HLRS
High-Performance Computing Center | Stuttgart

2019

ANNUAL REPORT

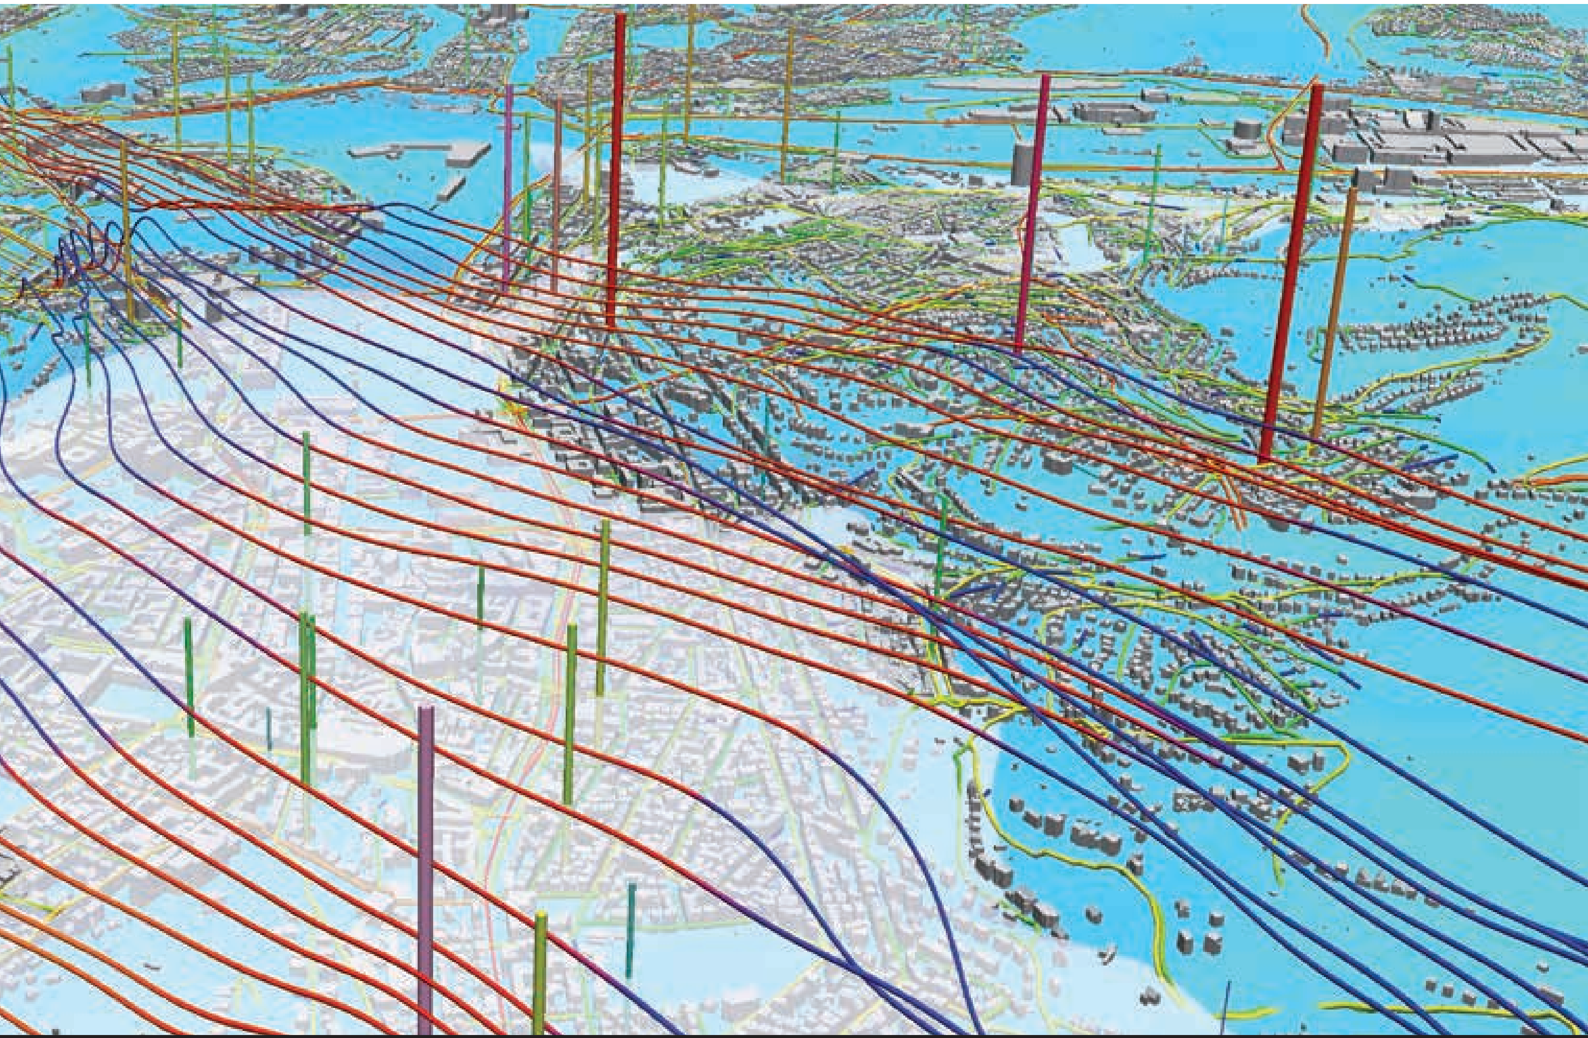

ENERGY HEALTH CLIMATE MOBILITY PHILOSOPHY

2019 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 2019 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 most important trend of the last year — and one that will keep us busy for the foreseeable future — is the convergence of high-performance computing and artificial intelligence. I am proud to say that HLRS has seized upon this development early, and in 2019 we started several novel collaborations in the field.

One example is a project we undertook with broadcaster SWR to develop a new tool for analyzing music and recommending related titles. The partnership not only enabled HLRS to help develop a practical application of new AI technologies, but will also support the development of other media-related projects.

This year also saw the start of a partnership with LandesCloud to provide a secure, cloud-based platform for AI as a service for industry. Considering that LandesCloud works closely with the Landesbank Baden-Württemberg, one of Germany's largest financial institutions, the collaboration will enable HLRS to bring AI and HPC resources into new domains.

In order to provide our users with access to AI-optimized computing resources, HLRS acquired a GPU-based Cray CS-Storm system late in 2019, making it possible to run AI jobs on a computing cluster specifically designed for these applications. In the coming years we also plan to explore ways to more efficiently merge AI and HPC systems into a single platform. The foundation for this will be our new Hewlett Packard Enterprise Apollo system, called Hawk. We began Hawk's installation this year and look forward to the

Willkommen beim Jahresbericht des Höchstleistungsrechenzentrums Stuttgart (HLRS) für 2019. Es freut uns, Ihnen einige Höhepunkte aus dem letzten Jahr und spannende Expansionsvorhaben zu präsentieren.

Der wichtigste Trend des letzten Jahres, der uns auch auf absehbare Zeit beschäftigen wird, ist die Annäherung von Hoch- und Höchstleistungsrechnen (HPC) und Künstliche Intelligenz (KI). Das HLRS hat diese Entwicklung früh aufgegriffen, was mich stolz macht, und 2019 gingen wir mehrere Partnerschaften in diesem Bereich ein.

So begannen wir ein Projekt mit dem SWR zur Entwicklung eines neuen Tools zur Musikanalyse und Empfehlung ähnlicher Titel. Diese Partnerschaft ermöglichte es dem HLRS, die Entwicklung einer praktischen Anwendung neuer KI-Technologien zu unterstützen, trägt aber auch zu anderen Medienprojekten bei.

In diesem Jahr begann auch eine Partnerschaft mit der LandesCloud zur Bereitstellung einer sicheren, cloud-basierten KI-Plattform für die Industrie. Da LandesCloud eng mit der Landesbank Baden-Württemberg zusammenarbeitet, kann das HLRS durch dieses Projekt KI- und HPC-Ressourcen in neue Bereiche einführen.

Um unseren Nutzern Zugang zu KI-optimierten Rechnersystemen zu ermöglichen, erwarb das HLRS Ende 2019 ein GPU-basiertes Cray CS-Storm-System; damit können KI-Aufgaben auf einem Rechen-Cluster zugeschnitten auf diese Anwendungen ausgeführt werden. In den nächsten Jahren streben wir nach einer effizienteren Zusammenführung von KI- und HPC-Systemen auf einer einzigen Plattform. Grundlage hierfür ist unser neues Hewlett Packard Enterprise Apollo 9000-System (Hawk). In diesem Jahr begann die Installation new possibilities the 26-petaflop supercomputer will offer for academic and industrial research once it begins operation in 2020.

In 2019 we were also able to bring our efforts to improve sustainability to a new level. Over the last several years HLRS has intensively analyzed its energy supply, cooling activities, and other activities that affect our environment. This has enabled us to implement many substantial improvements that have reduced cooling costs and other environmental impacts. After intensive preparation, HLRS passed a demanding audit to receive ISO 14001 and ISO 50001 certifications in November, and became the first European high-performance computing center to qualify for the EU's Eco-Management and Audit Scheme (EMAS), the world's most stringent framework for environmental management in organizations.

The past year also witnessed an extension of our international collaborations. We entered formal partnerships with both the National Supercomputing Center at the Sun Yat-Sen University in Guangzhou and the Supercomputing Center of the University of Science and Technology of China in Hefei. These partnerships will enable us to work together on topics of shared interest, including the operation of future exascale systems as well as the convergence of AI and HPC. We were also pleased to extend our long-time collaboration with the National Center for High Performance Computing in Hsinchu, Taiwan.

Two workshops should be mentioned here, as they highlight important areas in which simulation and visualization technologies are being used at HLRS. In May, experts in the field of urban planning joined us for a two-day meeting focusing on new applications of simulation and artificial intelligence in city planning processes. In September, high-performance computing centers from all over the world also joined us for the International Industrial Supercomputing Workshop. Considering HLRS's support of the dynamic high-tech industry in Baden-Württemberg, the topic is of utmost importance and the workshop strengthened our ties with the international community active in this field.

HLRS's key performance indicators highlight a successful year. We were able to increase the number of scientists on staff by about 20% and increased our scientific output. We were also delighted to see six doctoral students complete their PhD's. In comparison to 2018 we saw a decline in third-party funding, although the spike in the previous year was due to a one-time, large-scale federal government project. Industrial usage is high with an increase of companies using our systems. We saw a slight drop in the number of participants in our trainee program, though this is due to some ongoing adjustments to improve our curriculum.

Scientific users of our supercomputing resources also achieved great things this year. We were proud to have played a role in supporting the landmark achievement of the Event Horizon Telescope consortium, which produced the first ever image of a black hole. Simulations performed at HLRS by researchers at the Goethe University Frankfurt led to insights about plasma dynamics under extreme gravity conditions that were used in creating the image. This annual report also features reports on research by HLRS users to develop high-resolution models of combustion and methods for optimizing wind turbine blade designs to generate renewable energy more efficiently.

We are also pleased to present an interview with Matthias Hauser, who this year joined us as managing director of the Media Solution Center. In the conversation he describes HLRS's one-of-a-kind approach to supporting the arts and media industry, and the unique synergies that emerge when simulation scientists and artists work together.

I would also like to use this occasion to thank the supporters and funders who have made HLRS's successes in 2019 possible. At the same time, we look forward to continuing in our efforts to find innovative ways of using high-performance computing and other advanced digital technologies to address the most pressing challenges facing science, industry, the HPC community, and society at large.

With best regards,

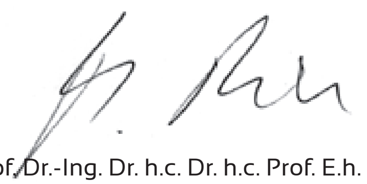

Prof. Dr.-Ing. Dr. h.c. Dr. h.c. Prof. E.h. Michael M. Resch
Direktor, HLRS

des Hawk und wir freuen uns auf die neuen Möglichkeiten dieses 26-Petaflop-Supercomputers für die Forschung und Industrie ab Inbetriebnahme 2020.

2019 setzten wir mehr denn je auf Nachhaltigkeit. In den letzten Jahren hat das HLRS seine Stromversorgung, Kühlsysteme und andere Umweltfaktoren ausgiebig analysiert. Dies führte zu wesentlichen Verbesserungen, welche den Kühlaufwand und andere Umweltauswirkungen verringert haben. Nach intensiver Vorbereitung bestand das HLRS eine Umweltprüfung und ist seit November gemäß ISO 14001 und ISO 50001 zertifiziert. Es qualifizierte sich auch als erstes europäisches Höchstleistungsrechenzentrum für das Eco-Management and Audit Scheme (EMAS) der EU, dem weltweit anspruchsvollsten Rahmen für das Umweltmanagement in Organisationen.

Im vergangenen Jahr weiteten wir auch unsere internationale Zusammenarbeit aus. Wir begannen formelle Partnerschaften mit dem National Supercomputing Center an der Sun-Yat-sen-Universität in Guangzhou und dem Supercomputing Center der Chinese University for Science and Technology in Hefei. Durch diese können wir bei Themen von gemeinsamem Interesse zusammenarbeiten, wie dem Betrieb zukünftiger Exascale-Systeme und der Verknüpfung von KI und HPC. Erfreulicherweise verlängerten wir auch unsere dauerhafte Zusammenarbeit mit dem National Center for High-Performance Computing in Hsinchu, Taiwan.

Zwei Workshops gilt es hier hervorzuheben, da sie wichtige Einsatzbereiche von Simulations- und Visualisierungstechnologien am HLRS aufzeigen. Im Mai trafen wir uns zwei Tage mit Stadtplanungs-Experten, um neue Anwendungsmöglichkeiten von Simulationen und KI auf diesem Gebiet zu erörtern. Im September nahmen zudem Höchstleistungsrechenzentren aus der ganzen Welt am International Industrial Supercomputing Workshop teil. Da das HLRS den dynamischen Hightech-Sektor Baden-Württembergs unterstützt, ist dieses Thema äußerst wichtig, und der Workshop hat unsere Beziehungen zu internationalen Experten in diesem Bereich gestärkt.

Die Leistungskennzahlen des HLRS sprechen für ein sehr erfolgreiches Jahr. Wir konnten die Anzahl an Wissenschaftlern um 20% erhöhen und erhöhten auch unsere wissenschaftliche Leistung. Es freute uns auch sehr, dass sechs Doktoranden ihren Abschluss machten. Im Vergleich zum Vorjahr verzeichneten wir einen leichten Rückgang an Einnahmen aus Drittmitteln, weil ein einmaliges Large-Scale-Bundesprojekt in 2018 finanziell abgewickelt wurde. Die industrielle Nutzung unserer Systeme war auch durch die gestiegene Anzahl an Industriekunden hoch. Die Besucherzahlen unserer HPC-Schulungen gingen leicht zurück, verursacht durch derzeit laufende Neuausrichtung zur Verbesserung unseres Lehrplans.

Auch die wissenschaftlichen Nutzer unserer Supercomputer haben in 2019 Großes geleistet. Wir sind stolz, unseren Teil zum Spitzenerfolg des Event Horizon Telescope Consortium beigetragen zu haben, welches das erste Bild eines Schwarzen Lochs überhaupt erstellt hat. Dieser Jahresbericht enthält auch Berichte über Forschungsarbeiten von HLRS-Nutzern zur Entwicklung von hochauflösenden Verbrennungsmodellen und Methoden zur Optimierung des Rotorblattdesigns für Windkraftanlagen zur effizienteren Erzeugung erneuerbarer Energie.

Wir freuen uns auch über ein Interview mit Matthias Hauser, der 2019 als Leiter des Media Solution Center zu uns kam. Darin beschreibt er den einmaligen Beitrag des HLRS zur Kunst und zu den Medien und die einzigartigen Synergieeffekte aus der Zusammenarbeit von Simulationswissenschaftlern und Künstlern.

Bei dieser Gelegenheit möchte ich allen Unterstützern und Geldgebern danken, welche die Erfolge des HLRS 2019 ermöglicht haben. Wir freuen uns darauf, auch in Zukunft Höchstleistungsrechnen und andere fortschrittliche digitale Technologien innovativ einzusetzen, um die drängendsten Herausforderungen für Wissenschaft, Industrie, die HPC-Gemeinde und die Gesellschaft insgesamt anzugehen.

CONTENTS

SPOTLIGHT

As part of the EOPEN project, HLRS staff helped implement an unsupervised method for automatically detecting rice pads in South Korea (shown in green). The findings make it possible to estimate the total area of rice pads, enabling experts to predict grain yield in the following year.

Preparing for the AI Future

The convergence of high-performance computing and artificial intelligence is beginning to offer exciting new opportunities for accelerating science and technology development. New programs and infrastructure at HLRS have begun laying the foundation to capitalize on this promise.

Although the idea of artificial intelligence (AI) has been around since at least the 1950s, recent advances in computer science seem to indicate that its time has finally arrived. Cheaper and more powerful computers, new software, the availability of vast collections of data, and increasingly ubiquitous sensor devices have made it possible to begin developing new kinds of sophisticated automated systems that can rapidly integrate and analyze data to make predictions, deliver insights, arrive at decisions, and even take actions without the need for human intervention. In response to the recent growth in AI investments in the United States and China, Europe has recently set out to develop its infrastructure and know-how for AI, a step that will be important for bolstering its research capabilities, industrial productivity, and global economic competitiveness. Growing awareness of AI has also led increasing numbers of scientists and engineers to begin exploring how it could enhance their research and technology development. Meanwhile, AI is becoming a key component in everyday technologies for fields as diverse as agriculture, finance, manufacturing, and the operation of self-driving vehicles, among many others.

Recognizing the growing demand from its community of academic and industrial HPC users, HLRS took steps in 2019 to expand its capabilities and resources for artificial intelligence, including procuring a new supercomputing system that is optimized for AI applications and initiating projects that are testing practical applications of AI. At the same time, philosophers and social scientists at the supercomputing center have been working to better understand and address questions related to the trustworthiness of AI-based decision-making.

HLRS is still early in its expansion into AI, but already it has been laying the foundation for advances at the convergence of HPC, high-performance data analytics, and artificical intelligence approaches such as deep learning. Focusing on real-world scenarios where they can have an impact, HLRS is exploring how such new methods and technologies can be combined to best support the center's users and open up new ways to address global challenges.

Why AI needs HPC

In the hype surrounding the recent AI "gold rush," some have speculated that artificial intelligence applications running on modestly sized computers could soon replace the more established field of simulation using supercomputers. Research methods being explored at HLRS, however, suggest that for the foreseeable future the two approaches will deliver the greatest benefit by being used in a complementary way.

Explained simply, artificial intelligence is a type of automated pattern recognition. Using deep neural networks — a computer science method modeled on how neurons in the human brain extract meaning and make decisions based on sensory input — applications for deep learning identify and compare differences in large collections of training data, revealing characteristics in the data that carry significance. In automated driving, for example, this might mean image-processing algorithms that distinguish pedestrians in the car's field of vision. Once these patterns have been clearly defined, AI systems use them as a model for decision-making — for example, how to avoid collisions during driving.

Developing a training model that an AI algorithm can use to "learn," however, requires the generation and processing of an enormous amount of data — a task that is ideally suited to high-performance computers. Simulating a car's field of vision as it moves through a city, for example, means simultaneously representing millions of features, and not just once but continuously over split-second time steps. Developing such a model is impossible without high-performance computers.

After an AI application combs through large datasets to develop a predictive model, it is then possible to convert it into a simplified AI tool that does not require the entire dataset, but only an understanding of the key parameters within it. In autonomous driving, for example, it is impossible for a car to carry a supercomputer onboard or rely on cloud technologies to communicate with a remote supercomputer fast enough to navigate rapidly changing road conditions. Instead, the model derived using HPC is converted into a low-power AI tool whose job is merely to observe and react based on key features in the data.

Such an approach is not only relevant for autonomous driving, but also for many kinds of AI applications. In agriculture, for example, AI connected to new kinds of sensors could help farmers optimize the timing of their planting and harvests. In engineering, inexpensive digital tools could monitor bridges on highways to alert maintenance staff when repairs are necessary. For any of these applications, however, researchers will need both supercomputing and deep learning to develop the models on which they are based.

How AI could accelerate traditional HPC research

While HPC is necessary to develop AI applications, artificial intelligence also holds promise to accelerate research based in traditional simulation methods. As HLRS Managing Director Dr. Bastian Koller explains, "HPC systems will continue to grow, which is good because it means we can tackle bigger problems. The troubling issue, though, is that the amount of data these larger systems produce is becoming unmanageable and it's becoming harder and harder to know what data is of value. We think that AI could be a natural fit to solve this problem."

In the field of computational fluid dynamics, for example — an area in which many users of HLRS's computing resources specialize — scientists have developed high-precision computational models of complex physical phenomena like turbulence in air- and water flow. As supercomputers have grown in speed and power, they have made it possible to mathematically represent those properties at extremely high resolution. Producing these precise models, however, is computationally expensive, making it impractical for certain design optimization problems where it would be desirable to quickly compare many different simulations to find the best solution

The arrival of AI offers a different approach. As more data has become available, scientists have begun exploring synthetic approaches using deep neural networks. Instead of trying to meticulously model an entire system based on physical principles, researchers have begun using deep learning to develop surrogate models. Here, they use a bottom-up, data-driven approach to create models that reproduce the relationship between the input and output data on which the model is based. Such surrogate models closely mimic the bahavior of a traditional simulation model, and even if they do not have the same precision or degree of explainability of traditional simulations, they can approximate those results much more quickly. For some problems, this offers huge advantages.


Dennis Hoppe

According to Dennis Hoppe, an HLRS scientist focused on artificial intelligence, "Using neural networks can enable performance gains, even if there is a tradeoff due to the loss of accuracy. This is something that some of our users, particularly our engineers, are not fond of. Still, for certain kinds of problems this approach makes it possible to get much faster results. And when you find something interesting, you're already dealing with a reduced parameter space, making it easier to recompute specific parameter sets using the 'real' simulation in a much more targeted way."

Koller also suggests that this data-driven approach could lead to novel insights. "One of the great opportunties for using AI to support traditional applications of high-performance computing is for identifying what I like to call the unknown unknowns," he explains. "There are sometimes dependencies in large data sets that we as humans don't see because we aren't looking for them. A machine methodically comparing data can often identify parameters within systems that have important effects that we would have missed, and that turn out to offer new opportunities we never would have thought to explore."

This perspective suggests that in the coming years AI and HPC are on a collision course in which new kinds of hybrid approaches combining both disciplines will offer new opportunities for research and technology development. HPC will produce the large datasets needed to develop robust models, deep learning will help researchers to analyze that data more efficiently, and HPC will use the results to create more robust models. Using these two approaches in an iterative way could thus help researchers make sense of complex systems more quickly.

Toward hybrid systems and workflows

Realizing this goal, however, presents technical challenges, particularly because these very different computational disciplines require different types of computer hardware and software.

Supercomputers optimized for simulation have grown to consist of many thousands of central processing units (CPUs), a type of computer processor that is good at quickly breaking complex calculations into parts, performing those smaller calculations rapidly in a parallel manner across many CPUs, and then recombining them to produce a result. HLRS's supercomputer — as of 2020 a 26 petaflop system called Hawk — is based on this model because it provides the best infrastructure for solving the most common simulation problems facing the center's users.

Artificial intelligence applications, on the other hand, have taken off due to the arrival of another type of

processor called a graphics processing unit (GPU). Originially designed for video gaming, these processors work fastest for iterative programs in which the same operation is performed repeatedly with slight variations, such as in a deep learning algorithm using neural networks.
In addition to different hardware, HPC and AI utilize different programming languages. Whereas simulation on large HPC systems often uses Fortran or C++, for example, artificial intelligence applications are more commonly developed in Python or TensorFlow. Developing workflows that can easily communicate between different languages will thus also be important.
In 2016 as interest grew in the possibilities offered by big data, HLRS began exploring opportunities to better integrate HPC and high-performance data analytics. Working together with supercomputer manfacturer Cray Inc. in a project called CATALYST, HLRS has been testing a combined hardware and software data analytics system called Urika-GX for its ability to support engineering applications. The researchers already investigated the system's ability to integrate with Hazel Hen, HLRS's flagship supercomputer at that time. By 2019, demand had been growing among HLRS's users not just for conventional data mining and machine learning, but also for deep learning using GPUs. In response, the center procured a new supercomputing platform, the Cray CS-Storm. A GPU-based system that comes with Cray's AI programming suite installed, the new system now makes it possible for HLRS to support research involving both HPC and AI under one roof.
Currently, HLRS has separate, dedicated systems for HPC and deep learning, each of which with its own architecture, programming frameworks, and job scheduler. In the future, Hoppe says, "the goal is to develop an integrated, holistic system for resource management and workflow execution that connects everything together in ways that make it possible to run hybrid HPC/AI workflows on a single platform."

Ethical implications of AI

At the same time that HLRS has been developing its technical infrastructure for AI, philosophers and social scientists in its Department of Philosophy of Science and Technology have begun addressing a very different set of questions facing the field: namely, when is it appropriate to rely on machine learning systems and how can we ensure that decisions produced by AI reflect our values?
Deep learning algorithms that use neural networks to develop surrogate models, for example, function like "black boxes," making it extremely difficult to reconstruct how they arrive at their findings. Although the results are often useful in practice, this presents an epistemological problem for philosophers, forcing us to ask what we can say we actually know when presented with such results. As AI becomes more pervasive, playing roles in fields such as loan risk evaluation or prison sentencing that have direct impacts on people's lives, this limitation is disconcerting, particularly for nonscientists who are affected by AI but do not have the context to understand how it works.
During an HLRS conference in November 2019, a multidisciplinary gathering of researchers explored such questions. (see page 23). Meanwhile, because existing ethical guidelines have not had widespread impact, members of the HLRS philosophy department joined together with VDE, the Bertelsmann Foundation, the Institute for Technology Assessment and Systems Analysis at Karsruhe Institute of Technology, and the International Center for Ethics in the Sciences and Humanities at Tübingen University to develop a more systematic framework for putting AI ethics principles into practice. As in standardization frameworks for other fields, this AI Ethics Impact Group is developing an enforceable set of ethical guidelines to orient scientists, developers, consumers, and others affected by AI.
"Because AI applications are becoming more common, the time is ripe for careful consideration of when it is appropriate to rely on machine learning," says Dr. Andreas Kaminski, leader of the HLRS's philosophy department. Doing so will require collaboration between computer scientists and other kinds of researchers, like philosophers, social scientists, and historians of science, who can provide context and methods for understanding these new technologies criticially."

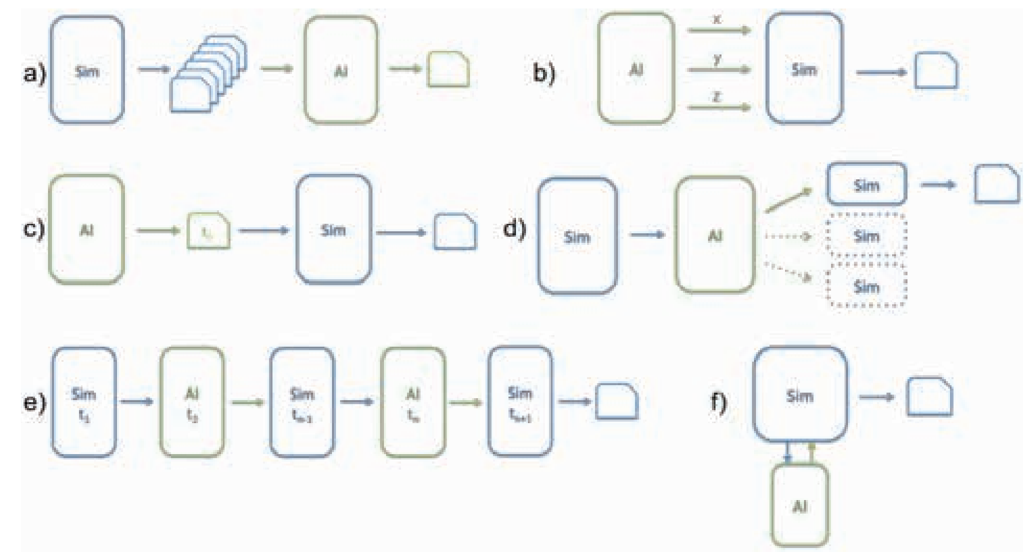

Various interaction patterns between AI and simulations: a) generation of data for AI training from simulations; b) finding simulation parameters by AI; c) generation of initial solutions by AI; d) choosing the right path by AI; e) replacing iterations in simulation by AI; f) replacing specific functions/equations by AI.

Experimenting with new applications of AI

As Koller explains, HLRS's engagement with AI will continue to evolve with the field: "With the CATALYST project our goal was to produce success stories and to understand for ourselves what we need to do in high-performance data analytics and deep learning. We learned a lot about the challenges and traps, and now we have gone a step further. We are now moving forward and identifying gaps that could be addressed by bringing in AI solutions. It's a process of learning by doing."
Although HLRS is still at an early stage in applying artificial intelligence methods, projects such as HiDALGO (see page 32) have begun providing a framework for exploring how AI could help scientists to identify the most meaningful parameters within large data sets. By focusing on pilot studies related to migration prediction, air pollution forecasting, and the tracking of misinformation over social media, researchers are working not only on the theoretical dimensions of this computer science challenge, but also directly on complex global challenges where high-performance computing and AI could have an impact.
At the same time, HLRS looks forward to making its new AI computing platform available for other researchers interested in exploring how these new approaches could offer new opportunities in their own fields.
"For many people AI is a Swiss Army Knife that will solve every problem, but it's not like that," Koller explains. "The interesting thing now is for people to find out what kind of problems can be solved using AI and how HPC can support these advances. We want to focus on real challenges from real life to show how you can best apply these new kinds of synergies." *(CW)*

NEWS HIGHLIGHTS

© IZKT, Nora Heinzelmann

HLRS Participates in Stuttgart's First-Ever "Smart and Clever" Science Festival

HLRS had an active role in Stuttgart's first Smart and Clever science festival, a weeklong event that spotlighted the importance of advanced scientific research and technology development for the city and surrounding region. Prof. Michael Resch, Director of HLRS, joined Baden-Württemberg Minister for Science, Research and Art Theresia Bauer and experts in the field for a public discussion about artificial intelligence. Dr. Andreas Kaminski, leader of HLRS's Department of Philosophy of Science and Technology, participated in a panel discussion after a public viewing of Friedrich Dürrenmatt's classic play *The Physicists* at the Stuttgart State Theater, focusing on themes surrounding ethical responsibilities of scientific research and disruptive technologies. The HLRS visualization department also presented an exhibit in Stuttgart's city hall that demonstrated how they are using virtual reality to improve public participation in city planning. *(EG)*

Mobility Living Lab Wins Award for Emissions-Free Campus Concept

In response to a competition organized by the State of Baden-Württemberg to develop concepts for an emissions-free campus, HLRS visualization staff partnered with researchers from the University of Stuttgart's Institute for Road and Transport Science and Institute for Internal Combustion Engines and Automotive Engineering in a project called MobiLab. The visualization team developed a 3D model of the university's Vaihingen campus and set virtual modes of transportation — including cars, buses, and foot traffic — in motion. The software allows users to adjust parameters and observe the effects of their changes in real time. The team presented the technology in September at the Mercedes Benz Museum in Stuttgart as part of an event titled "Vision Smart City: Future Mobility, Today." In December MobiLab was named a winner in the competition, for which the University of Stuttgart received a prize of € 300,000. *(CW)*

HLRS Welcomes Students from the Kulturakademie Baden-Württemberg

On March 6, 19 students visited HLRS in conjunction with the Kulturakademie Baden-Württemberg (Baden-Württemberg Cultural Academy). The event, organized in cooperation with the Association of German Engineers (VDI) and the University of Stuttgart, provided students with an introduction to supercomputing. During a workshop, participants worked with HLRS experts to computationally optimize airflow over an airplane wing and then visualized their results in HLRS's immersive 3D visualization facility, the CAVE. Since 2010, the Kulturakademie Baden-Württemberg has offered talented students enrichment programming in which they can dive deeper into topics in the arts, literature, science, and technology. The goal of the program is to give motivated students the opportunity to cultivate their interests, and to learn about potential career directions they may want to pursue. *(EG)*

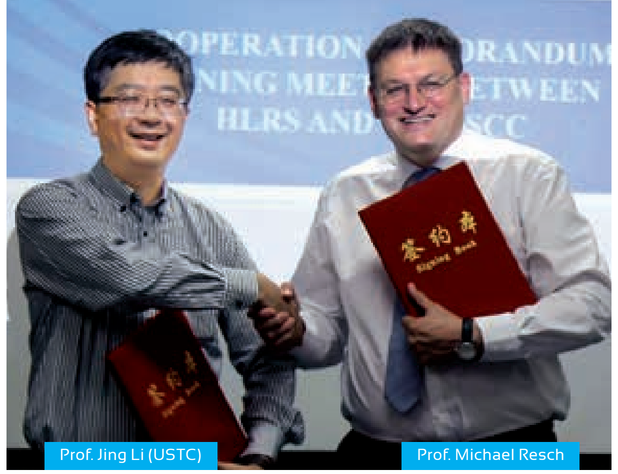

Prof. Jing Li (USTC) Prof. Michael Resch

© USTC-SCC

HLRS Strengthens Collaborations in Asia

HLRS deepened its collaborations with Chinese research institutions, signing memoranda of cooperation with the supercomputing center of the University of Science and Technology of China in Hefei (USTC) and the National Supercomputing Center Guangzhou at Sun Yat-sen University (NSCC-GZ). The agreements will facilitate information sharing on topics of common interest through exchange of scientists, periodic meetings, and collaborative research projects. Prof. Ming-Jyh Chen, dean of mechanical engineering of the National Taiwan University of Science and Technology (NTUST) also visited HLRS in May to renew the existing cooperation between HLRS and his institution. NTUST and HLRS anticipate an exchange of students and scientists in the coming years. *(CW)*

HLRS Raises Funds to Support Children's Hospice

In July, a team organized by HLRS and partners participated in the annual Hand In Hand Benefit Run, raising €860 to support the Stuttgart Hospice for Children and Youth. The team placed 11th out of 197 teams, completing a total of 860 laps. High-performance computer manufacturer Cray sponsored Team HLRS, and the money raised will contribute to facilities and programs that support severely ill children and their families. *(CW)*

EXCELLERAT Presents Data Management System at SC19 Conference

In conjunction with EXCELLERAT, the European Center of Excellence for Engineering Applications, HLRS and its partners are developing a new HPC platform for secure data exchange and management. The ultimate goal is to simplify data transfer for HPC users in industry, particularly for small and medium-sized enterprises. The platform includes an easy-to-use graphic web interface for executing calculations, tools for visualizing results, mechanisms for encrypted data transfer, techniques for reducing data quantity to accelerate data transfer, a visual dashboard for cluster management, and a queuing system for production runs, among other features. Members of the EXCELLERAT team presented a prototype of the system and workflow at the SC19 Conference in Denver, Colorado, USA, focusing on a use case from RWTH-Aachen University. *(CW)*

Supercomputing-Akademie Launches New Training Opportunity for IT Administrators

In an effort to promote understanding of HPC use and administration, the Supercomputing-Akademie launched its third course in September: "HPC Clusters: Plan, Build, Run." The course provided essential knowledge concerning operation of cluster systems for IT managers, administrators, computer scientists, and others involved in the coordination of HPC resources in industry. Designed using a "blended learning" approach, it was also organized to enable working professionals to complete the work in parallel with their normal day-to-day activities. Started in 2018, the Supercomputing-Akademie has already hosted other modules on parallel programming and simulation. Other forthcoming modules will address performance optimization, visualization, and data management. *(CW)*

Award-Winning Art Installation Considers Historical Context of Supercomputing

As part of their installation *Image Capital* at the Kunstmuseum Stuttgart, photographer Armin Linke and historian of photography Dr. Estelle Blaschke included large-format scientific visualizations generated at HLRS as well as video interviews with HLRS staffers, who explained how high-performance computing, data storage systems, and visualization are used in research and technology today. Linke and Blaschke positioned HLRS in the context of other examples from the history of data storage and visualization, including ancestors of today's digital tools. The exhibit was presented as part of the competition for the 2019 Kubus.Sparda Art Prize, an award given each year to a distinguished artist who either lives in or has a close relationship to the state of Baden-Württemberg. Linke was named winner of the 2019 prize in May. *(CW)*

Simulations Using LEGO Test Car Crash Modeling

Simulation experts from software company DYNAmore recently collaborated with automotive safety experts from ADAC and visualization experts at HLRS to conduct an experiment to predict the outcome of a real-world crash between two cars built from LEGO. Although the experiment focused on toy cars, the experimental and simulation methods were similar to those used in developing crash tests for real automobiles. The researchers created high-resolution digital models of a LEGO-built Porsche and Bugatti in which each block was composed of thousands of reference points. Researchers compared the results of the virtual and actual crash tests, finding very good agreement between the simulation and experiment, although the degree of destruction was higher in the physical crash test than the virtual one. *(CW)*

3rd iHURT Meeting Spotlights HPC Usage in Industry

To better understand and address the needs of industrial HPC users, SICOS BW and HLRS hosted the third annual Industrial HPC User Roundtable (iHURT) on December 3. The event offered a forum for HLRS and its industrial users to exchange perspectives on the state of the art in high-performance computing, specific computing challenges industrial users face, and how HLRS could help address them. The program included presentations by industrial users of HLRS's HPC systems, considered new approaches for using artificial intelligence in industry, and discussed recent developments at HLRS. Capping the meeting was an open discussion focusing on practical challenges that industrial HPC users face, including training, data management, workflow development and data transfer, access to software licenses, and the integration of HPC and AI. *(CW)*

Collaboration with German Literature Archive to Improve Digital Research

HLRS partnered with the German Literature Archive and the University of Stuttgart's Institute for Literary Studies and Institute for Natural Language Processing to develop a new platform for the management and analysis of data for literary research. Focusing on "born-digital" literature — such as writers' digital estates and texts published on blogs and other websites — the project aims to provide a sustainable and scalable archive for storing, organizing, and accessing digital data; an integrated pipeline of tools for digital literary and linguistic analysis; and a framework for disseminating research findings. The project focuses on the needs of scholars in the digital humanities, a quickly evolving field that uses computational methods to ask new kinds of questions in humanistic research. HLRS will lead the implementation and deployment of the archive's hardware and software infrastructure and will provide expertise to support the adaptation of digital text processing methods to HPC resources. The project is one of four new Science Data Centers sponsored by the Baden-Württemberg Ministry of Science, Research, and Art. *(CW)*

Virtual Reality Applications for Architecture in Rome

HLRS participated in a virtual reality workshop at Sapienza University of Rome also involving students from RheinMain University in Wiesbaden. The meeting continued a collaboration begun a few years ago with fashion company FENDI and Sapienza University. Students developed architectural interventions at locations around the city, including the Palazzo della Civiltà Italiana, the courtyard of the Ministry of Infrastructure and Transportation, and a cloister at Sapienza's Department of Engineering. HLRS visualization specialists made laser scans of the locations and, using a mobile VR projection system, helped the students to present, discuss, and evaluate their concepts with representatives of the participating organizations. As a result of the workshop, a future collaboration will focus on the historic excavation of the Portus Romae, the city's harbor in antiquity. *(CW)*

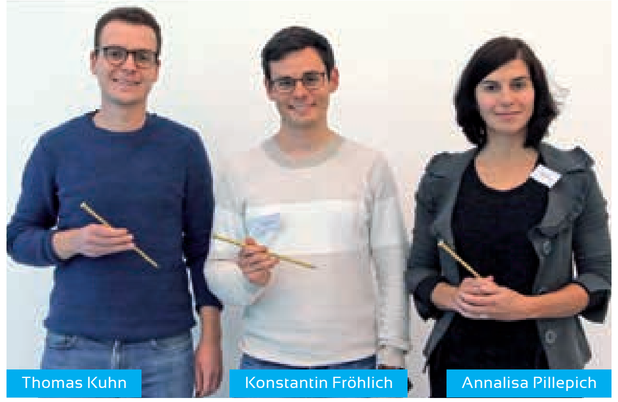

Thomas Kuhn | Konstantin Fröhlich | Annalisa Pillepich

Golden Spike Awards Presented at 22nd Annual Results and Review Workshop

At the conclusion of HLRS's 22nd annual Results and Review Workshop, Dr. Dietmar Kröner, a professor at the University of Freiburg and vice-chairman of the HLRS steering committee, announced the winners of the 2019 Golden Spike Awards, which recognize excellence in research and innovative applications of HPC resources at HLRS. Konstantin Fröhlich of RWTH Aachen won for his work simulating particulate flows, focused on making combustion processes cleaner and more efficient. Thomas Kuhn from the University of Stuttgart won for his work studying fluid dynamics related to airflow over cavities in airplane wings, aiming to make air travel more efficient and quieter. Annalisa Pillepich of the Max Planck Institute for Astronomy and Dylan Nelson from the Max Planck Institute for Astrophysics won for their simulation of galaxy formation in the Illustris project. *(EG)*

Tag der Wissenschaft 2019

On June 29, HLRS staff opened the center's doors to the public as part of the University of Stuttgart's annual open house event. Visitors had an opportunity to tour the HLRS computer room and learn about the capabilities of the center's Hazel Hen supercomputer. Guests could also visit the CAVE, HLRS's immersive 3D virtual reality environment, to take a virtual test drive in the streets of Stuttgart or hang-glide above a virtual Black Forest. Children participated in a scavenger hunt, zigzagging through the building to collect facts about HLRS's research. In addition, HLRS hosted a Chinese-language lecture introducing small children to basic concepts about computing. *(EG)*

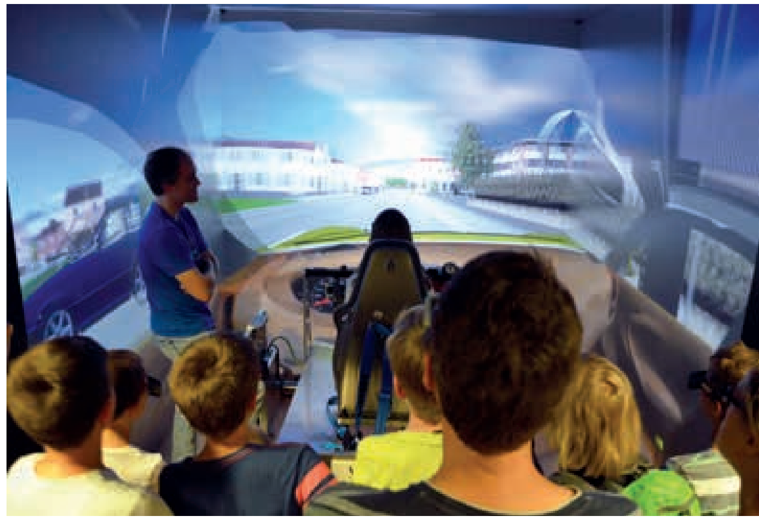

HLRS Certified for Environmental Responsibility

Certifications for environmental management and energy management recognize comprehensive steps the supercomputing center has taken to control its environmental impact.

In November, HLRS received certification for environmental management under the ISO 14001 norm and for energy management under the ISO 50001 framework. The achievement is the result of a multiyear effort to develop and implement a comprehensive sustainability concept that guides the center's operation and future development.

The internationally recognized certifications attest to the steps that HLRS has taken to minimize its environmental impact across the entire organization, and provide frameworks that will guide its efforts to improve its sustainability even more in the coming years.

"Because supercomputing requires significant resources," said Prof. Dr. Michael Resch, Director of HLRS, "we decided years ago that we need to do everything we can to minimize the environmental impact of our activities. These certifications show that we are on the right track to making sure that HLRS remains a sustainable resource for science, engineering, and all sectors of society."

ISO 14001 is an international standard for the development of an environmental management system in organizations. At HLRS, these measures include considering environmental impact in purchasing decisions, minimizing waste, reusing resources, supporting sustainability-oriented scientific research, and promoting sustainability improvements among its peers.

ISO 50001 certification recognizes HLRS's efforts to maximize its energy efficiency. This has involved setting targets for energy usage, tracking energy consumption, and making infrastructure improvements to optimize energy efficiency.

Following an environmental audit, HLRS was also approved to receive certification under the Eco-Management and Audit Scheme (EMAS), the world's most demanding system for environmental management in organizations. In early 2020, HLRS will become the first HPC center in Europe to receive EMAS certification. *(CW)*

Members of the HLRS sustainability team celebrate award of ISO certification.

A New Supercomputer for Artificial Intelligence Applications

A new Cray CS-Storm GPU-accelerated supercomputer brings AI and high-performance computing capabilities together under one roof.

In November, HLRS welcomed an important new addition to its computing profile with the installation of a new Cray CS-Storm supercomputer. The machine will complement the center's other high-performance computers — including its forthcoming flagship system, named Hawk — by addressing the growing demand among its users for a system that is optimized for machine learning and deep learning.

Many supercomputers, including Hawk, are based on an architecture that uses a type of processor called a central processing unit (CPU), which remains the workhorse for computationally heavy tasks such as simulation and modeling. In recent years, however, demand has been growing for systems like the CS-Storm that utilize graphics processing units (GPUs). Originally developed for video gaming and digital graphics, GPUs accelerate certain kinds of iterative parallel computing operations, such as those used in algorithms involving artificial neural networks, which evaluate thousands or even millions of parameters simultaneously.

By bringing these different types of systems under the same roof, HLRS now offers its user community the opportunity to select from and combine a wider range of computing tools and approaches, providing faster solutions for today's most pressing problems

"As we extend our service portfolio with AI, we require an infrastructure that can support the convergence of traditional high-performance computing applications and AI workloads to better support our users and customers," said HLRS Director Michael Resch. "We are now at a point where AI and deep learning have become even more important as a set of methods and workflows for the HPC community. Our researchers will use the new CS-Storm system to power AI applications to achieve much faster results and gain new insights into traditional types of simulation results." *(CW)*

© courtesy of Cray

The Cray CS-Storm is optimized for deep learning and AI methods.

LandesCloud to Offer Secure, Cloud-Based AI as a Service for Industry

A new startup called LandesCloud will give companies access to a state-of-the-art platform for artificial intelligence and high-performance data analytics.

In November, HLRS began a partnership with a new startup called LandesCloud to host data storage and computing infrastructure for a secure, cloud-based platform for artificial intelligence. LandesCloud will address the needs of clients across all industries — from small and medium-sized enterprises (SME's) to large corporations — that have a limited in-house computing infrastructure for AI but require secure data storage and high-performance computing for machine learning applications.

LandesCloud prioritizes data safety — a key concern for industry — during all stages of data transfer, storage, processing, and analysis. Once a company delivers its data to LandesCloud, it is moved to a dedicated, secure LandesCloud server hosted at HLRS that is assigned exclusively to one particular client. Each server is strictly separated from others, providing maximum security.

© Simon Sommer for HLRS

LandesCloud is designed to support collaborative AI projects, including teamworking, crowd-working, and even data science competitions. It could also provide a useful platform for sharing data among organisations, particularly — but not limited to — in the area of manufacturing and supply chain.

LandesCloud's cloud-based solutions include commonly used software packages for machine learning and data analysis, saving users from needing to install, maintain, and operate complex software and costly servers. It also offers workflows for key AI application areas, and can incorporate AI applications for machines and manufacturing processes related to the industrial Internet of Things.

"LandesCloud is unique in that we will be the first to work with a German HPC center to provide such a wide range of AI services," said LandesCloud Managing Director and Founder Stefan Weingärtner. "Our clients will also benefit directly from HLRS's expertise in scalable and complex computing operations." *(CW)*

HiDALGO: High-Performance Data Analytics for Global Challenges

An EU-funded project is developing methods for efficiently analyzing large datasets in an HPC framework, focusing on problems related to migration, air pollution, and the spread of information through social media.

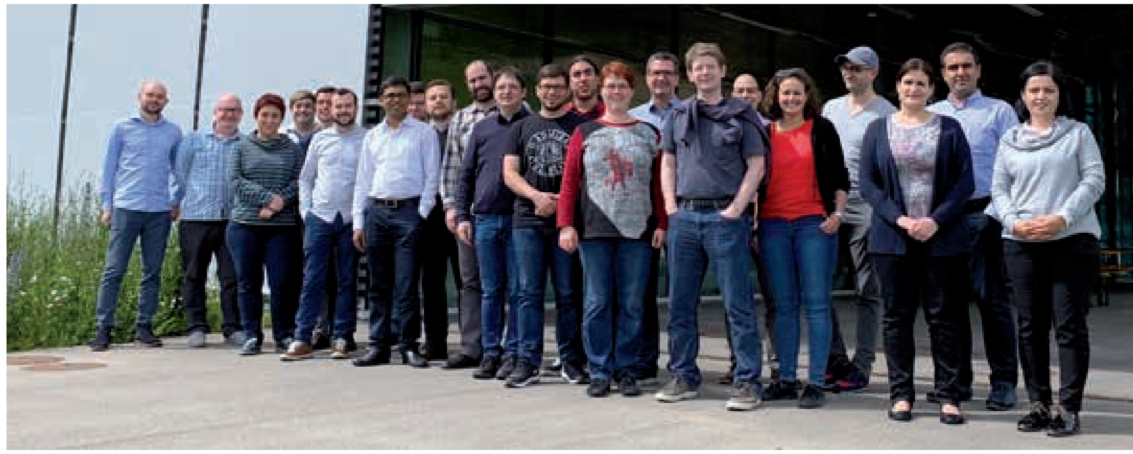

Members of the HiDALGO team during a meeting at HLRS.

As transportation and communications networks have grown, many regional challenges that societies face have become global. This has led to rising demand among governments for computational tools able to provide real-time forecasts they can use to anticipate and manage such challenges. Making such tools available, however, will require new methods for efficiently managing and analyzing the enormous, multidimensional data sets needed to represent such problems.

In 2019 HLRS, in collaboration with ATOS and 11 other institutions from seven countries, began a new research project called HiDALGO aimed at achieving this goal. Funded under the Horizon 2020 Framework Programme of the European Union, HiDALGO is developing computational methods, algorithms, and software for modeling complex, data-rich processes. This includes exploring how artificial intelligence could help accelerate the development of more realistic models of these challenges. HLRS is the technical lead for the project. Simultaneously, HiDALGO supports pilot projects in which these methods could have a practical impact on the management of specific global challenges: forced migration during war, air pollution, and the spread of information through online social networks. In each case, the goal is not only to address a specific problem, but also to produce advances in the science of high-performance computing that will improve the use of high-performance data analytics in an HPC framework.

"HiDALGO is showing that high-performance computing and high-performance data analytics are not only useful for scientific research or optimizing engineering designs," says Bastian Koller, Managing Director of HLRS. "HPC also has an important role to play in addressing some of the most difficult challenges that we as a society are facing." *(CW)*

New Visions for HPC and the Media Arts: An Interview with Matthias Hauser

The Media Solution Center supports collaborations between experts in high-performance computing and leaders in the arts and media industries.

Founded in 2018 by HLRS, together with the Hochschule der Medien and the Center for Art and Media (ZKM), the Media Solution Center Baden-Württemberg (MSC) promotes innovation in the arts and the media industries by facilitating access to technologies and expertise in high-performance computing and visualization.

In 2019 the MSC welcomed Matthias Hauser as its first director. Bringing 25 years of experience in culture management, Hauser has big plans for the center, seeing opportunities to establish new connections between the sciences and technology, the media, cultural organizations, industry, and other sectors. As more arts organizations and companies have begun to become involved, the MSC shows enormous potential to promote innovation across the media and cultural landscapes by bringing together the people and resources necessary to realize visionary interdisciplinary projects utilizing advanced digital tools.

We spoke with Hauser about his plans for the Media Solution Center and his perspectives on the exciting opportunities that exist at the intersection of the sciences, technology, and the arts.

? You've spent most of your career in the field of culture management. How did you end up working at a high-performance computing center?

› In 2015 HLRS approached my artist agency to ask if we would be interested in working together to to realize a project combining technology and the arts. Bringing together HLRS and the Itaú Cultural in São Paulo, we suggested the artist Regina Silveira for the collaboration. She visited HLRS in 2016 and together with the center's scientists and technologists created a virtual reality artwork called *Infinities*. The work attracted interest from the art scene in Stuttgart and in 2018 a large exhibit looking at virtual reality and augmented reality in contemporary art called *Mixed Realities* was presented at the Kunstmuseum Stuttgart under the curation of Dr. Eva-Marina Froitzheim.

Over the course of these activities, Prof. Michael Resch and I had the chance to meet often. From the very beginning we asked ourselves, why couldn't we establish the first eCulture festival in Stuttgart? We developed this idea very intensively over the course of various meetings and conferences, and in collaboration with the city of Stuttgart and the Kunstmuseum Stuttgart, we developed a large project. A convention will take place in conjunction with HLRS's 25th anniversary in 2021, followed by theGATE Festival (Global Arts Technology Environment) in 2022. theGATE is truly a 21st-century initiative that will explore new opportunities at the convergence of science, art, and society.

Personally, I found our collaboration on these innovative projects extremely exciting. When Prof. Resch asked me if I would be interested in taking over the management of the Media Solution Center I agreed immediately, because I found the idea of working onsite at HLRS very interesting. There is no other high-performance computing center in the world that has this sort of a concept attached to it. It's really one-of-a-kind.

? How would you describe the mission of the Media Solution Center?

› Our goal is to make high-performance computing accessible throughout the media, art, and cultural landscape. When it started, the Media Solution Center was focused on making production processes in the film and animation industries more efficient, but since our official founding we are finding that there is interest in many other areas. The Media Solution Center is now working on projects with, for example, the Stuttgart State Opera, the radio and television broadcaster SWR, the Stuttgart Chamber Orchestra, Theaterkunst in Berlin, the ZKM, the Ludwigsburg Film Academy, M.A.R.K. 13 film production, and other partners. Our door is open for anyone who is interested.

From the very beginning I recognized that the Media Solution Center represents the future, because digitalization has put artists and cultural institutions in a position where they are constantly facing new kinds of questions. Artists can sense the pulse of contemporary society, and they engage with subjects that are extremely relevant to how we live today and how society is changing. Many are also already taking advantage of the new opportunities that digital technologies offer, and integrate them into their art practices.

Such trends suggest that we are experiencing the development of a new branch of culture: eCulture. Because of this, the Media Solution Center and the possibilities that it offers are indispensible, even if we are just getting started. New movements in industry, in culture, and in society require the technologies and know-how that scientists can offer in order to bring their concepts to life.

Media Solution Center Managing Director Matthias Hauser

At the same time, it is exciting for scientists at HLRS to become networked with the arts, because it helps them to find better solutions to their own problems. This kind of exchange is becoming increasingly intense because the arts and digital media increasingly flow together, and because new developments are constantly happening in fields like augmented reality and virtual reality. Important topics in the sciences benefit from new kinds of inputs based on the perspectives of artists and others who are active in cultural fields.

The Media Solution Center is in a unique position, because we have exactly the right resources to promote this networking. At HLRS we have a powerful supercomputer, scientists, and experts in visualization technology. At the same time we work closely with collaboration partners like the Hochschule der Medien and the ZKM, one of the most important art museums worldwide. We're not an isolated institution, but apply the competencies of all of our partners.

? How does the Media Solution Center work in practice?

› If you look at the MSC logo, you see a circle. We want to provide a round table where scientists, technologists, visionaries, philosophers, artists, museum directors, and artists can meet and work on solving problems together.
Everyone is welcome to approach the Media Solution Center with his or her particular problem. Then everyone comes together around this table to share their individual knowledge and the questions that they also face. Together, we then work together to determine what resources are necessary to bring something exciting into being. People from different disciplines speak different languages and think about problems differently. Our hope is that such projects will reveal productive, symbiotic relationships that can emerge through the Media Solution Center. At that point we can also provide support and coordination, and determine what can be done.

? How do specialists from such different disciplines experience this kind of cooperation?

› In the first rounds of our meetings with various partners, we heard the same comment several times: "WOW, I would never have thought to contact a high-performance computing center." But when they approach us, they find that a door has opened to completely new possibilities: brainstorming at the highest level, and the opportunity to turn their ideas into reality with us. Through this process new kinds of approaches and pathways emerge.
When we have the courage to break out of our daily working environments and step into a situation where we have more freedom, where everything isn't bogged down in its normal professional context and where we don't quite know what might happen, we enter a free space where new things can come into being. This is what we create at the Media Solution Center and at other locations. We begin to experiment, and sometimes we even run into a brick wall. But out of this experience grows a different kind of self-confidence, a new perspective on what is unique about each of us, and in the end this leads to progress. Mistakes are essential for progress and so we should give ourselves the free space to make mistakes.
If you take a person out of his or her normal context — for example in a collaboration between physicians and a performance group like La Fura dels Baus, Barcelona — it is extremely exciting to watch what happens. The technology and the scientists end up working hand in hand on the development of the performance.
This is also what happens when experts in augmented reality, virtual reality, and simulation apply their technologies and methods outside of their normal work context. It leads to new ideas and completely new perspectives. And this is what the Media Solution Center is there for — to support this networking and to make this vision possible for everyone. *(CW)*

The MSC is coordinating a collaboration between HLRS and the Stuttgart State Opera to explore how use of virtual reality could support an upcoming major renovation.

International Workshop Looks at Trends in Industrial HPC Usage

Supercomputing centers are taking a wide range of measures to increase access to supercomputing resources for commercial R&D.

As digital technologies make it easier to collect data, industry is increasingly looking for solutions that require high-performance computing to analyze it. In particular, the movement toward Industry 4.0 means that, in addition to the ongoing need for classical applications of HPC for simulation, there is growing demand for data analytics, visualization, and artificial intelligence across a wide range of industries.

How industry's needs are evolving, and how supercomputing centers can help meet them, was the focus of the International Industrial Supercomputing Workshop 2019, hosted by HLRS on September 25–26. The meeting — the seventh in this series — attracted senior managers for industrial partnerships from leading supercomputing centers in Europe, Asia, and the United States.

Although many companies own modest parallel computing systems, having access to resources at supercomputing centers offers them unique capabilities. These include the ability to test new, more computationally intensive methods, gain insights about more advanced computing systems, and try out pilot projects that can justify the economic benefits of expanding their own internal computing platforms.

At the same time, companies both large and small benefit from the personal support that supercomputing centers offer. At some centers, this includes discipline-specific competence centers staffed with scientists with expertise in particular HPC application areas and the computer science methods that can best support them. Also important is making the systems easy to use and providing excellent technical support.

During the meeting, speakers presented examples of how HPC-enabled technologies have had an impact on industrial research at their respective centers, and described programs they have implemented to increase industrial access and productivity on their systems. *(CW)*

HLRS Helps SWR Develop Music Recommendation Software

HLRS Director Michael Resch
SWR Intendant Kai Gniffke
Baden-Württemberg Minister of Science, Research, and Art Theresia Bauer
SWR Verwaltungsdirektor Jan Büttner

Audio data analysis enabled the broadcasting company to develop a new software tool for identifying similarity between music titles.

Background music is commonly used in film and television to produce a certain mood, but for editors at broadcasting companies like Südwestrundfunk (SWR), finding the perfect selection can be extremely time consuming. Recently, SWR realized that computational methods utilizing artificial intelligence could help their editors to more quickly identify suitable music for specific situations. At a press conference held at HLRS on October 24, the company presented a software tool it created with HLRS's help to automatically identify similarity between pieces of music.

The software, called AIR, offers an easy-to-use graphical interface in which an editor can enter a music title that suits a situation. The tool then searches the SWR music archive for other music titles that are similar. Whereas commercial music streaming platforms like Amazon or Spotify rely heavily on user activity and "likes" for their recommendation algorithms, the new SWR tool analyzes the digital content of the music files themselves, investigating approximately 700 audio characteristics of each file on a mathematical basis. This includes quantifying music characteristics such as tempo, loudness, complexity, silence, and other factors that give each piece its unique character. Once these characteristics are defined for each clip, the algorithm performs a nearest neighbor analysis, ranking other files according to the similarity of their profiles to that of the clip being queried. Because the SWR music archive contains approximately 2 million files, running such a complex analysis on its own computing infrastructure would have required its full capacity for months. HLRS completed a test run of some 230,000 music titles in about a week and a half, a major improvement. Once the analysis data were returned to SWR, the software developers used the results as the basis for its artificial intelligence algorithm. *(CW)*

The Society of Learning Algorithms

At a conference organized at HLRS, philosophers, social scientists, and historians considered implications of the spread of artificial intelligence.

As use of artificial intelligence becomes more widespread, it raises many questions. What do we want AI to do for us and how should society react to the changes it brings? Coming to terms with such questions will require a clear understanding of what learning algorithms are, how they differ from human intelligence, and what opportunities, limitations, and risks they bring.

In a three-day conference organized by HLRS's program in the Philosophy of Science & Technology of Computer Simulation, philosophers, social scientists, and historians of science discussed how the perspectives their disciplines offer could help elucidate such issues. The event — the fourth in a series called the Science & Art of Simulation — focused on epistemological and ethical implications of artificial intelligence, historical perspectives on the rise of AI, and the political consequences of AI models, among other themes.

Learning algorithms increasingly inform decision-making about issues that have major impacts on people's lives. In banks they help evaluate loan applications, for example, while courts have used them during sentencing proceedings to predict risk of future criminality. Conference participants addressed two important questions about this trend. Considering that the precise mechanism through which a machine learning algorithm makes a prediction is impossible to observe or reconstruct, how can we know that its output can be trusted? Moreover, if bias is unavoidable in the construction of learning algorithms, how do we know that the decisions it suggests are in line with our values as a society?

In addition to providing a forum for an informed, critical consideration of learning algorithms, the conference demonstrated that experts based in other disciplines outside of computer science have important roles to play in the development of more reliable and trustworthy AI tools. *(CW)*

New Digital Tools for Urban Planning

Digital twins, interactive apps, and virtual reality tools could improve the design of cities.

As digital technologies have become more powerful and ubiquitous, cities face important questions of how to adapt to new ways of interacting, and of how digital tools could improve city planning and management. On May 29–30, HLRS hosted an international symposium titled "Urban Systems, Global Challenges, Digital Tools" to exchange insights surrounding these key questions.

Held in cooperation with the bwHPC-S5 Competence Center for Global Systems Science, the workshop brought together academic researchers, representatives of city governments and public utilities, citizen activists, and others with experience using new digital tools for city planning.

The event took place in conjunction with the conclusion of the project Reallabor Stadt:Quartiere 4.0, an experiment in using digital technologies to improve citizen engagement in city planning. HLRS's contribution included creating and testing a "digital twin" of Herrenberg, a city south of Stuttgart. Digital twins are digital representations of real-world objects or environments that contain models, simulations, and algorithms describing their physical counterparts, often implemented in virtual reality.

The digital twin of Herrenberg facilitated communication among city managers, architects, community members, and other stakeholders, turning abstract facts and figures into observable virtual activity. During their presentation at the symposium, Reallabor participants shared some practical lessons they learned about the challenges of communicating with city governments and community members.

Other talks focused on other examples of efforts to use digital tools to model urban environments, interactive technologies for gathering data about how people experience their surroundings, and applications for modeling population movements, transportation networks, air quality, and other features of cities. *(CW)*

Simulation of particulate matter concentration and wind flow in the city of Herrenberg as visualized in the HLRS CAVE.

Shining Light on Dark Data

In his PhD research at HLRS, computer scientist Björn Schembera proposes strategies to make valuable research data more productive and long-term data management at HPC centers more efficient.

Simulations running at high-performance computing (HPC) centers produce massive amounts of data. Once a research project is finished, however, potentially valuable data too often ends up abandoned and forgotten, taking up space on long-term storage servers as researchers move on to other topics.

In a 2019 publication, HLRS computer scientist Björn Schembera and philosopher of science Juan Durán characterized such data as "dark data." Just as astrophysicists know that dark matter must comprise a sizable proportion of the universe's mass even if it can't be observed, dark data can fill countless petabytes of storage — unlabeled, unorganized, and unusable by researchers.

The accumulation of dark data at HPC centers presents several problems. For one, creating, storing, and curating large data sets requires sizable funding, including the costs of building ever larger data storage systems and supplying them with power. From a scientific perspective, the virtual disappearance of dark data also means lost opportunities for computational scientists and engineers working on research that would benefit from access to it. Dark data can also pose security or legal risks, particularly in relation to personally identifiable data and data ownership.

"The concept of dark data has been discussed in other contexts," Schembera says, "but we wanted to better understand its unique features within a high-performance computing context. The paper was the first step toward identifying strategies that could minimize its accumulation." In his recently completed doctoral thesis, Schembera proposes several potential solutions for this problem.

Causes and effects of dark data

While pursuing his PhD at HLRS, Schembera simultaneously worked as a member of the Project User Management & Accounting department, which oversees data management at the center. This position has given him a first-hand perspective on how data is produced, saved, and used at the center. Based on this experience, two principal sources of dark data came to light.

In many cases, data goes dark due to missing or difficult to interpret metadata, standardized information about datasets that provide structure. Scientists typically have no time or incentive to tag their data carefully, and often use individual, ad hoc filing systems for organizing data without systematically annotating it. Although this might suffice in the midst of an active simulation project, it often later becomes extremely difficult to reconstruct what the data represent, or to identify connections between it and other related data.

A second source of dark data results when users of HPC systems become inactive. Once a simulation is complete at HLRS, for example, HPC systems store data in a data server, later moving it to long-term tape storage. However, when scientists disengage from the center — for example, when graduate students leave a university and find other jobs — the data remains unclaimed.

Schembera points out that the accumulation of dark data has a number of implications: It costs the center financially to maintain and operate the computing hardware needed to save the data; it raises legal and security risks when personal data is included; and it is inconsistent with FAIR principles (findability, accessibility, interoperability, and reusability), which govern best practices in data management and reuse. Eliminating dark data could therefore improve the operation and scientific productivity of HPC centers in multiple ways.

Björn Schembera

The Scientific Data Officer

Because academic users of HPC lack incentives to avoid producing dark data, Schembera argues, high-performance computing centers need to take responsibility for managing the problem.

The paper Schembera published with Durán proposes addressing this issue through the creation of a new kind of administrative position within high-performance computing centers: the scientific data officer (SDO).

Specifically, the SDO would be an expert in data management and HPC tools who moderates among researchers, administrators, and an HPC center's management to ensure that best practices are followed in data management. The SDO's responsibilities would include implementing and maintaining a standardized metadata framework for labeling data that is consistent with FAIR standards, and assisting in the management and retrieval of stored data.

Moreover, the SDO would work to reduce the amount of dark data saved at an HPC center. This could include identifying data connected to inactive or deleted users that could be eliminated from the system, evaluating the value of left-behind data to determine whether it should be preserved, and making decisions with regard to data stewardship. To ensure that the position's authority is not abused, they therefore also recommend a code of conduct governing professional behavior to ensure that data is managed ethically.

Automating metadata curation

Considering the enormous amount of data being generated at an HPC center like HLRS, organizing it through metadata is a formidable task for researchers or for a potential SDO.

In his dissertation, Schembera addresses this challenge by introducing a metadata model called EngMeta, which specifies a standardized framework for categorizing and organizing research data in computational engineering. He also extends this framework by developing software to automate metadata extraction. Although currently such a tool would need the support of an SDO or researcher to identify significant discipline-specific keywords, he suggests that it could simplify the often tedious process of metadata management as an automatic part of simulation workflows.

Ultimately, Schembera sees potential in these proposals to improve the productivity and efficiency of HPC centers on multiple levels. Reducing the amount of dark data that is produced and stored could make computing centers more economically efficient and, considering the power requirements of keeping large computer servers running, more environmentally sustainable. Archiving the right kinds of data from past simulations in a more organized and accessible way also holds the potential to make data more scientifically productive. *(CW)*

User Support Staffers Get Closer to Codes

More personalized support and mentoring positively impacts users' productivity.

As part of its most recent funding renewal in 2017, the Gauss Centre for Supercomputing embarked on a project to unify and enhance user support activities at Germany's three national supercomputing centers. In November, HLRS and its partner centers formally rolled out this new framework, which is designed to promote hands-on mentorship by HPC center user support staff.

At HLRS, this has meant improving frameworks for enabling staff members to work closely with users and develop long-term relationships with them. In one example of this close cooperation, long-time users from the Institute for Aerodynamics and Gas Dynamics (IAG) at the University of Stuttgart significantly improved the performance of their code on the Hazel Hen supercomputer, and have been preparing for the arrival of Hawk, HLRS's next-generation supercomputer, which will come online in 2020.

The SEAL project at IAG, funded by the Baden-Württemberg-Stiftung, is focused on understanding how to best quantify uncertainties that arise in simulations of turbulence and acoustics. To do so, the researchers must calculate their simulations at a huge range of scales, running many iterations of a similar simulation with slightly different inputs. This approach means that they not only need access to HPC resources but also need their simulations to run as efficiently as possible.

"The courses and other programs that HLRS offers to support users — such as its annual optimization workshop and MPI courses — allow us to bring our own code to be analyzed and optimized, which is important for us," said Thomas Kuhn, a graduate researcher in the SEAL group at IAG. When the team brought its FLEXI code to a recent workshop, they worked with HLRS staff to achieve significant performance increases in the team's 1D (40 percent speedup) and 3D (two-fold speedup) simulations.

This kind of success is a result of the team's personal relationship with HLRS staff. Specifically, they have long worked with HLRS user support staffer Philipp Offenhäuser. "We have a close collaboration with Philipp for optimizing the FLEXI code," said Jakob Dürwächter, another graduate researcher working on the SEAL project. "He has supported us directly in optimization efforts, particularly during these workshops, and if he cannot help us, he knows where to send us to get the answers we need."

In addition, the team was offered access to the early-access test nodes for Hawk at last year's Results and Review Workshop, giving them a head start on programming their code on the next-generation architecture at HLRS.

Offenhäuser and the rest of HLRS's user support staff want to build out their capacity to better personalize the assistance they are giving to users, especially as HLRS transitions to Hawk.

"We are always trying to provide users with individualized support and mentoring to make the most of their time on our resources, and as Hawk comes online, we have been working hard to ensure that our user base can take advantage of the machine from day one," Offenhäuser said. *(EG)*

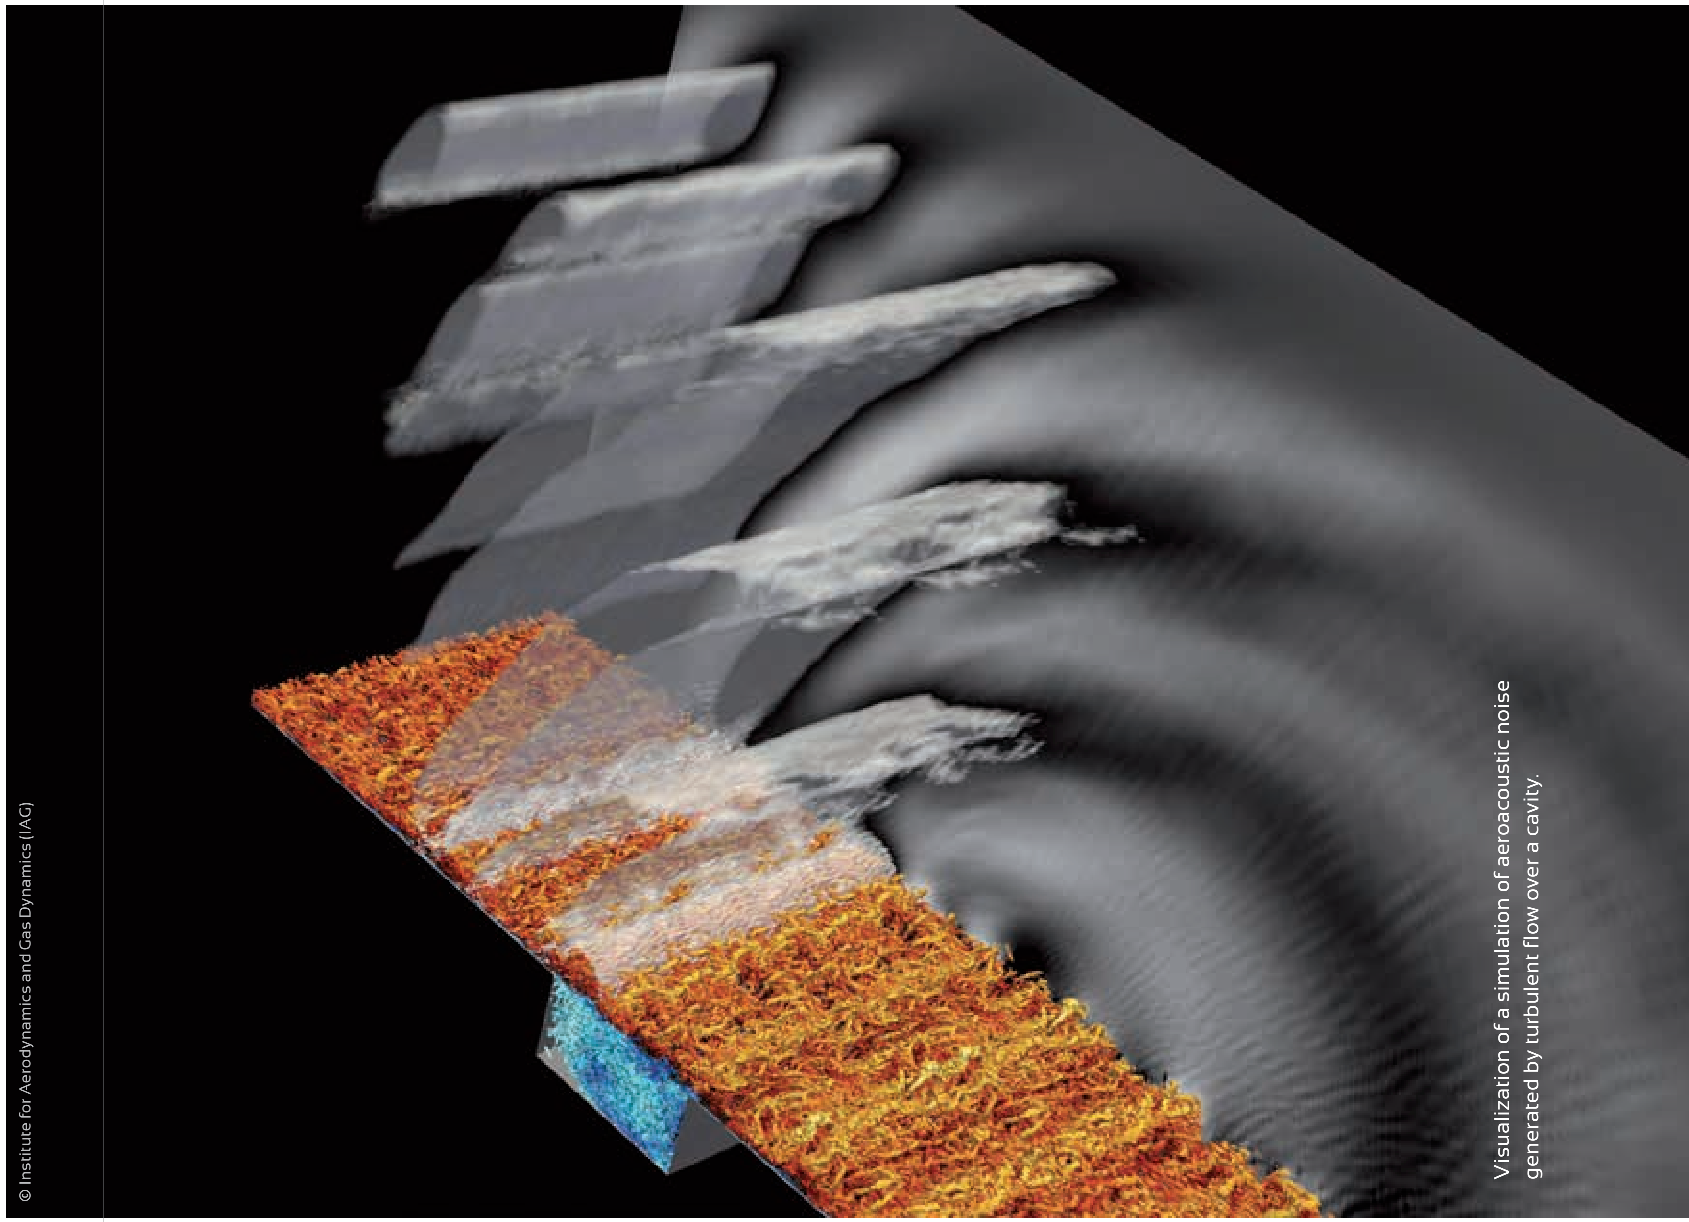

Visualization of a simulation of aeroacoustic noise generated by turbulent flow over a cavity.

HLRS by the Numbers

Staff

150

Visiting Researchers: 28

- 93 SCIENTISTS
- 41 NONSCIENTISTS
- 13 STUDENT ASSISTANTS
- 3 RESEARCH ASSISTANTS

Talks

55

Staff Publications

- 27 PAPERS IN JOURNALS, BOOKS, AND CONFERENCE PROCEEDINGS
- 3 BOOKS

System Usage

- 1.36 BILLION — CORE HOURS PRODUCED
- 127 — USER PROJECTS
- 46 — INDUSTRIAL CUSTOMERS
- 56 — USER PUBLICATIONS

HL RS

Education and Training

36 CONTINUING EDUCATION COURSES
- 838 Participants
- 125 Course-Days

9 SCIENTIFIC WORKSHOPS
- 410 Participants
- 19 Days

16 UNIVERSITY LECTURES
- 2444 Participants
- 34 SWS

Visitors at HLRS

1,928

Third-Party Funds

6,821,698

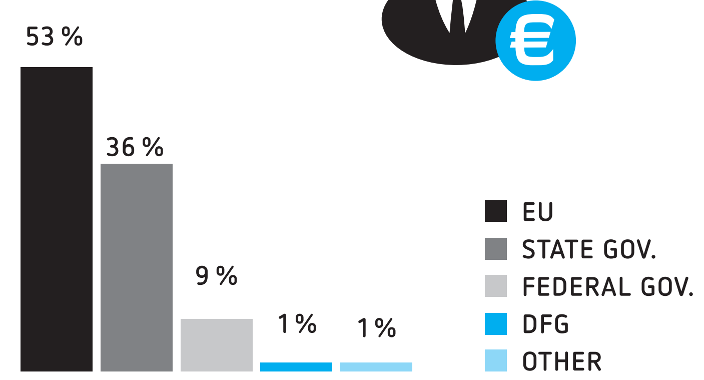

Inspiring Students to Explore Simulated Worlds

High school students learn about computer modeling and collaborate with HLRS researchers for hands-on training.

Simulated Worlds 2019 participants during the award ceremony at HLRS.

Computer modeling and simulation play an important role in our day-to-day lives, but how these tools help make our lives safer, easier, or more efficient is not always clear to those outside the research community.

To bring students closer to the world of modeling and simulation and help inspire the next generation of scientists and programmers, the Baden-Württemberg Ministry of Science, Research, and Art (MWK) started the "Simulated Worlds" program (German: Simulierte Welten) in 2011. From the beginning, HLRS has played an active role in the project, hosting student workshops and providing mentoring and tutoring for students interested in exploring how simulation can help solve problems they encounter in their daily lives.

By participating in the project, students get a better sense of the promises and challenges associated with simulation — where simulation hits its limits, how programmers can avoid or build in bias in simulations, and how simulation has improved technologies and processes students encounter at home or at school.

This year, 10 students worked with HLRS staff to develop projects ranging from using machine learning to predict a film's success before its opening night to computationally modeling the microstructure of human bones. At the year-end event in July, the students presented their projects and answered questions from attendees.

The program's continued success speaks to the growing interest in including more advanced technologies and digital tools in elementary and high school curricula. "We are proud to play an active role in supporting students' interests in modeling and simulation," said Doris Lindner, project coordinator for Simulated Worlds at HLRS. "By introducing high school students to a world of computing they may have been unfamiliar with, we are promoting HPC's role in our everyday lives and encouraging a new generation of students to pursue careers in computational science and information technology."

Simulated Worlds continues in 2020, with several outreach events already planned, including presenting at the educational trade fair, didacta 2020. *(EG)*

Collaboration with Dräger to Design Safer Operating Rooms

Building Information Modeling simulations help researchers better understand airflow in sterile environments.

For roughly 50 years, researchers have used Building Information Modeling (BIM) to make buildings safer, more efficient, or more comfortable for their occupants. Such models enable planners and engineers to better understand a building's characteristics and can reveal strategies for optimizing its various systems.

One common area of study is central air conditioning systems. While air conditioning is nice to have in offices during the summer, it plays a far more important role in specialized environments such as surgical operating rooms in hospitals, where the airflow around a patient can potentially raise the risk of infection.

To address this issue, HLRS visualization experts teamed up with medical device company Dräger to run computational fluid dynamics (CFD) simulations modeling airflow in an operating room.

"In this collaboration we are using a suite of simulations," said HLRS Visualization Department leader Uwe Wössner. Augmented reality allows us to manipulate the physical architecture model in the simulation, and those modifications modify the BIM model, which ultimately controls the input for the simulation. We call this bi-directional BIM, and it allows us to have an immersive simulation that we can modify in near real-time."

To couple these different models, the team uses the Revit software package in conjunction with the open-source CFD code OpenFOAM. The researchers then couple the CFD simulation with fiducial markers — markers to create points of reference in AR simulation — at the doors and air ducts of a room to observe patterns in airflow.

While the research is still in its beginning phases, the team has received positive feedback from its initial test case, working with the University Hospital at Aachen University to optimize the air conditioning systems in their operating rooms. *(EG)*

HLRS demonstrated its collaboration with Dräger at the SC19 Conference in Denver, Colorado, USA, in November.

PhD Graduates 2019

Six young investigators successfully completed their doctoral studies at HLRS in 2019.

Michael Gienger

Optimized Scheduling Strategies for Virtual Machine Deployment in Cloud Computing Infrastructures

Cloud computing is a flexible and efficient approach for offering information and communications technology services. Despite its advantages, however, the baseline technology comes with certain drawbacks. Gienger's work addresses these shortcomings by avoiding overloaded cloud computing environments in general. The results show that significant performance improvements can be achieved by applying optimized scheduling strategies. Methods described by Gienger increased the overall performance of a private cloud computing environment by more than 20%.

Steffen Hagmann

Adaptive Mesh Refinement Using the Finite-Pointset Method / Adaptive Refinement Using the Gridless Finite-Pointset Method

Hagmann worked on the Finite-Pointset-Method, a method suitable for the numerical simulations of vehicles crossing water, making it useful for simulating water management in vehicles virtually. The mesh-free method can numerically solve complex tasks with the necessary degree of physical detail. In the relevant case of a complete vehicle water crossing, the computing time was successfully reduced from just under one week to just over three days and thus by about 40% to 50%. In addition, features of the Finite-Pointset Method were extended by this work, making new and useful functionalities available to FPM users.

Dmitry Khabi

Energy Efficiency of Processors in High-Performance Computing for Engineering Applications

Khabi's work centers on the question of energy efficiency in high-performance computing. His work focused on processor and algorithmic efficiency, ultimately helping improve existing methods for measuring electrical power consumption of various hardware components. Among other things, this improvement helped explore correlations among CPU frequency, degree of parallelization, performance, and power dissipation. The experience gained will help development of a model to describe both computational and electrical characteristics, taking into account the different operating frequencies of the hardware components.

Jens Kouros

Modelling and Simulations of Systematic Risks Through the Use of Probabilistic Programming on the Basis of an Integrative Framework

Kouros investigated systemic risk from multiple perspectives. The first part of his work reviews different conceptions and definitions of systemic risk from different fields, which served as the basis for an overarching framework. He then identified three examples for systemic risk based on this framework. For each, a computer simulation was developed. The first two case studies deal with influenza and epidemics, and with social diffusion and the spreading of fake news. The third develops a rudimentary model of social inequality in which graph structure changes as a result of interactions of two different mechanisms.

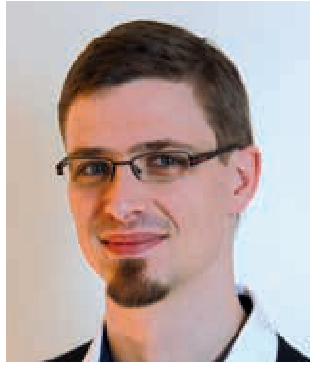

Tim Küstner

The Mixed Mock-Up in Tolerance Management in the Automobile Industry

Through the use of augmented reality in automobile manufacturing, companies created a new, third development platform: the mixed mock-up (MMU). In this work, Küstner develops a methodology for applications of new MMU application fields. The MMU concept of tolerance management, derived from this work as an example, consists both in the expansion of existing processes using MMU as well as developing new methods. The concept behind this work has been extensively checked and evaluated using examples across the development process of an automobile manufacturer.

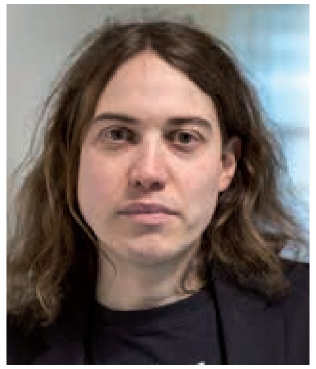

Björn Schembera

Research Data Management in Context: Dark Data in the Simulation Sciences

Schembera extends the concept of dark data to HPC. Dark data comes from missing metadata or inactive users, and his dissertation presents concepts for minimizing such data. He developed a metadata model, EngMeta, to enable data documentation for computational engineering applications, and designed and implemented an automated metadata extraction method to minimize manual metadata tagging. Because such metadata products are useless without corresponding processes, Schembera introduces the role of the Scientific Data Officer. The position applies criteria developed within this dissertation that provide guidance on which data should be stored and for how long.

HLRS Supercomputer Helps Generate First Image of a Black Hole

Researchers from the Goethe University Frankfurt simulated plasma dynamics near the black hole at the center of the M87 galaxy. Their findings supported a breakthrough international achievement.

In April 2019 an international team of researchers shared the first images ever created of a black hole. This landmark accomplishment, involving hundreds of scientists collaborating over many years as part of the Event Horizon Telescope (EHT) consortium, was considered by many to be the top science story of the year. Among the coauthors of the first results was Prof. Dr. Luciano Rezzolla at the Goethe University Frankfurt. Using high-performance computing (HPC) resources at HLRS and the Leibniz Supercomputing Centre (LRZ), his team created a model to describe plasma around the black hole in the middle of the Meissner 87 (M87) galaxy. The investigators also developed a database of synthetic images of a black hole under different conditions that were compared with experimental observations to test their accuracy.

Because black holes reflect no light and M87 is nearly 55 million light years away from Earth, simply pointing a camera and snapping a photo is not an option. Instead, researchers in the EHT collaboration combined observation data from radio telescopes with HPC simulations of phenomena in the vicinity of the M87 black hole — such as the properties of plasma and other materials — to generate images of physical phenomena that are impossible to see with the naked eye.

Integrating this radio astronomy data was a major undertaking. Considering the distance to M87, the black hole's massive size, and the complex fluid and particle dynamics resulting from its gravitational field, HPC was essential in turning the raw data into something that researchers — and in turn, the world — could better understand.

First impressions

One of the Rezzolla team's goals is to understand how the laws of gravity change near the extremely strong gravitational pulls surrounding a black hole. In addition to being strong enough to prevent light from escaping, these forces can distort the behaviour of astrophysical plasmas or other materials as they approach the center. The team developed three in-house computer codes that enable them to model how materials orbit around a black hole as they are being pulled inwards.

The team knew it would need to recreate images by studying how light and plasma bend as they approach a black hole's event horizon, the threshold where gravitational pull is so strong that no object can escape it. This phenomenon creates a "shadow" that envelops the surrounding area.

While the black hole's shadow makes it difficult for researchers to take a picture in the traditional sense, it helps reveal the profound impact that black holes have on their surroundings, giving researchers the opportunity to learn more about black hole properties.

"Modeling plasma near a black hole is a highly non-linear problem that includes a number of instabilities and turbulent flows," Rezzolla said. "Such phenomena are difficult to model under normal circumstances, and these conditions are only amplified near a black hole. You are essentially studying motion happening close to the speed of light in an environment being distorted by an extreme gravitational pull."

Using the Hazel Hen supercomputer at HLRS and the SuperMUC supercomputer at LRZ, as well as its own

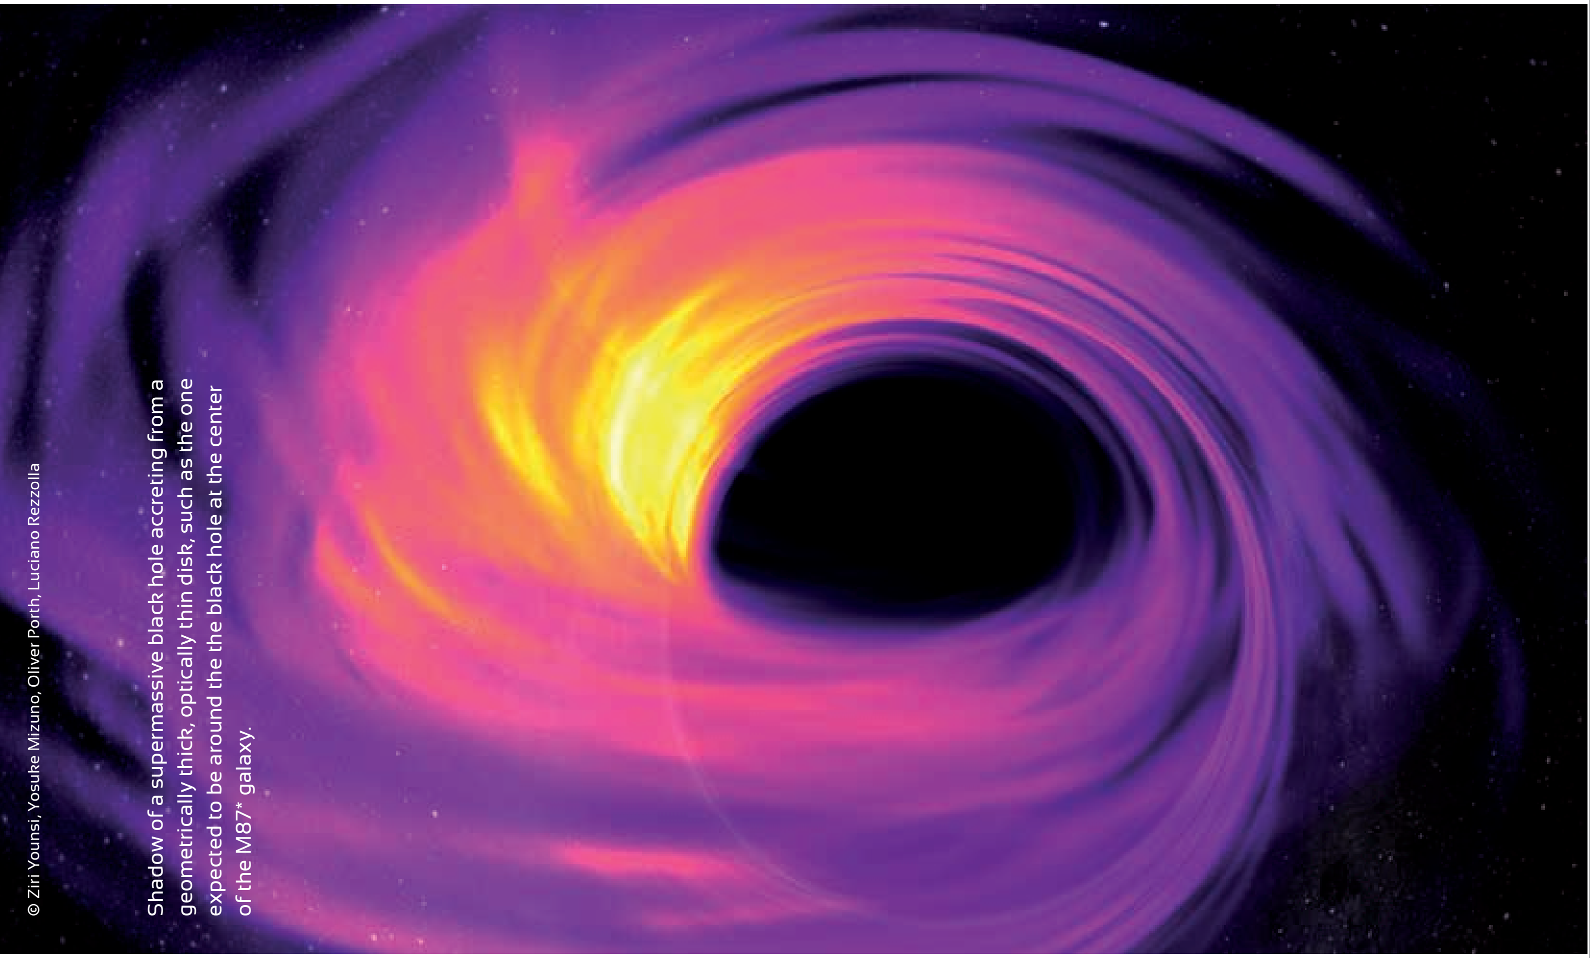

© Ziri Younsi, Yosuke Mizuno, Oliver Porth, Luciano Rezzolla

Shadow of a supermassive black hole accreting from a geometrically thick, optically thin disk, such as the one expected to be around the the black hole at the center of the M87* galaxy.

in-house cluster at Goethe University Frankfurt, Rezzolla's team successfully modelled the plasma dynamics near the center of the M87. About half of the simulations used by the EHT were computed by the Frankfurt group using HPC resources.

The researchers also used HPC to run magneto-hydrodynamic (MHD) simulations that can accurately model electromagnetic properties in materials, such as plasma, to develop a large synthetic image database — imitation black hole images based on their simulations. The 60,000 images in the database represent what a black hole would look like under a wide variety of conditions. By comparing them to the relatively few available observation images, researchers are able to differentiate properties that are unique to the M87 black hole from more general black hole phenomena.

To this end, the team developed GENA, a code based on a genetic algorithm (a class of algorithms inspired by evolutionary processes) that, when comparing the synthetic images with observations, identifies commonalities between the two and evolves them to a new "generation" that contains only the best "genes," ultimately providing a good match between the simulated image and the observed one. Researchers repeat this process over several generations until they find the best match and isolate the synthetic images that are closest to the observations. "It is like entering a stadium with one blurred photo and needing to find that person in the crowd," Rezzolla said. "We are improving our modelling of plasma, but also getting better at our ability to differentiate between stable and fluctuating features in these images."

Beyond the horizon

The team credits its important contribution to the EHT project in part to its successful collaborations with both HLRS and LRZ staff. "This level of staff support is important value added for HPC centers," Rezzolla said.

In the future, the team is looking forward to using next-generation supercomputers to better capture and understand images of other black holes, including the galactic core at the center of our own Milky Way galaxy. The challenge here comes from its proximity and the speed at which galactic phenomena change.

"Looking at our galactic center, there is an additional complication, because the timescale with respect to the image changing is shorter than the time scale that we can record the data," Rezzolla said. "It is like trying to take a picture of something moving all over the place very quickly. All of our physical and technological expertise will be required to handle this new challenge."

(EG)

Simulations Help Researchers Peer Inside Combustion Processes

University of Duisburg-Essen researchers use HPC to model fuel jet flames in unprecedented detail, verifying experiments done by the German Aerospace Agency.

A research group at the University of Duisburg-Essen (UDE) has been using HLRS supercomputing resources to study the critical milliseconds to microseconds when fuel is injected into an engine, observing how the flame ignites and changes over time. Ultimately, the team wants to understand the chemical interactions taking place in molecular, nanosecond detail.

“When we look at combustion, we have several different aims,” said Prof. Dr. Andreas Kempf, head of UDE’s Fluid Dynamics Chair and principal investigator on the project. “We want to minimize the amount of fuel needed, ensure that there are no unburned hydrocarbons in the combustion reaction, and then, of course, minimize the amount of nitrous oxide and carbon monoxide emissions.”

The full picture

By its nature, the small, controlled explosions happening during combustion in an engine are hard to observe — they happen very quickly in an extremely hot, sooty, volatile environment, making it difficult to record or photograph the process.

However, researchers at the German Aerospace Agency (DLR) who collaborate with the UDE team use a method called laser-induced fluorescence to get a sharper image. Essentially, researchers expand a laser beam into a laser “sheet” or “light blade.” At ultraviolet frequencies, these light blades can slice through a combustion reaction, illuminating the many individual particles making up the chaotic, turbulent ignition process. For the most recent collaborative work, the UDE and DLR teams have been studying methane, a relatively well-understood fuel that has relatively simple reactions in comparison with more complex fuels.

Even when taking snapshots at microsecond intervals, though, experimentalists are not able to see the whole picture of the ignition reaction. Think of it like a photographer trying to capture how a building looks from every side, but only getting snapshots of each side at different points in the construction and demolition processes.

Simulation, however, can recreate ignition conditions in 3D and follow the many individual particles at nanosecond intervals, allowing researchers to put their model in motion and observe the many different particles simultaneously. In order to truly optimize combustion processes, researchers must not only see ignition processes playing out in high resolution; they also have to be able to chart how individual fluid particles and reactions in different chemical species influence the combustion reaction as a whole. In fact, when running their most recent simulations, the researchers observed the influence of formaldehyde, a byproduct of the methane ignition reaction, on the process.

“With simulation we have access to every major and minor chemical species that plays a role in the combustion reaction, which is very important for understanding the reaction,” said Eray Inanc, a doctoral candidate at UDE and research leader on the team’s most recent paper. “We can correlate velocities, strain, heat transfer, reactions, and species transfer of major quantities of particular materials in the reaction with minor quantities, such as formaldehyde, which plays a significant role in the overall reaction. Rather than just following one chemical species at a time, simulation allows us to see everything.”

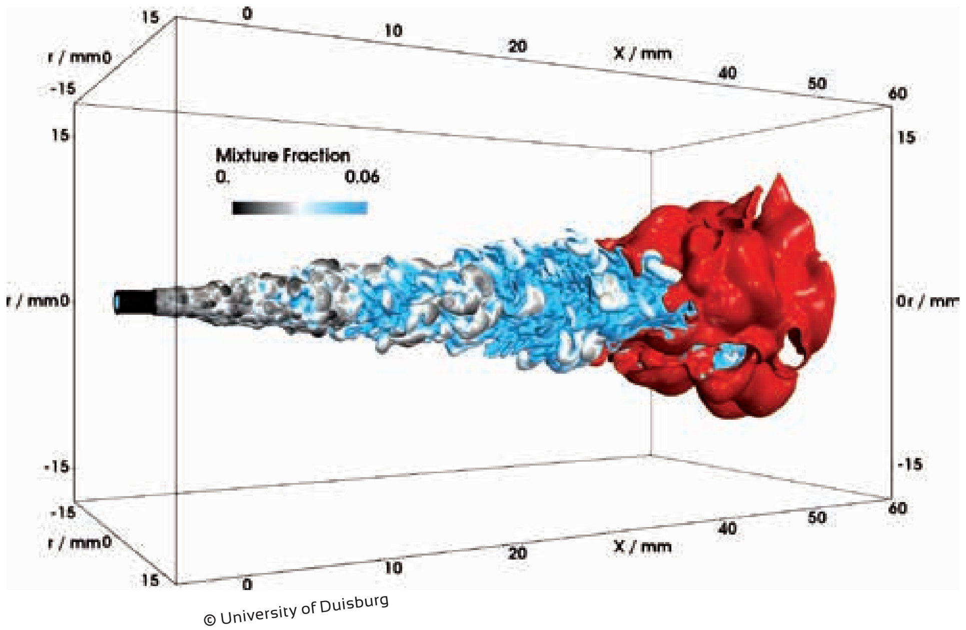

Using direct numerical simulations, the University of Duisburg team simulated a transient jet flame with 40 million compute hours on Hazel Hen.

The team ran two sets of simulations to compare the accuracy and computational costs of two different modelling approaches. The first relied on “tabulated” chemistry, meaning that the researchers generate a table describing the different thermochemical states at a given point in the ignition reaction, such as the amount of fuel or the amount of oxidizer (particles capable of taking new electrons in a chemical reaction). While this approach is computationally cheap, researchers introduce assumptions about the physics in the reaction, making the simulation less accurate.

The second approach, direct chemistry, tracks the many individual reactions occurring at each point in time. While this requires the computationally demanding task of solving transport equations for the chemical species in the simulation, it results in a much more accurate picture of the process. The team found that the difference in accuracy was worth the additional computational cost.

Kempf noted that without access to leading HPC resources such as those at HLRS, his team would be unable to make the same kind of advancements in its field. “In the field of turbulent combustion, you have a big transition happening toward new topics, and it is also a field where most researchers are coming from the world’s top universities,” he said. “To compete with our international contemporaries and competitors, we need access to truly high-end supercomputing power, and we are lucky this is possible in Germany.”

Gaining ground

With the team’s methane research completed, the researchers are looking to next-generation supercomputers to apply their concept to more complex fuels. Kempf and Inanc indicated that the current simulations provide the resolution necessary to get an accurate model of fuel ignition in simple fuels such as methane, but with next-generation computers, the researchers could study things like biofuels and diesel fuels. For methane, the team needed to run about 60 transport equations over the course of its simulations, but more complex fuels would require hundreds of transport equations due to the additional chemical complexity.

The team was able to effectively scale its code to take full advantage of the Jülich Supercomputing Centre’s (JSC’s) old JUQUEEN supercomputer, and effectively scaled on Hazel Hen. As HLRS’s next-generation supercomputer, Hawk, comes online in the first half of 2020, Kempf is confident the team looks forward to taking advantage of the additional computational muscle.

(EG)

Using HPC to Improve Wind Turbine Design

Using a machine learning algorithm and supercomputing, scientists are modelling wind turbine designs to improve energy efficiency.

Over the past several decades, green energy technologies have played an increasing role in nations' energy production. With the growing emphasis on sustainability and the need to fight climate change, green energy coming from solar panels, wind turbines, and geothermal sources will only become more important.

To increase clean energy production, researchers and companies in the green energy sector would like to be able to build larger, more efficient turbines that generate more power.

Until recently, engineers have designed relatively modest wind turbines. Typical turbines are anywhere from 50–150 meters tall, have roughly 120-meter blade diameters, and generate roughly 3 megawatts (MW) of power, or about enough power for 2,000 homes. New designs are getting bigger, though. Engineers are designing wind turbines that have 200-meter blade diameters and are capable of generating 10–20 MW. At such large scales, designers have to make sure these large investments are generating energy as efficiently as possible, including mitigating inefficiency introduced by environmental factors.

To that end, a group of researchers at the University of Stuttgart has been using HLRS's high-performance computing resources to help design more energy efficient wind turbines. "When we are talking about more than 10 megawatts of power, even a one-percent increase in efficiency means a lot of additional energy and a lot of money saved," said Dr. Galih Bangga, a postdoctoral researcher at the University of Stuttgart's Institute for Aerodynamics and Gas Dynamics (IAG). The team's recent research was published in the *Journal of Renewable and Sustainable Energy*.

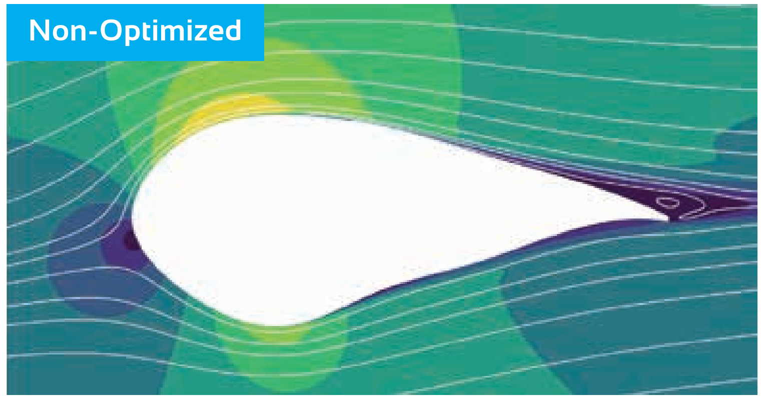

© J. Renewable Sustainable Energy

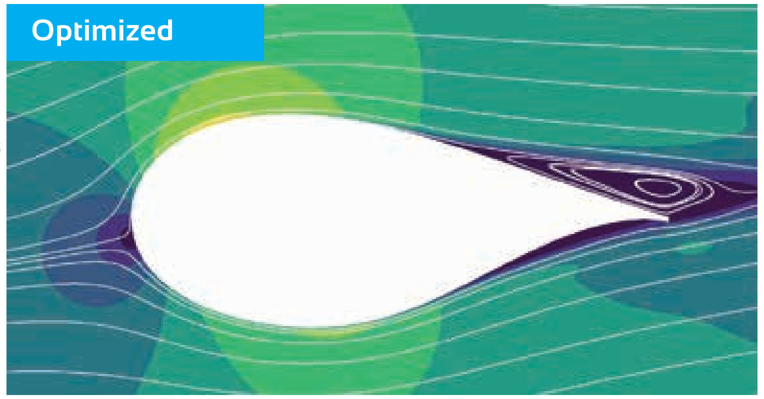

© J. Renewable Sustainable Energy

Total velocity field comparison between the baseline and optimized airfoil. The optimized airfoil shows weaker separation with improved aerodynamic performance.

Survival of the sleekest

As wind turbines get larger in order to generate more electricity, so too must their constituent parts. Specifically, wind turbine blades need to have thicker bases, or airfoils, that attach to the main body and ultimately ensure the turbine's structural integrity. While these thicker airfoils result in a safe, stable wind turbine, they also reduce efficiency due to reduced aerodynamic performance.

The IAG team wanted to learn how to make blades more aerodynamic without compromising a turbine's structural integrity. Unfortunately, building prototypes of many different blade designs and then running experimental tests on all models would be prohibitively expensive and time consuming.

Computer simulation offers a much more efficient and cost-effective way to optimize blade designs. In this case, the researchers virtually created many different variations of airfoils, then ran them through a genetic algorithm — an algorithm based on the same genetic laws that, for example, an agricultural researcher might use to maximize crop yields and resilience by breeding plants with the best traits.

Much like Gregor Mendel detailed how cross pollinating pea plants with the best traits would lead to better pea plants, the team's genetic algorithm takes dozens of turbine blade designs and runs rough turbulence simulations to compare how the models perform, continuing the process until the most optimized candidate emerges. Once the algorithm helps the team identify the best candidate, they turn to HPC to run a higher resolution computational fluid dynamics (CFD) simulation to verify that the algorithm was correct. In the team's paper, they were able to identify and improve the aerodynamic performance on a very thick blade-root airfoil by anywhere from 2.5–7 percent. In practical terms, Bangga noted that the optimization would lead to roughly €195,000 per year in energy savings on each 10-megawatt wind turbine using this design. While the team can run its genetic algorithm on personal computers, the high-resolution CFD simulations needed to verify the model would be impossible without HPC. "HPC is absolutely needed to verify the huge number of databases through the use of high-fidelity simulations," Bangga said. "The access to HLRS resources is a huge benefit to us and our work."

Engineering the future

Bangga has been presenting the team's airfoil findings, as well as other recent simulation work they have done related to controlling airflow at the blade level, at several conferences and workshops. In addition to discussing the airfoil, the team is also discussing the role that active flow controls (AFC) can have in improving wind turbine efficiency.

While making subtle design changes to wind turbine blades' airfoils bring modest, but noticeable changes in energy efficiency, AFC is a more involved and expensive process that can lead to even greater energy efficiency. Active flow controls are similar to the flaps influencing the airflow on an airplane wing — they are controllable parts that influence how air flows around a structure or machine at a local level.

"We want our work to help engineering communities at a variety of scales, not just those that can afford large research and development budgets," Bangga said. "We want to be able to find ways that improve energy efficiency that smaller companies or local governments can afford, but we also want to find ways to maximize energy efficiency as much as possible."

In the next phase of its work, the team wants to expand the level of detail while running the genetic algorithm, allowing them to achieve a higher level of confidence in their optimization recommendations. *(EG)*

Selected Publications 2019

Alon OE. 2019. **Analysis of a trapped Bose-Einstein condensate in terms of position, momentum, and angular-momentum variance**. Symmetry 11:1344.

Babb JF, Smyth RT, McLaughlin BM. 2019. **Radiative association of atomic and ionic carbon.** Astrophys J. 884: 2.

Bangga G, Dessoky A, Lutz T, Krämer E. 2019. **Improved double-multiple-streamtube approach for H-Darrieus vertical axis wind turbine computations.** Energy 182:673-688.

Beck A, Flad D, Munz CD. 2019. **Deep neural networks for data-driven LES closure models.** J Comp Phys. 39:108910.

Beck A, Ortwein P, Kopper P, et al. 2019**. Towards high-fidelity erosion prediction: on time-accurate particle tracking in turbomachinery.** Int. J. Heat Fluid Flow 79:108457.

Bocchini A, Neufeld S, Gerstmann U, Schmidt WG. 2019. **Oxygen and potassium vacancies in KTP calculated from first principles.** J Phys Condens Matter. 31:385401.

Boscheri W, Balsara DS. 2019. **High order direct Arbitrary-Lagrangian-Eulerian (ALE) PNPM schemes with WENO adaptive-order reconstruction on unstructured meshes.** J Comp Phys. 398:108899.

Breuer T, Bussieck M, ... Khabi D*, et al. 2019. **BEAM-ME: ein interdisziplinärer Beitrag zur Erreichung der Klimaziele.** OR-News: das Magazin der GOR. Juli, 6-8.

Cloutier B, Muite BK, Parsani M. 2019. **Fully implicit time stepping can be efficient on parallel computers.** Supercomputing Frontiers and Innovations. 6(2):75-85.

Colmenarez L, McClarty PA, Haque M, Luitz DJ. 2019. **Statistics of correlation functions in the random Heisenberg chain.** SciPost Phys. 7:064.

Daskalakis V, Maity S, Hart CL, et al. 2019. **Structural basis for allosteric regulation in the major antenna trimer of photosystem II.** J Phys Chem B. 123:9609-9615.

Dembski F*, Wössner U*, Yamu C. 2019. **Digital twin, virtual reality and space systems: civic engagement and decision support for smart, sustainable cities.** Proceedings of the 12th International Space Syntax Symposium. 316.1-316.13.

Dembski F*, Wössner U*, Letzgus M. 2019. **The digital twin: tackling urban challenges with models, spatial analysis and numerical simulations in immersive virtual environments.** Architecture in the Age of the 4th Industrial Revolution: Proceedings of the 37th eCAADe and 23rd SIGraDi Conference, Volume 1. 795-804.

Enugala SN, Kellner M, Kobold R, et al. 2019. **Theoretical and numerical investigations of rod growth of an Ni–Zr eutectic alloy.** J Mater Sci. 54(9):12605-12622.

Event Horizon Telescope Collaboration. 2019. **First M87 Event Horizon Telescope Results. I. The Shadow of the Supermassive Black Hole.** Astrophys J Lett. 875:L1.

Fasoulas S, Munz CD, Pfeiffer M, et al. 2019. **Combining particle-in-cell and direct simulation Monte Carlo for the simulation of reactive plasma flows.** Phys Fluids. 31(7):10.1063/1.5097638.

Friedrich A, Gehring P, Hubig C, Kaminski A*, Nordmann A, eds. 2019. **Steuern und Regeln.** Jahrbuch Technikphilosophie 5. Baden-Baden: Nomos.

Fröhlich K, Schneiders L, Meinke M, Schröder W. 2019. **Direct particle-fluid simulation of Kolmogorov-length-scale ellipsoidal particles in isotropic decaying turbulence.** Proceedings of the 11th International Symposium on Turbulence and Shear Flow Phenomena (TSFP11).

Grandin M, Battarbee M, Osmane A, et al. 2019. **Hybrid-Vlasov modelling of nightside auroral proton preciptation during southward interplanetary magnetic field conditions.** Ann Geophys. 37:791-806.

Guevara-Carrion G, Ancherbak S, Mialdun A, et al. 2019. **Diffusion of methane in supercritical carbon dioxide across the Widom line.** Sci Rep-UK. 9:8466.

Ham J, Wössner U*, Kieferle J, Lawrence H. **Exploring the affordances and musico-spatial performance opportunities of a virtual drumming environment.** Architecture in the Age of the 4th Industrial Revolution: Proceedings of the 37th eCAADe and 23rd SIGraDi Conference, Volume 2. 441-448.

Heinen M, Vrabec J. 2019. **Evaporation sampled by stationary molecular dynamics simulation.** J Chem Phys. 151:044704.

Herrmann J, Bangga G. 2019. **Multi-objective optimization of a thick blade root airfoil to improve the energy production of large wind turbines.** J Renew Sust Energy. 11:043304.

Hitz T, Heinen M, Vrabek J, Munz CD. 2019. **Comparison of macro- and microscopic solutions of the Riemann problem I. Supercritical shock tube and expansion into vacuum.** J Comput Phys. (ePub Oct 31).

Horsch MT, Niethammer C*, Boccardo G, et al. 2019. **Semantic interoperability and characterization of data provenance in compuational molecular engineering.** J Chem Eng Data. (ePub Dec 19).

Hummel M, Böhm C, Verestek W, Schmauder S. 2019. **Introducing a method of constructing realistic closed cell nano-porous iron crystals and MD simulations to investigate the influence of the system size on the stability and the mechanical properties.** Comput Materials Sci. 166:150-154.

Jain K. 2019. **Transition to turbulence in an oscillatory flow through stenosis.** Biomech Model Mechan. 19:113-131. (ePub 30 Jul).

Jordis H, Bangga G. 2019. **Multi-objective optimization of a thick blade root airfoil to improve the energy production of large wind turbines.** J Renew Sustain Energy. 11:043304.

Joseph T, Ghorbani-Asl M, Kvashnin AG, et al. 2019. **Nonstoichiometric phases of two-dimensional transition-metal dichalcogenides: from chalcogen vacancies to pure metal membranes.** J Phys Chem Lett. 10:6492-6498.

Kaliszan D, Meyer N, Petruczynik S, Gienger M*, Gogolenko S*, et al. **Comparative benchmarking of HPC systems for GSS applications.** HPC Asia 2019: Proceedings of the International Conference on High Performance Computing in Asia-Pacific Region. 43-52.

Kaliszan D, Meyer N, Petruczynik, Gienger M*, Gogolenko S*. 2019. **HPC processors benchmarking assessment for global system science applications.** Supercomputing Frontiers and Innovations. 6(2):12-28.

Kaminski A*. 2019. **Begriffe in Modellen: die Modellierung von Vertrauen in Computersimulation und maschinellem Lernen im Spiegel der Theoriegeschichte von Vertrauen.** Simulieren und Entscheiden: Entscheidungsmodellierung, Modellierungsentscheidungen, Entscheidungsunterstützung. Springer.

Kaminski A. 2019. **Cyborgisierungen.** Mensch-Maschine-Interaktion. Handbuch zu Geschichte – Kultur – Ethik. J.B. Metzler.

Kaminski A, Glass C. 2019. **Das Lernen der Maschinen.** Mensch-Maschine-Interaktion. Handbuch zu Geschichte – Kultur – Ethik. J.B. Metzler.

Kaminski A*, Wackerbarth A*. 2019. **The assemblage of the skull form. Parental decision, surgery and the normalization of the baby skull.** Monsters, Monstrosities, and the Monstrous in Culture and Society. Vernon Press.

Kempf D, Munz CD. 2019. **Towards an efficient detection of hydrodynamic-acoustic feedback mechanisms in an industrial context.** INTER-NOISE and NOISE-CON Congress and Conference Proceedings. 4112-4123(12).

Kieferle J, Wössner U*. 2019. **Virtual reality in early phases of architectural studies: experiments with first-year students in immersive rear projection based virtual environments.** Architecture in the Age of the 4th Industrial Revolution: Proceedings of the 37th eCAADe and 23rd SIGraDi Conference, Volume 2. 99-106.

Knoch TA. 2019. **A consistent systems mechanics model of the 3D architecture and dynamics of genomes.** IntechOpen. DOI: 10.5772/intechopen.89836.

Lampe H*, Kaminski A*. 2019**. Verlässlichkeit und Vertrauenswürdigkeit von Computersimulationen.** Mensch-Maschine-Interaktion. Handbuch zu Geschichte – Kultur – Ethik. J.B. Metzler.

Lautenschlaeger MP, Hasse H. 2019. **Shear-rate dependence of thermodynamic properties of the Lennard-Jones truncated and shifted fluid by molecular dynamics simulations.** Phys Fluids. 31:063103.

Malviya A, Vrabec J. 2019. **Henry's Law constant of nitrogen, oxygen, and argon in ternary aqueous alcoholic solvent mixtures.** J Chem Eng Data. 65(3):1189-1196.

Mamiyev Z, Sanna S, Ziese F, et al. 2019. **Plasmon localization by H-induced band switching.** J Phys Chem C. 124(1):958-967.

Marquard P, Smirnov AV, Smirnov VA, et al. 2019. **(g-2)SUB-MU at four loops in QED.** EPJ Web Conf. 218(3):01004.

Most ER, Papenfort LJ, Rezzolla L. 2019. **Beyond second-order convergence in simulations of magnetised binary neutron stars with realistic microphysics.** Mon Not R Astron Soc. 490(3):3588-3600.

Müller A, Schippers S, Hellhund J, et al. 2019. **Photoionization of tungsten ions: experiment and theory for W5+.** J Phys B-At Mol Opt. 52(19).

Nägel A, Logashenko D, Schroder JB, Yang UM. 2019. **Aspects of solvers for large-scale coupled problems in porous media.** Transport Porous Med. 130(1):363-390.

Nelson D, Pillepich A, Springel V, et al. 2019. **First results from the TNG50 simulation: galactic outflows driven by supernovae and black hole feedback.** Mon Not R Astron Soc. 490(3):3234-3261.

Neufeld S, Bocchini A, Gerstmann U, et al. 2019. **Potassium titanyl phosphate (KTP) quasiparticle energies and optical response.** J Phys Mater. 2:045003.

Olivares H, Porth O, Davelaar J, et al. 2019. **Constrained transport and adaptive mesh refinement in the Black Hole Accretion Code.** Astron Astrophys. 629:A61.

Pausch K, Herff S, Zhang F, et al. 2019. **Noise sources of lean premixed flames.** Flow Turbul Combust. 103(3):773-796.

Pfeiffer M, Nizenkov P, Fasoulas S. 2019. **Extension of particle-based BGK models to polyatomic species in hypersonic flow around a flat-faced cylinder.** AIP Conf Proc. 2312:100001.

Pillepich A, Nelson D, Springel V, et al. **First results from the TNG50 simulation: the evolution of stellar and gaseous discs across cosmic time.** Mon Not R Astron Soc. 490(3):3196-3233.

Pindzola MS, Loch SD. 2019. **Electron and photon ionization of the BeH molecule.** J Phys B–At Mol Opt. 52(19).

Reitzle M, Ruberto S, Stierle R, et al. 2019. **Direct numerical simulation of sublimating ice particles.** Int J Therm Sci. 145: 105953.

Resch M*. **High performance computing and machine learning.** Proceedings of the Sixth International Conference on Parallel, Distributed, GPU and Cloud Computing for Engineering. Civil-Comp Press.

Resch M*, Boenisch T*, Gienger M*, Koller B*.. 2019. **High performance computing: challenges and risks for the future.** Advances in Mathematical Methods and High Performance Computing. Springer.

Resch M*, Kaminski A*, eds. 2019. **Special Issue: The epistemological significance of methods in computer simulation.** Minds & Machines 29(1).

Saam NJ, Kaminski A*, Resch M*. 2019. **Simulieren und entscheiden.** Simulieren und Entscheiden: Entscheidungsmodellierung, Modellierungsentscheidungen, Entscheidungsunterstützung. Springer.

Sanchez PA, Vögele M, Smiatek J, et al. 2019. **Atomistic simulation of PDADMAC/PSS oligoelectrolyte multilayers: overall comparison of tri- and tetra-layer systems.** Soft Matter. 46.

Satcunanathan S, Zamponi R, Meinke M, et al. 2019. **Validation of a model for acoustic absorption in porous media.** INTER-NOISE and NOISE-CON Congress and Conference Proceedings. 4329-4344.

Schembera B*, Durán JM*. 2019. **Dark data as the new challenge for big data science and the introduction of the Chief Data Officer.** Philos Tech. 33:93-115.

Schembera B*, Iglezakis D. 2019. **The genesis of EngMeta: a metadata model for research data in computational engineering.** Metadata and Semantic Research — 12th International Conference. Springer.

Schneider R*, Kouros J. 2019. **Effiziente Programmierung sozialwissenschaftlicher Modelle: ein Vergleich von NetLogo, Anglican und C++ am Beispiel der Simulation sozialer Diffusion.** Simulieren und Entscheiden: Entscheidungsmodellierung, Modellierungsentscheidungen, Entscheidungsunterstützung. Springer.

Schneiders L, Fröhlich K, Meinke M, Schröder W. 2019. **The decay of isotropic turbulence carrying non-spherical finite-size particles.** J Fluid Mech. 875:520-542.

Schuchart J*, Gracia J*. 2019. **Global task data-dependencies in PGAS applications.** Proceedings of ISC High Performance 2019, 312-329. (Hans Meuer Award finalist).

Schwitalla T, Branch O, Wulfmeyer V. 2019. **Sensitivity study of the planetary boundary layer and microphysical schemes to the initialization of convection over the Arabian Peninsula.** Q J Roy Meteor Soc. (ePub Nov 28).

Shamakina A*, Tsoutsanis P. 2019. **Optimisation of the higher-order finite-volume unstructured code enhancement for compressible turbulent flows.** PRACE White Paper.

Shen B, Ye Q, Guettler N, et al. 2019. **Primary of a non-Newtonian liquid using a high-speed rotary bell atomizer for spray-painting processes.** J Coatings Technol Res. 16(6):1581-1596.

Stephan S, Dyga M, Urbassek HM, Hasse H. 2019. **The influence of lubrication and the solid–fluid interaction of thermodynamic properties in a nanoscopic scratching process.** Langmuir 35: 16948-16960.

Straub S, Forooghi P, Marocco L, et al. 2019**. The influence of thermal boundary conditions on turbulent forced convection pipe flow at two Prandtl numbers.** Int. J. Heat Mass Trans. 144:118601.

Tavelli M, Boscheri W. 2019. **A high-order parallel Eulerian-Lagrangian algorithm for advection diffusion problems on unstructured meshes.** Numerical Methods Fluids. (ePub 19 Jun).

Tchipev N, Seckler S, Heinen M, ... Bernreuther M*, ... Niethammer C*, ... Resch M*, et al. 2019. **TweTriS: Twenty trillion-atom simulation.** Int J High Perform C. (ePub Jan 6).

Waldmann A, Konrath R, Lutz T, Krämer E. 2019. **Unsteady wake and tailplane loads of the common research model in low speed stall.** New Results in Numerical and Experimental Fluid Mechanics XII. Springer.

Wang B, Kronenberg A, Stein OT. 2019. **A new perspective on modelling passive scalar conditional mixing statistics in turbulent spray flames.** Combustion Flame 208:376-387.

Yabansu YC, Rehn V, Hötzer J, et al. 2019. **Application of Gaussian process autoregressive models for capturing the time evolution of mictostructure statistics from phase-field simulations for sintering of polycrystalline ceramics.** Model Simul Mater Sci. 27(8):084006.

Zeifang J, Schütz J, Kaiser K, et al. 2019. **A novel full-Euler low Mach number IMEX splitting.** Commun Comput Phys. 27(1):292-230.

Zhou H*, Gracia J*, Schneider R*. 2019. **MPI collectives for multi-core clusters: optimized performance of the hybrid MPI+MPI parallel codes.** ICPP 2019: Proceedings of the 48th International Conference on Parallel Processing. 1-10.

Zirwes T, Zhang F, Habisreuther P, et al. 2019**. Large-scale quasi-DNS of mixed-mode turbulent combustion.** Proceedings in Applied Mathematics and Mechanics. 19(1). (ePub 18 Nov).

Zirwes T, Zhang F, Habisreuther P, et al. 2019. **Quasi-DNS dataset of a piloted flame with inhomogeneous inlet conditions.** Flow Turbulence Combust. (ePub 8 Nov).

* indicates HLRS staff member

HLRS Books

High Performance Computing in Science and Engineering '18

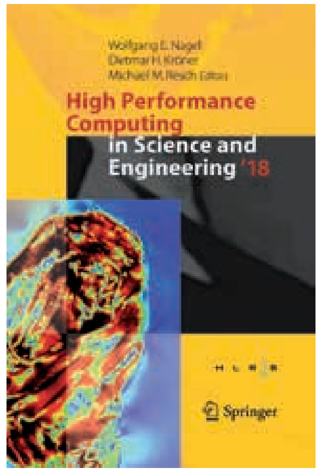

Editors:
Wolfgang E. Nagel
Dietmar H. Kröner
Michael M. Resch

This book presents the state-of-the-art in supercomputer simulation. It includes the latest findings from leading researchers using systems from the High Performance Computing Center Stuttgart (HLRS) in 2018. The reports cover all fields of computational science and engineering, ranging from CFD to computational physics and from chemistry to computer science with a special emphasis on industrially relevant applications. Presenting findings of one of Europe's leading systems, this volume covers a wide variety of applications that deliver a high level of sustained performance. The book covers the main methods in high-performance computing. Its outstanding results in achieving the best performance for production codes are of particular interest for both scientists and engineers. The book comes with a wealth of color illustrations and tables of results.

Tools for High Performance Computing 2017

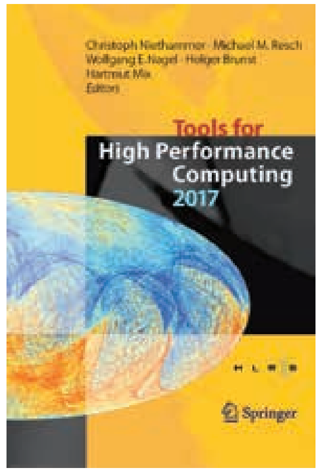

Editors:
Christoph Niethammer
Michael M. Resch
Wolfgang E. Nagel
Holger Brunst
Hartmut Mix

This book presents the proceedings of the 11th International Parallel Tools Workshop, held September 11-12, 2017 in Dresden, Germany. High-performance computing plays an increasingly important role for numerical simulation and modeling in academic and industrial research. At the same time, using large-scale parallel systems efficiently is becoming more difficult. A number of tools addressing parallel program development and analysis has emerged from the high-performance computing community over the last decade, and what may have started as a collection of a small helper scripts has now matured into production-grade frameworks. Powerful user interfaces and an extensive body of documentation together create a user-friendly environment for parallel tools.

Simulieren und Entscheiden: Entscheidungsmodellierung, Modellierungsentscheidungen, Entscheidungsunterstützung

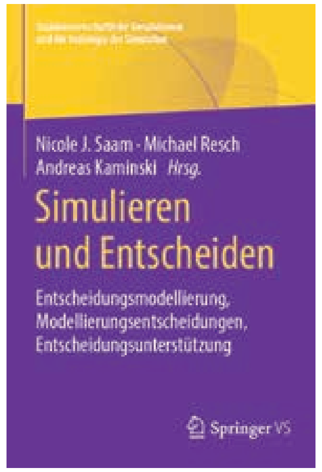

Editors:
Nicole J. Saam
Michael M. Resch
Andreas Kaminski

Simulation and Decision-Making: Decision Modeling, Decisions in Modeling, Decision Support
Simulations increasingly provide orientation for decision making in the context of politics, economics, ecology, or medicine. Doing so means that models must first be developed in which the decisions that go into making a model have a large influence on the simulation results. This book explores this relationship and illuminates computer simulation in the context of decisions from the perspective of sociology, simulation science, and philosophy. Three questions sit in the center of this investigation: How are decisions modeled or simulated? What metatheoretical and methodological decisions must simulation scientists make over the course of a simulation study? How are organizational or political decisions supported by modeling and simulation?

(Book published in German)

Inside Our Computing Room

Hewlett Packard Enterprise Apollo (Hawk)

Installation began in 2019 of HLRS's new flagship supercomputer, Hawk, which will replace its previous system, Hazel Hen. With an expected peak performance of 26 Petaflops, Hawk will at time of installation be among the fastest high-performance computers worldwide and the fastest general purpose system for scientific and industrial computing in Europe. Based on second-generation EPYC processors from AMD, the new system is optimized for the sustained application performance and high scalability required for large-scale simulations, particularly for engineering and the applied sciences. Hawk is due to begin operation in the first half of 2020.

CPU	AMD EPYC Rome 7742, 64 core, 2.25 GHz
Number of compute nodes	5,632
Number of compute cores	720,896
Peak performance	26 petaflops
Memory/node	256 GB
Disk storage capacity	25 PB

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.

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 Memory for pre- and postprocessing: 1 TB
Intel Xeon Gold 6138 @2.0GHz (SkyLake)	Number of nodes: 100 Memory per node: 192 GB Memory for pre- and postprocessing: 1.5 TB
Intel Xeon E5-2660 v3@ 2.6 GHz (Haswell)	Number of nodes: 88 Memory per node: 256 GB Memory for pre- and postprocessing: 2 TB
Intel Xeon E5-2680 v3 @ 2.5 GHz (Haswell)	Number of nodes: 168 Memory per node: 384 GB Memory for pre- and postprocessing: 4 TB
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 2019 the Gauss Centre for Supercomputing approved 7 new large-scale projects (each project requiring more than 35 million core hours) for HLRS's current flagship supercomputer, Hazel Hen, for a total of 878 million core-hours. The Partnership for Advanced Computing in Europe (PRACE) also approved 3 international simulation projects for HLRS, for a total of 145 million core-hours. In total, 127 projects were active on Hazel Hen in 2019 with more than 1.36 billion core-hours used.

System Usage by Scientific Discipline

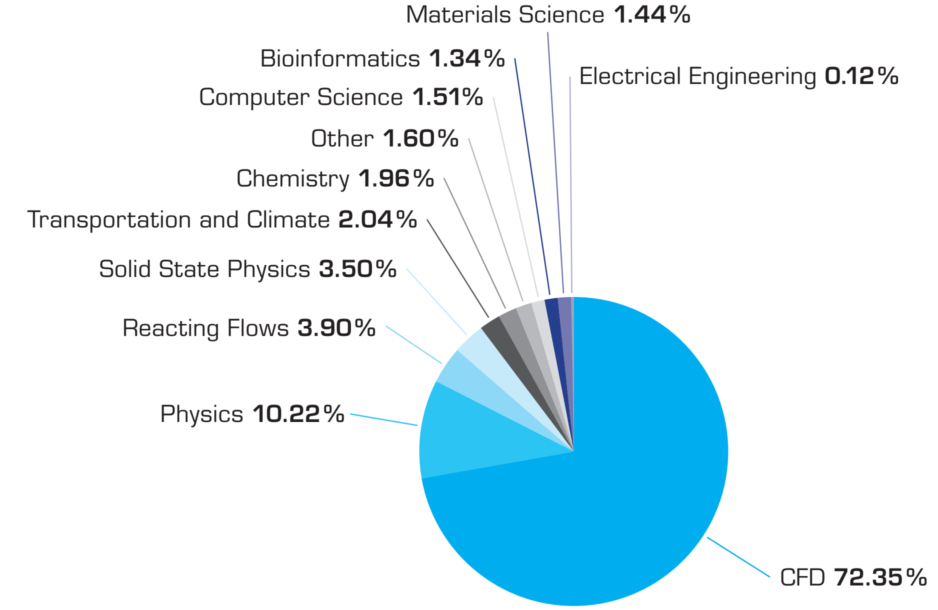

System Usage by State

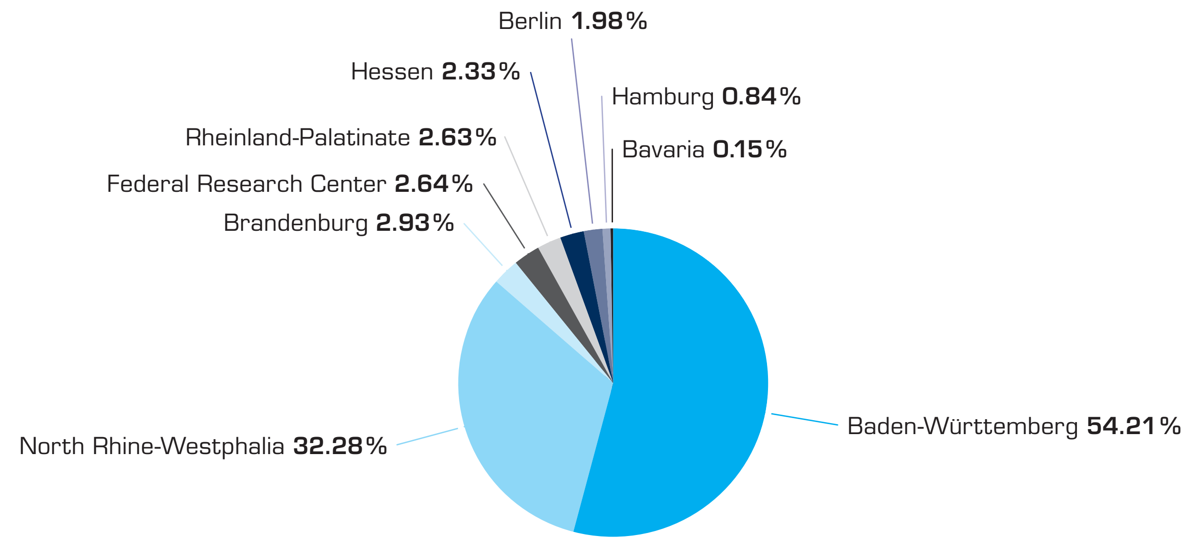

Third-Party Funded Research Projects

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. These activities, many of which are conducted in collaboration with investigators at other institutes and in industry, are designed to address key challenges and opportunities in the field. The following is a list of funded projects that operated in 2019.

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.		
BEAM-ME	**December 2015 - June 2019**	**BMBF**
→ Exploiting the potential of parallel and high-performance computing using distributed memory for high-resolution optimization models in energy system analyses.		
bwHPC-S5	**July 2018 - December 2020**	**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.		
bw Naha 2	**January 2017 - December 2019**	**MWK**
→ Supported the implementation of environmental and energy management systems that enabled HLRS to achieve ISO 14001, ISO 50001, and EMAS certifications, reducing resource usage and improving environmental sustainability.		
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.		
CATALYST	**October 2016 - December 2021**	**MWK**
→ Researches methods for analyzing modelling and simulation data with the goal of implementing a framework that combines HPC and data analytics.		
ChEESE	**November 2018 - October 2021**	**EU**
→ Brings leading European HPC centers, academics, and hardware developers together with SMEs, industry, and public governance bodies focusing on civil protection to prepare 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.		
DIPL-ING	**April 2017 - June 2019**	**BMBF**
→ Researches solutions for efficiently managing the high amounts of data emerging from engineering education programs at the University of Stuttgart.		
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.		
EuroLab-4-HPC 2	**May 2018 - April 2020**	**EU**
→ Aims to establish a European Research Center of Excellence for HPC systems.		
EUXDAT	**November 2017 - October 2020**	**EU**
→ Provides a platform that unites HPC and cloud infrastructures to manage and process high amounts of heterogeneous data. Its focus is to support sustainable development in agriculture.		

Exasolvers 2 May 2016 - April 2019 DFG

➜ The exascale computers of the future are characterized by extreme parallelism. Exasolvers is combining crucial aspects of extreme scale solving, developing methods that scale perfectly and have optimal complexity.

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.

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 - April 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 October 2018 - September 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 - September 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 - December 2020 ESF, MWK

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

OpenForecast September 2019 - August 2020 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.

OSCCAR June 2018 - May 2019 EU

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

PetaGCS January 2010 - December 2019 BMBF / MWK

➜ Supported the procurement and operation of next-generation supercomputers at HLRS from 2011 to 2019. Acquisitions were coordinated by the Gauss Centre for Supercomputing.

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.

Reallabor Stadt:quartiere 4.0	**January 2016 - June 2019**	**MWK**

➜ Explored new methods and technologies for participatory town planning and sustainable development, including developing a digital twin of an urban environment and a novel tool for planning and decision support.

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**

➜ Continues 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.

TalPas	**January 2017 - December 2019**	**BMBF**

➜ Developed a self-optimizing, task-based approach to high-performance particle simulations.

TranSim	**January 2016 - February 2020**	**MWK**

➜ The project Transforming Society – Transforming Simulation explores how computer simulation is transforming science and the worlds of work, knowledge, and values.

VirMuDeKo	**March 2018 - July 2019**	**MWK**

➜ Created a pilot and demonstration framework through which virtual reality and augmented reality can support scientific collaboration in spatially distributed projects.

Funder Abbreviations

BMBF	Federal Ministry of Education and Research
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

HPC Training Courses in 2019

HLRS offered 36 courses in 2019, providing continuing professional education on a wide range of topics relevant for high-performance computing. The courses took place over 125 course-days in Stuttgart and at other locations in Germany and internationally. A total of 838 trainees participated in these activities. For a current listing of upcoming courses, please visit **www.hlrs.de/training**.

Date	Location	Topic
Jan 28-29	Garching	Introduction to Hybrid Programming in HPC *
Feb 4-8	Dresden	Parallel Programming (MPI, OpenMP) and Tools
Feb 18-22	Siegen	Introduction to Computational Fluid Dynamics
Mar 7	Stuttgart	Introduction to Cluster Filesystems
Mar 11-15	Stuttgart	CFD with openFOAM
Mar 25-29	Stuttgart	Iterative Linear Solvers and Parallelization
Apr 1-5	Stuttgart	Cray XC40 Workshop on Scaling and Node-Level Performance
Apr 8-10	Frankfurt	Parallelization with MPI and OpenMP
May 6-10	Stuttgart	Fortran for Scientific Computing *
May 13-15	Vienna	Parallelization with MPI (TtT)
May 14-17	Stuttgart	Advanced C++ with Focus on Software Engineering
May 16-17	Vienna	Shared Memory Parallelization with Open MP (TtT)
May 20-21	Stuttgart	Scientific Visualization
May 22-23	Stuttgart	OpenMP GPU Directives for Parallel Accelerated Supercomputing *
Jun 12-13	Vienna	Introduction to Hybrid Programming in HPC
Jun 25-26	Stuttgart	Cluster Workshop
Jun 27-28	Stuttgart	Node-Level Performance Engineering *
Jul 1-2	Stuttgart	Efficient Parallel Programming with GASPI *
Jul 4-5	Stuttgart	Parallel Programming with HPX *
Jul 9-12	Stuttgart	Advanced C++ with Focus on Software Engineering
Jul 15-17	Stuttgart	Deep Learning and GPU Programming using OpenACC
Aug 19-22	Zurich	Parallel Programming with MPI / OpenMP
Sep 2-6	Siegen	CFD with OpenFOAM
Sep 9-13	Stuttgart	Introduction to Computational Fluid Dynamics
Sep 9-20	Novosibirsk	Parallel Programming with MPI / OpenMP
Sep 16-20	Garching	Iterative Linear Solvers and Parallelization
Oct 14-18	Stuttgart	Parallel Programming Workshop (MPI, OpenMP & advanced topics) *
Oct 24-25	Stuttgart	Scientific Visualization
Nov 6-8	Vienna	Parallelization with MPI (TtT)
Nov 20-21	Vienna	Shared memory parallelization with OpenMP
Nov 25-27	Göttingen	Parallelization with MPI and OpenMP (TtT)
Nov 26-29	Stuttgart	Advanced C++ with Focus on Software Engineering
Nov 28	Heverlee	Shared Memory Parallelization with OpenMP (TtT)
Dec 2-4	Jülich	Advanced Parallel Programming with MPI and OpenMP
Dec 9-13	Stuttgart	Fortran for Scientific Computing
Dec 11-12	Heverlee	Message Passing Interface (MPI) (TtT)

- Parallel Programming
- Computational Fluid Dynamics (CFD)
- Performance Optimization and Debugging
- Data in HPC
- Programming Languages for Scientific Comput
- Scientific Visualization
- Compute Cluster – Usage and Administration

* HLRS is a member of the Gauss Centre for Supercomputing, an official PRACE Training Centre in the EU.

TtT: Train the Trainer Courses

Workshops and Conferences 2019

Date	Location	Partners	Topic
Mar 14-15	Stuttgart	FAU Institute for Sociology	Sozialwissenschaftliche Simulationen & die Soziologie der Simulation
Apr 8-10	Stuttgart		18th HLRS/hww Workshop on Scalable Global Parallel File Systems: "Convergence"
May 28-29	Stuttgart	Reallabor Stadt:Quartiere 4.0	International Symposium: Urban Systems / Global Challenges / Digital Tools
Sep 2-3	Dresden	ZIH, TU Dresden	13th International Parallel Tools Workshop
Sep 25-26	Stuttgart		International Industrial Supercomputing Workshop 2019
Oct 7-8	Stuttgart		High-Performance Computing in Science & Engineering: 22nd Results and Review Workshop
Oct 9-10	Stuttgart	Tohoku University	30th Workshop on Sustained Simulation Performance
Nov 11-13	Stuttgart		Science and Art of Simulation Workshop 2019: The Society of Learning Algorithms
Dec 3	Stuttgart	SICOS-BW	3rd Industrial HPC User Roundtable (iHURT)

Structure

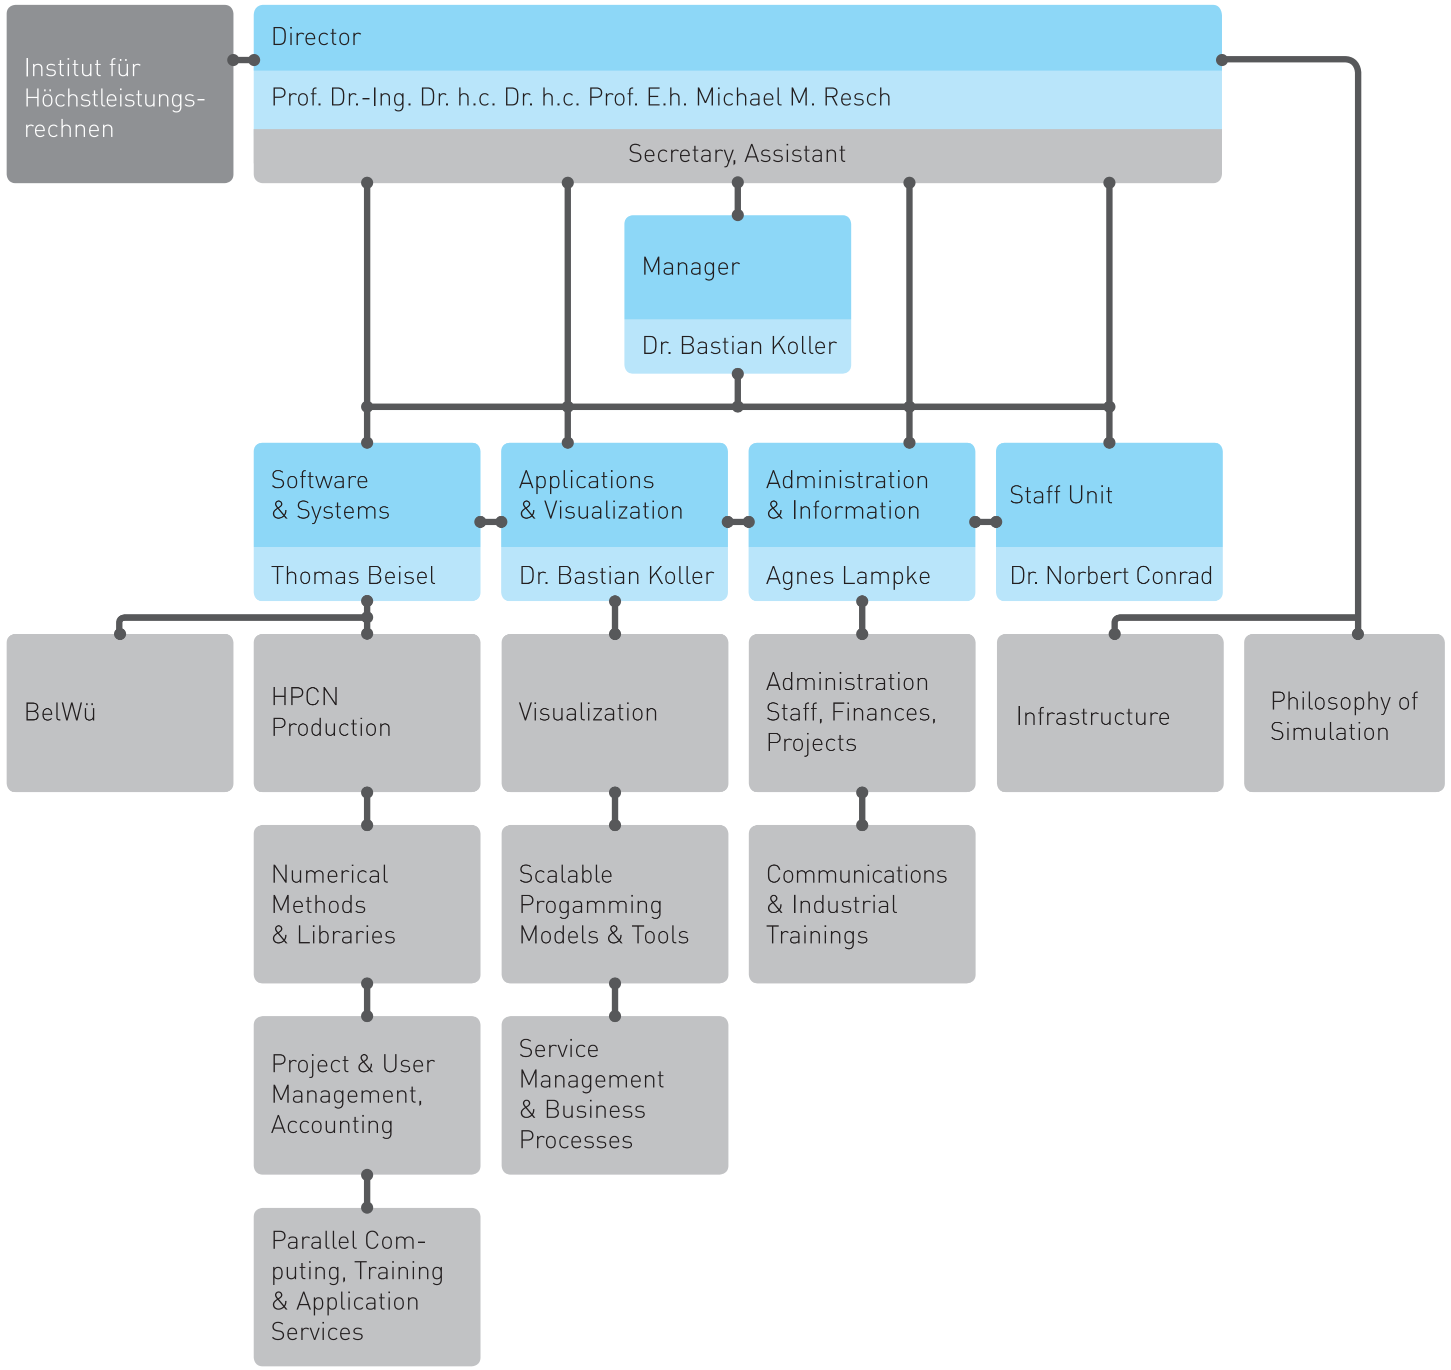

Divisions and Departments

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.

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

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

HLRS developed an immersive, virtual reality-based visualization of the spread of particulate emissions in Stuttgart that makes it possible to interact with simulations at multiple scales, even down to the level of an individual street.

HLRS

High-Performance Computing Center | Stuttgart

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