safer Automation

Industrial controls manufacturer Festo is using HLRS's high-performance computing resources to design safer, more efficient human-robot collaborations in manufacturing processes. Using a mixture of input data from the Festo R&D lab as well as video and sensor data of robots working in real-world manufacturing environments, the researchers train a reinforcement learning algorithm to replicate good behaviors by providing positive feedback. The resulting models can achieve roughly 80 percent of the necessary performance in a robot, saving the automation designers time. By tailoring these solutions to specific scenarios, Festo can help clients develop more complex automation workflows that enable safe collaboration between humans and robots.

AI for Small and Midsize Companies, "Made in Germany"

Stuttgart-based AI agency Seedbox.ai is using HLRS's supercomputer to accelerate the uptake of artificial intelligence in European industry.

In August 2024, HLRS launched a new partnership with Seedbox.ai, a local startup that develops custom AI solutions for German small and medium-sized enterprises (SMEs). Under the agreement, Seedbox.ai has begun using HLRS's supercomputing systems and is working with the center to plan a service package that supports companies in using supercomputing infrastructures for artificial intelligence.

"In many companies today, artificial intelligence has already become synonymous with innovation," said Kai Kölsch, co-founder and Managing Director of Seedbox.ai. "Using generative AI, we can achieve significant efficiencies in core business activities and tap into additional, sustainable sources of revenue by developing new business models. With the help of HLRS and access to its supercomputing infrastructure, we want to support companies across Germany in taking their innovation projects to the next level."

One focus at Seedbox.ai is to accelerate the adoption of generative artificial intelligence, a technology that has been exploding since the arrival of large language models (LLMs) like ChatGPT. In industry, LLMs can be used, for example, to create co-pilots and chatbots, for document processing and management, to analyze complex data sources, or to create editorial content. AI tools could also significantly reduce the time spent on structured, data-driven, or repetitive tasks, enabling major increases in efficiency at a time when many companies are suffering from staff shortages. As the German workforce ages and companies lose decades-old domain expertise to retirement, AI models could also potentially capture knowledge and serve as expert systems for younger professionals. The integration of AI will become even more relevant when used in existing and new systems, such as interfaces for customer relationship management (CRM) or enterprise resource planning (ERP), knowledge databases, or as an integral component of new, innovative software solutions.

Despite the rapid growth of interest in generative AI among German SMEs, more speed is needed in the implementation of practical applications. Doing so will require solutions that reliably tackle challenges such as data security, data sovereignty, performance, and cost efficiency. The collaboration between Seedbox.ai and HLRS aims to solve these challenges, enabling German industry to utilize the full potential of large language models.

Hunter powers creation of European large language models

Two additional factors that currently limit the advance of AI in Europe are that the most widely used LLMs have been trained using English-language texts and that they are operated by companies based outside the European Union. In such a linguistically diverse body as the EU, developing LLMs that can function across national borders will be important. The ability to use local infrastructure to store and process local, critical proprietary data in a secure way is also essential from both business and strategic points of view.

Even before HLRS's new Hunter supercomputer (see page 29) went into service, Seedbox.ai was the first company to demonstrate its capabilities for AI applications, training new large language models in 24 European languages. Seedbox.ai is making these models available on an open-source basis to other AI application developers.[1] This will promote AI-based innovation — for example, by accelerating the implementation of multilingual autonomous agents in everyday working life.

Although commercial vendors could also handle such a task, using AI supercomputing resources at a German, publicly funded HPC center like HLRS offers Seedbox.ai and its clients major advantages. Dennis Dickmann, who has worked intensively in the field of artificial intelligence since 2008 and is responsible for technology development at Seedbox.ai as its co-founder and CTO, welcomed the start of its partnership with HLRS, saying, "Such a powerful, secure and regional supercomputing infrastructure is a unique and extremely valuable resource. By cooperating with HLRS, we can guarantee the security of our clients' data while providing them with the AI technologies they need to remain innovative and competitive."

AI at the right scale

The agreement with Seedbox.ai marks the first time that HLRS has entered into a partnership with an AI startup. As HLRS Managing Director Dr. Bastian Koller explained, it is an important step in the high-performance computing center's strategy: "With Seedbox.ai as a strategic partner, we can not only already show what is possible with our current system but also what will be possible with our future systems. This way, we will lose no time in the race to develop the best possible uses of AI for science and industry."

At a time when large, international IT companies have recently announced enormous investments in AI infrastructure, the partnership with Seedbox.ai also demonstrates that oversized systems aren't necessary for many applications that can have significant, real-world impact for small and mid-sized businesses. "Large-scale computing power is certainly important for applications of artificial intelligence like training large-language models from scratch," remarked Dennis Hoppe, who leads the HLRS Department of Converged Computing. "However, what we are finding is that for many typical applications in science and industry, you can still get very interesting results with fewer compute resources. Many organizations are focused on fine-tuning existing, pre-trained models, a process that requires a much shorter training time and, thus, requires only modest resources. In the future, HLRS expects to open many more new opportunities for our growing user communities."

1 Information about Seedbox.ai's multilingual large language models is available at www.kafkalm.com.

Working together to bring AI to German SMEs (l-r): Kai Kölsch (Founder and CEO, Seedbox.ai), Dennis Dickmann (Founder and CTO, Seedbox.ai), Dennis Hoppe (Head, Converged Computing, HLRS), Bastian Koller (Managing Director, HLRS)

The Changing Face of Fortissimo: An Interview with Guy Lonsdale

Over more than 10 years the Fortissimo projects have supported hundreds of small- and medium-sized enterprises (SMEs) across Europe in conducting business experiments exploring how high-performance computing (HPC) can be used to solve business challenges. The High-Performance Computing Center Stuttgart (HLRS) has been a partner in the projects since their beginning, and has been lead coordinator for the most recent two iterations of the project, called FF4EuroHPC and FFplus.

Another key figure in the Fortissimo story is Guy Lonsdale, CEO of scapos AG, a company that specializes in the marketing of software products. Originally trained in numerical mathematics, he has led the open calls for proposals to all of the Fortissimo projects, and is an evangelist for the opportunities that high-performance computing can offer to SMEs.

The launch of FFplus in 2024 marked a new chapter in the history of Fortissimo, as it invited SMEs for the first time to propose "Innovation Studies" focused on applications artificial intelligence, with an emphasis on generative AI. The project's first two open calls received applications not only from traditional users of high-performance computing in the engineering and manufacturing industries, but also from a much wider range of industries, including finance, law, medicine, and journalism. In the following interview we talked with Lonsdale about the evolution and impact of the Fortissimo projects, as well as how artificial intelligence is opening the door to new user communities in high-performance computing centers.

Looking back to the beginning, how did Fortissimo start and what was it trying to accomplish?

Shortly after the financial crisis of 2008, the European Commission identified the need for a funding program to support European industry, and the manufacturing industry in particular. This led to the start of the Factories of the Future project. At that time they also realized that while large companies in the automotive and aerospace industries, for example, routinely use high-performance computing, SMEs face unique challenges and typically do not. They decided to create a program that

Under the FF4EuroHPC project, a team in Montenegro developed an artificial intelligence tool for tracking the health of chickens on poultry farms. They demonstrated that the solution can reduce farm operating costs and chicken mortality rates.

reaches out specifically to manufacturing SMEs, which became the focus of the first two Fortissimo projects.

The Fortissimo concept was to support experiments in which SMEs could use HPC address a business challenge. Although the concept has a research aspect, it's absolutely not about doing research, but rather demonstrating the advantages of using HPC in a commercial context. SMEs need access to support, and so FFplus gives them what is called cascading funding. The Joint Undertaking gives a grant to FFplus, and we run competitive open calls and select the best proposals using external experts. The winning companies receive funding and support, although FFplus doesn't provide computing resources. The expectation is that the SMEs apply for computing time on large-scale systems through the JU's access scheme.

In the end, the projects have an impact for the participants. At the same time the resulting success stories give Fortissimo ways to reach out to the wider ecosystem of potential HPC users. As I always say to small companies, if a similar SME used HPC successfully, it could be useful for you too.

Over the history of Fortissimo, have you noticed any changes in how SMEs use high-performance computing?

In the beginning of Fortissimo, the primary use of high-performance computing in manufacturing was for modelling and simulation — areas like computational fluid dynamics, structural mechanics, and molecular modelling. In the project that immediately preceded FFplus, called FF4EuroHPC, we still saw a lot of modelling and simulation. At the same time, however, we also received applications for projects involving machine learning, artificial intelligence, and big data analytics. FF4EuroHPC still had an emphasis on manufacturing, but we included a special priority in the second call to say we're really reaching out to SMEs doing manufacturing and we saw a balance between the two.

With FFplus two important things have changed. One is that we now have two very different types of sub-projects. The first type is classical Fortissimo, where SMEs that are new to HPC perform an experiment to show that they can use it to solve a business challenge. The other type focuses on technology development for artificial intelligence, specifically for generative AI, including the development or tuning of foundation models or large language models. Here, we're looking at SMEs that are already proficient and have existing AI-driven business models, but they need access to large-scale European HPC resources to achieve the targets of their LLMs. In the first FFplus open call in 2024, we saw that the SMEs that are new to HPC are also mainly focused on AI. The majority of applications were looking at AI alone or in combination with digital twins, which are driven by AI, combining machine learning and simulation.

Has the emphasis on artificial intelligence in FFplus led to changes in the kinds of companies that apply to participate?

The exciting thing about artificial intelligence is that it enables applications that are relevant for many sectors in industry and business. In our first open call we received applications from a wide range of fields and from a high number of European countries. For the Innovation Studies we targeted SMEs that have an existing business model linked to AI use or deployment. Together it makes for a colorful mixture. We have companies, for example, that are developing medical applications, and a number of companies from the worlds of finance and environment applied. We also have projects from the fields of information and communication technology and the media industry.

Some companies have AI know-how but haven't used HPC before, and recognize that access to large-scale computing resources would enable them to do more. Others are sticking their toe in the water with AI. They've heard that it could have an impact and know what they want to do. They get support partners who look at what they're doing so far and suggest potential next steps. For example, this could mean adding machine learning to their modelling and simulation approach to get faster or broader results.

You mentioned the importance of success stories for demonstrating to other companies what can be done using HPC. It seems like an important part of all of the Fortissimo projects has been communication and outreach.

Yes, absolutely. The National Competence Centers (NCCs) within the CASTIEL 2 and EuroCC 2 projects — projects that HLRS has been coordinating — have also contributed to this. In 33 European countries, each NCC is typically based at a high-performance computing center and connects HPC users with resources that can help them. The NCCs have a mission to reach out to SMEs in their home countries and are helping to generate a lot of the proposals for FFplus. Using the printed collections of success stories that have come out of the Fortissimo projects or videos that we have produced, they go to companies and say, "Hey, look at this. It is relevant for your business." The success stories focus on things like speeding up a design or business process, or reducing costs, or employing more people, or growing turnover. These are the business arguments for using HPC that are critical for SMEs.

Considering that the community surrounding Fortissimo is expanding due to the rise of AI and data-driven methods, is there anything significantly different about users' needs now that they might not have had in the days when they were focused on traditional simulation?

When SMEs use HPC to do CFD simulations, for example, they need support them in moving to a large cluster or high-performance computer. Their questions include things like: How do I use this remote system? If I am running software remotely, how do I manage my jobs? What are the relevant licensing issues for using commercial software in this way? In a sense this is mostly IT system management more than HPC. In addition, SMEs that haven't done simulation in the past need someone with HPC experience, someone who understands the software that could be used, to understand their design challenge and take them by the hand. This also meant that after an SME's Fortissimo business experiment was over it often needed continued support.

"Success stories give Fortissimo ways to reach out to the wider ecosystem of potential HPC users. As I always say to small companies, if a similar SME used HPC successfully, it could be useful for you too."

What we're now seeing is that many SMEs don't need much support. Many AI companies know exactly what they want to do and just need access to larger systems and help in finding the necessary resources. This is partly what FFplus does. We put them in touch with the right people to provide the necessary advice and support. And then in many cases the challenge is accessing the data that are necessary to train the software.

What would you say is the key achievement of Fortissimo?

I do think that at this point we're very well known, not just at the level of the European Commission and in the EuroHPC Joint Undertaking, but also in the user community. The NCCs were very keen to get SMEs in their countries involved in the program because it's something that provides a lot of value. The SMEs that are applying are also clear about the potential benefits. The result is that we now have a large collection of success stories that explain the impact for a wide range of applications. And they are no longer just focused on doing crash simulations or aerodynamics, or classic big science like weather forecasting. To take just one example, we had a project in Montenegro that uses artificial intelligence on an HPC system to monitor the health of the animals on a poultry farm. Showing the impact and encouraging people to explore applications in new areas like this is probably our biggest success.

How do you anticipate Fortissimo growing or evolving in the next couple of years?

The winning sub-projects from the first FFplus open call just started, and we have two tranches of calls coming up, for a total of four more sub-calls. The current expectation is that the approach will not significantly change, but at the same time the rise of artificial intelligence means that high-performance computing is evolving very quickly. At the recent ETP4HPC conference there was a presentation by somebody from Argonne National Laboratory explaining how large language models could be used as a research assistant in the future. The research would shift from doing analysis to using these things to provide data analytics. He also explained the capabilities of current LLMs for solving complex technical questions and, frankly, it was shocking. They are using LLMs to solve problems for which only a PhD student or a postdoc with expertise in that particular field could have matched their performance. On the one hand, it's easy to think that we've got another three years of FFplus before we have to think about what's next. When I hear stories like these, however, three years seems like a long time. A lot could change.

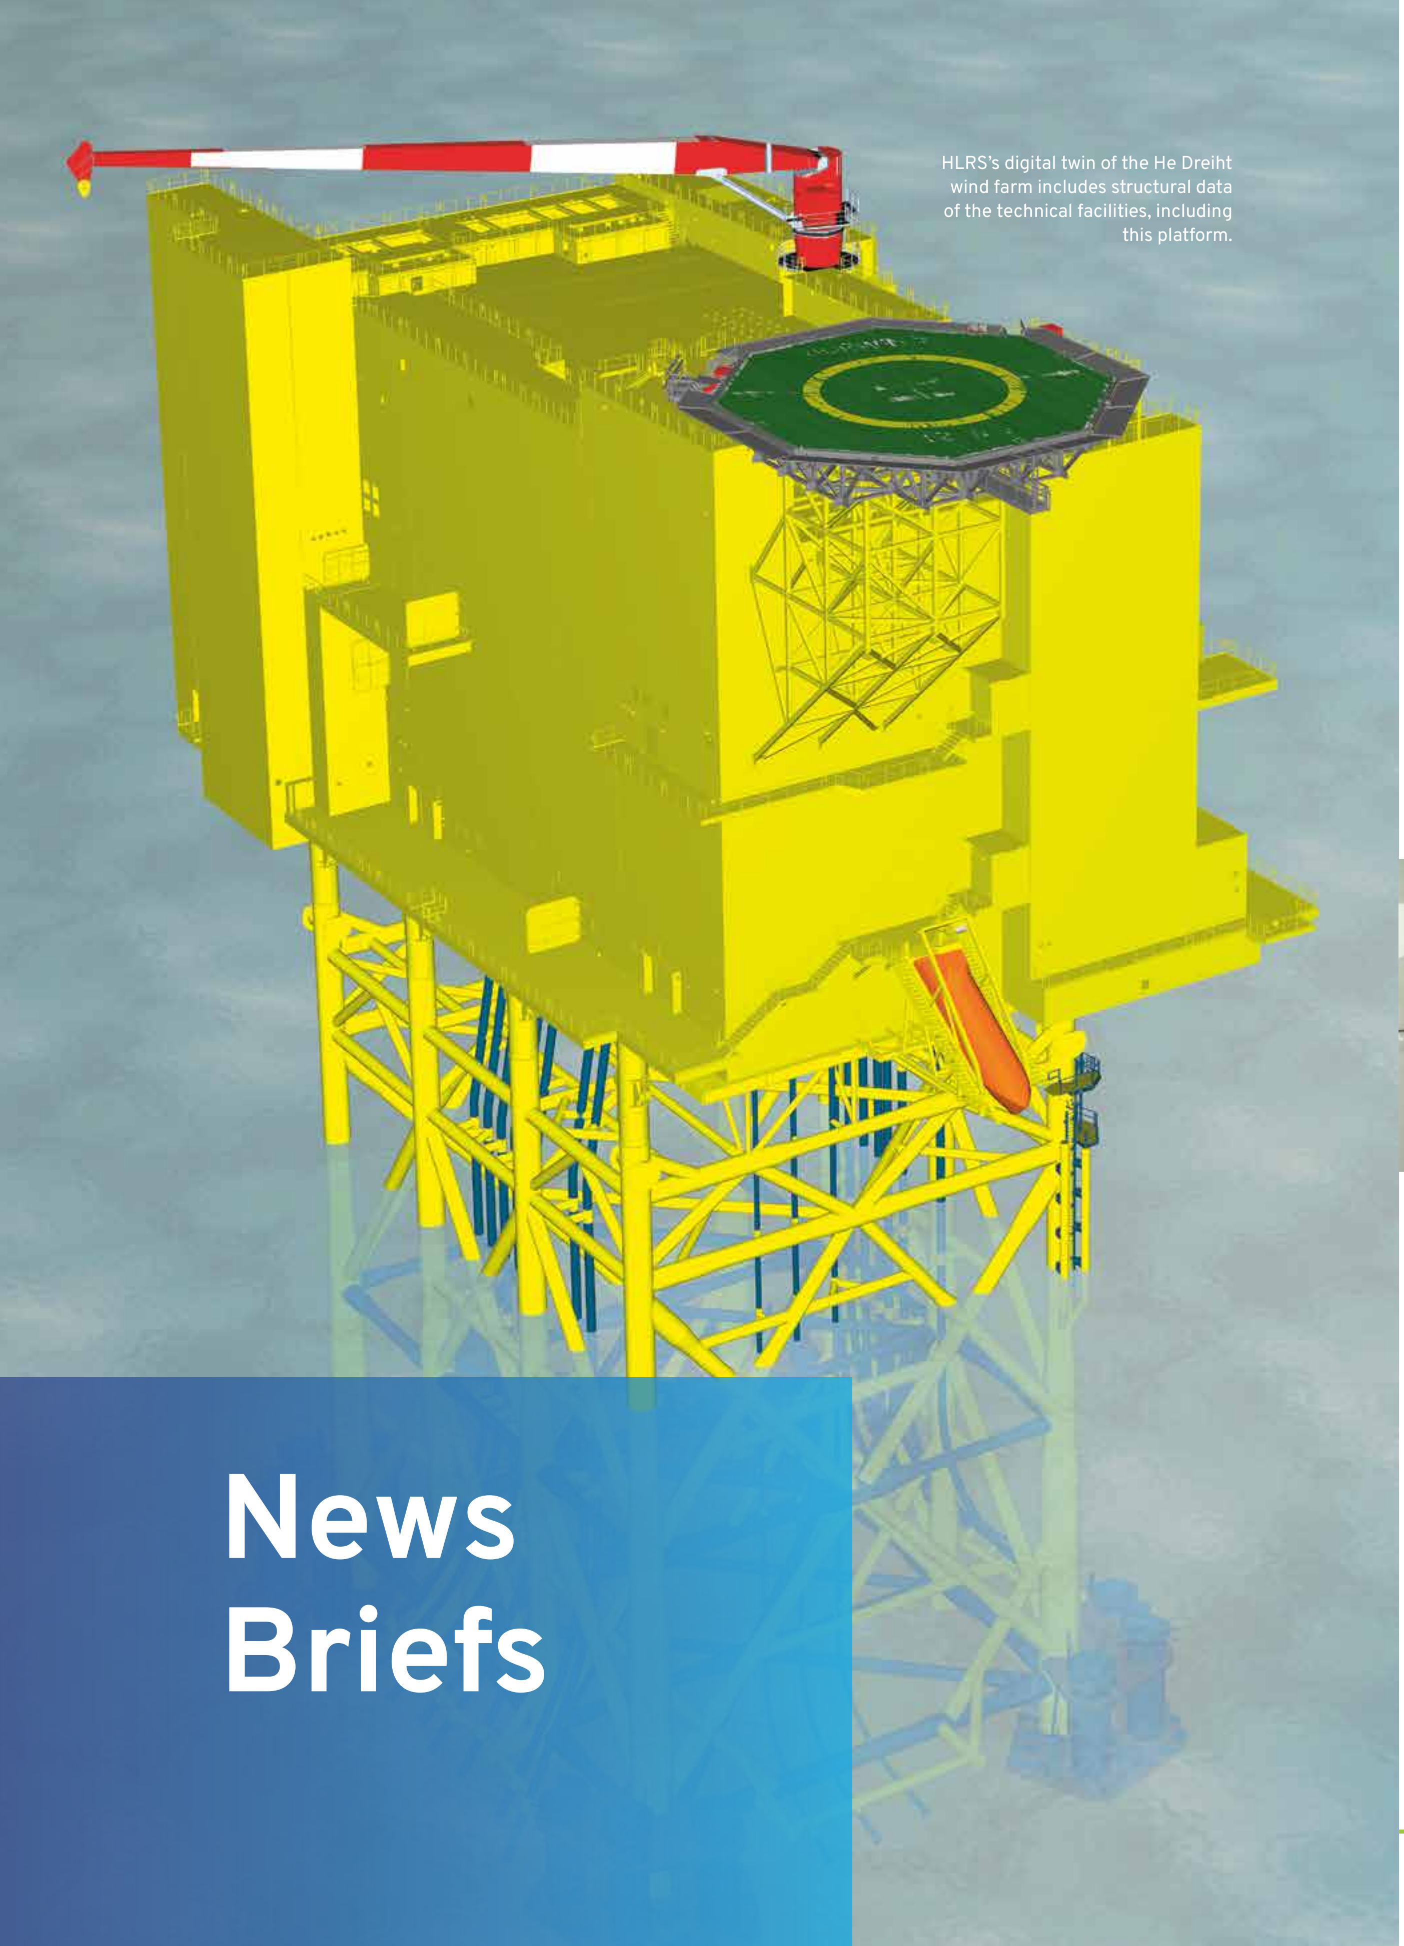

HLRS's digital twin of the He Dreiht wind farm includes structural data of the technical facilities, including this platform.

News Briefs

HLRS Completes TISAX Level 3 Assessment for Information Security

Since 2021 the University of Stuttgart has been registered as a participant in the Trusted Information Security Assessment Exchange (TISAX) on behalf of HLRS. In 2024, audit provider TÜV Nord CERT GmbH conducted a new evaluation of HLRS's information security management system, this time at TISAX Level 3. Governed by the ENX Association on behalf of the German Association of the Automotive Industry (VDA), the TISAX framework closely follows the international standard ISO 27001 for information security management systems, while also prescribing additional requirements for the implementation of controls that must be followed. TISAX Level 3 comprises the most comprehensive auditing approach within the information security framework audit, ensuring HLRS's adherence to information security standards that are appropriate for handling data with "very high protection needs." The result is available on the ENX Portal under the Scope-ID SP9M8V and the Assessment-ID ANCTTM. HLRS is also certified for information security under the ISO 27001 standard.

Berthold Schmidt (CTO, TRUMPF) and Anna Steiger (Chancellor, University of Stuttgart) signed the collaboration agrement at TRUMPF headquarters in Ditzingen.

New Collaboration Agreement with TRUMPF

In March 2024, industrial tools manufacturing company TRUMPF signed a collaboration agreement with HLRS to make large-scale computing capacity available for TRUMPF employees. "This cooperation demonstrates that Germany is an industrial center that also has a high-performance digital ecosystem. Using supercomputers, our developers can virtually fine tune machine functions even before the first prototype is created and train AI solutions for our production facilities much faster. This will enable us to innovate more sustainably and efficiently," said TRUMPF CTO Berthold Schmidt. Commenting on the agreement, Prof. Michael Resch, Director of HLRS, said that the center's high-performance computing systems "will support continuing improvements in technologies at TRUMPF. We are proud that our computing power will enable us to continue to support the strength and competitiveness of the Stuttgart economic region." The participants in the partnership also hope to identify new applications of HPC in industry. TRUMPF uses its own high-performance computers for simpler simulations. More complex tasks that require higher precision, however, are only possible using supercomputers like those at HLRS. One potential application is the simulation of quantum computers, which is so computationally demanding that in the future it will benefit from the acceleration offered by HLRS's supercomputers. For several years, TRUMPF has offered its customers machine tools that use artificial intelligence to make their work faster and more effective. In the future, the company will expand this range of offerings with new solutions.

Björn Klose, Simon Homes, and Tim Niklas Uhl were named winners of the 2024 HLRS Golden Spike Awards.

Winners of the 2024 Golden Spike Awards

On October 10-11, 2024, users of the high-performance computing systems at HLRS and the Scientific Computing Center (SCC) at the Karlsruhe Institute of Technology (KIT) met to present and discuss their recent research at the 27th Annual Results and Review Workshop. For the first time, this year's Workshop took place at KIT, where the SCC partners with HLRS in providing HPC resources for researchers across Germany. In 21 talks and a poster session, visitors learned about advanced studies across many scientific fields. Presentations described not only applications of traditional simulation methods using HPC, but also uses of machine learning methods and strategies for improving energy efficiency. HLRS steering committee chairman Professor Thomas Ludwig (Director, German Climate Computing Center) also announced the winners of the 2024 HLRS Golden Spike Awards, which recognize excellence in computational research and the use of HPC. Representing their respective projects, Golden Spike Awards went to Björn Klose (DLR) for "Numerical investigation of shock-boundary layer interaction on a highly-loaded transonic compressor cascade to compare with experimental results"; Simon Homes (TU Berlin) for "Molecular simulations: A thermodynamic study of bulk and interface"; and Tim Niklas Uhl (KIT) for "Scalable discrete algorithms for big data applications."

New Collaboration Agreement with Kumoh National Institute of Technology

HLRS signed a new Memorandum of Understanding for collaboration with the Kumoh National Institute of Technology (KIT). Located in Gumi, the birthplace of South Korea's electronics industry, KIT is an engineering-focused university focused on preparing professionals for careers in high-tech industries. HLRS Director Michael Resch signed the collaboration agreement together with KIT President Ho Sang Kwak during a visit to KIT in February 2024. The collaboration will promote communication and the exchange of expertise between Stuttgart and Gumi, particularly with a focus on improving access to and usage of high-performance computing in industry. "HLRS and KIT share a vision that high-performance computing can build bridges between university research and industrial development," Prof. Resch said. HLRS maintains collaboration agreements with more than a dozen leading centers for computational sciences and high-performance computing in Europe, Asia, and the Americas.

Kumoh National Institute of Technology President Ho Sang Kwak and HLRS Director Michael Resch signed the collaboration agreement during a visit to KIT.

HPC SPECTRA Promotes Skill Development across Europe

Since 2018, the EuroHPC Joint Undertaking (EuroHPC JU) has been coordinating the development of a world-class ecosystem for high-performance computing across Europe. In addition to supporting the procurement of new supercomputers, a critical component of the JU's mission is to ensure that European scientific and industrial research communities have the knowledge and skills they need to use these powerful resources effectively. A EuroHPC JU-funded project launched in 2024 called HPC SPECTRA is addressing this goal by developing tools for increasing awareness of HPC training opportunities across Europe and implementing standards that make it easier for trainees to find courses that fit their interests and needs. As coordinator of the EuroHPC JU projects CASTIEL and CASTIEL 2, HLRS has been leading efforts to develop the HPC in Europe (https://hpc-portal.eu/), which provides a one-stop online resource for finding HPC resources and expertise. As a project partner in HPC SPECTRA, the center is helping to develop a comprehensive online gateway that will be integrated into the portal, gathering announcements of all HPC training opportunities across Europe in one place. HPC SPECTRA also now manages the International HPC Summer School (IHPCSS). Started in 2010, the IHPCSS has brought together HPC experts from Europe, the United States, Canada, Australia, and Japan to focus on scientific and technical issues facing high-performance computing.

HPC SPECTRA

Analysts from Hyperion Research highlighted current trends in the use of HPC, AI, and quantum computing.

Hyperion HPC User Forum Returns to Stuttgart

Established in 1999 to promote the health of the global high-performance computing industry and address issues of common concern, the biannual HPC User Forum facilitates networking among users, vendors, government agencies, and private industry. At a two-day meeting organized by HLRS in collaboration with Hyperion Research in late October 2024, approximately 70 attendees gained insights about the latest trends in the fields of HPC, artificial intelligence, and quantum computing. An industry panel including representatives of TRUMPF, the National Center for Supercomputing Applications, Seedbox.ai, and SSC-Services offered insights into the use of HPC and AI in industry. Representatives of leading HPC centers, including the National Center for Atmospheric Research, the University of Bristol, Argonne National Laboratory, the European Centre for Medium-Range Weather Forecasts, the DiRAC High Performance Computing Facility, and the University of Innsbruck offered updates on recent developments in HPC and AI infrastructure and applications. Representatives of technology vendors including Lenovo, QuEra, KAYTUS, and Hammerspace also gave presentations describing their latest technologies and services.

A full-scale, functional replica of the historic Junkers A50 airplane was a highlight of HLRS's SC24 booth.

HLRS Takes Off at SC24

"Flying" was HLRS's motto at SC24, a gathering of the international high-performance computing community in Atlanta, Georgia. The highlight of the center's exhibit was a full-sized Junkers A50 airplane, an eye-catcher that attracted many conference attendees to the booth. This functional replica of a classic aircraft was the result of a collaboration with Junkers Aircraft and Kasaero GmbH, who used visualization systems at HLRS during optimization its design. HLRS also made multiple contributions to the SC24 conference program, including sessions focusing on the HANAMI collaboration between Europe and Japan, training and skill development within the EuroHPC ecosystem, activities in HLRS's Department of Philosophy of Computational Sciences, and research concerning the ethics of using AI in medicine. Dr. Theresa Pollinger, a recent University of Stuttgart graduate and user of HLRS's Hawk supercomputer, also discussed a novel experiment in which she simultaneously used the flagship supercomputers at HLRS, the Leibniz Supercomputing Center, and the Jülich Supercomputing Center to run a single large-scale simulation related to fusion research. As demand for compute, memory, and storage capacity in extreme-scale, grid-based simulations grows, this work demonstrated a strategy for running them at a much higher resolution than would be possible on a single supercomputer.

Modelling and Public Policy

When faced with complex questions, policy makers often turn to scientists to synthesize information, conduct analyses, and provide predictions that they can use as a basis for making difficult decisions. What happens when these two worlds meet, however, can be anything but straightforward. Differences in expectations, motivations, values, and scientific literacy between scientists and decision makers can lead to misunderstanding, while social, cultural, political, and practical considerations can influence how science-based models are created, interpreted, and used. On November 25-27, 2024 the HLRS Department of Philosophy of Computational Sciences organized and hosted an international conference titled Modeling for Policy. Bringing together philosophers, social scientists, historians of science, and users of computing technologies for modeling, the event reflected on the practice of modeling and the challenges it can face in the context of public policy making. This latest event in the Science and Art of Simulation conference series covered a wide range of domains in which models are used, including public health, energy transformation, climate modeling, employment programs, and value chain analysis, among others. The event sought to define the capabilities and limits of simulation more clearly, explore how trust in models could be improved among policy makers and the general public, and make recommendations that could lead to better interactions between simulation scientists and society at large.

HANAMI Fosters Collaboration between Europe and Japan

Launched in 2024, the HPC Alliance for Applications and Supercomputing Innovation (HANAMI) is developing a strategic collaboration between Europe and Japan to advance the use of simulation in science. This initiative, aligned with the agenda of the EuroHPC Joint Undertaking, focuses on porting existing code and testing and evaluating application performance on new computing architectures. It will also facilitate access for European researchers to advanced Japanese supercomputing resources like Fugaku and, in turn, will enable Japanese partners to access EuroHPC systems. In the framework of HANAMI European and Japanese institutions are exploring synergies in CPU architectures and strengthening partnerships in scientific software development, focusing on scientific topics in environmental science, biomedicine, and materials science that will benefit from closer collaboration. The project will also organize annual high-level symposia and develop a digital roadmap that will enable more sustainable collaborations between Europe and Japan, both in academia and industry. Led by the French Alternative Energies and Atomic Energy Commission (CEA), HANAMI involves 14 European organizations, including three German partners, working together closely with ten Japanese institutions. HLRS is involved in a work package focused on communication, dissemination, and exploitation of results, increasing HANAMI's impact by promoting its visibility and achievements.

A panel at the SC24 conference in Atlanta discussed challenges and opportunities for international collaboration, with a particular focus on the European and Japanese HPC communities.

8th iHURT Symposium

On December 3, 2024 HLRS hosted the eighth Industrial HPC User Round Table (iHURT), an annual gathering for HPC users in companies both large and small to exchange perspectives on new technologies and the unique challenges that they face. Organized and hosted by HPC consultancy SICOS BW, the event welcomed approximately 40 industry representatives. In a keynote talk, Gerd Büttner (Airbus Operation) described emerging challenges for industry, including certification of artificial intelligence applications in tightly regulated industries such as aerospace, legal issues surrounding data ownership and availability, limitations in portability of applications across different hardware platforms, the shortage of trained personnel in the computational sciences, and the environmental sustainability of large HPC systems. User presentations by Maxmilian Ehrle (MANN+HUMMEL), Pavel Apanasevich (hydrograv), and Erich Jehle-Graf (Mercedes-Benz) explored how their respective companies use HPC to develop products and services. Dr. Thomas Bönisch of HLRS surveyed the state of the technology market for high-performance computing, as well as trends that will shape the field in the coming years. In a closing open discussion, participants focused on other challenges they see for the future of the field, such as a perceived stagnation in the improvement of HPC system performance and the limitations posed by the high costs associated with commercial software.

HPC-AI Convergence at AISys 2024

The convergence of high-performance computing and artificial intelligence has begun to unlock exciting new opportunities for science, industrial R&D, and many other fields. Realizing the potential of this convergence, however, will require that HPC architectures satisfy the needs of the AI community, and that solutions are available to integrate AI applications and HPC workflows. This will only come through close collaboration and knowledge transfer between experts in HPC and AI. A special track at the First International Conference on AI-Based Systems and Services (AISyS 2024) organized by Prof. Michael Resch and Dennis Hoppe of HLRS was a step in this direction. Held in Venice, Italy, the meeting explored both how HPC can leverage AI to enhance system performance and efficiency, and how AI can harness the immense computational capabilities of HPC to tackle increasingly complex modelling and simulation challenges. In another contribution from HLRS, Rishabh Saxena presented a paper offering a comprehensive survey of cybersecurity concerns when using AI applications on HPC systems. Written in collaboration with other HLRS scientists, the work was honored with a best paper award at the conference.

New Collaboration Agreement Signed with TalTech

In September 2024, HLRS and the School of Information Technologies at the Tallinn University of Technology (TalTech) launched a new formal research collaboration. Signed in Tallinn, Estonia, a five-year memorandum of agreement (MoA) pledges to build on the organizations' complementary strengths in high-performance computing, applications of simulation for automotive research and industry, global systems science, and climate research, as well as professional training. The MoA will facilitate staff exchange programs, joint research projects, professional services and continuing education programs, technology transfer, and collaboration with industry. HLRS and TalTech also identified preliminary scientific areas for collaboration, including the development of numerical algorithms for engineering applications, the application of emerging technologies for parallel programming and networking computing to models in engineering applications, and the use of immersive visualization technologies in numerical simulation for engineering challenges. Tallinn University of Technology is the leading university for engineering and IT science and education in Estonia. The agreement was signed by Prof. Dr. Michael Resch, director of HLRS, and Dr. Erik Puura, Vice-Rector for Entrepreneurship at TalTech.

TalTech Vice-Rector for Entrepreneurship Erik Puura and HLRS Director Michael Resch signed the memorandum of agreement in Tallinn.

l-r: Prof. Dr. Bernd Eberhardt (Hochschule der Medien), Sven Meyer (Staatstheater Stuttgart), Matthias Hauser (Media Solution Center Baden-Württemberg), Alexander Bouquet (MACK One), Michael Mack (Europa-Park and MACK One).

Supercomputing for the Theme Park Industry

Originally founded as an in-house agency at the Europa-Park theme park, MACK One has evolved to become an international creator and consultant in the fields of theme park design, media-based entertainment, and media content. In July the company announced a new partnership with the Media Solution Center Baden-Württemberg (MSC), involving its members HLRS and the Hochschule der Medien (HdM). Through interdisciplinary collaboration that uses state-of-the-art digital tools for simulation, artificial intelligence, high-performance data analytics, and data visualization — including virtual and augmented reality — the new initiative, called MACK Research, will extend beyond the existing daily activities of the Mack Group to pursue innovative solutions for challenges facing the theme park industry. The initiative is intended to be exploratory, building experience and developing prototypes that could later be translated into marketable solutions. MSC General Manager Matthias Hauser sees great potential in MACK Research, remarking, "Europa-Park and MACK One are world leaders in the theme park and entertainment fields. The knowledge transfer and access to computing power that it will gain through the establishment of MACK Research will lead to exciting discoveries that could eventually benefit the entire industry."

News Highlights

Hunter Goes into Service

HLRS's newest flagship supercomputer combines combines higher performance with greater energy efficiency, offering a more powerful, state-of-the-art infrastructure for simulation, artificial intelligence, and converged computing.

In early October 2024 passersby at HLRS stared as a forklift heaved heavy crates through a hole in the side of the center's headquarters on Nobelstraβe. Inside the crates were the components of Hunter, Stuttgart's newest and most powerful supercomputer.

Designed and manufactured by Hewlett Packard Enterprise (HPE), Hunter offers a world-class infrastructure for large-scale simulation, artificial intelligence, and data analytics applications in science, industry, and the public sector. With a theoretical peak performance of 48.1 Petaflops (48.1 quadrillion floating point operations per second), Hunter is nearly twice as fast as HLRS's previous flagship supercomputer, called Hawk. At the same time, its use of next-generation processors means that it consumes 80% less energy at peak performance than its predecessor.

Hunter is based on the same HPE Cray Supercomputing EX4000 architecture used in the world's three verified exascale systems. The new supercomputer also constitutes an important technology shift at HLRS, marking the beginning of a new chapter in the history of supercomputing in Stuttgart. Whereas past systems have relied primarily on large numbers of CPUs to achieve high performance, Hunter is based on the AMD Instinct MI300A accelerated processing unit (APU), which combines CPUs, GPU accelerators, and high bandwidth memory in a single package.

Among the guests celebrating Hunter's inauguration were (l-r): Berthold Schmidt (TRUMPF), Michael Resch (HLRS), Peter Middendorf (University of Stuttgart), Matthias Lederer (Porsche), Petra Olschowski (Baden-Württemberg Minister of Science, Research and Art), Tom Schneider (TRUMPF), Brad McCredie (AMD), Anna Christmann (Member of German Parliament), Heiko Meyer (HPE), Trish Damkroger (HPE), Marc Fischer (HPE), and Michael Rafii (Federal Ministry for Education and Research).

Hunter was also conceived with economic and environmental sustainability in mind. The system makes use of 100% fanless direct liquid cooling system architecture from HPE, which greatly reduces the energy required to keep a supercomputer of this size from overheating.

Hunter nearly doubles the performance of its predecessor, Hawk, while using 80% less energy at peak performance.

In addition, the system will use a dynamic power capping software functionality developed by HPE in collaboration with HLRS to maximize computational productivity and energy efficiency. This functionality continuously monitors the currently running applications, optimizes power distribution to each application based on its individual power requirements, and ensures that overall energy consumption does not exceed a predetermined power limit (see page 43).

More power for artificial intelligence

The adoption of an architecture based on APUs will make Hunter a powerful system for artificial intelligence. Although Hunter will continue to support traditional applications of high-performance computing, it also offers the potential to expand HLRS's user community to include data scientists and artificial intelligence specialists interested in developing customized large language models, deep learning projects, and complex data analytics.

Moreover, the combination of CPUs and GPUs in a single package will make it easier to develop and run new kinds of converged computing workflows that combine simulation, data analysis, and artificial intelligence in innovative ways. Simulation could be used on Hunter, for example, to produce synthetic data sets for the training of AI algorithms. Conversely, artificial intelligence could be integrated into simulation workflows to accelerate computationally expensive codes.

A stepping stone to the next level of performance

Hunter is conceived as a transitional system that will prepare the way for HLRS's next supercomputer, called Herder, which is being planned for installation by HPE in 2027. With a top speed of several hundred petaflops, Herder will constitute a major jump in peak performance over Hunter. Because Hunter will leverage a similar GPU-accelerated approach to Herder, it will offer HLRS's user community the opportunity to prepare their application codes to take full advantage of the massive increase in computational power that will soon be available.

Funding of 15 million Euros

The total cost of Hunter was 15 million Euros. The Baden-Württemberg Ministry for Science, Research, and Art provided half of the funding in conjunction with its high-performance computing/data intensive computing (HPC/DIC) strategy. Germany's Federal Ministry for Education and Research provided the second half in the context of the SiVeGCS project. Financing was facilitated by the Gauss Centre for Supercomputing (GCS), the alliance of Germany's three national supercomputing centers.

HLRS celebrated the inauguration of Hunter on January 16, 2025 in the presence of Baden Württemberg Minister of Science Petra Olschowski, University of Stuttgart Rector Peter Middendorf, HPE Executive Vice President and Chief Sales Officer Heiko Meyer, AMD Senior Vice President for Data Center Engineering Brad McCredie, and many other honored guests and members of the HLRS community.

HLRS Wins Datacenter Strategy Award for "Transformation"

The prize recognizes HLRS for its sustainable computing center strategy and its visionary new construction project, HLRS III.

Dr. Bastian Koller, Managing Director of HLRS, accepted the award at the 2024 Datacenter Strategy Summit.

A major expansion is on the horizon for the High-Performance Computing Center Stuttgart. Beginning in 2027, HLRS will begin installing Herder, a supercomputer that will enable HLRS to make a large leap in performance. To provide a suitable home for this new system when it arrives, construction of a new building, called HLRS III, began in early 2025.

Already the planning of HLRS's future is making waves. In a ceremony at the Datacenter Strategy Summit in October 2024, HLRS was named winner of the Datacenter Strategy Award in the category "Transformation." The prize recognizes HLRS for innovative approaches to energy efficiency, sustainable computing center infrastructure, and dynamic energy management that are all integrated into the planning of HLRS III.

"I am very pleased about this award, which marks the success of our efforts in implementing sustainable strategies for high-performance computing," said Dr. Bastian Koller, Managing Director of HLRS. "The concept for HLRS III shows that we not only provide high-performance computing resources, but are also aware of our responsibility to society. This award recognizes that technology and sustainability can go hand in hand." Koller accepted the trophy on behalf of HLRS.

Herder will offer much higher performance in comparison to the center's recently installed flagship supercomputer, known as Hunter. But this means that it will also have much higher energy demands – up to 8 MW of IT power. In planning HLRS III, the goal has been to make the data center future-proof and energy-efficient. Design has prioritized the use of sustainable materials, while the surface of the building will be covered with a photovoltaic system to produce clean energy. The waste heat generated by Herder will be captured and distributed across the University of Stuttgart's Vaihingen campus, contributing to its decarbonization. Combining waste heat reuse with a systems operation approach called dynamic power capping (see page 43), HLRS III will enable a more efficient distribution of unused energy that responds continually to power supply and demand.

The Vogel IT Academy has presented the Datacenter Strategy Awards since 2022 in the categories Transformation, Sustainability, and Innovation. The award in the Transformation category honors a project or strategy that decisively advances digital transformation and offers an organization concrete added value.

A preliminary architectural rendering of HLRS III shows the future computing center and power facility behind HLRS's current building.

HPC for the Public Sector

Simulation and digital twins offer communities new approaches for improving public safety, adapting to climate change, and preserving fragile landscapes.

High-performance computing (HPC) supports powerful methods with the potential to make communities more sustainable, resilient, and efficient. Existing technologies for simulation, artificial intelligence, and visualization, for example, offer strategies for accelerating energy transformation, planning healthier and more livable cities, improving administrative efficiency, managing effects of climate change, and predicting risk of natural disasters, among countless other applications.

In recent years, HLRS has been developing strategies for helping public sector organizations in addressing the challenges they face. The effort gained new urgency during the COVID-19 pandemic, when HLRS collaborated with the German Federal Institute for Population Research to develop a tool for monitoring and predicting demand for intensive care unit beds. This was followed in 2021 with the launch of CIRCE, a project funded by the German Federal Ministry of Education and Research, and the Baden-Württemberg Ministry of Science, Research and Art that assessed the need and requirements for a potential Computational Immediate Response Center for Emergencies.

Simulation using supercomputers could help public authorities prepare for and manage emergency situations like floods, pandemics, or migration events.

HLRS expanded upon these initiatives in 2024 by conducting outreach to increase awareness of the capabilities of high-performance computing in public agencies. "HLRS is a publicly funded institute that operates very valuable machines … and society should profit from them," said Dr. Bastian Koller, Managing Director of HLRS, at a June 2024 symposium organized by CIRCE in Berlin. "For us the question is, how can we improve interactions with the public sector so that they better understand the resources that are available?"

CIRCE assesses HPC for crisis prevention and management

In a qualitative, interview-based survey of representatives of federal, state, and local agencies, as well as crisis response organizations, CIRCE found that many recognize opportunities for using simulation to address challenges they face. Among the potential application areas that respondents mentioned were migration, pandemics, floods, wildfires, chemical or nuclear accidents, and terrorist attacks. Dialogue between CIRCE and public sector representatives also identified technical and administrative considerations that will need to be addressed to ensure that supercomputing resources and expertise are readily available when crisis situations arise.

In an event co-organized by HLRS, approximately 120 representatives of local communities met to discuss strategies for adaptating to climate change.

In June 2024, the CIRCE team held a full-day symposium in Berlin that presented its preliminary findings to representatives of public administration and experts in disaster response. In addition to offering general insights into the requirements for establishing a computational emergency response center, the talks highlighted applications of simulation, AI, and data visualization for crisis management. Data scientists at German federal governmental institutes enriched the meeting by describing how data is currently used in government planning and decision-making, and offered insights into the practical challenges of navigating Germany's complex, federalized system of government.

One important takeaway from the meeting was that close collaboration between HPC centers and committed partners at the local, state, and federal levels is crucial for developing effective solutions. Another was that advance preparation for emergency situations is essential. "We can't simply expect to be able to manage crises on an ad hoc basis, but need to be active even before crises arise," remarked Dr. Juliane Braun, Chief Data Scientist of the Federal Ministry of the Interior and Community (BMI). This means not only developing the right simulation software, but also ensuring that monitoring procedures, communication systems, and guidelines for decision making are in place before a crisis computing application goes into service.

With Germany's 294 regional districts and an additional 107 self-organized cities, it will not be practical or efficient for each individual community to develop its own specialized solutions. Instead, increased collaboration is needed to design applications that are widely usable. As a result of CIRCE's outreach activities, Christian von Spiczak-Brzezinski of the Duisburg fire department has been working with HLRS to develop a simulation tool for predicting flooding along the Rhine. He observed that for simulation to have a wider impact, public organizations need a better understanding about what HPC can offer and how to work with high-performance computing centers to implement HPC-based solutions most effectively.

Tools for addressing climate change at the local level

A public outreach event co-organized by HLRS in April 2024 in the municipality of Sersheim (Ludwigsburg district) welcomed approximately 120 visitors from across Baden-Württemberg to discuss the potential effects of climate change at the local level. In the event, called "Science Goes Society," renowned experts from the fields of science, urban and regional planning, and municipal practice, including the president of the Baden-Württemberg Municipal Association, explored the challenges that municipalities face. Also present were representatives of public administrations; district, city, and municipal councils; and citizens concerned about their climate future.

As part of the event, a large LED wall with 3D simulations showed potential effects of climate change on local villages, cities, forests, and fields. They also showed where precautions can be taken for sustainable development and disaster prevention. In hands-on demonstrations and at exhibition stands, participants presented examples of how municipalities are already preparing for the expected effects of climate change, including droughts, heat waves, heavy rainfall, flooding, and erosion.

At "Science Goes Society," representatives of HLRS presented recent simulation and visualization tools

Virtual reality is supporting planning and preservation efforts for the vineyards on the banks of the Neckar River.

focused on modeling pandemic spread, flooding, and forest fires. As HLRS's capabilities for artificial intelligence grow in the coming years, high-performance tools for data analysis could also offer new approaches for simplifying administrative processes and real-time decision making in communities.

Digital twins help to preserve landscape and culture

Digital twins offer detailed models of physical locations in virtual reality, integrating data in ways that make them easier to interpret. Over the past several years, visualization scientists at HLRS have been demonstrating how digital twins can support decision-making in cities and communities. The benefits of digital twins include providing a data-rich, realistic model of a location; supporting more effective, data-based discussion and participatory planning processes; and making it possible to simulate scenarios virtually.

In April 2024, HLRS scientists met leaders of Ludwigsburg community governments for a boat ride along the Neckar River, whose terraced cliffs and hills are home to vineyards operated by independent and cooperative wine growers. At several locations along the river, the terraces and their associated agriculture now find themselves in a precarious state, and communities are seeking solutions to preserve the landscape and its unique winegrowing culture.

At the meeting, representatives of HLRS presented a digital twin of the Neckar terraces between Stuttgart and Besigheim, anticipating a number of potential applications. These could include, for example, simulating heavy rainfalls or landslides to identify potential risks. As new land survey data becomes available and is fed into the digital twin over time, methods based in artificial intelligence could also be used to document erosion and other deterioration on the cliffs.

According to Claus-Peter Hutter, president of the foundation NatureLife-International and the initiator of the project, "The goal is to gain a unique perspective on changes that threaten the landscape we have come to love."

The digital twin will also make it possible to simulate planning scenarios, providing information that will help stakeholders in local communities to understand the potential effects of specific decisions. In some cases, the digital twin could help to evaluate new uses of the large, exposed surfaces, such as for solar power generation. Ideally, the knowledge gained could also help to attract a new generation of wine growing projects that are optimally suitable for local conditions. In this way, HLRS's digital twin will support efforts aimed at retaining local heritage, while also preparing for the future.

At the CIRCE conference in Berlin, Bastian Koller said that HLRS cannot be alone in building out Germany's simulation capabilities for the public sector, but will need to find more data-oriented allies in public administration. "From where I sit, I see the beginnings of a change in thinking," he remarked. "We will have to continue to try to expand our activities and to communicate their results to the public to show them that something new is really happening here. As more people see what we are doing and find it valuable, my hope is that interest will snowball."

Professional Certification in High-Performance Computing

Focused tracks in the fields of simulation, machine learning and AI, parallel programming, and HPC administration offer new opportunities for career development.

The high-performance computing training program at the High-Performance Computing Center Stuttgart is among Europe's most experienced and comprehensive programs for continuing professional education in HPC, AI, and related technologies. Through a new initiative launched in 2024, HLRS now offers trainees the opportunity to achieve professional certification in key high-performance computing occupations.

HLRS offers four professional certification tracks: simulation development, machine learning and artificial intelligence development, parallel programming, and HPC administration. Each track offers a flexible curriculum, enabling trainees to select from a menu of courses and to customize their training experience in a way that aligns with their professional interests and career goals.

HLRS's professional certification tracks build on the center's extensive training program, which organizes more than 40 courses each year.

According to Dr. Lorenzo Zanon, who leads HLRS's Department of Training and Scalable Algorithms, the professional certification program will enable trainees to broaden and deepen their expertise in essential disciplines from across the HPC field. "Our new professional certification lets you develop proficiency on topics like machine learning and AI algorithm development, performance optimization, hardware accelerators, scientific visualization, cluster operation, or data management," he explained. "Combining multiple elected courses within a structured framework will help those who complete certification to gather well-rounded experience in the latest technologies and methods."

Certification builds on the foundation of HLRS's extensive calendar of approximately 40 training courses each year. Participants in the certification program must complete at least three courses selected from key topic areas and pass an exam. Once the requirements are fulfilled, participants can request a certificate that can be shared with current or future employers. Ultimately, professional certification at HLRS will enable trainees to demonstrate proficiency that makes them stand out in today's competitive technology job market.

Information about professional certification, including course requirements and related courses, is available at www.hlrs.de/training/professional-certification.

User Support for Code Porting to GPUs

HPC experts at HLRS helped University of Stuttgart scientists in preparing the FS3D code to move from Hawk to Hunter.

Dr. Karthrin Schulte (center) and graduate students David Gösele (left) and Jonathan Wurst (right) of the University of Stuttgart's Institute of Aerospace Thermodynamics are using HLRS's Hunter supercomputer to simulate droplet dynamics.

Most of the world's fastest supercomputers now use graphics processing units (GPUs) as accelerators. Initially developed for video games, GPUs can compute faster and more efficiently than central processing units (CPUs), the traditional workhorses of high-performance computing (HPC). The catch, however, is that codes written for CPUs will not typically run on GPU-based systems.

Recent generations of supercomputers at HLRS — including Hazel Hen and Hawk — primarily relied on CPUs. With the arrival of Hunter in 2024, however, longtime users of HLRS's systems began navigating the transition to GPUs in earnest. Although the new system is nearly twice as fast as Hawk when running at peak performance, users will only truly realize the benefit by successfully rewriting their codes for GPUs.

A research team led by Dr. Kathrin Schulte at the University of Stuttgart's Institute of Aerospace Thermodynamics (ITLR) is among the user groups that have been navigating this transition. For more than 25 years, ITLR researchers have used HLRS supercomputers to run a code called FS3D (Free Surface 3D). A large, modular code with many applications, FS3D uses direct numerical simulation to study fundamental physical processes involving droplets. Such simulations are relevant in numerous fields, from weather prediction, to the design of fuel injectors in engines, to the spraying of paint during manufacturing.

In mid-2024 HLRS held a hands-on code porting workshop for its user community. During the workshop, the ITLR team met with HLRS user support staff as well as hardware experts from Hewlett Packard Enterprise (HPE, Hunter's manufacturer) and AMD to identify the simplest and most effective strategies for offloading parts of FS3D to GPUs.

When it was initially developed many years ago, FS3D was optimized for a vector computing architecture. Thus, the scientists had experience using the programming model OpenMP for parallel computing using shared memory. This offered the team certain advantages in preparing for Hunter, as array-based vector computing has technical similarities to the concept behind GPU processors. After discussing the advantages and disadvantages of different programming models with HPC experts at HLRS, the team opted to use OpenMP in the GPU implementation of FS3D.

"It would have been possible to use any of a number of other models, although choosing unwisely would have meant needing to rewrite even more code," Schulte explained. "Because we had no experience with GPUs, we were not really qualified to make a decision. HLRS supported us in choosing a path."

Despite the experience with vector computing and OpenMP, significant work was still needed to prepare FS3D for Hunter. The team met regularly throughout 2024, communicating often with HLRS user support staff and HPE technical experts for guidance. When they encountered bugs or identified features in FS3D that were not supported in the new programming environment, having quick access to system experts helped them to resolve problems and develop workaround strategies. "When a problem arose with a compiler, for example, it wasn't like we got stuck for weeks. They helped us to find solutions quickly. We appreciated the constant support from all sides — HLRS, HPE, and AMD," Schulte remarked.

Before Hunter was fully installed, HPE provided a small cluster at HLRS that is based on the same architecture. This enabled the ITLR team and others in HLRS's user community to begin testing the ported versions of their codes. With these resources FS3D was ready for GPUs by the time Hunter went online in February 2025.

Additional optimization of FS3D is still needed to improve its performance, and so the ITLR team continues to work with staff at HLRS, HPE, and AMD to develop it further. By doing so now, the ITLR team will be in a good position to use HLRS's next-generation system — called Herder — which will offer a dramatic jump in performance when it arrives in 2027. These capabilities will enable Schulte and her lab to achieve results faster and to investigate new kinds of phenomena in droplet dynamics.

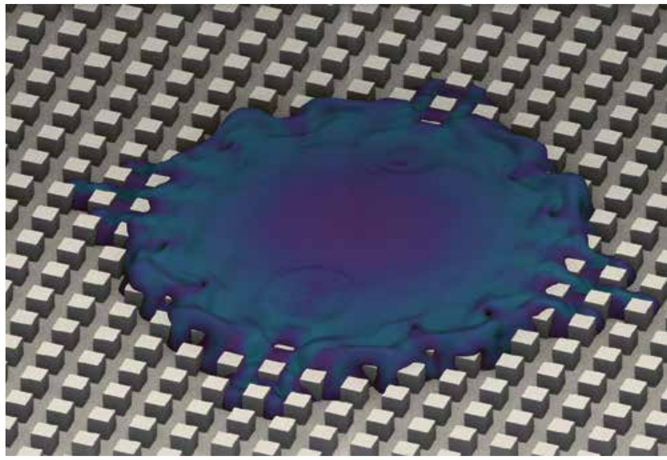

In recent research, the Schulte Lab has been using FS3D for high-resolution simulations of how droplets behave when impacting onto a structured surface.

Flow Simulations Reduce Operating Costs in Water Resources Management

Using HLRS supercomputers, Dresden-based engineering firm hydrograv GmbH runs CFD simulations to optimize efficiency in water management facilities.

Providing sufficient clean water is essential in any community, and the pumps and other machinery needed to do so require large amounts of power. Indeed, for many cities and towns across Germany, the cost of electricity to operate water infrastructure is the single largest expense in their annual budget.

Water treatment plants are typically designed and constructed by civil engineers, working closely with mechanical engineers who oversee the facilities' mechanical and process engineering components. Other team members include environmental and process engineers, as well as biotechnology engineers, who are responsible for biological wastewater management. They ensure the health of the microorganisms that break down organic materials in sewage, and ultimately the sanitary performance of the facility.

One highly relevant discipline that is typically missing in such projects, however, is computational fluid dynamics (CFD), the science of predicting how fluids and the materials of which they are made flow. According to Dr. Martin Armbruster, Managing Director of hydrograv GmbH, this is too often a missed opportunity. "Mercedes-Benz or Porsche would never put an automobile on the road without optimizing the aerodynamics extensively using simulation," he observes. "In our field, however, it is still exotic to optimize an engineer's design using CFD. Because this rarely happens, we often build too large and inefficient."

Hydrograv was founded in Dresden in 2004 as a spinoff company from the Karlsruhe Institute of Technology and the Technical University Dresden, and specializes in CFD simulation for water resource management. Using customized open source software, hydrograv advises facility operators on how to eliminate inefficiencies and improve water quality. The company also assists engineers during the planning of new facilities, using simulation to optimize efficiency during the design phase. This makes it possible to identify and rectify potential deficiencies in processes involved in mechanical or biological water purification early — long before the first cubic meter of concrete is poured.

CFD simulations of water management facilities model complex multiphase flows, an approach that is only practical using high-performance computing (HPC). Although hydrograv operates its own computer servers for smaller simulation tasks, running simulations at this scale would not be possible on its own. This is why in 2019 the company began using supercomputing resources at the High-Performance Computing Center Stuttgart (HLRS) as an integral part of its simulation pipeline.

"Considering the amount of parallel simulation that we now do, our work wouldn't be possible if we didn't have HLRS as a partner," Dr. Armbruster remarked. The partnership is an impressive example of the opportunities that small companies can derive from HLRS's high-performance computing systems, as well as how HPC can contribute to sustainability.

Designing more efficient aeration tanks

One key component of a water treatment facility is the aeration tank, where microorganisms break down organic pollutants in waste water. For these microorganisms to work optimally, it is important to provide evenly distributed, dissolved oxygen throughout the tank. This is usually accomplished using blowers and fine-bubble aeration systems that efficiently introduce oxygen into the tank. The principle is similar to oxygen enrichment in aquariums, but on a much larger scale and with specially designed technical systems for energy efficiency and process control.

The placement of the aeration elements (aerators) in the aeration tank is of decisive importance for an efficient distribution of oxygen. "Although the most inexpensive and easiest to maintain solution might be to place aerators in a single area, this is not optimal from a flow perspective," Dr. Armbruster explained. A poorly chosen distribution of aerators around a tank can cause air to reach the surface of the water too quickly and escape into the atmosphere. As a result, not enough oxygen dissolves in the water. This can require pumping more air into a tank than is really needed, an inefficient approach that leads to unnecessarily high energy usage.

Using simulation, hydrograv can realistically predict how oxygen will be distributed in a tank based on the distribution of aerators and different operating scenarios. A study based on engineering practice showed that optimizing aeration could increase the amount of oxygen introduced into the tank by more than 40% using the same amount of power. "When you consider that running aerators is the largest expense in a water treatment plant, which itself is often the largest energy consumer in a city, you can clearly see why optimization using CFD can be extremely important," Dr. Armbruster said.

In addition to improving aeration, CFD can also be used to optimize the placement of pumps and the geometries of water tanks and canals. When circulation is not optimized, dead zones can develop where waste water is not evenly distributed, leading to inconsistent water treatment. Suboptimal flow conditions can also lead to deposits that impede flow and increase the risk that a tank could overflow. This, in turn, forces pumps to work harder, raising energy usage and operating costs. In addition, an unfavorable inflow can reduce the service life of a pump, leading to increased wear and mechanical stress.

Using CFD simulations, hydrograv can make targeted suggestions about structural improvements that increase circulation, such as the ideal placement of guide walls to direct water flow and promote mixing. Such optimizations help to design complex water treatment processes to be more efficient and to reduce their energy usage.

Operating water management facilities is the largest expense that many communities face. Simulation can be used to optimize aeration and circulation, reducing the energy required to provide clean water.

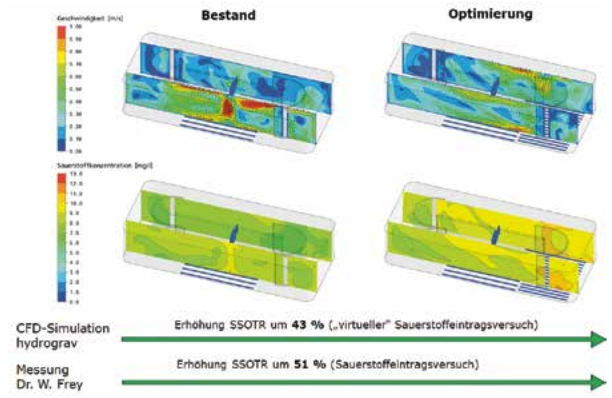

In this simulation project, the introduction of oxygen in an aeration tank was increased by 43 percent through flow optimization using CFD simulations. The simulation results were confirmed by measurements carried out by W. Frey.

In a project with the city of Erlangen, for example, hydrograv ran CFD simulations at HLRS that enabled its water management facility to make major improvements in the secondary clarification step. The approach was so successful that it was possible to completely shut off a filtration system that was previously needed to reach water quality goals. This optimization has reduced energy usage by 500,000 kWh per year.

In another project, hydrograv supported the city of Hanau in optimizing its secondary clarification tank. Avoiding a need for the city to spend 35 million Euros to reconstruct the facility, hydrograv provided an efficient solution for adjusting inflow that could be seamlessly integrated into the existing tank. At a cost of just 1.5 million Euros, the solution not only saved the city money but also for the first time enabled it to reach demanding targets for reducing phosphorous concentration in its water. Hydrograv GmbH was awarded the Saxon Environmental Award for its outstanding contribution to environmental protection.

HLRS offers powerful, flexible capabilities for simulation

Since its founding, hydrograv has completed more than 1,000 simulation projects, successfully serving clients across Germany, Europe, and around the world. Its need to run many simulations now far outstrips the capabilities of its in-house computing servers. Establishing a contract with HLRS to use its supercomputer has thus become absolutely essential for its work.

Using a node on HLRS's supercomputer, hydrograv's simulations can often take 24 hours, with large-scale simulation projects lasting up to a week. For some projects, the company must also run multiple simulations of this size to arrive at an optimized solution. With its software installed at HLRS, hydrograv can access the system from Dresden, giving it a powerful and flexible asset that is now integrated into its simulation workflows.

"At times we might have 3 or 4 clients in a month, each of which requires running 5 or 10 simulations. If next month we don't have the same computing needs, however, it would be difficult to decide how many computers to put in our server room," Dr. Armbruster observed. "This is the advantage of running these simulations at HLRS. We can access exactly the amount of large-scale computing power we need, when we need it."

For Hydrograv, knowing that its software and data are protected in a secure environment is also extremely important. "It is very advantageous to be able to do this in a German national high-performance computing center, as we and our clients can be sure that our know-how stays in the country," Dr. Armbruster said.

The use of CFD in the design of water treatment infrastructure is still unusual. Considering the large number of facilities across Germany and Europe, and that the cost of simulation is much lower than that of wasting energy, Dr. Armbruster hopes that the approach will be used more often in the future. Doing so could not only help communities in reducing costs, but also make sanitation and energy infrastructure more sustainable.

Supercomputing in the Arts

In partnership with the Media Solution Center Baden-Württemberg, HLRS hosts artist residencies exploring creative applications of HPC technologies.

Founded by HLRS, the ZKM, and the Hochschule der Medien in 2018, the Media Solution Center Baden-Württemberg (MSC) facilitates collaboration among artists, scientists, and experts in high-performance computing. It has become an important node within European culture and creative industry networks, giving HLRS a role in EU initiatives promoting digital innovation in the arts.

In cooperation with the MSC, HLRS has been contributing as a partner in S+T+ARTS (Science, Technology, and the Arts), an initiative launched by the European Commission in 2016 that organizes residency programs to support multidisciplinary collaboration. The MSC coordinated an open call and residency program called S+T+ARTS AIR, which awarded ten residencies, two of which were hosted at HLRS in 2024. The results of these and other S+T+ARTS AIR projects were presented in Barcelona in October in a two-day festival organized by the MSC in collaboration with the Fundación Épica La Fura dels Baus. The event was a centerpiece of the nearly two-month-long "Future Boundaries" festival organized by theGATE (Global Art-Science and Technology Development).

HLRS's participation in S+T+ARTS continues to grow. In a new residency program called S+T+ARTS EC(H)O, HLRS is hosting four new residencies in 2025. The artists' projects are exploring digital technologies for exploring European cultural heritage, technological solutions inspired by nature, immersive visual narratives based on data, and how digital twins can bring socio-cultural and ecological dimensions of urban environments to life.

In addition to its contributions to the S+T+ARTS program, HLRS hosted two ReACH Workshops (Research and Creation Center for eCulture ant the Humanities) in April and September 2024. The three-day meetings, also organized by the Media Solution Center, invited artists and other representatives of the creative industries to learn about new applications of digital technologies in the arts and to gain a better practical understanding of how to develop an artistic concept into a successful scientific collaboration.

Results of the S+T+ARTS AIR residency program were presented at TheGATE Festival in Barcelona.

Connecting the Simulation Sciences and the Humanities

Two multidisciplinary conferences explored how the computational sciences and fields such as theater, literary studies, and architecture can complement one another, leading to new kinds of insights.

Simulation is not just a tool for science and engineering, but also a way of seeing and interacting with the real world. At HLRS, investigating the concept of simulation from different perspectives is therefore an important part of gaining a better understanding of what the computational sciences can tell us.

Since 2023, for example, HLRS has participated in the Stuttgart Research Focus "(Re-)Producing Realities" (Re^2). Gathering experts in the digital humanities, social sciences, production technologies, and the simulation sciences from the University of Stuttgart and the surrounding cultural community, Re^2 uses theater as a structural model for investigating how representation shapes perception and understanding of the world around us.

A kickoff symposium held at the university on June 26, 2024 introduced the goals of Re^2. In one approach, Re^2 is studying how digital technologies can reproduce historical realities. At HLRS, visualization scientists have created digital twins of the Ludwigsburg palace theater and the Théâtre des Tuileries in Paris. The simulations offer compelling methods for preserving historically important stage machineries and opening new avenues of research. In another approach, Re^2 will look at how performance produces realities. Considering that both theater and simulation require artifice in creating representations, looking at the digital world through the lens of theatrical performance could lead to a deeper understanding of the assumptions that go into technology development and usage. At the Re^2 kickoff event, three keynote talks presented examples of the types of research the project aims to facilitate.

In an international conference titled Reality in Science, Art, and the Humanities, held at HLRS in October 2024, scholars sought to identify other kinds of similarities and differences in conceptions of reality across the sciences and humanities. Researchers from Germany, Japan, and Saudi Arabia explored concepts of simulation in a wide range of applications, focusing on the inherent assumptions and gaps that problematize their relationships to reality. Talks probed scientific and technical applications such as computational fluid dynamics and an early warning system for natural disaster management, as well as cultural applications like a virtual reconstruction of the Palace of Versailles and textual reconstruction of the writings of Franz Kafka.

HLRS created a digital twin of the Ludwigsburg palace theater, providing a tool for researching its historically important stage machinery using virtual reality.

Dynamic Power Capping Improves Energy Efficiency in HPC

An intelligent power management solution regulates power distribution to optimize system performance within a power budget.

In February 2024, HLRS implemented a new solution on its Hawk supercomputer that produced major gains in energy efficiency. Using an approach developed by Hewlett Packard Enterprise in collaboration with HLRS called dynamic power capping, the center reduced the overall power consumption of applications by about 20% without causing significant losses in performance. The collected power savings are comparable to the annual power usage of approximately 1,500 single family homes.

Conventional power capping approaches limit power usage by throttling processor speed, running the system slower than its full capabilities. Although this method can reduce the absolute amount of power consumed, it also limits the ability of supercomputers to fulfill their key mission: running massively parallel simulations as fast as possible.

Dynamic power capping, in contrast, exploits the fact that different types of codes have different power requirements. In compute-bound codes, an algorithm simply performs better when the system speed is faster. In memory-bound codes, however, the time it takes for a code to run depends less on processor speed and more on an HPC system's memory and data transfer capabilities. Here, maximizing CPU speed does not increase overall code performance, because the algorithm must wait for data transfer before performing its next calculation.

"HPE's dynamic power capping approach is unique in that it balances the different power requirements of these two categories of codes within a given available power budget," explained Dr. Christian Simmendinger, an HPC performance engineer at HPE. "For memory-bound codes the available power can be capped significantly, leading to energy savings without negatively affecting application performance. The power capping level is periodically optimized in an automated way, reacting to changing phases in the running of an application."

Since implementing the approach, the team found that balancing available power between compute-bound and memory-bound codes reduces sudden spikes and drops in overall system power usage, facilitating a consistent, steady-state power level that adheres to HLRS's power consumption goals. The framework can also respond dynamically if HLRS changes its desired power limit. HLRS and HPE plan to extend the capabilities of dynamic power capping for HLRS's next-generation supercomputers, Hunter and Herder.

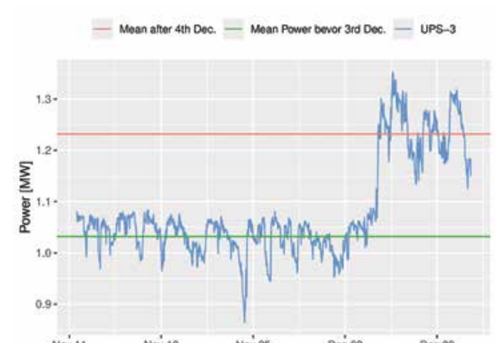

The green line indicates the mean power usage using dynamic power capping. The red line shows the mean power usage during an experimental operation in uncapped mode in December 2024. The experimental operation demonstrated that dynamic power capping enabled an energy savings of approximately 20% without recognizable performance degradation.

User Research

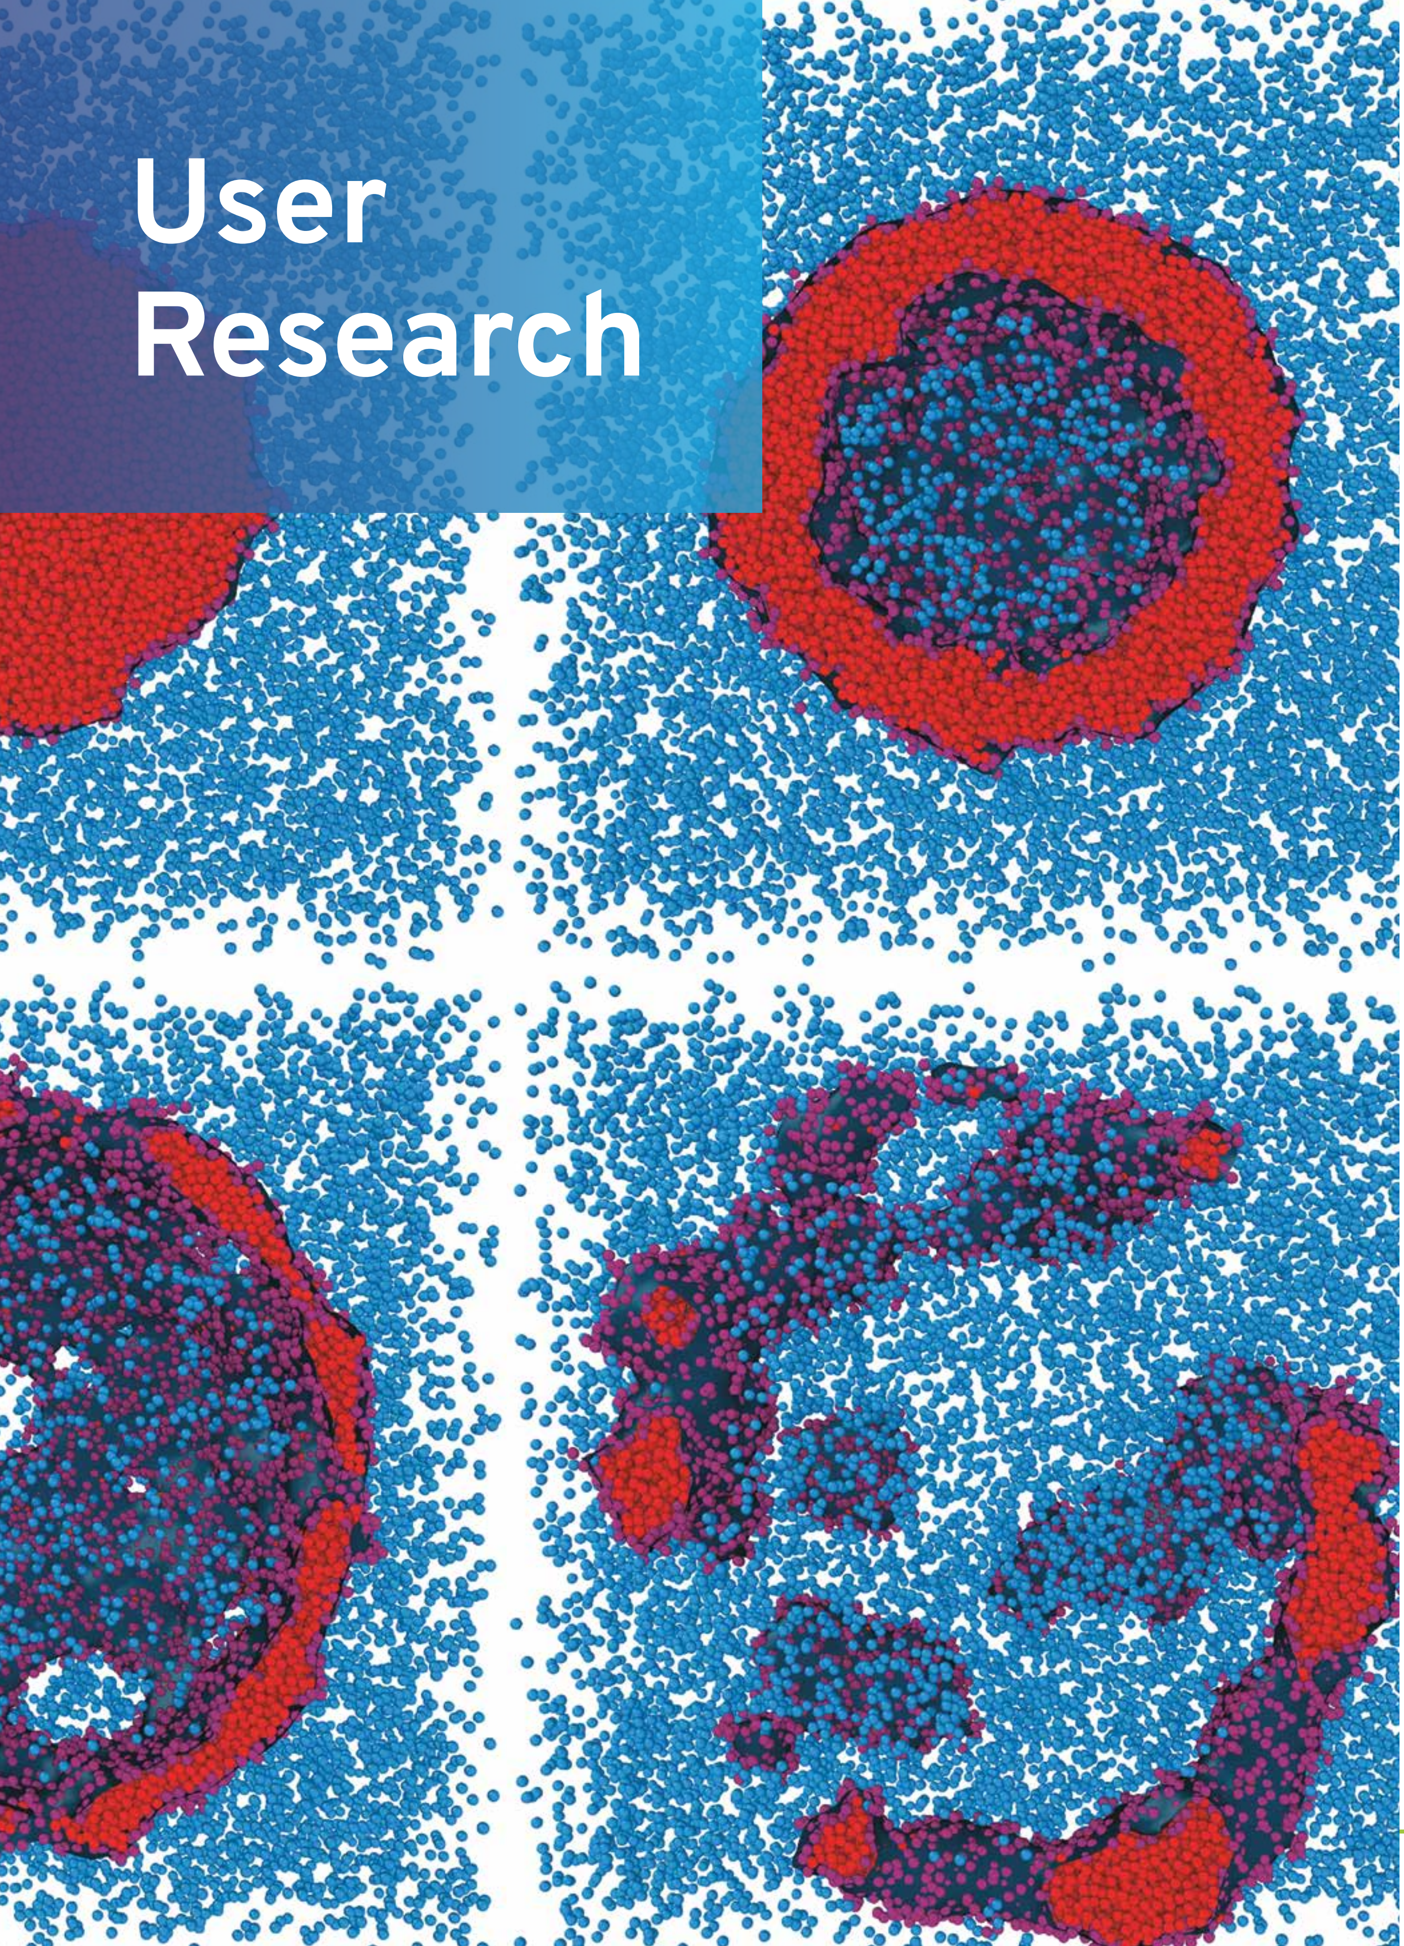

◀ Researchers in the Laboratory of Engineering Thermodynamics at RPTU Kaiserslautern ran molecular dynamics simulations at HLRS to investigate how droplets explode, a process that is important, for example, in combustion applications.

Designing Quieter, Safer Helicopters with HPC

Using HLRS supercomputers, University of Stuttgart researchers supported Airbus in the development of the next-generation RACER helicopter.

With the advent of air taxis and drone-based delivery services, demand is growing for a new generation of helicopters that are as safe, quiet, and energy-efficient as possible. A research group led by Prof. Manuel Keßler at the University of Stuttgart's Institute of Aerodynamics and Gas Dynamics (IAG) has been busy answering this need. Keßler's team has long focused on improving helicopters by using supercomputers at the High-Performance Computing Center Stuttgart (HLRS) to understand the complex aerodynamics surrounding helicopter rotors.

"My group has been focused on this topic for two and a half decades, and throughout that period, we have had a close cooperation with industry," Keßler said. "They come to us with questions they need support with, and we provide them with knowledge and insights, as well as simulation tools for answering these questions. We've had a long-running collaboration with Airbus Helicopters, and they are using many of the same tools in their research as we are. So, the outcome of our research finds its way directly into industrial products." Recently, Keßler and his collaborators worked closely with Airbus on modelling and simulation efforts to reduce noise and better understand flight dynamics in the company's new RACER prototype. RACER, which took its first test flight in April 2024, is unusual in that it has both helicopter rotor blades and airplane propellers on wings. These features will enable the new aircraft to respond more rapidly in emergency situations. The collaboration to improve RACER with modelling and simulation is an example of the ways in which public HPC centers contribute to European economic development and competitiveness.

Airbus Helicopter's newest RACER helicopter has both traditional rotor blades as well as airplane propellers on wings. The new aircraft will be able to respond more rapidly in emergency situations.

Modelling and simulation expedite time to market

While helicopters are not a recent invention, making them more efficient, safer, or quieter requires understanding a complex mix of physical phenomena at a fundamental level. To develop models of these phenomena, the researchers must break down the helicopter's complex geometry into a very fine computational grid of roughly 200 million grid cells. In a standard simulation, the researchers include 10 rotations of the helicopter rotor. To capture the small-scale changes that can affect flight dynamics or acoustics, they typically divide each rotation into 720 small time steps. This means completing, tracking, and understanding relationships among more than a quadrillion individual calculations. "These simulations are too large for us to attempt on a workstation or smaller cluster; we need to have access to supercomputers to perform these calculations in an efficient manner," Keβler said.

In their collaboration with Airbus and as part of the project CA³TCH within the framework of the European Union's Clean Sky 2 programme, Keβler's team modelled the new RACER design both to reduce the noise that it generates and to gain a more complete understanding of flight dynamics. In this way they can lower risk before the craft's first flight. The team ran a suite of simulations under a variety of circumstances, such as hovering during a crosswind and determining the so-called "roll angle," or the maximum sideways angle at which the craft can tilt before the propellers touch the ground. Because RACER has a traditional rotor as well as wings, there is very little prior experimental data to rely on, making the need for modelling and simulation even greater. "Based on our well-established workflow and the rapport that we have built with our industry partners, they take our simulations seriously if we indicate that a model could have a stability problem," Keβler said. "I've always said that our duty is to ensure that after an aircraft's first test flight, the pilots exit with a grin on their faces and were not sweating while they were in flight."

Challenges and opportunities for designing next-generation aircraft

As HLRS pivots from Hawk — until recently the center's flagship supercomputer — to Hunter — a next-generation system that went into service in early 2025 — research teams like Keβler's see both opportunities and challenges. Hunter is the first system at HLRS that includes a large number of GPU accelerators. While GPUs have become a common solution for accelerating performance and improving energy efficiency in the world's fastest computers, they are also a new technology for some scientific users of HPC. For researchers who have not yet used GPU-centric supercomputers, this transition will mean modifying — sometimes heavily — their code packages to take full advantage of the system architecture. With a strong command of Hunter's system, the team should be able to run more high-resolution simulations, further shortening production timelines and improving safety in the process.

Keβler pointed out that HLRS's user support staff has proactively helped make sure his team is ready to use the center's next-generation system as soon as it comes online. "The center has already been very supportive during this transition," Keβler said. "We were recently invited to a meeting where we discussed the different options with staff from different sides of this installation, including from the hardware manufacturers. HLRS came to us actively and said, 'We will be making this transition in the future, so how can we start helping you prepare for it now?'"

Keβler indicated that while HLRS's staff has neither detailed knowledge of his team's code nor the capacity to rewrite codes for the center's dozens of research projects, he considers this early collaboration essential. It will help to ensure that his team can continue to deliver simulations to industry partners, providing insight for next-generation helicopter designs and ultimately ensuring that safe, reliable, and efficient European helicopters continue to come off the production line.

EG

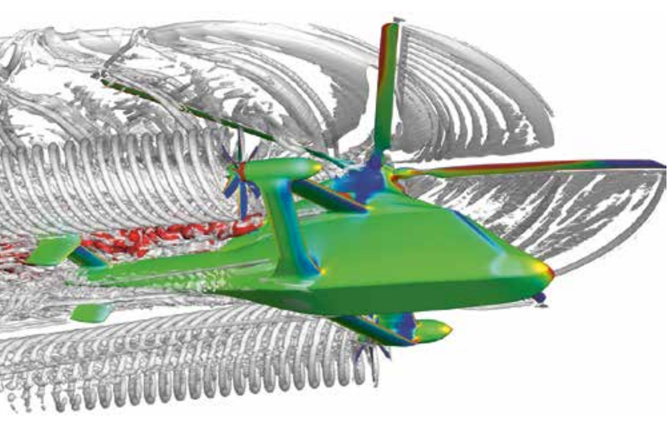

Using supercomputing resources at HLRS, Keβler's team supported Airbus Helicopters in simulating their new RACER helicopter under a variety of flight conditions.

Simulation Supports Search for Personalized Cancer Treatments

Using high-performance computing at HLRS, researchers ran large-scale molecular dynamics simulations to investigate hard-to-treat genetic mutations linked to aggressive cancers.

In recent decades, high-performance computing (HPC) has become an indispensable tool for biomedical science. In cancer research, for example, computational methods using supercomputers can quickly focus researchers' attention on details at the molecular level that hold the keys to understanding and potentially treating disease. Such methods have been particularly important in the field of personalized medicine, which aims to develop more effective drugs that target specific molecular features of individual patients' tumors.

"With the refinement of computational tools that has been going on in the last 40 years, computations have become an important asset for biophysics and pharmaceutical research," says Dr. Giovanni Settanni, who conducts research in computational biophysics at Ruhr University Bochum and Johannes Gutenberg University of Mainz (JGU). "Simulations can be seen as a sort of *in silico* microscope with molecular level resolution, and they help interpret experimental data that can hardly reach the same level of resolution."

Settanni is an expert in the field of molecular dynamics (MD) simulation, an approach that uses principles from biophysics to develop atomic-level models of molecules. One way this method is being used in biomedical research is to predict changes in protein structure, alterations that affect a protein's ability to interact with and bind to other molecules. Because MD simulations offer extremely high resolution and reveal how protein shape can change over time, they can only be done using supercomputers. Recent work in the Settanni Lab using the Hawk supercomputer at the High-Performance Computing Center Stuttgart (HLRS) demonstrates how this approach can help identify new approaches to treating cancer.

Understanding p53, an important protein in cancer

Dr. Settanni has been collaborating with the group of structural biologist Dr. Andreas Joerger of the Goethe University Frankfurt to study an important protein called p53. In a healthy human body, p53 serves as a bulwark against cancer by regulating how cells divide and repair damage to DNA. When a mutation occurs in p53, however, an alteration in its genetic code can produce a change in amino acid sequence that also changes the protein's 3D structure. This can cause it to lose its normal regulatory abilities, leaving cancer cells free to divide and spread uncontrollably. In fact, roughly half of all cancer cases are linked to mutations that inactivate the p53 protein.

One approach in cancer research involves searching for drugs that can neutralize harmful mutations by binding to proteins. In the case of p53, this means looking for small molecules that can stabilize the mutants and return the protein to its original shape so that it can function normally. Until recently, however, researchers have considered p53 "undruggable," as its surface appeared to be virtually impervious to binding with small molecule drugs.

In prior research, Joerger showed that it is possible to target a specific mutation in p53 called Y220C. The Y220C mutation causes a specific physical change in p53, creating a destabilizing surface crevice that results in a loss of the protein's normal shape at body temperature. His lab showed that it is possible to find small molecules that bind to the surface of the Y220C cavity, restoring p53 to its proper shape and its ability to perform its anti-cancer function. That previous result led to a clinical trial to test the safety and effectiveness of this approach in humans.

With this encouraging success, Joerger is now looking for other potentially druggable mutations in p53. He would also like to understand more generally whether or not mutation-specific small molecules offer an effective way to restore the critical protein's function. It is here that high-performance computing is now making an important contribution.

MD simulations help to investigate other p53 mutations

In earlier work, Joerger primarily used a laboratory approach called X-ray crystallography to understand protein structure. This approach works like a high-resolution microscope, beaming a crystallized protein with X-rays to create an atomic level snapshot of it. More recently, however, he has worked with Settanni to incorporate molecular dynamics simulation into his research. MD simulations complement X-ray crystallography by delivering a more comprehensive, high-resolution understanding of how protein shape can change over time. The approach can deliver new insights into a protein's behavior and how its susceptibility to drugging might alter under different conditions.

As reported in a recent paper in the *Nature* journal *Cell Death & Disease*, the team investigated a specific category of structurally unstable cancer-causing mutations in p53. As part of this effort, Settanni used HLRS's Hawk supercomputer to run multiple simulations of several mutants of p53. The simulations, also including the surrounding solvent, contain between approximately 36,000 and 39,000 atoms, and make it possible to follow the time evolution of each mutant cumulatively for 0.8 microseconds.

"Simulations helped characterize the behavior of the less stable mutants of p53, which are difficult to handle experimentally," Settanni said. "Simulations are able to identify the presence of cavities that can form on the surface of the mutant protein. These cavities can be targeted by drugs to stabilize the protein."

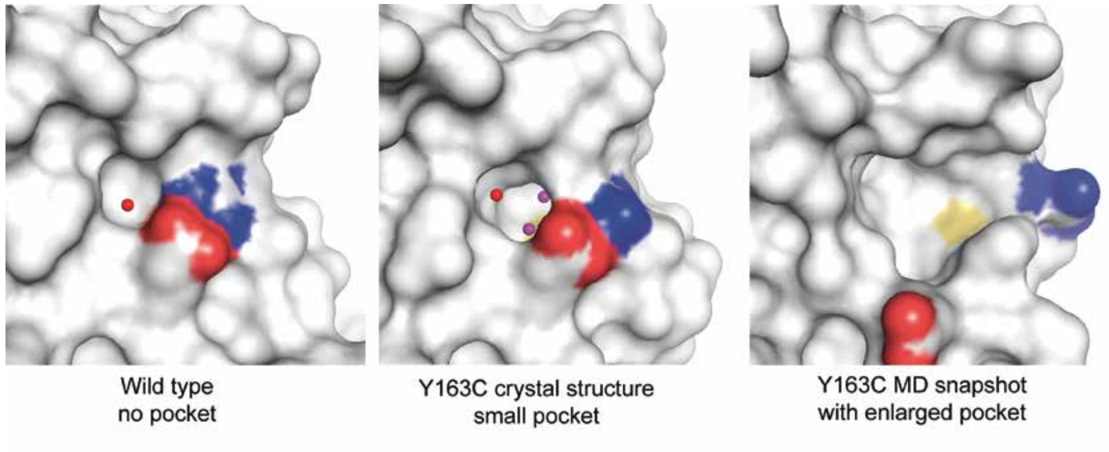

Three images of the same location on the p53 protein. The first image shows the location without a mutation. The center image shows the crystal structure of the Y163C mutant using X-ray crystallography. A representative snapshot using molecular dynamics simulation (right) reveals that the cavity can become significantly wider, offering a potential binding site for small molecules.

According to the paper, a mutation called Y163C turned out to be especially interesting. Although the static image generated using X-ray crystallography seemed to indicate that Y163C would not be a good target because the cavity created by the mutation is very small, MD simulations revealed that the shape of this pocket can become drastically enlarged, providing ample room for a small molecule to bind with the protein. This finding has helped the researchers to focus their energies on Y163C in favor of other mutations that MD and other approaches showed would be more resistant to binding.

Molecular dynamics simulation also provided structural details of the Y163C cavity that could help to identify specific small molecules capable of binding to p53. This information can support virtual drug screening, another approach in computational biology in which catalogs of data describing many thousands of different kinds of small molecules are compared to molecular and structural features of a potential binding site. Such a virtual screen can quickly identify which small molecules have chemical structures that are most likely to be compatible, removing the prohibitively labor-intensive need to test every potential interaction in the laboratory. For those relatively small number of small molecules that show high likelihood of interacting with p53, laboratory experiments can then quickly investigate that much more limited number of computational predictions. And if these experiments are ultimately successful, the findings could then be used to justify clinical trials exploring whether using a small molecule to target Y163C is safe and effective in humans.

There is no guarantee that Y163C will turn out to be an effective target for stopping cancer. The collaboration between Joerger and Settanni demonstrates, however, how molecular dynamics simulation using high-performance computing can make unique contributions to cancer research, accelerating the advance of personalized cancer medicine.

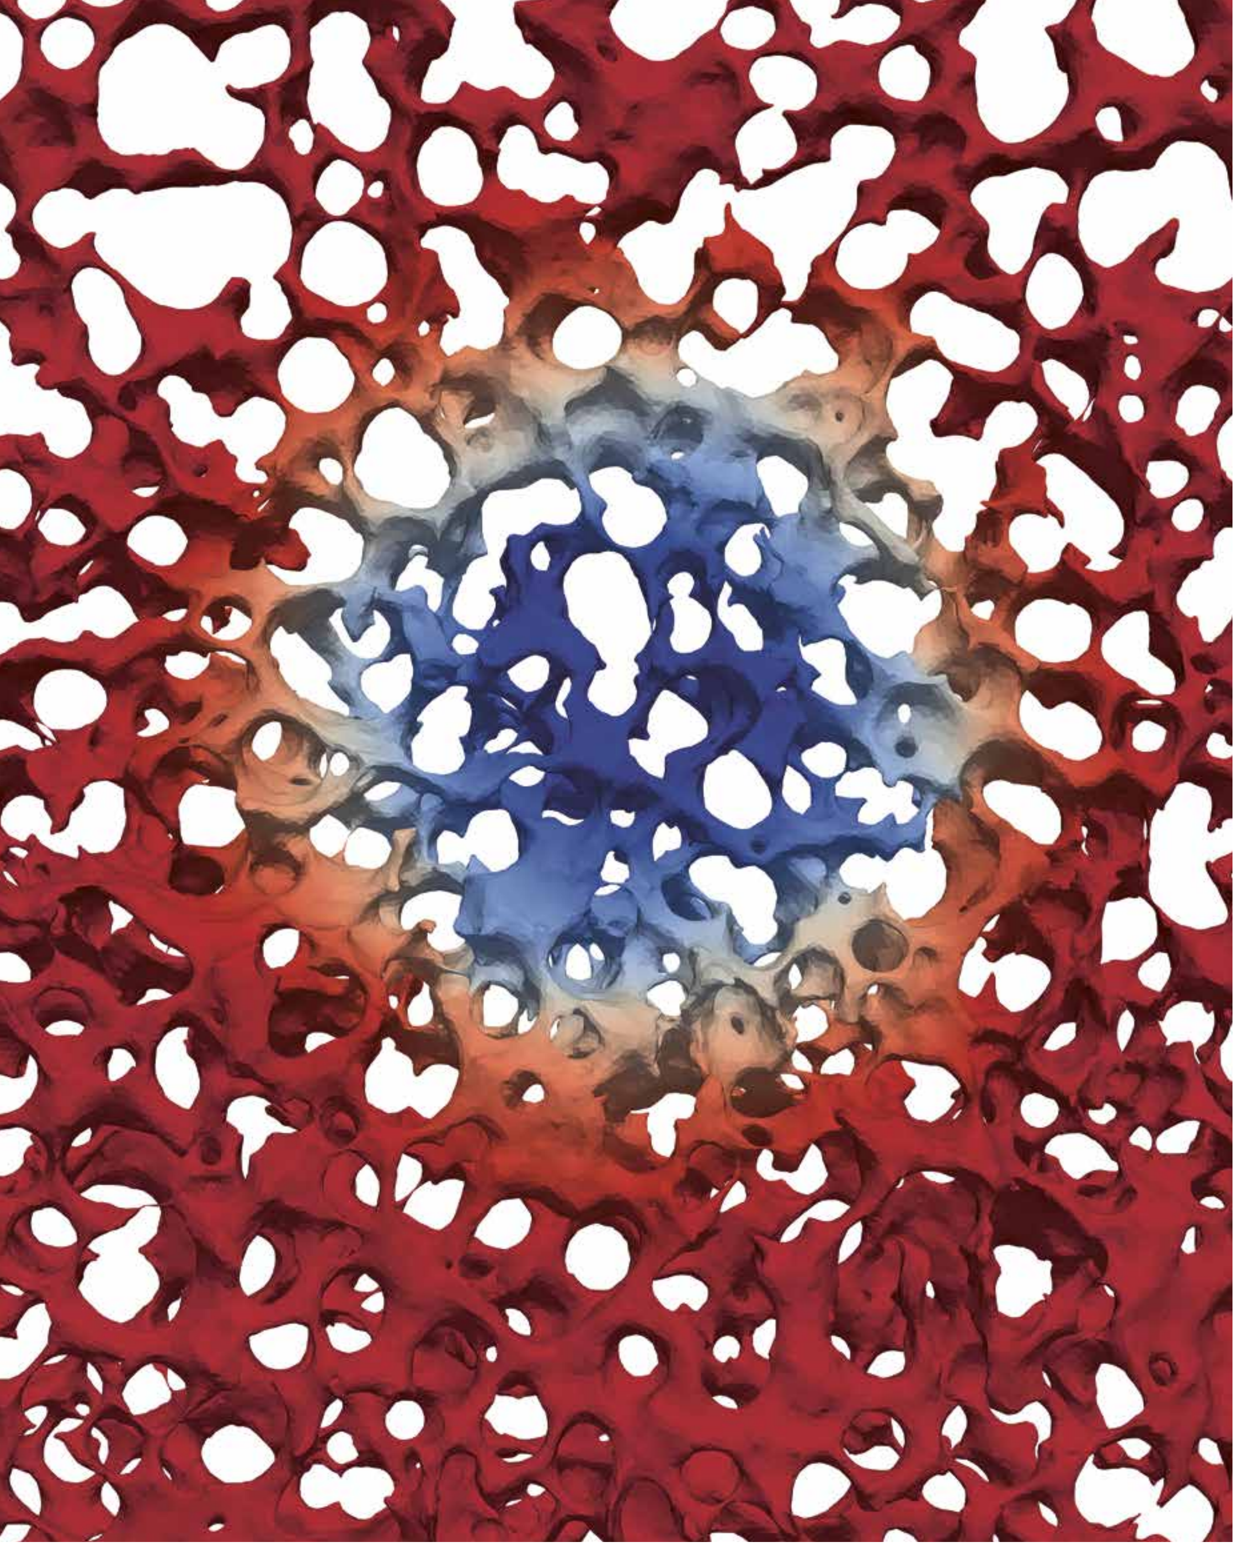

Using HLRS's supercomputer, researchers are developing an "in silico microscope" that would compare experimental results and simulations of biophysical models to better understand how viruses replicate inside cells. This simulation includes color coding that predicts the density of hepatitis C virus surface proteins at the endoplasmic reticulum during infection.

Selected Publications by Our Users in 2024

Albanesi S, Bernuzzi S, Daszuta B, et al. 2024. **INTRHYGUE: Simulations of hyperbolic binary black-hole mergers**. In: Nagel WE, Kröner DH, Resch MM, eds. High Performance Computing in Science and Engineering '22. Springer. 35-48.

Albers M, Shao X, Schröder W. 2024. **Energy efficient actuated drag reduced compressible turbulent flat plate flow**. Int J Heat Fluid Flow. 106: 109314.

Aleksa P, Ghorbani-Asl M, Iqbal S, et al. 2024. **Transition from fractal-dendritic to compact islands for the 2D-ferroelectric SnSe on graphene/Ir(111)**. Nanotechnology. 35(17): 175707.

Ansorge C, Kostelecky J. 2024. **Closing the scale gap for resolved-turbulence simulations in meteorology**. In: Nagel WE, Kröner DH, Resch MM, eds. High Performance Computing in Science and Engineering '22. Springer. 315-335.

Antolovic I, Vrabec J, Klajmon M. 2024. **COSMO-Pharm: Drug–polymer compatibility of pharmaceutical amorphous solid dispersions from COSMO-SAC**. Mol Pharmaceutics. 21(9). 4395–4415.

Appelbaum J, Kloker M, Wenzel C. 2025. **A systematic DNS approach to isolate wall-curvature effects in spatially developing boundary layers**. Theor Comput Fluid Dyn. 39: 10.

Balourdas DI, Markl AM, Krämer A, et al. 2024. **Structural basis of p53 inactivation by cavity-creating cancer mutations and its implications for the development of mutant p53 reactivators.** Cell Death Disease. 15: 408.

Bernhardt F, Gharat S, Kapp A, et al. 2024. **Lattice dynamics of $LiNb_{1-x}Ta_xI_3$ solid solutions: theory and experiment**. Phys Status Solidi A. 2300968.

Bernhardt F, Sanna S. 2024. **Modeling WO_3 beyond the harmonic approximation: theoretical characterization and thermal expansion**. J Phys Chem C. 129(1): 580-590.

Beyer J, Nizenkov P, Fasoulas S, Pfeiffer M. 2024. **Simulation of radiating non-equilibrium flows around a capsule entering Titan's atmosphere**. AIP Conf Proc. 2996: 2000002.

Bhowmik A, Alon OE. 2024. **Interference of longitudinal and transversal fragmentations in the Josephson tunneling dynamics of Bose-Einstein condensates**. New J Phys. 26: 123035.

Bilinskaya Y, Hughes M, Molignini P. 2024. **Realizing multiband states with ultracold dipolar quantum simulators**. Phys Rev Research. 6: L042024.

Blind M, Kopper P, Kempf D, et al. 2024. **Performance improvements for large simulations using the discontinuous Galerkin framework FLEXI**. In: Nagel WE, Kröner DH, Resch MM, eds. High Performance Computing in Science and Engineering '22. Springer. 249-264.

Borsanyi S, Fodor Z, Günther J, et al. 2024. **Bulk features of the quark gluon plasma at finite density**. In: Nagel WE, Kröner DH, Resch MM, eds. High Performance Computing in Science and Engineering '22. Springer. 77-86.

Branch O, Bauer HS. 2024. **WRF simulations to investigate processes across scales (WRFSCALE)**. In: Nagel WE, Kröner DH, Resch MM, eds. High Performance Computing in Science and Engineering '22. Springer. 337-352.

Branch O, Jach L, Schwitalla T, et al. 2024. **Scaling artificial heat islands to enhance precipitation in the United Arab Emirates**. Earth Syst Dynam. 15: 198-129.

Chabanov M, Cruz-Osorio A, Ecker C, et al. 2024. **Microphysical aspects of binary neutron star mergers**. In: Nagel WE, Kröner DH, Resch MM, eds. High Performance Computing in Science and Engineering '22. Springer. 19-34.

Chu X, Wang W, Weigand B. 2024. **An investigation of information flux between turbulent boundary layer and porous medium**. In: Nagel WE, Kröner DH, Resch MM, eds. High Performance Computing in Science and Engineering '22. Springer. 183-196.

Depta PN, Dosta M, Heinrich S. 2024. **Data-driven multiscale modeling of self-assembly and hierarchical structural formation in biological macro-molecular systems: Pyruvate dehydrogenase complex**. In: Nagel WE, Kröner DH, Resch MM, eds. High Performance Computing in Science and Engineering '22. Springer. 355-370.

Depta PN, Dosta M, Heinrich S. 2024. **Multiscale model-based investigation of functional macromolecular agglomerates for biotechnological applications**. In: Kwade A, Kampen I. Dispersity, Structure and Phase Changes of Proteins and Bio Agglomerates in Biotechnological Processes. Springer Nature Switzerland.

Diederich J, Velazquez Rojas J, Paszuk A, et al. 2024. **Ultrafast electron dynamics at the P-rich indium phosphide/TiO_2 interface**. Adv Func Materi. 2409455.

Diederich J, Velasquez Rojas J, Zare Pour MA, et al. 2024. **Unraveling electron dynamics in p-type indium phosphide (100): a time-resolved two-photon photoemission study**. J Am Chem Soc. 146(13): 8649-8960.

Dietrich T, Biswas P, Brügmann B, et al. 2024. **Simulating binary neutron star mergers**. In: Nagel WE, Kröner DH, Resch MM, eds. High Performance Computing in Science and Engineering '22. Springer. 5-18.

El Azzouzi FE, Klimm D, Kapp A, et al. 2024. **Evolution of the electrical conductivity of $LiNb_{1-x}Ta_xO_3$ solid solutions across the ferroelectric phase transformation**. Phys Status Solidi A: 2300966.

Erisis S, Hörning M. 2024. **Self-organization of PIP_3 waves is controlled by the topology and curvature of cell membranes**. Biophys J. 123(9): P1058-1068.

Fekri Z, Chava P, Hlawacek G, et al. 2024. **Tuning the electronic characteristics of monolayer MoS_2-based transistors by ion irradiation: The role of the substrate**. Adv Electron Mater. ePub May 13.

Feldmann H, Hundhausen M, Kohlhepp R, Breil M. 2024. **Impact of land-use change and user-tailored climate change information from a high-resolution climate simulation ensemble**. In: Nagel WE, Kröner DH, Resch MM, eds. High Performance Computing in Science and Engineering '22. Springer. 299-314.

Franzke KL, Schmidt WG, Gerstmann U. 2024. **Relativistic calculation of the orbital hyperfine splitting in complex microscopic structures**. J Phys Conf Ser. 2710: 012094.

Fröde F, Desjardins O, Bieber M, et al. 2024. **Multi-scale simulation of spray and mixture formation for a coaxial atomizer**. Int J Multiphase Flow. ePub Sep 3.

Gagnon L, Donners L, Zimmer F, Keßler M. 2024. **Aerodynamic performance of a hovering cycloidal rotor: a CFD study with experimental validation**. J Amer Helicopter Soc. 69(3): 1-19.

Genuit F, Keßler M, Krämer E. 2024. **Applications of a discontinuous Galerkin chimera method on 3D flow problems**. In: Nagel WE, Kröner DH, Resch MM, eds. High Performance Computing in Science and Engineering '22. Springer. 265-280.

Geppert AK, Stober JL, Steigerwald J, et al. 2024. **Analyzing the early impact dynamics of single droplets impacting onto wetted surfaces**. Phys Fluids. 36: 012005.

Gibis T, Sciacovelli L, Kloker M, Wenzel C. 2024. **Heat-transfer effects in compressible turbulent boundary layers — a regime diagram**. J Fluid Mech. 995: A14.

Gondrum M, Meinke M, Schroeder W, et al. 2024. **Porous fairings for landing gear noise mitigation**. 30th AIAA/CEAS Aeroacoustics Conference. ePub May 30.

Groll M, Bürger J, Caltzidis I, et al. 2024. **DFT-assisted investigation of the electric field and charge density distribution of pristine and defective 2D WSe_2 by differential phase contrast imaging**. Small. 2311635.

Herz MA, Hübner JM, Sichelschmidt J, et al. 2024. **Heavy atom cluster chains with strong spin–orbit coupling and magnetic cations in $Mn(PtBi_6I_{12})$**. Inorg Chem. 63(59): 23867-23876.

Hillebrand M, Breitenstein C, Lutz T. 2024. **Aeroelastic effects of a load alleviation system on a high aspect ratio wing based on CFD-CSM simulations**. AIAA Aviation Forum and Ascend 2024. ePub Jul 27.

Homes S, Antolović I, Fingerhut R, et al. 2024. **High-performance computing as a key to new insights into thermodynamics**. In: Nagel WE, Kröner DH, Resch MM, eds. High Performance Computing in Science and Engineering '22. Springer. 399-413.

Homes S, Vrabec J. 2024. **Resistivities across the vapor–liquid interface of a simple fluid: an assessment of methods**. Phys Fluids. 36: 022122.

Hösgen T, Meinke M, Schröder W. 2024. **Analysis of ingress into the downstream wheel space of a 1.5-stage turbine by full circumference and single blade passage large-eddy simulation**. J Turbomach. 146(9): 091006.

Hristozova L, Sebastian R, Schreyer AM. 2024. **Flow fields around asymmetrical micro vortex generators in supersonic flow**. Aerospace Sci Tech. 145: 108838.

Iskakov S, Katsnelson MI, Lichtenstein AI. 2024. **Perturbative solution of fermionic sign problem in quantum Monte Carlo computations**. npj Comput Mater. 10: 36.

Jafari V, Jarmatz P, Wittenberg H, et al. 2024. **Fault tolerant molecular-continuum flow simulation**. In: Nagel WE, Kröner DH, Resch MM, eds. High Performance Computing in Science and Engineering '22. Springer. 463-475.

Jacobi M, Guercilena FM, Huth S, et al. 2024. **Effects of nuclear matter properties in neutron star mergers**. Mon Not R Astron Soc. 527(3): 8812-8828.

Jain M, Kretschmer S, Meyer J, Krasheninnikov AV. 2024. **Adatom-mediated damage of two-dimensional materials under the electron beam in a transmission electron microscope**. Phys Rev Materials 8: 054004.

Jones T. 2024. **Electro-catalysis for H_2O oxidation and chlorine evolution**. In: Nagel WE, Kröner DH, Resch MM, eds. High Performance Computing in Science and Engineering '22. Springer. 89-100.

Knodel MM, Wittum G, Vollmer J. 2024. **Efficient estimates of surface diffusion parameters for spatio-temporally resolved virus replication dynamics**. Int J Mol Sci. 25(5): 2993.

Koehn H, Wouters T, Rose H, et al. 2024. **Classification of compact objects and model comparison using EOS knowledge**. Phys Rev D. 110: 103015.

Kofahl C, Ganschow S, Bernhardt F, et al. 2024. **Li-diffusion in lithium niobate – tantalate solid solutions**. Solid State Ionics 409: 116514.

Kofahl C, Uhlendorf J, Muscutt BA, et al. 2024. **Oxygen diffusion in Li(Nb,Ta)O_3 single crystals**. Phys Status Solidi A. 2300959.

Kostelecky J, Ansorge C. 2024. **Simulation and scaling analysis of periodic surfaces with small-scale roughness in turbulent Ekman flow**. J Fluid Mech. 992: A8.

Kostelecky J, Ansorge C. 2024. **Turbulent Ekman flow with cubic small-scale roughness under stable stratification (Re=1000). Dataset. Turbulent wall-bounded flows.** Repositorium der Freien Universität Berlin.

Krenz M, Bocchini A, Biktagirov T, et al. 2024. **Polaron formation dynamics in lithium niobite from massively parallel ab-initio simulations**. In: Nagel WE, Kröner DH, Resch MM, eds. High Performance Computing in Science and Engineering '22. Springer. 115-127.

Krenz M, Gerstmann U, Schmidt WG. 2024. **Defect-assisted exciton transfer across the tetracene-Si(111):H interface**. Phys Rev Lett. 132: 076201.

Kuberski S. 2024. Muon g−2: **Lattice calculations of the hadronic vacuum polarization**. Proceedings of Science, Vol 453. The 40^{th} International Symposium on Lattice Field Theory (LATTICE2023).

Kuberski S, Cè M, von Hippel G, et al. 2024. **Hadronic vacuum polarization in the muon g – 2: the short-distance contribution from lattice QCD**. J High Energy Phys. 2024: 172.

Kunert N, Gair J, Pang PTH, Dietrich T. 2024. **Impact of gravitational waveform model systematics on the measurement of the Hubble constant**. Phys Rev D. 110: 043520.

Lagemann C, Lagemann K, Mukherjee S, Schröder W. 2024. **Challenges of deep unsupervised optical flow estimation for particle–image velocimetry data**. Exp Fluids. 65: 30.

Laguarda L, Hickel S, Schrijer FFJ, van Oudheusden BW. 2024. **Reynolds number effects in shock-wave/turbulent boundary-layer interactions**. J Fluid Mech 989: A20.

Lautenschlaeger MP, Weinmiller J, Kellers B, et al. 2024. **Lattice Boltzmann simulation of flow, transport, and reactions in battery components**. In: Nagel WE, Kröner DH, Resch MM, eds. High Performance Computing in Science and Engineering '22. Springer. 449-462.

Li Y, Börrnert F, Ghorbani-Asl M, et al. 2024. **In situ TEM investigation of the lithiation and delitiation process between graphene sheets in the presence of atomic defects**. Adv Funct Mater. 2406034.

Lode AUJ, Alon OE, Bhowmik A, et al. 2024. **Correlations, shapes and fragmentations of ultracold matter**. In: Nagel WE, Kröner DH, Resch MM, eds. High Performance Computing in Science and Engineering '22. Springer. 63-75.

Magistrelli F, Bernuzzi S, Perego A, Radice D. 2024. **Element formation in radiation-hydrodynamics simulations of kilonovae**. Astrophys J Lett. 974(1): L5.

Marx J, Langenbach K, Kohns M. 2024. **Assessment of approaches towards the relative permittivity of mixtures.** J Chem Eng Data. ePub Mar 28.

Menczer A, Legeza Ö. 2024. **Tensor network state algorithms on AI accelerators**. J Chem Theory Comput. 20(20): 8897-8910.

Molignini P, Chakrabarti B. 2024. **Unbounded entropy production and violent fragmentation for repulsive-to-attractive interaction quench in long-range interacting systems.** New J Phys. ePub Sep 27.

Mondal B, Khan RU, Kreuter F, et al. 2024. **Organic functionalization on solid surfaces.** In: Nagel WE, Kröner DH, Resch MM, eds. High Performance Computing in Science and Engineering '22. Springer. 101-113.

Moriaux OK, Zamponi R, Satcunanathan S, et al. 2024. **Experimental and numerical investigation of the turbulent boundary layer of a grazing flow over porous materials**. 30^{th} AIAA/CEAS Aeroacoustics Conference. ePub May 30.

Munoz Lopez EJ, Hergt A, Klinner J, et al. 2024. **The unsteady shock-boundary layer interaction in a compressor cascade – part 3: mechanisms of shock oscillation**. Proceedings of ASME Turbo Expo 2024, June 24-28, 2024. London, England.

Müller J, Lutz T. 2024. **Numerical studies on the impact of atmospheric turbulence on aircraft loads**. In: Heinrich R, ed. Advanced Aircraft Understanding via the Virtual Aircraft Model. Springer. 99-119.

Neuweiler A, Dietrich T, Brügmann B, et al. 2024. **General relativistic magnetohydrodynamic simulations with BAM: implementation and code comparison**. Phys Rev D. 110: 084046.

Palmetshofer P, Geppert AK, Steigerwald J, et al. 2024. **Thermocapillary central lamella recess during droplet impacts onto a heated wall.** Sci Reports. 14: 1102.

Palmetshofer P, Wurst J, Geppert AK, et al. 2025. **Wetting behavior in the inertial phase of droplet impacts onto sub-millimeter microstructured surfaces.** J Colloid Interface Sci. 682: 413-422.

Panos E, Hassan A. 2024. **Accelerating the performance of large-scale TIMES models in the modelling of Sustainable Development Goals**. In Labriet et al, eds. Aligning the Energy Transition with the Sustainable Development Goals: Key Insights from Energy System Modelling. Springer. 67-95.

Paraschos GF, Kim JY, Wielgus M, et al. 2024. **Ordered magnetic fields around the 3C 84 central black hole.** Astron Astrophys. 682: L3.

Peter JMF, Gibis T, Kloker MJ. 2024. **Favorable-pressure-gradient influence on supersonic film cooling with turbulent main flow.** In: Nagel WE, Kröner DH, Resch MM, eds. High Performance Computing in Science and Engineering '22. Springer. 167-182.

Pfannstiel A, Hehemann T, Schäfer NA, et al. 2024. **Small electron polarons bound to interstitial tantalum defects in lithium tantalate**. J Phys Condens. Matter 36(35): 355701.

Pinzon J, Siebenborn M, Vogel A. 2024. **Geometric constrained scalable algorithm for PDE-constrained shape optimization.** In: Nagel WE, Kröner DH, Resch MM, eds. High Performance Computing in Science and Engineering '22. Springer. 415-428.

Pionteck MN, Bernhardt F, Dues C, et al. 2024. **Hyperpolarizabilities of $LiNbO_3$, $LiTaO_3$ and $KNbO_3$ calculated from first principles.** In: Nagel WE, Kröner DH, Resch MM, eds. High Performance Computing in Science and Engineering '22. Springer. 129-143.

Pionteck MN, Sanna S. 2024. **Photoelastic properties of stoichiometric lithium niobate from first-principles calculations**. Phys Status Solidi A. ePub Mar 11.

Pollinger T, Van Craen A, Offenhäuser P, Pflüger D. 2024. **Realizing joint extreme-scale simulations on multiple supercomputers – two superfacility case studies**. In: 2024 SC24: International Conference for High Performance Computing, Networking, Storage and Analysis.

Potyka J, Schulte K. 2024. **A volume of fluid method for three dimensional direct numerical simulations of immiscible droplet collisions**. Int J Multiphase Flow. 170: 104654.

Potyka J, Stober J, Wurst J, et al. 2024. **Towards DNS of droplet-jet collisions of immiscible liquids with FS3D**. In: Nagel WE, Kröner DH, Resch MM, eds. High Performance Computing in Science and Engineering '22. Springer. 197-212.

Rentschler T, Berkemeier MB, Fraas S, et al. 2024. **Multi-criteria hydraulic turbine optimization using a genetic algorithm and trust-region postprocessing**. Proc Appl Math Mech. 202400126.

Roccon A, Zonta F, Soldati A. 2024. **Turbulent drag reduction in water-lubricated channel flow of highly viscous oil**. Phys Rev Fluids. 9: 054611.

Rodekamp M, Engelhardt M, Green JR, et al. 2024. **Moments of nucleon unpolarized, polarized and transversity patron distribution functions from lattice QCD at the physical point**. Phys Rev D. 109: 074508.

Ruiz Alvarado IA, Dreßler C, Schmidt WG. 2025. **Band alignment at InP/$PiPO_2$ interfaces from density-functional theory**. J Phys Condens Matter. 37: 075001. ePub Dec 5, 2024.

Santra P, Ghaderzadeh S, Ghorbani-Asl M, et al. 2024. **Strain-modulated defect engineering of two-dimensional materials**. npj 2D Mater Appl. 8: 33.

Schaefer D, Kunstmann B, Schmitt S, et al. 2024. **Explosions of nanodroplets studied with molecular dynamics simulations**. Phys Fluids. 36: 037129.

Schianchi F, Ujevic M, Neuweiler A, et al. 2024. **Black-hole formation in binary neutron star mergers: the impact of spin on the prompt-collapse scenario**. Phys Rev D. 109; 123011.

Schubert S, Steigerwald J, Geppert AK, et al. 2024. **Micro-PIV study on the influence of viscosity on the dynamics of droplet impact onto a thin film**. Exp Fluids. 65: 69.

Sciotto R, Ruiz Alvarado IA, Schmidt WG. 2024. **Substrate doping and defect influence on P-rich InP(001):H surface properties**. Surfaces. 7(1): 79-87.

Sebastian R, Schreyer AM. 2024. **Design considerations for efficient spanwise-inclined air-jet vortex generators for separation control in supersonic and hypersonic flows.** Aerosp Sci Technol. 147: 109033.

Seitz T, Gerlinger P. 2024. **Investigation and assessment of production- and dissipation-limited DDES methods for rocket combustion chamber simulations**. Comput Fluids. 275: 106252.

Seiz M, Hierl H, Nestler B, Rheinheimer W. 2024. **Revealing process and material parameter effects on densification via phase-field studies**. Sci Rep-UK. 14: 5350.

Settanni G, Schmid F. 2024. **Molecular dynamics simulations of the structure of lipid-based nanomaterials**. In: Nagel WE, Kröner DH, Resch MM, eds. High Performance Computing in Science and Engineering '22. Springer. 49-51.

Sokolov M, Mastrikov YA, Bocharov D, et al. 2024. **Computational study of oxygen evolution reaction on flat and stepped surfaces of strontium titanate**. Catalysis Today. 114609, ePub Mar 2.

Strobel KR, Schlegel M, Jain M, et al. 2024. **Temperature-dependence of beam-driven dynamics in graphene-fullerene sandwiches.** Micron. 184: 103666.

Su Y, Daric D, Guevara-Carrion G, et al. 2025. **Fick and Maxwell-Stefan diffusion of the liquid mixture cyclohexane + toluene + acetone + methanol and its subsystems**. Chem Eng Sci. 301: 120662.

Tan Q, Hosseini SA, Seidel-Morgenstern, et al. 2024. **Thermal effects connected to crystallization dynamics: a lattice Boltzmann study**. Int J Multiphase Flow. 171: 104669.

Tan Q, Hosseini SA, Thévenin D. 2024. **Simulations of crystal growth using lattice Boltzmann formulation**. In: Nagel WE, Kröner DH, Resch MM, eds. High Performance Computing in Science and Engineering '22. Springer. 387-398.

Tandiana R, Barletta GP, Soler MA, et al. 2024. **Computational mutagenesis of antibody fragments: disentangling sidechains from ΔΔG predictions.** J Chem Theory Comput. 20(6): 2630-2642.

Topolski K, Tootle SD, Rezzolla L. 2024. **Post-merger gravitational-wave signal from neutron-star binaries: a new look at an old problem.** Astrophys J. 960(1): 86.

Tuppurainen V, Fleitmann L, Kangas J, et al. 2024. **Conceptual design of furfural extraction, oxidative upgrading and product recovery: COSMO-RS-based process-level solvent screening**. Comput Chem Eng. 191: 108835.

Ullah J, Hillebrand M, Ehrle M, et al. 2024. **On the effects of wing-gust interactions and wing flap deflections on the HTP aerodynamics.** In: Nagel WE, Kröner DH, Resch MM, eds. High Performance Computing in Science and Engineering '22. Springer. 213-232.

Verhoff LM, Pionteck MN, Rüsing M, et al. 2024. **Two-dimensional electronic conductiity in insulating ferroelectrics: peculiar properties of domain walls.** Phys Rev Research. 6: L042015.

Wegmann T, Meinke M, Schröder W. 2024. **Dynamic load balancing of a coupled Lagrange particle tracking solver for direct injection engine application**. In: Resch M, Gebert J, Kobayashi H, Takizawa H, Baz W, eds. Sustained Simulation Performance 2022: Proceedings of the Joint Workshop on Sustained Simulation Performance. Springer. 41-60.

Xu X, Zhang X, Ruban A, et al. 2024. **Accurate complex-stacking-fault Gibbs energy in Ni_3Al at high temperatures**. Scripta Materialia 242: 115934.

Yakhnevych U, Sargsyan V, El Azzouzi F, et al. 2024. **Acoustic loss in $LiNb_{1-x}Ta_xO_3$ at temperatures up to 900 °C**. Physica status solidi.

Yang Z, Meinke M, Schroeder W. 2024. **Numerical analysis of a propeller-airfoil interaction in a distributed propulsion system using a hybrid LES and FW-H approach.** 30th AIAA/CEAS Aeroacoustics Conference. ePub May 30.

Ye Q, Shen B, Tiedje O, et al. 202. **On charging and breakup of paints using a high-speed rotary bell atomizer with internal charging system**. Int J Multiphase Flow. 180: 104955.

Zhang F, Tavakkol S, Dercho S, et al. 2024. **Assessment of dynamic characteristics of fluidized beds via numerical simulations**. Phys Fluids. 36: 023348.

Ziese F, Sanna S. 2024. **Nonlinear optical response of tetrel-modified tetraphenyl-adamantane clusters**. ACS Omega. 9(50): 49816-49824.

Ziese F, Wang J, Rojas León I, et al. 2024. **Origin of the nonlinear optical response in organotretel molecules, (hetero)adamantane-type clusters with organic substituents, and related species**. J Phys Chem A. 128(39): 8360-8372.

Zimmer DN, Schmid F, Settanni G. 2024. **Ionizable cationic lipids and helper lipids synergistically contribute to RNA packing and protection in lipid-based nanomaterials**. J Phys Chem B. 128(41).

HLRS Books

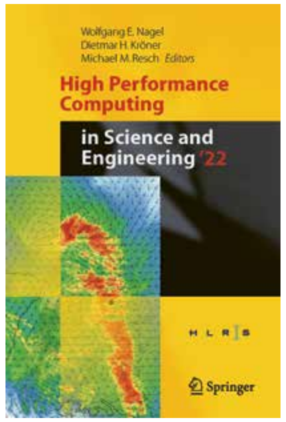

High Performance Computing in Science and Engineering ‘22

Editors: Wolfgang E. Nagel, Dietmar H. Kröner, Michael M. Resch

This book presents the state-of-the-art in supercomputer simulation, focusing on findings from leading researchers using systems at the High-Performance Computing Center Stuttgart (HLRS) in 2022. The reports cover all fields of computational science and engineering, ranging from CFD to computational physics and from chemistry to computer science, with a special emphasis on industrially relevant applications. Presenting findings using one of Europe’s leading high-performance computing systems, the book covers the main methods in high-performance computing. Outstanding results in achieving top performance for production codes will be of interest for both scientists and engineers.

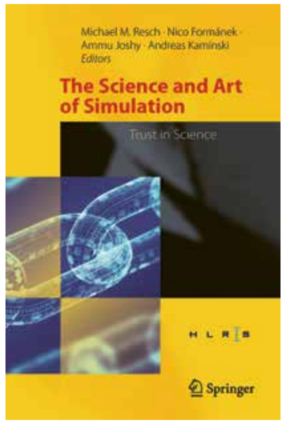

The Science and Art of Simulation: Trust in Science

Editors: Michael M. Resch, Nico Formanek, Ammu Joshy, Andreas Kaminski

Trust is a central pillar of the scientific enterprise. Much work in the philosophy of science can be seen as coping with the problem of establishing trust in a certain theory, a certain model, or even science as a whole. However, trust in science is threatened by various developments. With the advent of more complex models and the increasing usage of computer methods such as machine learning and computer simulation, it seems increasingly challenging to establish trust in science. How and on what basis can an appropriate trust in science be built? We are interested in how trust is established in such cases of increasing complexity (of models and communication) and what could be appropriate measures to alleviate doubt.

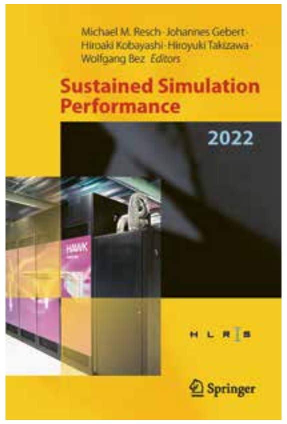

Sustained Simulation Performance 2022

Editors: Michael M. Resch, Johannes Gebert, Hiroaki Kobayashi, Wolfgang Bez

This book presents the state of the art in high-performance computing on modern supercomputer architectures. It addresses trends in hardware and software development in general. The contributions cover a broad range of topics, from performance evaluations in the context of power efficiency to computational fluid dynamics and high-performance data analytics. In addition, they explore new topics like the use of high-performance computers in the field of artificial intelligence and machine learning. Contributions are based on selected papers presented in 2022 at the 33rd Workshop on Sustained Simulation Performance (WSSP33) held at HLRS in Stuttgart and WSSP34, held at Tohoku University in Sendai, Japan.

HLRS by Numbers

150 Staff

112 Scientists

35 Nonscientists

2 Student Assistants

1 Research Assistant

System Usage

4.634 billion
Total Core Hours Produced

56
Talks by HLRS Staff

Research

4.275 billion
Core Hours

121 User Projects

115 User Publications

Industry

309 million
Core Hours

74 Industrial Customers

Staff Publications

34 Papers in Journals, Books, and Conference Proceedings

3 Books

Education and Training

46 Continuing Education Courses	11 Scientific Workshops	16 University Lectures
1,354 Participants 182 Course Days	518 Participants 25 Days	1,213 Participants 36 SWS

1,817
Vistors at HLRS

Third-Party Funds

7,100,000 €

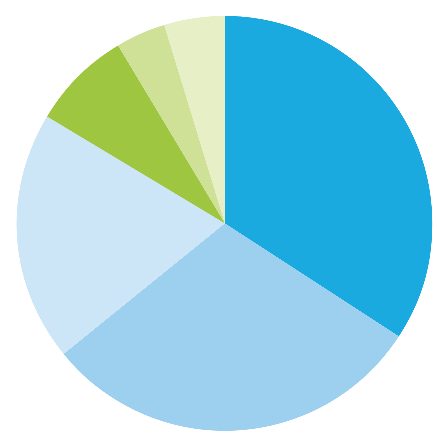

- EU 34.2 %
- State Gov./Land 30.1 %
- Federal Gov./Bund 19.5 %
- DFG 7.7 %
- Foundations 3.9 %
- Industry 4.6 %

About Us

Inside Our Computing Room

HPE Cray EX4000 (Hunter)

Hunter is HLRS's newest flagship supercomputer, going into service in February 2025. Based on the same technology as the world's three verified exascale systems, it enables large-scale, state-of-the-art applications of simulation, artificial intelligence, and data analytics, including hybrid computing approaches that combine diverse methods into powerful workflows. At its core is the AMD Instinct™ MI300A accelerated processing unit (APU), which combines CPUs, GPUs, and high-bandwidth memory in a single package. With a theoretical peak performance of 48.1 petaflops, Hunter's speed is nearly double that of its predecessor, Hawk, while slashing energy requirements at peak performance by approximately 80%. Hunter is conceived as a stepping-stone system that will enable system users to prepare for HLRS's next flagship supercomputer, Herder, which is scheduled to arrive in 2027.

Processors
APUs: 752 AMD Instinct MI300A
CPUs: 512 AMD EPYC 9374F

Memory technology:
APU: 512 GB HBM3 (~5.3 TB/s)
CPU: 768 GB DDR5-4800

Storage:
Cray ClusterStor E2000 (25 PB capacity on 2,120 disks)

Networking:
HPE Slingshot 11 Dragonfly (APU: 4 x 200 Gbps per node)

System peak performance:
48.1 petaflops

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

Baden-Württemberg Ministry of Science, Research and Arts

Federal Ministry of Research, Technology and Space

HPE Apollo 6500 Gen10 Plus (Hawk AI Expansion)

This partition was initially installed in 2021 as part of HLRS's previous flagship system, Hawk. Although Hawk has been taken out of service, the AI expansion continues to offer NVIDIA GPUs suitable for machine learning and artificial intelligence applications.

GPU Type: NVIDIA A100
Number of GPUs: 192
Performance: 120 petaflops AI performance

NEC Cluster (Vulcan)

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

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

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

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

AMD Genoa
CPU: AMD Genoa EPYC 9334 @ 2.7 GHz base, 32 cores
Number of nodes: 156
Memory per node: 768 GB

NVIDIA A30
Number of nodes: 24
CPU: 2 × AMD EPYC 9124 Genoa, 3.0 GHz base, 16 core
Memory per node: 768 GB DDR5
CPU: 1 × NVIDIA A30
CPU memory: 24 GB HBM2e

Large Memory Nodes
Number of nodes: 2
CPU: 2 × AMD EPYC 9124 Genoa, 3.0 GHz base, 16 core
Memory per node: 3,072 GB DDR5

Interconnects
Infiniband EDR / HDR

User Profile 2024

The following is a summary of usage of HLRS's supercomputer in calendar year 2024.

Gauss Centre for Supercomputing (GCS)

Awards for large-scale projects *	
Number of new large-scale projects (>100 million core hours)	11
Hawk (million core hours)	1,320
Hunter (thousand node hours)	394

GCS projects – system usage in 2024	
Number of active GCS projects	121
Core hours (millions)	4,275
AI node hours (thousands)	31

Other scientific usage	
Core hours (millions)	49.84

Usage by Industry

HLRS Industrial Clients	
Number of industry clients	74
Active users in 2024	43
Core hours (millions)	87.07

HWW Usage	
Core hours (millions)	222.14

Total Industrial Usage	
Core hours (millions)	309.21

Total System Usage in 2024	
Core hours (millions)	4,634
Percentage science	93,3 %
Percentage industry	6,7 %

System Usage by Scientific Discipline

CFD	53,8 %
Physics	38,6 %
Chemistry	3,6 %
Transportation and Climate	1,5 %
Reactive Flows	0,8 %
Solid State Physics	0,6 %
Materials Simulation & Materials Research	0,5 %
Computer Science	0,2 %
Other	0,2 %
Electrical Engineering	0,1 %

System Usage by State

North Rhine-Westphalia	45,91 %
Baden-Württemberg	30,58 %
Rheinland-Palatinate	7,91 %
Hesse	5,68 %
Brandenburg	3,84 %
Bavaria	2,52 %
Thuringia	1,54 %
Berlin	0,74 %
Federal Research Center	0,57 %
Saxony	0,33 %
Hamburg	0,23 %
Saxony-Anhalt	0,15 %

* Please note that only Hawk was in operation in 2024 and that grants for its usage were awarded in core hours. The second large-scale call in 2024 was for usage of Hunter, which would only go into service in 2025. Computing time grants for Hunter are calculated in node hours.

Third-Party Funded Research Projects

In addition to providing supercomputing resources for scientists and engineers in academia and industry, HLRS conducts its own funded research on important topics relevant for high-performance computing (HPC), artificial intelligence, visualization, and high-performance data analytics. These activities, many of which are conducted in collaboration with investigators at other institutes and in industry, address key problems facing supercomputing and are opening up new opportunities for addressing key German, European, and global challenges. The following is a list of funded projects in 2024.

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

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

AI Alliance BW
August 2023 – March 2026 MWK
The AI Alliance Baden-Württemberg aims to develop a data platform for AI-relevant data exchange among academic and industrial stakeholders.

BEGIN HPC+ ■
May 2024 – April 2027 MWK
Supplementary funding for the implementation of EuroCC 2 and CASTIEL 2 to support activities that have not been financed by the EU/BMBF, such as for dissemination events organized by the German National Competence Center and proof of concept studies of the German NCCs, as well as for additional personnel necessary for the technical and scientific management of the two projects.

bwHPC-S5
July 2018 – January 2026 MWK
Coordinates support for HPC users in Baden-Württemberg and the implementation of related measures and activities, including data intensive computing and large-scale scientific data management.

CASTIEL 2
January 2023 – December 2025 JU
CASTIEL 2 facilitates collaboration among the EuroCC 2 National Competence Centers and the EuroHPC Joint Undertaking Centers of Excellence, promoting the development of HPC expertise and the adoption of leading codes across Europe.

CEEC
January 2023 – December 2026 JU
The Center of Excellence in Exascale CFD will improve European state-of-the-art computational fluid dynamics algorithms to prepare them for efficient performance on exascale supercomputers.

ChEESE-2P
January 2023 – December 2026 JU
Focusing on critical applications for the prediction of geohazards, the Centre of Excellence for Exascale in Solid Earth aims to become a hub for HPC software within the solid earth community.

CIRCE
November 2021 – April 2025 BMBF, MWK
A study to assess potential applications of high-performance computing (HPC) in crisis situations, and what organizational procedures are needed to ensure that HPC resources are immediately available.

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

DEGREE
June 2021 – April 2024 DBU
DEGREE investigated a method for increasing energy efficiency in data centers by dynamically controlling cooling circuit temperatures, and developing guidelines for implementing the resulting concepts.

DimensionLab3
September 2022 – September 2024 SIH
By providing exciting, immersive experiences of current topics in engineering in virtual and augmented reality, DimensionLab3 aims to increase young students' interest in and motivation to study technical subjects.

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

EuroCC 2
January 2023 – December 2025 EU
Supported by the EuroHPC Joint Undertaking, EuroCC 2 manages a European network of National Competence Centers (NCC) for high-performance computing and related technologies, promoting a common level of expertise across the participating countries.

EuroHyPerCon
October 2023 – June 2024 EU
EuroHyPerCon completed a study aimed at defining a long-term hyperconnectivity specification and implementation roadmap to meet Europe's future ultra-high-speed network requirements.

EVITA ■
April 2025 – March 2029 JU
EVITA aims to establish a unified, high-quality training framework for HPC and emerging technologies across Europe, offering recognized qualifications and modular learning content to support academia, industry, and innovation.

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

EXCELLERAT P2
January 2023 – December 2026 JU
EXCELLERAT P2 is developing advanced applications for engineering in the manufacturing, energy, aeronautics, and automotive sectors, focusing on use cases that demonstrate the importance of HPC, HPDA, and AI for European competitiveness.

FFplus ■
May 2024 – April 2028 JU
FFplus supports European small and medium-sized enterprises (SMEs) and startups in testing new applications of high-performance computing and artificial intelligence.

■ New project ■ Grant awarded, starts in 2025

Funder Abbreviations:
BMBF – Federal Ministry of Education and Research (Effective May 2025, the BMBF is now the German Federal Ministry of Research, Technology and Space) | CZS – Carl Zeiss Foundation | DBU – German Federal Environmental Foundation | DFG – German Research Foundation | EC – European Commission | EU – European Union | ICM – InnovationsCampus Mobilität der Zukunft | JU – EuroHPC Joint Undertaking | MSA – Bavarian State Ministry of Science and the Arts | MSC – Lower Saxony Ministry of Science and Culture | MWK – Baden-Württemberg Ministry for Science, Research, and Art | SIH – Stiftung Innovation in der Hochschullehre | WAT –Baden-Württemberg Ministry of Economic Affairs, Labor and Tourism

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

HammerHAI ■
April 2025 – March 2028
JU, BMBF, MWK, MSA, MSC
HammerHAI is building a state-of-the-art "AI Factory" that will provide European businesses and researchers with secure, scalable, and easily accessible AI resources and support, lowering barriers needed to deploy AI-driven solutions.

HANAMI ■
March 2024 – February 2027   JU
HANAMI fosters collaboration between Europe and Japan to develop applications for future generations of supercomputers across diverse scientific fields, including environmental sciences, biomedicine, and materials science.

HiDALGO 2
January 2023 – December 2026   BMBF, JU
HiDALGO2 is addressing challenges caused by climate change, focusing on technical issues related to scalability on HPC and AI infrastructures, the use of computational fluid dynamics methods, and uncertainty analysis.

HPC SPECTRA ■
January 2024 – December 2025   JU, EU
HPC SPECTRA will promote the development of HPC expertise across Europe by building a comprehensive online platform of training opportunities, making it easy for trainees to find courses that fit their interests and needs.

IKILeUS
December 2021 – November 2024   BMBF
HLRS coordinated this project to integrate AI topics and technologies at the University of Stuttgart.

InHPC-DE
November 2017 – December 2026   BMBF
InHPC-DE furthers the federation of Germany's three national HPC centers, addresses new requirements such as security, and evaluates the Gaia-X ecosystem in the context of high-performance computing.

Inno4scale
July 2023 – June 2025   JU
Inno4scale will identify and provide funding to support the development of advanced algorithms and applications for upcoming European exascale systems.

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

MERIDIONAL
October 2022 – September 2026   EU
This project is developing a tool for assessing the performance and loads experienced by onshore, offshore, and airborne wind energy systems.

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

ORCHESTRA
December 2020 – November 2024   EU
ORCHESTRA developed a networked platform for sharing data and for creating a new large-scale, pan-European cohort for research on the SARS-CoV-2 pandemic, providing a model for addressing future public health threats.

POP3 ■
November 2024 – December 2026   EU, JU
The Performance Optimization and Productivity Centre of Excellence in HPC provides services for academic and industrial codes to assess and improve their efficiency on large-scale computing systems.

S+T+ARTS AIR
April 2023 – November 2024   EU, MWK
S+T+ARTS AIR made supercomputing technologies and expertise available to enable innovative collaborations involving the arts, science, and technology.

S+T+ARTS E(C)HO ■
January 2024 – December 2026   EU
This artist residency program supports collaboration between artists and experts in science and technology, focusing on promoting ecologically conscious and human-compatible digital technologies.

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

SimTech
July 2019 – June 2025   DFG
An interdisciplinary Excellence Cluster at the University of Stuttgart has been developing simulation technologies to enable integrative systems science. HLRS supported the development of efficient methods for uncertainty quantification and management.

Simulated Worlds
January 2011 – August 2028   MWK
Offers students opportunities to develop and execute simulation projects in collaboration with HLRS scientists.

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

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

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

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

Trust in Information
August 2020 – June 2024   MWK
Multidisciplinary research led by the HLRS Department of Philosophy that developed perspectives for assessing the trustworthiness of computational science and limiting the spread of misinformation.

URBANOME
March 2021 – August 2025   EU
URBANOME aims to promote health, well-being, and quality of life in cities by systematically integrating health considerations in city politics and the activities of city residents.

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

HPC Training Courses in 2024

HLRS offered 46 courses in 2024, providing continuing professional education on a wide range of topics relevant for high-performance computing. The courses took place over 182 course-days, online and in Stuttgart and in cooperation with other institutes in Germany and internationally. A total of 1,354 trainees participated in these activities.

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

Date	Location	Topic	Host
Jan 12–Feb 19	online/self-paced	HPC Cluster Design, Cost & Sustainability	HLRS (SCA)
Jan 23–25	Stuttgart/online	Hybrid Programming in HPC - MPI+X * NEW	HLRS
Feb 5–9	online	Introduction to Computational Fluid Dynamics	ZIH/HLRS/DLR
Feb 5–Mar 1	online/self-paced	Visualization - Basics & Application	HLRS (SCA)
Feb 19–23	online	Iterative Solvers and Parallelization	HLRS
Mar 5–8	Stuttgart/online	Modern C++ Software Design (Intermediate)	HLRS
Mar 11–15	Stuttgart	CFD with OpenFOAM(R)	HLRS
Mar 11–15	Dresden	Parallel Programming with MPI & OpenMP and Tools	ZIH/HLRS
Mar 26–28	Stuttgart	From Machine Learning to Deep Learning: a concise introduction	HLRS
Apr 3–5	online	N-Ways to GPU Programming Bootcamp *	EuroCC/NVIDIA
Apr 8–May 3	online/self-paced	Data Management	HLRS (SCA)
Apr 15–26	online/self-paced	Parallel Programming with OpenMP	HLRS (SCA)
Apr 16–19	Mainz	Parallelization with MPI and OpenMP	ZDV/HLRS
Apr 22–25	online	AMD Instinct GPU Training	HLRS/AMD
May 6–7	online	Multi-GPU Programming Bootcamp * NEW	EuroCC/NVIDIA
May 13–17	Stuttgart/online	Fortran for Scientific Computing	HLRS
May 21–Jun 14	online/self-paced	IKILeUS: Natural Language Processing NEW	HLRS (SCA)
May 27–28	online	Efficient Parallel Programming with GASPI	HLRS/F.ITWM
Jun 3–Jul 12	online/self-paced	HPC Cluster Setup & Operation	HLRS (SCA)
Jun 5–7	Stuttgart/online	Hackathon by HiDALGO2, CIRCE and SEAVEA projects NEW	HLRS/HiDALGO2/CIRCE/SEAVEA
Jun 10–19	online/self-paced	Simulation Basics & CFD	HLRS (SCA)
Jun 11–14	Stuttgart/online	Modern C++ Software Design (Advanced)	HLRS
Jun 18–21	online	Node-Level Performance Engineering	HLRS/NHR@FAU
Jun 25–26	online	AI for Science Bootcamp *	EuroCC/NVIDIA
Jul 2–5	Stuttgart	Multi-GPU Deep Learning NEW	HLRS
Jul 9–12	Stuttgart	Modern C++ Software Design (Intermediate)	HLRS
Jul 15–19	Stuttgart/online	Hunter Code Preparation Workshop NEW	HLRS
Aug 21–30	online	Six-day Course in Parallel Programming with MPI/OpenMP	ETH/HLRS
Sep 10–13	Stuttgart	Julia for High-Performance Computing	HLRS
Sep 16–20	Stuttgart	Introduction to Computational Fluid Dynamics	HLRS/DLR
Sep 17–18	Garching	Iterative Solvers for Linear Systems	LRZ/HLRS
Sep 23–31	online/self-paced	Simulation Basics & Structural Mechanics	HLRS (SCA)
Sep 23–25	online	Introduction to oneAPI, SYCL2020 and OpenMP offloading	HLRS/INTEL
Oct 7–11	Stuttgart	CFD with OpenFOAM(R)	HLRS
Oct 14–18	Stuttgart	Parallel Programming Workshop (with TtT)	HLRS
Oct 14–31	online/self-paced	IKILeUS: Time Series Forecasting NEW	HLRS (SCA)
Oct 22	online	Introduction to OpenMP Offloading with AMD GPUs NEW	HLRS/AMD
Oct 23–25	Dresden/online	alpaka and openPMD Workshop and Hackathon NEW	HZDR/Plasma-PEPSC/ENCCS/HLRS
Oct 28–29	Stuttgart	Scientific Visualization	HLRS
Nov 4–6	online/self-paced	Parallel Programming with MPI NEW	HLRS (SCA)
Nov 4–Dec 13	online/self-paced	Data Analysis with HPC	HLRS (SCA)
Nov 4–8	online	Introduction to GPU Programming using CUDA NEW	HLRS
Nov 11–15	Stuttgart/online	Hunter Code Preparation Workshop NEW	HLRS
Nov 19–22	Stuttgart/online	Modern C++ Software Design (Advanced)	HLRS
Dec 2–5	online	Advanced Parallel Programming with MPI and OpenMP	JSC
Dec 9–13	online	Fortran for Scientific Computing	HLRS

- Parallel Programming
- Simulation and Computational Fluid Dynamics (CFD)
- Performance Optimization and Debugging
- Data in HPC, Deep Learning, Machine Learning and AI
- Programming Languages for Scientific Computing
- Scientific Visualization
- Compute Cluster – Usage and Administration
- Community- and domain-specific content
- Hardware Accelerators

* EuroCC2 courses: HLRS is a member of the Gauss Centre for Supercomputing (GCS). EuroCC@GCS is the German National Competence Centre (NCC) for High-Performance Computing. This course is provided within the framework of EuroCC2..

TtT: Train the Trainer Courses | NEW: New or significantly modified course

Institute Abbreviations

AMD – Advanced Micro Devices, Inc. | CIRCE – Computational Immediate Response Center for Emergencies | DLR – German Aerospace Center | ENCCS – EuroCC National Competence Center Sweden | ETH – Scientific IT Services, ETH Zurich | EuroCC – EuroCC | F. ITWM – Fraunhofer Institute for Industrial Mathematics | HiDALGO 2 – HiDALGO Centre of Excellence | HLRS – High-Performance Computing Center Stuttgart | HLRS (SCA) – Supercomputing Academy of the HLRS | HZDR – Helmholtz-Zentrum Dresden-Rossendorf | INTEL – Intel Corporation | JSC – Jülich Supercomputing Centre | LRZ – Leibniz Supercomputing Centre | NHR@FAU – Erlangen National High Performance Computing Center | NVIDIA – Nvidia Corporation | Plasma-PEPSC – Plasma-PEPSC Centre of Excellence | SEAVEA – Software Environment for Actionable & VVUQ-Evaluated Exascale Applications | ZDV – Data Center, University of Mainz | ZIH – Center for Information Services and High Performance Computing (TU Dresden)

Workshops and Conferences in 2024

Apr 17–19 **ReACH Advanced School**
Organized by the Media Solution Center Baden-Württemberg, this kickoff event for the initiative Research and Creation Center for eCulture and the Humanities (ReACH) focused on applications of artificial intelligence and virtual reality in the arts and creative fields.

Apr 25 **Science Goes Society**
An event co-organized with the town of Sersheim focused on new approaches that communities can use for risk assessment, catastrophe prevention, and environmental preservation in the face of climate change.

Jun 11 **How Computer Simulation Can Help the Public Sector in Crisis Situations**
This conference focused on the findings of CIRCE, a project that investigated how HPC could support emergency prevention and management.

Jun 17–18 **37th Workshop on Sustained Simulation Performance**
Organized with NEC and Tohoku University, this annual meeting focuses on hardware architectures, programming styles, and strategies for achieving the highest sustained application performance.

Jul 16 **GCS Industry Event – Success in Innovation Using Supercomputers, Simulation, and AI**
Presenting multiple examples of how high-performance computing technologies are being used in industry, this outreach event offered representatives of companies the chance to learn more about HLRS's resources, to meet with HPC experts, and to discover how HPC can accelerate innovation.

Sep 25–27 **REACH II Workshop**
The second REACH Workshop offered representatives of the culture and creative industries insights into how to access and use HPC technologies and expertise.

Sep 29–Oct 4 **First International Conference on AI-Based Systems and Services**
Organized by HLRS, a special track at this conference in Venice, Italy focused on the synergistic integration of high-performance computing and artificial intelligence.

Oct 10–11 **27th Results and Review Workshop**
Scientists and engineers, including users of HLRS's computing infrastructure, presented and discussed research results as well as challenges and best practices in using HPC systems.

Oct 21–22 **Reality in the Sciences**
This interdisciplinary workshop explored conceptions of reality in the different branches of science, including similarities, differences, and interactions among them.

Oct 24–25 **HPC User Forum**
Co-organized by Hyperion Research and HLRS, the HPC User Forum brings together senior representatives of key HPC initiatives, internationally prominent high-performance computing centers, leading technology manufacturers, and other experts on new HPC, AI, and quantum computing technologies.

Nov 25–27 **SAS24 – Modeling for Policy**
Organized by the HLRS Philosophy of Computational Sciences group, this Science & Art of Simulation conference aimed to define the capabilities and limits of simulation, explore how trust in models could be improved, and make recommendations that could lead to better interactions between simulation scientists and society at large.

Dec 3 **8th Industrial HPC User Round Table**
The annual iHURT meeting facilitates dialogue between HLRS and its industrial user community, focusing on innovative applications of HPC for research and development as well as challenges that industry faces in using HPC.

Organization Chart

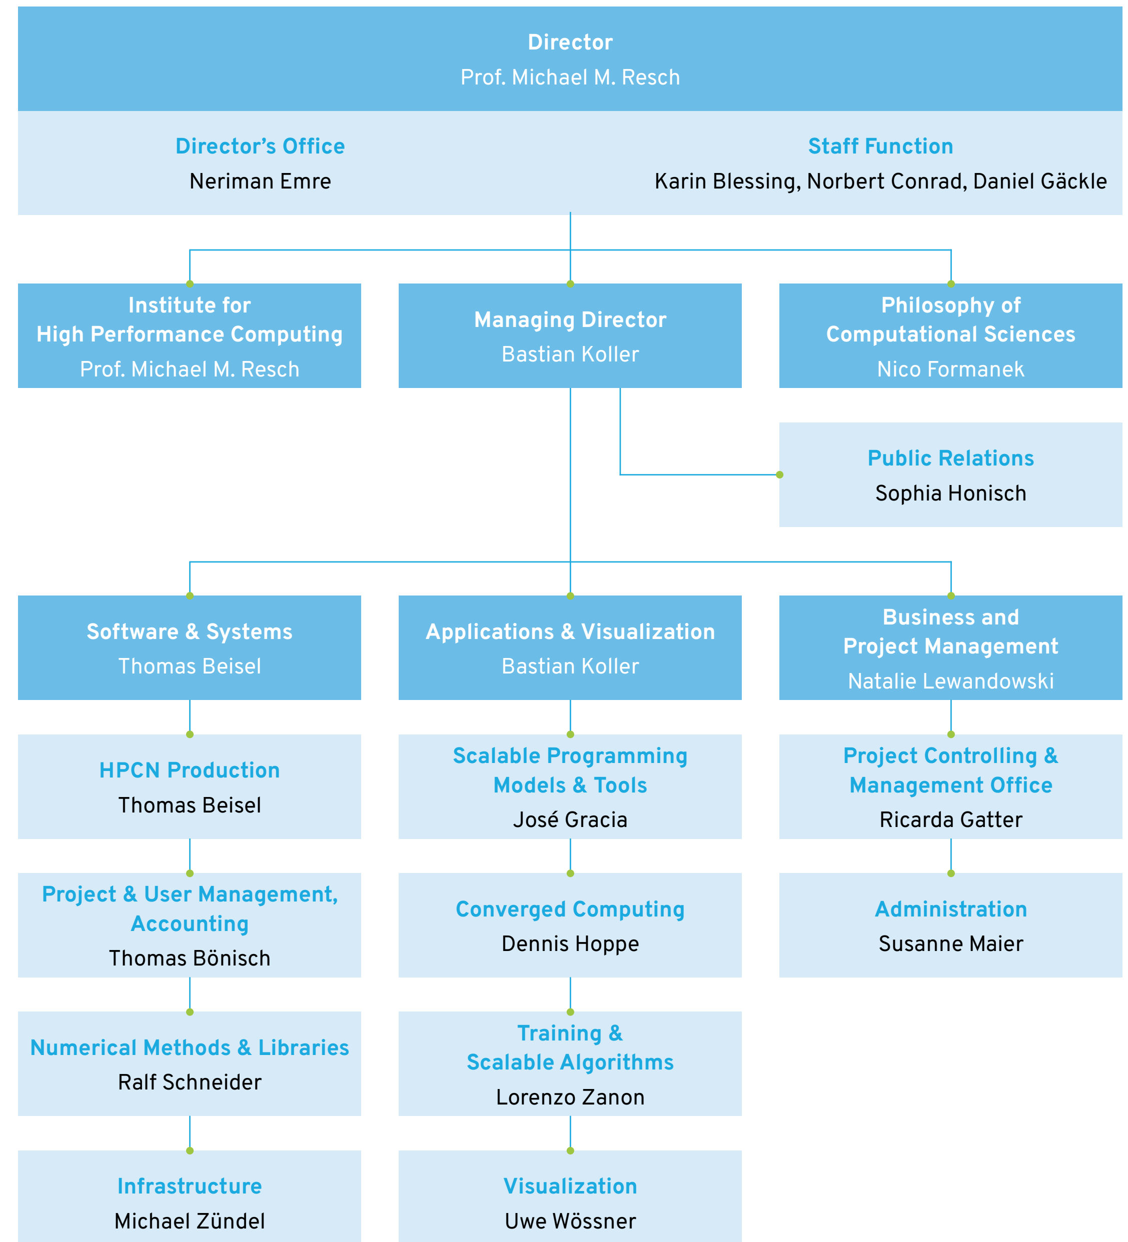

Departments

Administration

Leader: Susanne Maier

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

Converged Computing

Leader: Dennis Hoppe

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

High-Performance Computing Network – Production (HPCN Production)

Leader: Thomas Beisel

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

Infrastructure

Leader: Michael Zündel

Plans, develops, and operates the technical facilities and infrastructure of HLRS. This department ensures the reliable and efficient operation of HLRS's high-performance computing systems, provides a comfortable working environment for researchers and HLRS administration, and supports all aspects of energy-efficient HPC operations. The team also supports HLRS's sustainability program. It collects and monitors energy consumption data and drives continuous improvement through energy efficiency measures.

Numerical Methods and Libraries

Leader: Dr.-Ing. Ralf Schneider

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

Philosophy of Computational Sciences

Leader: Nico Formanek

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

Project Controlling and Management Office

Leader: Dr. Ricarda Gatter

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

Project and User Management, Accounting

Leader: Dr. Thomas Bönisch

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

Public Relations

Leader: Sophia Honisch

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

Scalable Programming Models and Tools

Leader: Dr. José Gracia

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

Training and Scalable Algorithms

Leader: Dr-Ing. Lorenzo Zanon

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

Visualization

Leader: Dr.-Ing. Uwe Wössner

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

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

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

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

Head, Department of Public Relations
Sophia Honisch

Writer and Editor
Christopher M. Williams

Contributing writer (where indicated)
Eric Gedenk (EG)

Photography and images
Unless otherwise indicated, all images property of HLRS.
p. 8, 10 stock.adobe.com / DZIKRA; p. 16 Courtesy of Guy Lonsdale; p. 17 FF4EuroHPC; p. 19 Sandra Ritschel; p. 22 Kumoh Institute of Technology (bottom); p. 26 Courtesy of Taltech; p. 27 MACK One; p. 28 Julian Holzwarth / HLRS; p. 29 Max Kovalenko / University of Stuttgart; p. 30 Julian Holzwarth / HLRS; p. 31 Manuel Emme Fotografie (top); Benthem Crouwel Architects, in collaboration with Birk Heilmeyer und Frenzel Architekten (bottom); p. 32 stock.adobe.com / mpix-foto; p. 37 ITLR, University of Stuttgart; p. 39 stock.adobe.com / Werner; p. 40 hydrograv GmbH; p. 41 Fundación Épica La Fura dels Baus.; p. 43 HPE / HLRS; p. 44 Schaefer et al. Phys Fluids. 2024; p. 45 Airbus Helicopters; p. 46 Manuel Keßler, University of Stuttgart; p. 49 Balourdas et al. 2024.; p. 50 Knodel et al. Int J Mol Sci. 2024; p. 59 Springer Verlag; p. 62 Julian Holzwarth / HLRS

Printing
gutenberg beuys, Feindruckerei GmbH, Langenhagen

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

Institutional affiliations
University of Stuttgart Germany
GCS Gauss Centre for Supercomputing

Funding for Hunterand Hawk provided by:
Baden-Württemberg Ministry of Science, Research and Arts
Federal Ministry of Research, Technology and Space

This magazine was printed climate-neutrally on paper that is certified according to the FSC® standard.

RECYCLED Paper made from recycled material FSC® C009051

Print product with financial climate contribution
ClimatePartner.com/10951-2505-1003

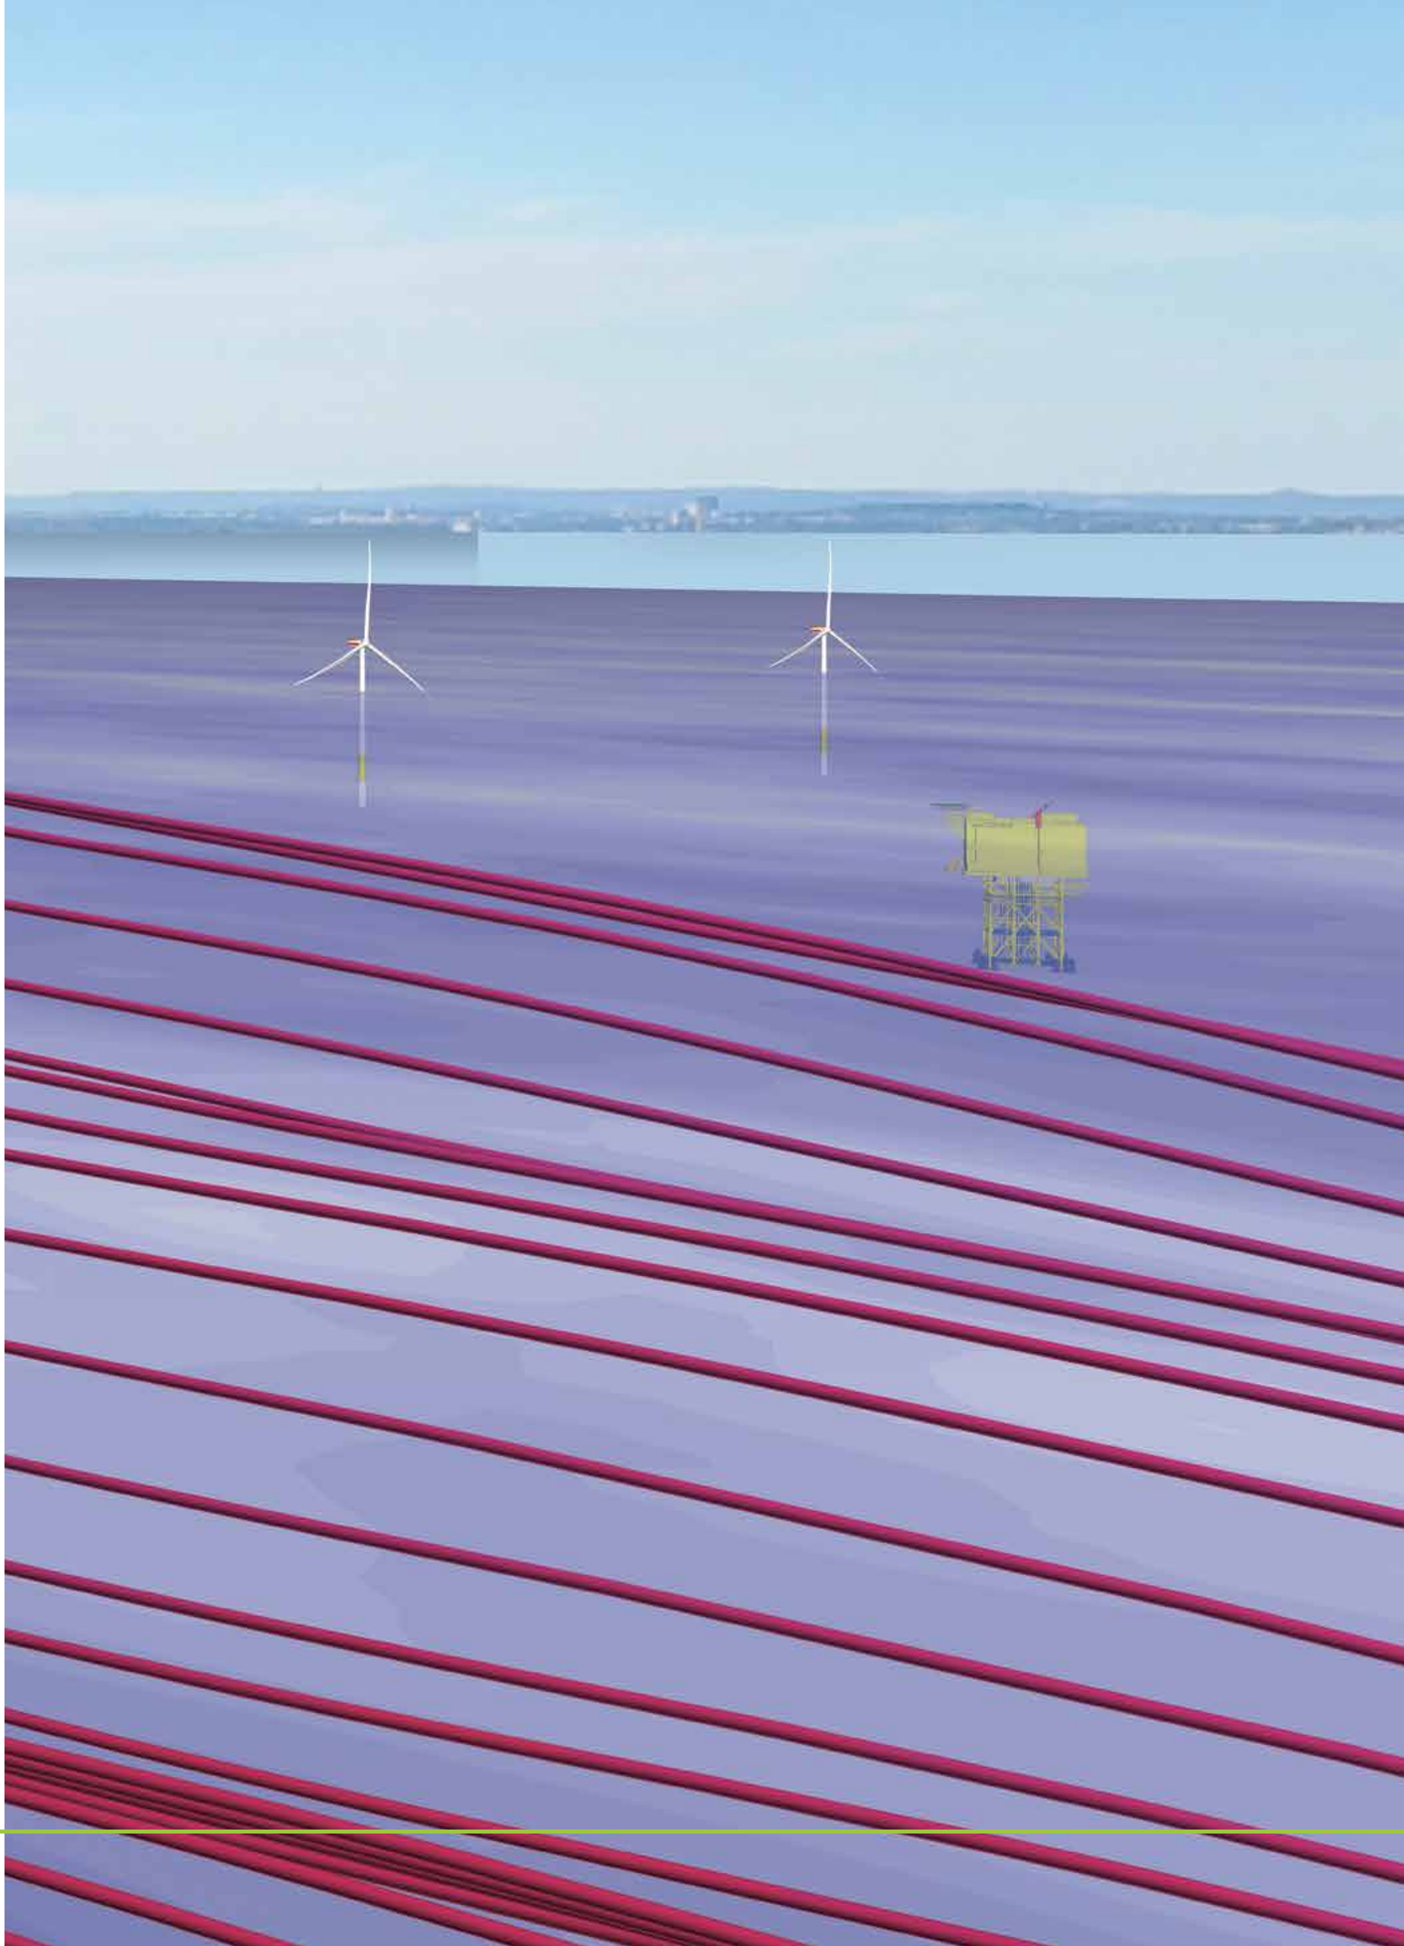

Follow us on:

LinkedIn Bluesky

High-Performance Computing Center Stuttgart
www.hlrs.de

Cover:
As participants in the EU-funded project MERIDIONAL, scientists at the High-Performance Computing Center Stuttgart have been developing interactive, high-fidelity visualizations of complex wind flow simulations to support the planning of Germany's largest offshore wind park, He Dreiht. Scheduled to go into operation in 2025, the facility is expected to provide renewable energy for the equivalent of 1.1 million households.

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