2020

HLRS

High-Performance Computing Center | Stuttgart

ANNUAL REPORT

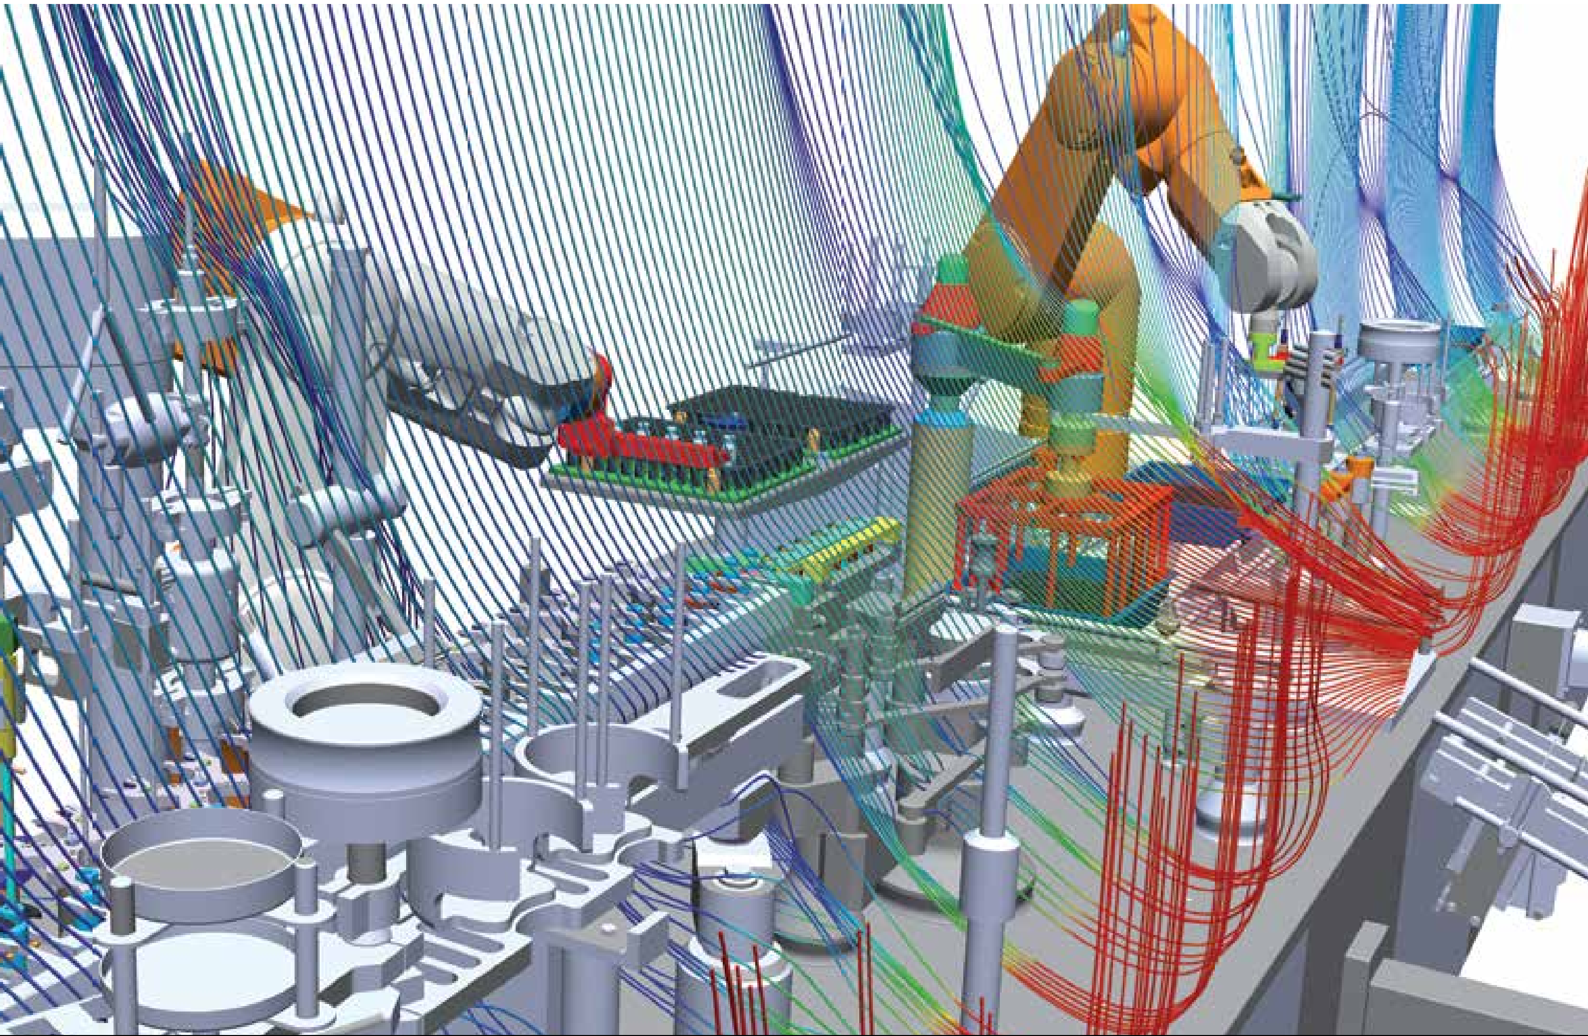

ENERGY HEALTH CLIMATE MOBILITY PHILOSOPHY

2020 HLRS ANNUAL REPORT

The High-Performance Computing Center Stuttgart (HLRS) was established in 1996 as the first German national high-performance computing (HPC) center. As a research institution affiliated with the University of Stuttgart and a founding member of the Gauss Centre for Supercomputing, HLRS provides HPC services to academic users and industry. We operate leading-edge HPC systems, provide training in HPC programming and simulation, and conduct research to address key problems facing society and the future of supercomputing. Among our areas of expertise are parallel programming, numerical methods for HPC, visualization, grid and cloud computing concepts, data analytics, and artificial intelligence. Our system users conduct research across a wide range of scientific disciplines, with an emphasis on computational engineering and applied science.

Director's Welcome

Grußwort

Welcome to the 2020 annual report of the High-Performance Computing Center Stuttgart (HLRS). We are pleased to share highlights from our activities over the past year with you, including information about some exciting expansions in our focus.

The global coronavirus pandemic meant that 2020 was a challenging year for institutions everywhere, and HLRS was no exception. Just two weeks after a celebratory inauguration of our new supercomputer, Hawk, in February, restrictions caused by COVID-19 hit HLRS. Interruption of supply chains and travel restrictions not only made the installation of our new system much more complicated than we had anticipated, but also created an extremely difficult situation for our users. In the end, it took six months longer than planned to get Hawk operational. During the installation period we were grateful for the support and understanding of our partners and users, and we thank all of you who have helped HLRS through these troubling and complex times.

COVID-19 did not only pose challenges, however. It also gave us an opportunity to demonstrate important ways in which computer simulation can help in addressing crisis situations. Two weeks after the government implemented lockdown measures to control spread of the virus, HLRS was approached by the Bundesinstitut für Bevölkerungsforschung (BIB, Federal Institute for Population Research) to support their simulations to predict how the pandemic would progress and to plan for the weeks and months ahead. Over a weekend HLRS made resources available to optimize, parallelize, and operationalize the BiB's code, and since then has been providing BIB with support and compute services. In addition, HLRS supplied scientific researchers compute time as part of both European and German

Herzlich willkommen zum Jahresbericht 2020 des Höchstleistungsrechenzentrums Stuttgart (HLRS). Wir freuen uns, Ihnen Höhepunkte aus dem vergangenen Jahr und einige spannende Projekte vorzustellen.

Die weltweite Coronavirus-Pandemie machte 2020 zu einem Jahr voller Herausforderungen für alle möglichen Einrichtungen. Auch das HLRS war keine Ausnahme. Nur zwei Wochen nach der feierlichen Einweihung unseres neuen Supercomputers Hawk im Februar trafen die COVID-19-Einschränkungen das HLRS. Die Unterbrechung von Lieferketten und Reisebeschränkungen machten die Installation unseres neuen Systems viel komplizierter als erwartet und unseren Nutzern das Leben extrem schwer. Letztlich dauerte es sechs Monate länger als geplant, Hawk in Betrieb zu nehmen. Während der Installationsphase waren wir dankbar für die Unterstützung und das Verständnis unserer Partner und Nutzer. Wir danken allen für ihre Hilfe in diesen unruhigen, komplizierten Zeiten.

Aber COVID-19 war nicht nur eine Herausforderung. Es war auch eine Chance, um aufzuzeigen, wie Computersimulationen bei der Bewältigung von Krisensituationen gut helfen können. Zwei Wochen nach der Lockdown-Einführung der Regierung zur Pandemiekontrolle wurde das HLRS vom Bundesinstitut für Bevölkerungsforschung (BiB) gebeten, dessen Simulationen zu unterstützen, um den Pandemieverlauf vorherzusagen und für die kommenden Wochen und Monate zu planen. An einem Wochenende stellte das HLRS Ressourcen zur Verfügung für die Optimierung, Parallelisierung und Operationalisierung des BiB-Codes. Seitdem bietet es dem BiB Support und Rechendienstleistungen. Zudem stellte das HLRS Forschern im Zuge der europäischen und deutschen initiatives to fight COVID-19, and received support through the AMD COVID-19 HPC Fund that will become available in 2021 and enable us to increase our capacity for crisis computing.

Despite the pandemic, HLRS maintained and initiated new activities to support the European supercomputing community. As a managing partner HLRS launched three European projects — EuroCC, CASTIEL, and FF4EuroHPC — to create a European network of HPC expertise in the framework of EuroHPC. Through these projects HLRS will take a leading role in creating a networked and world-class European HPC ecosystem. Within a few weeks of the COVID-19 lockdown, we also transformed our HPC training courses into a full-fledged online program and began providing intensive user support using web-based tools; both developments have expanded our ability to share our expertise. Furthermore, our efforts to create a sustainable high-performance computing center were honored by two environmental certifications (Blue Angel and EMAS).

We were pleased to see Hawk land in 16th place in its debut on the Top500 List of the world's fastest supercomputers. Also exciting was an agreement we reached with Hewlett Packard Enterprise to expand our CPU-based flagship system with graphic processing units (GPUs). This change in Hawk's system architecture will complement its power for simulation and open new opportunities for deep learning and artificial intelligence applications.

In 2020 HLRS also began research projects aimed at preparing for future technology developments and for future strategic settings. As part of the NFDI4CAT consortium HLRS became deeply involved in the creation of a German National Research Data Infrastructure. As a partner in the SEQUOIA project, we also acquired funding to begin exploring the possibilities of quantum computing, particularly for industry. Additionally, with the support of the State of Baden-Württemberg, we will extend our research at the intersection of the humanities, social sciences, simulation, and digitalization. These projects will help strengthen HLRS's position as one of the most innovative HPC centers in Europe.

Because travel was restricted in 2020, international collaborations were limited to giving talks and attending workshops virtually. However, international collaboration and solidarity with our partners were important for HLRS in 2020, and we would like to thank all who offered their support during the crisis. In Stuttgart we welcomed an extension of our collaboration with the Fraunhofer Society, entering into a formal collaboration agreement with the Fraunhofer Institute for Manufacturing, Engineering and Automation (IPA).

HLRS's key performance indicators highlight a very successful year. Our third-party funding saw an increase over previous years, while our income from industrial users of our HPC systems was again at a high level. We also maintained strong participation in our HPC training programs, including a large increase of international participation due to our online courses, and were delighted to see two doctoral students complete their PhD's at the center.

Finally, I would like to use this occasion to thank the supporters and funders who have made HLRS's successes in 2020 possible. We look forward to continuing working with you to find innovative ways of using HPC and other advanced digital technologies to address the most pressing challenges facing science, industry, the HPC community, and society at large.

With best regards,

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

Initiativen zur Bekämpfung von COVID-19 Rechenzeit zur Verfügung und wird durch den AMD COVID-19 HPC Fund unterstützt, der ab 2021 verfügbar ist und es uns ermöglicht, unsere Kapazitäten für Crisis Computing zu erhöhen.

Trotz der Pandemie setzte das HLRS seine Aktivitäten zur Unterstützung der europäischen HPC-Gemeinde fort und startete neue Projekte. Das HLRS ist Koordinator von drei europäischen Projekten – EuroCC, CASTIEL und FF4EuroHPC – um im Rahmen von EuroHPC ein europäisches HPC-Kompetenznetzwerk zu schaffen. Durch diese Projekte übernimmt das HLRS eine Führungsrolle bei der Schaffung eines erstklassigen vernetzten europäischen HPC-Ökosystems. Nach ein paar Wochen im COVID-19-Lockdown stellten wir unsere HPC-Schulungen auf ein vollwertiges Online-Programm um und begannen mit einem umfassenden Nutzer-Support mit webbasierten Tools. Durch beide Entwicklungen können wir unser Fachwissen besser weitergeben. Zudem wurden unsere Bemühungen um ein nachhaltiges Höchstleistungsrechenzentrum durch zwei Umweltzertifizierungen (Blauer Engel und EMAS) gewürdigt.

Erfreulicherweise schaffte es Hawk bei seinem Debüt in den Top500 der schnellsten Supercomputer der Welt auf Platz 16. Interessant ist auch unsere Partnerschaft mit Hewlett Packard Enterprise, die zur Erweiterung unseres CPU-Vorzeigesystem um Grafikprozessoren führen wird. Diese Änderung der Hawk-Systemarchitektur wird seine Simulationsleistung stärken und neue Deep-Learning- und KI-Anwendungen ermöglichen.

2020 lancierte das HLRS auch Forschungsprojekte zur Vorbereitung auf künftige Technologieentwicklungen und strategische Umfelder. Als Teil des NFDI4Cat-Konsortiums ist das HLRS maßgeblich am Aufbau einer nationalen Forschungsdateninfrastruktur beteiligt. Als Partner im SEQUOIA-Projekt haben wir auch Mittel eingeworben, um die Möglichkeiten von Quantencomputern auszuloten, insbesondere für die Industrie. Zudem werden wir mit Unterstützung des Landes Baden-Württemberg unsere Forschung an der Schnittstelle von Geistes- und Sozialwissenschaften, Simulation und Digitalisierung ausbauen. Diese Projekte werden die Stellung des HLRS als eines der innovativsten HPC-Zentren Europas weiter stärken.

Da die Reisefreiheit im Jahr 2020 eingeschränkt war, beschränkte sich die internationale Zusammenarbeit auf virtuelle Vorträge und Workshops. Internationale Kooperation und Solidarität mit unseren Partnern waren 2020 jedoch wichtig für das HLRS, und wir danken allen, die in der Krise ihre Unterstützung angeboten haben. In Stuttgart freuten wir uns über die Intensivierung der Zusammenarbeit mit der Fraunhofer-Gesellschaft über eine formelle Kooperationsvereinbarung mit dem Fraunhofer-Institut für Produktionstechnik und Automatisierung (IPA).

Die wichtigsten HLRS-Leistungsindikatoren bezeugen ein sehr erfolgreiches Jahr. Gegenüber den Vorjahren wurden mehr Drittmittel eingeworben, wobei die Einnahmen durch Nutzer unserer HPC-Systeme aus der Industrie erneut hoch waren. Unsere HPC-Schulungsprogramme stießen nach wie vor auf großes Interesse, wobei insbesondere die internationale Beteiligung durch unsere Online-Kurse stark anstieg. Erfreulicherweise schlossen auch zwei Doktoranden ihre Arbeit am Zentrum ab.

Abschließend möchte ich bei dieser Gelegenheit noch den Unterstützern und Förderern danken, die die Erfolge des HLRS im Jahr 2020 ermöglichten. Wir freuen uns darauf, weiterhin mit Ihnen gemeinsam innovative Einsatzmöglichkeiten für HPC und andere fortschrittliche digitale Technologien zu finden, um die drängendsten Herausforderungen für Wissenschaft, Industrie, die HPC-Gemeinde und die Gesellschaft allgemein anzugehen.

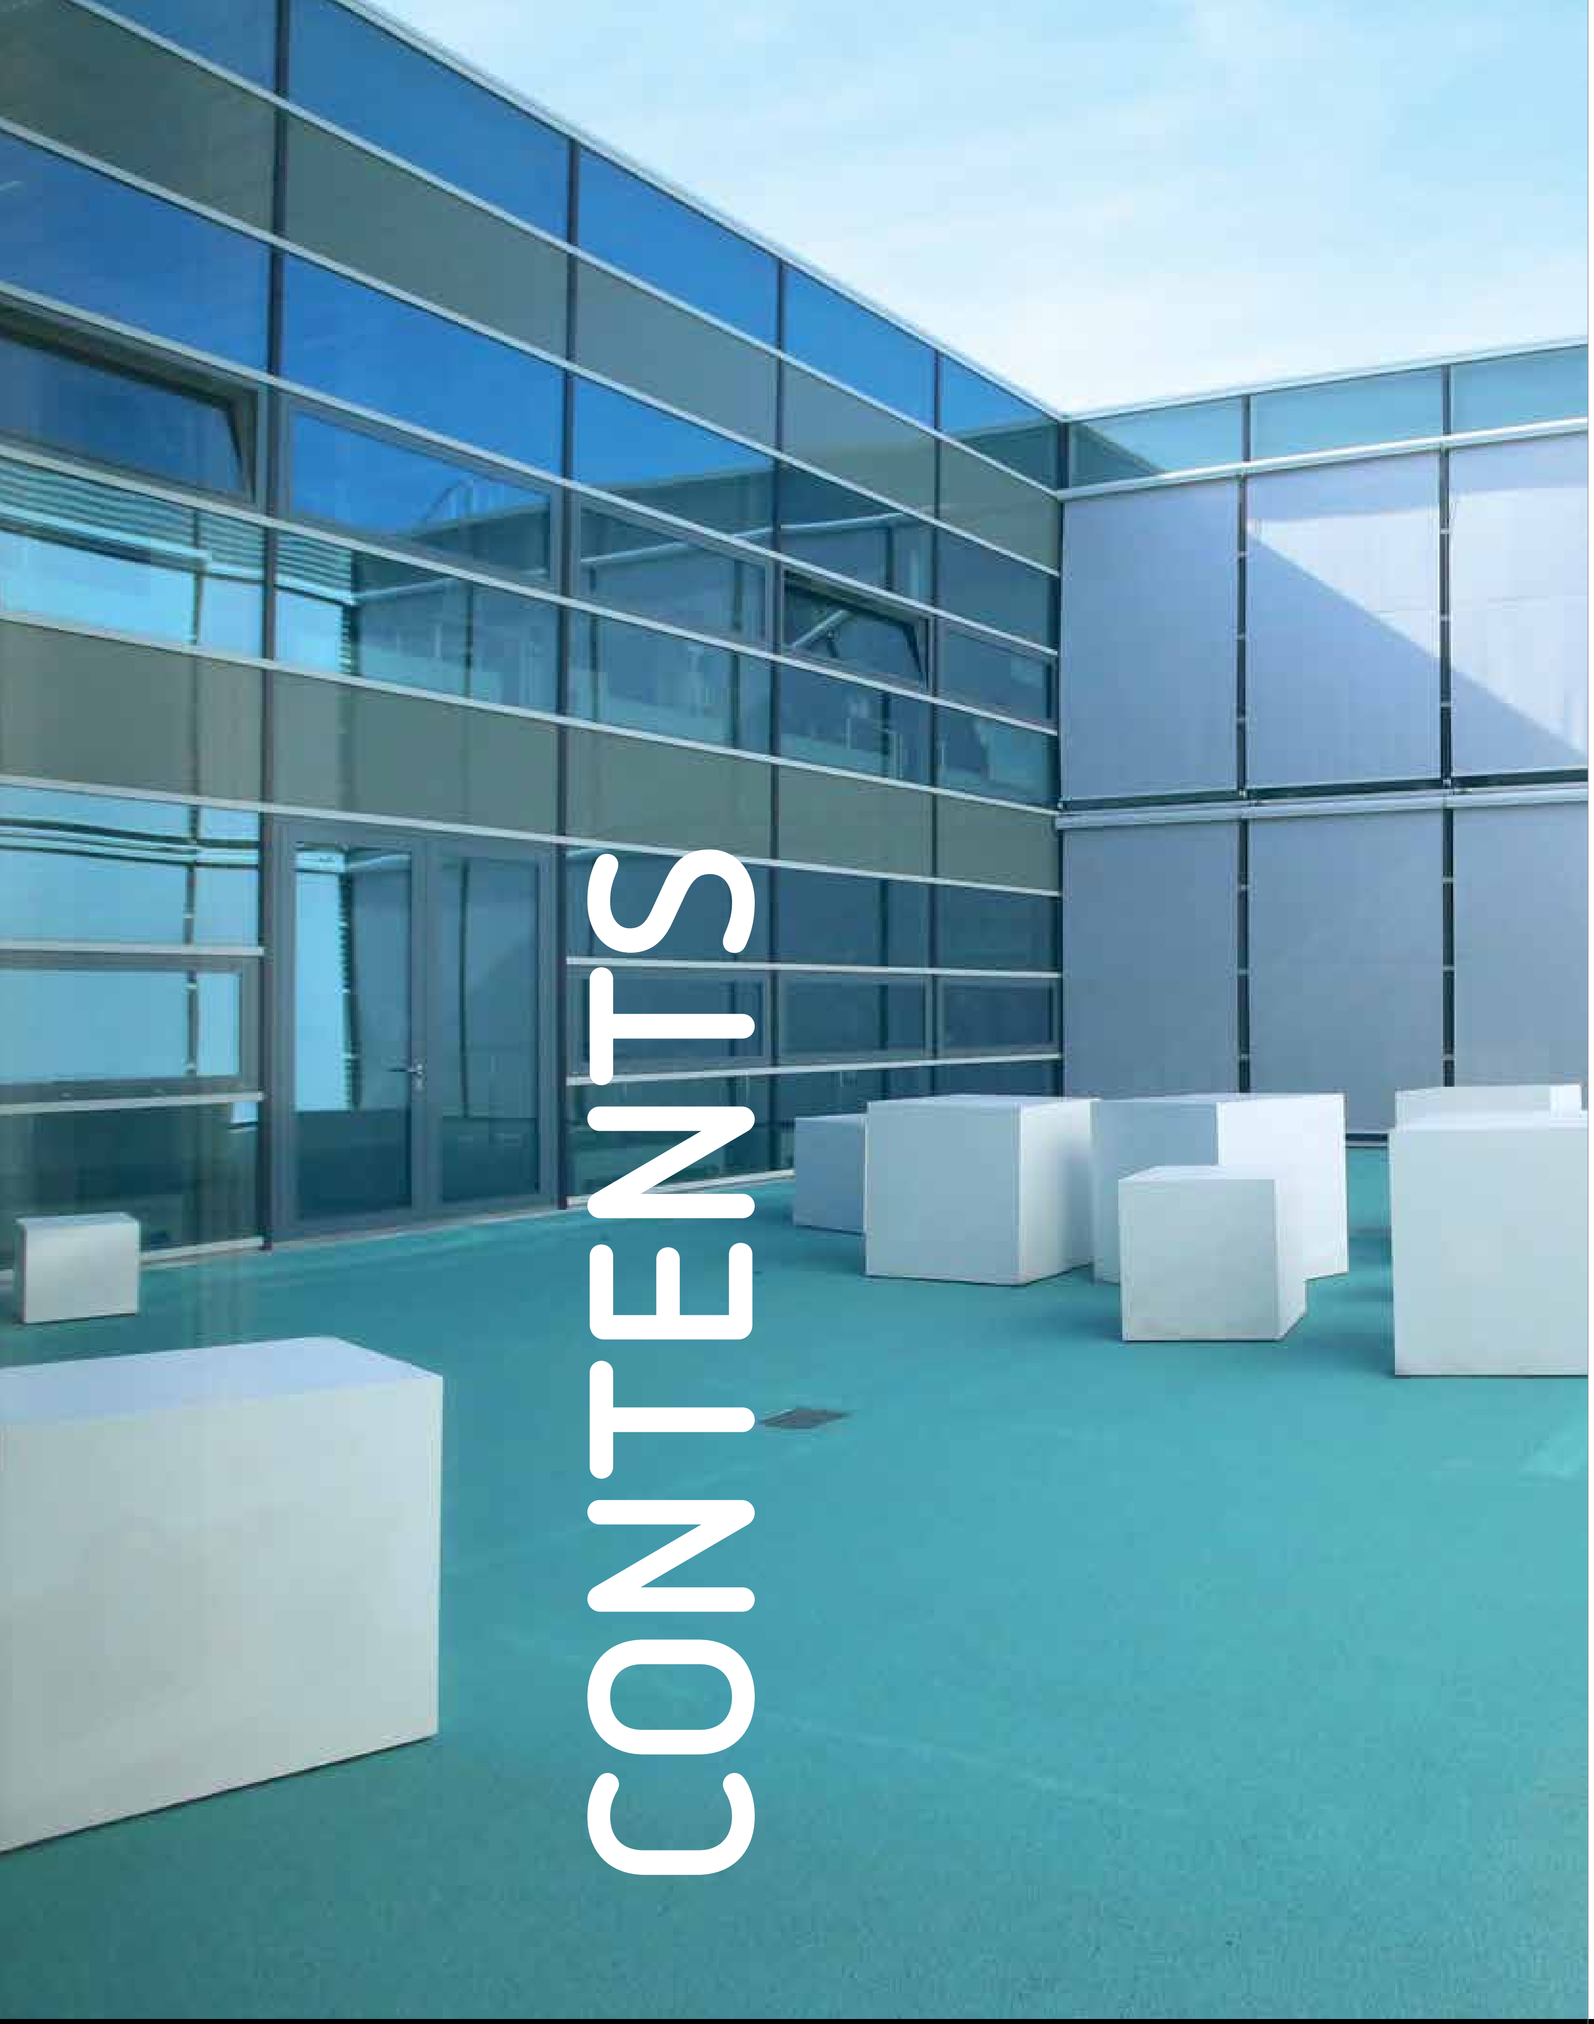

CONTENTS

8 Spotlight

14 News Highlights

36 HLRS by the Numbers

40 User Research

58 About HLRS

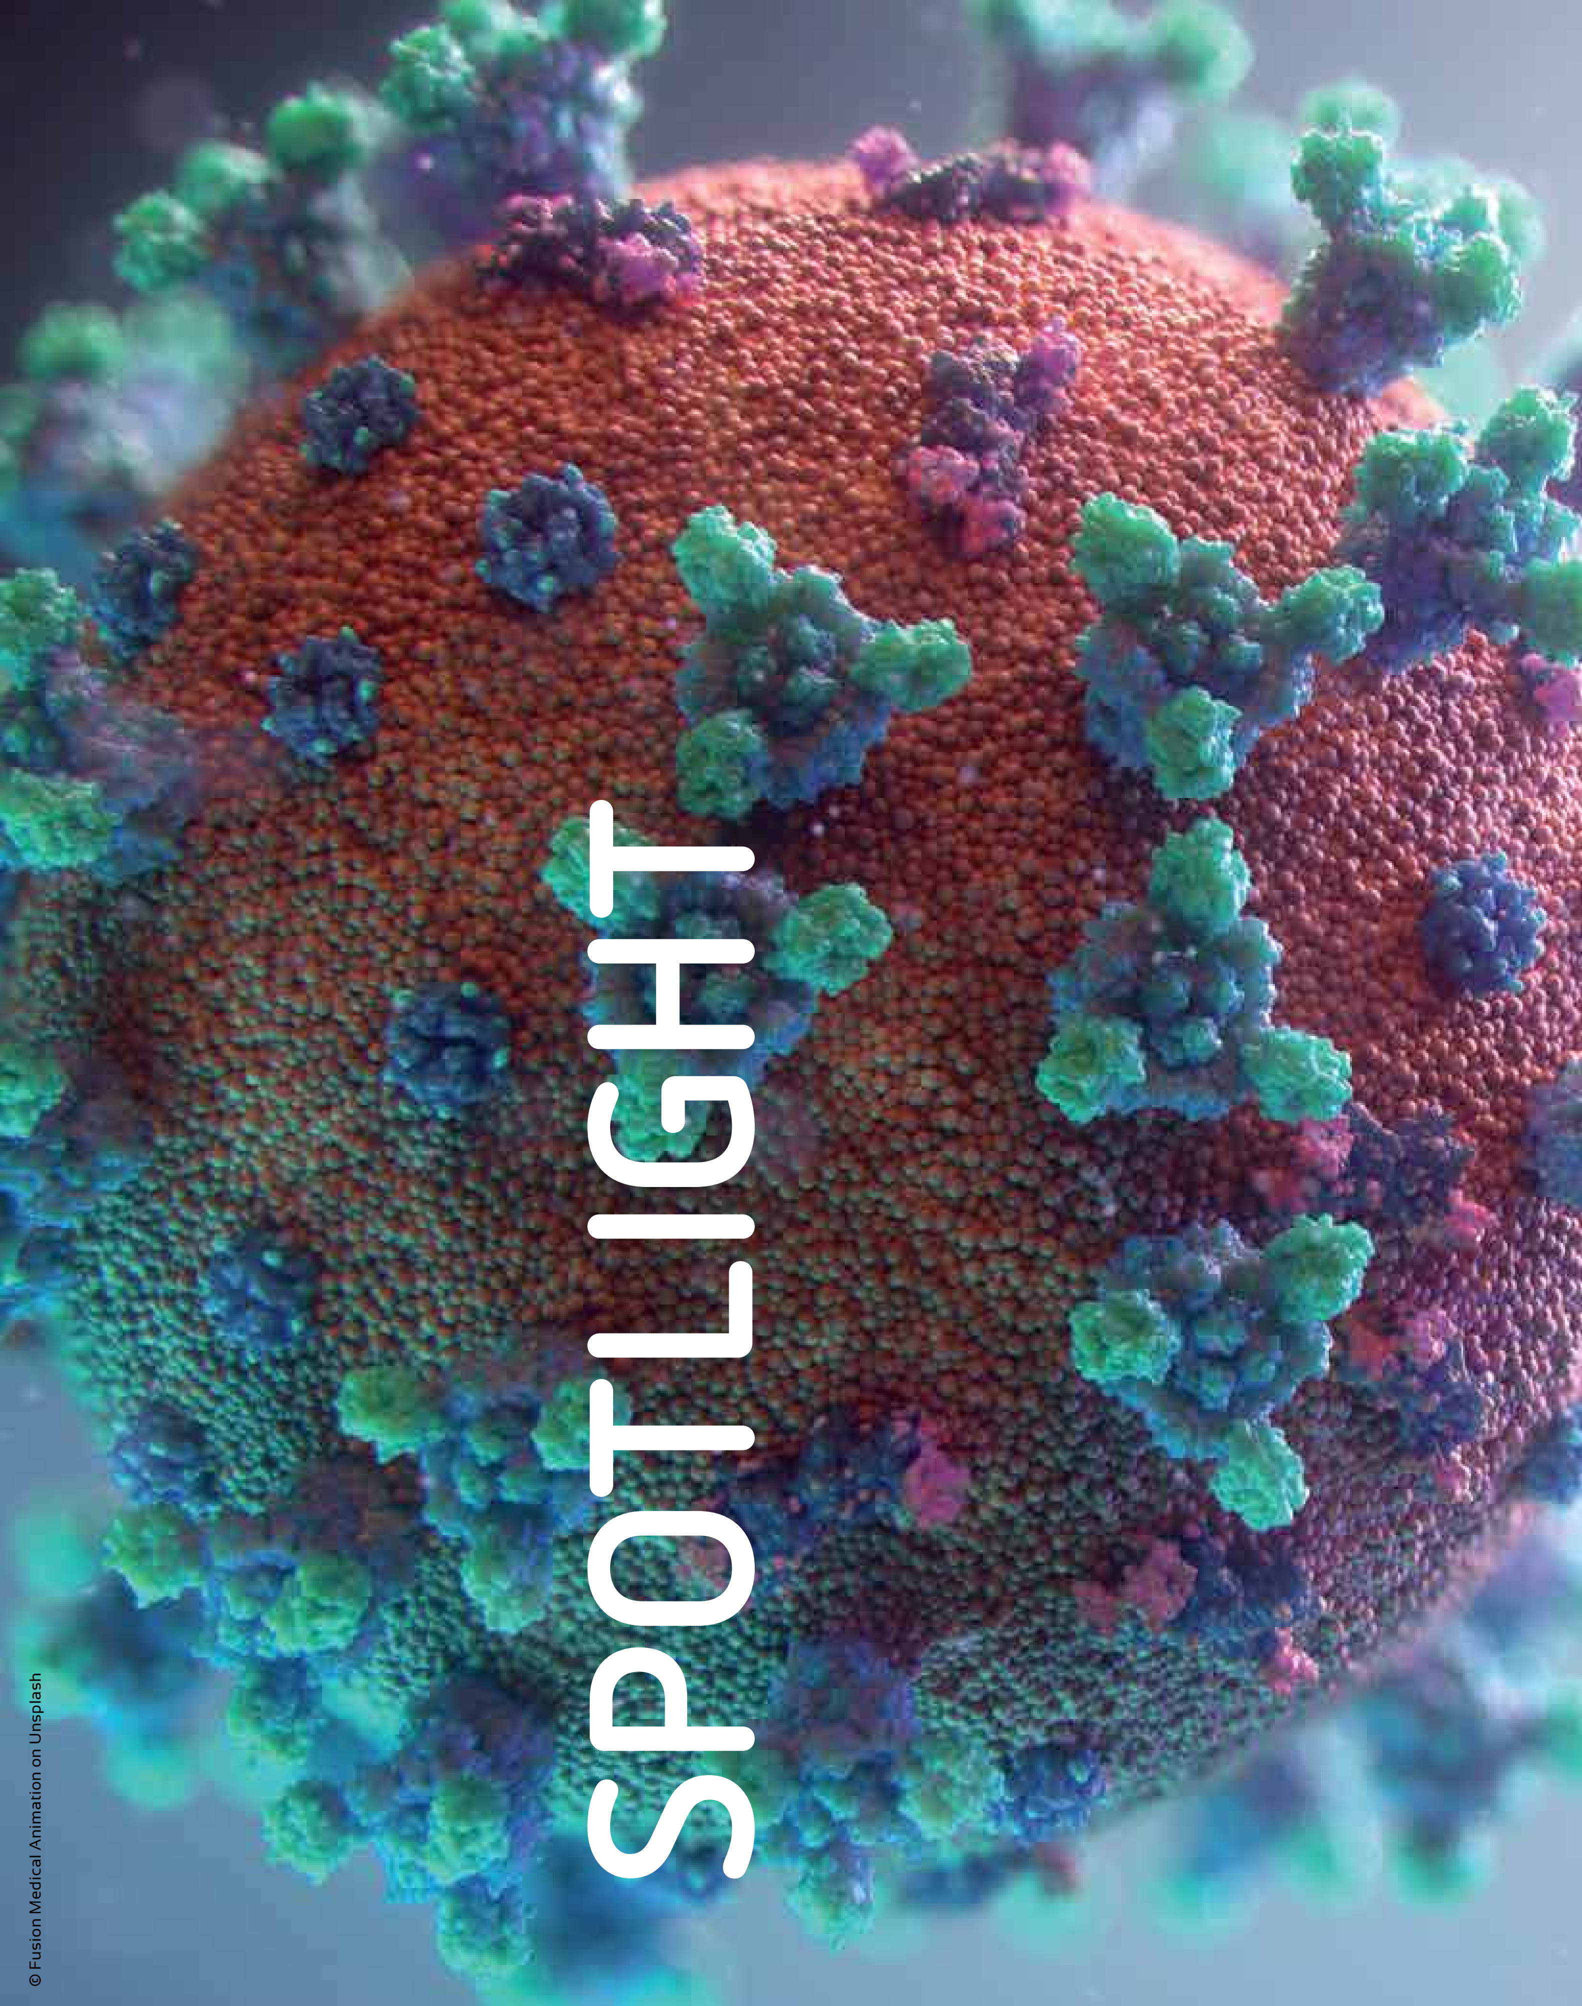

© Fusion Medical Animation on Unsplash

Supercomputing in the Year of COVID-19

HLRS's experiences during the exploding coronavirus pandemic demonstrated how advanced computing technologies can help respond to crisis situations.

In mid-February 2020, the High-Performance Computing Center Stuttgart (HLRS) celebrated the beginning of a new era in its history. At a standing-room-only ceremony, high-ranking government officials, industry representatives, researchers, and friends of HLRS gathered to celebrate the inauguration of the center's new 26-petaflop HPE Apollo supercomputer, called Hawk. Little did the festive crowd realize, however, that within just a few short weeks gatherings of this sort would become impossible — a new, unanticipated era of a very different kind was about to arrive.

By early March, the SARS-CoV-2 virus had begun spreading in Germany, and previously unimaginable measures were soon implemented to control a deadly new pandemic. On March 13, the state of Baden-Württemberg closed all schools to on-site learning, and on March 16, the University of Stuttgart recommended that all employees not essential for the operation of campus facilities work from home. Although case numbers in Germany fell over the summer, the absence of an effective vaccine and arrival of an even more dangerous second wave of disease in the fall meant that many restrictions remained in effect for the rest of the year. Even at the end of 2020, just a few employees were onsite at HLRS daily, with the majority working remotely from home, networked by email and online videoconferencing software.

The unprecedented disruption caused by the COVID-19 pandemic affected HLRS in many ways, forcing the center to quickly adapt to a new reality. At the same time, however, it highlighted the critical roles that high-performance computing (HPC) and other advanced computing technologies now play in addressing public health emergencies. For HLRS, 2020 was a year of rising to meet many new kinds of challenges, and foreshadowed how supercomputing could evolve in the coming years to prepare for other kinds of crises.

Providing computational resources for basic research

The computational sciences have come to play an essential role in disciplines such as physiology, immunology, virology, and drug discovery. Scientists now routinely use simulation, machine learning, and high-performance data analytics to study many kinds of pathogens and how they cause infection, gaining essential information for identifying potentially effective drugs, vaccines, and other strategies for fighting disease. Using algorithms designed to run on supercomputers, scientists generate testable predictions of biological activity that occurs at scales that are too small or that is too complex to study using experimental approaches alone. This, in turn, accelerates research, as it helps investigators to focus their wet lab experiments on specific hypotheses that have the highest likelihood of success.

In 2020 the international scientific community deployed computational approaches in many different ways to address the COVID-19 pandemic. Some, for example, used such approaches to gain a better understanding of molecular characteristics of

This figure shows a simulation run on HLRS's system of the degree to which ICU capacities would have been utilized by patients with COVID-19 during the first wave of the pandemic had efforts not been taken to contain disease spread. Levels above 50% are considered critical, as 50%-65% of ICU beds are typically occupied by patients with other illnesses. The upper graph is based on reported case numbers while the lower one assumes that during the early weeks of the pandemic, actual cases were undercounted by 50%.

Source: RKI, General Federal Employment Agency, Federal Statistical Office of Germany, own calculations.

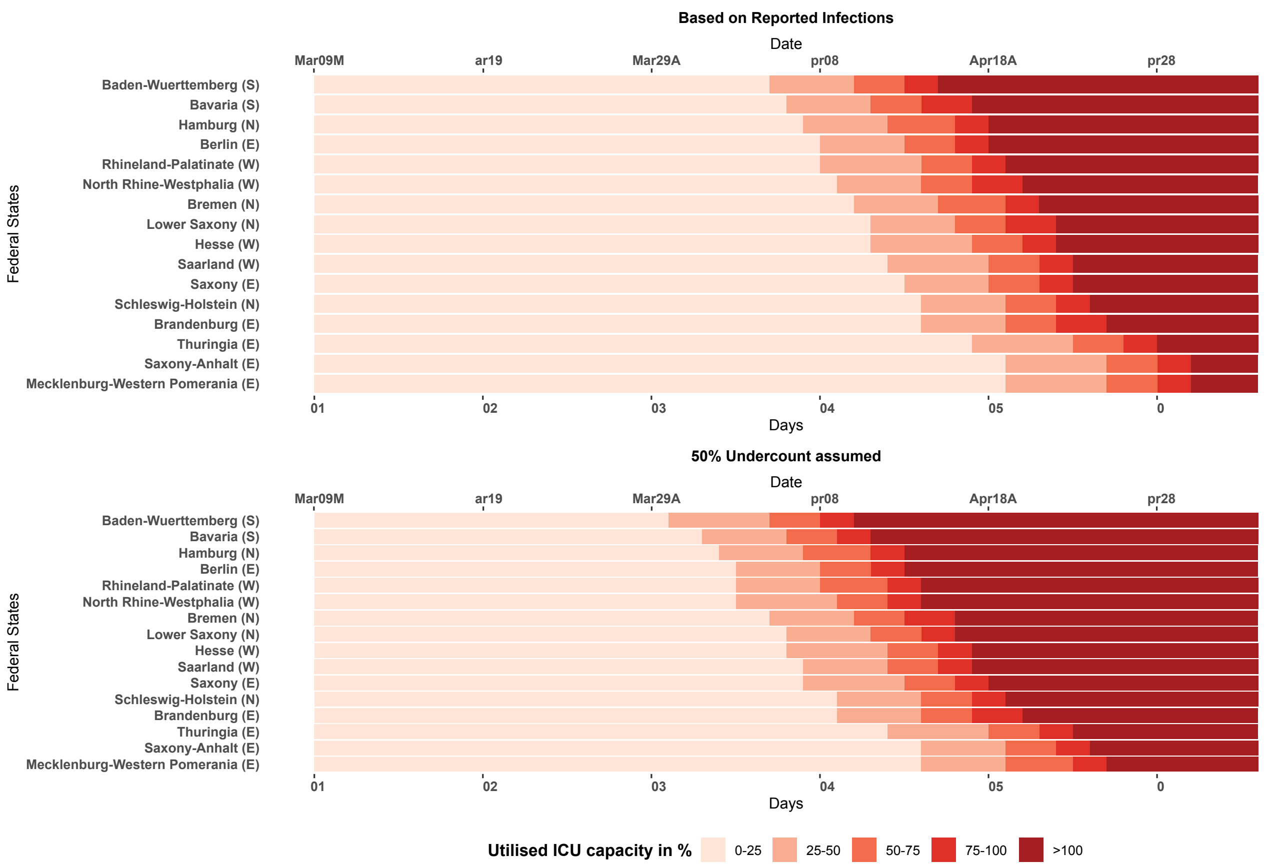

the virus. Others conducted virtual drug screenings, using large data sets to predict which existing drug compounds might be repurposed to incapacitate the virus. Clinical scientists began using computational approaches to analyze electronic patient records, looking for important features in the data that could reveal more personalized treatment strategies. Epidemiologists also used data to track and predict how disease spreads in populations, providing insights that public health officials can use to prepare for and contain an outbreak. An unprecedented global scientific response using these and other approaches rapidly produced a massive knowledge base that helped guide the search for vaccines, cures, and effective public health interventions.

At the same time that HLRS was going on lockdown in March, the center stepped forward to provide computing muscle for COVID-19 research. Stuttgart quickly joined its partners in the Gauss Centre for Supercomputing — the alliance of Germany's three national supercomputing centers — in announcing expedited access to its systems for scientists investigating the virus and its accompanying disease. Although HLRS is perhaps better known for its support of computational engineering, its HPC infrastructure is also suitable for health-related research, and in the following weeks and months, scientists came forward and began using Hawk to contribute to the fight against the novel coronavirus. Eight COVID-related projects made use of the HLRS system over the course of the year, while the center's invitation for additional users remains open.

In several of the approved projects, scientists focused on gaining a better understanding of the virus's physical structure, including features that are involved in the infection of human cells. This included research aimed at modeling the COVID-19 spike protein, which initiates infection when it interacts with the human ACE2 receptor on the surfaces of cells. Others used an approach called molecular dynamics to simulate at extremely high resolution how proteins on the surface of the virus change shape. Understanding these phenomena better could help identify drugs that could inhibit the interaction between the virus and its human host as it gains entry to cells. Another user of HLRS's systems received an allocation to investigate specific factors in the clinic that lead to negative outcomes in some patients. Because of the complexity of these kinds of problems, HLRS's computing systems provided essential tools for this research.

Predicting ICU demand at hospitals

Beginning with the arrival of the COVID-19 pandemic, Germany's national strategy focused on containing the spread of the virus to prevent overwhelming burdens on hospital capacity. This involved implementing, and at times relaxing, a wide range of interventions, including closing businesses, prohibiting public events, issuing stay-at-home orders, and requiring the wearing of protective masks in public places. Throughout 2020 this strategy was largely successful in preventing unmanageable demand for intensive care units, although the arrival of the pandemic's second wave once again forced the country to implement new restrictions to prevent a flood of critical cases. Until vaccines are effectively deployed, it is likely that such cycles will continue. In an effort to provide tools that can help public health agencies and governments to determine what protective interventions are needed, HLRS researcher Dr. Ralf Schneider collaborated with investigators at the German Federal Institute for Population Research (Bundesinstitut für Bevölkerungsforschung, BiB) to develop software for predicting spikes in demand for intensive care units (ICUs) across Germany up to eight weeks into the future. This COVID "weather report" could be used to help strike a sustainable balance between protecting public health and minimizing other social, economic, and personal impacts of the pandemic. A preprint publication describing the tool was published in December on the medRxiv preprint server for the health sciences.

The software was developed on HLRS's Hawk supercomputer, but is designed to be usable by data scientists anywhere. As Schneider explained, "The model is based on epidemiological parameters that correspond to available data in Germany. However, the code is structured so that other countries could potentially adapt it to their own situations, assuming the relevant data is available."

Also important to the success of this effort was EXCELLERAT, a European project led by HLRS that was created to support the engineering industry in the movement toward exascale supercomputing technologies. HLRS staff, together with EXCELLERAT partner SSC-Services, showed that technologies developed for the engineering branch could also be useful in addressing public health challenges. The collaborators implemented a secure data transfer tool that enabled the BiB researchers to quickly move data between Hawk and systems at their offices in Wuppertal, giving them access to HLRS's capabilities for large-scale data analysis without needing to be present onsite.

New digital infrastructure for COVID research

In addition to providing high-performance computing time to individual scientists for their research, HLRS also leads and contributes to dozens of funded research projects addressing technical challenges facing high-performance computing. In one new project approved in the summer by the European Commission's Horizon 2020 program, HLRS will help develop an important tool for managing public health during pandemics.

Since March 2020, medical centers in Europe and around the world have accumulated masses of COVID-19 patient data documenting patient characteristics, treatment plans, and disease progression. Initially, however, these data were largely gathered and saved by different healthcare providers in their own local databases, with few connections between them. Because modern biomedical and epidemiological studies gain statistical power by using larger datasets, bringing this information together would provide investigators with a valuable resource.

The new project, called ORCHESTRA, is being led by Prof. Evelina Tacconelli at the University of Verona, Italy, and includes 27 partner institutions from 15 countries in Europe, Africa, South America, and Asia. It will develop a multilayered data infrastructure for collecting and sharing COVID-related patient data from across Europe and other parts of the world. As part of this effort, HLRS will work together with scientists at the high-performance computing centers CINECA (Italy) and CINES (France), to develop the data management infrastructure.

According to the ORCHESTRA plan, national data hubs in participating countries will collate collections of data, while a centralized, cloud-based portal will enable researchers to access, share, and link together data stored in the various national hubs. The resulting virtual cohort, potentially including hundreds of thousands of patients, will enable a range of multidisciplinary studies investigating COVID-19 genetics, epigenetics, immunology, and epidemiology, among other topics.

In addition to helping design the infrastructure, HLRS will also provide expert support in standardizing data collection, storage, and curation across all of the national data hubs. Such harmonization is essential for ensuring that data sets originating in many different places can be seamlessly integrated in ways that maximize their utility and impact.

The ORCHESTRA team anticipates that this new platform will support studies to improve public health and vaccine strategies. It could also provide a model for collecting and sharing data in the event of future pandemics.

HPC training and user support teams discover new opportunities online

More locally, the COVID crisis affected HLRS in other practical ways. Because of "social distancing" regulations that prohibited large gatherings of people, it became impossible for the center to offer its busy calendar of HPC training courses in a normal classroom setting (see pages 68/69). Instead, within a few weeks of the beginning of the lockdown, instructors began revising courses for an online format, holding them using the videoconferencing platform Zoom.

Despite the challenge of preserving the hands-on, interactive nature of the courses, this digitalization of HLRS's training program had unique benefits. Not only did the center avoid a significant decrease in the total number of course participants it sees in a normal year, but the courses also reached an audience across a broader geographical range than normal, with 38% of all participants logging in from outside Germany.

"Although we obviously look forward to the time when we can once again offer courses in person, the experience over the last several months has shown how online learning could enhance our abilities to increase expertise in the use of high-performance computing," said Dr. Rolf Rabenseifner, who leads HLRS's training program. "At this point, we anticipate that HLRS could potentially continue to offer online courses in parallel with its core training program even after the current pandemic has passed." In this case, the necessity to adapt to COVID-related restrictions could have unforeseen future benefits.

Similarly, HLRS's user support team was forced to quickly adapt its April code porting and scaling workshop to an online format. Using videoconferencing software, scientific computing experts at HLRS worked one-on-one with individual teams of scientists in dedicated breakout rooms to address problems they encountered with their codes.

"In optimization workshops, we do not just hold a lecture or show a set of slides," said Björn Dick, a member of the HLRS user support team. "This new approach enables us to work collaboratively online, to really dig into the code together and make changes. By using remote control functionality, I can briefly take over control of a user's screen, write in some code to explain what I mean, and the person on the other side immediately understands it and can continue on his own."

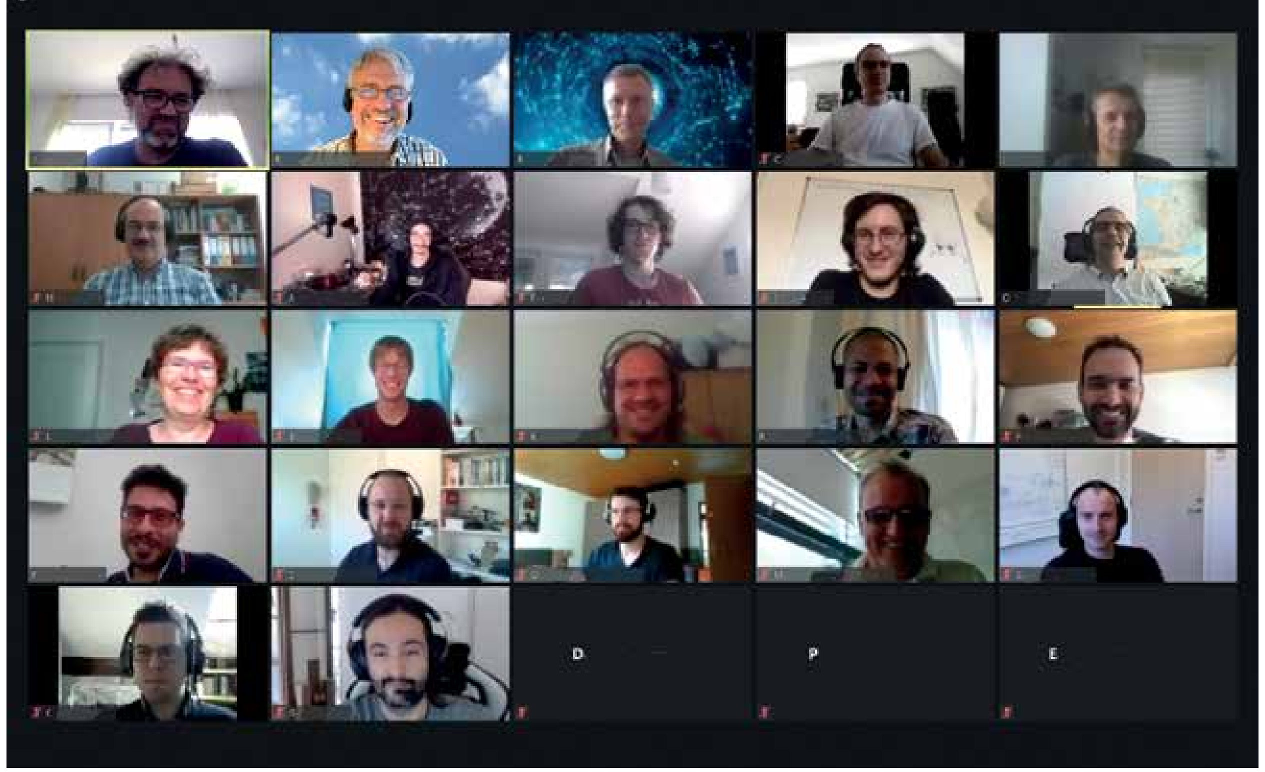

During the COVID-19 pandemic HLRS continued its training program using webconferencing tools.

The success of the event also inspired the HLRS team to begin using videoconferencing tools in other settings where in-person meetings weren't practical. "In one sense COVID forced us to change the way we work," said PUMA Department Leader Dr. Thomas Bönisch, "but it also enabled us to develop a user support solution we hadn't considered previously. This has improved our ability to help HLRS's system users even outside the context of the performance optimization workshop."

Crisis computing: preparing for an uncertain future

Unfortunately, COVID-19 is unlikely to be the last pandemic that humanity will face. Moreover, considering global atmospheric and social instabilities likely to created by climate change, ongoing risks of chemical or nuclear accidents, vulnerabilities to natural disasters, and other global threats, simulation and related fields such as artificial intelligence will likely be important in quickly responding to new kinds of problems.

To bolster HLRS's ability to meet these challenges, computer hardware manufacturer AMD, whose processors make up the core of Hawk, announced in September 2020 a donation of 10 new computing nodes to HLRS as part of its AMD COVID-19 High Performance Computing Fund. The additional computing power, consisting of AMD EPYC processors and Instinct accelerators and due to be delivered in early 2021, will be reserved for pandemic-related research. In the future, the new resources will also expand HLRS's capacity to address urgent computing needs when sudden crises arise.

As HLRS Director Prof. Dr. Michael Resch explains, "The COVID-19 crisis has been a big wakeup call in Germany and across Europe, showing that new challenges can develop very suddenly and have widespread impacts across societies. Because high-performance computing provides a backbone for tools that accelerate scientists' ability to address these challenges, it is clear that investments and preparation are needed now to be ready the next time a similar challenge comes along. At HLRS and at other HPC centers across Europe we anticipate that implementing new resources and strategies to be ready to handle emergency situations where simulation and data analysis can help will be a big focus in the coming years."

Supercomputers alone will never be the cure for threats like the COVID-19 pandemic, but they will continue have important roles to play in preparing for and successfully addressing them. *(CW)*

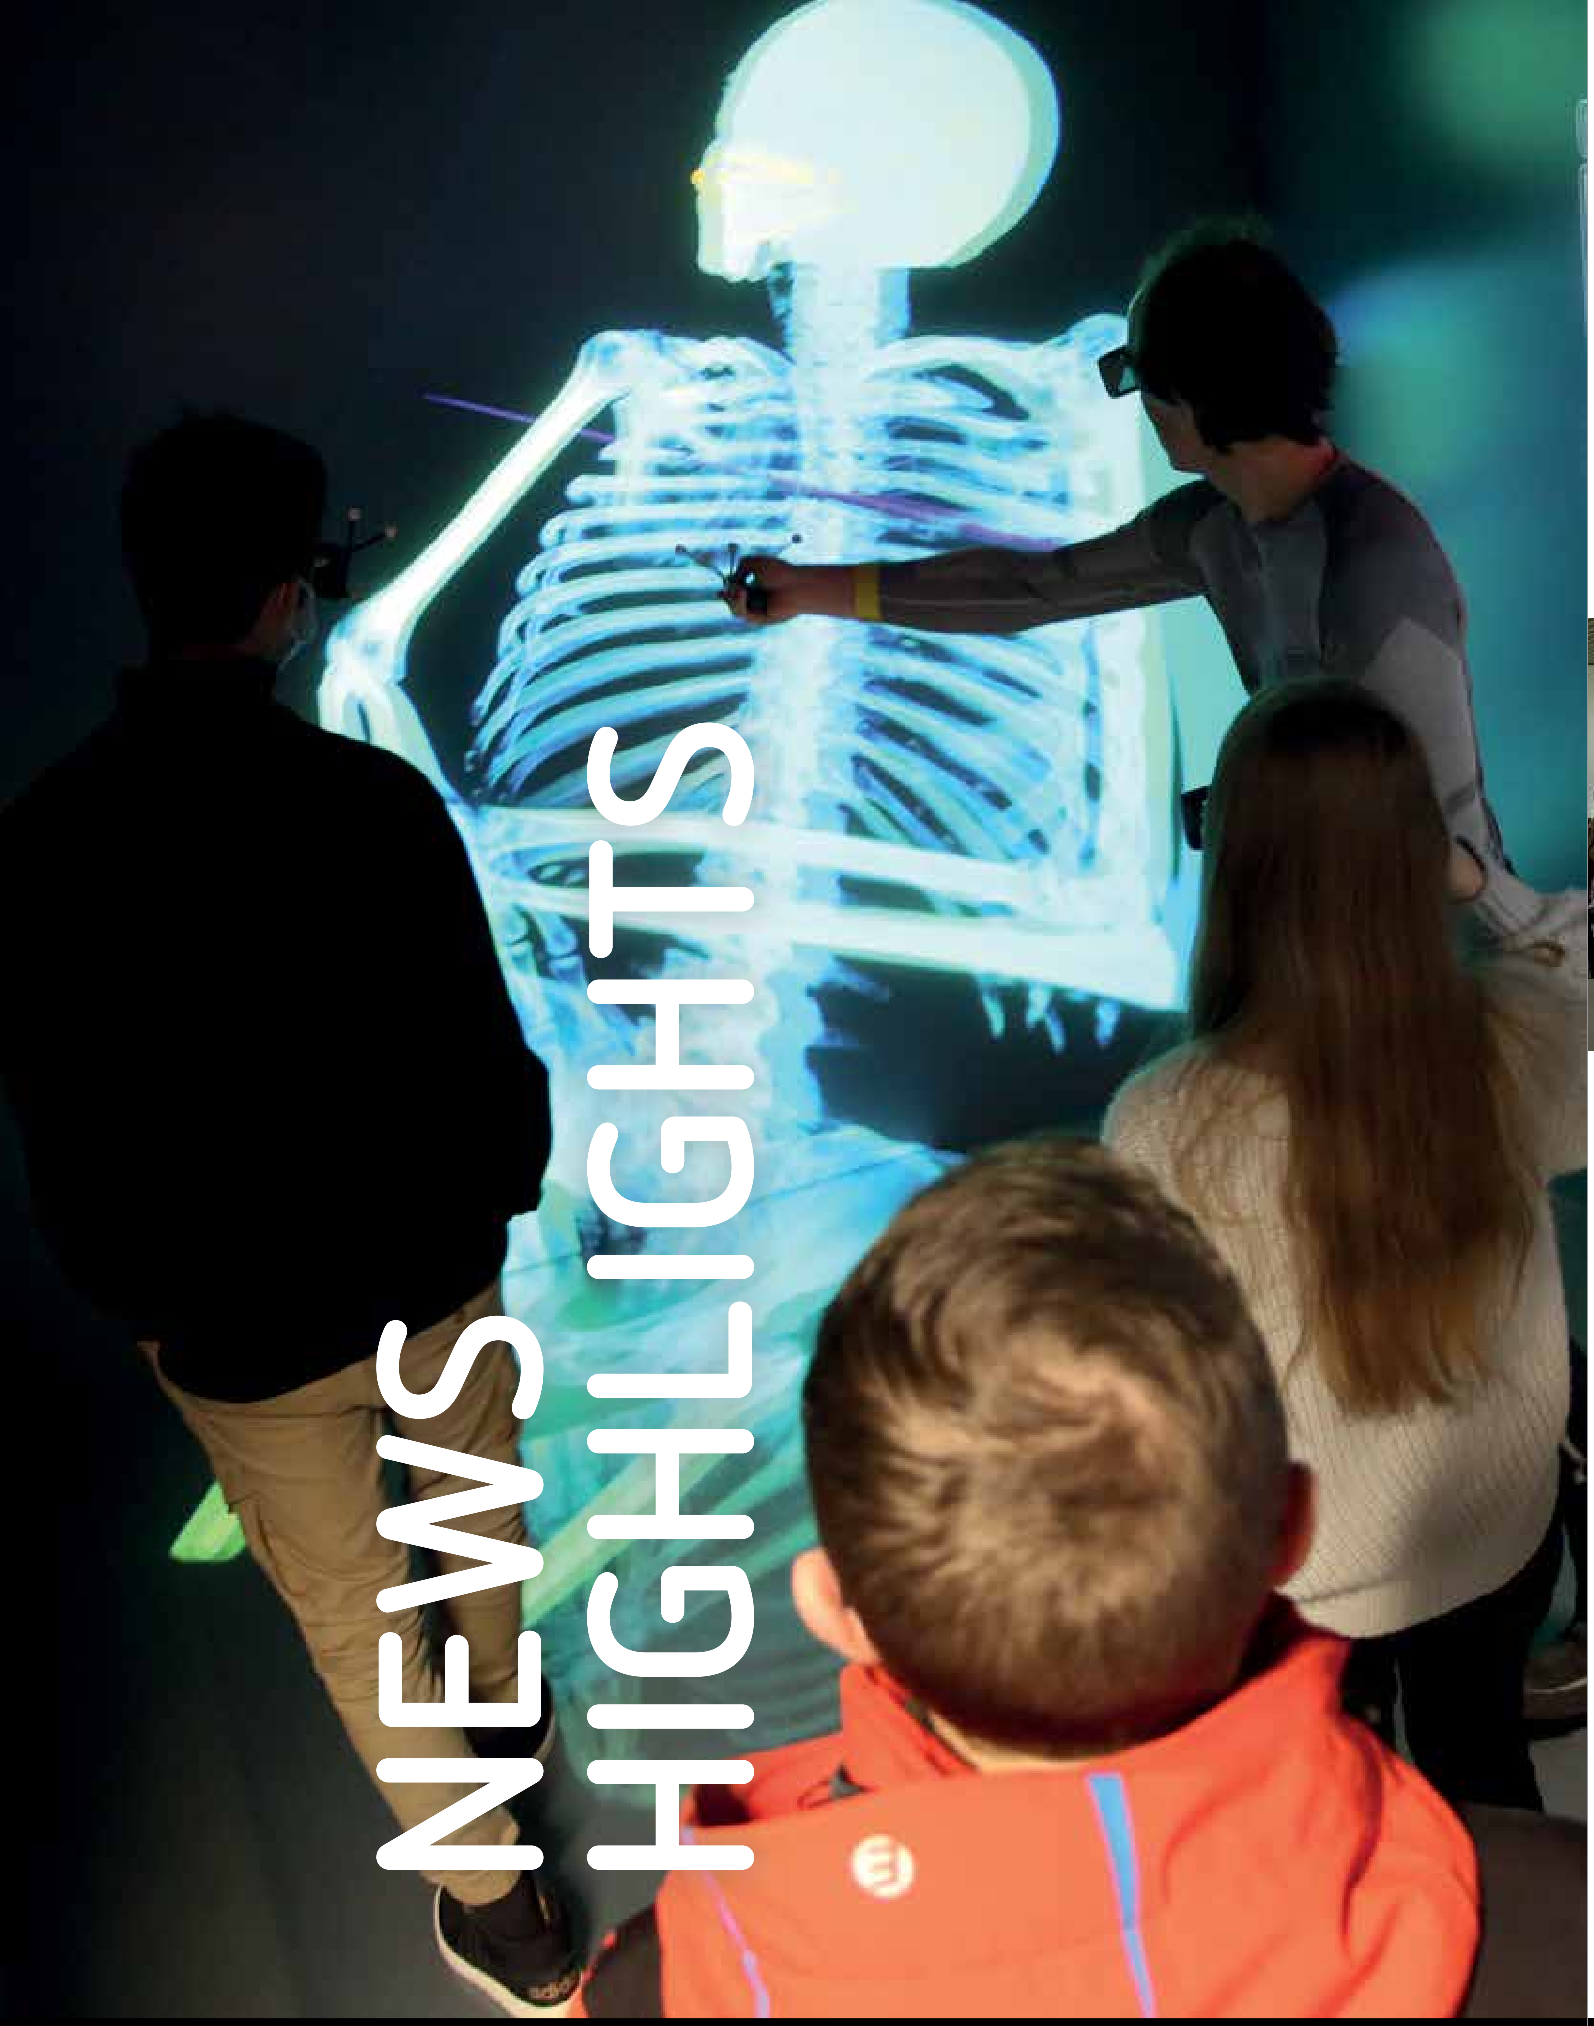

NEWS HIGHLIGHTS

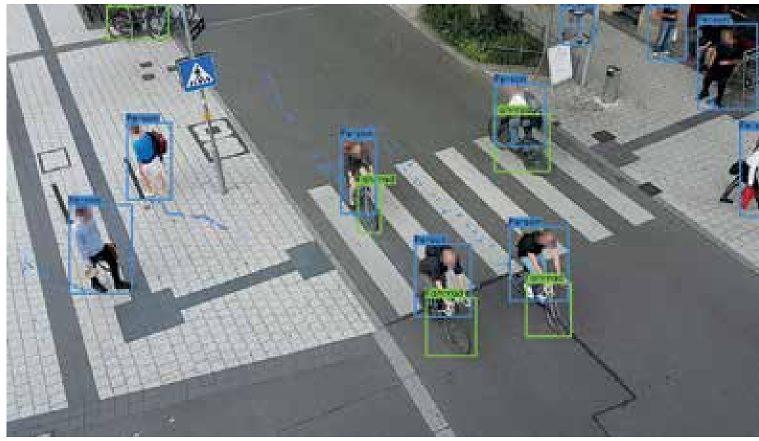

Supercomputing Academy Launches New Courses, Awards First Certifications to HPC Experts

The Supercomputing Academy provides continuing professional education in a variety of topics related high-performance computing for HPC users in industry. In 2020, the program launched four new course modules: Performance Optimization; Economy and Sustainability; Visualization; and Data Management and Analysis. It also celebrated an important milestone, for the first time awarding its highest certification, "HPC Expert," to two participants. Reflecting on his achievement, participant Sascha Michael Scholz commented, "Through the Supercomputing Academy I have gained first-hand insights into the key problems of supercomputing. I now feel like I am in a good position to work competently and professionally in the field of HPC." Following the conclusion of the MoeWE project (which developed and launched the Supercomputing Academy) in March 2021, the training program will continue in a self-sustaining manner. *(CW)*

Safer Cities for Cyclists and Pedestrians

When pedestrians and cyclists use the same spaces, stress often results. In a project sponsored by the German Ministry of Transport and Digital Infrastructure called Cape Reviso, HLRS and partners will investigate how simulation could help alleviate such problems. After conducting a long-term observation of pedestrian, bicycle, and motorized vehicle movements at a busy location, they will use machine learning to analyze the data and develop a model. After combining this model with urban emotions data gathered by investigators from the Karlsruhe Institute of Technology, they will create a digital twin — a comprehensive virtual reality simulation — of the location. Planners will then use the model to see how cyclists and pedestrians interact, and then virtually test the effects of interventions such as changes in signage. They could then investigate these scenarios further in the "Living Lab" at the physical location, comparing real-world behaviors with the simulation. *(CW)*

© Fundación Épica

Media Solution Center Baden-Württemberg Grows Network

Founded in 2018 by HLRS, the ZKM, and the Hochschule der Medien, the Media Solution Center Baden-Württemberg is a membership organization that connects arts organizations, media companies, and technologists to promote collaborative innovation at the intersection of the arts, science, and technology. In 2020, MSC membership grew dramatically, not only in Germany but across Europe. Prominent regional institutions such as the Kunstmuseum Stuttgart, Staatstheater Stuttgart, Stuttgarter Kammerorchester, and University of Tübingen were joined by internationally renowned organizations such as the Fundación Épica La Fura dels Baus (Spain), Festival Quartieri dell'Arte de Viterbo (Italy), Czech National Ballet (Czech Republic), and University of the Arts Utrecht (Netherlands). MSC membership is open to interested organizations. Read more at msc-bw.com *(CW)*

Marian Albers

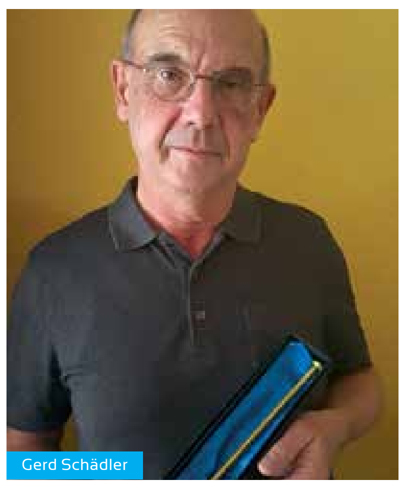

Gerd Schädler

Lars Schäfer

© Images courtesy of the Golden Spike Award winners.

Golden Spike Awards Announced at 23rd Annual Results and Review Workshop

Due to the COVID-19 pandemic, the 23rd Annual Results and Review Workshop took place online on October 8-9. The program featured 25 scientific talks and a virtual poster session with 20 posters made available as PDF files. Dr. Dietmar Kröner, a professor at the University of Freiburg and vice-chairman of the HLRS steering committee, also announced the winners of the 2020 Golden Spike Awards, which recognize excellence in research and innovative applications of HPC resources. This year's winners were Marian Albers of RWTH Aachen University for turbulence simulations that could help in designing more efficient airfoils, Gerd Schädler of the Karlsruhe Institute of Technology for his team's contributions to high-resolution models of climate change impacts in Africa, and Lars Schäfer of Ruhr University Bochum for molecular dynamics simulations of ABC transporters, a class of cell membrane proteins. *(CW)*

"Biedermeier" in the Digital Age

Speaking at the festival "Unter Beobachtung: Kunst des Rückzugs" (Under Observation: the Art of Retreat), HLRS Director Michael Resch presented a keynote lecture titled "Rückzug in die totale Öffentlichkeit - Das Paradox des digitalen Biedermeier" (Retreat into Total Publicity: the Paradox of the Digital Biedermeier). The talk compared the classic Biedermeier age in the nineteenth century, marked by a withdrawal into the private sphere, with contemporary trends in which we often use digital media to publicize details about our private lives. The lecture set the tone for the three-week festival, which through art projects and events presented across the greater Stuttgart region engaged with the effects of digital surveillance, data insecurity, and social media, as well as strategies and spaces for hiding from constant observation. *(CW)*

HLRS Runners Support Children's Hospice

On July 25, HLRS for the second time participated in the Hand in Hand Spendenlauf, a charity run organized to benefit the Stuttgart Hospice for Children and Youth. Due to the COVID-19 pandemic the event was held virtually, with 43 HLRS employees and their family members logging their distances covered at various locations using a mobile app. The HLRS team, High-Performance Runners, raised approximately 1,500 Euros by running the equivalent of 1,452 laps (580 km in total). The achievement nearly doubled last year's total of 860 laps. High-performance computer manufacturer Hewlett Packard Enterprise sponsored the team, and the money raised will contribute to facilities and programs that support severely ill children and their families. *(CW)*

High-Resolution Modelling of Air Quality

As members of the project Open Forecast, HLRS researchers have created a high-resolution tool for visualizing air quality forecasts for the city of Stuttgart. The model, developed by investigators at the University of Hohenheim and run on HLRS's supercomputer, combines the WRF weather and climate model with atmospheric chemistry calculations to model current and future concentrations of heavy particulate matter, nitrogen dioxide, and other pollutants. Visualization experts at HLRS transformed the simulations into an interactive 3D projection that integrates geodata and building information. The model simulates 24-hour cycles in Stuttgart at a resolution of as small as 50 meters, also accounting for the city's complex geography. The goal of Open Forecast is to make large, publicly available datasets more understandable and usable in decision making by policy makers, government officials, and the general public. *(EG)*

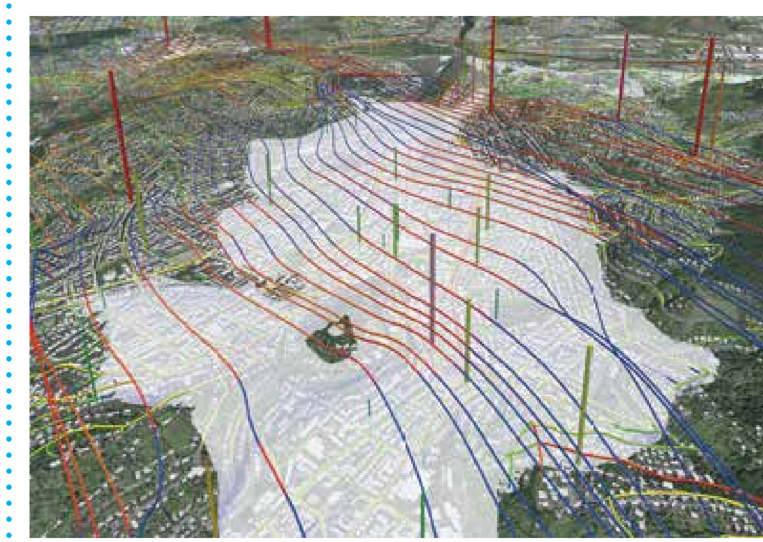

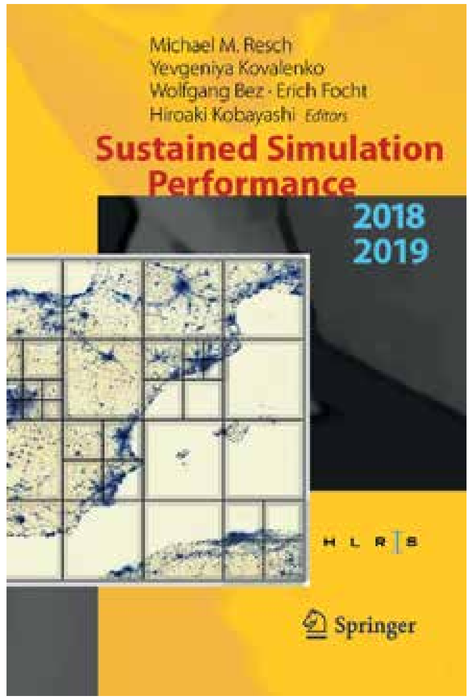

Book Publication: Sustained Simulation Performance 2018 and 2019

This book presents the state of the art in high-performance computing (HPC) on modern supercomputer architectures. It addresses trends in hardware and software development, as well as the future of HPC systems and heterogeneous architectures. Contributions cover a broad range of topics, from improved system management, to computational fluid dynamics, to high-performance data analytics and novel mathematical approaches for large-scale systems. In addition, articles explore innovative fields like coupled multi-physics and multi-scale simulations. Contributions are based on selected papers presented at the 26th and 28th Workshops on Sustained Simulation Performance, held at HLRS in October 2017 and 2018, and the 27th and 29th Workshops on Sustained Simulation Performance, held at the Cyberscience Center, Tohoku University, Japan, in March 2018 and 2019. *(CW)*

HLRS at 2020 ISC and SC Online Conventions

HLRS is a regular participant in the annual ISC High Performance (ISC) and International Conference for High Performance Computing, Networking, Storage, and Analysis (SC). Because of the COVID-19 pandemic, both 2020 meetings were held as online events. HLRS created virtual booths and presented talks with updates on several ongoing research projects. At ISC, presentations focused on the SODALITE, HIDALGO, and TalPas projects, while at SC, HLRS representatives and collaborators provided updates on the EuroHPC, CATALYST, SODALITE, HIDALGO, FF4EuroHPC, HPCWE, and EXCELLERAT projects, as well as new smart city approaches using digital urban twins. Together, the presentations offered a wide variety of insights into HLRS's research to develop technical solutions, address global challenges, and cultivate the growth of HPC expertise across Europe. *(CW)*

Fourth Industrial HPC User Roundtable (iHURT) Takes Place Online

SICOS BW and HLRS brought together representatives from large and small companies to discuss challenges they face when using high-performance computing. In his keynote presentation Dietmar Krase from Porsche AG focused on issues resulting from the different software architectures at the company and HLRS, and from the need to manage large data sets. In a virtual round table, participants exchanged insights concerning data management and remote post-processing, as well as software licensing, cloud computing, and compatibility between commercial software packages and MPI. HLRS's Dr. Thomas Bönisch reported on the latest developments in HPC technologies and support services, while Dr. Andreas Wierse and Nicole Prange of

SICOS BW introduced the Supercomputing-Akademie, a training program for industry. HLRS Director Michael Resch concluded with an outlook on how computational approaches combining HPC and AI could benefit industry. *(CW)*

Training Course Looks at Key Issues in Technology Transfer

This two-day course organized by HLRS in partnership with the program Eurolab4HPC equipped academic researchers with a basic understanding of the entrepreneurial process, strategies for generating business models, and managerial challenges that are a part of technology transfer. Focusing on typical difficulties encountered when founding and managing academic entrepreneurial ventures, the instructors analyzed the environments of academic entrepreneurship, presented potential opportunities, and provided advice on how to pursue those opportunities. The lecturers also addressed analytical and administrative domains of technology transfer, as well as managerial behavior. In this way, the course combined theoretical insights on entrepreneurship with practical considerations concerning the presentation of technologies and business ideas. *(CW)*

ChEESE Prototype for Simulation Workflows in Geoscience Research

As part of a project called ChEESE, HLRS is developing tools and workflows for predicting the effects of catastrophic events such as earthquakes, tsunamis, and volcanic eruptions. Models of potential risks could both identify vulnerable cities and regions and help responsible officials respond more quickly when catastrophic events occur. Because such models must process massive amounts of data, HLRS this year presented a prototype that enables geoscientists to integrate scientific workflows and move data between different computing platforms. This could make it possible, for example, to coordinate and run tasks that are distributed among their own workstations, institutional clusters, and supercomputers. In the context of ChEESE, HLRS is also performing code optimization for geoscience software and is developing visualization tools that decision makers could potentially use for disaster planning. *(CW)*

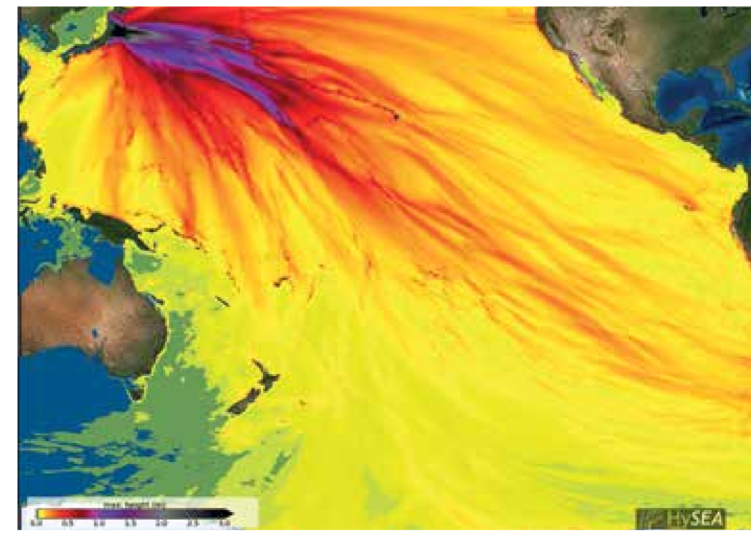

© EDANAY Group, University of Malaga

Putting Ethics into Practice in Artificial Intelligence

The AI Ethics Impact Group (AIEIG) proposed the first system for measuring and implementing European ethical principles in the development and use of artificial intelligence algorithms. Based on a similar model to ratings schemes for energy efficiency, the framework could be used to assess how well AI tools embody shared values such as transparency, accountability, privacy, social justice, reliability, and environmental sustainability. It defines observables for determining the extent to which such values are achieved, and assesses risks that an AI algorithm might introduce. The framework is intended to be flexible and to address the concerns of manufacturers of AI systems, consumers, regulators, and others. The AIEIG was initiated by the VDE with the support of the Bertelsmann Stifftung. Dr. Andreas Kaminski, leader of HLRS's Department of Philosophy was a core member of the group. *(CW)*

Hawk Supercomputer Takes Flight

HLRS welcomed the arrival of its new high-performance computing system, which will offer new opportunities for advanced computational research.

On February 19, 2020, HLRS celebrated the beginning of operation of its next-generation supercomputer, called Hawk. For the University of Stuttgart, the state of Baden-Württemberg, and Germany as a whole, the event marked the beginning of a new era for advanced academic and industrial research in the computational sciences, simulation, and artificial intelligence (AI), particularly with respect to their applications for engineering.

A Hewlett Packard Enterprise (HPE) Apollo system with a theoretical peak performance of 26 petaflops, Hawk replaced HLRS's previous flagship system, Hazel Hen, offering a 3.5-fold increase in speed. In test results announced in November, Hawk achieved a maximum LINPACK performance of 19.3 Petaflops, debuting at number 16 on the overall Top500 List of the world's fastest supercomputers.

Half of the funding for Hawk was provided by the Baden-Württemberg Ministry for Science, Research, and Art in conjunction with its high-performance computing/data intensive computing (HPC/DIC) strategy. The second half was provided by the Federal Ministry for Education and Research (BMBF) through the Gauss Centre for Supercomputing (GCS). Hawk is part of the GCS national supercomputing infrastructure.

Welcoming Hawk's arrival at the February inauguration ceremony were high-ranking representatives of the federal and state governments, including State of Baden-Württemberg Prime Minister Winfried Kretschmann, Parliamentary State Secretary of the German Federal Ministry for Education and Research Dr. Michael Meister, and Baden-Württemberg Minister for Science, Research and Art Theresia Bauer.

Also participating in the ceremony were University of Stuttgart Rector Dr. Wolfram Ressel, HPE Chief Sales Officer Heiko Meyer, AMD Corporate Vice-President of EMEA Sales Mario Silveira, SSC-Services Director Matthias Stroezel, HLRS Director Prof. Michael Resch, and Edinburgh Parallel Computing Centre Director Prof. Mark Parsons.

New tool for scientific and engineering excellence

Speakers at the event celebrated the new capabilities for academic and industrial research that Hawk will offer. "HLRS represents an essential component of the University of Stuttgart's comprehensive strategy with respect to our newly founded Excellence Cluster, Data-Integrated Simulation Sciences," explained Dr. Ressel in his opening remarks. "At the same time, however, it also offers scientists from around the world and experts from industry the opportunity to work together in an interdisciplinary way to use simulation to find solutions for complex problems."

In an accompanying press release, Secretary Meister explained, "Computers like Hawk are tools for advanced research in the sciences and in industry. They enable excellent science and innovation, and solidify Germany's international position as a top location for supercomputing. High-performance computers also contribute to technological sovereignty in the digital age; that is, the development of expertise in systems and applications that make science and industry competitive at the highest level. It is only in this way that we can guarantee a mutually reciprocal exchange with leading research centers and companies worldwide. Stuttgart's Hawk demonstrates this very clearly."

Prime Minister Kretschmann added, "With Hawk we can support key industries in the state of Baden-Württemberg, enabling fields like mobility, mechanical engineering, and health to take advantage of new opportunities using simulation. Having one of the fastest computers located at a public institution worldwide that is also open for industrial usage is a smart investment in our future as a location for science and industry. Additionally, it offers great promise for climate simulation. HLRS's work is not only of enormous importance for Baden-Württemberg: Hawk also strengthens Germany's national supercomputing infrastructure, as well as that of Europe."

Minister Bauer expanded on this, saying, "As the peak performance of supercomputers rises, it has an equally decisive impact for cutting-edge research and for the development of innovative products and processes in key sectors of our economy. In all things related to

Hawk debuted in 16th place on the Top500 List of the world's fastest supercomputers and will soon gain new capabilities for AI and machine learning.

At Hawk's inauguration ceremony, guests visited the new supercomputer.

supercomputing, Baden-Württemberg is a European leader and globally competitive. Having HLRS as a part of the University of Stuttgart is an important part of this. And it is not just the impressive computational performance of the supercomputer, but also the entire methodological expertise located here, that enables our leading computational research to achieve breathtaking results."

The beginning of Hawk's operation also marked the start of a new collaboration between HLRS and HPE to develop new technologies for supercomputing. This includes creating new software and tools for high-performance computing, for performance optimization, and for artificial intelligence that will be necessary to prepare for the arrival of new exascale supercomputers. "Hewlett Packard Enterprise and HLRS are a Dream Team," said Heiko Meyer at Hawk's inauguration. "For this reason we would like to develop a long-term development partnership with HLRS in which we optimize applications, test future technologies, and bring them to a mature, market-ready state."

Hawk is based on the EPYC Rome 7742 processor from Advanced Micro Devices (AMD). Mario Silveira, Corporate Vice President for Sales for Europe, the Middle East, and Africa at AMD, spoke about the opportunities his company gains through the partnership with HPE and HLRS in improving efficiency and performance, and making sustainability gains in its products. "We are proud to be part of this shared vision of using supercomputers to solve global challenges, and to be a partner in Hawk," he said.

Hawk expansion offers enhanced capabilities for artificial intelligence

When it was initially installed in February, Hawk was based exclusively on a type of processor called central processing units (CPUs). This is because CPUs offer the best architecture for many codes used in fields such as computational fluid dynamics, molecular dynamics, climate modeling, and other research areas in which HLRS's users are most active. In recent years, however, investigators have begun exploring how deep learning, high-performance data analytics, and AI could accelerate and simplify such research. These approaches benefit from a different type of accelerator called a graphic processing unit (GPU).

In December HLRS announced an important change to Hawk's architecture that reflects this trend. Through an agreement signed with HPE, HLRS decided to expand its world-class CPU-based computing system by adding 24 HPE Apollo 6500 Gen10 Plus systems with 192 NVIDIA A100 GPUs based on the NVIDIA Ampere architecture.

The addition of 120 petaflops of AI performance will dramatically expand HLRS's ability to support applications of deep learning, high-performance data analytics, and artificial intelligence, and will enable new kinds of hybrid computing workflows that combine traditional simulation methods with Big Data approaches.

The expansion offers a new AI platform with three times the number of NVIDIA processors found in HLRS's Cray CS-Storm system, its go-to system for AI applications. It will therefore enable larger-scale deep learning projects and expand the amount of computing power for AI that is available to HLRS's user community.

"At HLRS our mission has always been to provide systems that address the most important needs of our key user community," explained HLRS Director Michael Resch. "For many years this meant basing our flagship systems on CPUs. Recently, however, we have seen growing interest in deep learning and artificial intelligence. Adding GPUs to Hawk's architecture will thus improve our ability to support scientists in academia and industry who are working at the forefront of computational research."

As 2020 concludes, HLRS and its scientific and industrial user communities look forward to the discoveries and innovations that Hawk will enable in the coming years. *(CW)*

HLRS Begins Building European Network of HPC Expertise

In three new projects, HLRS is coordinating a landmark effort to enhance HPC competencies, promote international cooperation, and improve access to supercomputing in industry across the European Union.

The launch of the EuroHPC Joint Undertaking in 2018 marked the start of a new, Europe-wide approach to high-performance computing (HPC). By facilitating closer coordination at the level of infrastructure and resource allocation, technology development, and the creation of advanced software, EuroHPC set out to build a sustainable and globally competitive European HPC ecosystem.

As EuroHPC developed, however, it became clear that one essential component of a comprehensive European HPC strategy was missing: namely, a consistently high level of expertise across Europe in high-performance computing and related disciplines such as high-performance data analytics and artificial intelligence. In 2020 EuroHPC, under the auspices of the European Union's Horizon 2020 fund, awarded three grants aimed at addressing this need.

HLRS and its umbrella organization, the Gauss Centre for Supercomputing, will be at the center of this effort, leading the creation of a Europe-wide network focused on expanding HPC expertise. The three projects — EuroCC, CASTIEL, and FF4EuroHPC — began operation on September 1, 2020 and in the coming years will focus on different dimensions of this undertaking in a coordinated way.

According to HLRS General Manager Dr. Bastian Koller, who will be spearheading the project, "HLRS is thrilled to be leading the coordination of a project that holds enormous potential to elevate not only expertise but also cooperation and collaboration across the entire European HPC community. The effort will help to raise not only Europe's scientific capabilities but also its global industrial competitiveness."

EuroCC: establishing a European network of HPC competencies

At the center of this multifaceted effort is EuroCC, which has already begun building a pan-European network of national HPC competence centers.

As the project took shape, 33 participating member states each designated one HPC center as a national competence center. Approximately half of the funding to support the creation of each competence center is being provided by the EU, while the other half will come from the individual country.

Each competence center will undertake an audit of HPC capabilities at the national level, identifying available expertise and knowledge gaps across the country. Over time, the competence centers will become national resources for identifying and coordinating technical knowledge, training resources, industrial outreach, and HPC services and tools in their home countries.

CASTIEL: promoting collaboration across borders

Simultaneously, HLRS is leading a closely related coordination and support activity called CASTIEL (Coordination and Support for National Competence Centres on a European Level). While EuroCC coordinates HPC expertise within each member state, CASTIEL will promote interaction and the exchange of expertise across the EuroCC network.

CASTIEL will compile a Europe-wide competency map that will catalog resources and knowledge gaps across all EuroCC competence centers. It will then coordinate activities such as international workshops, mentoring and twinning partnerships, and topic-specific working groups to address issues of shared interest. Such activities should promote Europe-wide knowledge sharing and new kinds of productive synergies, potentially including new regional collaborations.

CASTIEL will also focus on providing expertise needed to address HPC needs for industrial R&D. This could include mentoring by centers with experience in supporting industry, as well as international meetings in which industrial HPC users present case studies of successful applications.

FF4EuroHPC: improving access to HPC for SMEs

The third pillar of this pan-European effort will involve focusing on the needs of the small and medium-sized enterprises (SMEs) that form an essential component of Europe's economy. A consortium called FF4EuroHPC, also coordinated by HLRS, will undertake targeted outreach to such companies, smoothing access to HPC technologies and solutions that will make them more competitive.

The team leading FF4EuroHPC includes several centers that were also involved in the EU-funded Fortissimo and Fortissimo 2 projects, which between 2013 and 2020 organized 92 successful business experiments involving SMEs and high-performance computing. FF4EuroHPC will extend the lessons learned during Fortissimo to identify and test new applications of HPC in industry across Europe.

FF4EuroHPC will organize two open calls for proposals from SMEs that can demonstrate compelling cases for using HPC technologies and expertise. Potential application areas include manufacturing, engineering, and other innovative growth sectors. Proposals may suggest traditional applications of high-performance computing for simulation and modeling or new kinds of applications involving data analytics, machine learning, and artificial intelligence.

FF4EuroHPC expects that approximately 40 SMEs — all first-time users of HPC — will undertake "application experiments," many of which will be paired with or receive support from national competence centers in the EuroCC network. This approach should promote industrial HPC usage among SMEs in regions where it has previously been slow to develop.

HLRS at center of European HPC strategy

As these projects proceed, such experiments will lead to the development of a portfolio of success stories demonstrating the innovation potential of HPC. Koller anticipates that this will help not only to attract interest from other SMEs that could benefit from HPC, but also to gain support from governments and other stakeholders for the development and maintenance of robust HPC resources and services.

"EuroCC, CASTIEL, and FF4EuroHPC have complementary goals," Koller says. "On the one hand, we hope that we will empower the participating countries to develop the expertise they need to become more technologically independent. On the other, countries that have already developed a broad base of expertise will also learn from these international interactions. By raising the knowledge level and improving communication across borders, we see a lot of potential to raise the productivity and impact of HPC all across Europe for academic research, industry, and public administration." *(CW)*

EURO

CASTIEL

ff4 EuroHPC

Building a National Data Research Infrastructure for Catalysis Research

HLRS will help create a basis for interdisciplinary computational research in chemistry and chemical engineering.

Catalysis research, like many scientific fields, increasingly relies on computational methods that support a continual dialogue between theory, simulation, and experiment. To promote the development of this field, a consortium including HLRS was awarded a grant of more than €10 million from the Deutsche Forschungsgemeinschaft (DFG) to create a National Research Data Infrastructure for Catalysis-Related Sciences (NFDI4Cat). The consortium, led by the nonprofit chemical society DECHEMA (Gesellschaft für Chemische Technik und Biotechnologie e.V.) and involving representatives from 15 additional partner institutions, will develop infrastructure, software, and data management standards to empower the next generation of chemical engineering research. NDFI4Cat is one of nine new consortia that will contribute to the construction of a German National Research Data Infrastructure.

As one of four members of the NFDI4Cat coordination group, HLRS will create and host a data repository for catalysis-related research, including a portal for sharing and accessing data stored at multiple locations. In addition, HLRS will play a major role in an effort to establish standardized metadata and ontologies for catalysis research that will ensure compatibility among different data sets, increasing their usability and amplifying their potential impact for scientific progress.

"We are very pleased that HLRS will be participating in the development of a National Research Data Infrastructure," said HLRS Director Prof. Michael Resch. "Working together with partners in the catalysis research community, this project should offer outstanding opportunities to accelerate research in a field that is not only of great economic importance, but that also holds keys to addressing some of our greatest global challenges." *(CW)*

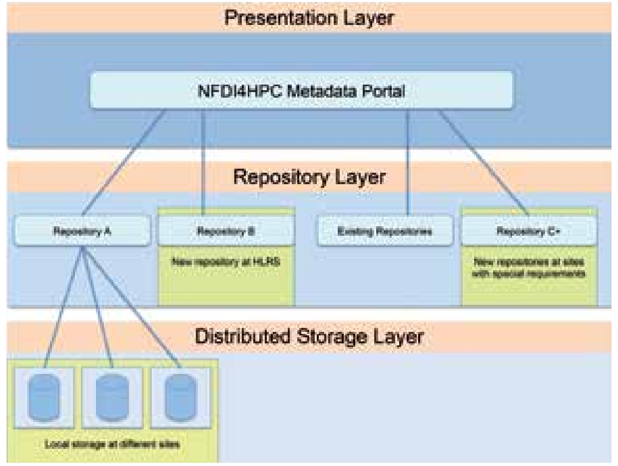

A schematic representation of the proposed NFDI4Cat data infrastructure.

Research Program Will Focus on Trust and Information

The HLRS Department of Philosophy will help develop perspectives for assessing the trustworthiness of computational science and limiting the spread of misinformation.

At the same time when the computational sciences are becoming more important than ever, concerns about how personal data and artificial intelligence can be used and misused, or about the corrosive effects of social media disinformation campaigns, have led to a growing anxiety about how to know that digital information can be trusted. How can scientists ensure, for example, that the models they develop are reliable and form a basis for public decision making? And how can people who access digital information be in a better position to distinguish between trustworthy information and misleading propaganda?

A new three-year project at HLRS aims to address such questions. With the support of a grant from the Baden-Württemberg Ministry of Science, Research and Art, a team led by Dr. Andreas Kaminski of the Philosophy of Science & Technology of Computer Simulation Department will bring together philosophers, social scientists, technologists, and other experts to investigate trust in the context of information technology. The project will produce insights for improving trustworthiness in computational research, developing AI-based approaches for judging reliability of information, and fighting deception in digital media.

The new project assumes that improving the trustworthiness of digital information is not just a technical challenge, but involves questions of how humans perceive information, as well as how trust is built between individuals and in communities. Kaminski and his team members will thus incorporate multidisciplinary perspectives, bringing together experts in psychology, sociology, political science, economics, pedagogy, and history. Through collaborative research, they will develop a theoretical basis for improving trustworthiness in simulation and AI technologies. *(CW)*

One of these faces is real and the other was generated by artificial intelligence. Can you tell which is which? Applications such as these highlight problems of trustworthiness that can arise with digital information.

Collaborative Virtual Reality Environments for the Home Office

New software developed by HLRS and the KoLab project enables interaction and cooperation in virtual 3D spaces.

Simulations using supercomputers or construction data from CAD systems are normally visualized and analyzed in a CAVE 3D facility. In an era of "social distancing" resulting from the coronavirus pandemic, however, this is not always possible.
To address this problem, members of the HLRS Visualization Department and researchers at five additional universities in Baden-Württemberg — working together in a project titled KoLab (Virtual Collaboration Laboratories Baden-Württemberg) — developed a new software platform that makes it possible for persons located in different places to meet and collaborate in virtual reality (VR).
Using new VR-software developed by KoLab, combined with inexpensive, commercially available VR-headsets and controllers, teams of scientists and developers can gather in a virtual meeting room. Meeting participants are represented by avatars that communicate in virtual reality, making it possible to collaboratively observe, analyze, and interact with 3D visualizations from their individual workplaces or home offices.
In situations where spontaneous in-person meetings are difficult — for example, when colleagues work at different locations or challenging conditions make it difficult to meet in person — this technology will empower collaboration and simplify complicated work processes. Even after the COVID-19 pandemic is past, it promises to accelerate scientific discovery and shorten the time it takes to bring new products to market.
In fall 2020, members of HLRS's Visualization Department used the technology to conduct a course at the Technical University of Vienna. Students in Vienna could interact in virtual reality with avatars representing the instructors, who remained in Stuttgart. *(CW)*

© HLRS/Fabian Dembski

Collaborative VR visualization of a multi-generational community center.

HLRS and Fraunhofer IPA Begin Cooperation

The partnership will explore new applications of high-performance computing, high-performance data analytics, machine learning, and quantum computing in the manufacturing sector.

© Fraunhofer

At the signing ceremony in November 2020.

HLRS and the Fraunhofer Institute for Manufacturing Engineering and Automation (Fraunhofer IPA) have long been neighbors on Stuttgart's Nobelstraße. Under a new cooperation agreement, the two organizations will grow even closer as Fraunhofer will gain access to HLRS's HPC systems and expertise. The partners will also work together on issues of mutual interest related to applications of advanced computing technologies in industry.
Fraunhofer IPA conducts research to address organizational and technological issues across the entire manufacturing sector. This includes new technologies to improve methods, components, machines, and production lines used in manufacturing. Considering the increasing availability of large data sets and powerful tools for simulation and artificial intelligence (AI), the use of advanced computing technologies operated by HLRS could help accelerate the development of innovations that make manufacturing more efficient and sustainable.
The new cooperation agreement will focus on applications of HPC that have become important for Fraunhofer IPA's research in recent years. This could include performing analyses of large data sets to improve manufacturing processes or using HPC to verify the accuracy of machine learning applications for human-robot collaboration. Under the agreement, HLRS and Fraunhofer also plan to explore new challenges and opportunities related to quantum computing.
As HLRS General Manager Bastian Koller explained, the agreement also opens new opportunities for HLRS: "The partnership will enable us to develop and improve solutions that integrate HPC, high-performance data analytics, and AI in ways that best support our users. It could also provide new impulses for HLRS's efforts to support the adoption of HPC and related technologies into industry." *(CW)*

How Artificial Intelligence Could Improve Robotics: An Interview with Marco Huber

The increasing availability of machine learning and artificial intelligence (AI) applications has made it possible to develop new kinds of automated manufacturing systems. As Professor for Cognitive Production Systems in the Faculty of Engineering Design, Production Engineering and Automotive Engineering at the University of Stuttgart and Head of the Center for Cyber Cognitive Intelligence at the Fraunhofer Institute for Manufacturing Engineering and Automation (Fraunhofer IPA), Professor Marco Huber is investigating innovative applications of machine learning that could make these systems more flexible.

This combination of roles places Huber in an ideal situation to facilitate productive interactions between research and industry. As a university professor, he investigates state-of-the-art topics not only from a theoretical perspective, but also with a focus on practical requirements in the manufacturing industry. At Fraunhofer IPA, he advises companies interested in using AI to improve their productivity. "The work at Fraunhofer is less about developing new methods," Huber explains, "than about applying methods developed for basic research in ways that will help industry to address the problems it faces."

In November 2020, Fraunhofer IPA and HLRS signed a cooperation agreement that will enable Huber and his colleagues to access resources for high-performance computing (HPC) and simulation at HLRS. Together, the two organizations also plan to undertake research on issues of shared interest — specifically, quantum computing and artificial intelligence. In this interview, Professor Huber describes how the partnership came into being and how HLRS's computing resources will enhance Fraunhofer IPA's research capabilities.

? Professor Huber, what new opportunities does artificial intelligence offer for robotics?

❯ In industry today, you need to have a certain level of expertise to program a robot. This is a situation we want to move beyond. Such a process should not only be accessible to experts, but everyone should have the ability to program robots when needed.

Instead of programming a robot directly, our approach involves formulating a task that it needs to complete — for example, grasping an object and removing it from a box. In a simulation, a robot can teach itself the best way to do this. At first, it does a pretty bad job, but because this is only a simulation, that's ok — it can't break anything. With the help of a paradigm from machine learning called reinforcement learning, it gradually improves until it reaches a point where we can determine that it is good enough for the application. The nice thing is that as soon as we reach this point, we also automatically have a robot program that we can implement on the factory floor. In a sense, the robot programs itself through this process.

The advantage of using simulation in robotics is that you don't need access to a physical robot. Normally, a robot must be taken out of production for the entire time it is being programmed. This isn't just a matter of a few hours, but can require a significant amount of time until everything functions correctly. By using simulation, we can avoid such interruptions in production.

At the same time that a machine manufactures a particular product, we believe that through simulation it should also be able to learn how it will need to go about manufacturing the next product. This is the vision that motivates our research. This goal isn't something we will be able to implement in the real world today or tomorrow, but we think that it should be possible in the coming years.

? What advantages could this approach offer for industry?

❯ The vision that we're pursuing is to facilitate the transformation of production to so-called "lot-size-1" manufacturing. There is growing interest in moving away from mass production towards a model that people in the industrial community around Stuttgart refer to as "mass personalization." Companies would like to be able to produce high-quality, individualized products, while at the same time keeping costs low. To make this possible, production machinery must be able to adapt itself efficiently to every new product.

You can see one example of this need in the manufacture of electrical circuit boxes, which are usually assembled in a highly customized way. Nowadays, nearly every circuit box is unique and built by hand. They contain, for example, so-called "top-hat rails" (also called DIN rails), on which components need to be mounted and cabled together. Because of how difficult and expensive this process is, we would like to use simulation to get to a point where a robot could learn how to manufacture customized circuit boxes efficiently.

? Why is high-performance computing necessary in these kinds of methods?

❯ Reinforcement learning and other kinds of machine learning rely on data that often doesn't exist in the amount or in the form that is needed. Therefore, we use highly precise simulation to generate the vast majority of the data that is needed to train the algorithm. This way, you only need a very small amount of real-world data in order to fine-tune the algorithms.

At Fraunhofer IPA, we are already doing this to a small degree, but we don't have the necessary computing capacity to carry out large-scale simulations. Right now, it isn't a problem for us to program an individual robot, but as soon as we begin thinking in larger scales — such as several robots or a production process involving multiple interrelated steps — we get to a point where our computing capacity isn't sufficient. This is why having access to the supercomputer at HLRS is a cornerstone of our partnership. The collaboration offers us new opportunities for simulating complex systems, having the machine learn through simulation, and transferring the resulting programs to industry, and to do it on a large scale.

Professor Marco Huber

© Fraunhofer IPA

? Are there other ways in which HPC could support machine learning at Fraunhofer?

❯ One other topic that we would like to address as part of this collaboration is the verification of neural networks. Today, researchers often face the problem that when they want to use neural networks in applications where safety is critical, it is difficult to guarantee that the network will do what it is supposed to. This is where verification comes into the picture. You formulate a set of requirements that the network should satisfy and then mathematically test how well the algorithm meets them. We have been working with a company to investigate this in a case study and in our test run we showed that it should be possible.

Nevertheless, this verification step is extremely computationally expensive, and at Fraunhofer IPA we unfortunately don't have the computing capacity to solve this problem in an acceptable amount of time. In our conversations with the company and with HLRS, we agreed

Through simulation robots can learn to improve their ability to manage tasks such as grasping an item in a box.

© Fraunhofer IPA

that we will undertake a demonstration project together in the first quarter of 2021. HLRS will make the necessary computing resources available to complete this verification exercise.

Fraunhofer IPA does have a GPU computing cluster available for research using artificial intelligence, but because researchers from many different parts of the organization are now using it, the demands are continually growing and we often run into bottlenecks. Having access to HLRS's computing resources should also help to address this problem.

? The collaboration agreement with HLRS also includes a focus on quantum computing. What are your plans in this area?

› Fraunhofer IPA, together with HLRS and several other partners, was successful in a proposal to start a new project called SEQUOIA (Software Engineering for Industrial Hybrid Quantum Applications and Algorithms), which will conduct research to understand the future opportunities that quantum computing offers. This kind of computing is still quite new, and so we need to be asking some important questions: What kinds of practical problems could quantum computing solve faster than a classical computer can? And which kinds of problems will, for scientific reasons, not benefit from so-called quantum supremacy?

Through this project, we will gain access to the quantum computer that IBM is making available to the Fraunhofer Society beginning in 2021. The system offers only a relatively small number of qubits, and so although it will not necessarily enable us to solve real-world problems, it will allow us to develop experience and undertake small-scale case studies so that we can gain a better understanding of what might be possible. Currently, quantum computing — in a similar way to the programming of robots — is still something that only experts can do. Worldwide, there are currently only a small number of experts who are able to program quantum computers. For this reason, we will also be looking at ways to simplify software development for such systems. By doing so, we hope to make quantum computing more accessible for both scientists and industry. *(Interview: CW)*

Performance Optimization Workshop Opens New Scientific Opportunities

HLRS's personalized user support benefits computational fluid dynamics research in many ways.

HLRS's performance optimization workshops offer users of the center's HPC systems the chance to work closely with support staff to improve the codes that are essential for their research. With a recent increase in funding through the SiVeGCS project, this meeting format enables HLRS to provide increasingly personalized consulting services.
Twice each year, the workshop invites approximately 10 teams of researchers that use HLRS's computing systems, pairing each with an individual HLRS expert in scientific computing. Meeting together for a full day, they dive deep into the researchers' computer code and in a compressed amount of time make important modifications that improve its performance on HLRS's high-performance computing systems. Representatives of hardware manufacturers HPE and AMD are also present to consult on detailed technical matters related to their products. Bringing this multidisciplinary expertise together at one time makes it possible to focus on codes intensively and leads to improvements that can not only get the scientists their results faster but also help them to address new kinds of scientific questions.

At a workshop held in early 2020, a team of researchers from the University of Stuttgart's Institute of Aerospace Thermodynamics (ITLR) brought a problem related to Free Surface 3D (FS3D), a code for computational fluid dynamics research. Initially developed over 20 years ago, FS3D has undergone continual improvements, though not always with a focus on performance optimization.
FS3D computationally subdivides a droplet, for example of water, into a mesh of billions of boxes, each of which contains a high-resolution simulation of the fluid in that volume. This means that the code must coordinate a massive number of parallel calculations.
When the team and HLRS's Dr. Martin Bernreuther looked at the FS3D code together, they realized that the multigrid solver they used was sending large numbers of calculations to a single processor, while other processors that could have assisted in the calculation sat idle. During the workshop they implemented a method that, in a treelike manner, coordinates communication among the solutions to the system of equations during different cycles of the multigrid solver. The results make much better use of the parallel architecture of HLRS's supercomputer.
Following performance testing, the team discovered that the improved code ran up to 20 times faster when scaled up to a large number of computing cores. Commenting about the speed-up, ITLR scientist Jonas Steigerwald remarked, "Previously we saw a dramatic drop in performance when we ran our code on larger numbers of processors. With the way the communication is structured now, we don't see that any more."
"At 64 processors, the change isn't so noticeable, but when you go to 4,000 or even higher numbers of processors, this makes an enormous difference," added his ITLR colleague Matthias Ibach.
For the scientists, the personal contact that the workshop offered made the difference. According to FS3D team member Jonathan Reutzsch, "In many cases being able to identify and fix a problem is something that can only happen during such workshops and through this kind of direct mentorship. Every year it brings us a step farther."
Ibach also credited HLRS's expertise in scientific computing for its impact on his institute's research, saying, "The ability to parallelize our computing on this scale makes it possible to move our research into new areas."

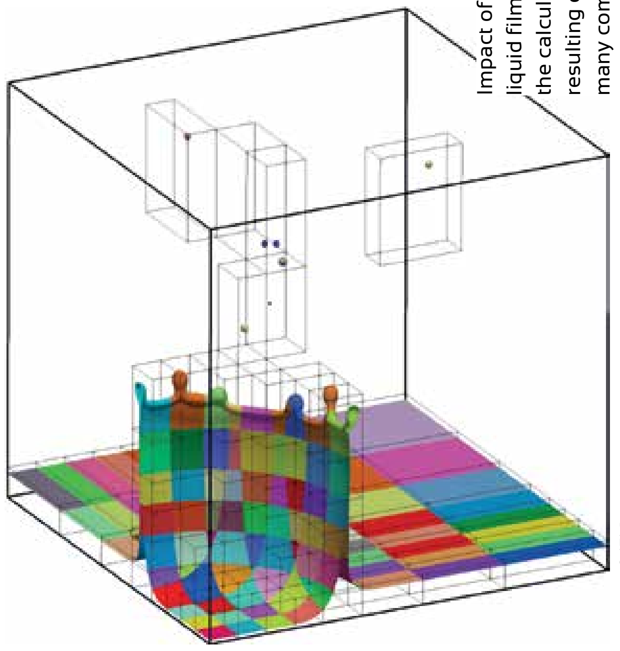

Impact of a droplet as it falls on a thin liquid film. The colored image shows how the calculations required to simulate the resulting crown are distributed among many computer processors.

© ITLR

The scientists also value the contacts the HLRS workshop facilitates with researchers from other institutes. As Reutzsch explained, "In the past we have met many people who face similar problems whom we wouldn't have met otherwise. HLRS is a unifying institution that brings together people who are facing similar challenges."
As a result of the COVID pandemic, HLRS held its second performance optimization workshop of 2020 using videoconferencing tools (See page 12). Although the scientists missed the social component of the live event, the online format was successful and also enabled HLRS to develop new approaches for providing personalized support for users outside the Stuttgart area through the Internet.
In addition to improving the research of its scientific users, HLRS's expert user support also helps to ensure more efficient operation of its computing systems and to maximize their availability for all users. *(CW)*

HLRS by the Numbers

System Usage

- 129 ACTIVE USER PROJECTS
- 51 INDUSTRIAL CUSTOMERS
- 128 USER PUBLICATIONS

Staff

145

Visiting Researchers

19

- 95 SCIENTISTS
- 34 NONSCIENTISTS
- 13 STUDENT ASSISTANTS
- 3 RESEARCH ASSISTANTS

Talks

55

Staff Publications

- 44 PAPERS IN JOURNALS, BOOKS, AND CONFERENCE PROCEEDINGS
- 2 BOOKS

HL RS

Education and Training

30 CONTINUING EDUCATION COURSES
- 1,005 Participants
- 100 Course-Days

4 SCIENTIFIC WORKSHOPS
- 171 Participants
- 7 Days

14 UNIVERSITY LECTURES
- 2,287 Participants
- 32 SWS

Visitors at HLRS

249

Third-Party Funds

6,925,747

- EU 49 %
- STATE GOV. 29 %
- FEDERAL GOV. 20 %
- DFG 2 %

HLRS Achieves Two Environmental Certifications

EMAS and Blue Angel awards recognize the center's environmental and energy management strategies.

In March, HLRS became the first major high-performance computing center to be certified under the Eco-Management and Audit Scheme (EMAS), the most demanding system for environmental management worldwide. In October it also received certification for Energy-Efficient Data Center Operation under the Blue Angel, the eco-label of the federal government of Germany.

"HLRS is home to one of Europe's most powerful supercomputers, which means high energy requirements," said HLRS Director Michael Resch. "We recognized years ago that we have to do everything we can to minimize our ecological footprint, especially with respect to carbon emissions. Receiving these certifications shows that HLRS takes its environmental responsibility very seriously and that our system users can be confident that we support their efforts toward sustainability."

HLRS's environmental management system covers all levels of the organization, addressing not just energy usage, but also resource consumption and disposal, and activities that can negatively impact the environment or human health. HLRS considers environmental impact in all purchasing, and has formalized systems for tracking and reducing energy consumption, minimizing waste, and reusing resources.

The EMAS and Blue Angel awards complement HLRS's previous certifications under the international standards ISO 50001 (energy management) and ISO 14001 (environmental management). Development of the center's environmental management plan was supported by the Baden-Württemberg Ministry for Science, Research, and Art.

To assist other computing centers interested improving their own environmental performance, HLRS also published a guidebook called "Sustainability in Computing Centers: A Practical Guide."

It is available on our website at the following link: www.hlrs.de/about-us/sustainability-in-hpc

(CW)

Heat generated by the Hawk supercomputer is used to heat the buildings, while the use of free cooling requires significantly less energy than traditional air conditioning systems

PhD Graduates 2020

Two young investigators successfully completed their doctoral studies at HLRS in 2020.

Fabian Dembski

Energy Conscious Urban Inward Development. Analytical Design Strategies for the Post-Oil City: The Case Study of Greater Paris.

Limiting individual automobile transport in cities could not only greatly reduce energy consumption, but also make urban environments more healthy and humane. Positive benefits could, for example, include reducing energy consumption and fossil fuel emissions, repurposing transportation infrastructure, regenerating urban districts, upgrading public spaces, and better integrating isolated districts or suburbs into the urban fabric. In his dissertation, Dembski uses a variety of methods, including space syntax theory, geographic information, strollology, and scenario development and testing to conduct a case study of Paris in which he demonstrates myriad ways in which method sets could help pave the way toward the post-oil city of the future. In addition, discussing his creation of the first European Digital Urban Twin, he explores how cutting-edge digital technologies can offer strategies for solving problems, attaining goals, and evaluating success in the planning of individual cities.

Joseph Schuchart

Global Task Data Dependencies in the Partitioned Global Address Space

The dominating programming model driving today's parallel computing applications is a two-level approach consisting of message-based communication between processes using MPI and static loop-level thread-parallel execution using OpenMP constructs. However, two programming models have tried to challenge this status quo. First, the PGAS model is an attempt to elevate shared memory programming to the level of distributed systems and to directly expose modern network hardware features to the application developer. Second, task-based programming aims at providing abstractions that help discover a greater amount of concurrency in parallel applications, which in turn can be used to better exploit the computational resources at hand. In his dissertation, Schuchart proposes a novel way of orchestrating the execution of tasks at a global scale by using distributed task graph discovery and data dependencies in the global memory space. The results demonstrate that applications exhibiting concurrency beyond single loop parallelism may use this new model to significantly improve performance and scalability by combining the benefits of task based programming and one-sided communication in the PGAS model.

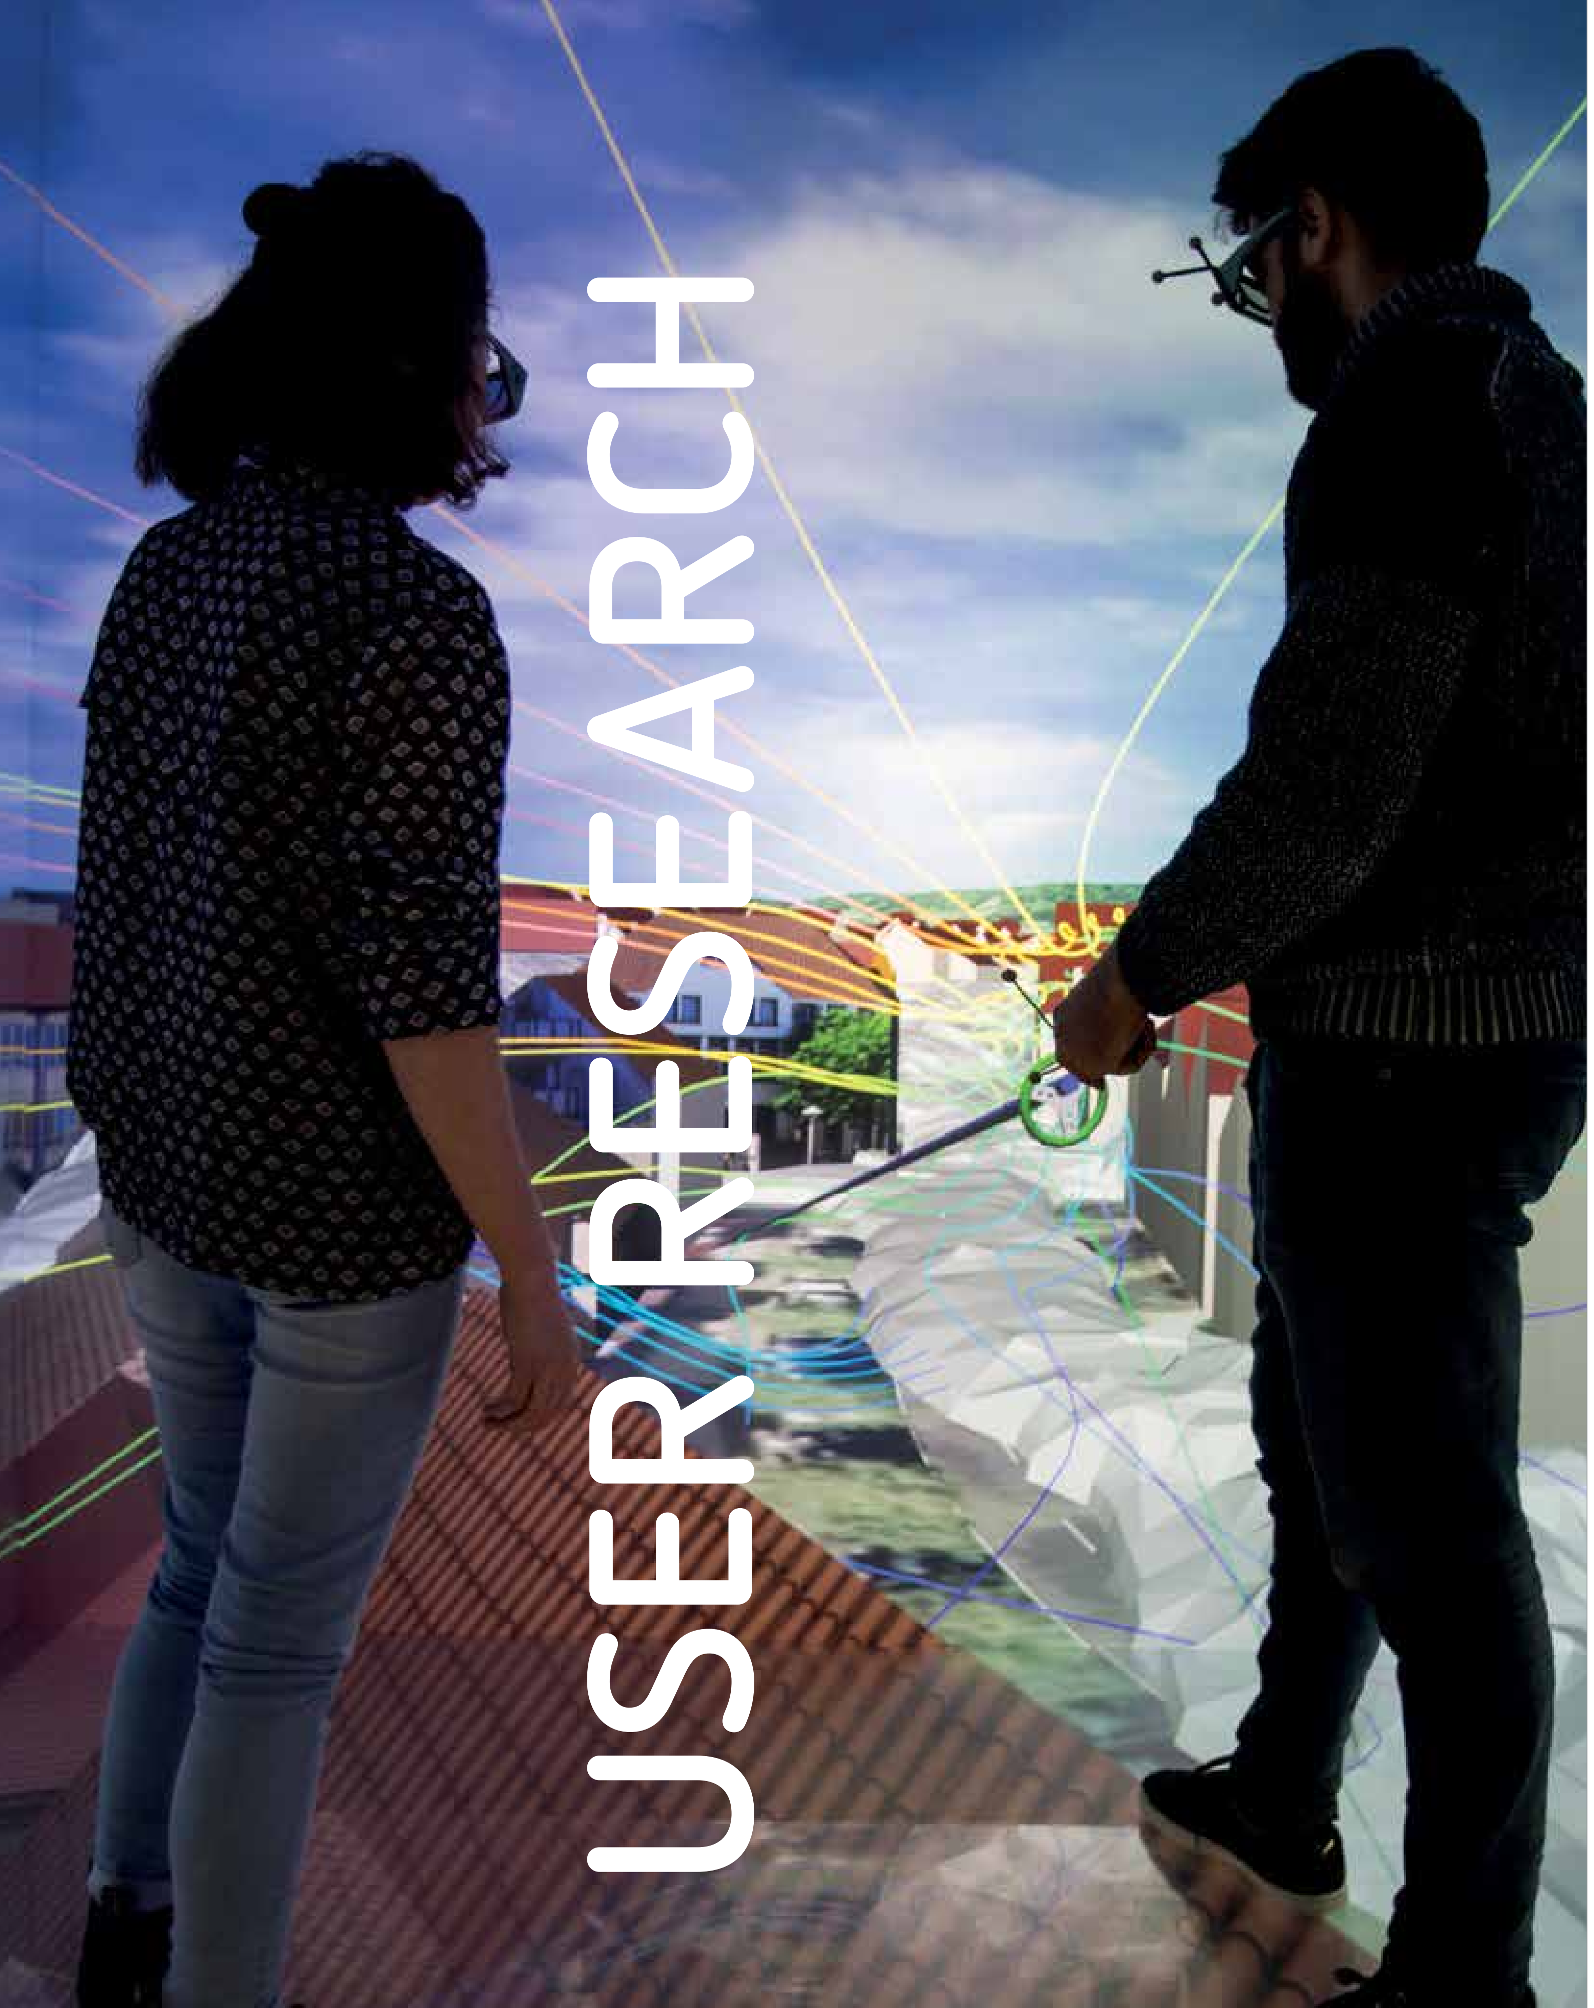

Machine Learning Could Improve Fluid Dynamics Research

Improving the accuracy of large-eddy simulations could lead to more accessible methods for boosting wind turbine efficiency or making aircraft more environmentally friendly.

For decades, engineers and scientists have used increasingly powerful high-performance computing (HPC) resources to better understand the turbulent interactions that influence fluid flows. With applications ranging from making industrial chemical processes safer to improving fuel efficiency and lowering noise pollution in commercial air travel, understanding how fluids interact at a fundamental level has improved countless processes that impact our daily lives.

Fluid dynamics research is well-suited to modeling and simulation. Many HPC simulations begin with a computational grid or mesh that allows researchers to break a large, complex problem into small equations that can be calculated in parallel and then quickly pieced back together. As supercomputers have gotten more powerful, researchers have been able to make these meshes increasingly finer.

Now that supercomputers like HLRS's Hawk system are capable of quadrillions of calculations per second, scientists have gained the ability to create direct numerical simulations (DNS) that model a fluid's movements "from scratch" without the need for simplifying models or assumptions. When computing turbulence in fluid flows, DNS allows researchers to accurately simulate the tiniest swirling motions, or eddies, that influence large-scale movements in a fluid.

Doing this work, however, requires regular access to large supercomputers. When this isn't possible — for example when researchers from industry want to get a rougher view of turbulence — they often run large-eddy simulations (LES), which calculate only the behavior of the larger swirling motions happening in a fluid. This approach also requires making assumptions based on empirical data or other models about the behavior of the smallest eddies. While this approach requires less computational power, researchers know that even the smallest eddies have a large impact on how the whole system behaves.

Recently, researchers led by Dr. Andrea Beck at the University of Stuttgart's Institute for Aerodynamics and Gas Dynamics (IAG) have been exploring how machine learning could make the power of direct numerical simulations accessible to researchers with more modest computational resources. A long-time user of HLRS's HPC resources for traditional modelling and simulation work, she has turned to using Hawk to generate training data for neural networks to improve LES models. Her experiments have already begun showing potential to accelerate fluid dynamics research.

Searching for closure

The difference between modelling turbulence using DNS and LES is similar to photographs. One can think of DNS as a high-resolution photograph, whereas LES would be a rougher, more pixelated version of that same image. However, unlike a photograph, researchers do not always have access to the high-resolution version of a given system from which to develop an accurate representation in lower resolution. To fill in this gap in fluid modeling, researchers have to look for what they call "closure terms."

"To use a photograph analogy, the closure term is the expression for what is missing between the coarse grained image and the full image," Beck said. "It is a term you are trying to replace, in a sense. A closure tells you how this information from the full image influences the coarse one." Just like a processed image, though, there are many ways in which LES turbulent models can diverge from a DNS, and selecting the correct closure term is essential for the accuracy of LES.

The team started training an artificial neural network, a machine learning algorithm that functions similarly to how the human brain processes and synthesizes new information. These methods involve training algorithms with a collection of large datasets, often generated by HPC systems, to forge connections similar to those formed in the human brain during language acquisition or object recognition. These types of algorithms are trained in a "supervised" environment, meaning that researchers feed in large amounts of data where they already have the solution, ensuring that the algorithm is always given the correct answer or classification for a given set of data.

"For supervised learning, it is like giving the algorithm 1,000 pictures of cats and 1,000 pictures of dogs," Beck said. "Eventually, when the algorithm has seen enough examples of each, it can see a new picture of a cat or a dog and be able to tell the difference."

For its work, the team did a series of DNS calculations on HLRS HPC resources. They did roughly 40 runs on the Hawk and Hazel Hen supercomputers, using about 20,000 cores per run, and then used this data to help train its neural network. Because this dataset could be useful for other researchers, the team is making it publicly available in 2021. The team compared two different training approaches, and found that one approach achieved 99% accuracy in correctly pairing the correct closure term with a given LES filter function.

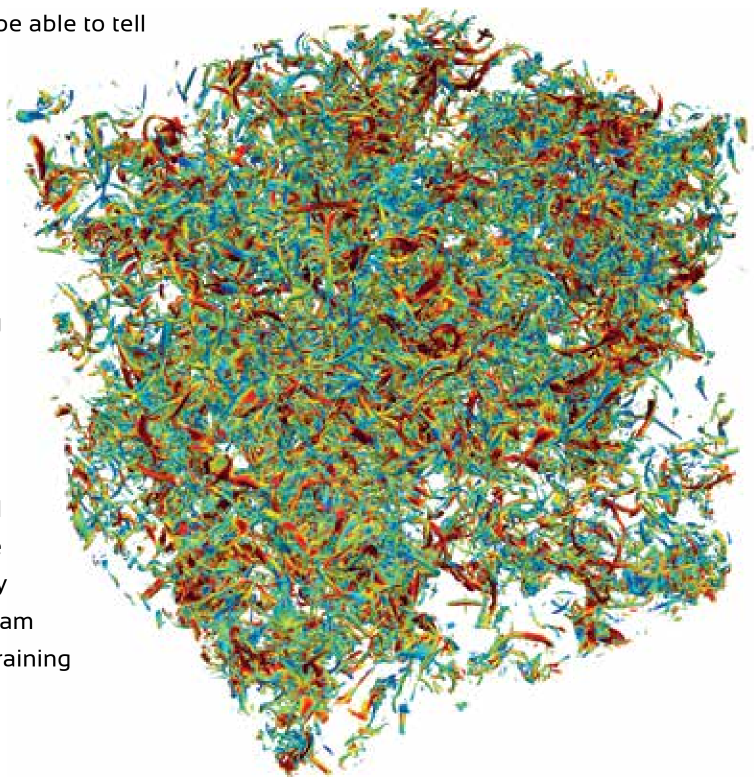

© Marius Kurz, University of Stuttgart

Vortical structures of decaying homogeneous turbulence from DNS. Simulations such as these provided the data used by Beck and her collaborators for training their machine learning application.

Learning more all the time

With such promising results, the team feels confident that these data-driven modeling approaches can be further refined to help bring the accuracy of HPC-driven DNS results to more modest simulation approaches. The team's work serves as an example of an emerging trend in science and engineering; merging and complementing traditional modeling and simulation done in HPC with artificial intelligence applications.

The addition of 192 NVIDIA graphic processing units to Hawk will support this convergence, benefiting users who need to generate large datasets to train neural networks by eliminating delays that would otherwise result from moving data between different computing systems (See page 22).

"This is a great step forward in helping us to augment traditional HPC codes with new, data-driven methods and fills a definite gap," Beck said. "It will not only help speed up our development and research processes, but provide us with the opportunity to deploy them at scale on Hawk." *(EG)*

Taking Quantum Dots for a Spin

A team from TU Dortmund University is using high-performance computing to model how lasers could regulate spin dynamics in quantum dots. These small structures could have big implications for improving quantum computers and other advanced electronics.

Modern computing technologies enable researchers across many scientific domains to simulate phenomena that are too large, small, dangerous, or difficult to observe through experiment. In fact, supercomputing initiated a new golden era of particle physics research, playing an indispensable role in helping illuminate interactions at the atomic level and below.

Since the 1920s, researchers have documented the laws that govern the atomic and subatomic world in painstaking detail. The quantum world, as it became known, operates according to a different set of laws than those first posited by Isaac Newton. As modern-day researchers have gained a deeper understanding of quantum mechanics, they have identified opportunities to manipulate and control subatomic particles in ways that could lead to new kinds of electronics or other technologies. Among these promising advances, quantum computing – which, through the application of quantum mechanics, processes information in a fundamentally different way than traditional computers – promises to accelerate certain kinds of research and the development of new applications.

While traditional computers rely on extensive patterns of 1s and 0s to transmit information in bits, quantum computers utilize so-called qubits, particles operating in accordance with the laws of quantum mechanics. Each qubit can be facing up, down, or in its "superposition," meaning it simultaneously represents both positions. Although the existence of this kind of state – where electrons can represent multiple positions simultaneously – could benefit data-intensive modelling and simulation, further understanding is needed of these counterintuitive phenomena. Scientists must also develop methods for reliably manipulating and controlling quantum particles, specifically with regard to how individual electrons "spin," or orient themselves under specific conditions.

In order to better understand how subatomic particles behave and interact at a fundamental level, a multidisciplinary research collaboration based at the TU Dortmund University is using high-performance computing (HPC) resources at HLRS to simulate some of these complex interactions. The team is partnered with experimentalists in the international collaborative research center 160 (ICRC 160) established at St. Petersburg and Dortmund to study how electrons' spins (and the spins of larger nuclei nearby in a given system) interact under certain conditions, and how laser technologies could help manipulate these systems.

In its most recent work, the team has focused on quantum dots. Developed within the context of semiconductor technologies, researchers found that quantum dots could control electrons' spins under the right conditions, making them good candidates to serve as qubits in future quantum computers.

"A quantum dot can be seen as a trap for a single electron and therefore for its spin," said Prof. Dr. Götz Uhrig, Professor at TU Dortmund and lead researcher on the project. "If we are looking at a solid state device, there are as many as 10^{20} electrons and their spins behave in

Prof. Dr. Götz Uhrig

such a way that no net effect can be seen for the outside; the spin of an excess electron, however, can be detected and manipulated."

Turning toward a better understanding of spins in quantum dots

When physicists study atomic systems, one approach involves strategically adding or removing an electron to determine the properties that it imparts on the system. As they learned more about how to modify these systems, scientists and engineers came up with relatively inexpensive ways to create silicon-based semiconductors that today are nearly ubiquitous in consumer electronics.

Now, as scientists turn their attention to developing new electronics such as quantum computers, they have employed specialized semiconductor nanostructures – so-called quantum dots – in order to better direct electrons and their spins, as is needed to follow certain algorithms.

Within a quantum dot, researchers can use lasers or other technologies to fix the positions of excess electrons in space, making it easier to manipulate their spins. Experimentally, researchers must account for the strange, novel physics that govern quantum spins of electrons, in their "up" and "down" states, as mentioned previously.

Accounting for superpositions ultimately means that researchers must calculate millions of possible states of electron spins for each quantum dot. And while simulating the spin of an individual quantum dot as it interacts with roughly thousands of nuclear spins may not be too computationally demanding, a meaningful simulation or experiment will have to deal with thousands of quantum dots at once.

One common experimental approach is to use laser pulses to "train" quantum dots so that their spins act synchronously. "My colleagues send pulses to the system that orient a quantum dot's spin," Uhrig said. "When they do that for quite a long time, they get some response before the next pulse. So you see some polarization, then it dies out, but when you have done that for a long time, there is already a signal before the next pulse comes – that shows that the system has been trained successfully."

In experiments, researchers will shoot a laser pulse every few nanoseconds and training the dots can take seconds or even minutes. This means that a realistic simulation must cover an extremely wide range of time scales. Performing the necessary calculations in a reasonable amount of time would not be possible without large-scale supercomputing resources.

Advancement through iteration

Uhrig and his collaborators use HPC primarily to help understand the data seen in experiments quantitatively. This approach allows the team not only to efficiently verify how well a given quantum system has been trained, but also to make predictions using the physical model they develop. If the team has designed the model correctly, new phenomena shown in models should then also be found in experiment. "This back-and-forth is a general feature in the dynamics between theory and experiment," Uhrig said. "When theory tries to be close to experiment, many parameters need to be taken into account, meaning you have to do heavy numerics, and that's where HPC comes in."

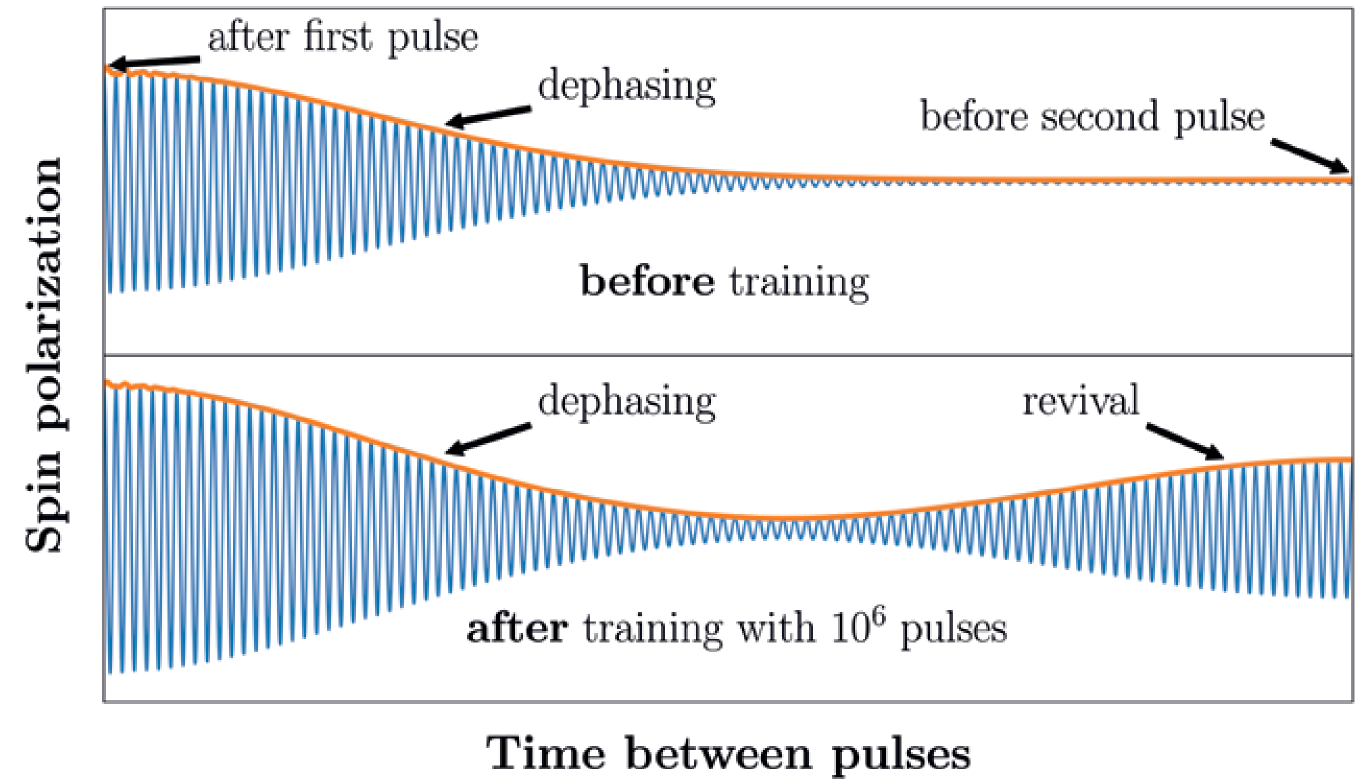

This image shows the evolution of electronic spin polarization in time. Each pulse generates a finite signal that strongly oscillates as the spin precesses around the applied magnetic field. Due to the interaction with the disordered nuclear spins, the signal quickly dies out (upper panel).
After long series of pulses, the nuclear spins are trained and a coherent revival of the signal occurs.

© Philipp Schering, TU Dortmund

Using this iterative approach, the team uncovered strategies for amplifying signals in the data that indicate how electron spins respond to laser pulses. Ultimately, this knowledge could improve researchers' ability to predict and rely on quantum-scale behavior when executing tasks in quantum computers and other advanced electronics.

The team was able to access large core counts to run larger calculations on HLRS's Hawk supercomputer during its acceptance phase last year. "On short notice, we were offered short queues for large core counts, and that made a big difference on how quickly we were able to do our research," Uhrig said. Now that the team has been able to further optimize its code on Hawk, the researchers feel confident that its future allocations will enable them to take on more complex calculations for larger systems. *(EG)*

Modeling Muscles with the Help of HPC

An interdisciplinary group of University of Stuttgart researchers is developing increasingly detailed models of parts of the human body. Starting from a cellular level, the team is scaling simulations up to model individual muscles in the human arm.

In recent decades, the frontier of understanding the human body lies in the cellular and sub-cellular world. The human body is made up of billions of cells, and while physiologists have gained an understanding of many processes controlled by muscle and nerve cells – from piano playing to weightlifting, for instance – the fundamental processes that govern these motions have not been adequately explained.

Acknowledging the importance of such questions, the German Research Foundation (DFG) in 2020 initiated a priority program supporting research that "couples" simulations and clinical studies of multiple anatomical systems. The outcomes could have major implications for personalized medicine and the development of more effective diagnostic methods and treatments.

A research group jointly led by University of Stuttgart Professors Dominik Göddeke, Miriam Mehl, and Oliver Röhrle has in recent years focused on the computational modelling of the human musculoskeletal system. For Göddeke, the complex and integrated nature of the research offers a lot of promise but also presents profound challenges. "Many scientific fields are getting so specialized that no one can keep up with all of the mathematics, computing, and biomechanical models that are needed to solve these problems," Göddeke said. As a consequence, working in interdisciplinary teams is essential.

The team has been using high-performance computing (HPC) resources at the High-Performance Computing Center Stuttgart (HLRS) to create high-resolution simulations of how our muscles, bones, and nervous system interact at a fundamental level. While experimental data still guide development of most medical treatments, the Uni Stuttgart team believes that computation could provide a better way forward.

"Understanding the musculoskeletal system in detail is unfortunately a field characterized by poor access to important information, and making measurements can corrupt datasets, lead to bodily harm, or be too imprecise," said Aaron Krämer, a graduate researcher at the University of Stuttgart involved in the project. "An alternative way of gaining insight in this field, which is becoming common, is to simulate the process of interest. In our case, we are investigating the full activation process from the nervous system to muscle contraction."

Flexing computational muscle

While researchers can create rough simulations of the human body based on generic input data from experiments or other simulations, creating a "first-principles" model quickly becomes computationally expensive. So expensive, in fact, that researchers have generally chosen to start with one small part of the body – think of a group of cells or tissue – and then expand simulations when more computational power is available.

Such simulations are computationally demanding due to the large range of scales that researchers must consider. Recently, for example, the Stuttgart researchers have focused on simulating a human biceps muscle in motion. Simulating only one muscle of many in the arm requires accurately calculating the motions for all of the "fascicles" that make up the muscle. Each muscle consists of anywhere from 10 to 100 fascicles, each of which contains between 10,000 and 250,000 muscle fibers. Ultimately, this requires more than 5 billion equations to be solved.

"Our model is such a multi-scale and multi-physics problem that in order to resolve all the processes – from the sub-cellular level to what we see when the human body is in motion–supercomputing is essential," Mehl said. "We really want to use these highly detailed models, because phenomenological models based on smaller amounts of input data do not give us the same degree of insight or the ability to generalize the observations we are able to see in our models."

In order to do these simulations in a timely manner, the team uses 7,000 cores on HLRS's Hawk supercomputer, which allows the team to model all 180,000 biceps brachii muscle fibers in an individual simulation.

Supercomputing supports modelling efforts on the path toward personalized medicine

Combining computer modelling and experiments in an iterative way is helping scientists to efficiently develop biomechanical models that could enable more individualized approaches to improving health. Together with the team and as part of a DFG-funded Cluster of Excellence at the University of Stuttgart called SimTech, Professor Oliver Röhrle leads several efforts to integrate datasets from experiments to improve biomechanical modelling. Using delicate sensor

© Benjamin Maier, University of Stuttgart

The biceps muscle is used to flex the elbow. It is attached to the ulna (right) and humerus bones (bottom left).

Simulating muscle contraction and the direction of internal forces.

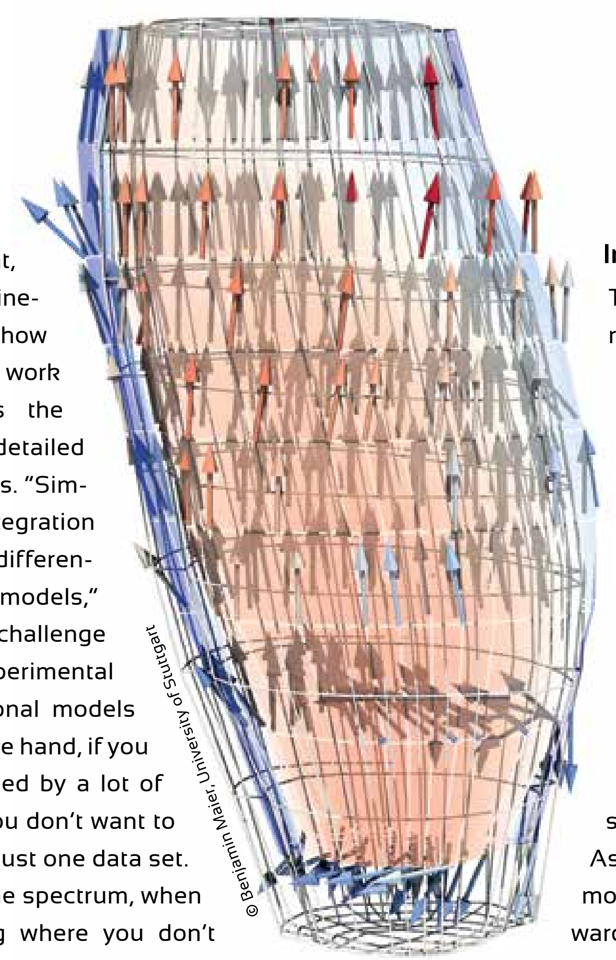

© Benjamin Maier, University of Stuttgart

and camera equipment, Röhrle extracts fine-grained details of how muscle groups work together, then uses the data as inputs for detailed computational models. "SimTech is about the integration of data into classical, differential equation-based models," Röhrle said. "The challenge with integrating experimental data and computational models is two-fold. On the one hand, if you have a model informed by a lot of experimental data, you don't want to over-fit the model to just one data set. At the other end of the spectrum, when modelling something where you don't have a lot of data, you have to have some kind of model because data can't tell you the entire story."

Recently, the team added realistic tendons to its muscle models. Although this expands the number of physics calculations that are needed, it also makes the models more realistic. Further, the team also began simulating electromyography signals, which allow medical professionals to identify disorders that impact motor skills. "This is one of the links to the experimental world in our work," said Benjamin Maier, a graduate researcher at Uni Stuttgart. "These EMG signals, as they are called, measure electric potential on the skin. To model this accurately, we need to be able to simulate a large number of muscle fibers in a highly parallel setting."

Integrating Inversions

The interdisciplinary biomechanical research collaboration at the University of Stuttgart is now approaching its next big challenge – going from primarily doing simulations that help confirm experimentalists' research findings to using modeling and simulation to help guide experiments. "A big thing in simulation science is going from solving problems in a 'forward' manner to solving inverse problems," Göddeke said. "We have just now reached the place where we can start to focus on inversions."

As supercomputing resources grow more powerful, the team also looks forward to coupling individual simulations of muscles to one another. Modelling a biceps and triceps in motion together in a human arm, for example, could improve understanding of how muscle systems work. *(EG)*

Selected User Publications 2020

Albers M, Meysonnat PS, Fernex D, et al. 2020. **Drag reduction and energy saving by spanwise traveling transversal surface waves for flat plate flow.** Flow Turbul Combust. 105: 125-157.

Atzori M, Vinuesa R, Fahland G, et al. 2020. **Aerodynamic effects of uniform blowing and suction on a NACA4412 airfoil.** Flow Turbulence Combust. 105: 735-759.

Badalov SV, Wilhelm R, Schmidt WG. 2020. **Photocatalytic properties of graphene-supported titania clusters from density-functional theory.** J Comput Chem. 41(21): 1921-1930.

Beck A, Dürrwächter J, Kuhn T, et al. 2020. **hp-Multilevel Monte Carlo methods for uncertainty quantification of compressible Navier-Stokes equations.** SIAM J Sci Comput. 42(4): B1067-B1091.

Bellwied R, Borsanyi S, Fodor Z, et al. 2020. **Off-diagonal correlators of conserved charges form lattice QCD and experiment.** Phys Rev D. 101: 034506.

Bhowmik A, Halder SK, Alon OE. 2020. **Impact of the transverse direction on the many-body tunneling dynamics in a two-dimensional bosonic Josephson junction.** Sci Rep-UK. 10: 21476.

Biktagirov T, Schmidt WG, Gerstmann U. 2020. **Spin decontamination for magnetic dipolar coupling calculations: application to high-spin molecules and solid-state spin qubits.** Phys Rev Res. 2: 022024(R).

Bocchini A, Eigner C, Silberhorn C, et al. 2020. **Understanding gray track formation in KTP: Ti3+ centers studied from first principles.** Phys Rev Materials. 4: 124402.

Borsanyi S, Fodor Z, Guenther JN, et al. 2020. **The QCD crossover at finite chemical potential from lattice simulations.** Phys Rev Lett. 125: 052001.

Cappellini G, Bosin A, Serra G, et al. 2020. **Electronic and optical properties of small metal fluoride clusters.** ACS Omega 5(22): 13268-13277.

Chatwell RS, Vrabec J. 2020. **Bulk viscosity of liquid noble gases.** J Chem Phys. 152: 094503.

Chepkasov IV, Ghorbani-Asl M, Popov ZI, et al. 2020. **Alkali metals inside bi-layer graphene and MoS2: insights from first-principles calculations.** Nano Energy. 75: 104927.

Chiarini A, Quadrio M, Auteri F. 2020. **A direction-splitting Navier–Stokes solver on co-located grids.** J Comp Phys. ePub Nov 26.

Chu X, Liu Y, Wang W, et al. 2020. **Turbulence, pseudo-turbulence, and local flow topology in dispersed bubbly flow.** Phys Fluids. 32: 083310.

Chu X, Wang W, Yang G, et al. 2020. **Transport of turbulence across permeable interface in a turbulent channel flow: interface-resolved direct numerical simulation.** Transport Porous Media. ePub Nov. 16.

Conroy AT, Ballance CP, Ramsbottom CA, Keenan FP. 2020. **Electron-impact excitation of Fe i. Astrophys J.** 902(1): 68.

Cormier M, Bühler M, Mauz M, et al. 2020. **CFD prediction of tip vortex aging in the wake of a multi-MW wind turbine.** J Phys Conf Ser. 1618: 06029.

Dehnhardt N, Luy JN, Klement P, et al. 2020. **Mixed group 14-15 metalates as model compounds for doped lead halide perovskites.** Angew Chem Int Ed. 10:1002/anie.202014696.

Diewald F, Lautenschlaeger MP, Stephan S, et al. 2020. **Molecular dynamics and phase field simulations of droplets on surfaces with wettability gradient.** Comp Methods Appl Mechan Eng. 361: 112773.

Donnari M, Pillepich A, Nelson, et al. 2021. **Quenched fractions in the IllustrisTNG simulations: comparison with observations and other theoretical models.** Mon Not R Astron Soc. 599(3): 4004-4024. (ePub Oct 10, 2020).

Dovesi R, Pascale F Civalleri B, et al. 2020. **The CRYSTAL code, 1976-2020 and beyond, a long story.** J Chem Phys. 152: 204111.

Dürrwächter J, Kurz M, Kopper P, et al. 2020. **An efficient sliding mesh interface method for high-order discontinuous Galerkin schemes.** Comput Fluids. ePub Dec 30.

Ehrle M, Waldmann, A, Lutz T, Krämer E. 2020. **Simulation of transonic buffet with an automated zonal DES approach.** CEAS Aeronautical J. 11: 1025-1036.

Encinar JA, Menendez JA. 2020. **Potential drugs targeting early innate immune evasion of SARS-coronavirus 2 via 2'-O-methylation of viral RNA.** Viruses. 12(5): 525.

Evrim C, Chu X, Laurien E. 2020. **Analysis of thermal mixing characteristics in different T-junction configurations.** Int J Heat Mass Tran. 158: 120019.

Evrim C, Laurien E. 2020. **Numerical study of thermal mixing mechanisms in T-junctions.** Appl Therm Eng. 116155.

Fernex D, Semaan R, Albers M, et al. 2020. **Actuation response model from sparse data for wall turbulence drag reduction.** Phys Rev Fluids. 5: 073901.

Flad D, Beck A, Guthke P. 2020. **A large eddy simulation method for DGSEM using non-linearly optimized relaxation filters.** J Comput Phys. 408: 109303.

Frey F, Thiemeier J, Öhrle C, et al. 2020. **Aerodynamic interactions on Airbus Helicopters' compound helicopter RACER in cruise flight.** J Am Helicopter Soc. 65(4): 1-14.

Ghaderzadeh S, Ladygin V, Ghorbani-Asl M, et al. 2020. **Free-standing and supported MoS monolayers under cluster irradiation: insights from molecular dynamics simulations.** ACS Appl Mater Interfaces. 12(33): 37454-37463.

Grandin M, Turc L, Battarbee M, et al. 2020. **Hybrid-Vlasov simulation of auroral proton precipitation in the cusps: comparison of northward and southward interplanetary magnetic field driving.** J Space Weather Space Clim. 10: 51.

Guettler N, Knee P, Ye Q, Tiedje O. 2020. **Initial droplet conditions in numerical spray painting by electrostatic rotary bell sprazers.** J Coatings Technol Res. 17: 1091-1104.

Guevara-Carrion G, Fingerhut R, Vrabec J. 2020. **Fick diffusion coefficient matrix of a quaternary liquid mixture by molecular dynamics.** J Phys Chem B. 124(22): 4527-4535.

Guo Z, Kloker MJ. 2020. **Effects of low-frequency noise in crossflow transition control.** J Am Inst Aeronaut Astronaut. 58(3).

Hajian CSS, Haringa C, Noorman H, Takors R. 2020. **Predicting by-product gradients of baker's yeast production at industrial scale: a practical simulation approach.** Processes 8: 1554.

Hanau K, Schwan S, Schäfer MR, et al. 2020. **Towards understanding the reactivity and optical properties of organosilicon sulfide clusters.** Angew Chem Int Ed. 10.1002/anie.202011370.

He Q, Li H, Li R, et al. 2020. **Constraining the inner density slope of massive galaxy clusters.** Mon Not R Astron Soc. 496(4): 4717-4733.

Heilmann N, Wolf M, Kozlowska M, et al. 2020. **Sampling of the conformational landscape of small proteins with Monte Carlo methods.** Sci Rep-UK. 10: 18211.

Herff S, Pausch K, Nawroth H, et al. 2020. **Impact of burner plenum acoustics on the sound emission of a turbulent lean premixed open flame.** Int J Spray Combust. 12(1): 1-20.

Hitz T, Jöns S, Heinen M, et al. 2020. **Comparison of macro- and microscopic solutions of the Riemann problem II. Two-phase shock tube.** J Comput Phys. ePub Nov 27.

Hitz T, Keim J, Munz CD, Rohde C. 2020. **A parabolic relaxation model for the Navier-Stokes-Korteweg equations.** J Comput Phys. 421: 109714.

Holtgrewe K, Mahatha SK, Sheverdyaeva PM, et al. 2020. T**opologization of β-antimonene on Bi2Se3 via proximity effects.** Sci Rep-UK. 10: 14619.

Homes S, Heinen M, Vrabec J, Fischer J. 2020. **Evaporation driven by conductive heat transport.** Mol Phys. e1836410.

Huang X, Albers M, Meysonnat, et al. 2020. **Analysis of the effect of freestream turbulence on dynamic stall of wind turbine blades.** Int J Heat Fluid Flow. 85: 108668.

Kashi A. 2020. **Wind collision and accretion simulations of the massive binary system HD 166734.** Mon Not R Astron Soc. 492(4): 5261-5270.

Kijanski N, Krach D, Steeb H. 2020. **An SPH approach for non-spherical particles immersed in Newtonian fluids.** Materials. 13(10): 2324.

Kim Y, Madsen HA, Aparicio-Sanchez M, et al. 2020. **Assessment of blade element momentum codes under varying turbulence levels by comparing with blade resolved computational fluid dynamics**. Renew Energ. 160: 788-802.

Klein BP, Harman SE, Ruppenthal L, et al. 2020. **Enhanced bonding of pentagon–heptagon defects in graphene to metal surfaces: insights from the adsoprtion of azulene and naphthalene to Pt(111).** Chem Mater. 32(3): 1041-1053.

Kohns M, Marx J, Langenbach K. 2020. **Relative permittivity of Stockmayer-type model fluids from MD simulations and COFFEE.** J Chem Eng Data. 65(12): 5891-5896.

Krenz M, Gerstmann U, Schmidt WG. 2020. **Photochemical ring opening of oxirane modeled by constrained density functional theory.** ACS Omega 5(37): 24057-24063.

Kretschmer S, Lehnert T, Kaiser U, Krasheninnikov AV. 2020. **Formation of defects in two-dimensional MoS_2 in the transmission electron microscope at electron energies below the knock-on threshold: the role of electronic excitations.** Nano Lett. 20(4): 2865-2870.

Kulkarni A, Garcia EJ, Damone A, et al. 2020. **A force field for poly(oxymethylene) dimethyl ethers (OME*n*).** J Chem Theory Comput. 16(4): 2517-2528.

Kuschel M, Takors R. 2020. **Simulated oxygen and glucose gradients as a prerequisite for predicting industrial scale performance *a priori*.** Biotech Bioeng. 117(9): 2760-2770.

Lamoth M, Jones T, Plodinec M, et al. 2020. **Nanocatalysts unravel the selective state of Ag.** Chem Cat Chem. 12: 1-13.

Lechte C, Conway GD, Görler T, Happel T, and the ASDEX Upgrade team. 2020. **Fullwave Doppler reflectometry simulations for density turbulence spectra in ASDEX Upgrade using GENE and IPF-FD3D.** Plasma Sci Technol. 22: 064006.

Lesnicki D, Zhang Z, Bonn M, et al. 2020. **Oberflächenladungen an der CaF_2-Wasser-Grenzfläche erlauben eine sehr schnelle intermolekulare Übertragung von Schwingungsenergie.** Angew Chem Int Ed. 59: 13116-13121.

Letzgus J, Heßler M, Krämer E. 2020. **Simulation of dynamic stall on an elastic rotor in high-speed turn flight.** J Am Helicopter Soc. 65(2): 1-12.

Letzgus P, El Bahlouli A, Leukauf D, et al. 2020. **Microscale CFD simulations of a wind energy test site in the Swabian Alps with mesoscale based inflow data.** J Phys Conf Ser. 1618: 062021.

Lintermann A, Schröder W. 2020. **Lattice-Boltzmann simulations for complex geometries on high-performance computers.** CEAS Aeronautical J. 11: 745-766.

Lipkowicz JT, Nativel D, Cooper S, et al. 2020. **Numerical investigation of remote ignition in shock tubes.** Flow Turbul Combust. ePub Oct. 3.

Lode AUJ, Lévêque C, Madsen LB, et al. 2020. **Colloquium: multiconfigurational time-dependent Hartree approaches for indistinguishable particles.** Rev Mod Phys. 92: 011001

Loureiro DD, Reutzsch J, Kronenburg, et al. 2020. **Primary breakup regimes for cryogenic flash atomization.** Int J Multiphase Flow. 132: 103405.

Mausbach P, Fingerhut R, Vrabec J. 2020. **Structure and dynamics of the Lennard-Jones fcc-solid focusing on melting precursors.** J Chem Phys. 153: 104506.

Moghadam SMA, Loosen S, Meinke M, Schröder W. 2020. **Reduced-order analysis of the acoustic near field of a ducted axial fan.** Int J Heat Fluid Flow. 85: 108657.

Mohapatra CK, Schmidt DP, Sforozo BA, et al. 2020. **Collaborative investigation of the internal flow and near-nozzle flow of an eight-hole gasoline injector (Engine Combustion Network Spray G).** Int J Engine Res. ePub Jun 10.

Mori P, Schwitalla T, Ware M, et al. 2020. **Downscaling of seasonal ensemble forecasts to the convection-permitting scale over the Horn of Africa using the WRF model.** Int J Climatology. Epub Aug 22.

Mustafa G, Nandekar PP, Mukherjee G, et al. 2020. **The effect of force-field parameters on cytochrome P450-membrane interactions: structure and dynamics.** Sci Rep-UK. 10: 7284.

Nathanail A. 2020. **A toy model for the electromagnetic output of neutron star merger prompt collapse to a black hole: magnetised neutron star collisions.** Astrophys J. 892(1): 35.

Nathanail A, Fromm CM, Porth O, et al. 2020. **Plasmoid formation in global GRMHD simulations and AGN flares.** Mon Not R Astron Soc. 495(2): 1549-1565.

Nathanail A, Gill R, Porth O, et al. 2020. **On the opening angle of magnetised jets from neutron-star mergers: the case of GRB170817A.** Mon Not R Astron Soc. 495(4): 3780-3787.

Navickas M, Giriunas L, Kalendra V, et al. 2020. **Electron paramagnetic resonance study of ferroelectric phase transition and dynamic effects in a Mn^{2+} doped $(NH_4)(Zn(HCOO)_3)$ hybrid formate framework.** Phys Chem Chem Phys. 22: 8513-8521.

Oman KA, Bahe YM, Healy J, et al. 2020. **A homogeneous measurement of the delay between the onsets of gas stripping and star formation quenching in satellite galaxies of groups and clusters.** Mon Not R Astron Soc. ePub Dec 14.

Pausch K, Herff S, Schröder W. 2020. **Noise sources of an unconfined and a confined swirl burner.** J Sound Vibr. 475: 115293.

Pindzola MS, Colgan J, Ciappina MF. 2020. **Neutron ionization of helium near the neutron-alpha particle collision resonance.** J Phys B: At Mol Opt Phys. 53(29): 205201.

Pindzola MS, Colgan JP. 2020. **Electron-impact double ionization of He (1s2s 3S).** J Phys B: At Mol Phys. ePub Dec 23.

Pinzon J, Siebenborn M, Vogel A. 2020. **Parallel 3d shape optimization for cellular composites on large distributed-memory clusters.** J Adv Simul Sci Eng. 7(1): 117-135.

Ren W, Reutzsch J, Weigand B. 2020. **Direct numerical simulation of water droplets in turbulent flow.** Fluids 5(3): 158.

Repetto P. 2020. **Tensorial solution of the Poisson equation and the dark matter amount and distribution of UGC 8490 and UGC 9753.** Mon Not R Astron Soc. 499(3): 3381-3398.

Reutzsch J, Kieffer-Roth C, Weigand B. 2020. **A consistent method for direct numerical simulation of droplet evaporation.** J Comput Phys. 413: 109455.

Schäfer F, Bastarrachea-Magnani MA, Lode AUJ, et al. 2020. **Spectral structure and many-body dynamics of ultracold bosons in a double-well.** Entropy-Switz. 22(4): 382.

Shang Z, Hashemi A, Berencén Y, et al. 2020. **Local vibrational modes of Si vacancy spin qubits in SiC.** Phys Rev B. 101: 144109.

Scheufele K, Subramanian S, Mang A, et al. 2020. **Image-driven biophysical tumor growth model calibration.** SIAM J Sci Comput. 42(3): B549-B580.

Schnücke G, Krais N, Bolemann T, Gassner GJ. 2020. **Entropy stable discontinuous Galerkin schemes on moving meshes for hyperbolic conservation laws.** J Sci Comput. 82: 69.

Schröder M. 2020. **Transforming high-dimensional potential energy surfaces into a canonical polyadic decomposition using Monte Carlo methods.** J Chem Phys. 152: 024108.

Schwitalla T, Warrach-Sagi K, Wulfmeyer V, Resch M*. 2020. **Near global scale high-resolution seasonal simulations with WRF-NOAHMP v.3.8.1.** Geosci Model Dev. 13: 1959-1974.

Sebastian R, Lürkens T, Schreyer AM. 2020. **Flow field around a spanwise-inclined jet in supersonic crossflow.** Aerosp Sci Technol. 106: 106209.

Shen X, Vogelsberger M, Nelson D, et al. 2020. **High-redshift JWST predictions from IllustrisTNG: II. Galaxy line and continuum spectral indices and dust attenuation curves.** Mon Not R Astron Soc. 495(4): 4747-4768.

Stephan S, Hasse H. 2020. **Interfacial properties of binary mixtures of simple fluids and their relation to the phase diagram.** Phys Chem Chem Phys. 22: 12544-12564.

Svyantyy VA, Kushnarenko VG. 2020. **Ein Ansatz zur gleichmässigen Lastverteilung zwischen Prozessoren des Mind-Simulators der dynamischen Netzobjekte mit verteilten Parametern.** Problems of Modeling and Design Automatization 1(15): 5-14.

Théveniaut H, Lan Z, Meyer G, Alet F. 2020. **Transition to a many-body localized regime in a two-dimensional disordered quantum dimer model.** Phys Rev Res 2: 033154.

Thiele L, Villaescusa-Navarro F, Spergel DN, et al. 2020. **Teaching neural networks to generate Fast Sunyaev-Zel'dovich Maps.** Astrophys J. 902: 129.

Thundathil R, Schwitalla T, Berhrendt A, et al. 2020. **Assimilation of Lidar water vapour mixing ratio and temperature profiles into a convection-permitting model.** J Metereol Soc Japan. 98(5): 959-986.

Übler H, Genel S, Sternberg A, et al. 2021. **The kinematics and dark matter fractions of TNG50 galaxies at z=2 from an observational perspective.** Mon Not R Astron Soc. 500(4): 4597-4619. (ePub Nov 10, 2020).

von der Grün M, Zamre P, Chen Y, et al. 2020. **Numerical study and LiDAR based validation of the wind field in urban sites.** J Phys Conf Ser. 1618: 042009.

Wei Y, Ghorbani-Asi M, Krasheninnikov A. 2020. **Tailoring the electronic and magnetic properties of hematene by surface passivation: insights from first-principles calculations.** J Phys Chem C. 124(41): 22784-22792.

Weih LR, Olivares H, Rezzolla L. 2020. **Two-moment scheme for general-relativistic radiation hydrodynamics: a systematic description and new applications.** Mon Not R Astron Soc. 495(2): 2285-2304.

Weih LR, Gabbana A, Dimeoni D, et al. 2020. **Beyond moments: relativistic Lattice-Boltzmann methods for radiative transport in computational astrophysics.** Mon Not R Astron Soc. 498(3): 3374-3394.

Wenz F, Boorsma K, Lutz T, Krämer E. 2020. **Cross-correlation-based approach to align turbulent inflow between CFD and lower-fidelity-codes in wind turbine simulations.** J Phys Conf Ser. 1618: 062005.

Wenzel C, Gibis T, Kloker M, Rist U. 2020. **Reynolds analogy factor in self-similar compressible turbulent boundary layers with pressure gradients.** J Fluid Mech. ePub Dec 2.

Wittig H, Gèrardin A, Cé M, et al. 2020. **Lattice calculation of the hadronic leading order contribution to the muon g-2.** EPJ Web Conf. 234: 01016.

Wragg J, Ballance C, van der Hart H. 2020. **Breit–Pauli R-matrix approach for the time-dependent investigation of ultrafast processes.** Comput Phys Commun. 254: 107274.

Wu Z, Bangga G, Lutz T, et al. 2020. **Insights into airfoil response to sinusoidal gusty inflow by oscillating vanes.** Phys Fluids. 32: 125107.

Wu Z, Wang Q, Bangga G, et al. 2020. **Responses of vertical axis wind turbines to gusty winds.** P I Mech Eng A-J Pow. ePub March 12.

Zamre P, Dessoky A, von der Grün M, et al. 2020. **Numerical study of the impact of urban terrain on the loads and performance of a small vertical axis wind turbine.** J Phys Conf Ser. 1618: 042005.

Zeman J, Kondrat S, Holm C. 2020. **Bulk ionic screening lengths from extremely large-scale molecular dynamics simulations.** Chem Commun. 56: 15635-15638.

Zirwes T, Zhang F, Habisreuther P, et al. 2020. **Identification of flame regimes in partially premixed combustion from a quasi-DNS dataset.** Flow Turbul Combust. ePub Oct 29.

* indicates HLRS staff member

Inside Our Computing Room

Hewlett Packard Enterprise Apollo (Hawk)

In 2020, HLRS installed its new flagship supercomputer, called Hawk. With a maximum LINPACK benchmark performance of 19.3 Petaflops, 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 December 2020, HLRS announced the upcoming addition of HPE Apollo systems with NVIDIA graphic processing units (GPUs) to Hawk. The expansion will enhance the center's capacity for deep learning and artificial intelligence applications, and enable 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 (to be installed in 2021)

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

Funding for Hawk was provided by the Baden-Württemberg Ministry for Science, Research and Art, 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.

Cray CS-Storm

The Cray CS-Storm is HLRS's primary system 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

Cray Urika-GX

The Urika-GX is optimized for classical machine learning applications, and is used for the analysis of large datasets produced by HLRS's flagship supercomputer. It provides an ideal platform for frameworks such as Apache Spark and scikit-learn for data mining and clustering in large datasets. The system has also been a key component in HLRS's research to advance the field of high-performance data analytics. (Though operational througout 2020, the Urika-GX is scheduled to be retired in early 2021.)

Number of nodes	41
Processor compute nodes	2 x Intel BDW 18-core, 2.1 GHz
Memory per node	512 GB
Disk capacity per node	2 TB
Software stack	Spark, Hadoop, Cray Graph Engine

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
GPU: 1 x AMD Radeon Pro WX8200
GPU memory: 8 GB

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

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

Interconnects
Infiniband EDR/FDR/HDR/QDR

User Profile

In 2020 the Gauss Centre for Supercomputing approved 8 new large-scale projects (each project requiring more than 35 million core hours) for HLRS's flagship supercomputer, for a total of 2.72 billion core-hours. The Partnership for Advanced Computing in Europe (PRACE) also approved 8 international simulation projects for HLRS, for a total of 440 million core-hours. In total, 129 projects, including test projects, were active on Hazel Hen and Hawk in 2020.

System Usage by Scientific Discipline

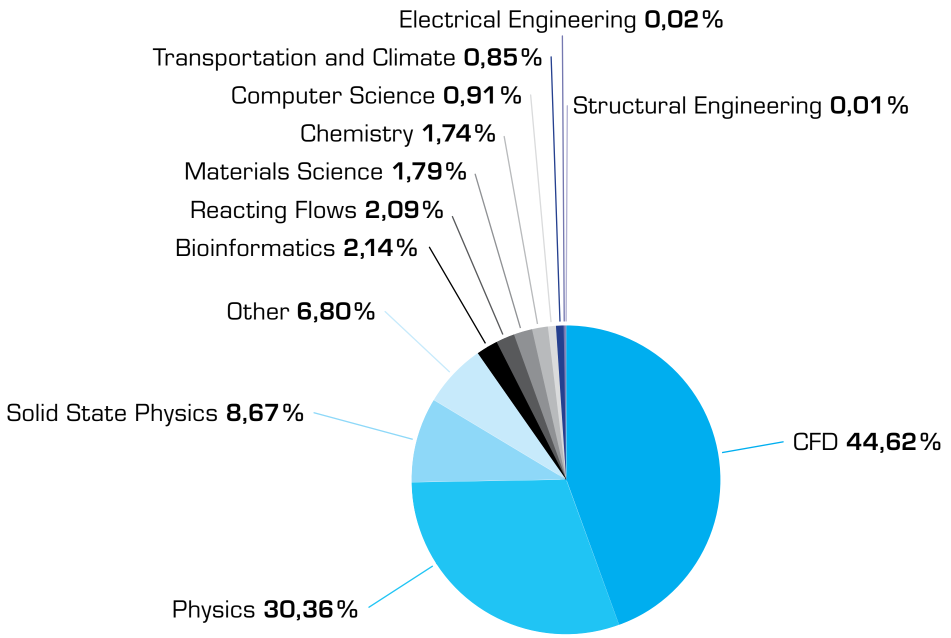

System Usage by State

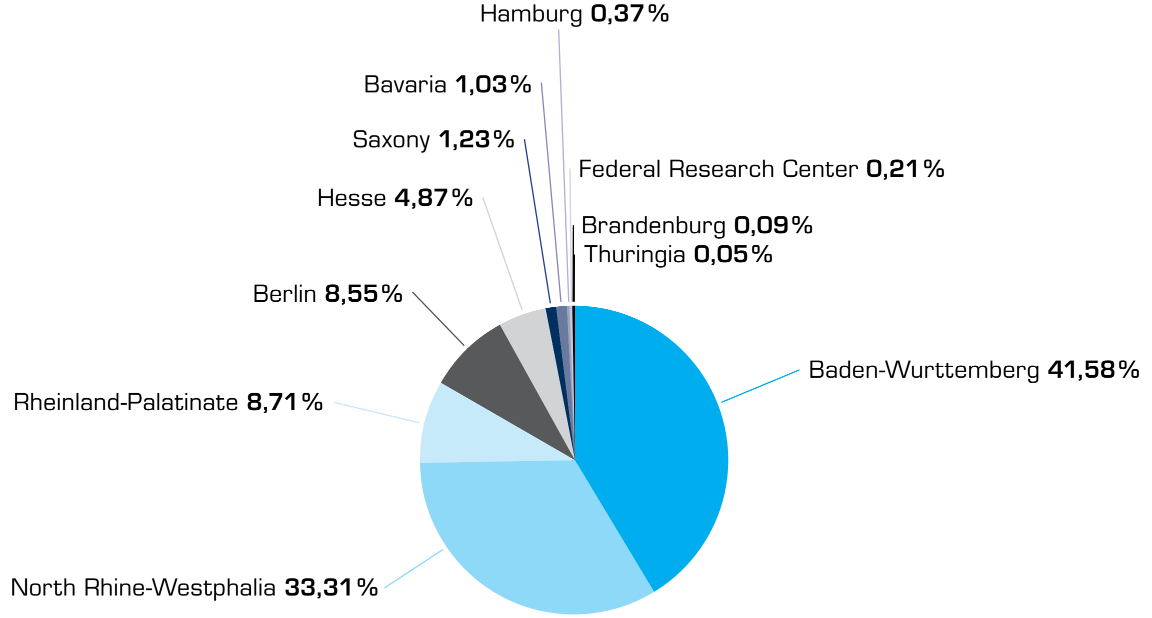

Third-Party Funded Research Projects

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 that operated in 2020.

For more information about our current projects, visit www.hlrs.de/about-us/research/current-projects/

Project	Duration	Funded by
aqua3S	**September 2019 - August 2022**	**EU**

➜ Developing a new system for detecting threats in drinking water safety and security, combining data from state-of-the-art sensors and other detection mechanisms.

bwHPC-S5	**July 2018 - March 2021**	**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.

bwVisu II	**August 2014 - October 2020**	**MWK**

➜ Developing a service for remote visualization of scientific data, particularly with respect to ensuring high scalability through cloud technologies.

Cape Reviso	**July 2020 - June 2023**	**BMVI**

➜ By combining machine learning, sensor technology, network analysis, and VR in digital twins, researchers are developing planning and decision support tools for conflict analysis and reduction between cyclists and pedestrians in urban locations.

CASTIEL	**September 2020 - August 2022**	**EU**

➜ This coordination and support action will enhance the activities of the EuroCC project by promoting collaboration and the exchange of knowledge and skills among HPC national competence centers across Europe.

CATALYST	**October 2016 - December 2021**	**MWK**

➜ Researches methods for analyzing large datasets produced by modelling and simulation with the goal of implementing a framework that combines HPC and data analytics.

ChEESE	**November 2018 - October 2021**	**EU**

➜ Preparing European flagship codes for upcoming pre-exascale and exascale supercomputing systems focusing on fields such as computational seismology, magnetohydrodynamics, physical volcanology, tsunamis, and the monitoring of earthquake activity.

CYBELE	**January 2019 - December 2021**	**EU**

➜ Integrates tools from high-performance computing, high-performance data analytics, and cloud computing to support the development of more productive, data driven methods for increasing agricultural productivity and reducing food scarcity.

EOPEN	**November 2017 - October 2020**	**EU**

➜ Tackles technical barriers that result from massive streams of Earth observation data and seeks to ensure that methods for data harmonization, standardization, fusion, and exchange are scalable.

EuroCC	**September 2020 - August 2022**	**EU**

➜ HLRS is the coordinating center of this Europe-wide project to establish national HPC competence centers and develop a shared, high level of expertise in high-performance computing, high-performance data analytics, and artificial intelligence.

EuroLab-4-HPC 2	**May 2018 - April 2020**	**EU**

➜ Laid the groundwork for establishing a European Research Center of Excellence for HPC systems.

EUXDAT	**November 2017 - October 2020**	**EU**

➜ Provided a platform for uniting HPC and cloud infrastructures to manage and process high amounts of heterogeneous data. Its focus was to support sustainable development in agriculture.

EXCELLERAT	**December 2018 - November 2021**	**EU**

➜ Facilitates the development of important codes for high-tech engineering, including maximizing their scalability to ever larger computing architectures and supporting the technology transfer that will enable their uptake in industry.

EXPERTISE	**March 2017 - February 2020**	**EU**

➜ A European training network for the next generation of mechanical and computer science engineers. Its objective is to develop advanced tools for analyzing fluid dynamics in large-scale models of turbine components and to eventually enable the virtual testing of an entire machine.

FF4EuroHPC | **September 2020 - August 2023** | **EU**

→ Conducts outreach and provides support to Europe's small and medium-sized enterprises (SMEs) to enable them to profit from the advantages offered by high-performance computing (HPC) technologies and services.

FocusCoE | **December 2018 - November 2021** | **EU**

→ Coordinates strategic collaboration and outreach among EU-funded Centres of Excellence to more effectively exploit the benefits of extreme scale applications for addressing scientific, industrial, or societal challenges.

HiDALGO | **December 2018 - November 2021** | **EU**

→ Develops novel methods, algorithms, and software for HPC and high-performance data analytics to accurately model and simulate the complex processes that arise in connection with major global challenges such as forced migration, air pollution, and the spread of disinformation through social media.

HPC-Europa 3 | **May 2017 - October 2021** | **EU**

→ Fosters transnational cooperation among EU scientists (especially junior researchers) who work on HPC-related topics such as applications, tools, and middleware.

HPCWE | **June 2019 - May 2021** | **EU**

→ A consortium of academic institutes, HPC centers, and industrial partners in Europe and Brazil that is developing novel algorithms and state-of-the-art codes to support the development of more efficient technologies for wind power.

HyForPV | **September 2018 - August 2021** | **BMWi**

→ Aims to develop and operationalize new prediction products for the integration of photovoltaics (PV) into the energy market and smart grids by delivering simulations of PV power output at high resolution.

InHPC-DE | **November 2017 - December 2021** | **BMBF**

→ Coordinates integration among Germany's three Tier-1 supercomputing centers to create a standardized and distributed HPC ecosystem. It provides funding for 100 Gbit networking and opportunities for high-speed data management and visualization.

MoeWe | **July 2016 - March 2021** | **ESF, MWK**

→ To address demand for supercomputing experts, particularly in industry, this project is developing a modular, flexible training program called the Supercomputing-Akademie.

NFDI4Cat | **October 2020 - September 2025** | **DFG**

→ Creating a German national research data infrastructure for catalysis and chemical engineering research.

OpenForecast | **September 2019 - May 2021** | **EU**

→ Develops approaches for combining freely available data and supercomputing resources to create a new generation of searchable data products for European citizens, public authorities, economic operators, and decision makers.

ORCHESTRA | **December 2020 - November 2023** | **EU**

→ Development of a networked platform for sharing data is enabling the creation of a new large-scale, pan-European cohort that will improve research and responses to the SARS-CoV-2 pandemic and provide a model for addressing future public health threats.

OSCCAR | **June 2018 - May 2021** | **EU**

→ Uses a novel, simulation-based approach to develop new systems for protecting vehicle occupants in accidents.

POP2 | **December 2018 - November 2021** | **EU**

→ A Center of Excellence that provides performance optimization and productivity services for academic and industrial users of HPC.

PRACE | **May 2019 - December 2021** | **EU**

→ Supports high-impact scientific discovery and engineering R&D to enhance European competitiveness for the benefit of society.

SDC4Lit | **May 2019 - April 2023** | **MWK**

→ The Science Data Center for Literature is an inter-disciplinary research project to sustainably organize the data-life cycle in digital literature. The resulting infrastructure will offer a data repository and research platform for the digital humanities.

SimTech | **October 2019 - March 2023** | **DFG**

→ This interdisciplinary Excellence Cluster at the University of Stuttgart is developing simulation technologies to enable integrative systems science. HLRS supports the development of efficient methods for uncertainty quantification and management.

Simulated Worlds | **February 2011 - June 2021** | **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.

Smart-DASH | **August 2016 - January 2020** | **DFG**
➜ Continued development of the C++ template library DASH, which offers distributed data structures with flexible data partitioning schemes and a set of parallel algorithms.

SODALITE | **February 2019 - January 2022** | **EU**
➜ Aims to provide an optimized, resilient, heterogeneous execution environment that enables operational transparency between cloud and HPC infrastructures.

TranSim | **January 2016 - February 2020** | **MWK**
➜ The project Transforming Society – Transforming Simulation explored how computer simulation is transforming science and the worlds of work, knowledge, and values.

Trust and Information | **August 2020 - July 2023** | **MWK**
➜ Multidisciplinary research led by the HLRS Department of Philosophy will develop perspectives for assessing the trustworthiness of computational science and limiting the spread of misinformation.

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
DFG	German Research Foundation
ESF	European Social Fund
EU	European Union
MWK	Baden-Württemberg Ministry for Science, Research, and Art

Structure

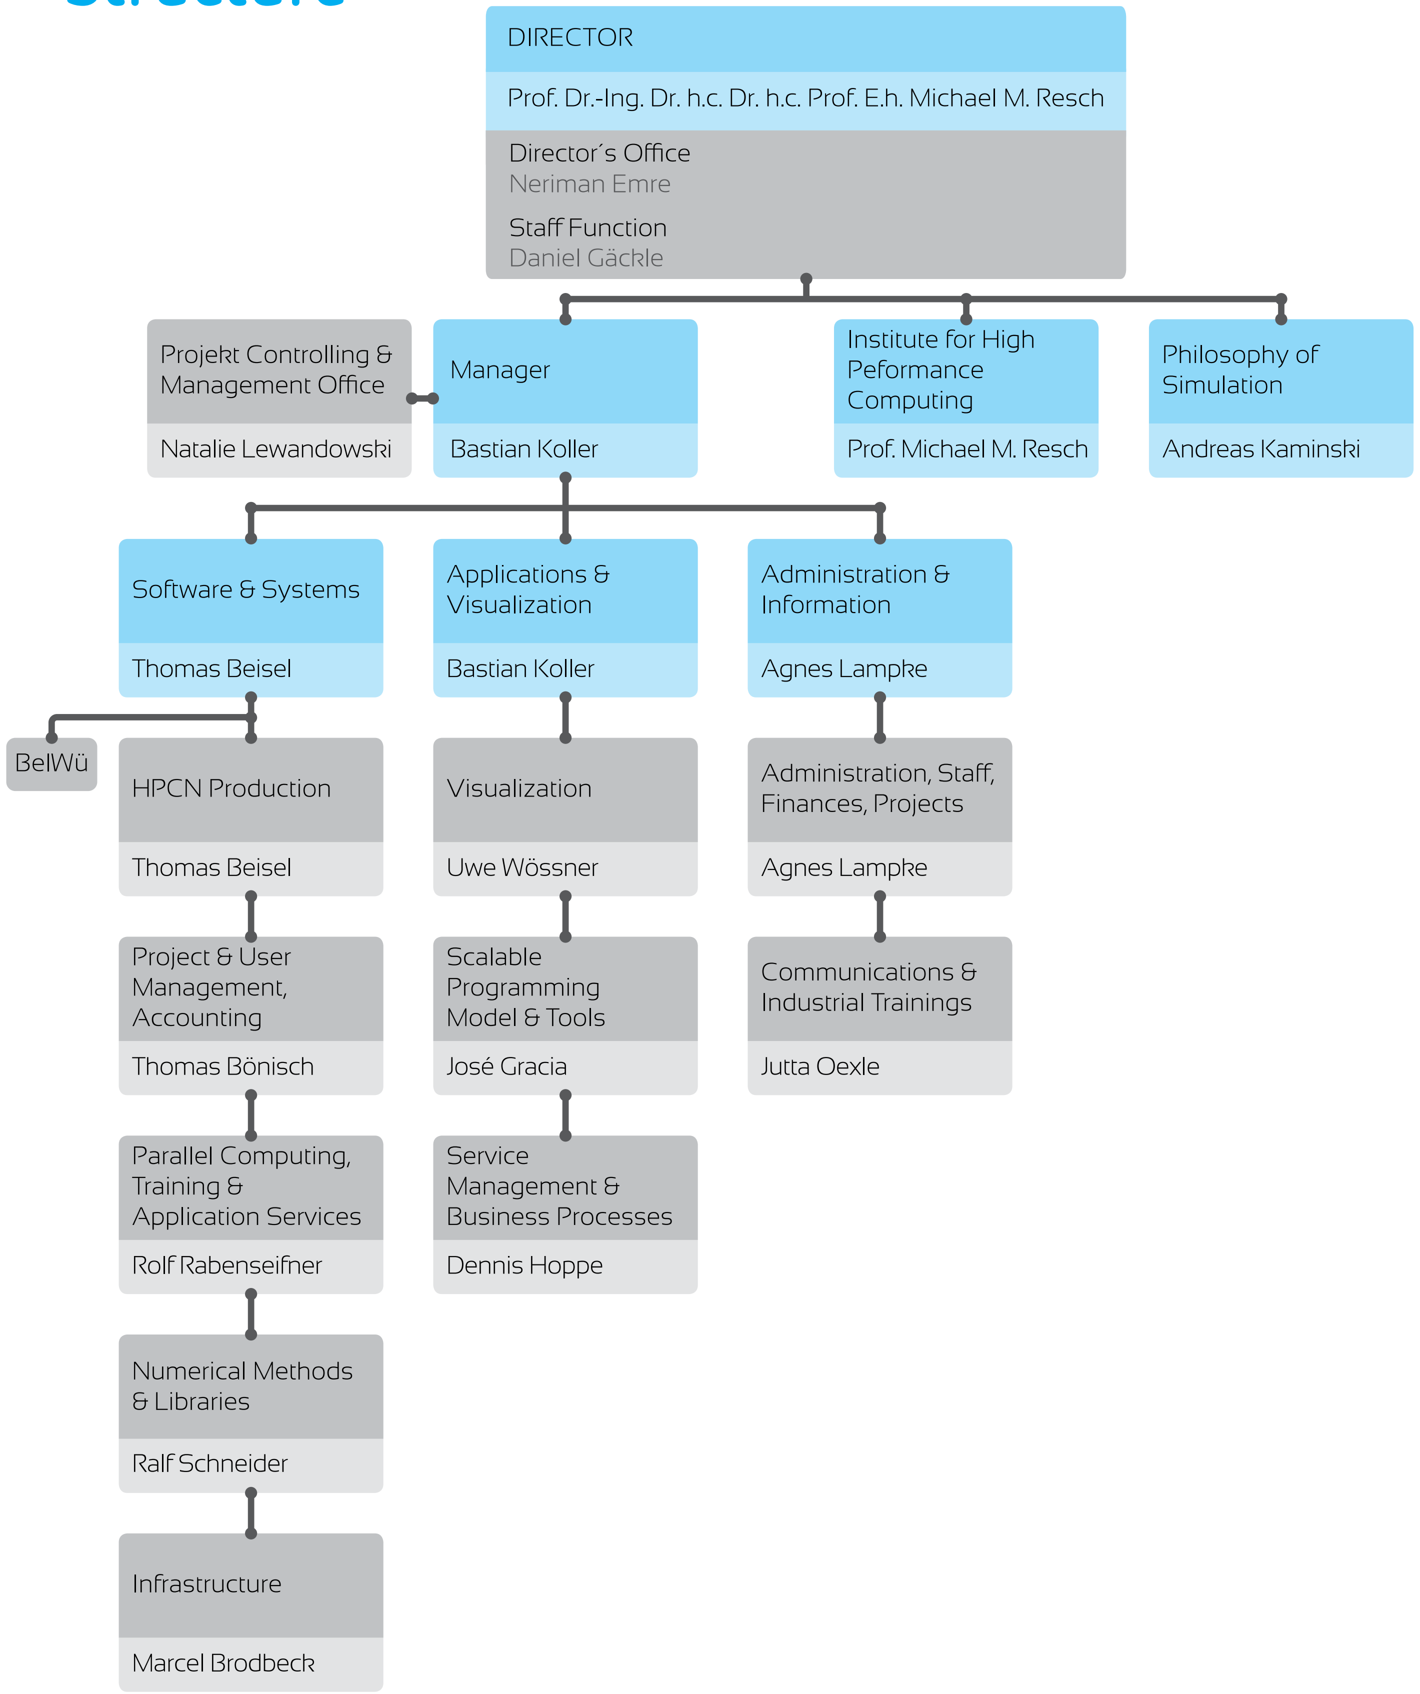

HPC Training Courses in 2020

HLRS offered 30 courses in 2020, providing continuing professional education on a wide range of topics relevant for high-performance computing. The courses took place over 100 course-days in Stuttgart and at other locations in Germany and internationally. A total of 1,005 trainees participated in these activities. For a current listing of upcoming courses, please visit www.hlrs.de/training.

Date	Location	Topic
Jan 27-28	Stuttgart	Introduction to Hybrid Programming in HPC
Jan 29-31	Stuttgart	HPE Porting and Optimization Workshop
Feb 10-14	Dresden	Parallel Programming (MPI, OpenMP) and Tools
Feb 24-28	Siegen	Introduction to Computational Fluid Dynamics
Mar 3-6	Stuttgart	Modern C++ Software Design (Intermediate) NEW
Mar 5-6	Stuttgart	HPC Technology Transfer Training Course NEW
Mar 9	Stuttgart	Introduction to NEC SX-Aurora TSUBASA Vector Platform NEW
Mar 13: Start of COVID-19 lockdown. The following courses took place online.		
Apr 20-24	Stuttgart	Fortran for Scientific Computing ((·)) *
Apr 20-30	Stuttgart	MOOC: A Short Introduction to MPI One-Sided Communication ((·)) * NEW
Apr 27-30	Stuttgart	Hawk Porting and Optimization Workshop ((·))
May 4-5	Vienna	Shared Memory Parallelization with OpenMP ((·)) TtT
May 5-8	Stuttgart	Modern C++ Software Design (Advanced) ((·))
May 6-8	Vienna	Parallelization with MPI ((·)) TtT *
Jun 9-12	Amsterdam	Parallel Programming with OpenMP and MPI ((·)) TtT *
Jun 17-19	Vienna/Stuttgart/Garching	Introduction to Hybrid Programming in HPC (MPI + X) ((·)) *
Jun 18-19	Stuttgart/Kaiserslautern	Efficient Parallel Programming with GASPI ((·)) *
Jun 29 - Jul 1	Stuttgart/Erlangen	Node-Level Performance Engineering ((·)) *
Jul 7-10	Stuttgart	Modern C++ Software Design (Intermediate) ((·)) NEW
Jul 14-17	Stuttgart/Garching	Deep Learning and GPU Programming Using Open ACC ((·)) NEW
Aug 24-27	Zurich	Parallel Programming with MPI / OpenMP ((·))
Sep 14-18	Stuttgart	35th VI-HPS Tuning Workshop ((·)) *
Sep 28 - Oct 2	Stuttgart	Computational Fluid Dynamics ((·))
Oct 12-16	Stuttgart	Parallel Programming Workshop (MPI, OpenMP & Advanced Topics) ((·)) *
Oct 21-23	Stuttgart	ChEESE Advanced Training on HPC for Computational Seismology ((·)) * NEW
Oct 26 - Nov 6	Stuttgart	MOOC: A Short Introduction to MPI One-Sided Communication ((·)) *
Nov 5-6	Vienna	Shared Memory Parallelization with OpenMP ((·)) TtT *
Nov 23-27	Vienna	Parallelization with MPI ((·)) TtT
Nov 24-27	Stuttgart	Modern C++ Software Design (Advanced) ((·))
Nov 30 - Dec 2	Jülich/Stuttgart	Advanced Parallel Programming with MPI / OpenMP ((·))
Dec 7-11	Stuttgart	Fortran for Scientific Computing ((·))

- Parallel Programming
- Computational Fluid Dynamics (CFD)
- Performance Optimization and Debugging
- Data in HPC / Deep Learning / Machine Learning
- Programming Languages for Scientific Computing
- Training for Special Communities

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

Divisions and Departments

Administration and Information
→ Leader: Agnes Lampke

Administration
Leader: Agnes Lampke

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.

Communications and Industrial Trainings
Leader: Dr. Jutta Oexle

Supervises and executes HLRS's communication to the general public and the media. It is the central point of contact for all questions regarding the center and its scientific work, and promotes new findings, achievements, and other news from around the center. In addition, the department designs and offers training courses and workshops for the industrial and service sectors, expanding interest in and accessibility of HPC technologies and solutions beyond its traditional community of scientific users.

Applications and Visualization
→ Leader: Dr. Bastian Koller

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 has expertise in tools such as virtual reality, augmented reality, and has designed a method for integrating processing steps spread across multiple hardware platforms into a seamless distributed software environment.

Scalable Programming Models and Tools
Leader: Dr. José Gracia

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

Service Management and Business Processes
Leader: Dennis Hoppe

Advances the convergence of high-performance computing and artificial intelligence, in particular with the goal of supporting hybrid HPC/AI workflows on a single infrastructure. This includes developing AI solutions, specifically in a business context, using cutting-edge technologies for big data, machine learning, and deep learning. The group also conducts research on related virtualization technologies such as containers, orchestration, and job scheduling. Leveraging synergies between virtualization and HPC, it has gained expertise in the development and operation of dynamic and scalable cloud computing services. The group efficiently applies performance and availability monitoring, elastic workflow management, and energy-efficient operation for federated cloud environments.

Project Controlling and Management
→ Leader: Dr. Natalie Lewandowski

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 offers coordinative assistance at the proposal planning and writing stage and acts as a supporting and coordinating entity between the HLRS management, the department heads, and the HLRS-administration in project-related matters.

Software and Systems
→ Leader: Thomas Beisel

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.

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.

Project and User Management, Accounting
Leader: Dr. Thomas Bönisch

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

Parallel Computing, Training and Application Services
Leader: Dr. Rolf Rabenseifner

Organizes HLRS's academic continuing education program in high-performance computing, with emphases on parallel programming, computational fluid dynamics, performance optimization, scientific visualization, programming languages for scientific computing and data in HPC. The department also organizes the review process for simulation projects running at the national supercomputing center and participates in service provision for industrial clients. Additionally, it provides installation and software support for academic researchers in structural mechanics and chemistry.

Staff Units: Related Research

Philosophy of Science and Technology of Computer Simulation
Leader: Dr. Andreas Kaminski

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?

Infrastructure
Leader: Marcel Brodbeck

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.

High-Performance Computing Center Stuttgart (HLRS)

University of Stuttgart
Nobelstrasse 19 | 70569 Stuttgart | Germany

phone ++49 (0)711 685-87269
fax ++49 (0)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

Leader, Department of Communications and Industrial Training: Dr. Jutta Oexle

Editor: Christopher M. Williams

Contributing Writers: Christopher M. Williams *(CW)*, Eric Gedenk *(EG)*

Translations: Proverb oHG

Production Manager: F. Rainer Klank

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

Printing: Nino Druck GmbH

Design: Zimmermann Visuelle Kommunikation
www.zimmermann-online.info

Institutional affiliations:

University of Stuttgart

GCS
Gauss Centre for Supercomputing

Funding for Hawk provided by:

Baden-Württemberg
MINISTRY OF SCIENCE, RESEARCH AND ARTS

SPONSORED BY THE
Federal Ministry of Education and Research

This magazine is printed on paper that has been certified by FSC, the EU Ecolabel, and the Blue Angel Ecolabel.

HLRS

High-Performance Computing Center | Stuttgart

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

Simulation of air circulation in an automated clean room for packaging pharmaceuticals.