On the way to Exascale

HLRS

High-Performance Computing Center Stuttgart

2023 Annual Report

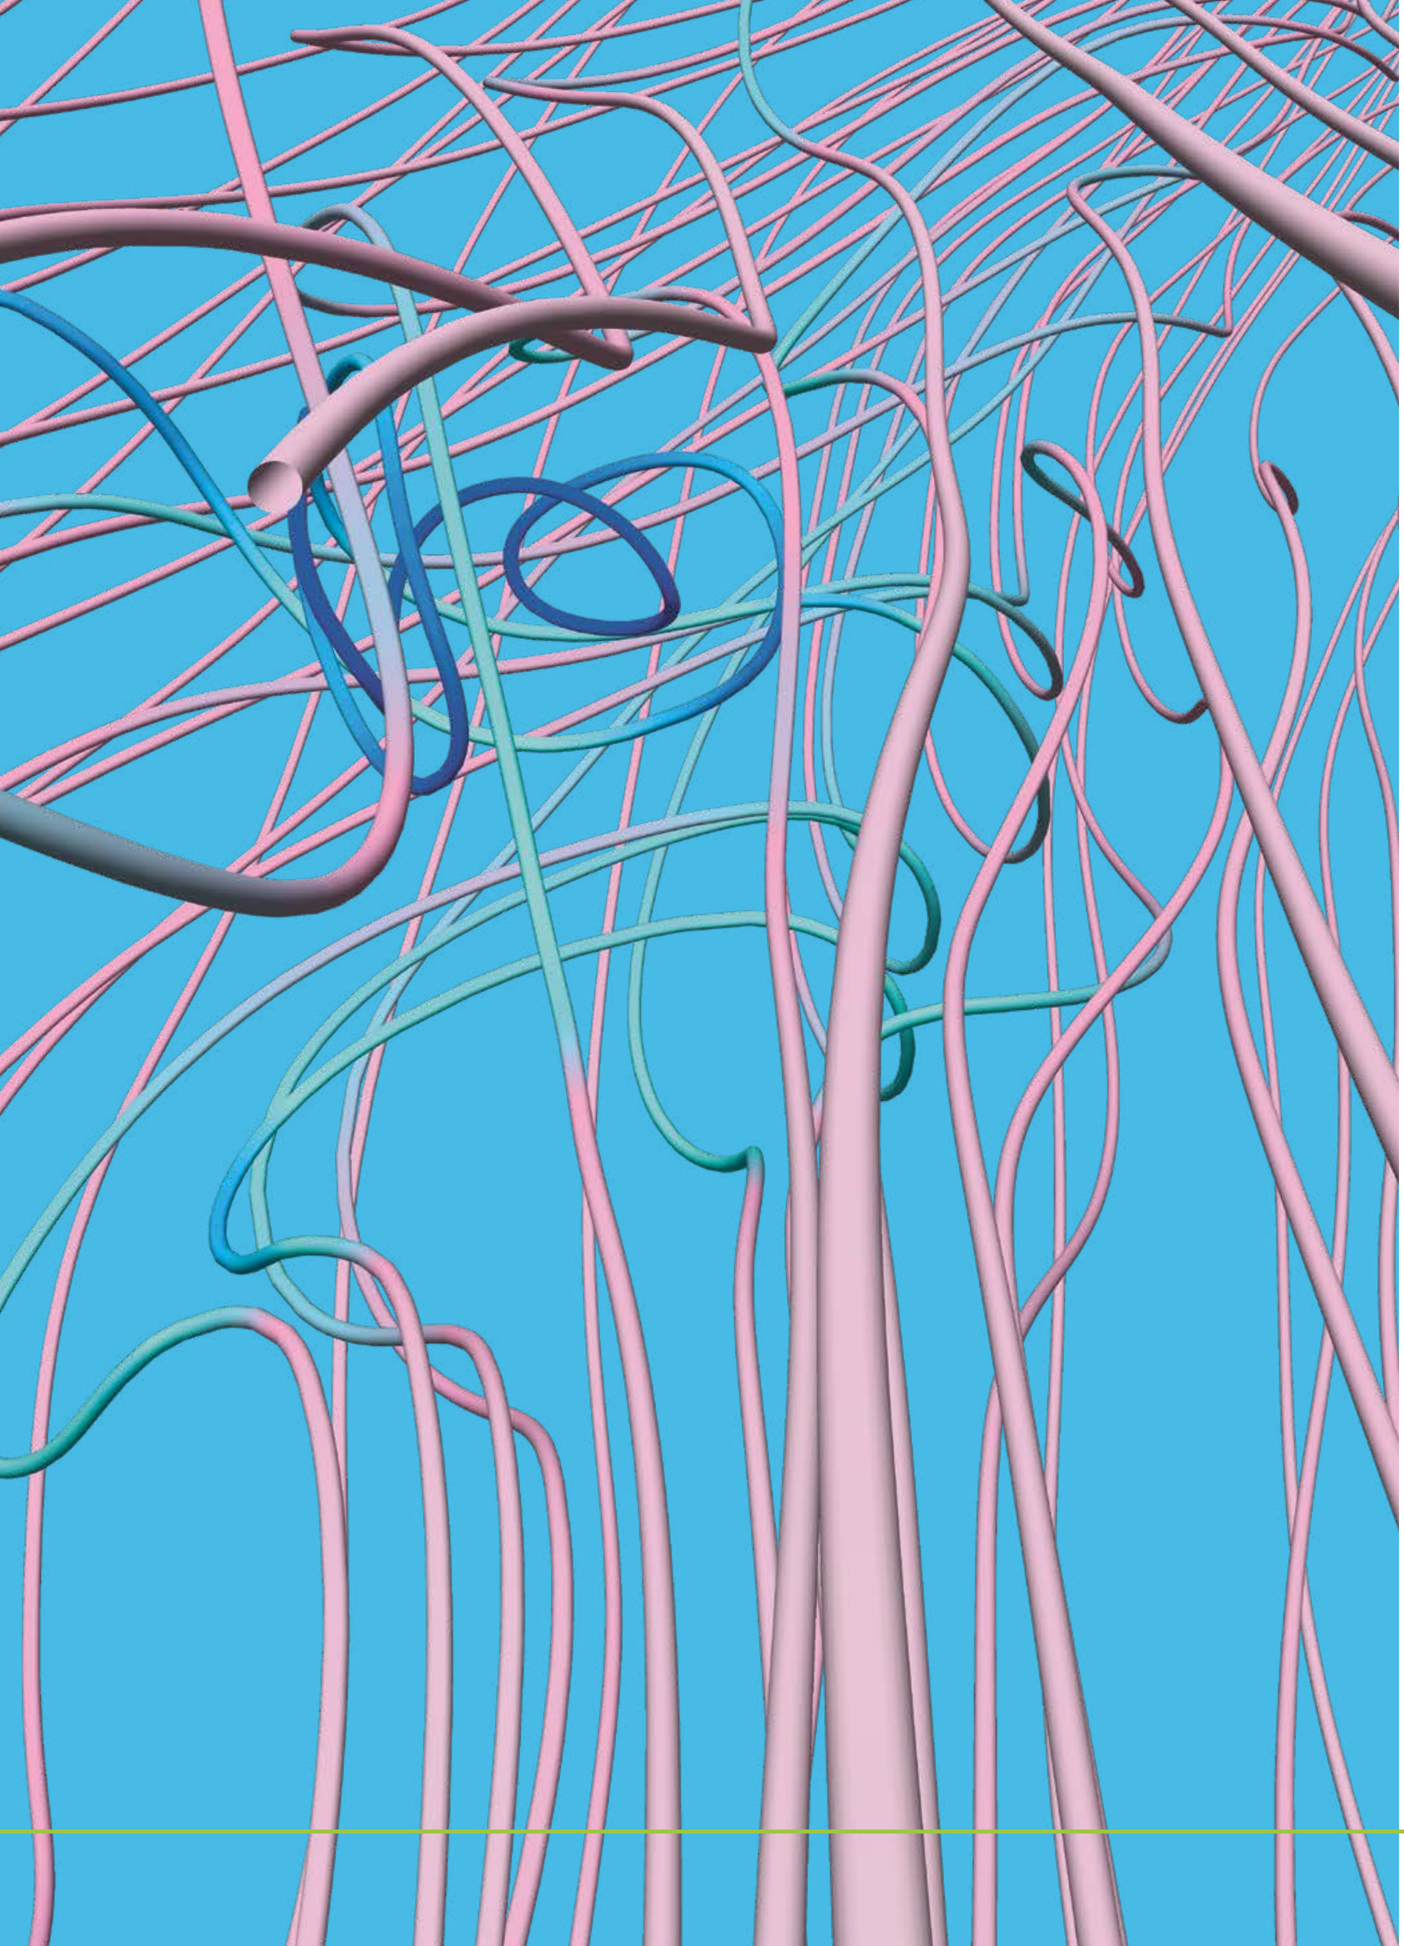

2023 Annual Report

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 and industry. HLRS operates one of Europe's most powerful supercomputers, provides advanced training in HPC programming and simulation, and conducts research to address key problems facing the future of supercomputing. Among HLRS's areas of expertise are parallel programming, numerical methods for HPC, visualization, grid and cloud computing concepts, data analytics, and artificial intelligence. Users of HLRS computing systems are active across a wide range of disciplines, with an emphasis on computational engineering and applied science.

Director's Welcome

Grußwort

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

Welcome to the 2023 Annual Report of the High-Performance Computing Center of the University of Stuttgart, presenting the results of an exciting year for HLRS, our users, and our partners.

This year the disruptions caused by the COVID-19 pandemic had largely passed, enabling our operations to stabilize. Our high-performance computing (HPC) training program has fully recovered, enhanced by a robust online curriculum launched during the pandemic that enables us to reach HPC users who cannot travel to Stuttgart. Our research project funding also returned to normal, pre-pandemic levels. Industrial usage was not affected by COVID-19, making 2023 a very good year with respect to industrial income. The number of companies working with HLRS and its systems is still growing, and industrial usage reached about 300 million core hours this year.

The highlight of 2023 was our procurement of two new supercomputers, culminating in December with the signing of a contract with Hewlett Packard Enterprise (HPE). The first, called Hunter, will arrive in 2024. It will serve as a stepping stone to a larger system, called Herder, which will enable HLRS to reach exascale performance in 2027. Hunter and Herder will be based on AMD technology and the Cray networking architecture, enabling our users to jump to the next stage in supercomputing technology. In this annual report we preview the new systems and some of the exciting new opportunities they will offer.

HLRS provided supercomputing resources and support for more than 130 scientific research projects in 2023. Outcomes of these activities were presented at the 26th Results & Review Workshop, and will also soon appear in the Transactions of the High-Performance Computing Center Stuttgart, published by Springer Verlag. This annual report highlights several of our users' applications, demonstrating how HLRS enables scientific discovery that is addressing today's largest challenges.

2023 saw HLRS's scientific team extend our activities across a variety of fields. As the field of artificial intelligence (AI) grows, it is gaining traction in both our user community and our research projects. Working together with Dr. Andrea Beck and staff at HPE, for example, we have been exploring how physics-informed machine learning approaches could enhance traditional simulation methods for CFD. In this annual report you can find an interview with Dr. Beck about this project. We are also proud that a collaboration between HLRS scientists and WIKKI GmbH was honored with an HPC Innovation Excellence Award. The team implemented an improvement to the open-source simulation software OpenFOAM that will make CFD simulations more accessible for industrial HPC users.

Dieser Jahresbericht des Höchstleistungsrechenzentrums der Universität Stuttgart umfasst die Ergebnisse eines aufregenden Jahres für das HLRS, seine Nutzerinnen und Nutzer sowie seine Partner.

Im Jahr 2023 waren die von der COVID-19-Pandemie verursachten Beeinträchtigungen weitgehend überwunden, sodass sich die Aktivitäten des Zentrums stabilisieren konnten. Unser Schulungsprogramm ist wieder vollständig im Normalbetrieb und wurde um ein umfangreiches Online-Kursangebot erweitert, das wir während der Pandemie eingeführt haben. Damit erreichen wir jetzt viele Nutzer von Höchstleistungsrechnern (HPC), die nicht nach Stuttgart reisen können. Auch die Einnahmen aus Drittmittelprojekten haben wieder das alte Niveau erreicht. Die industrielle Nutzung wurde nicht von der Pandemie beeinträchtigt. Etwa 300 Millionen Kernstunden wurden 2023 von Industrieunternehmen gerechnet und die Anzahl an industriellen Anwendern nimmt weiter zu.

Der Höhepunkt des Jahres 2023 war die Beschaffung von zwei neuen Supercomputern, die im Dezember bei der Unterzeichnung eines Vertrags mit Hewlett Packard Enterprise (HPE) besiegelt wurde. Das erste System, Hunter, wird 2024 geliefert und als Übergangssystem für das Exascale-System Herder dienen, das im Jahr 2027 eingeführt wird. Beide Systeme werden auf AMD-Technologie und Cray-Netzwerkarchitektur basieren und den Übergang zur nächsten Stufe der Supercomputing-Technologie ermöglichen. Hier geben wir einen Ausblick auf Hunter und Herder und einige der Chancen, die diese mit sich bringen werden.

Das HLRS unterstützte 2023 mehr als 130 Forschungsprojekte. Die Ergebnisse dieser Aktivitäten wurden auf dem 26. Results & Review Workshop vorgestellt und werden in den Transactions of the High-Performance Computing Center Stuttgart (Springer) veröffentlicht. Hier werden einige Anwendungen unserer User vorgestellt, die sich mit großen Herausforderungen der heutigen Zeit befassen.

Im Jahr 2023 haben wir die wissenschaftlichen Aktivitäten des HLRS auf diverse Gebiete ausgeweitet. Künstliche Intelligenz (KI) gewinnt sowohl in unserer User Community als auch in unseren Forschungsprojekten zunehmend an Bedeutung. In Zusammenarbeit mit Dr. Andrea Beck und Mitarbeitern von HPE haben wir beispielsweise untersucht, wie physikalisch informierte Ansätze des maschinellen Lernens traditionelle Simulationsmethoden für CFD verbessern könnten. Im Interview spricht Dr. Beck über diese neuen Möglichkeiten. Wir sind auch stolz darauf, dass eines unserer gemeinsamen Projekte mit der WIKKI GmbH mit einem HPC Innovation Excellence Award ausgezeichnet wurde. Das Team hat eine Verbesserung von OpenFOAM implementiert, die CFD-Simulationen für industrielle HPC-Anwender zugänglicher machen wird.

HLRS also contributed to projects that are bridging the gap between research and applications that are relevant for industry and public administration. A project called CIRCE, for example, has been investigating how digital twins could support crisis management. HLRS and its partners in the CASE4Med project have also been planning the establishment of a Medical Solution Center. This will bring HPC and AI to the medical technology community, which could benefit from HLRS's expertise and support.

Although HLRS's main focus is on supporting applied science and engineering, we are also interested in exploring how HPC and AI could support the humanities, the social sciences, and the arts. In collaboration with the University of Stuttgart's program in digital humanities, we helped to create a Stuttgart Research Focus on the (Re)production of Reality. Together with the Media Solution Center Baden-Württemberg we also hosted a meeting of European Institute of Innovation and Technology Culture & Creativity (EIT CC), a knowledge and information community that supports creatives in contributing to the development of a sustainable society.

HLRS continues to coordinate a Europe-wide effort to establish a network of national HPC competence centers within the scope of the EuroHPC Joint Undertaking (JU). The projects EuroCC and CASTIEL started their second project phases in 2023. An affiliated project called FF4EuroHPC held an industry summit in Berlin that highlighted new, successful applications of HPC and AI in small and medium-sized enterprises. Although the event marked the end of the project, we will likely implement a follow-up to this valuable effort in 2024. In the past year we also launched EuroHyPerCon, which will advise the JU on hyperconnectivity requirements across Europe.

The easing of the pandemic meant that we were once again able to visit many of our partners around the world, renewing and extending our cooperation. For example, we added a workshop on Japanese theatre to our annual visit in Japan, hoping to bring together science and art in ways that strengthen our long-term partnership with the Cyberscience Center of Tohoku University in Sendai. We were also pleased that collaborations with the Korean Institute of Science and Technology Information (KISTI) and Guangzhou National Supercomputing Center were energized by mutual visits.

In the midst of our many activities, HLRS remains at its core a national operations and service center for high-performance computing. In one important development, we completed certification under the ISO 27001 international standard for information security management. This should give our users and supporters the confidence that our systems are as secure as possible, and will enable our outstanding operations and infrastructure team to improve the quality of our services in the coming years.

The arrival of Herder in 2027 means that HLRS will need a new data center, and planning for this key infrastructure project is well under way. Working together with the University Construction Office for Stuttgart and Hohenheim, and the science and finance ministries of the State of Baden-Württemberg, our goal is to complete construction in 2027.

These many developments mean that in 2024 HLRS will face new challenges. With the help of our sponsors at the state, federal, and European levels; with our national and international partners; and thanks to the commitment of our staff, however, it is clear that HLRS is well prepared to meet them.

With best regards,

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

Darüber hinaus ist das HLRS an Projekten beteiligt, die die Kluft zwischen Forschung und Anwendungen überbrücken, die für die Industrie und Behörden relevant sind. So untersucht das CIRCE-Projekt beispielsweise, wie digitale Zwillinge das Krisenmanagement unterstützen könnten. Das HLRS und seine Partner im Projekt CASE4Med bauen außerdem ein Medical Solution Center auf. Dieses wird HPC und KI in die Medizintechnikbranche einbringen.

Obwohl der Schwerpunkt des HLRS auf der Unterstützung der angewandten Natur- und Ingenieurwissenschaften liegt, untersuchen wir auch, wie HPC und KI die Geistes- und Sozialwissenschaften sowie Kunst unterstützen können. In Zusammenarbeit mit den Digital Humanities der Universität Stuttgart haben wir einen Stuttgarter Forschungsschwerpunkt „Re/producing Realities" mitgeschaffen. Gemeinsam mit dem Media Solution Center Baden-Württemberg waren wir auch Gastgeber eines Treffens des European Institute of Innovation and Technology Culture & Creativity (EIT CC), einer Wissens- und Informationsgemeinschaft, die Kreative dabei unterstützt, zur Entwicklung einer nachhaltigen Gesellschaft beizutragen.

Das HLRS koordiniert weiterhin wichtige Projekte auf europäischer Ebene, die von der EuroHPC Joint Undertaking (JU) gefördert werden. Die Projekte EuroCC und CASTIEL gingen 2023 in ihre zweite Projektphase, in der das europäische Netzwerk nationaler HPC-Kompetenzzentren weiter ausgebaut wird. Ein weiteres JU-gefördertes Projekt, FF4EuroHPC, veranstaltete einen Industriegipfel in Berlin, auf dem erfolgreiche Anwendungen von HPC und KI in KMUs vorgestellt wurden. Obwohl der Gipfel anlässlich des Projektendes veranstaltet wurde, planen wir, diese Aktivitäten von 2024 an fortzusetzen. Im vergangenen Jahr haben wir auch die EuroHyPerCon-Studie mitinitiiert, die die JU zu den Anforderungen an die Hyperkonnektivität in Europa berät.

Das Ende der Pandemie ermöglichte es uns, viele unserer weltweiten Partner zu besuchen sowie unsere Kollaborationsvereinbarungen zu erneuern oder zu erweitern. So haben wir unseren jährlichen Besuch in Japan um einen Workshop über japanisches Theater erweitert, um Wissenschaft und Kunst miteinander zu verbinden. Darüber hinaus konnten wir die Zusammenarbeit mit dem Korean Institute of Science and Technology Information (KISTI) und dem Guangzhou National Supercomputing Center dank gegenseitiger Besuche kräftigen.

Inmitten zahlreicher Aktivitäten war die Zertifizierung des HLRS nach ISO 27001 eine weitere wichtige Entwicklung. Mit dieser Zertifizierung möchten wir gegenüber unseren Nutzerinnen und Nutzern sowie unseren Fördergebern die Sicherheit unserer Systeme hervorheben und unser Bestreben ausdrücken, die Qualität unserer Dienstleistungen laufend zu verbessern. Dies ist nur dank unserer ausgezeichneten Teams in Betrieb und Infrastruktur möglich.

Für die Installation von Herder im Jahr 2027 benötigt das HLRS ein neues Gebäude. Die Planung für dieses wichtige Infrastrukturprojekt ist bereits in vollem Gange. In Zusammenarbeit mit dem Universitätsbauamt Stuttgart und Hohenheim sowie dem Wissenschafts- und dem Finanzministerium des Landes Baden-Württemberg streben wir an, den Bau im Jahr 2027 abzuschließen.

Diese vielfältigen Entwicklungen stellen das HLRS im Jahr 2024 vor neue Herausforderungen. Mithilfe unserer Fördergeber auf Landes-, Bundes- und europäischer Ebene, mit unseren nationalen und internationalen Partnern und dank des Engagements unserer Mitarbeiterinnen und Mitarbeiter ist das HLRS jedoch gut darauf vorbereitet.

Mit freundlichen Grüßen

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

Exascale Supercomputing Is Coming to Stuttgart

The University of Stuttgart and Hewlett Packard Enterprise announced an agreement to build two new supercomputers at HLRS: Hunter and Herder.

The High-Performance Computing Center Stuttgart (HLRS) has set its path to exascale. Following a contact signing ceremony in December 2023, HLRS and Hewlett Packard Enterprise (HPE) announced the next two generations of supercomputers at HLRS.

In the first stage, a transitional supercomputer, called Hunter, will begin operation in 2025. This will be followed in 2027 with the installation of Herder, an exascale system that will provide a significant expansion of Germany's high-performance computing (HPC) capabilities. Hunter and Herder will offer researchers world-class infrastructure for simulation, artificial intelligence (AI), and high-performance data analytics (HPDA) to power cutting-edge academic and industrial research in computational engineering and the applied sciences.

The total combined cost for Hunter and Herder is €115 million. Funding will be provided through the Gauss Centre for Supercomputing (GCS), the alliance of Germany's three national supercomputing centers. Half of this funding will be provided by the German Federal Ministry of Education and Research (BMBF), and the second half by the State of Baden-Württemberg's Ministry of Science, Research, and Arts.

Hunter to Herder: a two-step climb to exascale

Hunter will replace HLRS's current flagship supercomputer, Hawk. It is conceived as a stepping stone to enable HLRS's user community to transition to the massively parallel, GPU-accelerated structure of Herder.

Hunter will be based on the HPE Cray EX4000 supercomputer, which is designed to deliver exascale performance to support large-scale workloads across modeling, simulation, AI, and HPDA. Each of the 136 HPE Cray EX4000 nodes will be equipped with four HPE Slingshot high-performance interconnects. Hunter will also leverage the next generation of Cray ClusterStor, a storage system purpose-engineered to meet the demanding input/output requirements of supercomputers, and the HPE Cray Programming Environment, which offers programmers a comprehensive set of tools for developing, porting, debugging, and tuning applications.

Hunter will raise HLRS's peak performance to 39 petaflops ($39*10^{15}$ floating point operations per second), an increase from the 26 petaflops possible with its current supercomputer, Hawk. More importantly, it will transition away from Hawk's emphasis on CPU processors to make greater use of more energy-efficient graphics processing units (GPUs).

Hunter will be based on the AMD Instinct™ MI300A accelerated processing unit (APU), which combines CPU and GPU processors and high-bandwidth memory into a single package. By reducing the physical distance between different types of processors and creating unified memory, the APU enables fast data transfer speeds, impressive HPC performance, easy programmability, and great energy efficiency. This will slash the energy required to operate Hunter in comparison to Hawk by approximately 80 % at peak performance.

Prof. Michael Resch (Director, HLRS), Anna Steiger (Chancellor, University of Stuttgart), and Heiko Meyer (Chief Sales Officer, HPE) signed the contract for Hunter and Herder in a ceremony at HLRS.

Herder will be designed as an exascale system capable of speeds on the order of one quintillion (10^{18}) flops, a major leap in power that will open exciting new opportunities for key applications run at HLRS. The final configuration, based on accelerator chips, will be determined by the end of 2025.

The combination of CPUs and accelerators in Hunter and Herder will require that current users of HLRS's supercomputer adapt existing code to run efficiently. For this reason, HPE will collaborate with HLRS to support its user community in adapting software to harness the full performance of the new systems.

Supporting scientific excellence in Stuttgart, Germany, and beyond

HLRS's leap to exascale is part of the Gauss Centre for Supercomputing's national strategy for the continuing development of the three GCS centers: The upcoming JUPITER supercomputer at the Jülich Supercomputing Centre will be designed for maximum performance and will be the first exascale system in Europe in 2025, while the Leibniz Supercomputing Centre is planning a system for widescale usage in 2026. The focus of HLRS's Hunter and Herder supercomputers will be on computational engineering and industrial applications. Together, these systems will be designed to ensure that GCS provides optimized resources of the highest performance class for the entire spectrum of cutting-edge computational research in Germany.

For researchers in Stuttgart, Hunter and Herder will open many new opportunities for research across a wide range of applications in engineering and the applied sciences. For example, they will enable the design of more fuel-efficient vehicles, more productive wind turbines, and new materials for electronics and other applications. New AI capabilities will open new opportunities for manufacturing and offer innovative approaches for making large-scale simulations faster and more energy efficient. The systems will also support research to address global challenges like climate change, and could offer data analytics resources that help public administration to prepare for and manage crisis situations. In addition, Hunter and Herder will be state-of-the-art computing resources for Baden-Württemberg's high-tech engineering community, including the small and medium-sized enterprises that form the backbone of the regional economy.

HPC in Transition: Hunter and Herder Will Bring New Opportunities, New Challenges

HLRS's coming supercomputers will not only enable traditional high-performance computing to reach new heights in performance, but will also better support new, complementary approaches involving artificial intelligence, deep learning, and high-performance data analytics.

The announcement of HLRS's next-generation supercomputers, Hunter and Herder, marks a significant moment in the history of high-performance computing (HPC) at the University of Stuttgart, charting a path for HLRS to reach exascale performance. This major leap in computing power will help to maintain the center's status as a leading institute for high-performance computing in Europe.

At the same time, the announcement marks a noteworthy technological shift. Whereas HLRS has offered a predominantly CPU-based architecture for many years, Hunter and Herder will achieve their speedup gains by including large numbers of graphics processing units (GPUs). The new systems will make it possible for scientists and engineers who use numerical methods to run larger simulations faster than ever before. At the same time, the shift to GPUs will increase HLRS's ability to support new approaches involving artificial intelligence, machine learning, deep learning, and high-performance data analytics.

This transitional moment at HLRS is emblematic of a range of changes that are taking place in high-performance computing across the world. As the center navigates this new terrain, users can expect both exciting new capabilities and new challenges.

Why HPC is transitioning to GPUs

When Gordon Moore articulated his eponymous Moore's Law in 1975, he predicted that the number of components that could be packed onto an integrated circuit would double approximately every two years. For decades the steady, rapid increase in computing power was evidence of the accuracy of this prediction. Today, however, is a different story. Moore's Law appears to be approaching its end.

In high-performance computing, most CPU-based supercomputers have been built using an architecture called x86, but this framework has reached its limit. Potential strategies such as adding more cores, shrinking component sizes, or increasing processor frequency have either become or will likely soon become impossible. Meanwhile, simply building larger supercomputers with even more CPUs would be financially and environmentally unsustainable due to the power and material resources they would require.

Computer manufacturers have been aware of this impending limit for some time, and have been developing new architectures to meet ever-increasing demand for higher computing speeds. One of the most popular involves shifting from CPUs to an architecture that uses graphics processing units (GPUs) as accelerators.

Most of the world's fastest supercomputers now include GPUs, including the Frontier and Aurora systems by Hewlett Packard Enterprise / Cray, currently in places 1 and 2 on the Top500 List.

GPUs are simpler in construction than CPUs, but provide a much faster and more energy-efficient way to execute large numbers of calculations in parallel. Because GPU clock speeds are lower, they require less power. Still, GPUs run faster than CPUs by packing many more cores on each computing node, making data transfer among cores much more efficient.

The upcoming Hunter supercomputer, arriving by 2025, will be based on the AMD MI300A accelerated processing unit (APU), which combines CPU and GPU processors on a single chip with a shared memory. Hunter will offer at least a 50 % increase in speed over HLRS's current flagship machine, Hawk. At the same time, the HPE Cray EX-4000 system will consume the equivalent of just 20 % of Hawk's power at peak performance. When it arrives in 2027, Herder will expand on Hunter's GPU-accelerated concept in the form of a much larger exascale machine.

Once they enter production, Hunter and Herder will enable HLRS to support more powerful simulations. For scientific users, this will make it possible to investigate complex phenomena at higher accuracy, to simulate larger systems than in the past, or to more efficiently run repeated simulations that reveal effects of specific parameters or provide greater statistical power for data analysis. The combination of CPUs and GPUs on APU chips will also make it possible to integrate simulation and data analytics more efficiently than was possible in the past, permitting a number of exciting potential applications.

Opportunities at the intersection of simulation and artificial intelligence

In conventional high-performance computing, researchers run large-scale simulations, receive results as data, and then analyze their results for interesting features and patterns by visualizing or processing the data further. In most cases, data post-processing can be very

"GPU systems will not only enable larger simulations, but also speed up data post-processing and analysis."
Xiang Xu, IMS, University of Stuttgart

time consuming, taking significantly longer than running the simulation itself. This challenge will become even more severe as larger supercomputers make it possible to generate even larger datasets.

Dr. Matthias Meinke at the RWTH Aachen University Institute of Aerodynamics (AIA) has used HLRS's supercomputers for many years to run numerical simulations of turbulent flows. "The problems we want to investigate are getting more complex and will produce larger and larger amounts of data," he says. "This means that it will be important to investigate how high-performance data analytics and artificial intelligence could help."

Using APUs, Hunter could enable researchers to integrate artificial intelligence methods into their applications and workflows. Here, AI could be utilized to browse simulation data in real-time, reducing the need for later data post-processing.

Graduate student Xiang Xu of the University of Stuttgart's Institute of Materials Science foresees advantages of this approach. Recently, he has been using Hawk to perform molecular dynamics simulations of nickel aluminide crystals. Ni-based superalloys are very popular materials, especially in aircraft engines, and Xu's research focuses on properties at the atomic level that can affect its mechanic behavior. In addition to being able to simulate larger systems of atoms than is possible today, he says that Hunter and Herder will help to save significant time in analysis.

"When doing molecular dynamics simulations we don't need all of the data for all of the atoms, but are looking for evidence of dislocations, vacancies, or defects that have already been measured at the macro scale," Xu says. Real-time data analytics tools, he anticipates, will help identify and isolate features of interest more quickly and make it possible to download and store just the data he needs.

In addition, approaches using artificial intelligence could help to synthesize and analyze the massive datasets that can accumulate at a lab or research institute over many years. AI could provide a global perspective that is often not relevant when researching specific problems or parameters, identifying patterns hidden in data that might not have been important in the original simulations but become interesting in the context of larger datasets. Using AI, data accumulated over many years could become a rich resource for generating new hypotheses and making new discoveries.

At the same time that Dr. Meinke looks forward to exploring possibilities offered by data analytics approaches — in developing surrogate models of complex simulations, for example — he also emphasizes that they should be seen as a complement to traditional methods and not their replacement. "We will need to be careful in terms of our expectations and to be critical in determining whether AI provides a general solution or only functions under specific conditions," he cautions. Research using Hunter and Herder could help in testing these strategies.

Surrogate models and their challenges

Most of HLRS's supercomputing resources go to scientists in academia with large grants of computing time. For engineers in industry, however, it is often not possible or even desirable to run simulations at such large scales. For this reason, researchers have been investigating how machine learning and artificial intelligence could be used to develop surrogate models of complex systems. Based on data generated from first-principles simulations, surrogate models accurately replicate relevant features of complex systems in a simplified manner that can be run on a conventional desktop computer. Such models provide engineers in industry with practical tools that are firmly grounded in the best possible research.

According to Prof. Bernhard Weigand, director of the University of Stuttgart's Institute of Aerospace Thermodynamics (ITLR), the convergence of supercomputing power, new experimental techniques, and machine learning offers exciting new opportunities to develop better surrogate models. "Measurement technologies have improved in parallel with computational methods and technologies, making it possible to see many de-

"Machine learning could help to assimilate experimental data and numerical methods to develop better models."
Bernhard Weigand, ITLR, University of Stuttgart

tails that we couldn't see numerically," he says. "The exciting challenge for the coming years is to synchronize these developments, using machine learning to assimilate experimental data and numerical methods to develop better models."

Such methods will also benefit from an emerging approach called physics-informed neural networks. Here, training algorithms incorporate fundamental physical laws to ensure that they are grounded in physical reality. This approach will ensure not only that the resulting surrogate models are faster than traditional simulations using HPC, but also that that the predictions they generate correspond to the real world and are reliable in a wide range of situations.

Achieving this goal does not mean that high-performance computing in its traditional form will disappear. Rather, GPU-accelerated systems will be more critical than ever for running highly accurate simulations and providing advanced AI capabilities. Prof. Wolfgang Schröder, who heads the RWTH Aachen University Institute of Aerodynamics (AIA) sees great potential in AI methods but emphasizes that numerical methods will continue to be the gold standard for computational research. Moreover, AI approaches will have a cost. "When people say that in the future everything will be done using machine learning and AI, they assume that the necessary data are available. In computational fluid dynamics, this isn't true," Schröder says. "We need large amounts of computing time to generate enough data — in some cases we might need to run thousands of simulations to come up with a sustainable surrogate model. This will not happen by itself, and it is one reason why physics disciplines will need exascale systems like Herder."

"We need large amounts of computing time to generate enough data to train AI models. This is one reason why physics disciplines will need exascale systems like Herder."
Wolfgang Schröder, AIA, RWTH Aachen

Training and user support will help in the transition to GPUs

For a variety of technical reasons, algorithms written for CPU-based architectures cannot simply be ported onto GPUs and be expected to run efficiently. Instead, users of HLRS's supercomputers will, in many cases, need to look closely at how their codes are structured and written in order to ensure that they operate at high performance on the new hardware.

"In recent years we have continually optimized our code to improve its performance CPUs," says ITLR scientist Matthias Ibach. "Having access to GPUs offers new opportunities and will be much more effective, but we need to make sure that we continue to utilize the full potential of the new system. If we continue to think in the same ways we have in the past, the performance of the new system could turn out to be worse than what we have now."

HLRS's user community will not be alone in making this transition. An important consideration in the center's decision to contract with Hewlett Packard Enterprise for its next-generation systems is the manufacturer's commitment to user support. In the coming years, HPE staff with expert knowledge of their systems, working in collaboration with the HLRS user support team, will assist HLRS's users in this effort.

HLRS Director Prof. Michael Resch points out that this is not the first time that the HLRS community has had to migrate to a new technology. "Since its founding in 1996, HLRS has periodically needed to upgrade to new supercomputing architectures to stay at the cutting edge of high-performance computing," he explains. "In this sense, the challenge that we are currently facing with GPUs is not new, although the specific challenge is real because many of our users have not utilized this type of system before. This is why Hunter is conceived as a transitional system. Supporting our users in making this jump is imperative for us."

"GPUs offer new opportunities and will be much more effective, but we need to make sure that we utilize the full potential of the new system."
Matthias Ibach, ITLR, University of Stuttgart

Over the past several years, HLRS has been laying the groundwork for the transition to GPUs by expanding its HPC training program to address skills that users will need. In collaboration with other HPC centers across Europe, HLRS has offered "bootcamp" courses introducing programming models that users can apply to port their codes onto GPUs. Additional courses have also focused on approaches for deep learning and artificial intelligence.

While some research groups are expected to find quickly that their codes can run much faster on GPUs, others might require more support. Methods that simply run clearly defined groups of repetitive calculations will have an easier time, for example, than algorithms that use object-oriented programming to track changing relationships among distinct objects, like particles. Looking toward the coming years, Dr. Schröder explains, "There is a tension we have to accept between what is expected with respect to our research in fluid dynamics and what is expected in numerical analysis. It will be difficult, but we like such a challenge because we know that writing the next generation of 'superpower code' is necessary not just for us but also for HLRS."

Reliable resources and new opportunities for industry

Whereas many academic scientists actively develop their own advanced codes, researchers in industry who use HPC for simulation are typically reliant on commercially produced software packages that they modify for specific applications. In the near term, this could present some challenges in transitioning to GPUs. In some cases software packages are already available and will work just fine, but this is not universally the case.

For this reason, HLRS plans to continue offering its industrial users access to x86 nodes on its Vulcan cluster. This long-running, heterogeneous system has been updated over the years to offer industry a multitude of CPU processor generations. It will continue to support HLRS's industrial user community, including small and medium-sized enterprises (SMEs) in running the simulations they need to innovate.

At the same time, the availability of GPUs on Hunter and Herder will open the doors to HLRS for new industrial user communities. Following the storm of public interest surrounding ChatGPT and generative AI, many companies are currently exploring how artificial intelligence and data analytics could support their activities, and some are already using applications designed to run on GPUs.

A preliminary architectural rendering of HLRS III shows the location of the new building behind HLRS's current HPC training facility.

According to HLRS Managing Director Dr. Bastian Koller, "We have been talking about artificial intelligence for a long time, but with Hunter and Herder HLRS we will gain powerful new platforms to support it. Many users from industry who know what they want to do say they don't need massive systems, that access even to a relatively modest number of processors could move their work forward. The benefit of HLRS's approach is that we can work together closely with our technology vendors on a small scale in order to find solutions for our users' specific problems."

Ready for the future

With Hunter and Herder, HLRS will first and foremost continue to offer its user community world-class tools and services for cutting edge computational research. Although the transition will present challenges, the ability to perform larger scale simulations while also having access to powerful AI and data analytics capabilities will offer a flexible platform for science, engineering, public administration, and other communities.

"Because so much is changing right now in high-performance computing it's hard to predict exactly how our users will utilize the new systems," says Prof. Resch. "We are very confident, though, that Hunter and Herder put HLRS on a good course for the future. As our users make the transition to exascale computing with us, we look forward to seeing the exciting new applications and discoveries will become possible."

Preparing for HLRS III

When Herder arrives in 2027, it will mean not just a major increase in computing power, but also an expansion of HLRS's facilities. Because the exascale system's larger physical and power requirements will exceed the capacity of HLRS's current infrastructure, the state of Baden-Württemberg will erect a new building, called HLRS III, and install a larger power supply. Planning has already begun to build the facilities on land behind HLRS that had previously been zoned for the center's expansion.

Sustainability has been a key factor in planning for Herder, and HLRS III will contribute to the University of Stuttgart's efforts to achieve an emissions-free campus. This includes capturing and reusing the heat that the exascale supercomputer will produce by connecting it to the University's heating network. Engineers estimate that Herder's waste heat could cover approximately 35 % of the university's heating needs in the winter months and would provide more than enough process heat for the entire campus during the summer. This could help to reduce the university's energy consumption and the associated heating costs.

In 2023 HLRS held three community outreach events in which local residents learned about the center's plans and offered comments and suggestions for integrating the new building into the neighborhood and surrounding landscape. Representatives of HLRS, the Buildings Department of the Universities of Stuttgart and Hohenheim, and HLRS III's project architects were on hand to present preliminary concepts for the design and engineering of the new facilities. Visitors could also experience the new building in virtual reality in HLRS's CAVE 3D-visualization facility.

Three events offered a chance for the local community to learn more about plans for HLRS III.

In the HLRS CAVE, design concepts for Stuttgart's new Rosensteinbrücke came to life in virtual reality.

News Briefs

German Minister of Education and Research Bettina Stark-Watzinger met with HLRS director Prof. Michael Resch.

Federal Science Minister Bettina Stark-Watzinger Tours HLRS

On July 28, as part of a five-day summer tour across Germany, Minister of Education and Research Bettina Stark-Watzinger paid a visit to HLRS. The summer tour highlighted some of Germany's most innovative research projects, start-ups, and companies, focusing on key technologies for the future, including artificial intelligence, supercomputing, and biotechnology. Prof. Dr. Michael Resch, director of HLRS, welcomed Minister Stark-Watzinger before escorting her on a tour of the center's computing room and its CAVE visualization facility. Members of the HLRS visualization department demonstrated several applications of simulation and virtual reality, including how urban digital twins can support city planning and the analysis of air quality. The German Ministry of Education and Research (BMBF) is a key funder of HLRS, providing core support for HLRS's facilities and operations together with the Baden-Württemberg Ministry for Science, Research and Art. BMBF funds HLRS through the Gauss Centre for Supercomputing (GCS), the alliance of Germany's three national supercomputing centers.

City Planning Hackathon Imagines Bridge of the Future

In 2022 the City of Stuttgart announced that the Rosensteinbrücke, a busy bridge crossing the Neckar River, needs to be replaced. As a side event of the Urban Future 2023 conference, an international event focusing on sustainable urban development, researchers in the HLRS Visualization Department hosted a full-day hackathon to develop ideas for the structure. The event was also organized in cooperation with CapeReviso, a project co-organized by HLRS that has developed planning and decision support tools for reducing conflicts between cyclists and pedestrians. During the hackathon, teams of students led by Dr. Peter Zeile (Karlsruhe Institute of Technology) presented preliminary design concepts. HLRS staff then helped to import them into the CAVE visualization environment, where the students experienced in virtual reality how their plans might look and discussed how they could be optimized. The resulting proposals offered attractive solutions both for the structure and for making better use of the city's underutilized waterfront.

Before entering the CAVE, the students presented posters and discussed seven different design concepts for the bridge.

User Survey Results Identify Future Trends

In 2023, HLRS conducted a survey of its user community to assess how well it is addressing its needs and to plan for future developments in the HPC resources and services that it provides. More than 75 % of respondents reported satisfaction with HLRS and more than 80 % replied that it is easy to make use of HLRS's resources. Questions about users' interests also identified clear trends that will guide HLRS in the coming years. Forty-six percent of respondents anticipate engaging with the shift to accelerators, while 43 % see machine learning as an important topic in their future research. In comparison, just 7 % expressed interest in high-performance data analytics, a valuable insight that will help guide the evolution of HLRS's future technology mix. Considering other infrastructure needs, between 60 % and 70 % of respondents expect their data storage and performance requirements to grow in the coming year. Interestingly, most respondents expected their bandwidth requirements to remain the same. HLRS will use these predictions to ensure that it continues to optimize its resources for high-performance computing in ways that best support research and technology development across its user community.

Trustworthiness and Reliability in AI

Because artificial intelligence is far from transparent, the question of whether we can trust it has become an important issue in using and assessing its capabilities. Until recently, however, there has been little discussion of what exactly we mean by trust in this context. Is it sufficient for AI to be reliable, producing results that correspond to what we observe? Or must we achieve a deeper level of trust based on the kinds of ethical considerations that govern human interactions? At a multidisciplinary conference organized by the HLRS Department of Philosophy of Computational Sciences called "Reliability or Trustworthiness?", researchers from the humanities, social sciences, and computer sciences explored what such concepts might mean for AI. Their talks offered novel perspectives on applied AI as well as insights that could inform future AI models and policy.

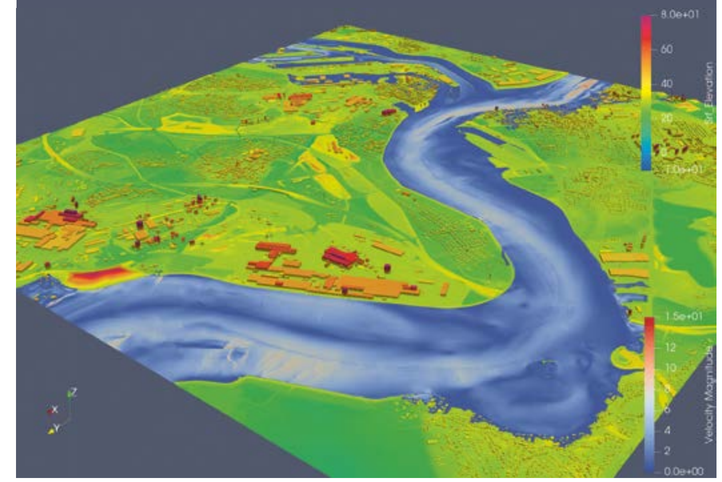

A visualization of flood vulnerabilities near the Rhine River in Duisburg.

CIRCE Collaboration Simulates Flood Risks in Duisburg

Building on experience that HLRS has gained in recent years, the project CIRCE has been conducting outreach to representatives of city governments and public administration to understand how high-performance computing and artificial intelligence could help them in managing crisis situations. In March 2023, CIRCE organized an event in which officials learned how HLRS has supported the development of tools for forecasting intensive care unit demand in the COVID-19 pandemic, predicting sudden migration events, or assessing tsunami risks. The discussion also invited attendees to consider potential applications of HPC and AI in their own municipalities, and addressed important logistical questions related to data availability and systems access that could determine the feasibility of this approach. In one result of this outreach, HLRS started a new collaboration with the fire department of the city of Duisburg focused on better understanding flood risks behind a railway embankment near the Rhine River. Although the city is aware that low-lying neighborhoods close to the water lie in a flood plain, simulation is helping it to address the complex question of how barriers located at underpasses below the rail line would affect the flow of water through the city's sewer network. The study will be presented as part of an event in Berlin in June 2024, where public officials can learn from CIRCE's findings and outcomes. The project is funded by the Federal Ministry of Education and Research (BMBF) and the Baden-Württemberg Ministry of Science, Research, and the Arts.

ORCHESTRA Organizes COVID-19 Science Challenge

Launched in 2020, ORCHESTRA is a multinational, EU-funded research project focused on COVID-19 and its long-term aftereffects. Coordinated by researchers at the University of Verona, Italy, the project has developed a data management infrastructure that integrates SARS-CoV-2 clinical data hosted in multiple countries, providing a large virtual cohort for research. In July 2023 the project announced a COVID-19 Science Challenge, which for the first time invited external researchers to use the data contained in this virtual cohort. HLRS scientists have helped to develop the federated data infrastructure, including implementing a software stack for national COVID-19 data hubs in Germany, Italy, and France. HLRS staff has also helped to ensure secure storage of sensitive health data. The ORCHESTRA team has been focused on harmonizing datasets across the entire network, implementing standardized variables and metadata frameworks that make it possible to analyze data hosted in multiple national hubs in an integrated way.

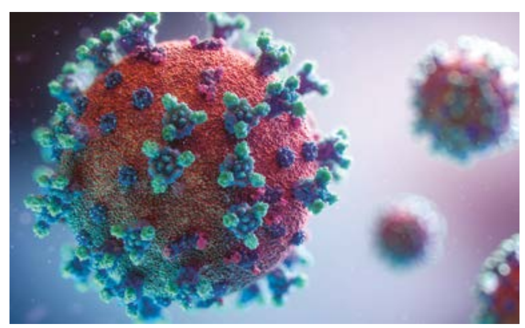

AI Alliance Baden-Württemberg to Create Cooperative Data Platform

In an effort to create an internationally competitive ecosystem for AI, a community of administrative organizations, universities, and private research centers in southwestern Germany united in 2023 to establish the AI Alliance Baden-Württemberg. Sponsored by the Baden-Württemberg Ministry of Economic Affairs, Labor and Tourism, the AI Alliance is creating regionally distributed, sector-specific hubs that drive the development of new AI applications. As a partner in one of the AI Alliance's sub-projects, HLRS is creating a technical infrastructure for networking existing data collections from industry, scientific research, and public administration into a searchable, cooperative data space. By implementing metadata standards and establishing a secure, decentralized architecture, the project will improve the findability of data resources. With the support of an additional grant from the Ministry, HLRS will also install hardware that will be used as a test environment for the development of simulation and AI applications. The results could support innovation across Baden-Württemberg in fields such as health, manufacturing, mobility, environment, and smart cities, particularly for the state's community of small and medium-sized enterprises. In addition, the AI Alliance will seek opportunities to link to related initiatives such as Gaia-X, a European initiative to facilitate data integration and sharing.

Baden-Württemberg Sustainability Weeks

On June 21, HLRS opened its doors as part of the statewide event Nachhaltigkeitswochen@Hochschulen BaWü 2023, offering the public a chance to learn about HLRS's approach to sustainability and steps it has taken to minimize its environmental impacts. As representatives of HLRS explained at the event, the basis for these efforts is its environmental and energy management systems, which are certified under the Eco-Management and Audit Scheme (EMAS). The event also highlighted projects that are investigating strategies for additional energy savings in its system operations. This includes developing plans to reuse the heat that its computers generate to warm other buildings on the University of Stuttgart campus. In a lecture titled "How Energy Efficient Software Can Support Sustainable Digitalization," HLRS research scientist Björn Dick explained how software design affects energy usage in HPC systems, and offered strategies for developing programs that consume less energy.

The digital twin of the Forbach facility included visualizations of data that predict construction noise.

Digital Twin Used in Design and Construction of Hydroelectric Facility

During the planning of a new pumped storage power plant in the Black Forest town of Forbach, energy company EnBW turned to HLRS for its expertise in digital twins. The collaboration began in 2011, when scientists in the HLRS Visualization Department supported a planning study aimed at identifying the best location for the project. The resulting digital twin of the final site incorporates data depicting not just the planned buildings and machinery, but also the surrounding landscape, local geology, and a physics-based model of construction noise. Engineers and other stakeholders viewed the immersive 3D visualization in HLRS's CAVE to discuss and optimize plans for the facility, and HLRS presented it as part of EnBW's community outreach efforts to nearby residents. Construction started in 2023, and the digital twin continues to be used to support construction management. The Forbach facility will expand on an existing hydropower plant to capture and pump water to a reservoir upstream of the turbines, providing energy storage that will help to ensure a stable power grid. After its expected completion in 2027, the project will support Baden-Württemberg's transition to renewable energy.

Prof. Yutong Lu and Prof. Michael Resch signed a memorandum of understanding between the National Supercomputer Center in Guangzhou and HLRS.

HLRS Renews Collaborations with NSCC-GZ and KISTI

In 2023, HLRS extended its collaboration agreements with the National Supercomputing Center Guangzhou (NSCC-GZ) and the Korea Institute of Science and Technology Information (KISTI). The agreements will enable continuing collaboration in research and education, and will facilitate joint workshops, courses, and the exchange of faculty, staff, and PhD students to increase expertise at HLRS and its partner institutions. The memoranda of understanding also identify scientific areas that will frame the cooperation. The agreement with NSCC-GZ highlights a shared interest in numerical algorithms using unstructured grids, emerging technologies in parallel programming and networking, and the immersive visualization of numerical simulations. The two centers will also explore how these approaches can support the development of applications for engineering and global systems science. Among the notable outcomes of the long-term collaboration between HLRS and KISTI has been the joint endeavor to develop tools aimed at enhancing the visualization of simulation results, which holds significant promise for industries reliant on simulation-based design and fabrication processes.

2023 Golden Spike Award Winners

On October 12–13, users of HLRS's systems presented and discussed their recent research at the 26th annual Results and Review Workshop. Each year, the conference highlights current applications of high-performance computing and strategies for optimizing their performance. In addition to covering traditional HPC simulation methods, several talks and posters also highlighted new approaches that combine machine learning with traditional simulation methods. At the conclusion of the meeting, Thomas Ludwig (German Climate Research Center), chair of the HLRS steering committee, announced the winners of the 2023 HLRS Golden Spike Awards, which recognize excellence in computational research and the use of HPC. Representing their respective projects, this year's Golden Spike Award winners were: Pascal Mossier (Institute for Aerodynamics and Gas Dynamics, University of Stuttgart) for "A High-Order Framework for Compressible Droplet Dynamics," Theresa Pollinger (Institute for Parallel and Distributed Systems, University of Stuttgart) for "The Sparse Grid Combination Technique for Higher-Dimensional Simulations at Extreme Scales," and Maximilian Jacobi (Institute for Nuclear Physics, Technical University of Darmstadt) for "Neutron Star Mergers: From Nuclei to Cosmic Explosions."

Pascal Mossier, Theresa Pollinger, and Maximilian Jacobi were named winners of the 2023 HLRS Golden Spike Awards.

News Highlights

Visualization of air quality in Sofia, Bulgaria. HLRS supported the development of CFD software for urban physics modeling.

Industry Summit Demonstrates Innovative HPC Applications in European SMEs

A two-day event in Berlin presented results of successful business experiments made possible by the FF4EuroHPC project.

Launched in 2020 with funding from the EuroHPC Joint Undertaking, FF4EuroHPC was created to promote usage of high-performance computing (HPC), artificial intelligence (AI), and high-performance data analytics (HPDA) in small and medium-sized enterprises (SMEs). Through two open calls and funding cycles, FF4EuroHPC executed 42 experiments in collaboration with 118 experiment partners from 22 European countries, covering industries including manufacturing, engineering, energy, environment, and healthcare. Selected SMEs received free computing access and user support to realize test projects at HPC centers across the EuroHPC network.

On October 18–19, 2023, the FF4EuroHPC community converged in Berlin to present results of these business experiments at its HPC Industry Summit. The meeting was jointly organized by FF4EuroHPC and EuroCC 2, a related project led by the High-Performance Computing Center Stuttgart (HLRS) that has been coordinating the development and networking of 32 National Competence Centers across Europe.

Held as a public event for representatives of industry and government, the HPC Industry Summit highlighted both the opportunities for innovation and the challenges that SMEs encounter when using high-performance computing. The event featured keynote talks by Daniel Opalka (Head of Sector Research & Innovation, EuroHPC Joint Undertaking), Javier Cordova Morey (Policy Officer, European Commission), and Michael Rafii (German Federal Ministry of Education and Research, BMBF), in addition to several expert talks and panel discussions looking at the future outlook for HPC, AI, and quantum computing in industry.

In an additional highlight of the event, participants in the FF4EuroHPC experiments presented their projects and achievements. Their success stories provided an inspiring look at several different ways in which HPC and related technologies can help to develop better products and services.

New applications for shipping, water management, manufacturing, and city planning

Paris-based engineering firm AYRO, for example, has been working with the support of FF4EuroHPC to improve the design of its Oceanwings wind-assisted ship propulsion technology. This new approach outfits ships, including those powered with conventional fossil fuel-driven engines, with sail-like structures. Artificial intelligence software monitors wind conditions continuously, and automatically adjusts the shapes of the sails to optimize aerodynamic performance. The company

At the HPC Industry Summit, leaders in small companies discussed successful applications of HPC, AI, and data analytics.

estimates that harnessing wind power can reduce fuel consumption in ships by up to 45 %, and AYRO's goal is to support the maritime industry in reducing its carbon footprint. With FF4EuroHPC support, it has been developing a tool that uses CFD simulations to optimize the installation of Oceanwings on all types of ships.

In another business experiment, ARESYS S.r.l. has been developing a new service that uses HPC and AI to provide real-time monitoring of water levels in reservoirs. Using remote sensing data gathered by the European Space Agency's Sentinel-2 satellite system, the approach creates an evolving, real-time 3D model of an entire reservoir, providing a low-cost method that water resource managers can use to respond to droughts or extreme weather events.

A collaboration between production machines manufacturer Gabler Engineering GmbH and software developer Kimoknow UG is focusing on how AI could accelerate product assembly. They have been using HPC to generate synthetic 3D CAD data and train an AI algorithm to visually recognize parts needed to build a machine. The results have helped Gabler to improve worker productivity, while Kimoknow now has a business model in which its software can be rapidly adapted for thousands of potential customers at low cost and sold on a subscription basis.

FF4EuroHPC has also supported a joint project involving SoftSim Consult and CFD software company Engys, who have been developing improved simulation workflows for urban physics modeling (UPM). The results will make it possible for SME service providers to offer a wider range of UPM solutions, and help engineers, architects, and construction planners to address more demanding regulations for sustainability and citizen well-being in cities.

Success stories will inspire uptake of HPC in industry

As the EuroHPC Joint Undertaking rolls out larger, more powerful supercomputers across Europe, increasing their usage and impact in industry is a high priority. "Success stories like the ones presented at the HPC Industry Summit are an important part of showing companies how HPC, AI, and HPDA could help them innovate and compete more effectively on the global market," said Dr. Bastian Koller, who is leading HLRS's coordination of FF4EuroHPC and EuroCC 2. "Our goal is that companies across Europe will find inspiration in these stories and begin to ask how they could use HPC to improve their own products and processes. And when they do, we and other centers are ready to provide services and support."

Learn more about FF4EuroHPC and find more success stories at ff4eurohpc.eu.

HLRS and WIKKI GmbH Win HPC Innovation Award

The award recognizes the researchers' development of the coherent file format, which will make simulation using OpenFOAM faster and more accessible for science and industry.

HLRS and WIKKI GmbH, a leading OpenFOAM consultancy and software development company, were named winners in the 19th HPC Innovation Awards. The award recognizes their development of a novel approach called the coherent file format, which enables OpenFOAM — a popular, open source programming framework for parallel computing — to run much larger simulations much more efficiently. Using the coherent file format, the team ran a large-scale OpenFOAM simulation on 4,096 nodes (524,288 CPU cores) of HLRS's Hawk supercomputer. This size exceeded the previous scaling record for a simulation using OpenFOAM by a factor of four.

In addition to setting this record, the achievement is notable because the coherent file format dramatically reduces workflow times for OpenFOAM users. This will make simulation more accessible to programmers in academia and industry who do not have access to large supercomputing systems. The code is available on an open source basis for the global OpenFOAM community at https://code.hlrs.de/exaFOAM.

The coherent file format was initially conceived by Dr. Henrik Rusche, CEO of WIKKI GmbH and one of OpenFOAM's lead developers. It accelerates simulation workflows by simplifying data management and eliminating several time-consuming steps that are typically required when using the programming framework. HLRS scientists Dr. Gregor Weiß, Dr. Andreas Ruopp, and Dr. Flavio Galeazzo worked closely with Dr. Sergey Lesnik and Dr. Rusche from WIKKI GmbH to develop the coherent file format for high-performance computing systems. The research was conducted within the project exaFOAM, which is sponsored by the EuroHPC Joint Undertaking and the German Federal Ministry of Education and Research.

Scientists at HLRS and WIKKI GmbH reached an unprecedented level of scalability using the popular CFD software OpenFOAM.

The HPC Innovation Awards recognize noteworthy achievements in the field of high-performance computing. Announced twice each year, the awards are coordinated by HPC market analysts Hyperion Research, Inc. This is the second year in a row that HLRS has received an HPC Innovation Award. In 2022, HLRS and the German Federal Institute of Population Research were named for their implementation of a simulation that predicted intensive care unit demand across Germany during the COVID-19 pandemic.

Defining a EuroHPC Hyperconnectivity Roadmap

EuroHyPerCon will identify and analyze the connectivity requirements of the EuroHPC ecosystem, delivering recommendations and specifications for an ultra-high-speed network infrastructure that connects European science and industry to the next generation of EuroHPC supercomputers.

As European supercomputers grow larger and artificial intelligence is used more widely, the amount of data that is produced and needs to be transferred between computing centers is likely to explode. Addressing this challenge will require an upgraded and agile end-to-end infrastructure that connects HPC users across Europe. Funded by the EuroHPC Joint Undertaking, a project called EuroHyPerCon aims to address this need. This collaboration — involving networking specialists Innov-Acts and Enomix together with HLRS — will develop plans for a new, ultra-high-speed network that connects EuroHPC systems with regional and national computing networks. At its conclusion, the project will deliver a comprehensive implementation roadmap to guide the JU in constructing the network.

Dr. Bastian Koller, General Manager at HLRS, is overseeing the center's contribution to the project. "EuroHyPerCon will bring stakeholders from around Europe to the table," he explains. "By gathering their insights, we will ensure that the implementation plan is realistic and feasible, and that it is capable of integrating changes in HPC technologies and usage that we anticipate over the next decade." Such planning means not just accommodating growing numbers of HPC users and larger quantities of data, but also incorporating the unique requirements of emerging technologies like artificial intelligence, edge computing, and quantum computing. The roadmap will both address technical features of the network and cover practical considerations that will be essential to the success of the project, including administration and cost planning.

"European science and industry are already highly networked, but as usage of high-performance computing, AI, and quantum computing grows, the risk is that bottlenecks in data transmission capacity will slow everyone down," says Dennis Hoppe, head of the HLRS Department of Converged Computing. "It is similar to what happens as cities grow and roadways become congested. At some point it becomes necessary to build an autobahn. We want to find out what it will take to do this for European HPC, and to make sure that users across Europe can easily access the onramp."

Machine Learning in Computational Fluid Dynamics: An Interview with Andrea Beck

Artificial intelligence offers great promise for computational engineering, but new data science approaches will need to be grounded in physical principles.

Prof. Andrea Beck is Deputy Director of the University of Stuttgart's Institute for Aerodynamics and Gas Dynamics, and serves as its head of numerical methods for computational fluid dynamics (CFD). With a background in aerospace engineering, she uses numerical methods to conduct basic research on compressible flows, which arise in extreme conditions such as when aircraft fly at supersonic speeds.

As a postdoctoral researcher, Beck was among the first to begin exploring how machine learning methods could be combined with large-scale simulation software to improve the simulation of compressible flows. In 2023, her lab released Relexi, a method developed with HLRS and Hewlett Packard Enterprise that seamlessly incorporates reinforcement learning into a computational fluid dynamics solver. In this interview, she discusses the possible applications and some challenges of using machine learning in CFD, as well the opportunities that HLRS's next-generation supercomputers will offer for her research. The interview was edited from a conversation for readability.

Professor Beck, your work focuses on numerical methods for computational fluid dynamics. How does one simulate a compressible flow?

Many engineering applications deal with multiscale problems. They are like the famous butterfly effect, where small, localized changes can have large impacts across distance and time. This means, for example, that when simulating an airplane in flight you would like to use a very fine computational mesh – ideally, surrounding the entire aircraft in virtual 1 mm cubes – to simulate at high precision how a compressible turbulent flow interacts with the plane. In practice, however, a simulation at that scale would take far too long, even with HLRS's Hawk supercomputer or next-generation systems like Hunter or Herder. Instead, we reformulate the problem and resolve it to an intermediate scale, sometimes called the mesoscale.

Approximating in this way means that information is lost, and so to make such a simulation useful, you need a model to account for the fine-scale effects that you don't capture. This is sometimes called scale bridging or closure modeling. To take a simple example, a colleague simulated how people evacuate a football stadium. You can essentially model people's movements like a flow of water, and although it's easy enough to model a smooth stream out of the building, it is hard to model individual people bumping into one another. This is important, though, because a single interaction between two people at a doorway could block an exit. To account for this, his closure model involved using Gaussian noise to recover these fine-scale effects. We do a similar thing when simulating compressible flows.

In your research, you have been exploring how machine learning could be used in identifying such closure models. What would the advantage of this approach be?

All engineering problems can be described by applying the laws of physics. To do so, we use complicated mathematical expressions that can only be solved using a supercomputer, which is why we need HLRS. Over the years, people have tried to find closure models for turbulence based on mathematical and physical reasoning. Essentially, we want to condense this complex phenomenon into a single, reliable number that can be input into a flow solver. However, there is still no clear and consistent model that represents turbulent flow in all cases. Instead, there are many models, and every researcher has his or her favorite.

Deriving such models is precisely what machine learning should be good at, and beginning a few years ago people started using it to extract low-dimensional features from high-dimensional data sets. The idea is that if we have enough data that have been generated by the process we are trying to model, then all of the necessary information about physical laws should be embedded in the data.

What do you see as the limits or frontiers of machine learning for CFD at the moment?
While it is true that nature will always obey physics, the problem with machine learning is that it does not directly replicate this data, but can only make approximations based on the data on which it is trained. For very specific cases, you can use machine learning to build astonishingly accurate models. For engineering applications, though, we want machine learning models that don't just reproduce what we measured, but also give us reliable answers in situations for which we did not provide data.

From my perspective, and I have been guilty of this as well, we've put a little bit too much emphasis on saying it's all in the data. Because machine learning uses a black box approach, there is currently no consistent way of making sure that if you use it, the answer you get doesn't violate a basic law of physics. In the case I described earlier, the fact that two people will not merge together when they bump into each other is a physical law, but a neural network doesn't know that and there is no guarantee that it will enforce the corresponding physical constraints. Or to take an example from aerospace engineering, it is important that your machine learning algorithm knows that gravity points straight down, and not at a slight angle. A naïve algorithm can't necessarily know this just from the data, though.

Prof. Andrea Beck has been a leader in exploring the potential of AI using HLRS's systems.

During training, all that machine learning models see is the neighborhood of their data. Often, however, we don't know what situations these models will actually be confronted with when we use them. When you need to apply your model to data from the next town, so to speak, we need models that contain physical constraints. With physical constraints we know that the model understands that gravity is always the same here as it is there.

A couple of years ago, people were optimistic that we wouldn't have to do simulation anymore, but just machine learning. What we see now, however, is that we have to shift from purely data-driven models to physics-consistent models that are data-augmented or data-informed. The models should be based in physics and the data should help us to nudge them in the right direction. The approach that my lab is using is to keep the original simulation methods, but to use machine learning as an augmentation device that replaces some parts of the equation, where we know what the physical properties are.

Can you give an example of how such a physics-consistent, data-driven model might work in practice?
In an approach called Relexi, which we developed at HLRS in collaboration with staff at Hewlett Packard Enterprise, we combine traditional simulation approaches and machine learning in a single method. This is because in an aircraft turbine, for example, there is very slow flow, hot flow, cold flow, and it could contain water or steam. The physics is very sensitive to tiny changes. The dimensionality of the problem is huge, making it nearly impossible to come up with long lists of data tables that tell the machine learning algorithm, "If you see this, output that." You would have to have so much data for so many different situations.

Instead, we continue to use a high-order flow solver called FLEXI that we've used for a long time, but it now also includes a model based on reinforcement learning. This requires running the flow solver with a machine learning model in the loop. The machine learning model recognizes what actions the solver took and what the results were. This requires an efficient HPC implementation – the solver runs on the CPU, the machine learning model on the GPU, and both have to communicate constantly and exchange data about the learning process, the current model, and the observed results. I run my solver with my untrained machine learning model in place and check the output to see if it is what I expect. If it's not, the machine learning model improves itself.

This process is similar to how someone learns to ride a bike. If you sit on a bike for the first time and try to pedal, you will fall over, but even with the first pedal stroke you sense that you did something correctly. You can then improve upon that strategy and repeat it until after several attempts you can ride without falling. To think about this in terms of a physical model, the simulation would include components of the system describing the bike's configuration, gravity, and your physiology. The machine learning part would use this basic information to learn how to move your joints. This is also the method that people usually use to train robots or self-driving cars.

Upcoming supercomputers at HLRS will include more graphics processing units, which could both accelerate simulations and make it easier to combine simulation and machine learning methods. What could the availability of GPUs mean for the future of your research?
The AI revolution in the last 10 years only came about because the mathematical operations they need to do are perfectly aligned for the GPU. Relexi has already been running on the AI components of Hawk, so we have already made good use of different architectures and are aware of the different strengths that they offer. For us, it will be great to have both code and machine learning on the GPU already.

The upcoming GPU-accelerated architecture at HLRS is also an exciting opportunity because the increase in computing power means that our code will run two orders of magnitude faster. This, in turn, means that the types of problems we will be able to tackle have grown by two orders of magnitude.

To make flying more sustainable, for example, the airplane of the future will look very different than what we are used to. Predicting the flow physics over such configurations to the level that we need is not possible, and we still work with approximation methods. Herder, or maybe even Hunter, will allow us to compute challenging flow situations across the whole wing of an aircraft to a degree that is currently not possible. This could also help to better predict flow situations when conditions become dangerous, for example when you fly too fast or things go wrong for physical reasons.

We are aware that porting to the new structure of Hunter and Herder will be a challenge. We have to refactor most of our code, rethink most of the data structures, and rethink all of the loops, which take effort and time. The benefit, though, is that it will dramatically increase the size of the computations we can do.

In any case, we are excited about the opportunities provided by the new generation of machines at HLRS. We will be able to achieve results faster and at a lower economic footprint. It is an exciting time for aerospace engineering at the moment and new ideas and architectures are being explored at a great pace; for example, NASA's X59 supersonic demonstrator and the push towards electrified flying. We need simulation tools that can keep pace!

HLRS Achieves Certification for Information Security Management

The ISO 27001 international standard defines requirements and best practices for the safe handling of proprietary data.

Dr. Martin Hecht (center) oversees HLRS's information security management system and led the ISO 27001 certification process. Also pictured (l–r): Thomas Beisel (Head, HLRS Division of Software and Systems), Prof. Dr. Michael Resch (Director, HLRS), Dr. Bastian Koller (Managing Director, HLRS), and Inna Wöckener (Finance and Project Administration, Sustainability, HLRS).

In June 2023 HLRS completed certification under the International Organization for Standardization (ISO) 27001 standard for information security management. ISO certification confirms that HLRS has implemented a comprehensive information security management system (ISMS), which details technical and organizational measures for identifying security threats to stored data, and for preventing and reacting to attacks. This includes protecting high-performance computing systems from unauthorized access and usage.

"Achieving ISO 27001 certification confirms via an external audit that HLRS follows industry-standard best practices for information security management," said HLRS Director Prof. Michael Resch. "It should give users of our high-performance computing systems the confidence that their data are well protected, and should reassure our funders that our supercomputer and other systems are not used in inappropriate ways. Finally, this certification will ensure that HLRS updates and improves its security measures on a continual basis to address new threats that might arise in the future."

The scope of HLRS's information security management system includes the provision of computing time on its high-performance computers as well as the supporting processes for operating the production environment. The ISMS establishes policies that address potential security risks at all levels of the organization, defines procedures for identifying and reacting to security breaches, and establishes clear roles and responsibilities for managing information security. It also provides a framework for tracking the effectiveness of security measures and requires that the entire HLRS staff be informed of and comply with information security policies. Moreover, the ISMS sets guidelines for HLRS's product suppliers and their subcontractors, ensuring that they adhere to stringent secrecy and procedural requirements.

Certification under ISO 27001 complements steps that HLRS took prior to 2021, when it completed a data security assessment in accordance with the TISAX (Trusted Information Security Assessment Exchange) framework. TÜV NORD CERT GmbH performed HLRS's ISO 27001 audit, which consisted of a thorough documentation review, a site visit, and an examination of the center's security processes.

Partnership with Ukrainian University Persists in Difficult Times

A conference at HLRS celebrated 50 years of scientific collaboration between the Donetsk National Technical University and the University of Stuttgart.

In 1973, at a time of rapprochement between East and West, young Ukrainian and German scientists first established contact between the Donetsk National Technical University (DonNTU) and the University of Stuttgart. 50 years later, during a time of war, the partnership continues. At a full-day conference held in November 2023, German and Ukrainian researchers met at HLRS to reflect on the history of their scientific cooperation, and to discuss ongoing collaborative research.

Following years of collaboration, the DonNTU and the University of Stuttgart signed a formal cooperation agreement for joint research and teaching in 2000. Since then, numerous reciprocal visits have taken place, a joint conference was held in 2013, and more than 100 researchers from Ukraine have received scholarships within the framework of the Leonhard Euler Scholarship Program of the German Academic Exchange Service (DAAD).

The transnational partnership was initiated by Dr. Volodymyr Sviatnyi, Professor of Simulation Science and Head of the Department of Computer Engineering at DonNTU, and a recipient of the Medal of Honor of the University of Stuttgart. Dr. Svjatnyj fled Ukraine following the Russian invasion in 2022, and at the invitation of HLRS Director Michael Resch has been working as a visiting scientist at the University of Stuttgart. Through personal support for Ukrainian colleagues, by accepting researchers who fled Ukraine, and, in particular, within the framework of the DAAD program "Ukraine digital", the University of Stuttgart has also been helping DonNTU maintain its teaching activities.

The conference at HLRS began with welcoming remarks by University of Stuttgart Rector Wolfram Ressel and DonNTU Rector Yaroslav Lyashok, and was followed by presentations by scientists from many disciplines focusing on recent work that has grown from this close relationship. Commenting on the event, Prof. Resch remarked, "We are proud that the partnership between the scientific communities in Stuttgart and Donetsk has endured over the years and continues to be strong despite crises and changing political situations. Science benefits immeasurably from trusting, international cooperation, and our partnership with DonNTU has greatly benefited the research in both Germany and Ukraine. In view of the ongoing war, it is very important to us that we continue."

Representatives of the Donetsk National Technical University and the University of Stuttgart met at HLRS 2023 to celebrate 50 years of collaboration. l–r: Yaroslav Lyashok (Rector, DonNTU), Prof. Volodymyr Sviatnyi (Head, Department of Computer Engineering, DonNTU), Prof. Wolfram Ressel (Rector, University of Stuttgart), and Prof. Michael Resch (Director, HLRS).

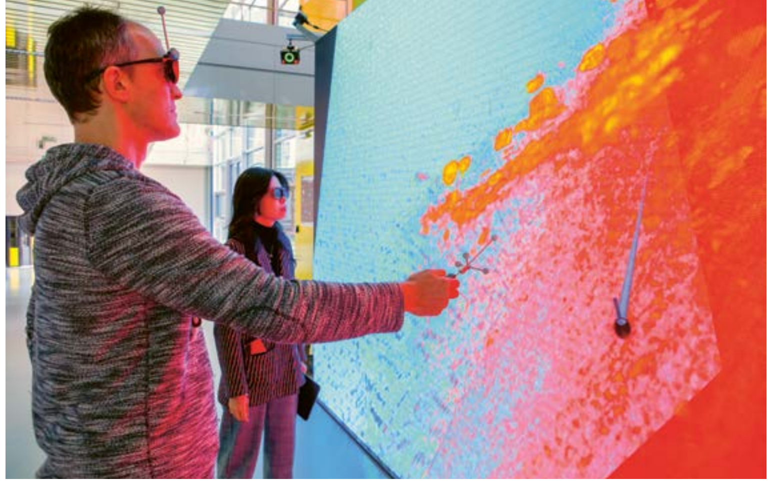

Artist Hiền Hoàng (r) is a resident in the S+T+ARTS AIR program. She is developing a new, immersive multimedia artwork in collaboration with HLRS.

HPC for European Culture and Creativity

As a lead coordinator in European initiatives, the Media Solution Center Baden-Württemberg helps to identify opportunities for supercomputing in the arts and creative industries.

Co-founded by HLRS, the Media Solution Center Baden-Württemberg (MSC) has been building a network of expertise to explore the potential of high-performance computing (HPC) and visualization in the arts. Recently, the MSC has expanded on this mission by assuming leading roles in several European projects aimed at promoting innovation and economic development through the culture and creative industries.

In one role, the MSC is a co-founder and lead partner of an EU-wide Knowledge and Innovation Community called EIT Culture & Creativity (EIT CC). Within this initiative it also co-chairs a Germany-wide consortium called Innovation by Creative Economy (ICE), one of 50 partners within EIT CC. As a founding partner of the MSC, this makes HLRS is the only high-performance computing center involved in EIT CC, positioning it to provide models for utilizing HPC in the arts.

In September 2023 the MSC organized a conference at HLRS that brought together the EIT CC directors and advisory board to discuss strategies for the initiative when it formally launches in 2024. MSC Managing Director Mathias Hauser explained, "Together with EIT, we want to network the know-how that exists at HLRS with that found across Europe," he says. "The support that HLRS provides is helping the culture and creative community to gain a better understanding of what is possible."

In parallel, the MSC coordinates S+T+ARTS AIR, an EU-wide residency program that for the first time made supercomputing resources at HLRS and at the Barcelona Supercomputing Center available to artists. In 2023 S+T+ARTS AIR completed an open call for proposals in which 10 artists and artist groups were selected from 127 entries. HLRS will participate in two S+T+ARTS AIR projects focusing on human-AI interaction and urban ecology.

The Media Solution Center was also involved in the development of the exhibit Renaissance 3.0, which was on view for most of 2023 at the Center for Art & Media (ZKM) in Karlsruhe. Curated by Peter Weibel, the exhibition explored new alliances between artists, scientists, and engineers based on a set of shared tools, including computing technologies. After closing at the ZKM, the show is scheduled to travel to Paris and Budapest.

Supercomputing Academy Offers Path to HPC Expertise

Two certified HPC Experts reflect on the benefits of HLRS's training programs for their education, research, and career development.

Maha Badri and Alaa Bejaoui were both born in Tunisia, but they also have much more in common. As young engineers who in 2023 completed their masters studies in aerospace engineering at the University of Stuttgart, they share a passion for mathematics, artificial intelligence, parallel programming, and high-performance computing. During their bachelor's studies, Badri and Bejaoui began to take advantage of the comprehensive high-performance computing (HPC) training program offered at the High-Performance Computing Center Stuttgart (HLRS). They became so intent on gaining a thorough understanding of how to program HPC systems that during their master's studies they became the first participants in HLRS's Supercomputing Academy to be awarded the certification "HPC Expert."

From the beginning of their university studies, the engineers recognized the importance of parallel programming, but saw that the necessary training would not take place within their academic program. "When we began to study in 2013, we quickly noticed that engineers often did not place a high value on programming," Bejaoui recalls. "Because of this gap we decided to take as many workshops at HLRS as we could." Initially, they took courses in HLRS's core HPC training program. After spending time in Canada to study artificial intelligence at the Polytechnique Montréal, they returned to Stuttgart in 2020 and discovered HLRS's newly founded Supercomputing Academy, and welcomed the opportunity to dive deeper. According to Badri, "We already had basic knowledge, but we wanted to expand on it."

A deep dive into HPC and AI

Whereas the HPC classes they had taken previously at HLRS were held over 3–5 days, courses offered by the Supercomputing Academy typically extend over approximately 6 weeks in a blended learning format, enabling participants to learn at home in parallel with their other professional responsibilities. It is possible to take individual courses, but for participants interested in a comprehensive training experience, the Supercomputing Academy offers certification as an "HPC Expert." The program's content is modular, enabling participants to select from a menu of courses to explore their individual needs and interests. Badri and Bejaoui elected to pursue certification in the category "HPC Developer." In the process, they gained knowledge about MPI, Open MP, communication, and node-level performance. Because of their interest in artificial intelligence, they also took courses in data management and data analysis.

"From the beginning our goal has been to complete doctorate degrees," Badri explained, "and the idea (to pursue Supercomputing Academy certification) was to prepare ourselves as well as possible… The fact that the program requires tests was important for us, be-

HLRS's training programs helped Alaa Bejaoui und Maha Badri to develop essential skills for parallel programming and artificial intelligence that they will continue to use in their doctoral studies.

cause taking a class without a test is not the same as when you study to pass a test. You go much deeper and learn more intensively."

Because Badri and Bejaoui were active in the Supercomputing Academy during the COVID-19 pandemic, they unfortunately missed the personal contacts with instructors and students that they enjoyed during earlier on-site courses at HLRS. Nevertheless, the design of the Supercomputing Academy was an ideal way for them to develop their HPC skills while also completing their academic program. Badri says, "I could concentrate on my master's thesis during the afternoon and at night I would watch the (Supercomputing Academy) course videos… If everything had taken place in person, I can't imagine how it would have worked."

A strong foundation for career development

Now that they have completed the Supercomputing Academy's training program, Badri and Bejaoui recognize that the experience has had many benefits. "It was only through the Supercomputing Academy that we gained a deep understanding of HPC hardware, and this has made us more deliberate in our programming, particularly with respect to parallelization," Badri says. In practice, they used their new skills to apply MPI for Python programming, making it possible to parallelize an approach that combines simulation and reinforcement learning.

During job interviews, they also sensed that the certification offered through the Supercomputing Academy made them more attractive job candidates. Badri recalls, "In all of my interviews, it was a feature in my resume that impressed potential employers. They always asked what kind of certificate it was. Was it for a single course, or for a week-long course? When I explained all that we had learned, they were very impressed at how comprehensive the program was."

Through the experience they have gained through their academic studies, job searches, and work at a startup company, Badri and Bejaoui sense that there is a growing demand for the training that the Supercomputing Academy offers. "In fields such as aerospace engineering that rely on simulation, HPC is well known, as simulations are not possible without an understanding of HPC. With the rise of AI, however, other working groups are realizing that they need expertise," Badri explains. Bejaoui agrees, and thinks that the training they received has positioned them well. "HLRS and the Supercomputing Academy are good partners for developing these skills," he says.

The two are currently looking forward to starting their doctoral studies, which will take them in exciting new directions. Badri will soon join the Potsdam Institute for Climate Impact Research, where she will use artificial intelligence to develop vegetation models that consider interactions between the plant kingdom and the atmosphere. Bejaoui has secured a doctoral position at the Charité University Hospital in Berlin, where he will collaborate with medical doctors to develop methods that use artificial intelligence to interpret CT scans, tools that could lead to better patient care. Both anticipate that the HPC expertise they have gained will continue to be necessary in their respective projects.

Over the longer term, Badri and Bejaoui also hope to carry their knowledge back to their home country, Tunisia. "We have good scientists, very smart minds… with expertise in mathematics and physics. And now computer science and HPC are coming. I feel like we now have the complete chain, from mathematical modeling, to implementation, to optimizing the implementation on hardware," Bejaoui says. "We want to build bridges between Germany and Tunisia, and Germany could also benefit."

New HLRS training concept integrates Supercomputing Academy and compact courses

HLRS continues to develop its training concept to offer course participants a menu of options that best address HPC-users' professional interests and needs. According to Lorenzo Zanon, head of the HLRS Department of Training Scalable Algorithms, the center has plans to combine successful elements of the Supercomputing Academy and its compact HPC training courses: "This could make it possible, for example, to complete exams at the end of compact training courses that will count toward the HPC Expert certification. Conversely, HPC users could register for blended learning courses in the Supercomputing Academy without the requirement of taking a final exam." In this way, HLRS aims to provide a flexible training program that is as accessible as possible. More news about this concept should be available in 2024.

Learn more about HLRS's training program at www.hlrs.de/training. Additional information about the Supercomputing Academy is available at www.supercomputing-akademie.de.

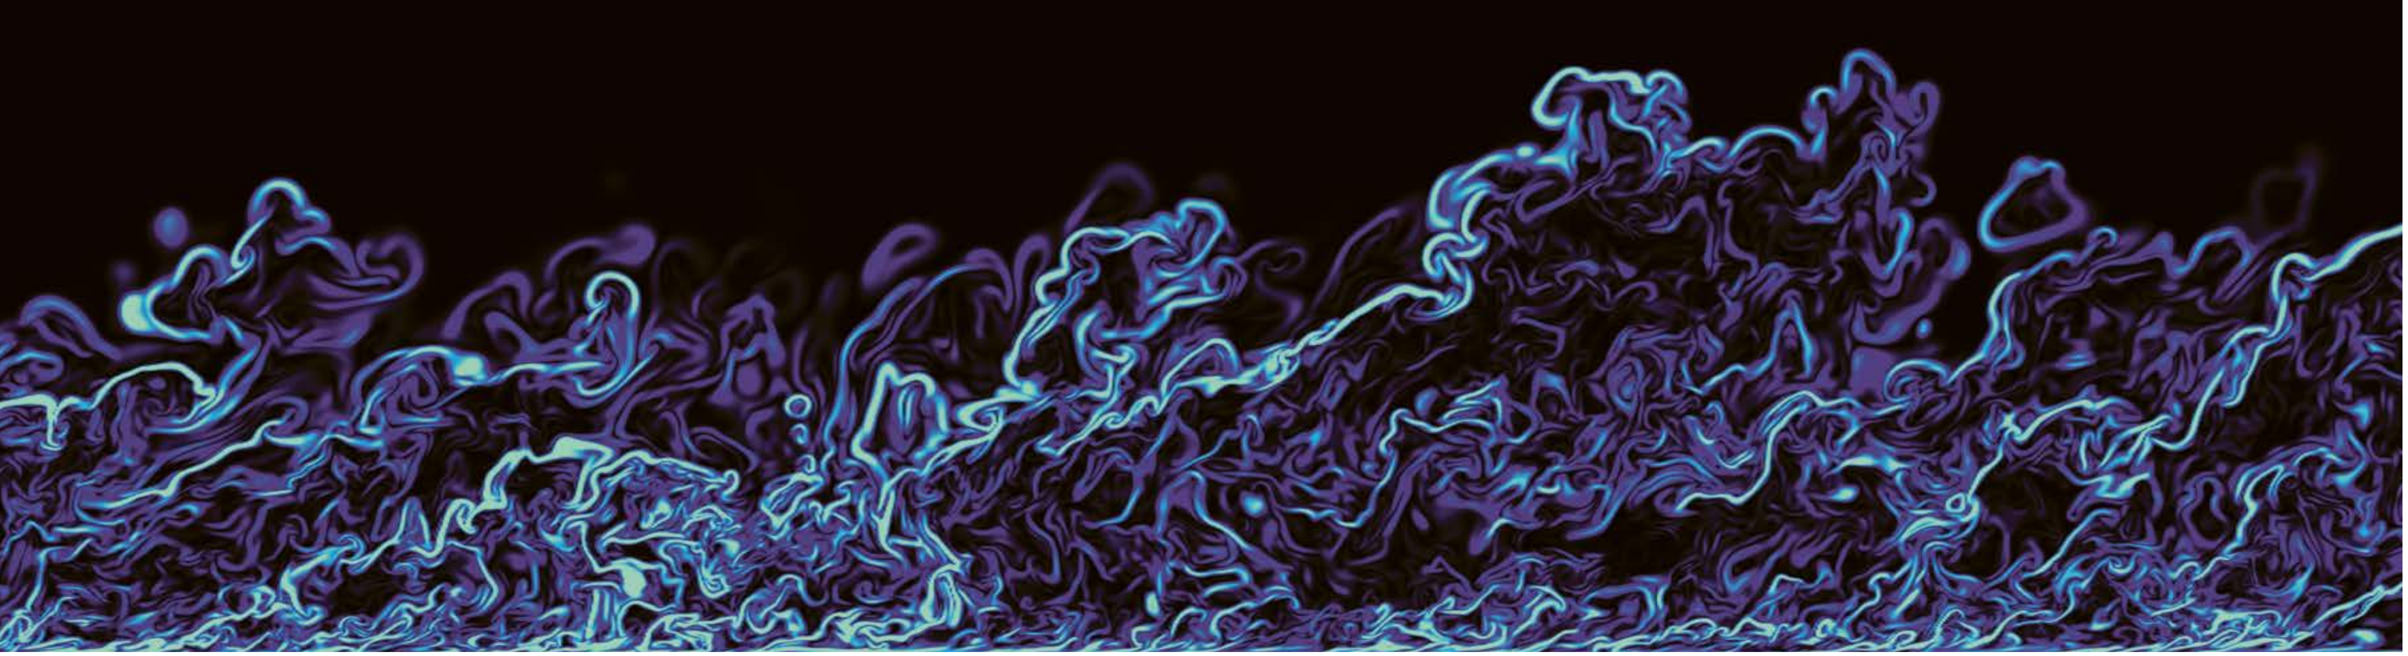

Using HLRS's Hawk supercomputer, scientists create high-resolution numerical simulations of turbulent flows within just 1 millimeter of a surface.

When a Code Goes "XXL"

With the help of HLRS user support staff, researchers ran their software efficiently at the largest possible scale on the Hawk supercomputer, opening new opportunities for scientific discovery.

Researchers in the University of Stuttgart's Institute of Aerodynamics and Gas Dynamics (IAG) have used HLRS's supercomputers since the center's founding. Since 2008 this has included running NS3D, a software code for simulating compressible turbulent flows. Using a computationally intensive approach called direct numerical simulation, NS3D provides high-resolution data needed to study fundamental physical processes in fluid dynamics.

As HLRS's supercomputers have become more powerful over the years, NS3D's capabilities have improved dramatically. When it was first developed, IAG scientists ran NS3D at HLRS to predict what happens when an orderly, smooth laminar flow breaks down into a complex turbulent flow. Using the center's Hawk supercomputer, the researchers can now study much finer-scaled phenomena in the turbulent boundary layer, an approximately 1-millimeter-thick region where free-flowing air moving at supersonic speeds is disturbed by its interaction with a surface.

With each new generation of supercomputer at HLRS, NS3D's programmers have adapted their code to run at the highest possible performance on a new architecture. The user support staff at HLRS has played an important role along the way, providing expertise in high-performance computing and in-depth knowledge of HLRS's systems. Through code optimization workshops and ongoing mentoring of NS3D's programmers, staff at HLRS have also gained a deep understanding of the software.

In the summer of 2023, this close collaboration helped the IAG researchers to reach a new height in scaling. During an "XXL day" at HLRS, they for the first time ran a simulation efficiently on 524,288 compute cores, the largest simulation that is possible on Hawk in regular production.

According to Björn Dick, a scientist in the HLRS user support team, facilitating scientific applications that utilize such massive power is an important part of the center's mission. "Most computers do not have the power that Hawk offers, and running extra-large-scale jobs is what the machine is built for, so we like to foster usage in this way whenever we can," he says.

Code optimization is important for simulations of any size, but it is essential when running codes like NS3D at large scales. Because they are especially resource-intensive, optimizing their performance is an important part of improving HLRS's energy efficiency. It also ensures that HLRS's shared HPC resources are available for as many users as possible, producing the greatest possible scientific impact. Researchers who optimize their codes for HLRS's systems also gain the greatest benefit from their allotments of computing time.

Recent optimization efforts involving the HLRS user support staff and the NS3D team have focused on improving the usage of MPI and on the optimization of I/O bandwidth for collective writing on the Lustre file system. HLRS has also been working closely with Hewlett Packard Enterprise and AMD to begin porting NS3D to graphics processing units (GPUs), an adaptation that will be necessary to program HLRS's next-generation supercomputers, Hunter and Herder.

The ability to optimize and scale NS3D to more powerful HPC systems promises to open new opportunities for research at IAG. As scientist Jason Appelbaum explains, running NS3D on Hawk reveals relationships among all of the important variables in a turbulent flow. Taking the software to the XXL scale will enable his team to achieve new insights. "These advances allow us to investigate turbulence at high Reynolds numbers over long time scales. The results will provide insight into fundamental statistical relationships of wall turbulence that are useful for researchers developing turbulence models," he remarked. In the long run, the knowledge they gain about the underlying physics could eventually help engineers in industry to design airplane wings with lower skin friction drag, improving fuel efficiency in air travel.

User Research

High-Resolution Simulations Enable Better Understanding of Extreme Weather

Atmospheric scientists use HLRS's Hawk supercomputer for research that could improve the modeling of risks associated with climate change.

In July 2021, a heavy rainstorm caused devastating flooding in the Ahr River valley in western Germany. The water rose up to seven times higher than estimated 100-year flood levels, leaving small towns in ruins and killing at least 135 people. It became a wakeup call about the need for better models of flood risks, and brought new urgency to an awareness that local communities need to prepare now for the future effects of climate change.

“We clearly expect in the future that with a warming climate, the intensity of events like the Ahr River flood will further increase,” says Hendrik Feldmann, a specialist in regional climate modeling at the Karlsruhe Institute of Technology. “What we cannot predict is exactly where such extreme events will happen in the next 100 years.”

Developing tools that help communities to prepare for and manage the emerging risks that climate change holds is the focus of Feldmann's research. “We aim to statistically assess the frequency, intensity, and spatial distribution of extreme weather events under climate change,” he explains. To do so, he is investigating how a powerful, high-resolution simulation approach called convection-permitting modeling (CPM) could improve scientists' ability to assess and predict such risks more precisely at the local level. Using the Hawk supercomputer at HLRS, he and his colleagues have been investigating whether CPM could offer benefits over other modeling approaches. Their work has also begun generating information that could help guide future adaptation.

Flooding in the Ahr River valley in the summer of 2021.

Climate modeling is becoming more precise

Although widely used regional climate models (RCM) can provide useful insights about climate and extreme events, they cannot accurately represent the characteristics of localized heavy rainstorms. The typical mesh size used in RCM simulations is 12 to 25 km, which does not provide sufficient resolution to capture important features that affect the impacts of extreme weather. For example, RCM cannot resolve localized details when simulating factors like temperature, precipitation, topography, or soil conditions, all of which can mean particular risks for specific communities.

Convection-permitting modeling could potentially represent the risk of extreme events more realistically. In CPM, atmospheric conditions are simulated at a scale of less than 4 km. This means that in addition to capturing a much more refined model of the atmosphere and its relation to the landscape, CPM provides scientists sufficient resolution to study convection, the atmospheric process in which warm, moist air rises to form rain clouds. Using this higher-resolution approach, they can simulate the evolution of convective processes without the parameterization necessary in coarser models.

Large supercomputers have made it feasible to use CPM to simulate the effects of extreme events in specific areas on climate timescales. The downside, however, is that CPM is much more computationally demanding that RCM. Currently, it is a huge effort to use this approach on global or long-term scales because of the time it would take. This means that as they look to the future, atmospheric scientists would like to better understand when and how CPM could be used most effectively.

What factors cause extreme precipitation?

In a recent paper published in the journal *Weather and Climate Dynamics*, Dr. Alberto Caldas-Alvarez, a former KIT postdoc who coauthored the publication with Feldmann, compares results of convection-permitting modeling to those of regional climate simulations. The team's goal was to determine whether CPM could, in fact, simulate local properties of extreme rainfall events better than regional models.

The research was conducted as part of the German Federal Ministry for Education and Research (BMBF)-funded ClimXtreme project, in which Feldmann coordinates the module "Physics and Processes." The project is focused on identifying atmospheric patterns associated with extreme weather events.

By analyzing historical weather data gathered in the European Alps between 2000 and 2015, Caldas-Alvarez first identified past extreme precipitation events. He then used a statistical approach called principal component analysis to determine the specific atmospheric conditions that were typically associated with them. In the next step, he analyzed regional climate and convection-permitting simulations of the same period to see how well their representation of extreme storms corresponded to actual historical measurements.

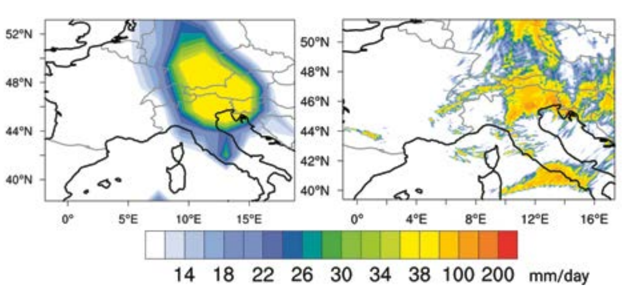

Comparison of two simulations of a historical extreme precipitation event at 200 km resolution (left) and at 3 km resolution (right). The simulation at convection-permitting scale provides a much more realistic picture of precipitation variation due to local atmospheric conditions and topography, including differences north and south of the Alps.

The team found that convection-permitting modeling was, indeed, better than regional climate modeling in identifying extreme precipitation events. At the same time, they also detected relevant differences in how the two modeling approaches simulate the processes responsible. The findings present open questions that future research will explore to better understand the strengths and limitations of each approach.

Despite these questions, CPM's higher accuracy in representing extreme precipitation events suggests that it could help to develop regional climate adaptation strategies. "Convection-permitting modeling provides a much more realistic distribution than traditional regional climate modeling when compared to observations," Feldmann says. In the second phase of the ClimXtreme project, he and his team aim to further increase understanding of extreme events such as heavy precipitation, heat waves, and wind storms.

Future heat waves in southern Germany

As a member of another BMBF-funded project called RegIKlim (Regional Information for Climate Action), Feldmann has been working to develop climate information for practical use that could help local communities across Germany develop better planning and risk management strategies. He co-coordinates the subproject NUKLEUS (Nutzbare Lokale Klimainformationen für Deutschland), whose goal is to generate, evaluate, and make available CPM climate simulations to support decision making. In addition, he is a principal investigator in a RegIKlim companion project called ISAP (Integrative city-regional adaptation strategies for a growing polycentric region: Region Stuttgart), which is using high-resolution climate information to support climate change adaptation in the area surrounding Baden-Württemberg's state capital.

Ensemble median of the number of days per year with strong human heat stress in southern Germany, defined by the Universal Thermal Climate Index (UTCI) for 1971-2000 (a), global warming of 2 °C (b), and global warming of 3 °C (c). Simulations show that the Rhine River valley will be warmer than other areas.

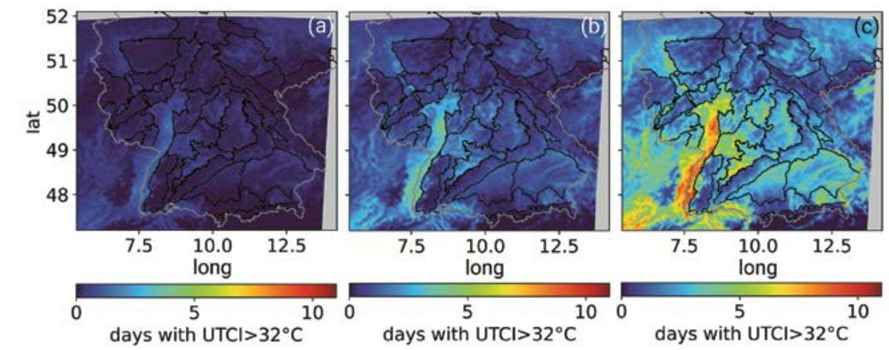

For a publication in the journal *Natural Hazards and Earth Systems Sciences*, Feldmann and the ISAP/NUKLEUS team generated an ensemble of regional climate simulations of future heat waves in southern Germany. By combining different regional convection-permitting climate simulations run on HLRS's Hawk supercomputer, they once again investigated the benefits of the CPM approach. Their findings provide a host of insights into what the future might hold under climate change.

The ISAP/NUKLEUS team found, for example, that the probability of heat waves in Germany will increase, particularly in the late summer. Their severity is also likely to become significantly worse in a warming climate. The resulting heat stress will depend on features in the landscape, with the most severe stresses likely in the Rhine River Valley. Convection-permitting modeling provided the regional specificity needed for this nuanced climate change projection.

Simulation provides necessary insights for planning

In another study, Feldmann and colleagues investigated the historical context of the Ahr flood, comparing it to observational data of past heavy precipitation events. They found that although historical events with comparable severity occurred in 1804 and 1910, they weren't recorded in the observational network used to determine current flood risks, which only dates back approximately 70 years. One lesson of their study is that risk modeling should begin to incorporate hydrological conditions and changes in land use that have occurred over the years. In the future, they suggest, it will be important to combine such information with better models of changing climate conditions to be able to make better predictions of risk.

"We need climate data to derive statistics of where and under which circumstances heavy precipitation events are more likely to occur, just as for other atmospheric extremes such as heat waves. This is valuable information for decision makers and industries such as urban planning or construction," Dr. Caldas-Alvarez observes. "Investing in high-performance computing and high-resolution modeling can help reduce the large uncertainty associated with such weather extremes. This aids in making decisions that could better protect human lives, infrastructure, and investments."

These days, Caldas-Alvarez is no longer at KIT but works in Innovation and Development at EDP, a multinational renewable energy company. In this role, he uses climate science in an applied way, developing analytical tools to improve solar and wind yield predictions. "Being in a renewable energy company every day you experience how important it is to have access to accurate and highly-resolved predictions and simulations," he reflects. "We need to get the best quality we are capable of, as sometimes biases smaller than 1 m/s make the difference between pursuing or dismissing wind projects."

By continuing to study and refine convection-permitting modeling, Feldmann and his partners aim to improve the scientific tools needed to make better projections in the future.

Molecular Simulations of Ammonia Mixtures Support Search for Renewable Fuels

Using HLRS's Hawk supercomputer, scientists at the TU Berlin generate valuable thermodynamic data for chemical engineering research.

Ammonia (NH_3) is an important molecule with many applications. The end product of the famed Haber-Bosch process, it is commonly synthesized to capture nitrogen for fertilizers, and is used for refrigeration, in cleaning products, and in the production of pharmaceuticals. Recently, this modest molecule has also attracted interest as a potential resource for addressing one of today's most pressing challenges – the need for reliable and abundant renewable fuels.

Ammonia is stable and safe to handle, is combustible, and contains the largest fraction of hydrogen of any molecule except for pure hydrogen itself. These factors promise to make it a feasible alternative to the carbon-based energy carriers that are driving climate change. Research has begun to explore how ammonia could be used to directly power engines, gas turbines, and hydrogen fuel cells, for example. It is also believed that ammonia could be used to store energy for times when other renewables like wind and solar power cannot meet demand.

Much is known about ammonia, but this interest in using it as a fuel has initiated a search for new ammonia technologies. This has, in turn, led to an increased need among chemical engineers for accurate data describing ammonia's fundamental thermodynamic properties. Such properties include a wide variety of measurable traits such as phase equilibria, density, or heat capacity, for example, that characterize physical systems and determine how chemical processes work. In the case of ammonia, engineers would also like to have better knowledge of how such properties change when mixing ammonia with other molecules. Such knowledge could help them to optimize processes and operating conditions.

Dr. Jadran Vrabec, currently the director of the Institute for Process Sciences at the Technical University of Berlin, has spent much of his career using high-performance computing (HPC) to investigate thermodynamic properties at the molecular level. "Thermodynamic properties are 100 % determined by molecular interactions," he explains. "And because these interactions happen so fast and at such a small scale, it is only possible to study them by performing large simulations using supercomputers."

In a recent paper published in the *Journal of Chemical & Engineering Data*, he and coauthor Erich Mace of the TU Berlin report on the results of simulations focused on the thermodynamic properties of mixtures containing ammonia. Produced using the Hawk supercomputer at the High-Performance Computing Center Stuttgart (HLRS), their results add valuable data that could support the development of new applications of ammonia. The results could also help to assess the accuracy of other existing data, ensuring that engineers have the best available information for working with the substance.

Large-scale simulations provide unique insights into thermodynamic properties

Vrabec is a longtime user of HLRS supercomputing resources for molecular dynamics and Monte Carlo simulations. His approach relies on concepts of thermodynamics that were first articulated by Ludwig Boltzmann in the 19th century but only became practical to apply in the 1950s with the arrival of the first computers. Since then, the field has advanced in parallel with the development of larger and faster supercomputers, to the point that Vrabec's simulations now track the individual motions and interactions of billions or even trillions of molecules simultaneously. Using software his lab developed to selectively capture data of interest, he can then study the molecules' thermodynamic properties.

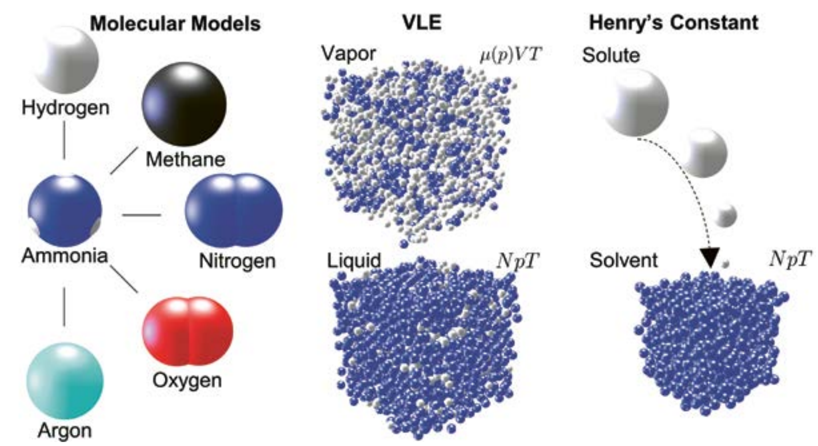

Simulations of thermodynamic properties of mixtures of ammonia and other molecules provided insights into their vapor–liquid equilibria and Henry's constants, important factors in determining how gases and liquids will mix in chemical processes.

Vrabec uses two simulation codes called ms2 and ls1, which he has developed and optimized over the course of a long and fruitful collaboration with HLRS staff members Martin Bernreuther and Christoph Niethammer. In 2019 the team even set a world record for the largest molecular system ever simulated using molecular dynamics methods. Using ls1, they efficiently scaled their code to a system of 21 trillion atoms in which every individual molecule and its interactions with other molecules could be tracked.

In the recent work on ammonia, Mace and Vrabec performed molecular dynamics and Monte Carlo simulations using ms2 to investigate five commonly used mixtures involving ammonia in chemical engineering processes: argon-ammonia, methane-ammonia, hydrogen-ammonia, nitrogen-ammonia, and oxygen-ammonia. For each mixture the simulations generated data describing the vapor-liquid equilibrium (VLE) – a measurement of the distribution of molecules in a system across the vapor or liquid phases – for a wide range of temperatures and pressures. In their paper Mace and Vrabec point out that VLE data is often used in developing equations of state for industrial fluids; that is, the data can be used to predict the state of matter under different physical conditions due to changes in temperature, pressure, volume, or composition. Such information is essential for determining optimal mixtures and working conditions in industrial applications.

Vrabec's molecular simulations are particularly valuable because they can be used to investigate a much wider range of scales than is possible using experimental approaches. "In our simulations, we provided measurements of thermodynamic properties even up to pressures of 50 megapascals. This is 500 times our ambient air pressure," Vrabec remarks. "Although data for ammonia mixtures have been gathered for more than a century, the data coverage is surprisingly narrow. The reason is that the effort to measure it experimentally is prohibitively huge. It would require expensive special equipment that would be dangerous to operate. In computer simulations, we can get results safely and relatively inexpensively." His methods also provide a comparable level of accuracy to that of experimental approaches in ranges where experimental data is available.

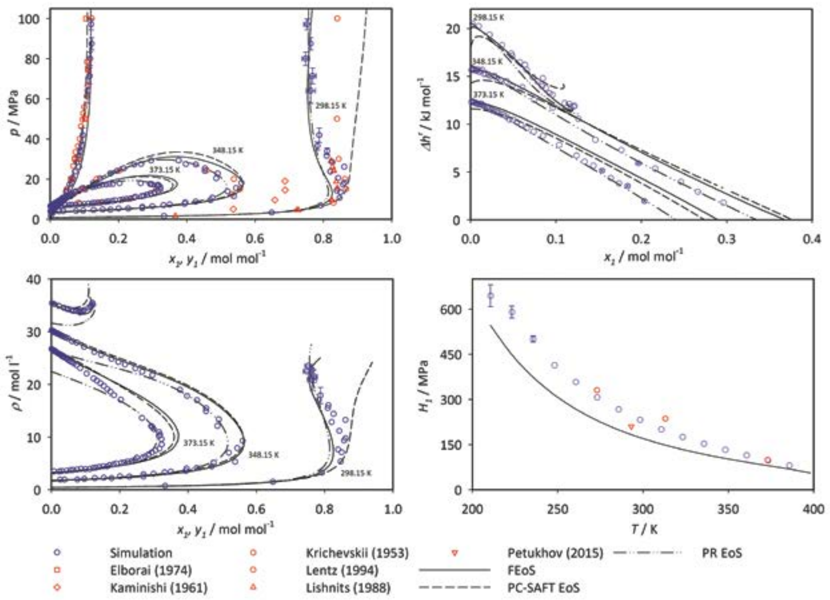

The graphs compare simulation and experimental data for mixtures of ammonia and methane at a wide range of concentrations, pressures, and temperatures. The simulation data (represented in blue circles) correspond well to other experimental data, and reveal outliers in experimental data (seen, for example, in the red diamonds for Kaminishi's 1961 results in the lower half of the top left figure) that are likely to be inaccurate.

Better data for ammonia research

When Mace and Vrabec analyzed their simulation data, they showed that although ammonia is a component in all five systems they studied, the resulting graphs of VLE values look dramatically different for different molecular mixtures. According to Vrabec, "The phase behavior of different mixtures is strongly determined by the interactions among the molecules in the system. You need to understand these properties if you are interested in working with ammonia mixtures."

The paper and its supplementary data offer more than 400 new data points for each mixture they studied. Using Hawk, they were able to produce the results of each mixture within just a few days of computing time. The results will be of particular value for extreme, difficult-to-study conditions for which little data is available, and could help engineers to identify sweet spots where conditions would be optimal for efficient ammonia processing.

The study included both new simulation data and previously published data from the scientific literature, enabling Mace and Vrabec to compare their results with other existing datasets of VLE values. In most situations, their results corresponded closely with those of previous studies. In some cases, however, they identified significant divergences between their results and other research groups' experimentally generated measurements and predictions. The authors attribute these discrepancies to limitations or inaccuracies in the corresponding experimental methods. They also suggest that specific experimental data sources should be used with caution in future research or chemical engineering applications.

Vrabec says that in recent work, he has focused primarily on simulating thermodynamic properties of molecular systems, generally at the sub-micrometer scale. Despite the many orders of magnitude that lie between this scale and the level of observable processes, accurate methods exist for translating these molecular-level insights into useful real-world predictions. As supercomputers grow larger, however, he anticipates that it might also become possible to simulate not just properties but also thermodynamic processes using boundary conditions that are close to real-world applications. Increased HPC performance could produce more accurate results about dynamic phenomena with a better signal-to-noise ratio.

In the meantime, though, his team's results demonstrate the value of molecular dynamics and Monte Carlo simulation using high-performance computing, and will provide new understanding of phase behavior that engineers can use to develop new ammonia-based technologies.

Vlasiator Provides Global Model of Earth's Magnetosphere

Using HLRS supercomputing resources, scientists led by University of Helsinki physicist Minna Palmroth are exploring phenomena in near-Earth space that could never be investigated before.

With the development of reusable rockets and smaller, cheaper satellites, the space industry is booming. As a consequence, society is increasingly dependent on satellites for critical functions such as communication and navigation. Ensuring that satellites remain safe in space is thus not only a question of protecting investments, but also of ensuring that human life on Earth runs smoothly.

Unfortunately, however, scientists still do not have a satisfactory understanding of important phenomena in near-Earth space that affect satellite function. Dynamics in the solar wind as it approaches Earth, for example, can wreak havoc on orbiting satellites, and despite many decades of research it has been challenging to develop detailed models of such conditions. Obstacles include the fact that space is obviously much larger than anything found on Earth, space weather is physically more complex than terrestrial weather, and there are limits on the number of satellites that can be shot into space to gather observational data. In the future, scientists would like to be able to predict weather in space as well as they do on Earth. For this to be possible, though, simulation of space physics using supercomputers is needed now.

For more than 10 years, Dr. Minna Palmroth, a professor in the Space Physics Research Group at the University of Helsinki, has been leading the development of a model called Vlasiator that holds the potential to improve scientists' understanding of the magnetosphere, the region surrounding Earth in which the planet's electromagnetic field interacts with the solar wind. Using several generations of supercomputers at HLRS – Hermit, Hazel Hen, and now Hawk – she and her team have been improving Vlasiator's capabilities so that it now provides a six-dimensional, global model of near-Earth space. In what she considers the culmination of her career so far, a 2023 publication in the journal *Nature Geoscience* demonstrates that Vlasiator can provide unique insights into phenomena that have been impossible to study using other simulation methods.

Vlasiator reduces approximation in particle modeling

Physicists have been using computers to simulate conditions in near-Earth space since the 1970s, but existing algorithms have had no choice but to approximate some of its most important features. In magnetohydrodynamics (MHD), an established approach for simulating

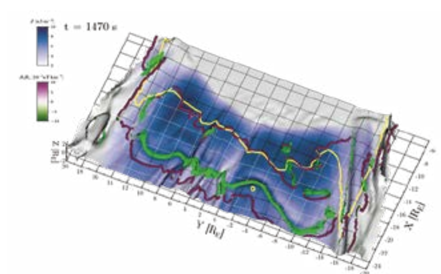

Snapshot of the plasma sheet surface in the Vlasiator simulation, showing large tail-wide plasma eruptions in current density and magnetic field topology.

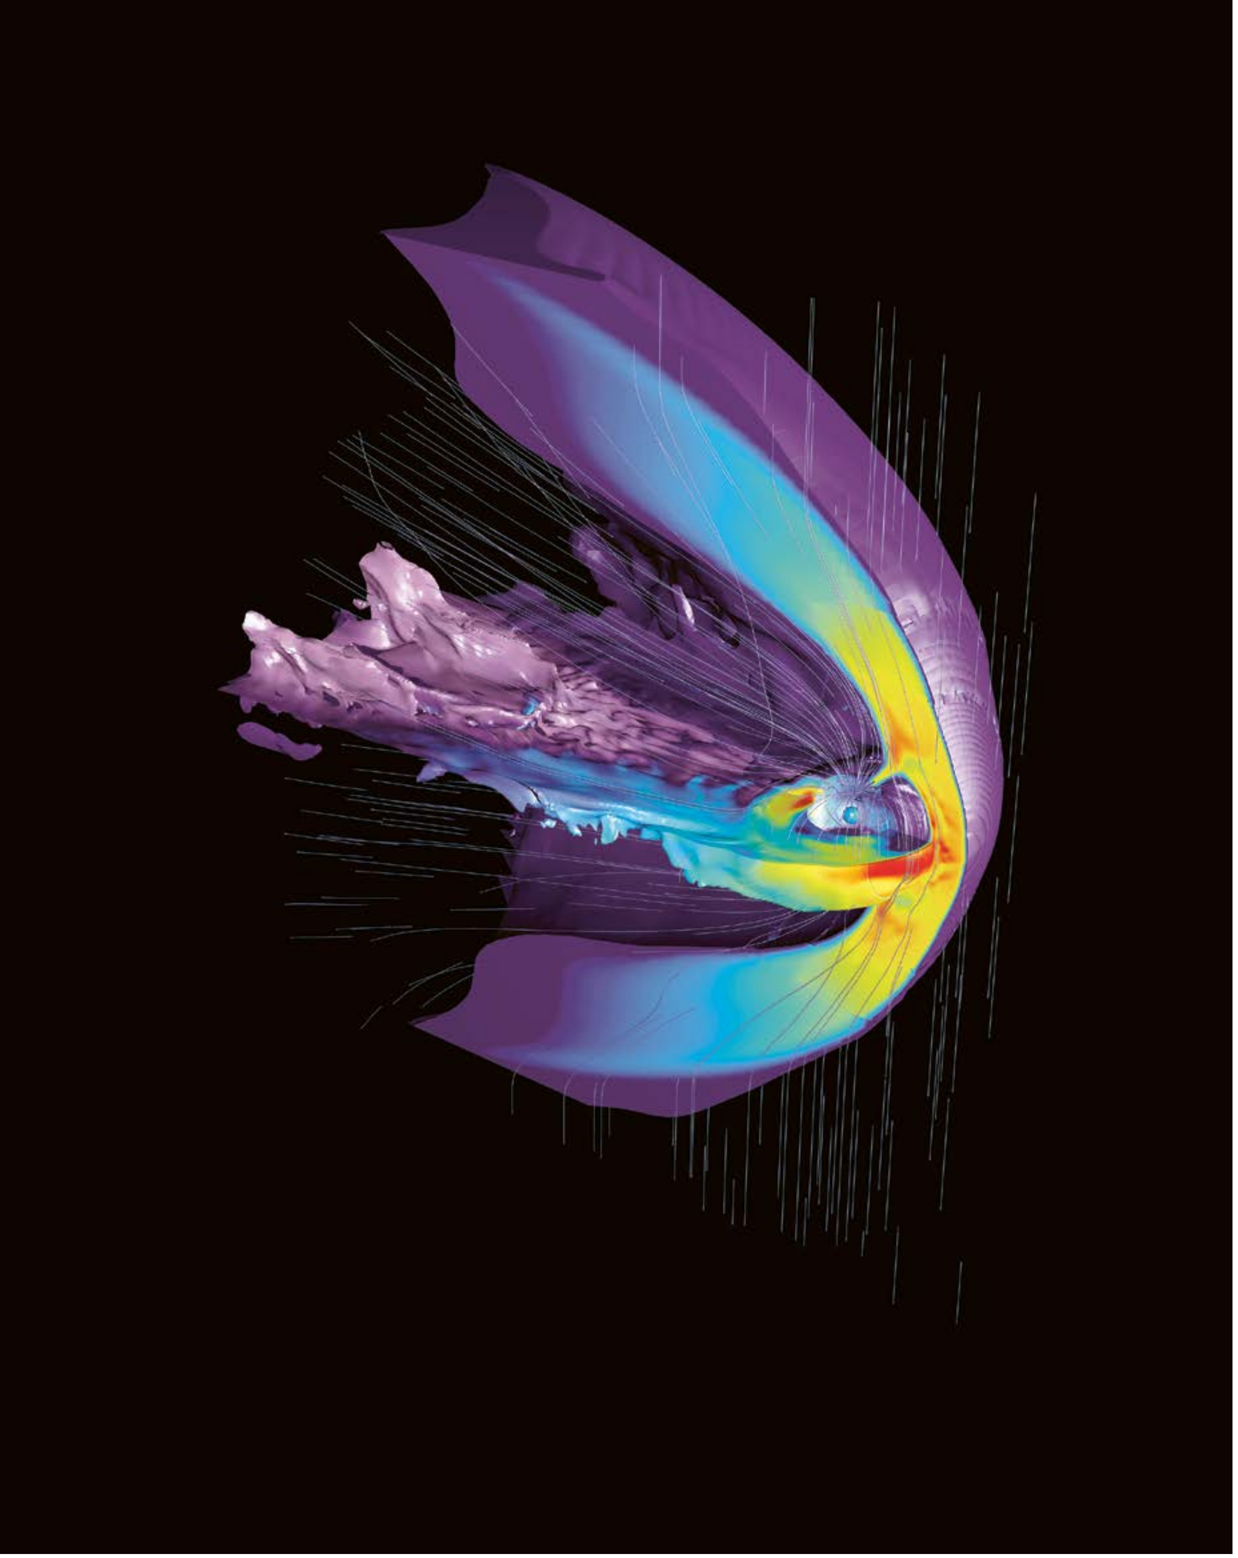

Researchers in the Vlasiator project used HLRS's Hawk supercomputer to complete the first 6-dimensional simulation of ion-scale dynamics within near-Earth space. In this visualization, the solar wind encounters the Earth's magnetic field, resulting in a bullet-shaped magnetosphere.

the magnetosphere, codes contain a built-in assumption that particle velocities in space are distributed in the same way they are on Earth; that is, following a normal distribution. In the terrestrial world, many particles in a given 3-dimensional cube cluster around an average velocity, with a relatively small number of outliers having either a very low or very high velocity. Because temperature is the result of moving particles, this distribution gives a good representation of the temperature of the air in the cube.

One long-recognized problem in space physics is that proton velocities in space do not follow a normal distribution, but are much more variable. Accounting for protons' actual behavior within MHD codes would require prohibitively large amounts of computing power, however. This means that modelers of space have had no choice but to treat particles in space in the same way as particles on Earth.

For Dr. Palmroth and her colleagues, improvements in algorithm development and the arrival of more powerful supercomputers like those at HLRS offer the opportunity to develop a model that doesn't require this compromise. "Running Vlasiator at HLRS, we use a different approach that enables us to model all proton-related phenomena in near-Earth space as they are, based on fundamental physical principles, without the need to approximate," she says.

To produce global simulations of the Earth's magnetosphere and ionosphere, Vlasiator considers six different dimensions – three spatial dimensions plus three dimensions that quantify particle distributions. The code does not model the movement of individual particles, but rather represents how proton distribution changes in space and time. (The mathematical approach they use is called the Vlasov equation.) "The shape of the particle distribution is crucial in many decades-old mysteries that haven't been solved," Palmroth explains. "By focusing on this issue, Vlasiator can reveal processes that haven't been visible before." This capability has made it possible for her team to open new windows for studying space physics.

Simulation explains Earth plasma eruptions

One of the most unpredictable phenomena that takes place in the magnetosphere is called a substorm. Here, interactions between the solar wind and the Earth's electromagnetic field lead to a buildup and sudden ejection of plasma called a plasmoid in the magnetotail, the part of the magnetosphere on Earth's nighttime side. From Earth's surface, these ejections are visible as especially spectacular auroras, but they can also damage satellites or interfere with their operation.

Although there have been attempts to explain substorms, their causes have been little understood. One potential explanation, called magnetic reconnection, suggests that changes in the electromagnetic field in the magnetotail sever part of the magnetosphere, which is released as a plasmoid. Another, called kinetic instability, posits that instabilities in the magnetotail produce waves that disrupt the current that sustains the magnetospheric tail, leading to plasmoid ejection.

Running Vlasiator on HLRS's Hawk supercomputer, Palmroth and her team simulated the magnetosphere at a scale capable of revealing the physics underlying both of these hypotheses. Surprisingly, their results indicated that both magnetic reconnection and kinetic instabilities occur, but not in the way that space physicists understood them in the past. Their paper in *Nature Geoscience* reports that plasmoid ejection occurs when smaller plasmoids are unified into a large plasmoid as a result of disruption in the current, which is caused by a kinetic instability.

According to Palmroth, it is the global nature of Vlasiator that made these novel findings possible. "Vlasiator now includes the upstream solar wind, the magnetosphere, and the ionosphere," she explains. "Plasma phenomena occur at a much smaller scale, but because global features affect local features, and vice versa, you need a global model to provide the necessary context. Without the global model, it is impossible to understand the evolution of the whole system at the same time."

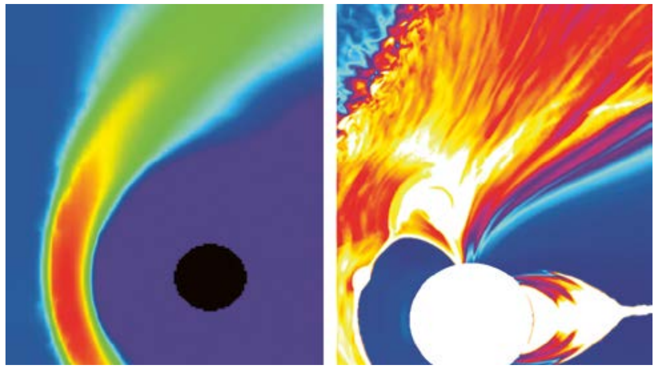

Comparison of modeling techniques using magnetohydrodynamics (left) and Vlasiator (right). Both methods model electrons as a fluid but Vlasiator uses a hybrid-Vlasov approach that simulates proton velocities as distribution functions. Vlasiator can capture ion kinetic effects at a smaller scale, and provides a global model that includes the solar wind, magnetosphere, and ionosphere.

Comparing Vlasiator's results to available observational data, Palmroth has found that the results provide a detailed and accurate picture. The code enables them to observe other features of the magnetosphere as well. Additional recent papers have reported on discoveries regarding the transmission of foreshock waves through Earth's bow shock, auroral proton precipitation, and the properties of Pc3 waves, for example.

Success through a long-time partnership with HLRS

Palmroth recalls that when she first proposed the idea of creating a global ion-kinetic model of the magnetosphere around 2004, many colleagues advised her that it would never be computationally practical. During Vlasiator's development, however, she has always developed code not for the architectures and capabilities of current HPC systems, but for future, more powerful HPC technologies.

"This philosophy only works if we are in close collaboration with people who look at the technological future," she explains. "This is why the High-Performance Computing Center Stuttgart is one of my favorite high-performance computing centers to work with. HLRS has an absolutely exquisite HPC knowledge and expertise, which is crucial when we are developing our algorithms. It enables a productive co-design process."

In addition to opening new fields for studying space science, Vlasiator is also remarkable for its efficiency when running on HLRS's Hawk supercomputer. Palmroth says that experts in computational modeling are often surprised at the code's performance, as it scales almost linearly when running on up to 200,000 compute cores. Palmroth credits this achievement to close collaboration with HLRS's user support staff, who have provided valuable advice on how to manage issues such as input-output, data storage, and other technical issues necessary for optimal usage of the center's supercomputers. "HPC experts at HLRS set high standards in application performance that we strive to realize in practice," she says.

As HPC approaches the exascale, Palmroth and her team see a number of opportunities for increasing Vlasiator's capabilities. In the future, they anticipate being able to improve the resolution of their model, simulate longer periods of physical time, and do multiple runs of the same simulation to produce even better models based on comparative statistical analysis. Vlasiator was originally designed for supercomputers based on CPU processors, and so the team is also currently completing a rewrite of Vlasiator's code so that its calculations of velocity space can run on GPU accelerators. Because the velocity calculations take up to 90 % of the computing time, she estimates that this effort could enable the code to run at least 20 times faster.

Such plans suggest that as high-performance computing continues to evolve, the Palmroth Lab will continue deliver new insights about the neighborhood surrounding our planetary home.

Selected Publications by Our Users in 2023

Agadjanov A, Djukanovic D, von Hippel G, et al. 2023. **Nucleon sigma terms with Nf = 2 + 1 flavors of O(a)-improved Wilson fermions.** Phys Rev Lett. 131: 261902.

Agadjanov A, Djukanovic D, von Hippel G, et al. 2023. **The pion-nucleon sigma term with Nf = 2 + 1 O(a)-improved Wilson fermions.** Proceedings of the 39th International Symposium on Lattice Field Theory, 9th–13th August 2022.

Appel D, Jöns S, Keim J, et al. 2023. **A narrow band-based dynamic load balancing scheme for the level-set ghost-fluid method.** In: Nägel WE, Kröner DH, Resch MM, eds. High Performance Computing in Science and Engineering '21. Springer.

Badalov SV, Bocchini A, Wilhelm R, et al. 2023. **Rutile, anatase, brookite and titania thin film from Hubbard corrected and hybrid DFT.** Mater Res Express. 10: 075501.

Barnowsky T, Krasheninnikov AV, Friedrich R. 2023. **A new group of 2D non-van der Waals materials with ultra low exfoliation energies.** Adv Electron Mater. 2201112.

Bashir U, Böttcher K, Klimm D, et al. 2023. **Solid solutions of lithium niobite and lithium tantalate: crystal growth and the ferroelectric transition.** Ferroelectrics. 613: 250-262.

Bauer HS, Schwitalla T, Branch O, Thundathil R. 2023. **WRF simulations to investigate processes across scales (WRFSCALE).** In: Nägel WE, Kröner DH, Resch MM, eds. High Performance Computing in Science and Engineering '21. Springer.

Beck A, Kurz M. 2023. **Toward discretization-consistent closure schemes for large eddy simulation using reinforcement learning.** Phys Fluids. 35: 125122.

Blind MP, Kahraman AB, Larsson J, Beck A. 2023. **Residual estimation for grid modification in wall-modeled large eddy simulation using unstructured high-order methods.** Comput Fluids. 254: 105796.

Blind MP, Kleinert J, Lutz T, Beck A. 2023. **A time-accurate inflow coupling for zonal LES.** CEAS Aeronaut J. ePub March 9.

Bocchini A, Xie Y, Schmidt WG, Gerstmann U. 2023. **Structural and electrochemical properties of F-doped RbTiOPO4 (RTP:F) predicted from first principles.** Crystals. 14(1): 5.

Borgelt J, Hösgen T, Meinke M, Schröder W. 2023. **Analysis of the hot gas ingress into the wheel space of an axial turbine stage.** In: Nägel WE, Kröner DH, Resch MM, eds. High Performance Computing in Science and Engineering '21. Springer.

Bromley SJ, McCann M, Loch SD, Ballance CP. 2023. **Electron-impact excitation of Pt I–III: the importance of metastable and collision processes in neutron star merger and laboratory plasmas.** Astrophys J Suppl Series. 268(1): 22.

Bui MN, Rost S, Auge M, et al. 2023. **Optical properties of $MoSe_2$ monolayer implanted with ultra-low-energy Cr ions.** ACS Appl Mater Interfaces. 15(29): 35321-35331.

Buijse R, Parnet M, Nägel A. 2023. **Large scale simulations of partial differential equation models**

on the Hawk supercomputer. In: Nägel WE, Kröner DH, Resch MM, eds. High Performance Computing in Science and Engineering ’21. Springer.

Caldas A, Feldmann H, Lucio-Eceiza E, Pinto JG. 2023. **Convection-parameterized and convection-permitting modelling of heavy precipitation in decadal simulations of the greater Alpine region with COSMO-CLM.** Weather Clim Dynam. 4: 543-565.

Carbonio EA, Sulzmann F, Klyzshin AY, et al. 2023. **Adjusting the chemical reactivity of oxygen for propylene epoxidation on silver by rational design: the use of an oxyanion and Cl.** ACS Catal. 13: 5906-5913.

Cè M, Chao EH, Gérardin A, et al. 2023. **Hadronic contributions to the anomalous magnetic moment of the muon from Lattice QCD.** In: Nägel WE, Kröner DH, Resch MM, eds. High Performance Computing in Science and Engineering ’21. Springer.

Chabanov M, Tootle SD, Most ER, Rezzolla L. 2023. **Crustal magnetic fields do not lead to large magnetic-field amplifications in binary neutron star mergers.** Astrophys J Lett. 945 (1): L14.

Chepkasov IV, Krashneninnikov AV. 2023. **Tetracene crystals as promising anode material for alkali metal ion batteries.** Carbon 213: 118190.

Denev JA, Zirwes T, Zhang F, Bockhorn H. 2023. **A low-pass filter for linear forcing in the open-source code OpenFOAM – Implementation and numerical performance.** In: Nägel WE, Kröner DH, Resch MM, eds. High Performance Computing in Science and Engineering ’21. Springer.

Depta PN, Dosta M, Heinrich S. 2023. **Data-driven multiscale modeling of self-assembly and hierarchical structural formation in biological macro-molecular systems.** In: Nägel WE, Kröner DH, Resch MM, eds. High Performance Computing in Science and Engineering ’21. Springer.

Engler C, Pillepich A, Joshi GD, et al. 2023. **Satellites of Milky Way and M31-like galaxies with TNG50: quenched fractions, gas content, and star formation histories.** Mon Not R Astron Soc. ePub May 8.

Feichi Z, Wachter S, Zirwes T, et al. 2023. **Effect of nozzle upscaling on coaxial, gas-assisted atomization.** Phys Fluids 35: 043302.

Fortuni F, Merlin E, Fontana A, et al. 2023. **FORECAST: a flexible software to forward model cosmological hydrodynamical simulations mimicking real observations.** Astron Astrophys. 677: A102.

Funke B, García-Cosmas M, Glatthor N, et al. 2023. **Michelson interferometer for passive atmospheric sounding Institute of Meterology and Climate Research/Instituto de Astrofísica de Andalucía version 8 retrieval of nitric oxide and lower-thermospheric temperature.** Atmos Meas Tech. 16(8): 2167-2196.

Gao M, Appel D, Beck A, Munz CD. 2023. **A high-order fluid–structure interaction framework with application to shock-wave/turbulent boundary-layer interaction over an elastic panel.** J Fluids Structures. 121: 103950.

Ganse U, Koskela T, Battarbee M, et al. 2023. **Enabling technology for global 3D + 3V hybrid-Vlasov simulations of near-Earth space.** Phys Plasmas. 30: 042902.

García-Comas M, Funke B, López-Puertas M, et al. 2023. **Version 8 IMK-IAA MIPAS temperatures from 12–14 μm spectra: middle and upper atmosphere modes.** Atmos Measurement Techniques. 16(21): 5357-5386.

Gondrum M, Satcunanathan S, Niemöller A, et al. 2023. **A direct-hybrid CFD/CAA method based on lattice Boltzmann and acoustic perturbation equations.** Acta Acust. 7(5).

Grandin M, Connor HK, Hoilijoki S, et al. 2024 (ePub Jul 10, 2023). **Hybrid-Vlasov simulation of soft X-ray emissions at the Earth’s dayside magnetospheric boundaries.** Earth Planetary Phys. 8: 70-88.

Grandin M, Luttikhuis T, Battarbee M, et al. 2023. **First 3D hybrid-Vlasov global simulation of auroral proton precipitation and comparison with satellite observations.** J Space Weather Space Clim. 13: 20.

Grauer M, Behar E. 2023. **Absorption of gamma-ray burst x-ray afterglows by the missing baryons: confronting observations with the cosmological simulations.** Astrophys J. 953: 158.

Guenther JN, Borsanyi S, Fodor Z, et al. 2023. **Resummed lattice QCD equation of state at finite baryon density: strangeness neutrality and beyond.** 39th International Symposium on Lattice Field Theory, LATTICE 2022 8th–13th August, 2022 Bonn, Germany.

Hespe D, Hübner L, Hübschle-Schneider L, et al. 2023. **Scalable discrete algorithms for big data applications.** In: Nägel WE, Kröner DH, Resch MM, eds. High Performance Computing in Science and Engineering ’21. Springer.

Homes S, Fingerhut R, Guevara-Carrion G, et al. 2023. **Scaling in the context of molecular dynamics simulations with ms2 and ls1 mardyn.** In: Nägel WE, Kröner DH, Resch MM, eds. High Performance Computing in Science and Engineering ’21. Springer.

Homes S, Heinen M, Vrabec J. 2023. **Influence of molecular anisotropy and quadrupolar moment on evaporation.** Phys Fluids. 35: 052111.

Hundhausen M, Feldmann H, Laube N, Pinto JG. 2023. **Future heat extremes and impacts in a convection-permitting climate ensemble over Germany.** Nat Hazards Earth Syst Sci. 23: 2873-2893.

Jones T. 2023. **Electro-catalysis for H_2O oxidation.** In: Nägel WE, Kröner DH, Resch MM, eds. High Performance Computing in Science and Engineering ’21. Springer.

Kellers B, Lautenschlaeger MP, Rigos N, et al. 2023. **Systematic workflow for efficient identification of local representative elementary volumes demonstrated with lithim-ion battery cathode microstructures.** Batteries 9(7): 390.

Kempf D, Gao M, Beck A, et al. 2023. **Development of turbulent inflow methods for the high order HPC framework FLEXI.** In: Nägel WE, Kröner DH, Resch MM, eds. High Performance Computing in Science and Engineering ’21. Springer.

Kiefer M, Funke B, García-Comas M. 2023. **Global long-term MIPAS data processing.** In: Nägel WE, Kröner DH, Resch MM, eds. High Performance Computing in Science and Engineering ’21. Springer.

Kleinert J, Ehrle M, Waldmann A, Lutz T. 2023. **Wake tail plane interactions for a tandem wing configuration in high-speed stall conditions.** CEAS Aeronautical J. ePub Jun 14.

Klimach H, Masilamani K, Roller S. 2023. **Exploiting hybrid parallelism in the LBM implementation of Musubi on Hawk.** In: Resch MM, Gebert J, Kobayashi H, Bez W, eds. Sustained Simulation Performance 2021. Springer.

Knodel MM, Di Stefano S, Nägel A, Grillo A. 2022. **An efficient algorithm for biomechanical problems based on a fully implicit nested Newton solver.** Theor Appl. 49(2): 183–221

Knodel MM, Nägel A, Herrmann E, Wittum G. 2023. **PDE models of virus replication coupling 2D manifold and 3D volume effects evaluated at realistic reconstructed cell geometries.** In: Franck E, Fuhrmann J, Michel-Dansac V, Navoret L, eds. Finite Volumes for Complex Applications X—Volume 1, Elliptic and Parabolic Problems. FVCA 2023. Springer Proceedings in Mathematics & Statistics, vol. 432.

Köster J, Kretschmer S, Storm A, et al. 2023. **Phase transformations in single-layer $MoTe_2$ stimulated by electron irradiation and annealing.** Nanotechnology. ePub Dec 14.

Kretschmer S, Joseph T, Ghaderzadeh S, et al. 2023. **Production of defects in two-dimensional materials under ion and electron irradiation: insights from advanced first-principles calculations annual report on 'Bundesprojekt' ACID 44172 (Def-2-Dim).** In: Nägel WE, Kröner DH, Resch MM, eds. High Performance Computing in Science and Engineering '21. Springer.

Kromer J, Potyka J, Shulte K, Bothe D. 2023. **Efficient sequential PLIC interface positioning for enhanced performance of the three-phase VoF method.** Comput Fluids. ePub Sep 1.

Lewis JSW, Pillepich A, Nelson D, et al. 2023. **The signature of galaxy formation models in the power spectrum of the hydrogen 21cm line during reionization.** Mon Not R Astron Soc. ePub Dec 19.

Li J, Ghorbani-Asl M, Lasek K, et al. 2023. **A van der Waals heterostructure with an electronically textured moire pattern: $PtSe_2/PtTe_2$.** ACS Nano 17(6): 5913-5920.

Li Y, Ries F, Sun Y, et al. 2023. **Direct numerical simulation of atomization characteristics of ECN spray-G injector: in-nozzle fluid flow and breakup processes.** Flow Turbul Combust. ePub Dec 15.

Liu Q, Lin YC, Kretschmer S, et al. 2023. **Molybdenum chloride nanostructures with giant lattice distortions intercalated into bilayer graphene.** ACS Nano, ePub Nov 26.

Liu Y, Wang W, Yang G, et al. 2023. **The interfacial modes and modal causality in a dispersed bubbly turbulent flow.** Phys Fluids. 35: 083309.

Lode AUJ, Alon OE, Arnold J, et al. 2023. **Quantum simulators, phase transitions, resonant tunneling, and variances: a many-body perspective.** In: Nägel WE, Kröner DH, Resch MM, eds. High Performance Computing in Science and Engineering '21. Springer.

Long F, Ghorbani-Asl M, Mosina K, et al. 2023. **Ferromagnetic interlayer coupling in CrSBr crystals irradiated by ions.** Nano Lett. 23(18): 8468-8473.

Mace EJ, Vrabec J. 2023. **High-pressure fluid-phase equilibria and Henry's constants of supercritical gases in ammonia.** J Chem Eng Data. ePub Oct 2.

Maier B, Schneider D, Schulte M, Uekermann B. 2023. **Bridging scales with volume coupling – Scalable simulations of muscle contraction and electromyography.** In: Nägel WE, Kröner DH, Resch MM, eds. High Performance Computing in Science and Engineering '21. Springer.

Mallor F, Semprini-Cesari G, Mukha T, et al. 2023. **Bayesian optimization of wall-normal blowing and suction-based flow control of a NACA 4412 wing profile.** Flow Turbul Combust. ePub Sep 5.

Markin I, Neuweiler A, Abac A, et al. 2023. **General-relativistic hydrodynamics simulation of a neutron star–sub-solar-mass black hole merger.** Phys Rev D. 108: 064025.

Martins FF, Krewald V. 2023. **Cooperative dinitrogen activation: identifying the push-pull effect of transition metals and Lewis acids in molecular orbital diagrams.** Eur J Inorg Chem. 26(35): e202300268.

Marx J, Kohns M, Langenbach K. 2023. **Co-oriented fluid functional equation for electrostatic interactions (COFFEE) for mixtures: molecular orientations.** J Chem Eng Data. ePub Aug 25.

Marx J, Kohns M, Langenbach K. 2023. **Vapor-liquid equilibria of binary mixtures containing Stockmayer-type model fluids from Monte-Carlo simulations.** Fluid Phase Equilibr. 568: 113742.

Mäteling E, Albers M, Schröder W. 2023. **How spanwise travelling transversal surface waves change the near-wall flow.** J Fluid Mechan. 957: ePub Feb 23.

Meier L, Schmidt WG. 2023. **Adsorption of cyclic (alkyl) (amino) carbenes on monohydride Si(001) surfaces: interface bonding and electronic properties.** J Phys Chem C. 127(4): 1973-1980.

Meller D, Engelmann L, Baik S-J, Kempf AM. 2023. **A four mixture fraction FPV-LES for the co-firing of coal and ammonia.** In Marchioli C, Salvetti MV, Garcia-Villalba M, Schlatter P, eds. Direct and Large Eddy Simulation XIII. DLES 2023. ERCOFTAC Series, vol. 31. Springer.

Meller D, Engelmann L, Stein OT, Kempf AM. 2023. **Exhaust gas recirculation (EGR) analysis of a swirl-stabilized pulverized coal flame with focus on NOx release using FPV-LES.** Fuel 343: 127939.

Meringolo C, Cruz-Osorio A, Rezzolla L, Vervidio S. 2023. **Microphysical plasma relations from special-relativistic turbulence.** Astrophys J. 94(2): 122.

Mpisketzis V, Duqué R, Nathanail, et al. 2023. **Impact of anisotropic ejecta on jet dynamics and afterglow emission in binary neutron-star mergers.** Mon Not R Astron Soc. ePub Dec 6.

Mulone A, Awad S, Chiarugi D, Aldinucci M. 2023. **Porting the variant calling pipeline for NGS data in a cloud-HPC environment. 2**023 IEEE 47th Annual Computers, Software, and Applications Conference COMPSAC.

Nanni L, Thomas D, Trayford J, et al. 2023. **iMaNGA: mock MaNGA galaxies based on Illustris TNG and MaStar SSPs – II. The catalogue.** Mon Not R Astron Soc. 522(4): 5479-5499.

Nelson D, Byrohl C, Ogorzalek A, et al. 2023. **Resonant scattering of the OVII X-ray emission line in the circumgalactic medium of TNG50 galaxies.** Mon Not R Astron Soc. 522(3): 3665-3678.

Neumann P, Wittmer N, Jafari V, et al. 2023. **Dynamic molecular dynamics ensembles for multiscale simulation coupling.** In: Nägel WE, Kröner DH, Resch MM, eds. High Performance Computing in Science and Engineering '21. Springer.

Nitzke I, Vrabec J. 2023. **Numerical discrimimation of thermodynamic Monte Carlo simulations in all eight statistical ensembles.** J Chem Theory Comput. 19(12): 3460-3468.

Oelschläger F, Klein D, Müller S, Roth J. 2023. **Molecular dynamics simulation of selective laser melting.** In: Nägel WE, Kröner DH, Resch MM, eds. High Performance Computing in Science and Engineering '21. Springer.

Ohno D, Greiner M, Müller J, Rist U. 2023. **Numerical investigations of laminar separation bubbles under the influence of periodic gusts.** AIAA J. ePub Mar 16.

Ohno D, Selent B, Kloker MJ, Rist U. 2023. **Direct numerical simulation of bypass transition under free-stream turbulence for compressible flows.** In: Nägel WE, Kröner DH, Resch MM, eds. High Performance Computing in Science and Engineering '21. Springer.

Owen E, Han Q, Wu K. 2023. **Effects of large-scale magnetic fields on the observed composition of ultrahigh-energy cosmic rays.** Phys Rev D. 103027.

Palmroth M, Pulkkinen TI, Ganse U, et al. 2023. **Magnetotail plasma eruptions driven by magnetic reconnection and kinetic instabilities.** Nat Geosci. 16: 570-576.

Pathirage V, Lasek K, Krasheninnikov AV, et al. 2023. **Mirror twin boundaries in WSe_2 induced by vanadium doping.** Mater Today Nano. 22: 100314.

Paul S, Hanlon AD, Hörz B, et al. 2023. **The long-distance behavior of the vector-vector correlator from $\pi\pi$ scattering.** 39th International Symposium on Lattice Field Theory, LATTICE 2022 8th–13th August, 2022 Bonn, Germany.

Pfeiffer M, Beyer J, Vaubaillon J, et al. 2023. **Numerical simulation of an iron meteoroid entering into earth's**

atmosphere using DSMC and a radiation solver with comparison to ground testing data. Icarus. 115768.

Pieck F, Luy JN, Kreuter F, et al. 2023. **Reactivity of organic molecules on semiconductor surfaces revealed by density functional theory.** In: Nägel WE, Kröner DH, Resch MM, eds. High Performance Computing in Science and Engineering '21. Springer.

Pinzon J, Siebenborn M, Vogel A. 2023. **Scalable multigrid algorithm for fluid dynamic shape optimization.** In: Nägel WE, Kröner DH, Resch MM, eds. High Performance Computing in Science and Engineering '21. Springer.

Pionteck MN, Bernhardt F, Bilk J. 2023. **Dynamical properties of the Si(553)-Au nanowire system.** In: Nägel WE, Kröner DH, Resch MM, eds. High Performance Computing in Science and Engineering '21. Springer.

Pohl S, Fingerhut R, Thol M, et al. 2023. **Equation of state for the Mie (λr,6) fluid with a repulsive exponent from 11 to 13.** J Chem Phys. 158: 084506.

Potyka J, Schulte K. 2024 (ePub Oct 26, 2023). **A volume of fluid method for three dimensional direct numerical simulations of immiscible droplet collisions.** Int J Multiphase Flow. 170: 104654.

Potyka J, Schulte K, Planchette C. 2023. **Liquid distribution after head-on separation of two colliding immiscible liquid droplets.** Phys Fluids. 35: 102125.

Prather BS, Dexter J, Moscibrodzka M, et al. 2023. **Comparison of polarized radiative transfer codes used by the EHT collaboration.** Astrophys J. 950: 35.

Ramaswamy DP, Sebastian R, Schreyer AM. 2023. **Effects of jet-injection-pipe length on the flow-control effectiveness of spanwise-inclined jets in supersonic crossflow.** Phys Rev Fluids. 8: 113902.

Ramesh R, Nelson D, Heesen V, Brüggen M. 2023. **Azimuthal anisoptropy of magnetic fields in the circumgalactic medium driven by galactic feedback processes.** Mon Not R Astron Soc. 526(4): 5483-5493.

Rist U, Weinschenk M, Wenzel C. 2023. **Investigation of Lagrangian areas of minimal stretching (LAMS) in a turbulent boundary layer.** J Fluid Mechan. 970, A31.

Rose H, Kunert N, Dietrich T, et al. 2023. **Revealing the strength of three-nucleon interactions with the proposed Einstein Telescope.** Phys Rev C. 108: 025811.

Ruiz Alvarado IA, Zare Pour MA, Hannapel T, Schmidt WG. 2023. **Structural fingerprints in the reflectance anisoptropy of AlInP(001).** Phys Rev B. 108: 045410.

Ryu T, Krolik J, Piran T, et al. 2023. **Shocks power tidal disruption events.** Astrophys J. 957: 12.

Sabater E, Solà M, Salvador P, Andrala DM. 2023. **Cage-size effects on the encapsulation of P_2 by fullerenes.** Comput Chem. 44(3): 268-277.

Saeed MS, Gugliuzza JMJ, Leibl M, et al. 2023. **Parameter study and optimization of forming simlations for tape-based fiber layups.** In: Kiefl N, Wulle F, Ackermann C, Holder D, eds. Advances in Automotive Production Technology – Towards Software-Defined Manufacturing and Resilient Supply Chains. SCAP 2022. ARENA2036. Springer.

Saric D, Guevara-Carrion G, Gaponenko Y, et al. 2023. **Diffusion of hydrocarbons diluted in supercritical carbon dioxide.** Sci Rep. 13: 16107.

Schäfer T, Joerger AC, Spencer J, et al. 2023. **Ligand-induced protein stabilization and enhanced molecular dynamics sampling techniques.** In: Nägel WE, Kröner DH, Resch MM, eds. High Performance Computing in Science and Engineering '21. Springer.

Schaefer D, Kohns M. 2023. **Molecular dynamics study of ion clustering in concentrated electrolyte solutions for the estimation of salt solubilities.** Fluid Phase Equilibr. 571: 113802.

Schaefer D, Kohns M, Hasse H. 2023. **Molecular modeling and simulation of aqueous solutions of alkali nitrates.** J Chem Phys. 158: 134508.

Schlottke A, Ibach M, Steigerwald J, Weigand B. 2023. **Direct numerical simulation of a disintegrating liquid rivulet at a trailing edge.** In: Nägel WE, Kröner DH, Resch MM, eds. High Performance Computing in Science and Engineering '21. Springer.

Schneider M, Ertl B, Diekmann C, Khosrawi F. 2023. **The HPC project GLOMIR+ (GLObal MUSICA IASI Retrievals – plus).** In: Nägel WE, Kröner DH, Resch MM, eds. High Performance Computing in Science and Engineering '21. Springer.

Schollenberger M, Beckers MF, Lutz T. 2023. **Validation of ACD and ACL propeller simulation using blade element method based on airfoil characteristics.** In: Nägel WE, Kröner DH, Resch MM, eds. High Performance Computing in Science and Engineering '21. Springer.

Schoof E, Mittnacht T, Seiz M, et al. 2023. **High-performance multiphase-field simulations of solid-state phase transformations using Pace3D.** In: Nägel WE, Kröner DH, Resch MM, eds. High Performance Computing in Science and Engineering '21. Springer.

Schuchart J, Bosilca G. 2023. **MPI Continuations and how to invoke them.** In: Resch MM, Gebert J, Kobayashi H, Bez W, eds. Sustained Simulation Performance 2021. Springer.

Schwagerus A, Habisreuther P, Zarzalis N. 2023. **Numerical calculation of the lean-blow-out in a multi-jet burner.** In: Nägel WE, Kröner DH, Resch MM, eds. High Performance Computing in Science and Engineering '21. Springer.

Sebastian R, Ramaswamy DP, Schreyer AM. 2023. **Spanwise-inclined jets in supersonic crossflow: effects of injection pressure and separation-control effectiveness.** AIAA J. ePub Jun 5.

Sebastian R, Schreyer AM. 2023. **Forced injection of a spanwise-inclined jet in supersonic crossflow.** AIAA J. ePub Sep 25.

Seel F, Lutz T, Krämer E. 2023. **Numerical study of the unsteady blade root aerodynamics of a 2 MW wind turbine equipped with vortex generators.** Wind Energ Sci. 8: 1369-1385.

Seibold F, Weigand B. 2023. **Numerical investigation of the flow and heat transfer in convergent swirl chambers.** In: Nägel WE, Kröner DH, Resch MM, eds. High Performance Computing in Science and Engineering '21. Springer.

Seiz M, Hierl H, Nestler B. 2023. **An improved grand-potential phase-field model of solid-state sintering for many particles.** Modelling Simul Mater Sci Eng. 31: 055006.

Seiz M, Hierl H, Nestler B. 2023. **Unravelling densification during sintering by multiscale modelling of grain motion.** J Mater Sci. 58: 14051-14071.

Seiz M, Kellner M, Nestler B. 2023. **Simulation of dendritic–eutectic growth with the phase-field method.** Acta Mater. 254: 118965.

Settanni G. 2023. **Computational approaches to lipid-based nucleic acid delivery systems.** Eur Phys J E. 46: 127.

Sheina V, Lang G, Stolyarov V, et al. 2023. **Hydrogenic spin-valley states of the bromine donor in 2H-$MoTe_2$.** Communications Phys. 6: 135.

Singh E, Pionteck MN, Reitzig S, et al. 2023. **Vibrational properties of $LiNbO_3$ and $LiTaO_3$ under uniaxial stress.** Phys Rev Materials 7: 024420.

Skopinski L, Kretschmer S, Ernst P, et al. 2023. **Velocity distributions of particles sputtered from supported two-dimensional MoS_2 during highly charged ion irradiation.** Phys Rev B. 107: 075418.

Sofieva VF, Szelag M, Tamminen, et al. 2023. **Updated merged SAGE-CCI-OMPS+ dataset for the evaluation of ozone trends in the stratosphere.** Atmos Meas Tech. 16(7): 1881–1899.

Sotillo-Ramos D, Bergemann M, Friske JKS, Pillepich A. 2023. **On the likelihoods of finding very meal-poor (and old) stars in the Milky Way's disc, bulge, and halo.** Mon Not R Astron Soc: Lett. ePub Jul 27.

Storm A, Köster J, Ghorbani-Asl M, et al. 2023. **Electron-beam- and thermal-annealing-induced structural transformations in few-layer $MnPS_3$.** ACS Nano. ePub Feb 20.

Storm A, Köster J, Ghorbani-Asl M, et al. 2023. **Structural transformations in few-layer $MnPS_3$ stimulated by thermal annealing and electron irradiation.** J Phys Chem C. 127(51): 24713-24723

Subakti S, Daquqshirazi M, Wolf D, et al. 2023. **Electron holographic mapping of structural reconstruction at mono- and bilayer steps of h–BN.** Phys Rev Res. 5: 033137.

Suni J, Palmroth M, Turc L, et al. 2023. **Local bow shock environment during magnetosheath jet formation: results from a hybrid Vlasov simulation.** Annales Geophys. 41(2): 551-568.

Ujevic M, Gieg H, Schianchi F, et al. 2023. **Reverse phase transitions in binary neutron-star systems with exotic-matter cores.** Phys Rev D. 107: 024025.

Verestek W, Kaiser J, Bonten C, Schmauder S. 2023. **Molecular dynamics investigations on the influence of solutes on the tensile behavior of Polyamide6.** In: Nägel WE, Kröner DH, Resch MM, eds. High Performance Computing in Science and Engineering '21. Springer.

Walzer U, Hendel R. 2023. **Natural climate change and glaciations.** Earth Sci Rev. ePub Apr 29.

Wack J, Zorn M, Riedelbauch S. 2023. **Numerical simulation of vortex induced pressure fluctuations in the runner of a Francis turbine at deep part load conditions.** In: Nägel WE, Kröner DH, Resch MM, eds. High Performance Computing in Science and Engineering '21. Springer.

Wellinger P, Weigand B. 2023. **On the validity of the linear Boussinesq hypothesis for selected internal cooling features of gas turbine blades.** In: Nägel WE, Kröner DH, Resch MM, eds. High Performance Computing in Science and Engineering '21. Springer.

Wenz F, Maas O, Arnold M, et al. 2023. **Assessment of low-frequency aeroacoustic emissions of a wind turbine under rapidly changing wind conditions based on an aero-servo-elastic CFD simulation.** Wind Energy. ePub May 1.

Wenzel P, Radgen P. 2023. **Extending effectiveness to efficiency: comparing energy and environmental assessment methods for a wet cooling tower.** J Indust Ecol. 27(3): 693-706.

Wollny P, Engelmann L, Loewenich M, et al. 2023. **The role of phase transition by inception and surface reactions for the synthesis of silicon nanoparticles in a hot-wall reactor – simulation and experiment.** Chem Eng J. ePub Mar 9.

Ye Q, Dreher M, Shen B. 2023. **Numerical simulation of flake orientation during droplet impact on substrates in spray painting processes.** In: Nägel WE, Kröner DH, Resch MM, eds. High Performance Computing in Science and Engineering '21. Springer.

Zvejnieks G, Mastrikov Y, Gryaznov D. 2023. **Jahn-Teller distortion in Sr_2FeO_4: group-theoretical analysis and hybrid DFT calculations.** Sci Rep. 13: 16446.

HLRS Books

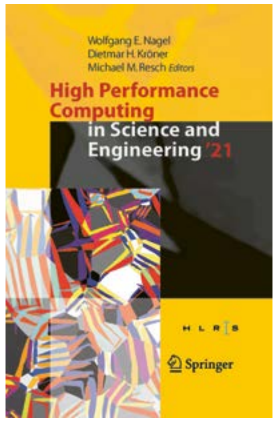

High Performance Computing in Science and Engineering '21

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

This book presents the state-of-the-art in supercomputer simulation. It includes the latest findings from leading researchers using systems from the High-Performance Computing Center Stuttgart (HLRS) in 2021. 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 of one of Europe's leading systems, this volume covers a wide variety of applications that deliver a high level of sustained performance. The book covers the main methods in high-performance computing. Its outstanding results in achieving the best performance for production codes are of particular interest for both scientists and engineers. It comes with a wealth of color illustrations and tables of results.

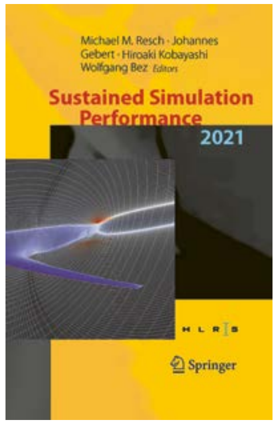

Sustained Simulation Performance 2021

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 context with 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. All contributions are based on selected papers presented in 2021 at the 31st Workshop on Sustained Simulation Performance, WSSP31, held at HLRS in Stuttgart, Germany, and WSSP32, held at Tohoku University in Sendai, Japan.

HLRS by Numbers

153 Staff

12 Visiting Researchers

113 Scientists

32 Nonscientists

7 Student Assistants

1 Research Assistant

System Usage

4.928 billion

Total Core Hours Produced

68

Talks by HLRS Staff

Research

4.662 billion

Core Hours

131 User Projects

132 User Publications

Industry

325 million

Core Hours

74 Industrial Customers

Staff Publications

35 Papers in Journals, Books, and Conference Proceedings

2 Books

Education and Training

41 Continuing Education Courses	12 Scientific Workshops	16 University Lectures
1,154 Participants	372 Participants	1,400 Participants
161 Course Days	19 Days	36 SWS

784

Vistors at HLRS

Third-Party Funds

7,345,882 €

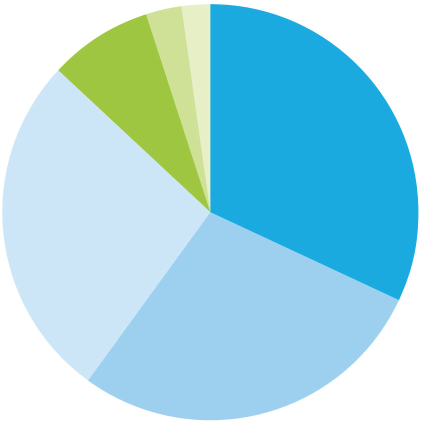

- EU 32 %
- State Gov./Land 28 %
- Federal Gov./Bund 27 %
- DFG 8 %
- Foundations 3 %
- Industry 2 %

About Us

Inside Our Computing Room

Hewlett Packard Enterprise Apollo (Hawk)

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

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

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

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

Baden-Württemberg
MINISTRY OF SCIENCE, RESEARCH AND ARTS

Federal Ministry of Education and Research

Cray CS-Storm

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

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

AMD GPU System

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

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

NEC Cluster (Vulcan)

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

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

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

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

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

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

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

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

Interconnects
Infiniband EDR/FDR/HDR/QDR

User Profile

In 2023 the Gauss Centre for Supercomputing approved 10 new large-scale projects (each project requiring more than 35 million core hours) for HLRS's flagship supercomputer, Hawk, for a total of 3.64 billion core-hours. In total, 131 scientific projects, including test projects, were active on Hawk in 2023, for a total of 4.454 billion core hours.

System Usage by Scientific Discipline

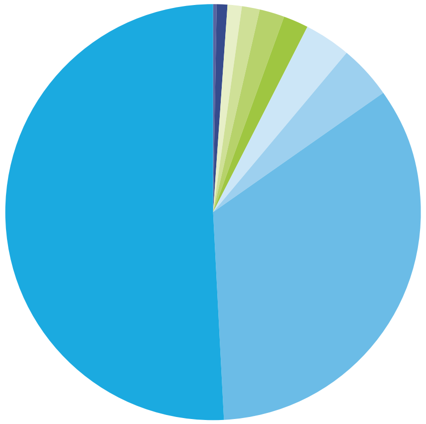

- CFD **50,85 %**
- Physics **34,10 %**
- Bioinformatics **4,06 %**
- Other **3,52 %**
- Chemistry **2,03 %**
- Festkörperphysik **1,94 %**
- Reactive Flows **1,30 %**
- Transportation and Climate **1,26 %**
- Materials Simulation & Materials Research **0,76 %**
- Electrical Engineering **0,16 %**
- Computer Science **0,02 %**

System Usage by State

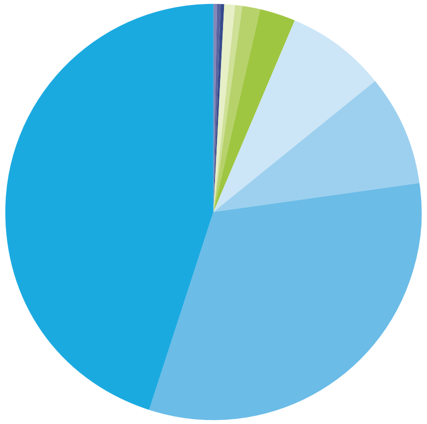

- North Rhine-Westphalia **45,19 %**
- Baden-Württemberg **32,30 %**
- Hesse **8,52 %**
- Brandenburg **7,77 %**
- Rheinland-Palatinate **2,64 %**
- Thuringia **1,42 %**
- Berlin **0,81 %**
- Bavaria **0,59 %**
- Federal Research Center **0,36 %**
- Saxony **0,30 %**
- Saxony-Anhalt **0,09 %**
- Hamburg **0,01 %**

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 2023.

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.

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.

Cape Reviso
July 2020 – December 2023 BMVI
By combining machine learning, sensor technology, network analysis and virtual reality in digital twins, HLRS is developing planning and decision support tools for conflict analysis and reduction between cyclists and pedestrians.

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.

■ **New in 2023** ■ **Grant awarded, starts in 2024**

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

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

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

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 is investigating a method for increasing energy efficiency in data centers by dynamically controlling cooling circuit temperatures, and developing guidelines for implementing the resulting concepts.

EE-HPC
September 2022 – August 2025 BMBF
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.

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

EuroCC 2 ■
January 2023 – December 2025 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 aims to shape the future of HPC in Europe by defining a long-term hyperconnectivity specification and implementation roadmap to meet Europe's future ultra-high-speed network requirements.

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

EXCELLERAT P2 ■
January 2023 – December 2026 JU
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.

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

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

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 is the coordinating center for this project to integrate artificial intelligence (AI) topics into curricula at the University of Stuttgart, and to implement AI technologies to improve instruction.

InHPC-DE
November 2017 – June 2024 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 – March 2025 JU, EU
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 is developing tools for meeting and collaborating from remote locations in three-dimensional virtual reality environments.

MERIDIONAL
October 2022 – September 2026 EU
This project is developing tools 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 is developing 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.

S+T+ARTS AIR ■
April 2023 – September 2024 EC, MWK
S+T+ARTS AIR is making supercomputing technologies and expertise available to enable innovative collaborations involving the arts, science, and technology.

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

SEQUOIA End-to-End ■
January 2023 – March 2024 WAT
SEQUOIA End-to-End aims 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.

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

SimTech
July 2023 – June 2025 DFG
This interdisciplinary Excellence Cluster at the University of Stuttgart is developing simulation technologies for integrative systems science. HLRS focuses on uncertainty quantification for heterogeneous hardware architectures in dispersion problems.

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

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

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 is developing perspectives for assessing the trustworthiness of computational science and limiting the spread of misinformation.

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

HPC Training Courses in 2023

HLRS offered 41 courses in 2023, providing continuing professional education on a wide range of topics relevant for high-performance computing. The courses took place over 161 course-days, online and in Stuttgart and in cooperation with other institutes in Germany and internationally. A total of 1,154 trainees participated in these activities.

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

Date	Location	Topic	Host
Jan 23–Feb 24	online/blended	HPC-Cluster – Auslegung, Kosten & Nachhaltigkeit	HLRS (SCA)
Feb 13–17	online	Fortran for Scientific Computing	HLRS
Feb 20–24	online	Introduction to Computational Fluid Dynamics	HLRS/DLR
Mar 7–10	online	Modern C++ Software Design (Intermediate)	HLRS
Mar 13–17	Stuttgart	CFD with OpenFOAM	HLRS
Mar 20–24	Stuttgart	Iterative Solvers and Parallelization	HLRS
Mar 27–31	Dresden	Parallel Programming with MPI & OpenMP and Tools	ZIH/HLRS
Apr 11–14	Mainz	Parallelization with MPI and OpenMP	ZDV/HLRS
Apr 17–18	online	AI for Science Bootcamp *	NVIDIA/HLRS/ JSC/LRZ/VSC
Apr 17–May 15	online/blended	Datenmanagement	HLRS (SCA)
Apr 24–May 5	online/blended	Paralleles Programmieren mit OpenMP	HLRS (SCA)
Apr 26–27	online	NVIDIA/HLRS SciML GPU Bootcamp	HLRS/NVIDIA
May 9–12	Stuttgart	Modern C++ Software Design (Advanced)	HLRS
May 15–16	online	N-Ways to GPU Programming Bootcamp*	NVIDIA/HLRS/ LRZ/JSC/VSC
May 22–26	Stuttgart/online	Optimization of Node-level Performance and Scaling on Hawk	HLRS
Jun 5–7	online	From Machine Learning to Deep Learning: a concise introduction	HLRS
Jun 5–Jul 17	online/blended	HPC-Cluster – Aufbau & Betrieb	HLRS (SCA)
Jun 12–13	online	Efficient Parallel Programming with GASPI	HLRS/Fraun-hofer-ITWM
Jun 12–14	online	Data analytics for engineering data using machine learning	Fraunhofer-SCAI/HLRS
Jun 15–16	online	Introduction to NEC SX-Aurora TSUBASA vector platform	HLRS/NEC
Jun 19–23	Stuttgart	Fortran for Scientific Computing	HLRS
Jun 27–30	online	Node-Level Performance Engineering	HLRS/NHR@ FAU
Jul 4–7	Stuttgart	Modern C++ Software Design (Intermediate)	HLRS
Jul 11–13	Stuttgart	Deep Learning and GPU programming using OpenACC NEW	HLRS/LRZ
Jul 26–28	Stuttgart	Summer School: Trust and Machine Learning NEW	HLRS
Aug 16–24	online	Six-day course in parallel programming with MPI/OpenMP	ETH/HLRS
Sep 11–Oct 13	online/blended	Performance Optimierung – Kommunikation	HLRS (SCA)
Sep 13–15	online	Introduction to oneAPI, SYCL2020 and OpenMP offloading*	HLRS/INTEL
Sep 18–22	Stuttgart	CFD with OpenFOAM	HLRS
Sep 25–28	online	AMD Instinct GPU Training NEW	HLRS/AMD
Sep 25–29	Stuttgart	Introduction to Computational Fluid Dynamics	HLRS/DLR
Sep 25–Oct 27	online/blended	Linguistische Datenverarbeitung & Zeitreihenvorhersage	HLRS (SCA)
Oct 9–10	Stuttgart	Scientific Visualization	HLRS
Oct 16–20	Stuttgart	Parallel Programming Workshop (with TtT)	HLRS
Oct 23–27	Stuttgart	Julia for High-Performance Computing	HLRS
Oct 30–Dec 4	online/blended	Datenanalyse mit HPC	HLRS (SCA)
Nov 6–10	Stuttgart/online	Optimization of Scaling, I/O and Node-level Performance on Hawk	HLRS
Nov 21–24	Stuttgart/online	Modern C++ Software Design (Advanced)	HLRS
Nov 27–30	online	Advanced parallel programming with MPI and OpenMP (online)	JSC/HLRS
Dec 4–7	online	Introduction to GPU Programming using CUDA NEW	HLRS
Dec 11–15	Stuttgart/online	Fortran for Scientific Computing	HLRS

- Parallel Programming
- 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
- GPU Programming

* EuroCC 2 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 EuroCC 2.

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

Institute Abbreviations

AMD – Advanced Micro Devices, Inc. | DLR – German Aerospace Center | ETH – Scientific IT Services, ETH Zurich | F. ITWM – Fraunhofer Institute for Industrial Mathematics | F. SCAI – Fraunhofer Institute for Algorithms and Scientific Computing | HLRS – High-Performance Computing Center Stuttgart | HLRS (SCA) – Supercomputing Academy of the HLRS | INTEL – Intel Corporation | JSC – Jülich Supercomputing Centre | LRZ – Leibniz Supercomputing Centre | NEC – Nippon Electric Company | NHR@FAU – Erlangen National High Performance Computing Center | NVIDIA – Nvidia Corporation | VSC Vienna – Vienna Scientific Cluster | ZDV – Data Center, University of Mainz | IH – Center for Information Services and High Performance Computing (TU Dresden)

Workshops and Conferences in 2023

Mar 29 **High-Performance Computing in Crisis Situations**
In an event organized for representatives of public administration and institutions, HLRS staff presented case studies on how HPC and AI can support crisis management and help address global challenges.

Mar 31 **The ENRICH Project**
This event presented the results of a two-year collaborative research project focusing on identifying new opportunities for improving sustainability in IT and computing centers.

Apr 13–14 **35th Workshop on Sustained Simulation Performance**
Organized in cooperation with NEC, this annual meeting brings scientists, application developers, and hardware designers from different continents together to discuss hardware architectures, programming styles, and strategies for achieving the highest possible sustained application performance.

Jul 17 **Explaining, Understanding, Optimizing: The Role of AI in High-Performance Computing**
This digital workshop organized by the CIRCE project explored issues of explainability in AI models and how AI can be used in simulation.

Jul 26–28 **Summer School: Trust and Machine Learning**
This three-day meeting discussed ethical, epistemic, and practical facets of machine learning, focusing on the challenges of assessing the trustworthiness of such methods.

Sep 27 **CIRCE Digital Workshop**
The project CIRCE is evaluating what measures are needed to make HPC available to support decision making in crisis situations. This workshop presented an introduction to data processing and simulation using HPC systems.

Oct 12–13 **26th 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 18–19 **HPC Industry Summit**
Coordinated by HLRS as part of the EuroCC 2 and FF4EuroHPC projects, the conference demonstrated how simluation, AI, and high-performance data analytics can catalyze innovation, streamline processes, and maximize productivity in SMEs.

Nov 8 **50 Years of Scientific Collaboration**
This international research conference celebrated the anniversary of a partnership between the University of Stuttgart and the Donetsk National Technical University in Ukraine.

Nov 27–28 **CapeReviso Abschlusssymposium**
The project CapeReviso developed methods for using digital twins, machine learning, sensors, and other digital technologies to reduce conflicts between cyclists and pedestrians in cities. At this symposium, research partners presented results to representatives of local communities, city planners, citizen activists, and scientists.

Nov 30–Dec 1 **SAS23 – Reliability or Trustworthiness?**
Recent debates on AI have pointed out trust as a seminal issue. The philosophy of trust has developed fundamentally different notions of what this means, however. The Science and Art of Simulation 23 conference explored whether reliability or trustworthiness should be the goal of AI and simulation models.

Dec 5 **7th Industrial HPC User Round Table (iHURT)**
The annual iHURT meeting facilitates dialogue between HLRS and its industrial user community, focusing on innovative applications of HPC for research and development as well as challenges that industry faces in using HPC.

Organization Chart

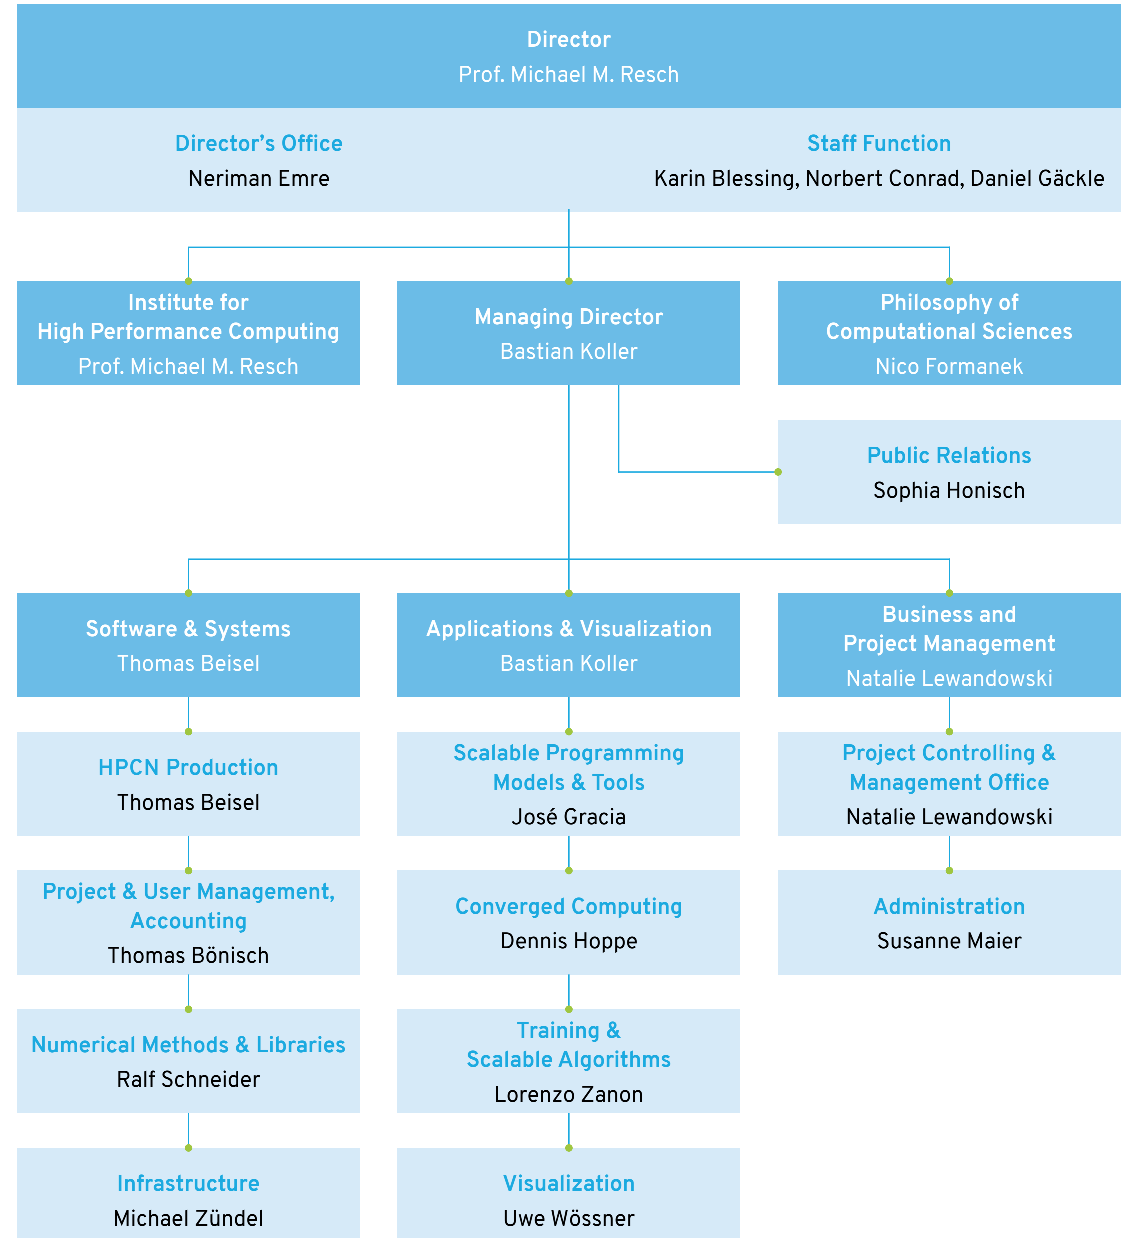

Departments

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
Advances the convergence of artificial intelligence (AI), cloud, edge, and quantum computing with high-performance computing (HPC). The group's goal is to make HPC more accessible by transforming operational models and promoting seamless workflows across the entire computing spectrum. This integration enables hybrid HPC/AI workflows that combine traditional simulations with AI techniques, including big data, machine learning, and deep learning. Additionally, the department works to lower the barriers to supercomputing and broaden the HPC ecosystem to new users through virtualization technologies like containers, workflow orchestration, and job scheduling. Consequently, we have developed expertise in creating and managing dynamic, scalable federated cloud computing services such as Gaia-X.

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

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

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

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

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

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

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

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.

Using high-performance computing, engineers gain high-resolution insights into how water flows through a pump, helping them to achieve more efficient designs more quickly.

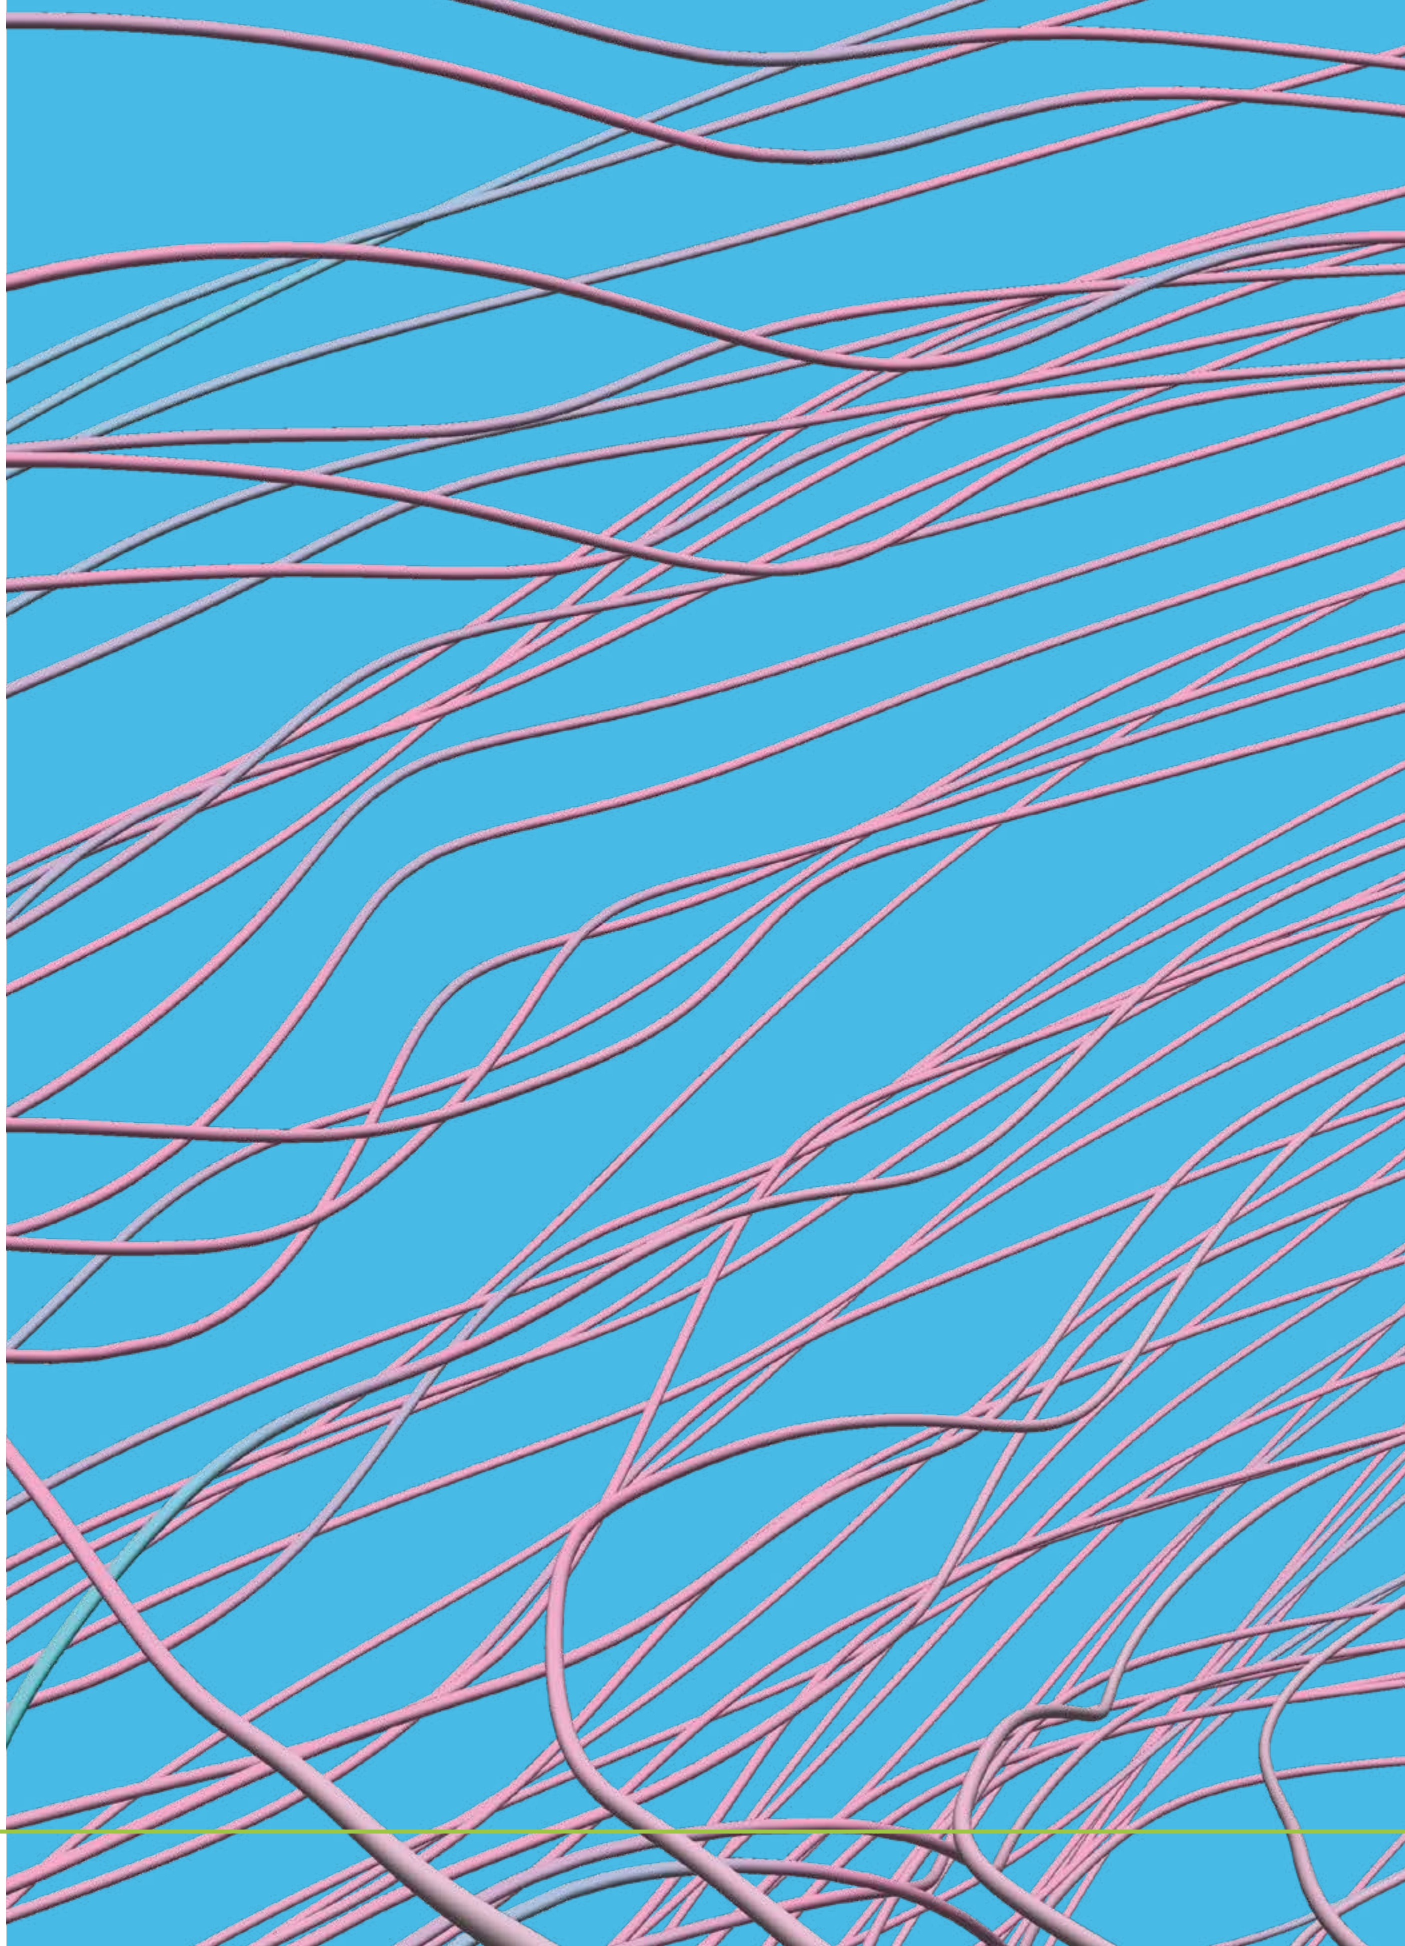

High-Performance Computing Center Stuttgart (HLRS)
University of Stuttgart
Nobelstraße 19
70569 Stuttgart, Germany

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

Translations
Sophia Honisch

Photography and images
Unless otherwise indicated, all images property of HLRS.
p. 8 ©HLRS/generated with Firefly by Groothuis; p. 12 Courtesy of Xiang Xu; p. 13 Courtesy of Bernhard Weigand; p. 14 Courtesy of Wolfgang Schröder; p. 15 Courtesy of Matthias Ibach; p. 16 Birk Heilmeyer and Frenzel Architects; p. 20 Fusion Medical Information / Unsplash; p. 22 National Supercomputer Center in Guangzhou; p. 24 Soft-Sim Consult; p. 26 Sandra Ritschel; p. 28 Florian Steciuk on Unsplash; p. 30 Courtesy of Andrea Beck; p. 36 Yüksel Patir; p. 38–39 IAG, University of Stuttgart; p. 40 stock.adobe.com / Christian; p. 42, 43 Karlsruhe Institute of Technology; p. 45–46 TU Berlin; p. 47, 48, 50 University of Helsinki; p. 59 Springer Verlag

Printing
gutenberg beuys, Feindruckerei GmbH, Langenhagen

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

Institutional affiliations

University of Stuttgart
Germany

GCS
Gauss Centre for Supercomputing

Funding for Hawk provided by:

Baden-Württemberg
MINISTRY OF SCIENCE, RESEARCH AND ARTS

Federal Ministry of Education and Research

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

FSC www.fsc.org
RECYCLED
Paper made from recycled material
FSC® C009051

Print product with financial climate contribution
ClimatePartner.com/10951-2404-1003

Follow us on:

LinkedIn

X (Twitter)

High-Performance Computing Center Stuttgart
www.hlrs.de

Cover:
Using vsualization technologies including digital twins, HLRS has helped engineers at Junkers Aircraft and Kasaero GmbH in designing a modernized version of the iconic A50 Junior airplane, and supported energy company EnBW in the planning and ongoing construction of a new pumped storage power plant in the Black Forest.

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®.