INNOVATIVES SUPERCOMPUTING IN DEUTSCHLAND
Nº 18-1 · SPRING 2020

INSIDE

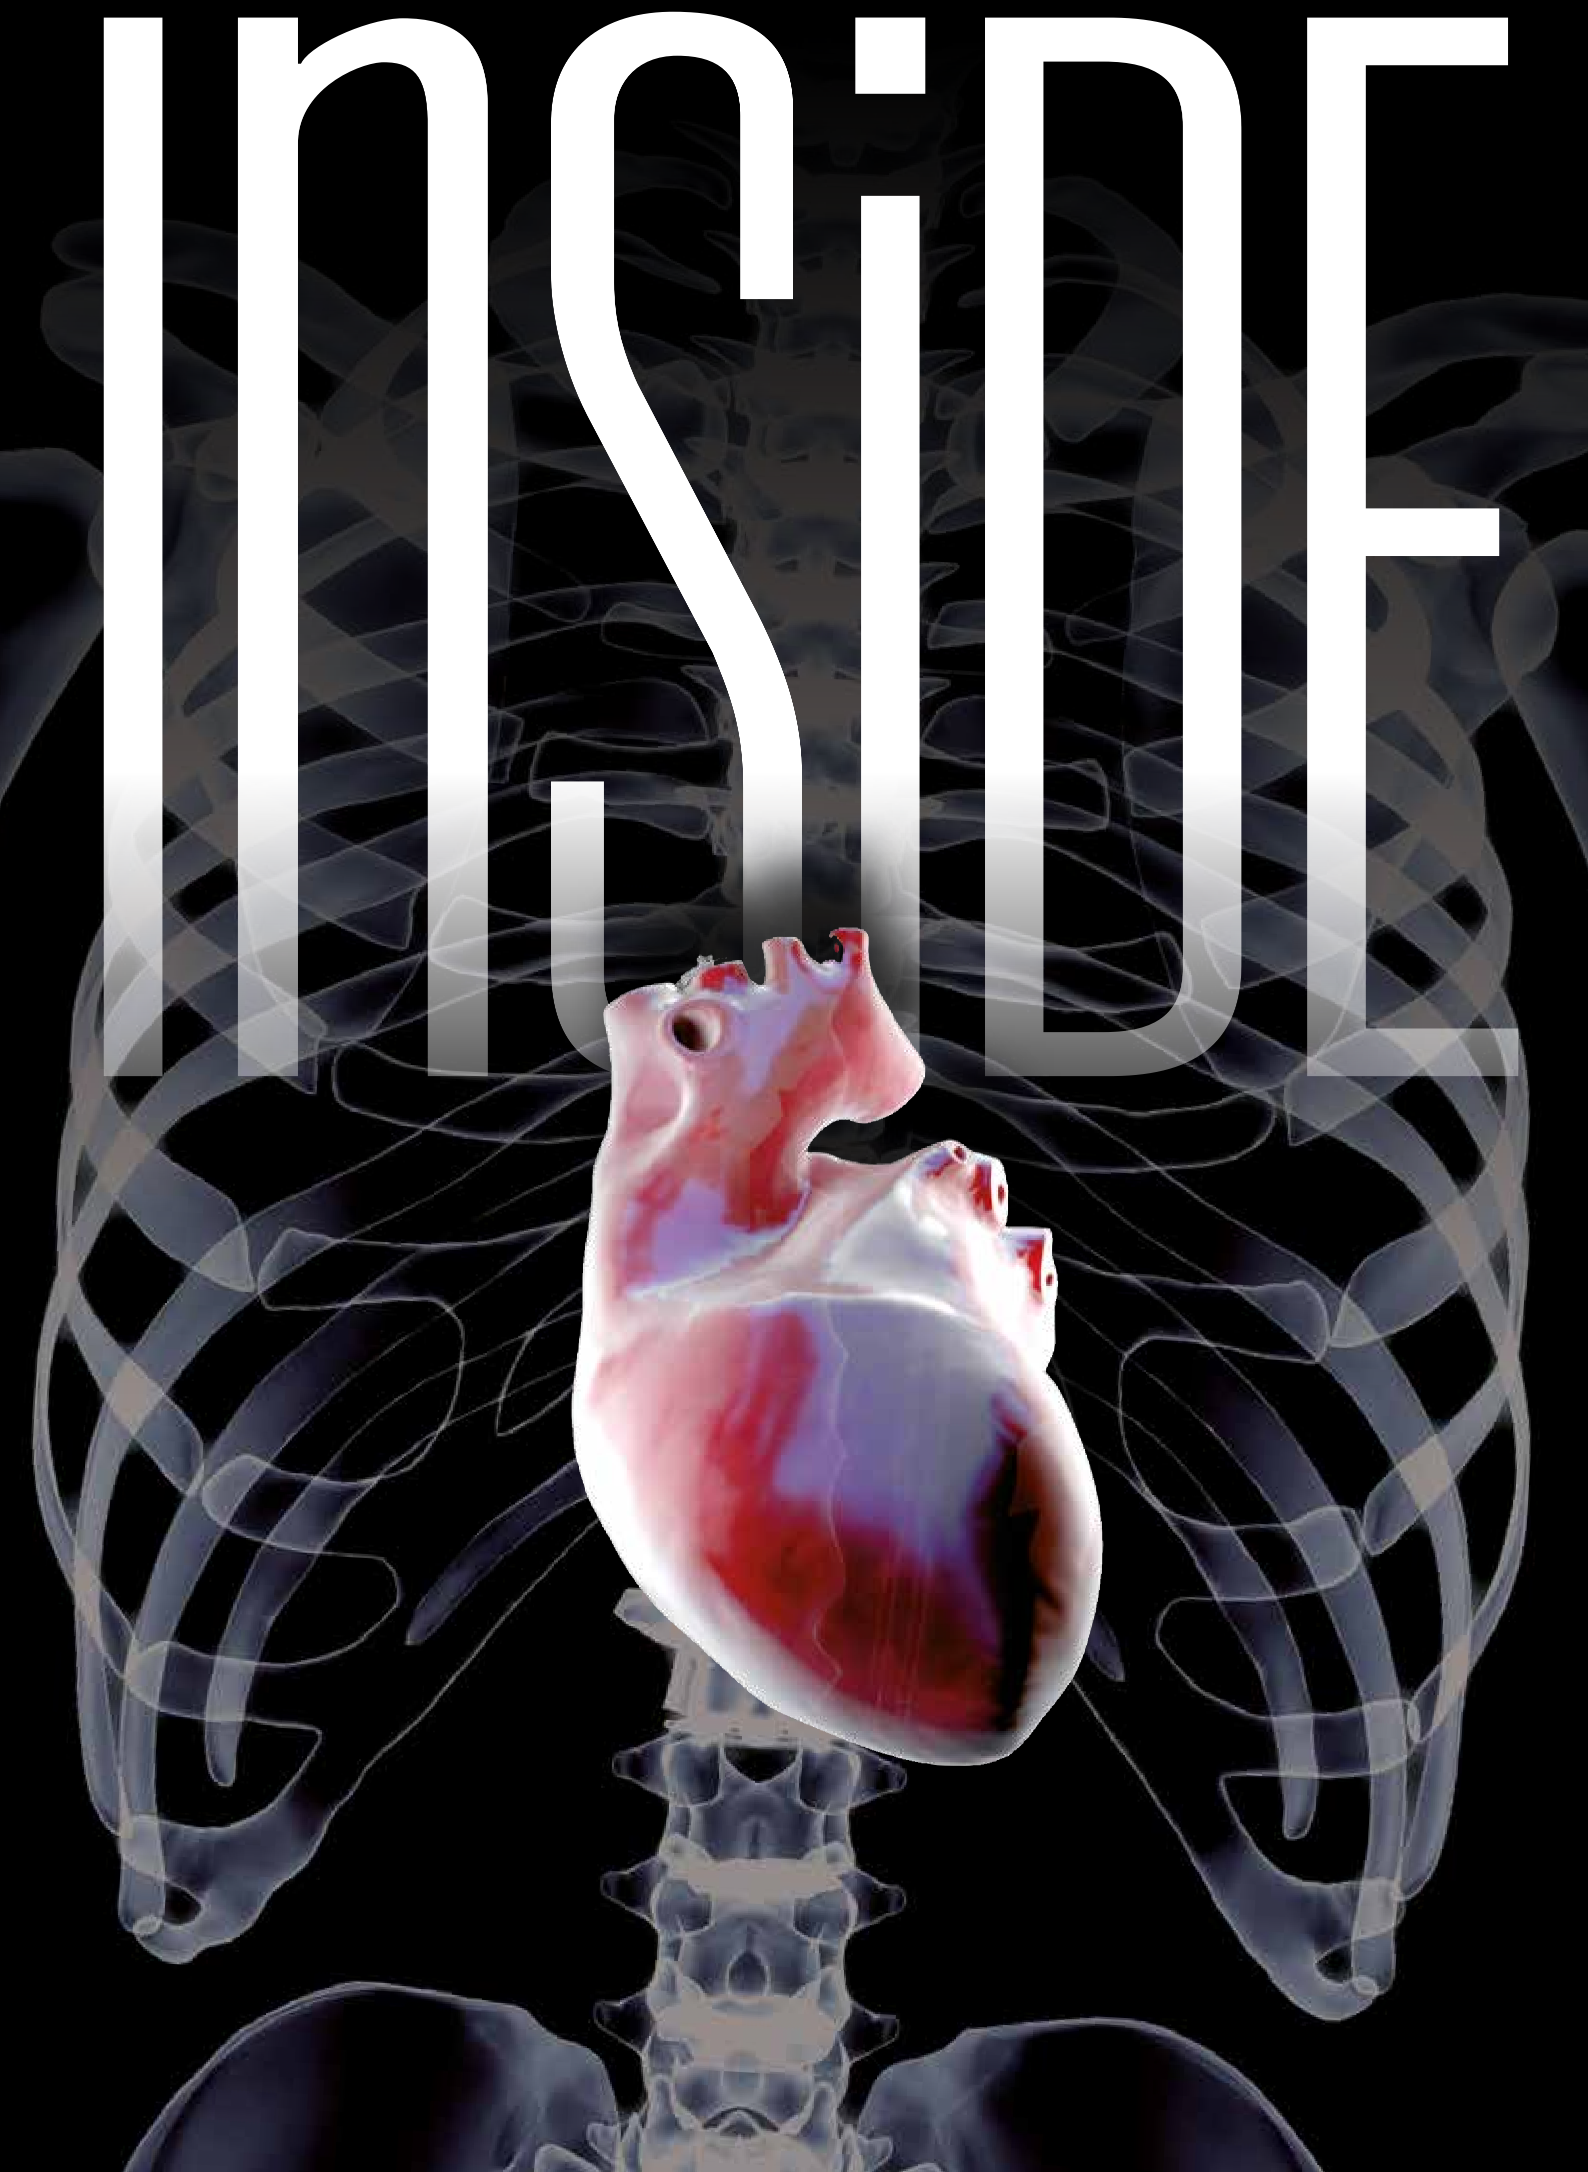

Publishers

Prof. Dr. D. Kranzlmüller
Director
Leibniz Supercomputing Centre

Prof. Dr. Dr. Th. Lippert
Director
Jülich Supercomputing Centre

Prof. Dr.-Ing. Dr. h.c. Dr. h.c. Hon.-Prof. M. M. Resch
Director
High-Performance Computing Center Stuttgart

InSiDE is published two times a year by the Gauss Centre for Supercomputing (HLRS, JSC, LRZ)

Editor-in-Chief
Michael Resch, HLRS
resch@hlrs.de

Editor
Eric Gedenk, GCS, HLRS
gedenk@hlrs.de
e.gedenk@gauss-centre.eu

Production Manager
F. Rainer Klank, HLRS
klank@hlrs.de

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.

If you would like to receive InSiDE regularly, please send an e-mail with your postal address to F. Rainer Klank: klank@hlrs.de

Contributing Authors

Eric Gedenk (*eg*), GCS, HLRS
gedenk@hlrs.de
e.gedenk@gauss-centre.eu

Christopher Williams (*cw*), HLRS
williams@hlrs.de

Dr. Alexander Trautmann (*at*), JSC
a.trautmann@fz-juelich.de

Dr. Olav Zimmermann (*oz*), JSC
olav.zimmermann@fz-juelich.de

Florian Berberich (*fb*), JSC
f.berberich@fz-juelich.de

Susanne Vieser (*sv*), LRZ
susanne.vieser@lrz.de

Sabrina Schulte (*ss*), LRZ
sabrina.schulte@lrz.de

Visit us at:
www.gauss-centre.eu

Welcome to the latest issue of InSiDE, the bi-annual Gauss Centre for Supercomputing (GCS) magazine highlighting innovative supercomputing in Germany. The first half of 2020 has been marred by the novel coronavirus pandemic sweeping across the globe. As providers of HPC resources, we saw our mission as providing researchers expedited access to our HPC systems to help accelerate the search for a cure, development of a vaccine, and epidemiological modelling to help track the spread of the virus. You can find a summary of these efforts on PAGE 10.

Before COVID-19 changed the way that we work and collaborate, we also welcomed the latest addition to the GCS family—HLRS installed Hawk, its next-generation system, in partnership with Hewlett Packard Enterprise, AMD, and DDN. The machine was inaugurated on February 19, 2020 (PAGE 4).

We also continue to support a wide range of research activities across many scientific disciplines, with researchers from the Technical University of Munich have been using LRZ's HPC resources in tandem with artificial intelligence workflows to optimize access and organization of protein databases (PAGE 16). Scientists at the University of Duisburg are using HLRS supercomputing resources in partnership with the German Aerospace Center resources to understand fuel jet flames in combustion processes at the molecular level (PAGE 20). In an effort to accelerate bioengineering techniques, researchers at FZ Jülich and Heinrich Heine Universität Düsseldorf have been using JSC resources to understand how to imbue enzymes with increased tolerance for detergents and solvents for a wide variety of biotechnological applications (PAGE 18).

While the world has shifted further towards video conferencing and email, our collaborations internationally continue to grow. HLRS and GCS were selected to lead a major effort to bring together pan-European HPC competencies as part of the EuroCC project. The project officially begins this fall, with HLRS leading CASTIEL, which will catalog HPC competencies and resources across Europe (PAGE 22). After an extensive collaboration as an associated partner, LRZ formally became a core partner of CompBioMed, a Horizon 2020-funded project aimed at accelerating drug discovery and personalized medicine (PAGE 24). JSC also joined the SANO project to further advanced personalized medicine computationally (PAGE 23).

As the world continues to grapple with the COVID-19 pandemic, the GCS centres remain committed to our core mission while simultaneously supporting research into accelerating a return to normal life. We hope that this issue of InSiDE provides our readers some interesting and valuable information, and hope that our staffs, users, and international partners remain healthy during these trying times.

Prof. Dieter Kranzlmüller
Prof. Thomas Lippert
Prof. Michael Resch

CONTENTS

HLRS CELEBRATES INAUGURATION OF NEW SUPERCOMPUTER “HAWK”

The arrival of a new 26-petaflop high-performance computing System marks the beginning of a new era for advanced computational research in Stuttgart.

Hawk is capable of 26 petaflops, a 3.5-fold increase over Hazel Hen. **© Ben Derzian, HLRS**

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

As it goes online, Hawk — a Hewlett Packard Enterprise (HPE) Apollo 9000 system with a peak performance of 26 petaflops — is among the fastest high-performance computers worldwide and the fastest general purpose system for scientific and industrial computing in Europe. It replaces HLRS's previous flagship system, Hazel Hen, offering a 3.5-fold increase in speed over its predecessor. The new power will enable scientists to tackle dramatically more complex scientific problems in the coming years.

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

Also participating in the ceremony were University of Stuttgart Rector Dr. Wolfram Ressel, HPE Chief Sales Officer Heiko Meyer, AMD Corporate Vice-President of EMEA

Sales Mario Silveira, SSC-Services Director Matthias Stroezel, HLRS Director Prof. Michael Resch, and Edinburgh Parallel Computing Centre Director Prof. Mark Parsons.

New tool for scientific and engineering excellence

Speakers during the event celebrated the new opportunities for academic and industrial research that Hawk will offer, and the important role that the new supercomputer will play in supporting research at the University of Stuttgart and at academic research centers in Germany and Europe, as well as in the surrounding industrial high-tech community in Baden-Württemberg.

In his welcoming remarks, Dr. Ressel explained, "HLRS represents an essential component of the University of Stuttgart's comprehensive strategy with respect to our newly founded Excellence Cluster, 'Data-Integrated Simulation Sciences.' At the same time, however, it is not only essential for our university, but also offers scientists from around the world and experts from industry the opportunity to work together in an interdisciplinary way to use simulation to find solutions for complex problems in a wide range of areas."

Prime Minister Kretschmann welcomed Hawk as the latest tool for cementing Baden-Württemberg as a pillar of the German technological landscape, ultimately ensuring that Germany remains competitive at the international level. "We do not just find ourselves in the midst of a deep technological transformation; we are in the midst of an intensive competition," he said. "This transformation is being driven in multiple technological sectors at once, including digitalization, quantum technologies, molecular biology, and green technologies. Our aspiration is not to run behind this technological revolution, but rather to help shape it in all its technological and cultural dimensions."

Providing an overview of Germany's national supercomputing strategy, State Secretary Meister emphasized the role that high-performance computing (HPC), at Stuttgart and at its partner national supercomputing centers in Jülich and Garching, plays in ensuring that Germany has the tools for cutting-edge research. "High-performance computing is an indispensible key technology in digitalization. It is enabling innovation in science and, what particularly impressed me, in industry." he said.

Baden-Württemberg Minister President Winfried Kretschmann delivers an address during the Hawk inauguration. . © Ben Derzian, HLRS

As director of SSC-Services, a company that facilitates industrial access to high-performance computing, Matthias Stroezel spoke to the importance of Hawk's arrival for the dynamic high-tech sector in Baden-Württemberg. "(Hawk) is the prerequisite for identifying and realizing new leaps forward in innovation. HLRS is one of a kind in that it provides access to such a resource for private industry, including clients of the SSC. What an opportunity!"

Collaborations to develop new technologies

The beginning of Hawk's operation also marks the beginning of a new collaboration between HLRS and HPE to develop new technologies for supercomputing. This includes the development of new software and tools for high-performance computing, for performance optimization, and for artificial intelligence that will be necessary in the coming years to prepare for the arrival of new exascale supercomputers.

"Hewlett Packard Enterprise and HLRS are a Dream Team," said Heiko Meyer. "For this reason we would like to develop

From left to right: State Secretary Gisela Splett, Member of the Baden-Württemberg Landtag Sabine Kurtz, Parliamentary State Secretary Michael Meister, Baden-Württemberg Minister of Science, Research and Art Theresia Bauer, Baden-Württemberg Minister President Winfried Kretschmann, HLRS Director Michael Resch, HPE Chief Sales Officer Heiko Meyer. © Ben Derzian, HLRS

a long-term development partnership with HLRS in which we optimize applications, test future technologies, and bring them to a mature, market-ready state."

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

A golden-age for high-performance computing

Edinburgh Parallel Computing Centre Director Mark Parsons provided a broader perspective about the recent explosion in high-performance computing power and the exciting new opportunities that it offers for science and technology development, looking at applications for weather prediction, genomics research, and engineering. "Many people say, surely everything that can be done in supercomputing has been done," he said. "Actually I would say we're at the end of the beginning. These big systems open up opportunities for science that were simply impossible before. We're going to be able to do things with them we've never managed to do before. We're at the beginning of a golden age of modeling and simulation."

Returning to future applications of Hawk's, HLRS Director Michael Resch also presented a brief overview of ways in which supercomputing can help to address key global challenges. Discussing topics including reducing air pollution, optimizing energy production in wind turbines, and developing new mobility solutions, he explained, "Hawk allows us to better understand the environment, and use these new technologies more effectively. We want to give every citizen an opportunity to be a part of these changes. With Hawk, we can digitalize the future in a sustainable way."

Following the talks, guests were able to visit the new supercomputer and learn more about how HLRS's supercomputing tools and services are having an impact. *cw & eg*

JÜLICH SUPERCOMPUTING CENTRE RESEARCHERS COLLABORATE AND INNOVATE AT THE FRONT LINE OF QUANTUM COMPUTING TECHNOLOGY

New partnerships, hardware acquisitions, and a knowledge base firmly at the intersection of high-performance computing and emerging technologies help Jülich researchers lead on the international stage.

Over the last year, staff members at the Jülich Supercomputing Centre (JSC) have been involved in international research collaborations with major implications for the way humanity thinks about computers. While JSC staff has decades of combined experience working with traditional high-performance computing (HPC) architectures, neither they nor any computer scientist has much experience with quantum computing—a paradigm shift in the computing world that has potential to unseat traditional HPC in certain realms of modelling a simulation.

While quantum computing offers the promise of faster insights into some complex problems with many variables, the technology is still a new, disruptive technology. In an effort to help quantum computing technologies reach their game-changing potential, international cooperation between researchers in academia and private industry has become pivotal—the need to combine expertise in computing technologies, leading HPC architectures, and effective use of quantum computer prototypes will all play a major role in advancing quantum computing technologies.

While quantum computing will ultimately offer advantages over traditional HPC in some areas of research, it is still reliant on HPC to help accelerate system design. To the end, the Quantum Information Processing group at JSC, led by Prof. Kristel Michielsen, has been partnering with leading organizations from industry to scale up quantum computers.

Partnerships and progress with Google

Last year, Michielsen, JSC Director Prof. Thomas Lippert, and Forschungszentrum Jülich Director Prof. Wolfgang Marquardt met with representatives from Google to sign an agreement, which is providing a basis for sharing research results, hosting joint training sessions for researchers interested in using quantum computing systems, and benchmarking results obtained from quantum computing systems. In addition, JSC would leverage its leading HPC systems to model quantum computing systems in order to better understand how these systems work under a range of conditions and configurations.

At the time, Michielsen said, "This is a symbiotic relationship. I think we at JSC benefit from their experience with quantum computing and annealing devices, as well as their connections with industries looking to use this technology. They benefit from us due to our in-house simulation expertise in quantum computing and having access to leading high-performance computing resources. Both organizations share a strong interest in training researchers on how to program for these new compute devices."

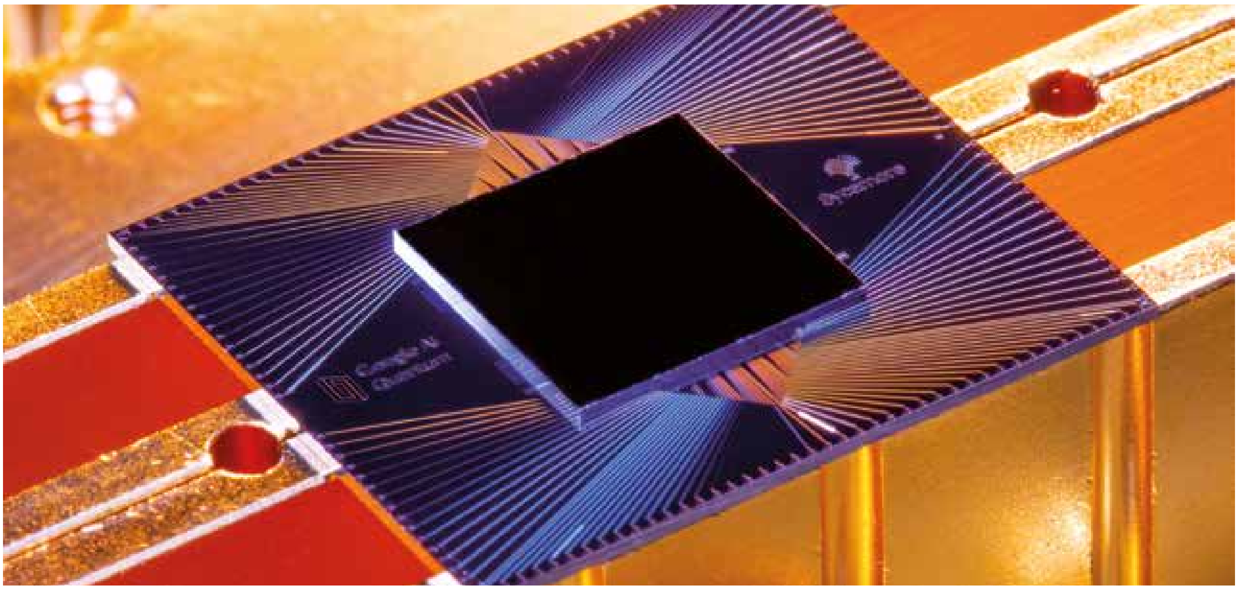

The Google Sycamore quantum processor. **© Google**

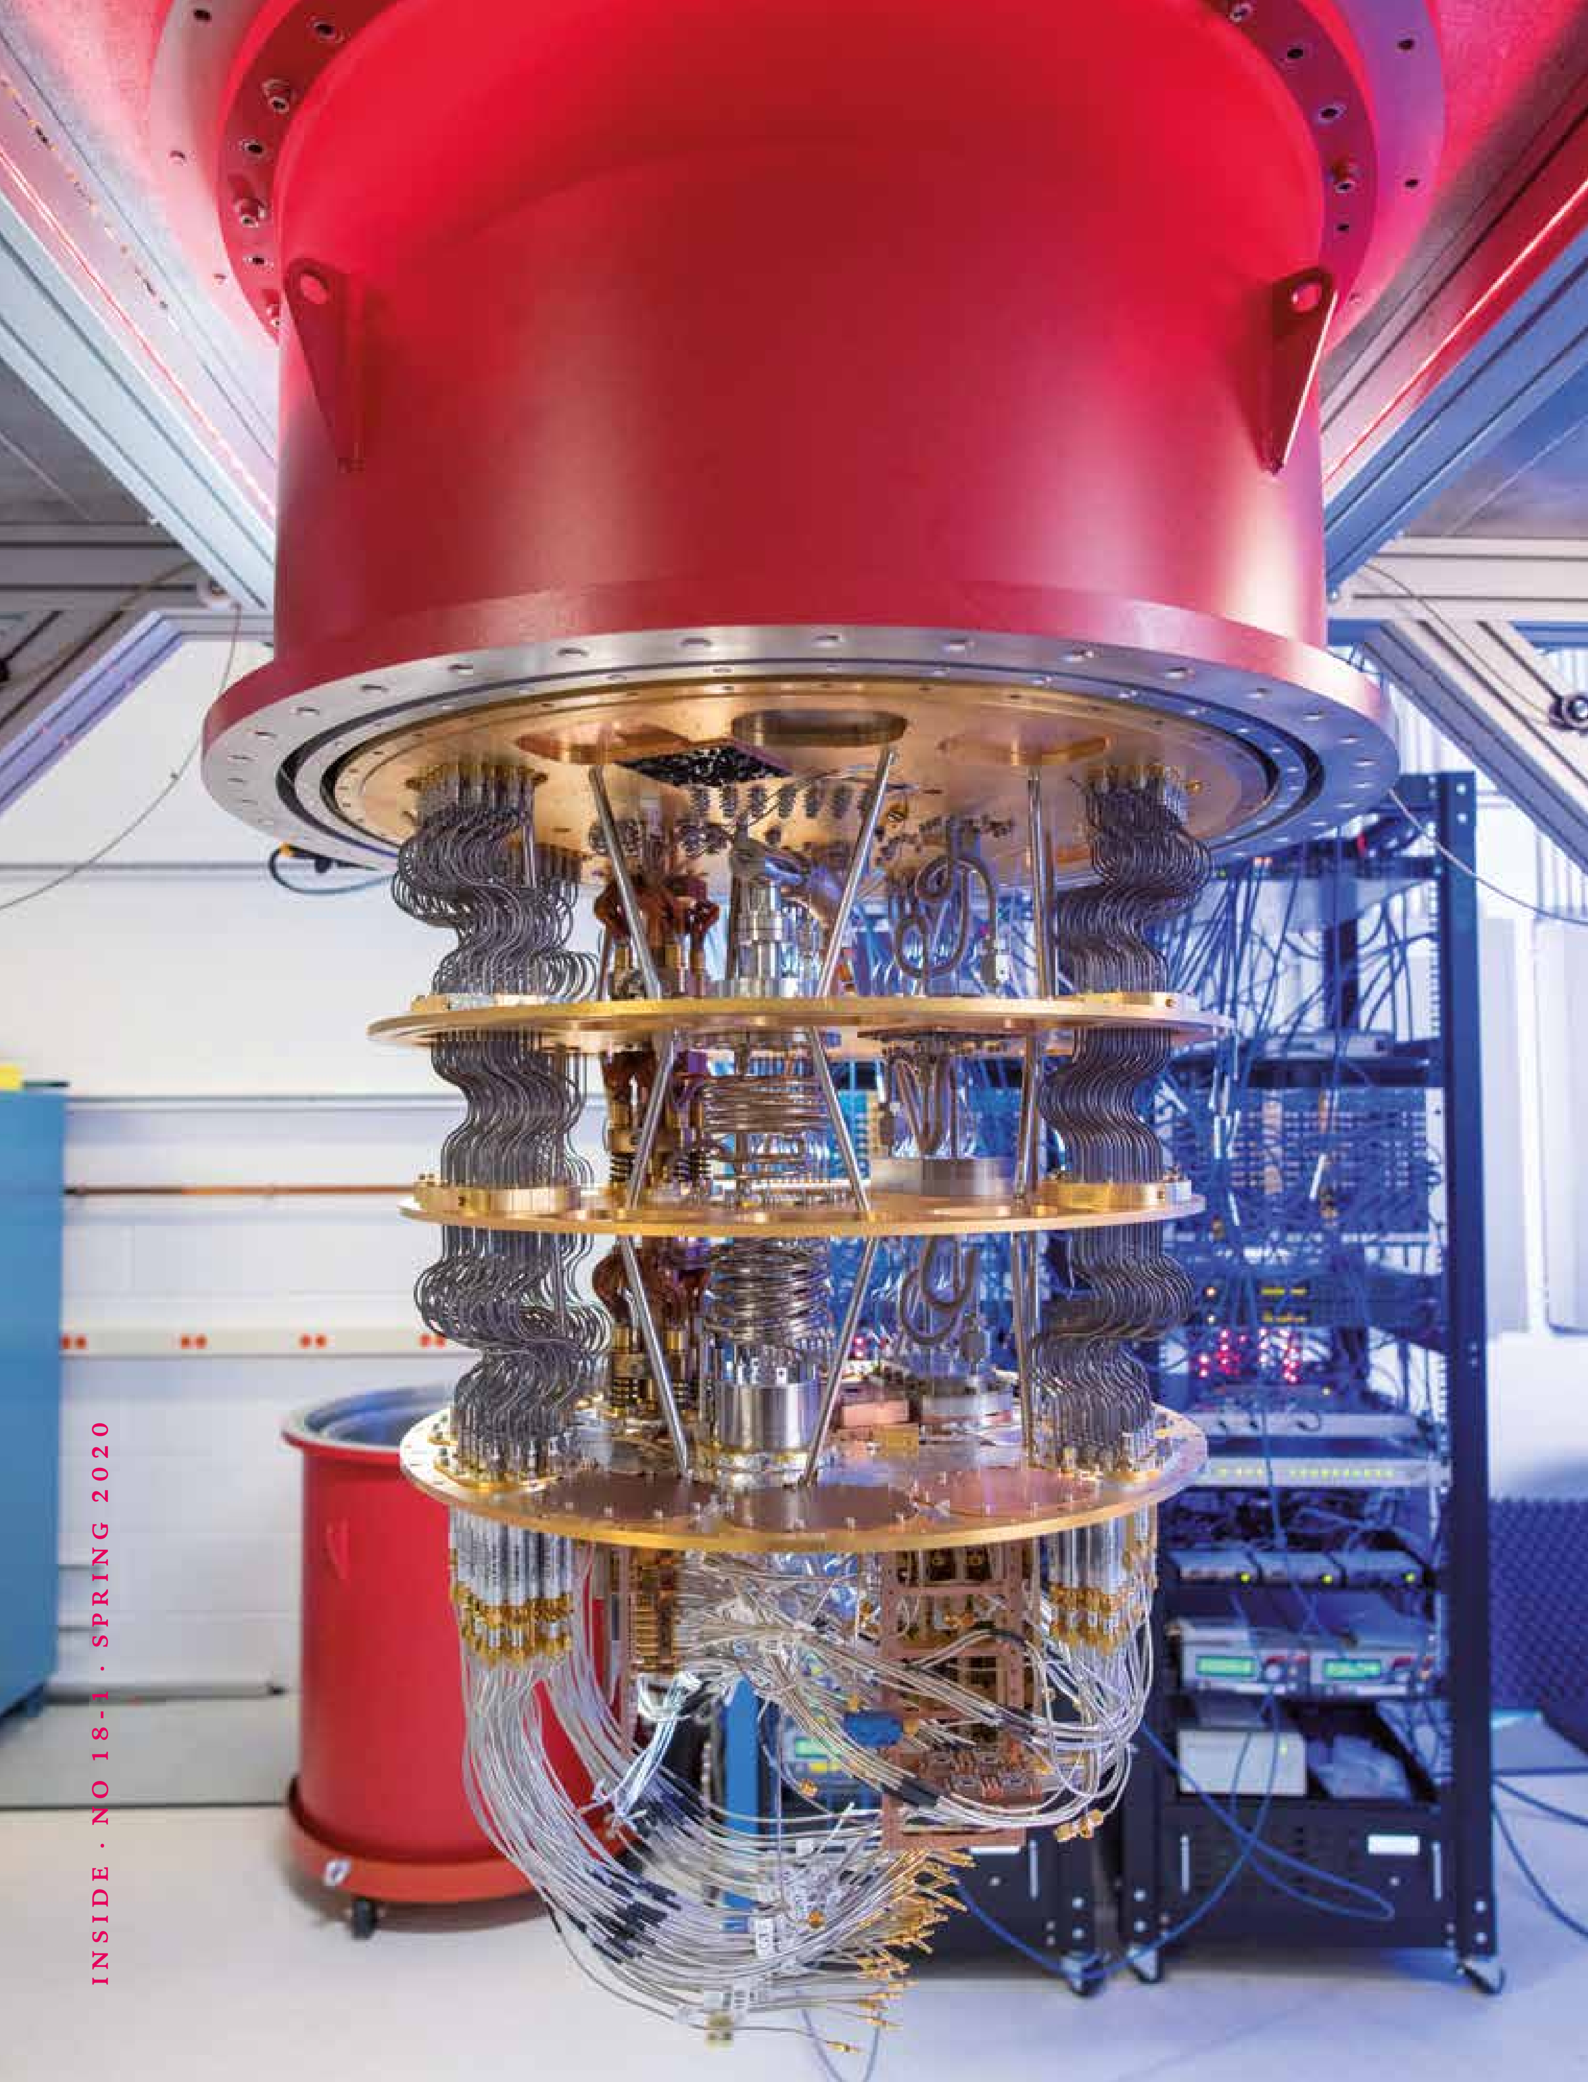
Google's quantum computer. © Google

Shortly afterward, Michielsen was a co-author on a *Nature* paper where Google demonstrated quantum supremacy; that is, the paper demonstrated that for certain problems, quantum computers could, in fact, achieve results far quicker than traditional HPC systems. Michielsen and the Quantum Information Processing group assisted with verifying and benchmarking the quantum machine used in the research group's calculations.

In November, 2019, JSC and Google also hosted their first joint "Cirq Bootcamp," inviting researchers from academia and industry to learn how to apply Google's Cirq language to write efficient codes for first-generation quantum computing machines.

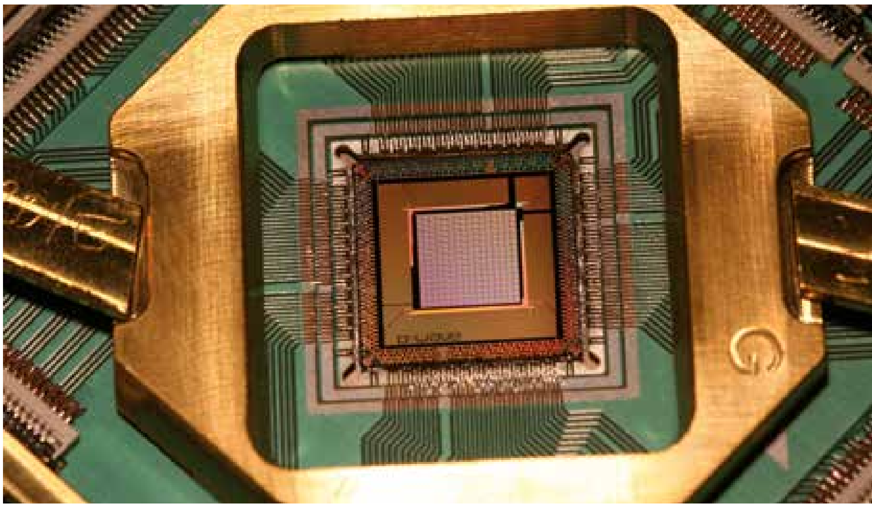

D-Wave 2000Q quantum chip. **© D-Wave Systems, Inc.**

JUNIQ Lab wins bid for first D-Wave system outside of North America

In addition to JSC's progress toward quantum supremacy with Google, the centre also secured a state-of-the-art quantum computing system developed by D-Wave. The Canadian quantum computing company selected JSC to serve as the first location to host an "Advantage" quantum annealer connected to its Leap cloud service outside of North America.

The new system will serve as the cornerstone piece of hardware comprising JSC's newest lab—the Jülich Unified Infrastructure for Quantum Computing (JUNIQ).

D-Wave 2000Q quantum annealer. **© D-Wave Systems, Inc.**

JUNIQ will serve as Europe's gateway to several different architectures of quantum computers, allowing German and European researchers efficient access to the hardware, including the company's current 2000Q system. FZJ and JSC leadership held the signing ceremony with D-Wave leadership last in October, 2019.

In additional to serving as Europe's flagship location for accessing D-Wave's quantum computing resources, the JUNIQ lab brings researchers closer together with researchers focused on developing and scaling quantum applications, offering support and perspective for industrial and academic researchers anxious to scale applications to take advantage of the promises that quantum computing offers. JSC also plans to integrate the D-Wave quantum annealer into its modular HPC environment—the centre's JUWELS system has a long-term upgrade plan that will give the machine "modules" with multiple computing architectures, allowing users to take advantage of the part of the machine most suited to their respective research tasks.

Michielsen indicated that JSC was a natural fit for this convergence of quantum computing technologies and expertise. "JSC not only operates an HPC infrastructure and makes it available to external users," she said. "it has also has demonstrated its competence in the field of system architectures, technology evaluation and co-development, and implementation of system software. In addition, the staff at JSC have long been deeply involved in research and development on key components for advanced computer technologies, such as quantum computing." *eg*

GCS CENTRES SUPPORT RESEARCH TO MITIGATE IMPACT OF COVID-19 PANDEMIC

High-performance computing provides essential tools for drug discovery and epidemiological modeling in the fight against the global pandemic.

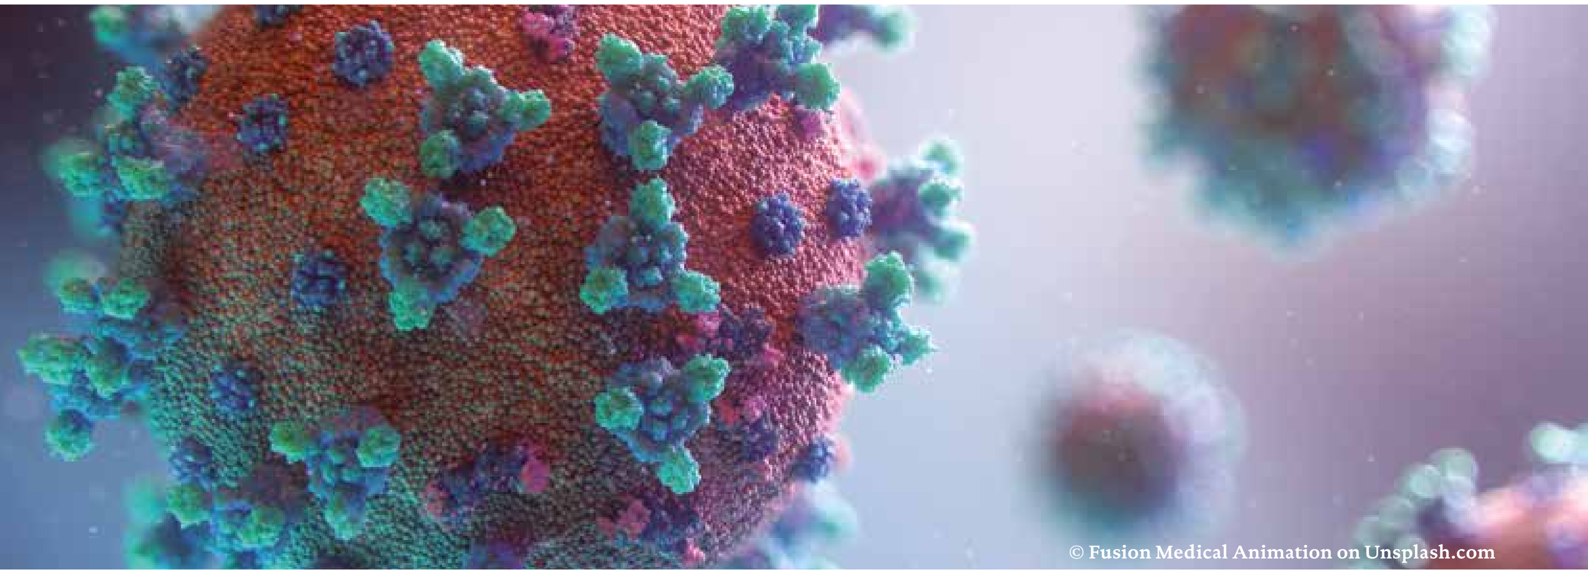

© Fusion Medical Animation on Unsplash.com

In December 2019, the world learned of a new and deadly pathogen. News coming out of Wuhan, China confirmed public health experts' worst fears—a novel coronavirus appeared to have jumped from animals to humans. It was extremely contagious, and its penchant for hospitalising and killing vulnerable individuals has led to sweeping and indefinite changes to daily life around the globe.

Molecular biologists, chemists, and epidemiologists responded quickly in a race to combat the pandemic. As the full extent of the threat became clear in early March, the Gauss Centre for Supercomputing (GCS) joined the effort, announcing that it would fast-track applications for computing time aimed at stopping the spread of the virus or developing new treatments. Since then, GCS has supported roughly a dozen projects focused on epidemiological and drug discovery research, and remains committed to supporting scientists around the globe who are working tirelessly to combat the world's worst pandemic in at least a generation.

Coronaviruses are a broad class of virus that cause illnesses ranging from the common cold to the severe acute respiratory syndrome (SARS) illness that first appeared in humans at the turn of the century. The pandemic coursing across the world over the last 6 months is also a coronavirus, known as SARS-CoV-2, which causes the illness 'coronavirus disease 2019' (COVID-19). As of June, 2020, the world has no proven course of treatment, and promising vaccine candidates are just beginning human trials.

Coronavirus spreads when droplets of infected individuals´ saliva are transmitted by coughing, sneezing or speaking to other individuals, who absorb them through the mucous membranes of the nose and mouth. Although evidence is not conclusive, the virus might also spread through contact with infected saliva droplets that land on surfaces. While medical experts largely understand how the virus spreads, humans have no effective immunity against emerging diseases stemming from novel viral strains like SARS-CoV-2. This means that containment and social isolation are the most effective tools for buying researchers time to study treatments, develop vaccines, and create tools for tracking disease spread.

While societies have shuttered businesses and populations have largely remained at home, scientists are doing everything possible to support medical professionals at the front lines of the pandemic. Computational biologists and chemists have been using high-performance computing (HPC) to understand the virus at a molecular level, in order to identify potential treatments and accelerate the development of an effective vaccine. Epidemiologists have turned to the power of supercomputers to model and predict how the disease spreads at local and regional levels in hopes of forecasting potential new hot spots and guiding policy makers' decisions in containing the disease's spread. GCS is supporting several projects focused on these goals.

Searching for the next outbreak: epidemiological modelling to track COVID-19

While researchers begin to understand how coronavirus spreads on a person-to-person level, modelling how it spreads in communities or regions requires significant amounts of computing power and access to quality data. Even before Germany began seeing its first COVID-19 cases, leadership at JSC started collaborating with researchers at the University of Heidelberg and the Frankfurt Institute for Advanced Studies (FIAS) who had been modelling the disease's spread in China. JSC offered its computational tools and expertise to digitalise epidemiological modelling and ultimately help predict how the virus would spread at state and local levels both in Germany.

"At the very beginning of this crisis, we were interested in how we could support early reaction and detection systems like computational scientists can do with tsunami or earthquake simulations," said Prof. Thomas Lippert, Director of JSC. "As this is a very dynamic situation, we began to model system changes and we try to predict developments."

With the pandemic still actively spreading around the globe, researchers knew that running quantitative, retrospective analyses of the situation was not yet appropriate. However, supercomputers could be used to combine datasets on infection growth rate, the so-called reproduction number (Rt), and the virus incubation time to create predictive models. With supercomputers, researchers begin to run suites of scenarios to predict the mortality rate on a local and national level based on the degree of social distancing measures and other actions meant to slow the virus's spread

"The qualitative validity of these models comes from the fact that one can play through the different assumptions and detailed interactions, so you can validate those methods with hard data," Lippert said. "Then you put these different measures into the model and see what it is doing. We can then ask, 'when we put these measures together, are they moving things in a positive or negative direction?'"

Lippert noted that such models became less accurate the farther into the future they tried to model, but that their early results were accurate enough to help guide policy makers.

"In a paper we published based on data up to the March 20, we predicted the situation in Germany for April 20 within a few percent," he said. "Because we already knew what measures were in place around the country, our work was pretty good at these predictions. Nevertheless, the model still underestimated the number of fatalities. At the policy and public health level, that means that if our data seems to overestimate the number of fatalities, it may not actually be doing that."

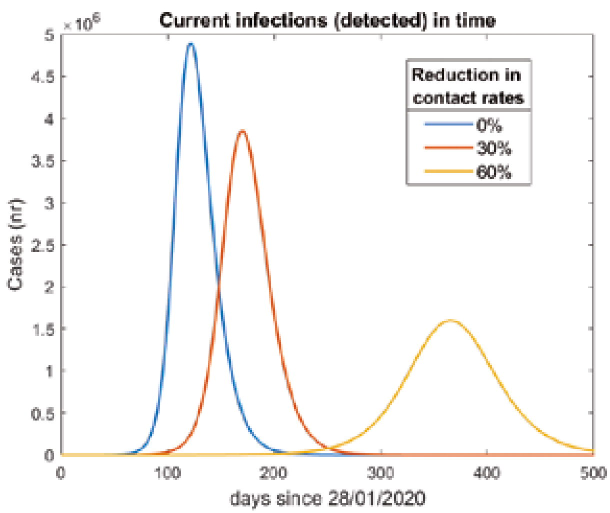

Different scenarios based on a reduction of contacts by 0, 30%, and 60%. This plot shows that the curve becomes flatter and wider as the number of contacts is reduced.

© Barbarossa, et al. DOI: 10.1101/2020.04.08.20056630

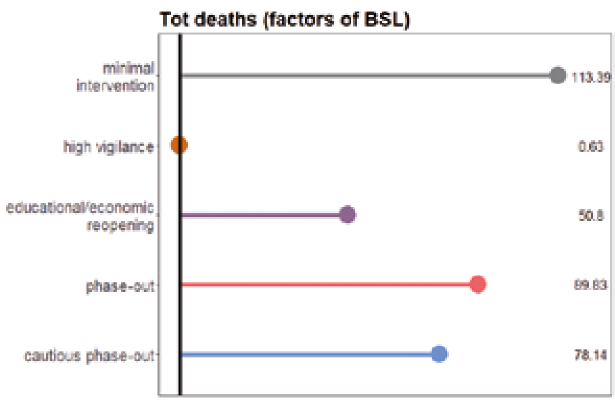

The effect of different measures on the total number of deaths. The baseline scenario (BSL) assumes that all measures had remained in place.

© Barbarossa, et al. DOI: 10.1101/2020.04.18.20069955

Lippert, JSC researchers Dr. Jan Meinke, Dr. Jan Fuhrmann, and Dr. Stefan Krieg, and principal investigator and University of Heidelberg / FIAS research group leader Dr. Maria Vittoria Barbarossa were all contributors to a position paper published April 13 by the Helmholtz Association of German Research Centres. The paper, which was signed by the leadership of the Helmholtz Association and co-authored by 17 researchers, presented 3 scenarios for German government officials with respect to easing restrictions imposed during the COVID-19 pandemic.

The team demonstrated that if contact restrictions were raised too quickly, the Rt value would quickly rise over (an Rt value of 1 representes that each infection will spawn 1 new infection), and Germany's healthcare system could become overburdened within several months. In the second scenario, researchers modelled easing restrictions gradually and adopting an aggressive "feedback-based" contact tracing model to help slow the disease spread throughout the country. While in principle this scenario seemed promising, it required that significant contact restrictions would remain in place for an extended period of time—think months rather than weeks. The third scenario had the most resonance with German policy makers—keeping strong contact restrictions in place several weeks longer to help the Rt drop well below 1, then beginning the gradual reopening process.

Prof. Peter Coveney of University College London. European lead on the Consortium for Coronavirus. **© UCL**

International collaborations converge on essential drug discovery work

While predicting the virus's spread over the initial weeks and months of the pandemic is essential, making it possible for society to return to normal will require development of effective treatments and scalable vaccinations to protect from infection.

University College London (UCL) Professor Dr. Peter Coveney has long leveraged supercomputers to understand drug interactions with pathogens and the human body. Since 2016, he has led the European Union's Horizon 2020-funded project CompBioMed, which stands for 'Computational Biomedicine', and its successor project, CompBioMed2 (for more information, visit www.compbiomed.eu or the project overview on page 24). Both projects focus on accelerating drug discovery by augmenting experimental validation with modelling and simulation.

In the face of the COVID-19 pandemic, Coveney and over a hundred of his peers jumped into action, in part focusing their knowledge and access to HPC resources on indentifying existing drug compounds that could turn the tide against the virus. Specifically, Coveney and his collaborators model the binding affinities of drug compounds and pathogens. A drug's binding affinity essentially means the strength of the interaction between, for instance, a protein in the lifecycle of a virus and active compounds in a medication—the stronger the binding affinity, the more effective the drug.

"We can compute binding affinities in a matter of hours on a supercomputer; the size of such machines means that we can reach the industrial scale of demand necessary to impact drug repurposing programs," Coveney said. "This can save us enormous amount of wall-clock time and resources, including person hours, which are very precious in such a crisis situation."

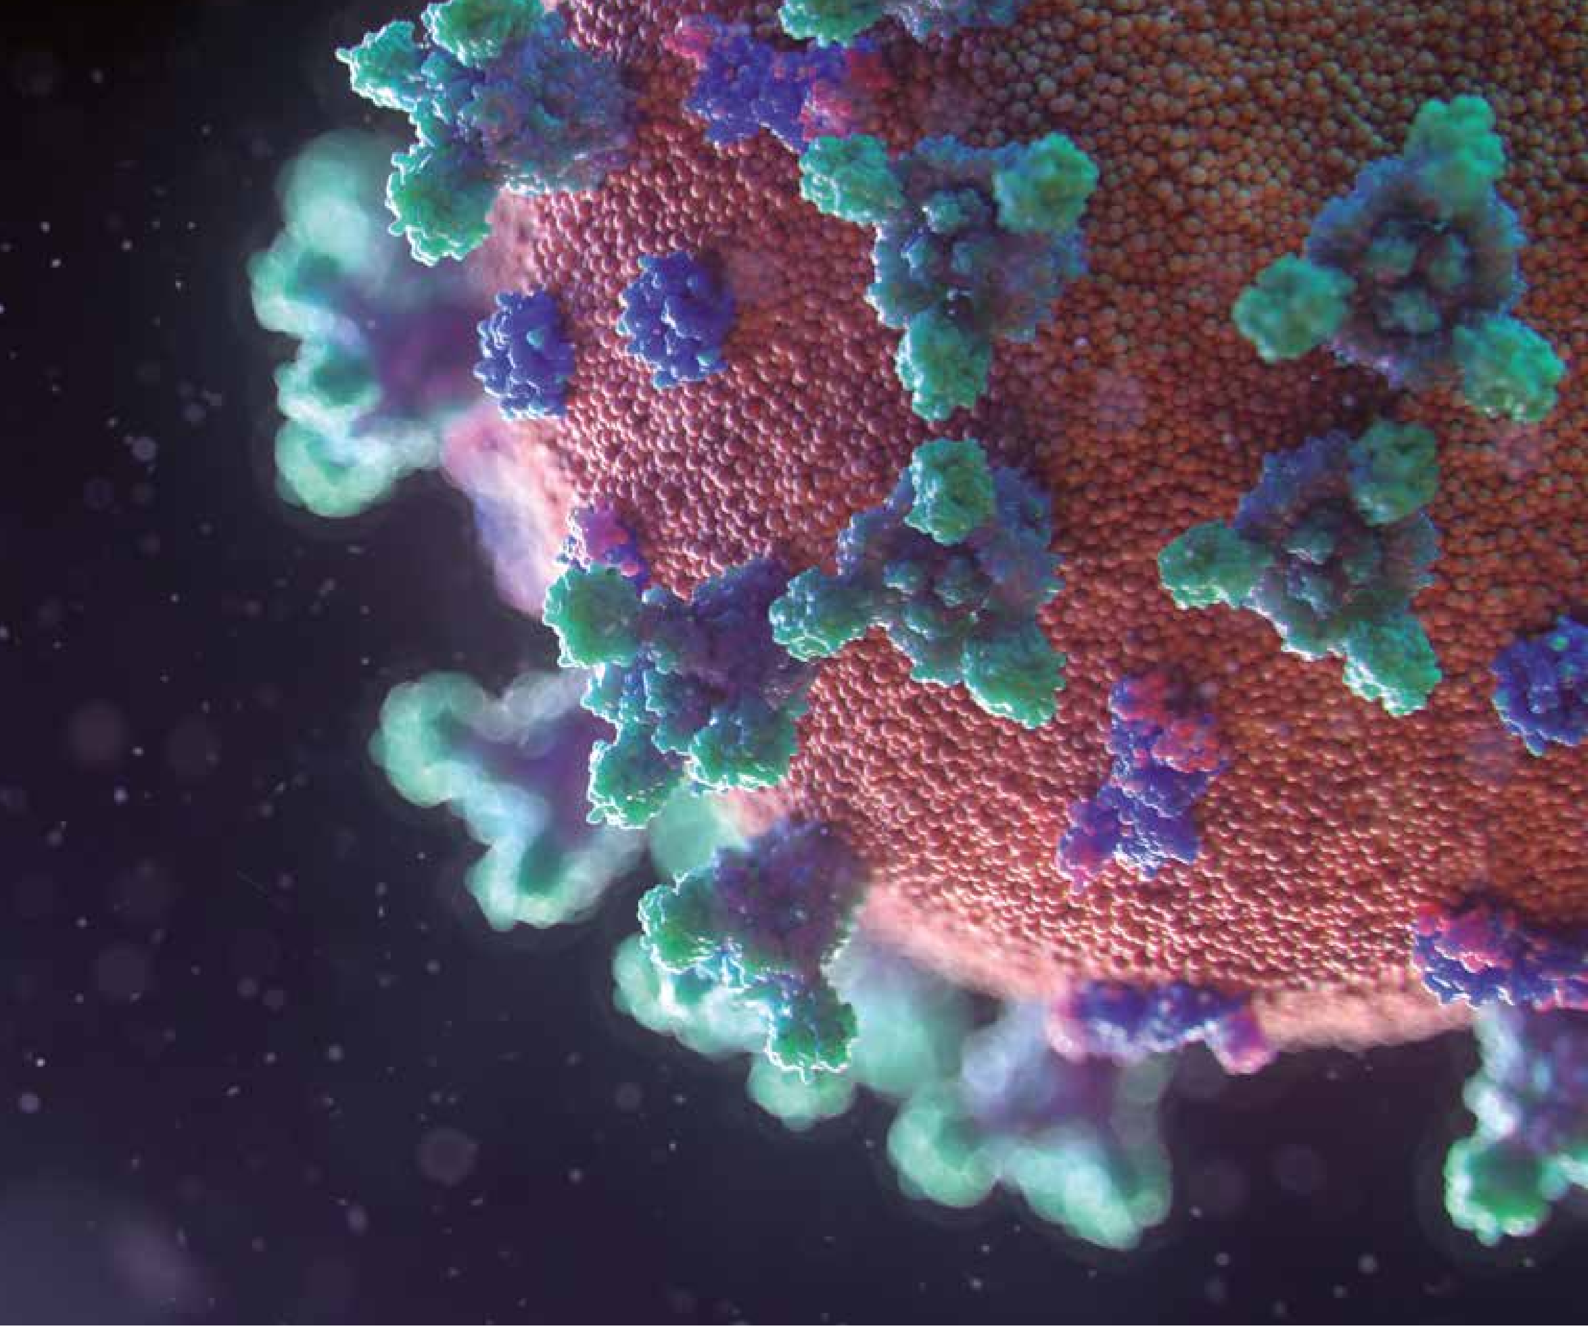

As part of the Consortium on Coronavirus, European researchers led by CompBioMed lead Prof. Peter Coveney are using some of the world's leading HPC resources to scan through billions of drug compounds to find promising candidates to treat COVID-19. **© Fusion Medical Animation on Unsplash.com**

Supercomputers allow researchers to run large numbers of binding affinity simulations in parallel. Here, they compare information about the structure of the virus with a database containing information about known drug compounds to identify those with a high likelihood of binding stringing. This computational approach enables researchers to investigate large numbers of potential drugs much more quickly than would be possible if they had to mix individual drug samples with actual viruses in a lab. Coveney has been using the SuperMUC-NG supercomputer at LRZ to run many of his binding calculations.

"SuperMUC-NG offers us an immense capability for performing a large number of binding affinity calculations using our precise, accurate and reproducible workflows—ESMACS (Enhanced Sampling of Molecular dynamics with Approximation of Continuum Solvent) and TIES (Thermodynamic Integration with Enhanced Sampling)," Coveney said. "So far, we have already performed a few hundred such calculations very quickly."

Coveney has long collaborated with LRZ, developing his workflow to scale effectively on multiple generations of the SuperMUC architectures at LRZ. LRZ Director Prof. Dieter Kranzlmüller saw the recent work as a continuation of Coveney's efforts. "Our long-term collaboration has enabled us to immediately identify and reach out to Peter to offer our assistance," he said. "By strongly supporting research in drug discovery activites for years, we were in a position to ensure that research toward identifying therapeutics could be accelerated right away."

Coveney has been performing his work as part of the Consortium on Coronavirus, an international effort involving researchers and resources from 9 universities, 5 United States Department of Energy national laboratories, and some of the world's fastest supercomputers, including SuperMUC-NG (currently number 9 in the Top500 list) and Summit at Oak Ridge National Laboratory in the United States (currently the world's fastest machine for open science). "This consortium is

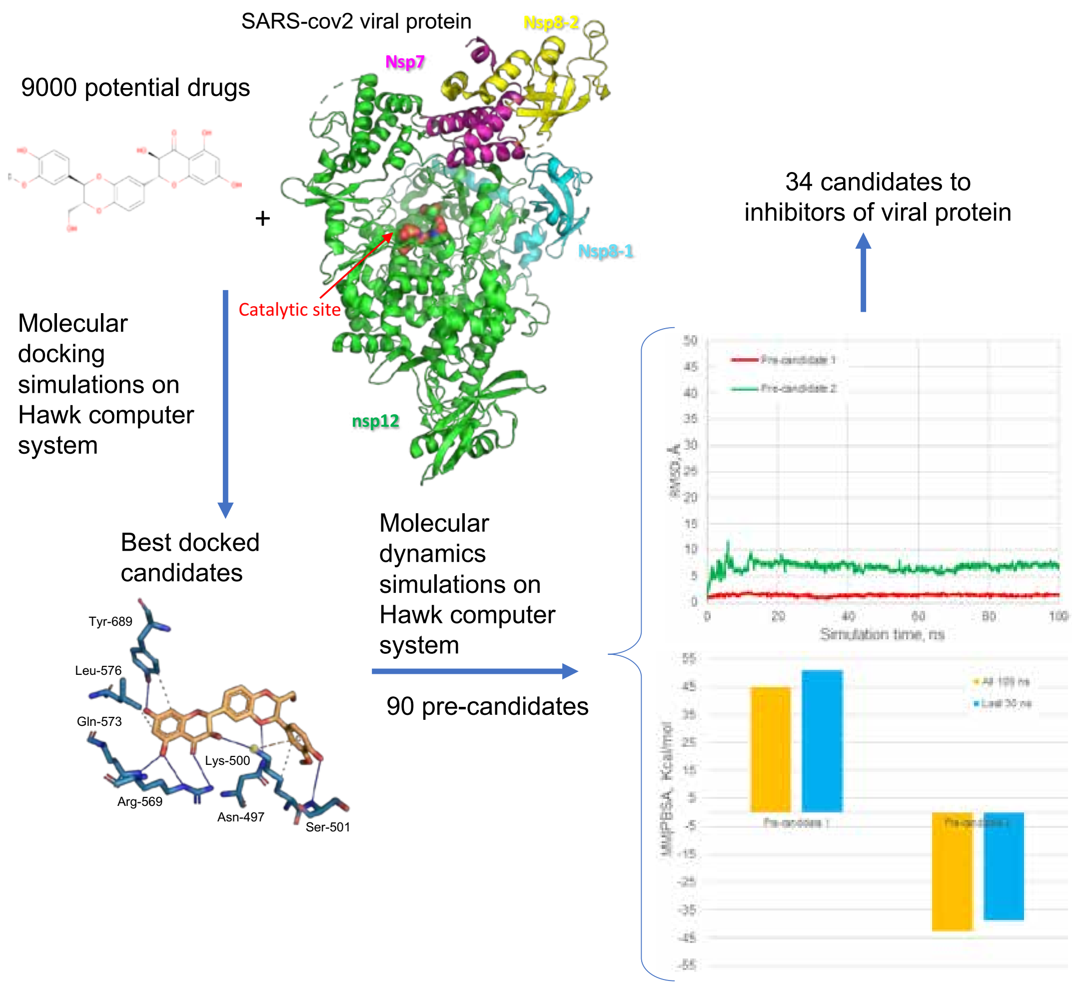

Using HLRS resources, researchers were able to narrow the list of 9,000 potential drug candidates to inhibt one of SARS-CoV-2's key proteins to 34 of the most promising candidates. **© Jose Antonio Encinar Hidalgo**

a vast effort, involving many people, supercomputers, synchrotron sources for experimental structural biology and protein structure determination, wet labs for assays, and synthetic chemists who can make new compounds," Coveney said. "In all, it's a massive 'one stop shop' to help fight COVID-19."

Considering the team's ability to use supercomputers to run many iterations of drug binding affinity calculations, Coveney, who leads the European side of the consortium, is grateful for as much access to world-leading supercomputers as he can get. "Our workflows are perfectly scalable in the sense that the number of calculations we can perform is directly proportional to the number of cores available," he said. "Thus, having access to multiple HPC systems speeds things up substantially for us. Time is of the essence right now."

With access to HPC resources in Europe and the United States, Coveney and his collaborators have narrowed a list of several hundred drug compounds and identified several dozen that have the potential to inhibit SARS-CoV-2's replication in the body. In total, Coveney and his colleagues have scanned millions to billions of possible compounds via machine learning, ultimately helping them narrow down existing and novel compounds to find the most promising candidates. Once machine learning helps identify the most promising candidates, these are then subjected to computationally intensive, physics based simulations, which provide more accurate calculations.

Molecules in motion: molecular dynamics simulations for observing drug-virus interactions

As a traditional leader in computational engineering, HLRS staff have extensive experience supporting molecular dynamics (MD) simulations. In the realm of engineering, MD allows researchers to understand how combustion processes happen from the moment of ignition, but in the realm of computational biology, researchers can turn to these computationally intensive simulations to investigate how molecular structures in proteins move and interact at extremely high resolution.

A team led by Prof. Dr. José Antonio Encinar Hidalgo at the Universidad Miguel Hernandez in Elche, Spain has been using HPC resources at HLRS to run molecular dynamics simulations and molecular docking models for 9,000 different drug candidates to fight COVID-19.

Proteins on human cells and viruses come in distinctive shapes, and designing effective treatments requires that researchers understand the molecular configurations most likely to bind to one another. Molecular docking simulations serve as a basis for determining drug binding affinities—by simulating the structures of panels of drug compounds in various molecular positions, researchers can assess their potential to bind to and inhibit the function of viral proteins.

Encinar noted that while some molecular docking simulations could be performed on more modest computing resources, HLRS's supercomputer enabled the team to put these snapshots of molecular docking configurations in motion through the use of molecular dynamics simulations. "Our calculations consisted of some 90 molecular dynamics simulations," Encinar said. "On Hawk, a simulation takes approximately 5 days to calculate. But Hawk also enables us to calculate about 50 simulations at a time. In two weeks, we have all the necessary data. This work is not approachable in adequate time without high-performance computing resources."

The team is preparing a journal manuscript reporting on its scan of 9,000 different drug compounds. It identified roughly 34 candidates that appear to have a high likelihood of inhibiting one of the key proteins of SARS-CoV-2.

Dreams of vaccines and hope for the future

In addition to the work described above, dozens of researchers focusing on other aspects of drug discovery and epidemiology related to COVID-19 have been granted access to HPC resources at the GCS centres through GCS's fast-track program as well as PRACE's calls for fast-tracked access to Europe's leading HPC resources. (For a full list of COVID-19 related projects running at the GCS centres, please visit: https://www.gauss-centre.eu/news/newsflashes/article/mitigating-the-impact-of-covid-19-pandemic/).

The ultimate goal for scientists, medical professionals, and government officials, though, lies in developing an effective vaccine and scaling up production on a global scale. Coveney indicated that supercomputers have already helped pave the way for vaccine trials, enabling researchers to comb through 30,000 DNA sequences and to design vaccine candidates that are currently entering the testing phase. There are some aspects of fighting a pandemic that supercomputing cannot accelerate, though, and as vaccine candidates enter clinical trials, societies around the globe can only hope that the foundational work done by computational scientists has helped make identifying and designing a vaccine as efficient as possible.

Coveney was encouraged by the degree of collaboration we are currently witnessing between researchers across the globe. "Drug design involves a long and tedious pipeline of tasks with a large number of steps that require a different type of expertise at each level," he said. "Working in a large consortium has obvious advantages for such projects. We are part of a well-organized project with each partner having a clear idea of his role leads to quick turnaround. Proper and clear communication is vital for the success of our project. We are using online repositories for sharing of codes as well as data or information. Weekly video conferencing allows us to make checks on progress and remains in sync, along with frequent chats between concerned subsets of people, and this has made it possible to successfully move forward in step."

For GCS leadership, this crisis has shown that making sure that computing resources are quickly and efficiently deployed for researchers in the midst of a crisis is of the utmost importance. "At LRZ, we have discussed the need for detailed plans to address the next crisis, not necessarily a pandemic," Kranzlmüller said. "We had an internal plan, and were able to send all staff into home office within a few days, but we also have a long tradition focusing on biological and virology research, computational climate science, and other research areas that could be relevant for future disasters or crises. We want to ensure when the next crisis comes, supercomputers are among the first technological resources capable of being deployed in support of front-line efforts."

Lippert, who has studied both computational science and quantum physics at the PhD level and is a vocal advocate for science, remains positive due to his trust in the international scientific community.

"Any vaccine will come from science, any measure to be validated for epidemiology will come from science, any pharmaceutical therapy will come from science, any understanding of the hygienic aspects needed to redesign or rebuild public places where people are gathering together, these are all things that are to be understood scientifically," he said. "And I believe we will be successful because of the strength of science in Germany, Europe, and around the world."

eg

HIGH-PERFORMANCE COMPUTING AND ARTIFICIAL INTELLIGENCE DECODE THE LANGUAGE OF LIFE IN PROTEINS

TUM researchers partner with LRZ HPC experts to improve access to and organization of protein databases.

For over seven decades, scientists have been trying to crack the code of proteins, the machinery of life, out of proteomics research—the biological field dedicated to describing and cataloging proteins and their functions in living things. In the process, scientists began collecting valuable information about some of life's most fundamental processes and how various stresses, mutations, or other changes can impact health.

However, there can be too much of a good thing. As technologies have made protein sequencing and cataloging cheaper and faster, medical professionals and researchers have rapidly increased the size and scope of protein databases. With terabytes (one thousand gigabytes) of information to sift through, those needing quick answers now encounter many difficulties in finding information that is pertinent and helpful to a specific diagnosis, treatment plan, or research goal.

To address this problem, researchers at the Technical University of Munich (TUM) partnered with staff at the Leibniz Supercomputing Centre (LRZ) to use high-performance computing (HPC) and a machine learning approach to help structure and organize these extremely valuable databases.

"What we've developed is certainly already useful for computer-savvy researchers. Over time, we hope this method will also enable medical professionals and researchers who don't have access to computer clusters to search through these databases faster," said Michael Heinzinger, member of the bioinformatics lab led by Prof. Burkhard Rost at TUM (Rostlab). "For example, when a patient comes to a hospital in an emergency, doctors may need to know what protein mutations that person has and what potential effect those mutations have so that they can better assess potential impacts on treatment options. Bringing this work to the commercial level is our goal."

Early returns are promising, and the team has already made protein databases more accessible. The team recently published its results in BMC Bioinformatics.

The language of life

Proteins are molecules that act as the unsung heroes of living things. Whether a protein is helping replicate DNA, automate processes like blinking and breathing, or processing and reacting to external stimuli, these molecules underpin most of the things that make us sentient. Proteins are long chains of amino acids, one of the fundamental building blocks of life. Much like how words connect to form sentences, paragraphs, or whole novels, the order of these amino acids determines the role a protein will play in the body, helping convey its "meaning" to the complex molecular inner workings of the body.

The analogy between how grammar and syntax impart meaning and a protein's amino acid sequence served as an inspiration behind the team's work. Specifically, the team began using an algorithm developed to study language in its protein database work. The team used ELMo, a natural language processing (NLP) algorithm for analyzing word use and variation in different contexts.

Ahmed Elnaggar, another member of the Rostlab, came across this algorithm in the context of machine learning and deep learning research, and brought it to the attention of the biology focused members of the lab. "I am working on these problems with a machine learning and deep learning focus, but we have many interdisciplinary discussions in the lab," he said. "I brought this to Michael's attention, and we started this collaboration. This kind of work would not be doable if people from the purely computational or the biological sides of this research worked alone on these problems."

This interdisciplinary collaboration led researchers to use ELMo for organizing proteins by their similarities and begin to shine light on a dark corner of proteomics research—the dark proteome. Essentially, the dark proteome came to be from the rapid expansion of protein sequencing.

"What happens often times is that we have no information other than the sequence information itself," Heinzinger said. "More than half of a terabyte of data is dedicated to just the amino acid sequences comprising proteins, but that doesn't mean we understand their origin, functions, or any other context. It is like having the words of an ancient language without grammar or syntax to organize them into something meaningful."

To shed light on this ancient language, the team trained ELMo to detect re-occurring patterns within millions of protein sequences. Ultimately, this compresses protein databases into a more organized, computer-readable format. Using the pre-trained ELMo model allows researchers to draw conclusions about their protein of interest without

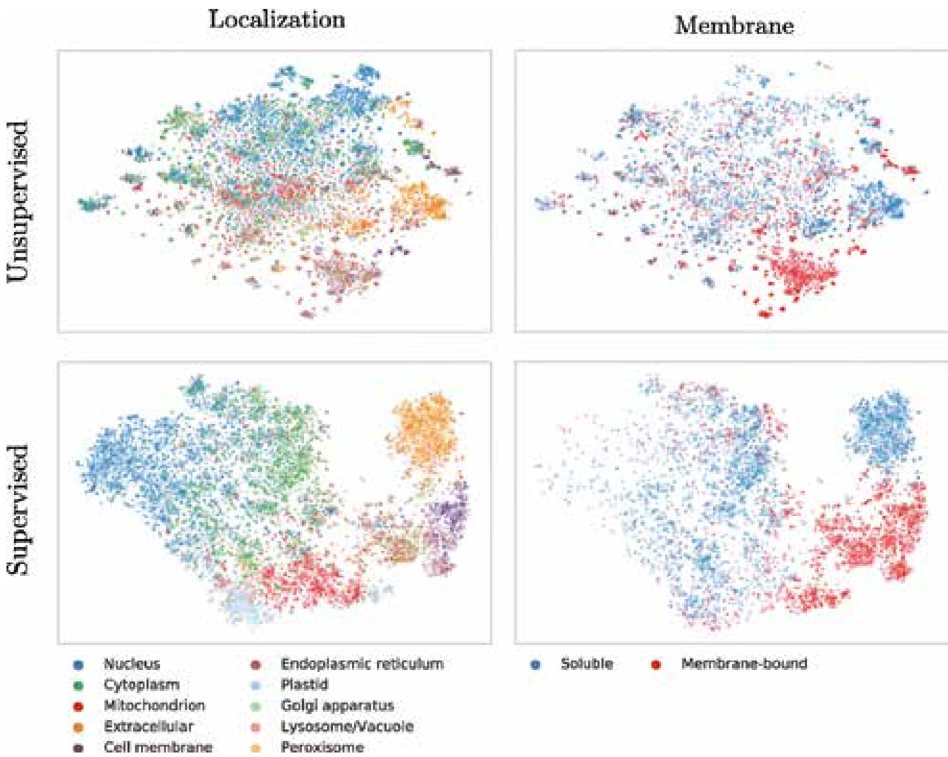

In these plots, every dot represents a protein. Proteins with the same color share an aspect of protein function, in this case, where it lies in the cell. A good model can pull similar colors together, and push dissimilar colors apart. In the upper row, "unsupervised" embeddings, or proteins represented by the raw output of the model. These results show that the model implicitly learned different aspects of a protein's function without having ever been taught to do so.
© Heinzinger, et al. BMC Bioinformatics 20 (723)DOI:10.1186/s12859-019-3220-8

the need to search hundreds of millions of proteins stored in today's databases. This shifts the hardware requirements from clusters needed for searching databases to a normal laptop needed to run the pre-trained model.

Working with LRZ's Big Data and Artificial Intelligence team, the team scaled its application using the SuperMUC supercomputer at LRZ and the centre's NVIDIA DGX-1 cluster. In the process, it was able to not only achieve a three-fold speedup; the researchers also got valuable experience scaling their application from a single node to hundreds of nodes, allowing them to take advantage of larger GPU-centric architectures.

"Training large deep learning models on large datasets requires tremendous amount of computing power and storage, which are usually not directly available at universities' chairs," Elnaggar said. "That's when LRZ shines because it provides the computing power and storage required by researchers for training large deep learning models or for simulations. Simply put, without having access to the supercomputers on LRZ, our research would not have been possible."

Accelerating and advancing

By proving their concept and method, the team now looks forward to continuing to scale the application in order to further compress and organize protein databases. The team still has a large task of filtering out "noise," or uncorrelated, disorganized data in databases, but with access to increasingly powerful GPU-powered clusters and HPC resources, the team is in a good position to bring this work to the people who need it most.

"The concept here was to make a representation of protein sequences in a space that makes sense computationally, and that is the gist of our recent paper," said team member Christian Dallago. "But the bigger point here is to shift this research to the people who need it and might just have a laptop. We are trying to democratize the use of these machine learning algorithms for people who might not have access to a supercomputer or powerful cluster, be it research institutions in the developing world or hospitals and other medical facilities who don't have a lot of resources." *eg*

RESEARCHERS BUILD BIOMOLECULE TOLERANCE WITH SUPERCOMPUTING SIMULATIONS

High-performance computing helps scientists better understand enzymes that are more resistant to detergents and solvents.

Enzymes are the tiny, unheralded heroes of the bioengineering world. From ensuring that drug compounds reach their intended place in the human body to accelerating specific chemical reactions for a variety of industrial uses, enzymes play a role in countless medical and industrial processes.

For decades, bioengineering research advanced almost exclusively through experimental work, but scientists on the leading edge want to more fully understand enzymes' biochemical interactions with one another and their environments at a fundamental level. To do this, researchers have increasingly turned to high-performance computing (HPC), which allows them to offer narrow predictions that can be tested experimentally as well as verify experimental results quickly and efficiently.

To that end, a group of researchers led by Dr. Holger Gohlke, Professor of Pharmaceutical and Medicinal Chemistry at Heinrich Heine Universität Düsseldorf and head of the John von Neumann Institute of Computing (NIC) research group Computational Biophysical Chemistry at the Jülich Supercomputing Centre (JSC) and Institute of Biological Information Processing (IBI-7: Structural Biochemistry) at Forschungszentrum Jülich (FZJ), is using HPC resources at JSC, one of the three centres comprising the Gauss Centre for Supercomputing (GCS), to learn how to imbue enzymes with increased tolerance to detergents and solvents encountered in a variety of environments—an important characteristic needed for drug design and other biotechnological applications.

"I would like to understand, predict, and modulate bimolecular properties and interactions of biomolecules with their environments," Gohlke said. "In this case of our current research, environments are different types of solvents, mixtures, or solvents containing detergents. In our latest study, we looked at an enzyme. Enzymes are biotechnologically interesting molecules, where they often have to function under conditions that do not comply with their natural environments. For that, such molecules need to be optimized."

To its knowledge, the team just completed the largest-ever study of variants of the enzyme Bacillus subtilis lipase A with respect to two enzyme properties, which has significant applications in the biotechnology sector. The work would be the largest computational analysis to date of a complete library of the enzyme's variants, or the sum total of all possible single mutations that could be made to a specific enzyme. The team's research was published in the Journal of Chemical Information and Modeling.

Creating a virtual test tube

Biomolecules are a broad class of "soft matter," meaning that they are easily influenced and deformed by changes in their environments. Even under normal temperature settings, the weak bonds holding biomolecules together can break and change quickly. This also means biomolecules can appear in different molecular shapes as they shift and change.

In the vast majority of contexts, a limited range of shapes is required for biomolecules to function, such as helping drug molecules reach their target destination and play their designed role in the human body or accelerating a certain (bio-)chemical reaction. However, there can be instances when biomolecules adopt shapes, or even become shapeless, which leads to a loss of their function. Gohlke and his collaborators use HPC to create large "ensembles" of biomolecular shapes that are post-processed to learn about their behavior in different environments.

Researchers see a two-fold benefit from using molecular simulations in bioengineering: In certain contexts, experimentalists send researchers large amounts of testing data with the hope that computational scientists can uncover the origin of their findings at the atomic level or more broadly help make predictions based on a given dataset; Inversely, computational scientists can also help generate hypotheses about certain systems, which can then be tested by experimentalists.

For the team's current research, it focused on the enzyme Lipase A in Bacillus subtilis, a well-studied enzyme and bacteria with applications in bioengineering. The team was able to do a complete "site-saturation mutagenesis" on the enzyme, meaning that every amino acid in the protein is mutated to every other natural one, leading to about 3,440 variants of the enzyme.

Gohlke pointed out that learning how to catalogue these mutations computationally helps experimentalists drastically reduce the amount of time needed to chart biomolecule mutations experimentally in the future. "What makes this system notable is that the mutation data available in public databases are usually limited to only a few positions of proteins, or limited to substitutions of simple amino acids," he

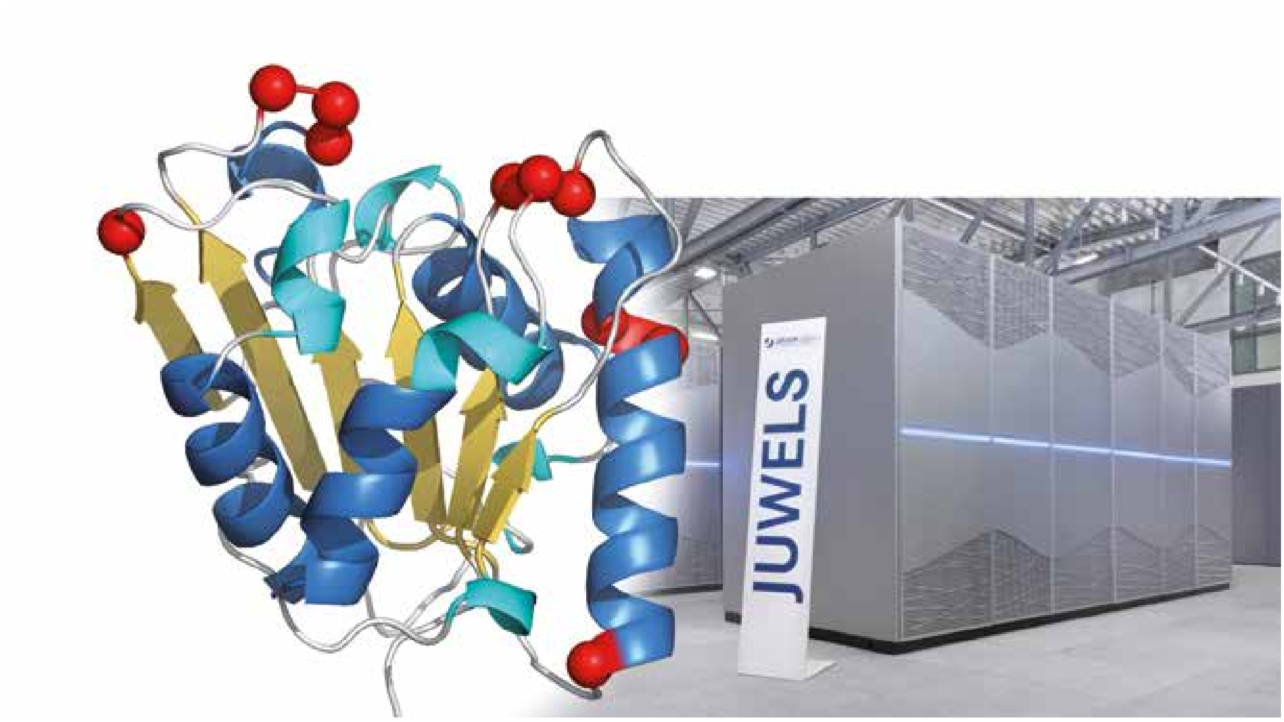

Using the JUWELS supercomputer at the Jülich Supercomputing Centre, a team lead by Prof. Holger Gohlke just completed the largest-ever study of variants of the enzyme Bacillus subtilis lipase A with respect to two enzyme properties, which has significant applications in the biotechnology sector. © Holger Gohlke/Forschungszentrum Jülich

said. "It is not possible to use this publically available data to exhaustively validate algorithms or completely understand what factors are influencing protein stability, because you can't see all the combinations that are possible."

Further, Christina Nutschel, a PhD student on this project, pointed out that simulations such as this help develop better data-driven approaches to protein engineering, particularly with respect to computational models. That means that researchers don't need to make perfect predictions with their simulations, but the insights point to more narrow research focuses for experimentalists to verify.

Finally, the team has been developing a computational approach to study enzymes it calls Constraint Network Analysis. This approach analyzes proteins' static properties like they were a bridge—it looks for hot spots and specific points in the enzyme particularly useful for improving structural stability of the biomolecule. When compared with scattershot predictions made at random, the team's approach saw a nine-fold gain in accuracy for identifying these points in the enzyme. To put that in perspective—an experimentalist is barely able to completely analyze an enzyme with 6 mutation points, because as there are 20 different possible mutations in each place, that leaves the researcher with 64 million different possibilities.

Bridging the gap in bioengineering

Moving forward, the team is working with several pharmaceutical and biotechnological companies to apply its approach to the companies' collections of enzymes, and they are ramping up the work to not only model predictions for various substitutions and mutations, but also predict enzyme structure at a fundamental level.

The team has been working with JSC's Dr. Olav Zimmermann very closely on the work, who also appeared as an author on the team's Lipase A paper. Zimmermann helped verify the team's predictions using the ProFASi tool developed at JSC. "There is a clear, tight interaction and collaboration with JSC," Gohlke said.

Further, the team looks forward to the upcoming booster module to be installed on JSC's current flagship supercomputer, JUWELS. The team's simulation code is optimized for GPU-based architectures, and with the booster module installed later this year, JUWELS peak performance will jump from 12 to over 70 petaflops.

Ultimately, hardware is only one part of the whole equation, though. Gohlke emphasized that his collaborations with Zimmermann and co-author Dr. Karl-Erich Jäger, Professor at Heinrich Heine Universität Düsseldorf and director at the Institute of Bio- and Geosciences IBG-1 at FZJ, would only be possible in a collaborative, multidisciplinary environment like FZJ. "We can all come together at JSC and do this kind of work, and in my view, for this field to continue to progress, we are going to have to continue these cycles of simulation and experiment," he said. *eg*

SUPERCOMPUTING 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.

At its core, combustion is one of humanity's oldest technologies. While this burning process helped the earliest humans cook food and defend themselves from predators, today's combustion processes move cars down the road, help aircraft take flight, and govern countless industrial processes for making everything from tires to the white pigment used in furniture and paints. Ever since the industrial revolution, humanity has continued to harness the power of controlled burning to incorporate into new technologies.

However, as researchers and the public at large begin to reckon with issues posed by climate change and seek to limit the amount of fossil fuels being used in our daily lives, researchers have been studying combustion on a molecular level in hopes of making the process safer, cleaner, and more efficient.

To that end, a research group at the University of Duisburg-Essen (UDE) has been using Gauss Centre for Supercomputing (GCS) resources in order to study the critical milliseconds to microseconds when fuel is injected into an engine, observing how the flame ignites and changes over time. Ultimately, it 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."

In order to develop combustion processes that do all three of these things, Kempf and his collaborators are using supercomputing resources at the High-Performance Computing Center Stuttgart (HLRS) to run high-resolution, four-dimensional simulations of ignition processes.

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 to diesel or other 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 a snapshot 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 process playing out in high resolution; they have to be able to chart how individual fluid particles and chemical "species" reactions influence the combustion reaction as a whole. In fact, when running their most recent simulations, the researchers observed how formaldehyde, a byproduct of the methane ignition reaction, has significant influence 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, 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."

The team ran two sets of simulation 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

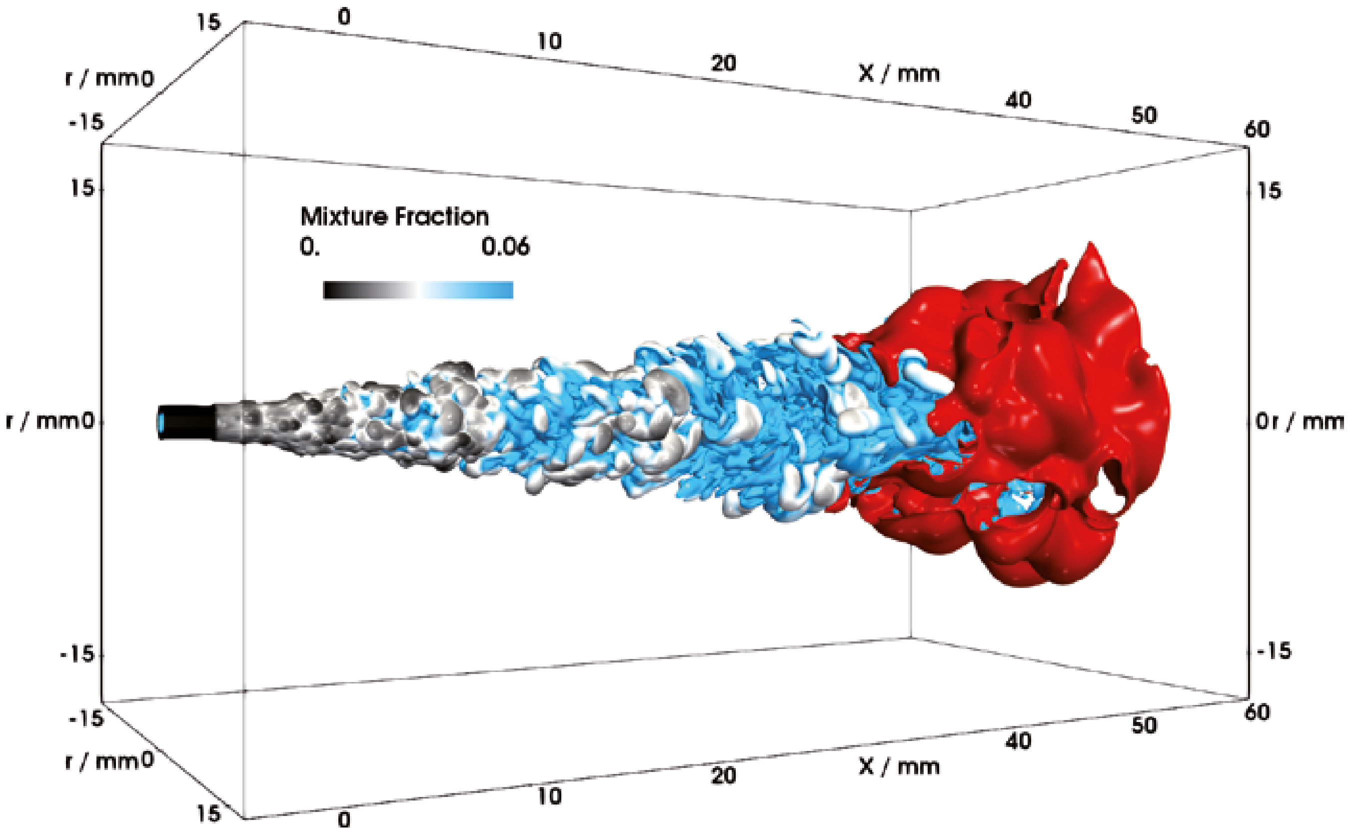

Using Hazel Hen, the team ran DNS simulations of to improve efficiency and reduce pollution in combustion engines.

© Eray Inanc, Uni Duisburg

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 in the GCS centres, 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 due to the GCS initiative."

Gaining ground

With the team's methane research completed, it looks to next-generation supercomputers to take their concept and apply it to more complex fuels. Kempf and Inanc indicated 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 throughout 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 prepares for its next-generation supercomputer, Hawk, to come online in the first half of 2020, Kempf is confident the team will once again be able to take advantage of the additional computational muscle.

"Our code has been ported to several different architectures, and we feel good about our abilities to port it to new architectures," he said. "Our impression of Hawk is that it is a well-rounded, balanced system design. Much like how the GCS centres upgrade their systems in a 'round robin' fashion, we port our code to all three centres to ensure it remains portable, and we are optimizing it all the time." *eg*

EUROCC AND CASTIEL WILL BUILD EUROPEAN NETWORK OF HPC EXPERTISE

In two new EU-wide projects, HLRS will coordinate an effort on behalf of GCS to enhance HPC competencies, promote cooperation, and support the implementation of best practices across Europe.

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

As EuroHPC developed, however, it became clear that one essential component of a comprehensive European strategy was missing: namely, a consistently high level of expertise in HPC and related disciplines such as high-performance data analytics and artificial intelligence.

To begin addressing this need, EuroHPC, under the auspices of the European Union's Horizon 2020 fund, has approved a proposal from the High-Performance Computing Center Stuttgart (HLRS) and the Gauss Centre for Supercomputing (GCS) to create a Europe-wide network of national HPC competence centers, called EuroCC.

To date, 33 participating member states have each designated one HPC center as a national competence center. The EU is providing approximately half of the funding to support the creation of each competence center, while the other half will come from the individual country. The first phase of the project will last two years, starting on September 1, 2020.

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

Simultaneously, HLRS, together with GCS, CINECA, TERATEC, the Barcelona Supercomputing Center, and PRACE, will lead a closely related coordination and support activity called CASTIEL (Coordination and Support for National Competence Centres on a European Level). While EuroCC coordinates HPC expertise within each member state, CASTIEL will promote the exchange of expertise across the entire EuroCC network.

Following the inventories of competencies at the national level, CASTIEL will develop a Europe-wide competency map that will catalog available resources and knowledge gaps across all EuroCC competence centers. It will then coordinate activities such as international workshops, mentoring and twinning partnerships, and topic-specific working groups to address issues of shared interest. Such activities might focus, for example, on advanced programming methods or best practices in HPC center operation. The improved international dialogue could also lead to regional collaborations between centers in which resources are shared to address complementary needs.

Because HPC centers in different countries have different levels of engagement with industry, CASTIEL will also facilitate the exchange of knowledge and expertise necessary to address HPC needs for industrial R&D. This could include mentoring by centers with experience supporting industry, as well as international meetings in which industrial HPC users present successful applications. Such meetings could not only enable national competence centers to better understand the specific needs of industrial HPC users, but also help them demonstrate to local stakeholders the benefits of investing in HPC technologies and resources.

"We hope that EuroCC and CASTIEL will empower the participating countries to develop the expertise they need to become more technologically independent," says HLRS Managing Director Bastian Koller, who is leading the coordination of the two projects. "Countries with a broad base of existing expertise will also learn from these international interactions. We see a lot of potential to raise the productivity and impact of HPC all across Europe for academic research, industry, and public administration. *cw*

EUROCC/CASTIEL

Runtime (Funding Period):
Sept. 2020 – Sept. 2022

Funding Source:
EuroHPC JU Program

Funding:
up to 27,942,672 Euro / up to 1,999,881 Euro

Partner Organizations:
33 national centers / 5 centers + PRACE'

EU PROJECT SANO ON COMPUTATIONAL MEDICINE

EU-funded project looks to employ supercomputing as part of establishing a centre for personalized medicine in Poland.

Forschungszentrum Jülich is a partner in the 7-year, European Union-funded project Sano. The project, which officially started in August 2019, received €15 million funding from the H2020-WIDESPREAD-2016-2017 Teaming Phase 2 program, with matching funds from Poland. The project aims at establishing, developing, and sustaining a Centre for Computational Personalised Medicine in Kraków, Poland.

The mission of Sano involves developing new computational methods, algorithms, models, and technologies for personalized medicine; introducing new diagnostic and therapeutic solutions based on computerized simulations into clinical practice; fostering creation and growth of enterprises which develop cutting-edge diagnostic and therapeutic technologies; and contributing to novel training and education curricula which meet the needs of modern personalized medicine.

During the first phase, the Polish Ministry of Science and Higher Education is working through its National Centre for Research and Development (NCBR) to coordinate the consortium of the six partners from Poland, the United Kingdom, and Germany:

- AGH University of Science and Technology, Kraków with its academic computing centre Cyfronet
- University of Sheffield with its Insigneo institute for in silico medicine
- Klaster LifeScience Kraków Foundation
- Fraunhofer Gesellschaft with its Fraunhofer-Institut für System- und Innovationsforschung ISI
- Forschungszentrum Jülich with the Jülich Supercomputing Centre
- Polish National Centre for Research and Development

Based on the partners' previous research, a rich assortment of complementary modelling approaches is already available, including models of fluid mechanics for cardiovascular and respiratory applications, as well as finite element analysis, growth and remodelling frameworks, and agent-based models of structural mechanics for cardiovascular, musculoskeletal and oncological applications.

Starting from these initial models, research at Sano will establish a number of research groups covering all fundamental topics related to computational medicine: modeling and simulation, high-performance computing (HPC), data science, in silico techniques, healthcare informatics, and algorithmic decision science. Close ties to regional hospitals also provide a unique environment for developing and testing novel computational biomarkers, in silico methods of patient subgrouping, treatment planning and clinical trials, as well as novel clinical decision support systems for diagnosis and therapy.

To these ends, Sano is closely cooperating with hospitals and enterprises from the growing health and life science sector in and around Kraków. Professional business development and project management staff allows the CoE to efficiently interact with entrepreneurs and large enterprises alike thereby covering all technology readiness levels from early research to regulatory approved clinical solutions.
Sano has already filled some initial staff positions. It is currently in the process of recruiting a director of the centre and plans to establish five new research groups by 2021. Research positions at Sano are filled by way of international competitions organized by the centre's International Scientific Committee, which consists of 15 experts including Prof. Dr. Dieter Kranzlmüller, Chairman of the Board of Directors of LRZ, and Prof. Dr. Holger Gohlke, NIC Research Group Leader at JSC. Jülich's contributions, in particular in the areas of data management, computational pipelines, and high-level services, are coordinated by the SimLab Biology at JSC. Further information can be found at https://sano.science. *oz*

SANO

Runtime: 2019 – 2026

Funding Source:
European Union's Horizon 2020 Program and Foundation for Polish Science

Funding:
30,000,000 Euro in total
(50% from H2020 + 50% from Poland)

Project Partners:
AGH University of Science and Technology, Kraków, University of Sheffield, Klaster LifeScience Kraków Foundation, Fraunhofer Gesellchaft, Forschungszentrum Jülich, Polish National Centre for Research and Development

COMPBIOMED 2 PROJECT HELPS SIMULATE THE HUMAN BEING

After a successful first funding period, the H2020-funded project CompBioMed enters phase 2.

CompBioMed2 is aiming to using computational tools to simulate "the virtual human." **© CompBioMed**

In the wake of the COVID-19 pandemic, SuperMUC-NG is helping researchers fight coronavirus. The supercomputer at the Leibniz Supercomputing Centre (LRZ) in Garching is currently busy filtering out active compounds from drug libraries and millions of chemical substances, natural remedies, or approved drugs that can be used to defeat the coronavirus, or SARS-CoV-2. "Supercomputers are a remarkable resource for the development of COVID-19 treatments," Peter Coveney, Director of the Centre for Computational Science (CCS) at the University College London (UCL), says. Coveney has extended experience working in this realm, and served as lead on the EU-funded projects CompBioMed and CompBioMed 2. "These machines help us identify possible treatments through a variety of ways, including machine learning, complex molecular dynamics, and artificial intelligence methods. Not only do we need to find molecules that bind to the spikes on the coronavirus, but we also need to model how well these bind when we know the spikes move around."

Digital twins for personalized medicine

The coronavirus crisis underscores the societal benefit of CompBioMed. Over the last five years, scientists and universities from Europe and the USA, under Coveney's leadership, have developed around 20 high-performance computing (HPC) programs, algorithms, and applications that can be used to create a virtual human model. These programs can be used to calculate and simulate the movement of muscles and bones, biochemical and molecular processes, as well as phenomena in the circulatory system. The codes and programs process big data sets to create a "virtual human being." This essentially means creating a "digital twin," or virtual representation of a dynamic, real-world process—such as the myriad functions happening inside the human body. These digital twins, when infused with quality data, serve as a starting point for personalized medicine.

In the fight against COVID-19, CompBioMed's tools are now paying off: the simulation programs help to model the coronavirus and its reproduction behavior, test molecules for their efficacy, and discover antidotes as quickly as possible. Thanks to HPC, detailed models of the coronavirus were available only weeks after the outbreak of COVID-19, and HPC helped quickly identify the first 40 active substances with which the virus could be fought.

By using HPC, researchers are working on true-to-scale models of the human body. Dr. Peter Coveney described the effort as "an expedition to Mount Everest, and we are moving toward Camp 2." © CompBioMed

Human model in HD resolution

CompBioMed's successor project, CompBioMed 2, started in 2019. While LRZ has supported the project since 2016, it became a core partner in 2019, and the organization is responsible for the analysis and management of large amounts of data.

In the second phase of the project, researchers will use supercomputers and artificial intelligence workflows to visualize three-dimensional, true-to-scale models of the human body in high resolution. The first milestones have already been reached: with data from the Swiss IT'IS Foundation and the CompBioMed blood flow program HemeLB, researchers used SuperMUC-NG to visualize the flow of red blood cells in the veins. "Compared to an expedition to Mount Everest, we are now moving towards Camp 2," Coveney said when describing the status of CompBioMed. HemeLB was the first application to use all 311,040 compute nodes of SuperMUC-NG.

Data management support

In addition to HPC consulting, training and support, LRZ will also help process ComBioMed data and catalogue it with metadata. This helps researchers more easily access relevant data sets online. The LRZ has already set up the storage system, LRZ Data Science Storage, for similar tasks, and offers tools for automated high-throughput analysis of huge amounts of data and for data security.

In CompBioMed 2, LRZ will also cooperate with EUDAT, the Open European Science Cloud. "This is an interesting development, because EUDAT offers a rich toolbox for data management, which we might be able to combine with our services," says Stephan Hachinger, head of the Research Data Management team at LRZ. CompBioMed 2 can benefit from the experience that the LRZ team is already gathering in data management projects such as LEXIS or Generic Research Data Infrastructure (GerDI), and vice versa.

Smarter through artificial intelligence

CompBioMed 2 also incorporates big data and artificial intelligence specialists at the LRZ. "Machine learning has arrived in research and medicine; it is used everywhere and complements classical computer simulation," says Peter Zinterhof, member of the Big Data and Artificial Intelligence team at LRZ, "By examining the functioning of trained networks, we can learn more about the inner workings of the problem." Zinterhof, who is also in charge of the DigiMed project on digitization in Bavarian medicine at the LRZ, expects that artificial intelligence and machine learning, image recognition and other smart systems will therefore become more widespread in medicine. *sv*

CompBioMed 2
(Cluster of Excellence)

Runtime:
October 1st, 2019 until September 30., 2023

Funding Organisation:
European Union's Horizon 2020 research programme.
Grant Agreement: 823712

Funding:
8,350,000 Euro

Core partners:
16 EU partners
+ 2 international partners

JSC HOSTS 10TH NIC SYMPOSIUM

Biannual review event brings together JSC users and staff for multi-day interactive conference of presentations and discussions.

JSC Director Prof. Thomas Lippert addresses the NIC Symposium, speaking to the importance of exascale computing.

The John von Neumann Institute for Computing (NIC) supports research projects from a broad scientific spectrum, including astrophysics, biology and biophysics, chemistry, elementary particle physics, materials sciences, condensed matter, soft matter sciences, earth and environmental sciences, computer science, numerical mathematics, fluid mechanics and plasma physics. On February 27–28, 2020, computational scientists presented their research results at the 10th NIC Symposium in Jülich. About 20 years after the first NIC Symposium took place in December, 2001, this series of biannual meetings has established a valued tradition, highlighting a diverse range of some of the best, modern, computational science at the NIC. Like the years before, the symposium was well- attended this year with over 200 participants.

Director of the Jülich Supercomputing Centre (JSC) and the NIC Prof. Thomas Lippert welcomed all symposium participants at the beginning of the event. In his address, Prof. Lippert emphasized the enormous importance of exascale computing, which will shape the high-performance computing (HPC) landscape in the coming years, as well as gave some insights into the development of exascale computing at JSC. Additionally, he also provided information on the latest status of the JUWELS Booster Module, planned to be installed and made available to the scientists during the course of this year.

Researchers across the scientific spectrum presented results and outcomes both during 13 detailed talks throughout the event as well as through an astounding 120 posters. The poster session served as fertile ground for discussion between researchers, providing plenty of time and space to exchange ideas and experiences on JSC resources in an interdisciplinary environment. Following tradition, researchers' recent activities and results were compiled in the NIC proceedings. The symposium and the proceedings address both computational scientists and practitioners as well as the general public that is interested in the advancement of computational science and its applications in diverse, contemporary research fields.

All accompanying materials such as the program, talks, posters, proceedings, and photographs are available at http://www.john-von-neumann-institut.de/nic/nic-symposium-2020. *at*

HLRS ADDRESSES SOCIAL AND ETHICAL IMPLICATIONS OF AI AND MACHINE LEARNING

Philosophers of science and technology contribute to development of new framework for evaluating design and effects of algorithms.

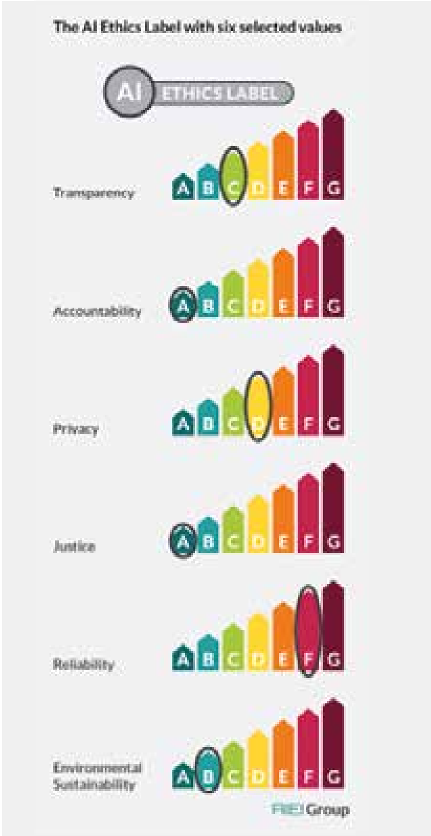

The AIEIG's framework uses a values scoring system similar to that seen in energy efficiency ratings. © AIEIG

Artificial intelligence (AI), deep learning, and machine learning offer opportunities for faster innovation. However, as they become more common, concerns have grown that a lack of transparency in their design and use could exacerbate societal problems and inequities. As AI infiltrates our daily lives, how could it reliably incorporate society's needs and values?

The High-Performance Computing Center Stuttgart (HLRS) Department of Philosophy of Science and Technology of Computer Simulation is well poised to investigate this question. By bringing philosophers, historians of science, and social scientists together with scientists specializing in high-performance computing, it facilitates critical, informed discussion about emerging digital technologies like AI within their broader cultural contexts.

Late last year, for example, the HLRS philosophy department hosted a conference called "The Society of Learning Algorithms," which brought together specialists across disciplines to discuss epistemological and ethical implications of artificial intelligence, historical perspectives on the rise of AI, and the political consequences of AI models.

The concerns of the HLRS philosophy department are not purely theoretical, however. More recently, department leader Dr. Andreas Kaminski and his colleagues participated in the AI Ethics Impact Group (AIEIG), which was created to develop a standardized framework for defining and rating ethical aspects of AI applications.

The AIEIG was organized by the VDE (Association for Electrical, Electronic, and Information Technologies) with the Bertelsmann Stiftung. HLRS was a key contributor, along with researchers from the Karlsruhe Institute of Technology, Tübingen University, the Technical University Darmstadt, the Technical University Kaiserslautern, and the iRights.Lab think tank.

VDE creates technical standards for electronics products of all kinds, and so the AIEIG's goal was to develop a universally applicable AI ethics ratings system similar to that used in promoting energy efficiency. In late March, the group published a working paper to provide technologists, regulators, and the general public with a rigorous framework for ensuring that AI applications' design and use are consistent with a society's values.

"In developing our recommendations, we tried to avoid several potential traps," said Kaminski. "One approach might be to implement rules in the algorithm that would govern how it should make decisions, but it would be impossible to find consensus that would be internationally acceptable for doing so across all application contexts. A second approach might be to create advisory boards inside companies, but it's not hard to imagine them quickly becoming pro forma and ineffective. We chose a third way that aims to provide a clear, consistent framework people can use to orient themselves, and that at the same time is measurable and enforceable."

The AIEIG based its approach on a model invented by TU-Darmstadt philosopher Christoph Hubig, a member of the working group. Once the group defined a set of desirable values for ethical AI applications—transparency, accountability, privacy, justice, reliability, and environmental sustainability—the model provided an easy-to-use system for assessing how well an algorithm incorporates them.

The report proposes a rating scale of A to G for each of the key values tracked in the framework. Once the evaluation is complete, an AI Ethics Label similar to that used in energy efficiency ratings would make it easy for others to quickly understand the ethical strengths and weaknesses of an AI tool.

In addition, the approach involves a comprehensive risk assessment, looking at the degree of potential harm that an algorithm could cause and how dependent humans could become on it. This perspective is intended to enhance the values assessment by helping potential users of the framework to identify ethical problems in an algorithm's particular social and cultural context.

The AIEIG plans to work with other stakeholders to refine the ethical criteria proposed in their new framework and to contribute to the implementation of universal AI standards. *cw*

BAVARIAN QUANTUM COMPUTING EXCHANGE BRINGS RESEARCHERS TOGETHER TO EXPLORE COMPUTERS OF THE FUTURE

Monthly meeting brings industry and academia together for a quantum computing ideas exchange.

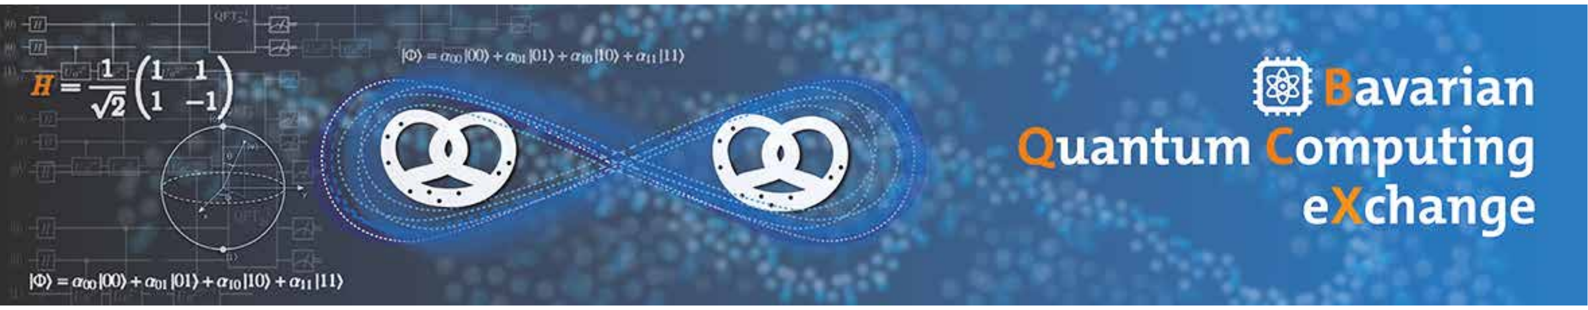

They are expected to process large amounts faster than ever. Although the technology still poses many questions and challenges, the future of quantum computing could both serve as a paradigm shift in the world of high-performance computing while simultaneously greatly strengthen supercomputers' abilities to solve humanity's most intractable problems. In an effort to bring stakeholders and computing experts together to discuss this seismic shift in computing, LRZ began the Bavarian Quantum Computing eXchange (BQCX) "We wanted to understand the interests of academia and industry with regard to quantum computing, and quickly realized we needed to facilitate an exchange of knowledge," said Laura Schulz, head of LRZ's strategic development. "We bring in top-notch experts for this community to gain knowledge, foster local networking for partnership building, and leverage the community's input to build an innovative applied quantum computing program for the Bavarian research ecosystem."

Exchange of methods and technology

Over the last several months, LRZ set up a monthly get-together that has attracted scientists from Munich and Bavarian universities as well as representatives of international companies like Intel, Google, D-Wave and Strangeworks. Participants provide insights into ongoing research and discuss technical approaches. "This helps us understand what options are available now and what the goals are moving forward," says Dr. Luigi Iapichino, physicist and co-founder of BQCX. "We are looking at hardware solutions that are already available to see which algorithms, applications and software are suitable for the near and mid-term".

Quantum computers make use of quantum mechanics effects like superposition and entanglement, which are extremely short-lived phenomena that are still difficult to control for macroscopic systems such as large supercomputers made up of many qubits. Prof. Rudolf Gross from the Walther-Meissner-Institute in Munich, Prof. Sabine Tornow from the University of Applied Sciences in Munich, and Dr. Fabio Baruffa from Intel are interested in the extraordinary computing power and in the possibilities of how qubits could be stabilized. Experiments have been conducted with chemical, electromagnetic or optical methods. The Australian startup Quantum Brilliance uses nitrogen vacancies in diamonds to build up qubits. Initial tests have shown that their systems can operate at room temperature. Temperatures around minus 270 degrees Celsius are otherwise common in quantum computing hardware. The most prominent application of this technology so far, resulting in the "quantum supremacy" paper released 2019, was presented by Alan Ho of Google.

Simulating quantum technology with supercomputers

Quantum computing is still in its infancy, but. Companies such as IQM, Cambridge Quantum Computer, IBM, and Intel are working on the first processors and tools. "The BQCX community helps us to advance our work on SuperMUC-NG now as well as prepare for exascale and even post-exascale systems," Iapichino says. "Just like with GPUs, quantum accelerators and their integration into extreme-scale HPC systems is a topic we need to think about now as we consider desigining future systems to meet our users' needs."

The exchange of knowledge works: the first courses for quantum computing were held in cooperation with Google CIRQ and Intel Quantum Simulator, and more workshops are in preparation. Together with the Walter-Meissner-Institute and the Max Planck Institute for Quantum Optics, LRZ is a strong part of the Bavarian quantum offensive. In

Dr. Luigi Iapichino, co-founder of BQCX and LRZ quantum team lead, introduces the January meeting's presenters from IQM (Finland and Germany), Quantum Brilliance (Australia), and Google (USA). **© LRZ**

addition, the PRACE activity "Future-Oriented Software Solutions" also benefits. Together with the Irish Centre for High-End Computing (ICHEC), LRZ works on the "QuantEx" project, which acts as an evaluation of tools for simulating quantum circuits running on heterogeneous computing platforms and are intended to help prepare the future of computers. "Participants at BQCX can present projects and find partners for collaboration", Schulz said. "We hope this nexus point proves greatly beneficial for the Bavarian quantum ecosystem." *sv*

UNSEEN PROJECT SEEKS TO FURTHER IMPROVE MODELLING ENERGY GRIDS OF THE FUTURE

Energy system optimization modelling has become a key ingredient in transitioning to decarbonized energy supply systems based mostly on renewables. Yet these systems reveal a growing complexity, such as decentralized infrastructures or an increasing variety of potential technologies capable of balancing energy demand and supply. These complications make traditional optimization modeling techniques impossible. In the project UNSEEN, several partners, including JSC, are engaged in developing model-oriented and algorithmic approaches tailored explicitly for the use of high-performance computing (HPC) resources. The prior project, BEAM-ME, confirmed the potential of this approach and demonstrated the need for further development. UNSEEN seeks to profit from artificial intelligence methods in order to further speed-up and facilitate the treatment of large numbers of scenarios. It will also address the crucial issue of reducing uncertainties when searching for adequate setups of a future energy system in Europe.

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, busses, 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 Euro.

HLRS CERTIFIED FOR ENVIRONMENTAL AND ENERGY MANAGEMENT

HLRS has received certification under the Eco-Management and Audit Scheme (EMAS), the most demanding system for sustainable environmental management worldwide. It is the first major high-performance computing center to receive EMAS certification. The accomplishment is the culmination of a multiyear effort to develop and implement a comprehensive sustainability concept that guides HLRS's operation and will help shape its future growth. Organizations that receive EMAS certification commit themselves to maintaining stringent environmental standards and to continually reduce their environmental impact over time. The new EMAS certificate constitutes an enhancement of HLRS's other formal commitments to environmental responsibility, including certification for environmental management under the ISO 14001 norm and for energy management under the ISO 50001 framework, which it received in November. HLRS's environmental guidelines include considering environmental impact in purchasing decisions, reducing energy consumption, minimizing waste, reusing resources, supporting research that leads to sustainability gains, and sharing its knowledge about environmental management with other HPC centers.

SHAPING THE FUTURE OF SHARED MEMORY PROGRAMMING

Since April 2019, Volker Weinberg, PhD physicist and coordinator of the education and training program at LRZ, is a member of the OpenMP Architecture Review Board (ARB), a standardisation committee for one of the two main paradigms in parallel programming. The ARB has more than 30 members, with representatives from some the biggest supercomputing centres of the world as well as manufacturers like AMD, ARM, Cray, IBM, Intel, and others. “If we actively participate in the standardisation of OpenMP, we help to determine the future of high performance computing,”. Weinberg said. Currently, the ARB is planning on finalising Version 5.1 of OpenMP. Due to the coronavirus pandemic, the next meeting, planned for May in Garching, will take place virtually. Final decisions will be taken at the end of July to modernise the standard.

ARKTIK – NEW ADSORPTION COOLING PROJECT AT LRZ

On March 1, LRZ started its new adsorption cooling project ArKtIK (Adsorptions-Kältetechnologie für Informations-Technologie-Kühlung, or Adsorption Cooling Technology for Information Technology Cooling). While adsorption cooling is a potential technology to re-use the excess heat of today's mostly water-cooled HPC systems, it still requires adoption to the particular usage scenarios in supercomputing centres. LRZ, a pioneer in hot water and adsorption cooling in supercomputing, is investigating on how to improve the efficiency of adsorption cooling and ease its deployment with ArKtIK. More details of the technology and the project can be found here (in German): https://www.lrz.de/forschung/projekte/forschung-hpc/arktik/

FIRST FENIX SERVICES OPERATIONAL AT JSC

A new cluster system, named JUSUF, has gone into operation at JSC. The system from Atos is part of the JSC's engagement in building up Fenix, a pan-European federated e-infrastructure providing computing, data, and cloud services to the Human Brain Project (HBP) and other science communities with similar requirements. JUSUF features 205 nodes with AMD EPYC Rome processors. A distinctive feature of the cluster is a setup that allows nodes to be either part of an HPC cluster or an OpenStack cluster. This enables the system to flexibly deliver different types of services within Fenix: Scalable computing services, interactive computing services, and Virtual Machine services. The distributed e-infrastructure Fenix is built by a consortium of the five European HPC centres BSC, CEA, CINECA, CSCS, and JSC within the Interactive Computing e-Infrastructure (ICEI) project, co-funded by the European Union. Fenix forms the basis of EBRAINS, a shared digital research infrastructure for brain research developed by the HBP. Fenix services, including JUSUF resources, are available to HBP members (through EBRAINS), European scientists from all domains (through PRACE), and users from Germany. Details on how to get access can be found at https://fenix-ri.eu/access.

PRACE ADA LOVELACE AWARD FOR DR. ALICE-AGNES GABRIEL

Dr. Alice-Agnes Gabriel won the PRACE Ada Lovelace Award for HPC for her outstanding contributions to and impact on HPC in Europe. Dr. Gabriel is a lecturer at the Chair of Seismology in the Institute of Geophysics at Ludwig-Maximilians-Universität in Munich. She uses the supercomputers at GCS SuperMUC and SuperMUC-NG to simulate earthquakes. "Dr. Alice-Agnes Gabriel uses numerical simulations coupled to experimental observations to increase our understanding of the underlying physics of earthquakes. The work includes wide scales and can improve our knowledge and safety against these natural phenomena." said Núria López, Chair of the PRACE Scientific Steering Committee.

DATA TRANSFER AT HIGH SPEED – EVEN ACROSS THE ATLANTIC

With the Data Science Storage (DSS) Service, LRZ offers an innovative service that allows researchers and staff to share very large amounts of data fast – not only within the LRZ ecosystem, but also worldwide via an easy-to-use graphical user interface on the web. This system was tested with LRZ staff sharing data with scientists in California, 9,000 kilometers away, at record-breaking speed . In conjunction with Globus Online, DSS enabled the transfer of 500 terabytes of data from Garching to NERSC at the Lawrence Berkeley National Laboratory in California at a speed of 4.5 GB per second. The high speed of data transfer over wires and across continents is possible because NERSC has similarly powerful network and storage systems.

LAUNCH OF JUNIQ AND CONTRACT WITH D-WAVE SYSTEMS

On 25 October 2019, the "Jülich Unified Infrastructure for Quantum Computing (JUNIQ) was put into operation. The go-ahead was given with the official signing of a user contract for a quantum annealer from the Canadian quantum computer manufacturer D Wave Systems, making Forschungszentrum Jülich the first D-Wave Leap™ cloud-based quantum site in Europe (for more information, please visit page 7). The event was opened with a welcoming address by Annette Storsberg, State Secretary from the North Rhine-Westphalian Ministry of Culture and Science, and a message from Thomas Rachel, Parliamentary State Secretary at the Federal Ministry of Education and Research. The state government and the Federal Ministry of Education and Research are each supporting the establishment of JUNIQ with 5 million Euro in funding. JUNIQ will be the unified portal to a number of different quantum computers—accessible via the cloud for German and European users. It will thus offer quantum computing services similar to those that have long been available for Jülich's supercomputers. The construction work for a new JUNIQ building that will house the D-Wave quantum system started in April 2020 in the vicinity of the Jülich Supercomputing Centre..

© Simon Sommer for HLRS.

LANDESCLOUD TO OFFER SECURE, CLOUD-BASED AI AS A SERVICE FOR INDUSTRY

HLRS has begun 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 during all stages of data transfer, storage, processing, and analysis. It is also designed to support collaborative AI projects, including teamworking, crowd-working, and data science competitions. The platform'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.


Universität Stuttgart
Nobelstraße 19
Höchstleistungsrechen-
zentrum Stuttgart (HLRS)
Anlieferung
HPC Training Center
← Eingang links

STAFF SPOTLIGHT: INTERDISCIPLINARY COLLABORATION FOR ACCELERATING APPLICATION PERFORMANCE

Dr. Huan Zhou, HLRS

As far back as middle school, Dr. Huan Zhou knew that she loved working with mathematics.

As she got older, though, she realized that she wanted to pursue a career where her passion for math would be applied to compelling real-world problems. To that end, she decided to major in computer science in Hefei. "I really preferred to learn something that combined theory with the practical side of things, so I wanted to choose computer science," she said. During her undergraduate studies, she was exposed to a wide variety of computer science topics, ranging from hardware and microcomputer principles to programming.

As she progressed through her masters' studies at the University of Science and Technology of China, though, an advisor introduced her to the world of parallel computing, and she immediately took an interest in working at the cutting edge of computer technologies. Shortly after graduating, she received an opportunity to work at Intel, and although the opportunity provided valuable experience in the computer industry, her thoughts always returned to parallel computing and the mystique of working with the world's fastest computers, ultimately helping to accelerate applications pushing the limits of computational science.
In her own time spent researching the topic, she was intrigued by some of the scientific literature published by staff members at the High-Performance Computing Center Stuttgart (HLRS), which ultimately led to her contacting centre director Prof. Michael Resch about opportunities to continue her education while also working at a world-leading high-performance computing (HPC) centre.

Specifically, Zhou took an interest in parallel programming models, such as the message passing interface (MPI) library—a set of rules governing communications across processors and nodes for programs running on parallel computers. Once in Stuttgart, she threw herself deeper into that work, and felt her confidence and passion grow by virtue of working in a diverse, interdisciplinary environment. "In all the projects I've been involved with since coming here in 2013, I've always collaborated with people from really diverse scientific fields of knowledge," Zhou said. "For me, this kind of collaboration facilitates mutual learning and common progress. I think that is a particularly special part of being at HLRS—all of the people I work with have different background and fields of knowledge."

Zhou's thesis focused on novel approach for developing a communications system that would accelerate the performance for researchers' applications and hide many of the more complex and difficult programming aspects of using MPI efficiently from users as well.

After getting her PhD, Zhou's work at HLRS expanded to take the lessons learned through her doctoral research and applying to users employing HLRS's then-flagship computer, Hazel Hen, to solve some of humanity's biggest research challenges in subjects ranging from industrial engineering to elementary particle physics. "The challenge in my post-doctoral work comes from working with things that I'm unfamiliar with, like physics for example," she said. "There is a physics problem, and we need HPC to help solve it. That was difficult for me because in physics, for example, this was work I had never touched before. But I like that kind of challenge, because it makes our collaborations very cross-disciplinary. It helped me realize that you can be an expert in a specific area but still need support."

Early in 2020, HLRS inaugurated its new flagship computer, a Hewlett Packard Apollo 9000 machine called Hawk (See page 4), presenting exciting new opportunities for Zhou and her collaborators. The machine offers users a massive upgrade in peak performance to Hazel Hen, but also comes with new-generation processors from AMD, much larger nodes than were present on Hazel Hen (128 cores per node instead of 24), and new interconnects between nodes on the system. While this does present new challenges for Zhou and HLRS developers and researchers, it also provides a new opportunity to bring different competencies together overcome challenges and ultimately support the HPC user base. "Our goal has been that first we need to learn the system inside and out so we are in a position to best teach the users," she said.

"While architectures are always changing, users' needs are never changing—they always want to see more performance," she said. "Now that we are advancing to Hawk, things are more complex, but for the users, the first and most important thing is to make sure they can port their existing codes to Hawk. We are rising to the challenge to ensure that we can deliver what users are counting on—performance ultimately being better than before." *eg*

GCS SUPPORT STAFF FORGE CLOSER CONNECTIONS TO USERS THROUGH MENTORING CONCEPT

With a common approach that still remains flexible for centres' individual needs, user support staff members are getting closer to user projects. We check in with work package leader Dr. Wolfgang Frings, user support specialist at JSC, to learn more about the advancements in supporting GCS users.

When the Gauss Centre for Supercomputing (GCS) secured another decade of funding in 2017, leadership at all three centres made sure that substantial investments would go toward enhanced user support. With the rise of more diverse high-performance computing (HPC) architectures and the complex software stack that runs on their systems, HPC centres want to ensure that users are making the best possible use of the allocations they have been given.

That commitment resulted in a blossoming collaboration between the three GCS centres' user support staffs and the advent of a personalized mentoring concept that intends to bring the centres closer to their user base, ultimately helping to make users' experience in porting and optimizing codes efficient and being the best possible stewards of these large resources by ensuring they are being used as efficiently as possible.

Now that user support staff is moving from conceptualization to implementation of a tailored GCS user support concept, we sat down with Dr. Wolfgang Frings (JSC), leader of the GCS user support work package, to check in on the development of GCS's more robust, comprehensive user support program.

? *GCS: Can you summarize some of the major changes that have taken place over the last two years?*

WF: The GCS support concept has been our first deliverable, and I think its key element is the assignment of a qualified mentor to each individual user project as a single, reliable contact point between the user and the centre. Another major aspect of the concept is enabling more proactive support. In the past, HPC centres generally operated with the assumption that users are doing their work, and we were mainly contacted by them if they had problems. With the new mentoring concept, we want to have more contact between the user support and the researchers in the project, ultimately forging a real collaboration. We intend the mentor to contact the PI at the beginning of the computing time period and really discuss what is going to be done in the project, any concerns they may have, and also describe the offers we have, including help with optimization and training courses. We want to have proactive contact throughout the whole project.

From our side, with all the software tools we have available, we are able to see which applications aren't running as good as they could. Many basic types of optimizations are done in the initial project or a dedicated preparatory phase. For example, we offer to run performance audits developed by the EU-funded Centre of Excellence POP (Performance Optimization and Productivity), which realistically estimates the performance quality of applications. We now try to do it the same way in collaboration with POP, so at the end, the user has a report giving hints which parts of the application are running well and which parts need to be improved.

For us, the goal is to run efficient user codes on our systems, and for our users, the goal is utilizing their allocation as efficiently as possible. However, for many users, as they are also under pressure to produce research results, they often do not have time to optimize their application. In some cases, users can take the approach, "if it runs, it runs." They may, rightfully, ask why they should feel obligated to spend six months optimizing a code for a small efficiency gain. The mentors help address this, as they can identify small-effort efficiency gains in the short term while also helping to create a roadmap for how we can continue to improve efficiency, for example, by running the initial performance audit. Typically, long projects run for several one-year periods. The same mentor will continue to work with these users and create a long-term collaboration on this kind of work.

Also, the mentor doesn't need to know everything about the project or how to solve every problem, but because the mentor is that single point of contact, he or she has the institutional knowledge to involve other people from the centre when necessary, such as with I/O or other specific optimization issues. Since the mentors understand who has specific skill and knowledge in their centres, they can help pointing the users to someone to help with a specific problem.

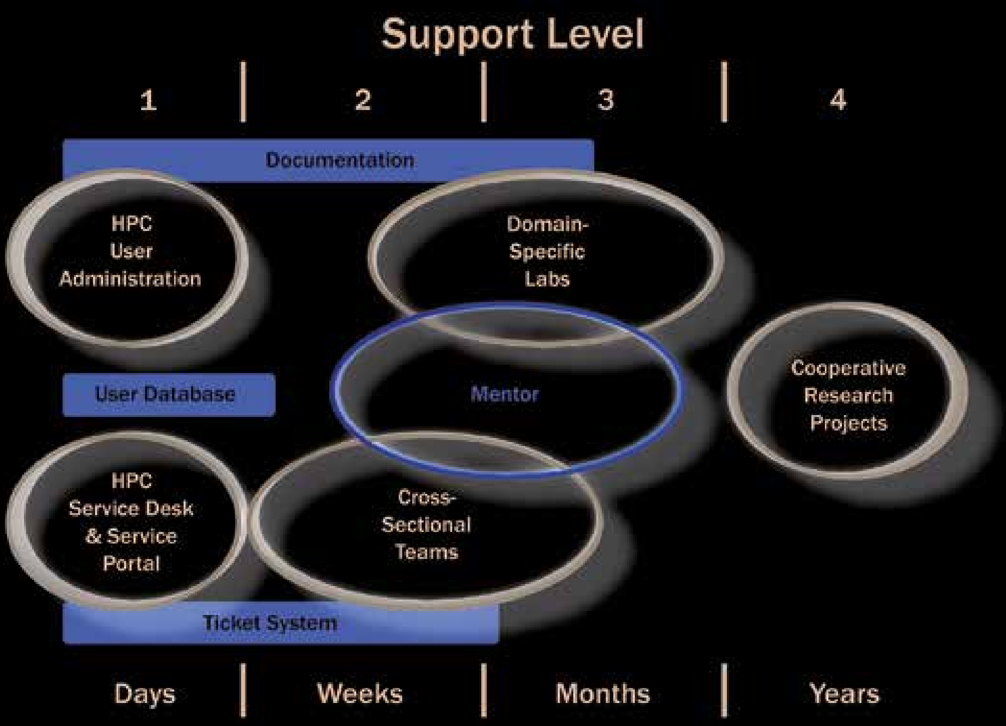

Building on well-established traditional user support services, GCS developed a new HPC application support structure that offers a series of several new HPC support levels such as mentoring—a new service available to any user with a large-scale project. Services are offered at 4 different levels. **© GCS**

? GCS: All three centres have their own approaches to user support, but GCS brings your working group together in a particularly close fashion. How has the relationship between the centres' respective user support staffs evolved over time?

WF: I think there are several aspects. First, there are the face-to-face meetings of the working group and staff associated with the GCS project across the three centres, which give us the chance to sit and talk together multiple times a year. Between these meetings, we also organize a lot of video conferences for detailed discussions on special topics. We started in our working group in 2018, and over the last two years, we have gotten closer together, both in getting to know each other on a personal level, but also using the same language and terminology with key aspects of support. Now we have more of a common understanding of how to discuss aspects of our work, and it makes it a lot easier to explain things to one another.

As a third part, we created a series of webinars where we identified topics that are essential for all three centres to exchange ideas for solutions, and these generally have around 20 people participating in them. We are now opening the webinars for partner institutions. It is an ongoing discussion about how to integrate these webinars with our training activities for the users. The next step we want to achieve with the webinars is to show what happens in practice with the mentoring concept. We want to have best-practice talks showing how we addressed specific problems and compare our results.

? GCS: What has changed with respect to users' needs over the last two years, and how did those needs inform our own decision making? How has GCS' architectural diversity across the three centres informed the user support concept?

WF: In recent years, there have been many innovations in HPC from the user's point of view, not only in terms of hardware, but also especially in programming of HPC system, the use of new types of application workflows as well as the increasing demand on data science solutions. This has been accompanied by increased demands for support. Based on the increased requirements, it is quite clear that cooperation between the three GCS centres is important. The working group in the context of the second round of GCS funding gave us the right framework for this kind of intensive cooperation. Support is important for both sides: It allows researchers to fully utilize their allocation and ensures the hosting centres to efficiently utilize their HPC systems.

When it comes to architectural changes, this is something that happens often, as we regularly have to integrate new hardware, or new software environments. This year for us at JSC, that change comes from installing the JUWELS booster module and getting a big performance jump with a large number of GPUs, but also the need to introduce users to this new approach. The mentoring concept can help here, because the mentors can get a better sense of the progress with respect to porting, and understanding their challenges. In terms of the three centres, we have to support different environments. For the support concept, that means we have to keep in mind having three "flavors" of the support, but thanks to the our work package as part of GCS, we have a common concept that helps us to address the needs of our users. *eg*

TRAINING CALENDAR
HPC COURSES AND TUTORIALS

Course / Workshop Title	Location	Date
Deep Learning and GPU programming using OpenACC	Vienna	(Jun 3-5) postp.
ONLINE: HPC code optimisation workshop (PRACE course)	Garching	Jun 8-10, 2020
High-performance computing with Python (PRACE course)	Jülich	(Jun 8-10) postp.
High-performance scientific computing in C++ (PRACE course)	Jülich	(Jun 15-17) postp.
ONLINE: Deep Learning and GPU programming Workshop (PRACE course)	Garching	Jun 15-18, 2020
Cluster Workshop	Stuttgart	(Jun 16-17) cancl.
ONLINE: Introduction to hybrid programming in HPC	Vienna	Jun 17-19, 2020
ONLINE: Efficient Parallel Programming with GASPI (PRACE course)	Stuttgart	Jun 18-19, 2020
ONLINE: Optimizing OpenCL Programs for Intel FPGAs	Garching	Jun 25, 2020
ONLINE: Intel OneAPI for FPGAs	Garching	Jun 19, 2020
ONLINE: Node-Level Performance Engineering (PRACE course)	Stuttgart	Jun 29-Jul 1, 2020
ONLINE: Optimization of Node-level Performance and Scaling on Hawk	Stuttgart	Jul 2-3, 2020
ONLINE: Intermediate C++ with Focus on Software Engineering	Stuttgart	Jul 7-10, 2020
Deep Learning and GPU programming using OpenACC	Stuttgart	Jul 14-17, 2020
Introduction to parallel programming with MPI and OpenMP	Jülich	Aug 10-14, 2020
Parallel Programming with MPI / OpenMP	Zurich	Aug 24-27, 2020
14th International Parallel Tools Workshop	Stuttgart	Sep 2020 (tbc)
Advanced Fortran Topics (PRACE course)	Garching	Sep 2020 (tbc)
Iterative Linear Solvers and Parallelization	Garching	Aug 31-Sep 4, 2020
Introduction to Computational Fluid Dynamics	Stuttgart	Sep 14-18, 2020
CFD with OpenFOAM®	Siegen	Sep 28 - Oct 2, 2020
Porting code from Matlab to Python	Jülich	Oct 5-6, 2020
Data Analytics / Machine Learning Courses	Garching	Oct 5-9, 2020
High Performance Computing in Science and Engineering	Stuttgart	Oct 8-9, 2020
Introduction to Python	Jülich	Oct 2020 (tbc)
Parallel Programming Workshop (MPI, OpenMP & advanced topics) (PRACE course)	Stuttgart	Oct 12-16, 2020
Scientific Visualization	Stuttgart	Oct 19-20, 2020
Directive-based GPU programming with OpenACC	Jülich	Oct 26-27, 2020
Intro. to ANSYS Fluid Dynamics (CFX, Fluent) on LRZ HPC Systems	Garching	Oct 26-30, 2020
Optimization of Node-level Performance and Scaling on Hawk	Stuttgart	Nov 2-6, 2020
From zero to hero, Part II: Understanding and fixing intra-node performance bottle-necks	Jülich	Nov 3-4, 2020
Software Development in Science	Jülich	Nov 10-11, 2020
Advanced C++ with Focus on Software Engineering	Garching	Nov 18-20, 2020
Introduction to the programming and usage of the supercomputer resources in Jülich	Jülich	Nov 23-24, 2020
Advanced C++ with Focus on Software Engineering	Stuttgart	Nov 24-27, 2020
Advanced Parallel Programming with MPI and OpenMP	Jülich	Nov 30-Dec 2
Node-Level Performance Engineering (PRACE course)	Garching	Dec 2020 (tbc)
Fortran for Scientific Computing	Stuttgart	Dec 7-11, 2020

VISIT INSIDE ONLINE FOR DETAILS

For a complete and updated list of all GCS courses, please visit:
http://www.gauss-centre.eu/training

The German HPC calendar (organized by the Gauss Allianz in cooperation with all German HPC centres) provides an extensive list of training all taking place German HPC centres. More information can be found at:
http://hpc-calendar.gauss-allianz.de/

Further training courses and events can be found on GCS member sites:
http://www.hlrs.de/training/
http://www.lrz.de/services/compute/courses/
http://www.fz-juelich.de/ias/jsc/events

The Rühle Saal at HLRS in Stuttgart **© HLRS**

JÜLICH SUPERCOMPUTING CENTRE FORSCHUNGSZENTRUM JÜLICH

The Jülich Supercomputing Centre (JSC) at Forschungszentrum Jülich is committed to enabling scientists and engineers to explore some of the most complex grand challenges facing science and society. Our research is performed through collaborative infrastructures, exploiting extreme-scale supercomputing, and federated data services.

Provision of supercomputer resources: JSC provides access to supercomputing resources of the highest performance for research projects coming from academia, research organizations, and industry. Users gain access for projects across the science and engineering spectrum in the fields of modelling and computer science.

- Supercomputer-oriented research and development in selected fields of physics and other natural sciences by research groups and in technology, e.g. by doing co-design together with leading HPC companies.
- Higher education for master and doctoral students in close coperation with neighbouring universities.
- Implementation of strategic support infrastructures including community-oriented simulation and data laboratories and cross-sectional teams, e.g. on mathematical methods and algorithms and parallel performance tools, enabling the effective usage of the supercomputer resources.

The Cluster module of JSC's Modular Supercomputer "JUWELS". © Forschungszentrum Jülich

Contact

Jülich Supercomputing Centre (JSC), Forschungszentrum Jülich
Prof. Dr. Dr. Thomas Lippert
Wilhelm-Johnen-Straße, 52425 Jülich, Germany
Phone +49 - 24 61 - 61 - 64 02
th.lippert@fz-juelich.de
www.fz-juelich.de/jsc

Compute servers currently operated by JSC

System	Size	Peak Performance (TFlop/s)	Purpose	User Community
Modular Supercomputer "JUWELS"	Cluster (Atos): 10 cells, 2,567 nodes 122,768 cores Intel Skylake 224 NVIDIA V100 GPUs 275 Tbyte memory	12,266	Capability Computing	European (through PRACE) and German Universities and Research Institutes
Modular Supercomputer "JURECA"	Cluster (T-Plattforms): 1,884 nodes, 45,216 cores Intel Haswell 150 NVIDIA K80 GPUs 281 TByte memory	2,245	Capacity and Capability Computing	German Universities, Research Institutes, and Industry
	Booster (Intel/Dell): 1,640 nodes 111,520 cores Intel Xeon Phi (KNL) 157 TByte memory	4,996		
Fujitsu Cluster "QPACE 3"	672 nodes, 43,008 cores Intel Xeon Phi (KNL) 48 TByte memory	1,789	Capability Computing	SFB TR55, Lattice QCD Applications
Atos Cluster "JUSUF"	205 nodes, 26,240 cores AMD EPYC Rome 61 NVIDIA V100 GPUs 52 TByte memory	1,372	Capacity Computing	European and German Universities and Research Institutes, Prace, Human Brain Project
Modular Supercomputer "DEEP-EST" (Prototype)	Cluster: 50 nodes, 1,200 cores Intel Xeon Gold 6146 9.6 TByte memory + 25.6 TByte NVM	45	Capacity Computing (low/medium-scalable code parts)	Partners of the EU-project DEEP-EST and interested users through Early Access Programme
	Booster: 75 nodes, 600 cores Intel Xeon Silver 4215 75 NVIDIA V100 GPUs 6 TByte memory	549	Capacity and Capability Computing (high-scalable code parts)	
	Data Analytics Module: 16 nodes, 768 cores Intel Xeon Platinum 8260 16 NVIDIA V100 GPUs 16 Intel Stratix10 FPGAs 7.1 TByte memory + 32 TByte NVM	170	Capacity and Capability Computing (data analytics codes)	

A detailed description can be found on JSC's web pages: www.fz-juelich.de/ias/jsc/systems

LEIBNIZ SUPERCOMPUTING CENTRE

Leibniz Supercomputing Centre
of the Bavarian Academy of Sciences and Humanities

For nearly six decades, the Leibniz Supercomputing Centre (Leibniz-Rechenzentrum, LRZ) has been at the forefront of its field as a world-class high performance computing centre dedicated to providing an optimal IT infrastructure to its clients throughout the scientific community–from students to postdocs to renowned scientists—and in a broad spectrum of disciplines–from astrophysics and engineering to life sciences and digital humanities.

Leadership in HPC and HPDA

Located on the research campus in Garching near Munich, the LRZ is a leadership-class HPC and HPDA facility delivering top-tier supercomputing resources and services on the national and European level. Top-notch specialists for HPC code portability and scalability support the LRZ' broad user base and ensure to run operations in the most energy-efficient way.

Future Computing at LRZ

The LRZ is leading the way forward in the field of Future Computing focusing on emerging technologies like quantum computing and integrating AI on large-scale HPC systems. A robust education program for HPC, machine learning, artificial intelligence and big data is complementing the LRZ offer.

IT backbone for Bavarian science

In addition to its role as national supercomputing centre, the LRZ is the IT service provider for all Munich universities as well as research organizations throughout Bavaria.

SuperMUC NG at the Leibniz Supercomputing Centre. **© LRZ**

Contact

Leibniz Supercomputing Centre (LRZ)
Prof. Dr. Dieter Kranzlmüller
Boltzmannstraße 1, 85748 Garching near Munich, Germany
Phone +49 - 89 - 358 - 31 - 80 00
kranzlmueller@lrz.de
www.lrz.de

Compute servers currently operated by LRZ

System	Size	Peak Performance (TFlop/s)	Purpose	User Community
"SuperMUC-NG" Intel/Lenovo ThinkSystem	6,336 nodes, 304,128 cores, Skylake 608 TByte, Omni-Path 100G	26,300	Capability Computing	German universities and research institutes, PRACE (Tier-0 System)
	144 nodes, 8,192 cores Skylake 111 TByte, Omni-Path 100G	600	Capability Computing	
"SuperMUC Phase 2" Lenovo Nextscale	3,072 nodes, 86,016 cores, Haswell EP 197 TByte, FDR 14 IB	3,580	Capability computing	German universities and research institutes, PRACE (Tier-0 System)
"CooLMUC-2" Lenovo Nextscale	384 nodes, 10,752 cores Haswell EP 24.6 TByte, FDR 14 IB	447	Capability computing	Bavarian Universities (Tier-2)
"CoolMUC-3" Megware Slide SX	148 nodes, 9,472 cores, Knights Landing, 17.2 TByte, Omnipath	459	Capability Computing	Bavarian Universities (Tier-2)
IvyMUC	Intel Xeon E5-2650 ("Ivy Bridge")	13	Capability Computing	Bavarian Universities (Tier-2)
Teramem	1 node, 96 cores, Intel Xeon E7-8890 v4 ("Broadwell"), 6 TByte RAM	13	Big Data	Bavarian Universities (Tier-2)
DGX-1, DGX-1v Machine Lerning Systems	2 nodes, Nvidia Tesla, 8 x P100, 8 x V100	1,130 (Mixed Precision)	Machine Learning	Bavarian Universities (Tier-2)
Compute Cloud for SuperMUC-NG	64 nodes, 3,072 cores, Intel Xeon ("Skylake"), 64 Nvidia V100	128, 8,000 (Mixed Precision)	Cloud	German Universities and Research Institutes, PRACE

A detailed description can be found on HLRS' web pages: https://doku.lrz.de/display/PUBLIC/Access+and+Overview+of+HPC+Systems

HIGH-PERFORMANCE COMPUTING CENTER STUTTGART

HLRS
High Performance Computing Center | Stuttgart

Based on a long tradition in supercomputing at University of Stuttgart, HLRS (Höchstleistungsrechenzentrum Stuttgart) was founded in 1996 as the first German federal centre for high-performance computing. HLRS serves researchers at universities and research laboratories in Europe and Germany and their external and industrial partners with high-end computing power for engineering and scientific applications.

Service for industry
Service provisioning for industry is done together with T-Systems, T-Systems sfr, and Porsche in the public-private joint venture hww (Höchstleistungsrechner für Wissenschaft und Wirtschaft). Through this cooperation, industry always has access to the most recent HPC technology.

Bundling competencies
In order to bundle service resources in the state of Baden-Württemberg HLRS has teamed up with the Steinbuch Centre for Computing of the Karlsruhe Institute of Technology. This collaboration has been implemented in the SICOS BW GmbH.

World class research
As one of the largest research centres for HPC, HLRS takes a leading role in research. Participation in the German national initiative of excellence makes HLRS an outstanding place in the field.

Hawk at the High-Performance Computing Center Stuttgart. **© Ben Derzian for HLRS**

Contact

High-Performance Computing Center Stuttgart (HLRS), University of Stuttgart
Prof. Dr.-Ing. Dr. h.c. Dr. h.c. Hon.-Prof. Michael M. Resch
Nobelstraße 19, 70569 Stuttgart, Germany
Phone +49 - 711 - 685 - 8 72 69
resch@hlrs.de
www.hlrs.de

Compute servers currently operated by HLRS

System	Size	Peak Performance (TFlop/s)	Purpose	User Community
HPE Apollo 9000 "Hawk"	5,632 nodes 720,896 cores 1,44 PB memory	26,000	Capability Computing	European (PRACE) and German Research Organizations and Industry
NEC Cluster (Vulcan, Vulcan 2)	622 nodes 18,736 cores 119 TB memory	1,012	Capacity Computing	German Universities, Research Institutes and Industry
NEC SX-Aurora TSUBASA	64 nodes 512 cores 3072 GB memory	137.6	Vector Computing	German Universities, Research Institutes and Industry
Cray CS-Storm	8 nodes 64 GPUs 2,048 GB memory	499.2	Machine Learning Deep Learning	German Universities, Research Institutes and Industry

A detailed description can be found on HLRS' web pages: www.hlrs.de/systems

InSiDE magazine (German: Innovatives Supercomputing in Deutschland) is the bi-annual publication of the Gauss Centre for Supercomputing, showcasing recent highlights and scientific accomplishments from users at Germany's three national supercomputing centres. GCS was founded in 2007 as a partnership between the High-Performance Computing Center Stuttgart, Jülich Supercomputing Centre, and the Leibniz Supercomputing Centre. It is jointly funded by the German Ministry of Education and Science (Bundesministerium für Bildung und Forschung – BMBF) and the corresponding ministries of the three states of Baden-Württemberg, North Rhine-Westphalia, and Bavaria.

Cover image: As part of the the CompBioMed 2 project, EU researchers are leveraging computational power at the Leibniz Supercomputing Centre and other HPC centres to create true-to-scale, high-resolution simulations of the human body. © CompBioMed

www.gauss-centre.eu

Gauss Centre for Supercomputing