InSiDE
Innovatives Supercomputing in Deutschland

Spring 2018 · Vol. 16 · No. 1

Understanding How the Universe Formed

IllustisTNG provides largest-ever galaxy formation simulation (p. 35 and p. 59)

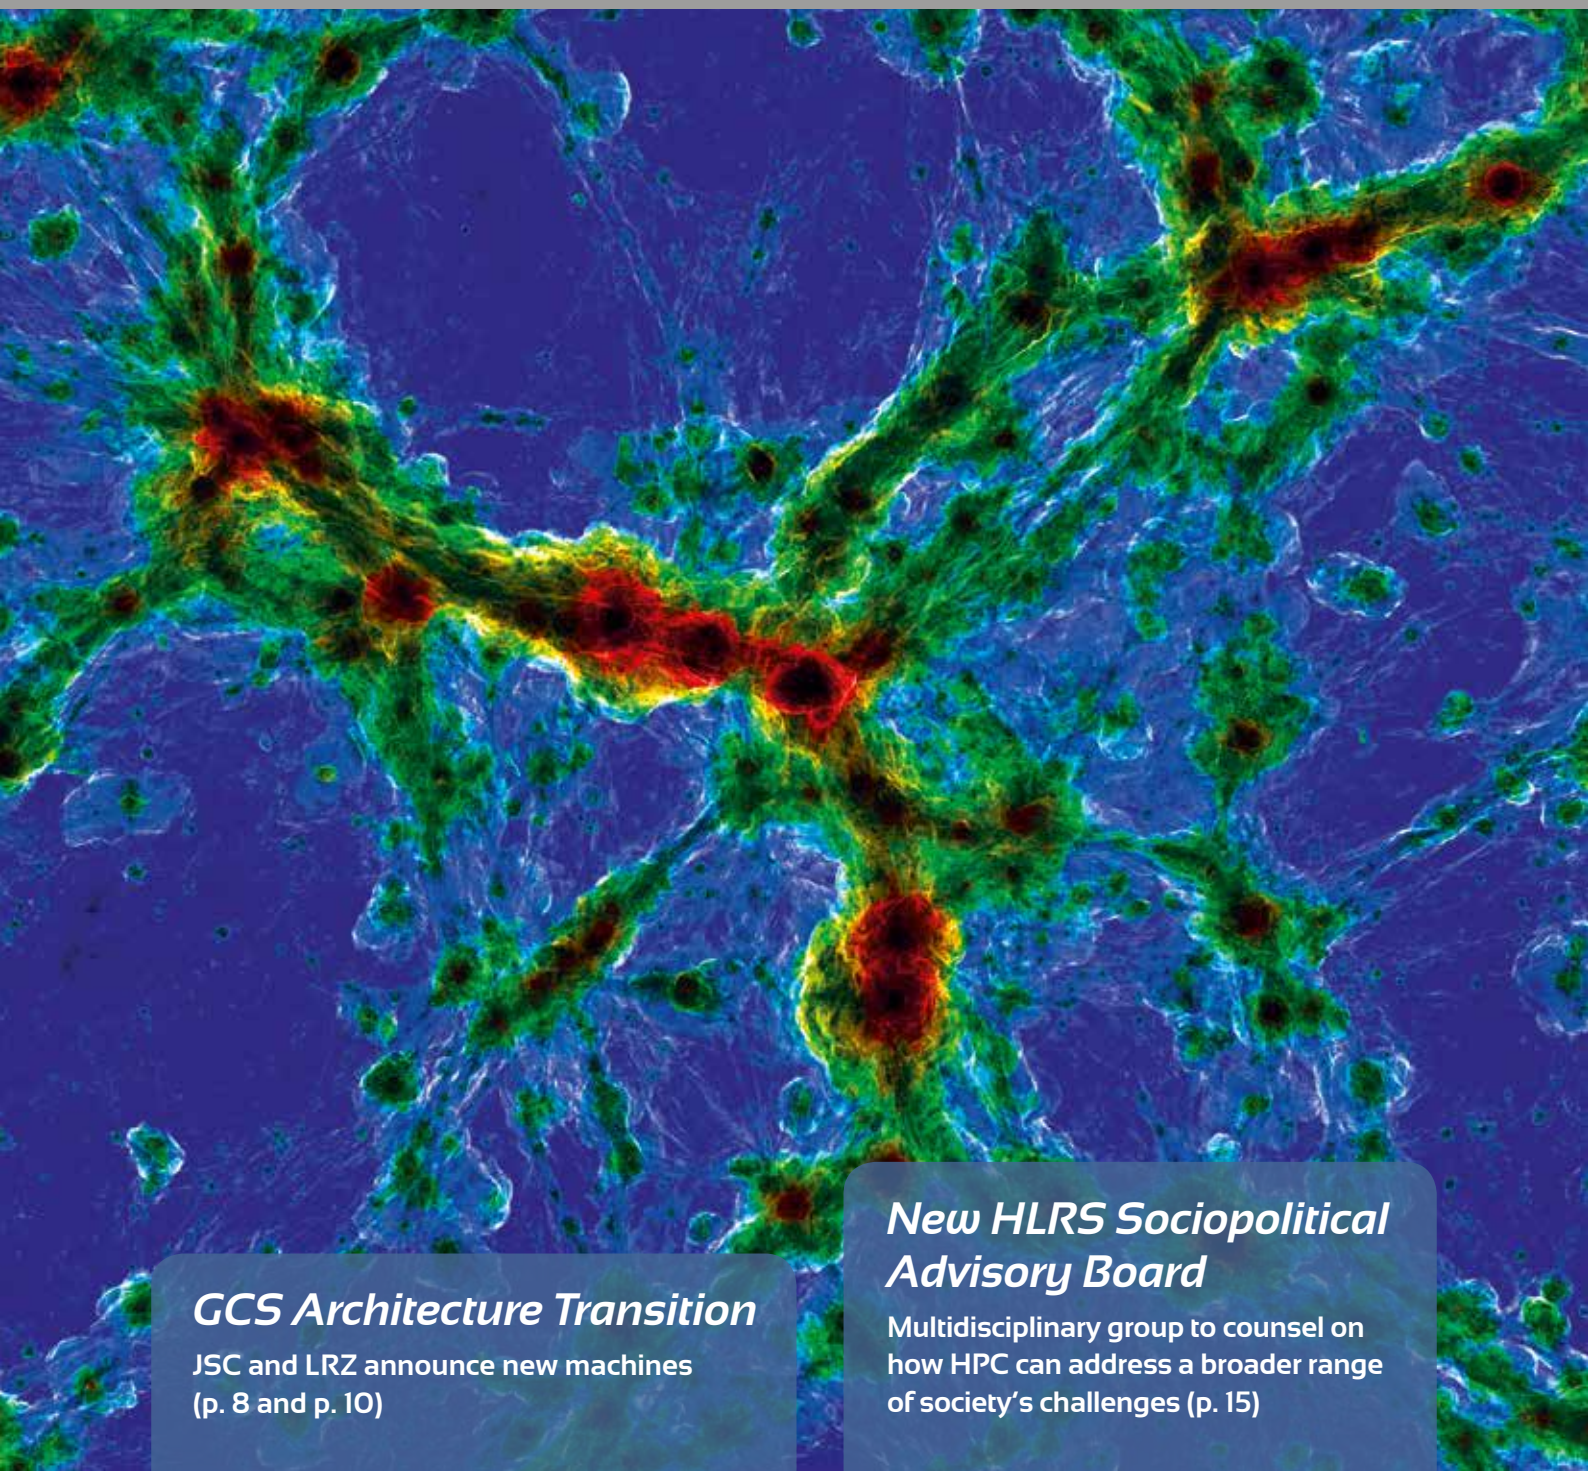

GCS Architecture Transition

JSC and LRZ announce new machines (p. 8 and p. 10)

New HLRS Sociopolitical Advisory Board

Multidisciplinary group to counsel on how HPC can address a broader range of society's challenges (p. 15)

Imprint

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 Center for Supercomputing (HLRS, JSC, LRZ)

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

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

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

Design
Zimmermann Visuelle Kommunikation
www.zimmermann-online.info

Printing
DCC Document Competence Center
Siegmar Kästl e.K., Ostfildern

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

Editorial

Welcome to this new issue of InSiDE, the publication about innovative supercomputing in Germany published by the Gauss Centre for Supercomputing (GCS). Over the last year, GCS has been taking the next steps in hardware upgrades. Jülich Supercomputing Centre will put new systems in operation by mid-2018. The next generation of LRZ's SuperMUC system will replace the current system in autumn this year. With these two installations, Germany is strengthening its position as Europe's HPC leader. 2019 will see a further step toward realizing Germany's HPC strategy when HLRS will start to install its new system to replace Hazel Hen.

These new systems serve as the first phase of the German smart scale strategy, announced last year at ISC 2017. The project SiVeGCS, which is supported by the federal government as well as the state governments of Baden-Württemberg, Bavaria, and North Rhein-Westphalia, provides the necessary funds for German leading-edge high-performance computing for the next 8 years. Funding will not only cover the initial investments and operational costs of these machines, but will also cover funding for additional scientists and researchers to get the absolute best from these systems.

The German approach for HPC is entirely application driven. It is the users that determine which way we are headed. To that end, this issue presents a variety of applications that were run on GCS systems. Researchers from the Technical University of Munich won the best paper award at last year's SC17 conference for presenting the largest and longest multiphysics earthquake simulation ever performed. The team used SuperMUC to recreate the devastating 2004 Sumatra-Andaman earthquake. A research consortium led by cosmologists from the Heidelberg Institute of Theoretical Studies used HLRS to break another simulation record—their own. The team's IllustrisTNG simulation is the largest simulation of galaxy formation, surpassing the scope and detail they achieved in their record-breaking 2015 Illustris simulation. In addition to massive simulation among the stars, University of Cologne researchers used JSC's JUQUEEN to study extreme inner-Earth processes. The team combined theory and experiment to gain greater understanding of the material and geological processes happening deep below the surface of the Earth.

GCS also has a clear focus on its own research as well as training activities. This is reflected in a number of reports on European and German projects. GCS is grateful for the funding received by the European Commission, the Federal Ministry of Science (BMBF) and the German Research Foundation (DFG). It is only by providing such a robust framework that HPC can have a true impact on national research and industry. This issue also presents some of GCS' training activities that are unique in Europe. About 2,000 scientists and industry experts go through our training programs every year, making this the largest HPC training program in Europe.

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

News Events

Applications

Projects

Centers Systems Trainings

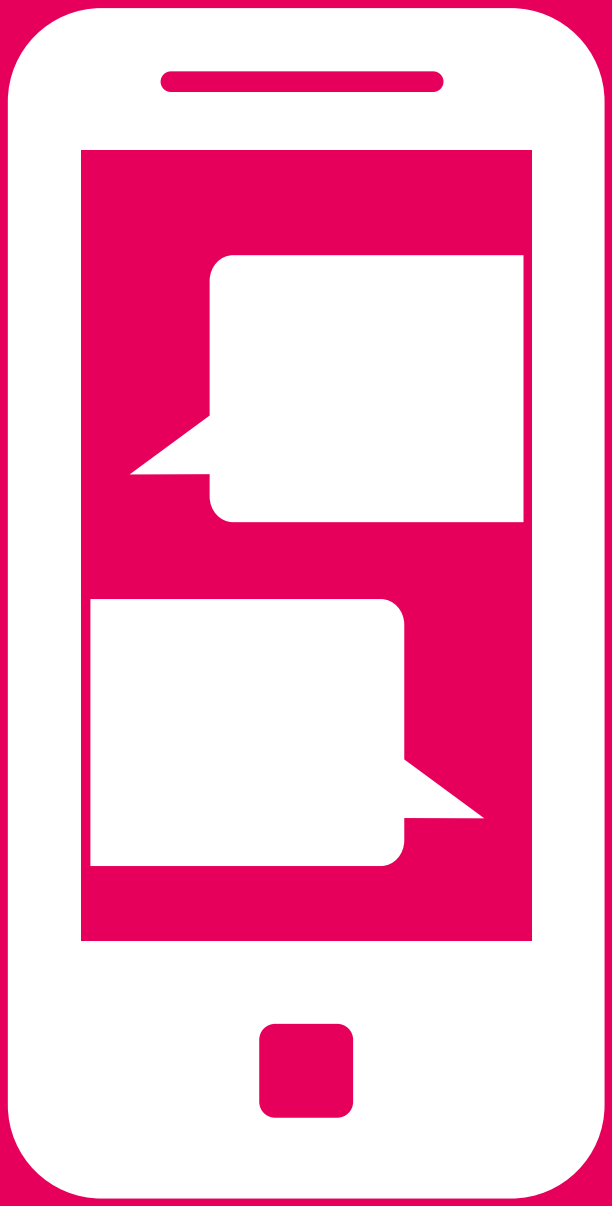

News/Events

In this section you will find general activities
of GCS in the fields of HPC and education.

End of an Era at Jülich Supercomputing Centre as JUQEEN End-of-Life approaches

In spring 2018, the Blue Gene/Q system JUQUEEN at Jülich Supercomputing Centre (JSC) is being decommissioned and dismantled. This event marks the end of a 13-year era at JSC in which the leadership-class system in Jülich was based on IBM's Blue Gene technology. The Blue Gene architectures were defined by a unique combination of energy-efficient embedded processor technology, a well-balanced node and network design, advanced packaging for industry-leading density, and an uncompromised focus on the highest scalability throughout the hardware and software layers.

JSC operated and, in cooperation with vendor IBM, shaped all three generations of the Blue Gene product line for the benefit of its user communities.

In 2005, the first Blue Gene/L system in Germany was installed in Jülich. The system, named JUBL (Jülich Blue Gene/L) consisted of one rack with 1,024 compute nodes and provided a peak performance of 5.7 TFLOP/s. JUBL immediately attracted a large number of users and, despite its specialized architecture, showed promise as an enabler of breakthrough science in a large number of different science disciplines. Consequently, already shortly after its commissioning date, JUBL was extended to an 8-rack system with 16,348 CPUs and a peak performance of 45.6 TFLOP/s. In June 2006, JUBL was the fastest system in Germany and eighth-fastest system worldwide according to the Top500.

The successor system, JUGENE [1], a second-generation Blue Gene/P system, was installed in 2007 with 16 racks, 65,536 compute cores and a peak performance of 222.8 TFLOP/s. In November 2007, JUGENE entered the Top500 list on the second spot and was the fastest open-science supercomputer worldwide. The system was formally inaugurated in February 2008. To accommodate the increasing demand from national and European users, JUGENE was later upgraded in 2009 to 72 racks and 294,912 compute cores, making it the third-fastest system worldwide. With a peak performance of more than one PFLOP/s, JUGENE was the first European petaflop system.

In 2012 and 2013, the current Blue Gene/Q system JUQUEEN [2] was deployed in two phases. JUQUEEN now consists of 28 racks with 458,752 compute cores, 1.8 million hardware threads and a peak performance of 5.9 PFLOP/s. In November 2012, JUQEEN was the 5th fastest supercomputer in the world and the fastest European supercomputer. Together with its user base, JSC tackled the challenge of exploiting the massive parallelism of this architecture. The continuous growth of the High-Q Club [3], a list of applications capable of efficiently using the whole JUQUEEN system, showed the success of these efforts. Currently 31 applications from numerous domains have reached High-Q Club status.

The road of Blue Gene systems' architectural evolution in Jülich is paved by numerous

high-grade scientific discoveries and publications, many of which were widely recognized for their breakthrough character. The list of major research fields using the systems included elementary particle physics, engineering, earth systems modeling, chemistry, condensed matter, plasma physics, and astrophysics. Only two of the breakthrough science examples from the fields of nuclear and particle physics [4] as well as chemistry [5] are referenced here.

Despite its age, JUQUEEN's popularity did not decrease. The system once more showed its superior stability and scalability in the last "Big Blue Gene Week" in early 2018. Between 29.01.2018 and 05.02.2018, the system was solely reserved for large-scale simulations. The third realization of this event was well-received by the users. A total of 14 projects took advantage of the possibility. A staggering 72% of the executed jobs used the full machine.

Since 2006, following JSC's dual architecture approach, the leadership class system was augmented by a general-purpose architecture targeting a broader user base with less scalable applications. With the JUQUEEN successor system JUWELS, JSC is now looking forward to bringing these complementary systems closer together and leveraging the centre's own computing architecture developments for its modular supercomputing infrastructure. This architecture will enable an even broader class of applications to use massively parallel architectures efficiently.

JUQUEEN at Forschungszentrum Jülich.

References

[1] N. Attig, et al., Blue Gene/P:
JUQEEN, Contemporary High Performance Computing: From Petascale toward Exascale, Computational Science Series, 153-188 (2013).

[2] M. Stephan, J. Docter, JUQUEEN:
IBM Blue Gene/Q Supercomputer System at Jülich Supercomputing Centre, Journal of large-scale research facilities, 1, A1 (2015).

[3] http://www.fz-juelich.de/ias/jsc/EN/Expertise/High-Q-Club/_node.html

[4] S. Dürr, et al.:
Ab Initio Determination of Light Hadron Masses, Science Vol. 322, no. 5905 (2008).

[5] D. Marx, et al.:
Unexpected mechanochemical complexity in the mechanistic scenarios of disulfide bond reduction in alkaline solution, Nature Chemistry 9 (2016).

Written by Dorian Krause and Daniel Rohe
Jülich Supercomputing Centre (JSC)
Contact: d.krause@fz-juelich.de

Jülich Supercomputing Centre starts deployment of JUWELS

Fig. 1: Bull Sequana X1000 blade with three compute nodes as used in JUWELS. Copyright: Atos SE

As one era ends, a new one begins: Jülich Supercomputing Centre (JSC) started the deployment of its new modular Tier-0/1 infrastructure. The new computing system, named JUWELS (Jülich Wizard for European Leadership Science) as a reference to its capabilities as a math wizard, will replace the Blue Gene/Q system JUQUEEN in summer, 2018. As a modular supercomputing system, JUWELS will consist of multiple architecturally specialized modules integrated in a unique homogeneous software architecture. The first JUWELS module—a Cluster architecture—provides the basis for future extensions. With a total peak performance of 12 PFLOP/s (10.6 PFLOP/s without accounting for GPU accelerators), the Cluster system already provides significant computing capability. Its most important trait, however, is the versatility of the architecture. With the later addition of a Booster module, architecturally targeted for highly scalable, floating-point-intense simulation and data-science workloads, the performance will be significantly extended. In contrast to the Cluster, however, applications will need to fulfill stronger requirements to take advantage of the Booster's strengths. The Booster installation is targeted for early 2020.

Already the JUWELS Cluster alone is a highly innovative and energy-efficient architecture. The system, supplied by Atos and based on its Bull Sequana X1000 architecture, is also equipped with ParTec's ParaStation Cluster Suite. The system was co-designed by JSC and

the vendors around the best available software and hardware components. The JUWELS Cluster in total consists of 2,559 compute nodes. 2,511 compute nodes are equipped with two latest generation Intel Xeon Skylake Platinum 8168 CPUs with 24 cores each (48 cores per node) and a base frequency of 2.7 GHz. The Skylake microarchitecture supports the AVX-512 instruction set architecture extension providing vector arithmetic operations with 512-bit vector length. 2,271 nodes are equipped with 96 GiB DDR4 main memory, while 240 feature the double main memory size. Further 48 compute nodes are each equipped with two Intel Xeon Skylake Gold 6148 20-core CPUs and four NVidia Volta V100 GPUs with 7.5 TFLOP/s peak performance per GPU and inter-GPU NVLink connection. The compute nodes themselves are interconnected with a 100 Gb/s Mellanox EDR InfiniBand interconnect organized in a fat tree topology. The system will connect to the central Jülich Storage Cluster, JUST, which is undergoing extension during the same installation timeframe.

The JUWELS Cluster consists of ten Sequana cells, each containing two cabinets and one switch rack. The Sequana architecture enables warm water cooling of up to 95% of the generated heat, significantly decreasing the operational cost since free outdoor cooling techniques based on dry coolers are applicable. The necessary outdoor cooling infrastructure will be deployed at Forschungszentrum Jülich in 2018.

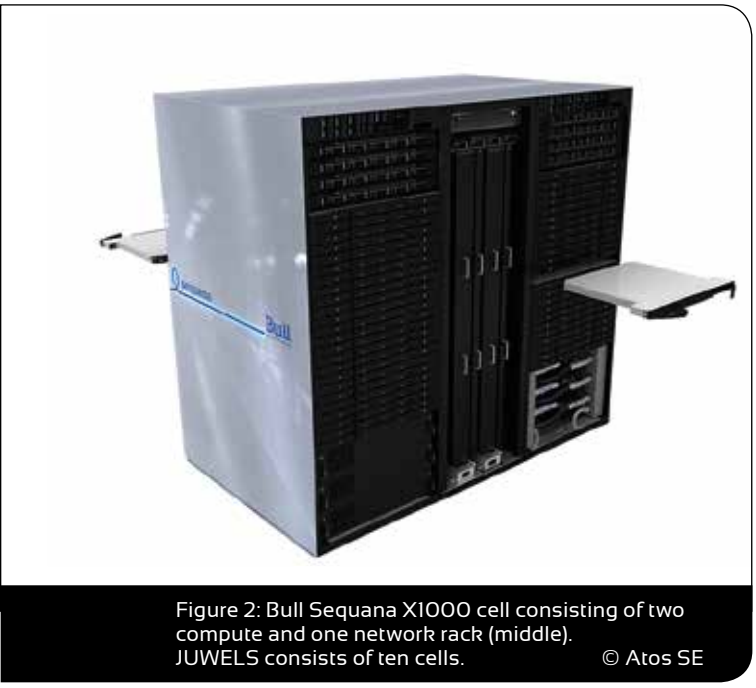

Figure 2: Bull Sequana X1000 cell consisting of two compute and one network rack (middle). JUWELS consists of ten cells. © Atos SE

References

[1] D. Krause, Jülich:
Supercomputing Centre starts deployment of a Booster for JURECA, Innovatives Supercomputing in Deutschland 15 (2017).

Written by Dorian Krause
Jülich Supercomputing Centre (JSC)
Contact: d.krause@fz-juelich.de

Michael Stephan
Jülich Supercomputing Centre (JSC)
Contact: m.stephan@fz-juelich.de

SuperMUC-NG: In the Service of Science

On December 14, 2017, the Leibniz Supercomputing Centre of the Bavarian Academy of Sciences and Humanities announced its next-generation world-class supercomputer—SuperMUC-NG.

The machine will provide 26.7 PLFOP/s of peak performance, and will cater to a wide-range of scientific communities, including astrophysics, life sciences, environmental computing, among others.

SuperMUC-NG will not only significantly improve LRZ's raw compute power, but also enable the handling of tremendous amounts of data (big data) accumulated in today's experiments and simulations. In addition, LRZ aims to give users the flexibility to deploy their own software and visualisation environments for analysing data generated by simulations, and then sharing the results.

For better integration with modern concepts of handling and visualising huge amounts of data, SuperMUC-NG will be linked to separately operated cloud components delivered with the system.

SuperMUC-NG is currently being installed and will start production in early 2019. It will

Fig. 1: Prof. Dr. Dieter Kranzlmüller, Chairman of the Board of Directors at LRZ, Prof. Dr. Thomas O. Höllmann, President of the Bavarian Academy of Sciences and Humanities, Charles Wuischpard, Vice President, Data Center Group General Manager, Scalable Data Center Solutions Group, Intel Corporation, Dr. Ludwig Spaenle, Bavaria's Minister of Science, Dr. Herbert Huber, Head of the High Performance Computing Group at LRZ, , Scott Tease, Executive Director, HPC and AI, Lenovo Data Center Group (from left to right).

SuperMUC-NG, Hardware overview		
Peak performance / Linpack performance	26.9 PFlop/s / 20.4 PFlop/s	
Main memory	718 TByte	
High performance parallel file system	50 PByte	
Data science storage	20 PByte	
Cooling	Direct warm water cooling	
Re-use of excess heat	Adsorption chiller	
	Thin Nodes	**Fat Nodes**
Processor type	Intel Skylake	Intel Skylake
Total number of nodes of this type	6,336	144
Number of cores per node	48	48
Total CPU Cores	304,128	6,912
Number of islands with this node type	8	1
Memory per Node	96 GByte	768 GByte
Interconnect	Intel Omni-Path 100G	Intel Omni-Path 100G
Topology	Pruned Fat Tree	Pruned Fat tree
Software		
Operating system / batch queueing system	Suse Linux / SLURM	
Parallel filesystem	IBM GPFS	
Cloud Components		
Nodes with Nvidia V100 GPUs	32	
Nodes without GPUs	32	

be equipped with more than 6,400 Lenovo ThinkSystem SD650 DWC compute nodes based on the Intel Xeon Scalable processor. This will make it the largest homogeneous supercomputer delivering this kind of performance while also ensuring it is easy to use for a wide spectrum of users. The system will feature 700 TByte of main memory and more than 70 PByte disk storage. For more technical details, please see the table.

Next-generation energy efficiency

Just like SuperMUC, SuperMUC-NG will be cooled using warm water. Lenovo, the system integrator, has developed a cooling concept that will further reduce power consumption while also reusing the supercomputer's waste heat to generate cold water. The Lenovo SD650 DWC servers were designed to be able to withstand up to 50°C inlet temperatures, enabling 85–90 percent heat recovery. Water is piped

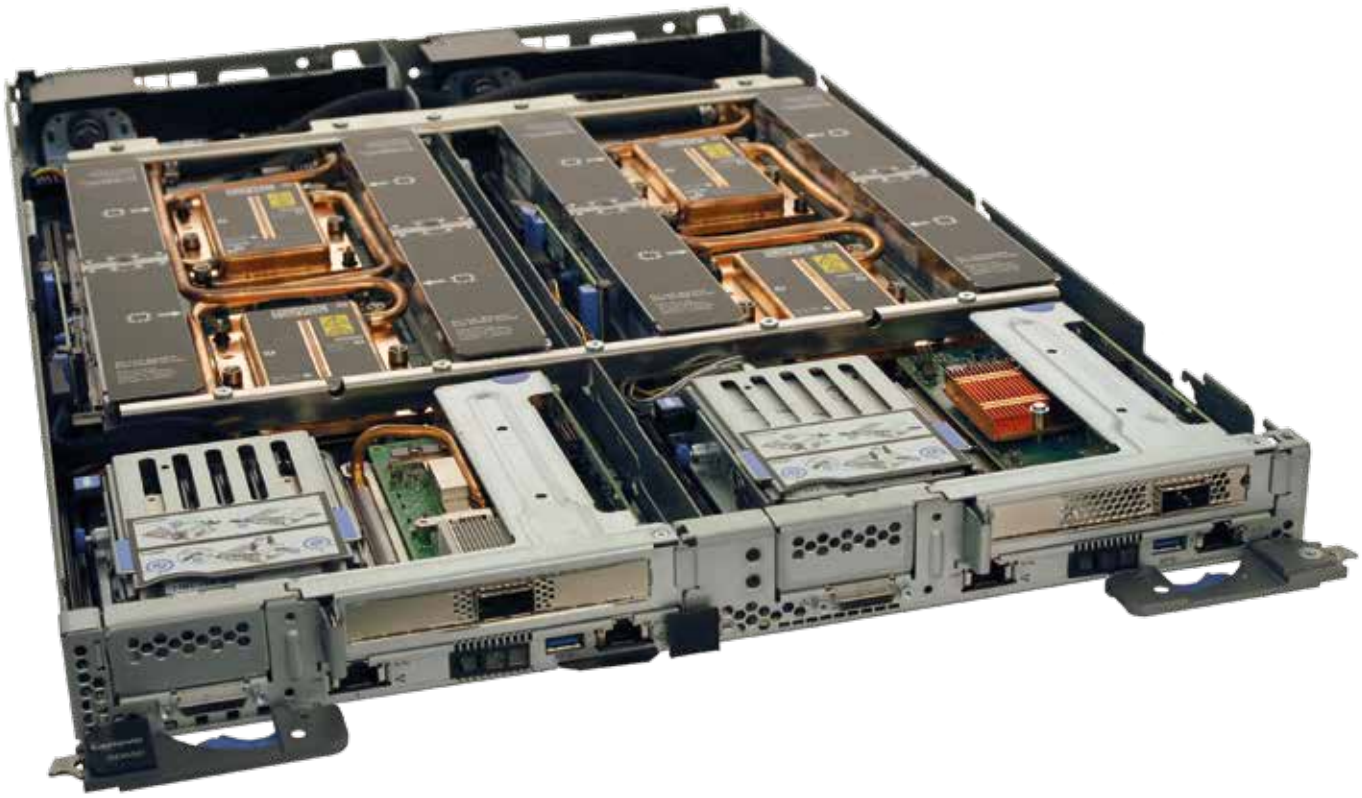

Fig. 2: One blade of Lenovo's SD650 DWC contains two compute nodes.

out of the servers at 58-60°C, depending on workload, and sent through an adsorption chiller, where it is converted to chilled 20°C water suitable for cooling storage and networking components. Adsorption chilling will generate about 600 kilowatts of chilled water capacity. This translates into energy savings of more than 100,000 Euros per year.

Joint funding from Berlin and Bavaria

The funding of SuperMUC-NG is shared in equal parts by the federal government of Germany and by the Free State of Bavaria through the Gauss Centre for Supercomputing's (GCS's) strategic plan. The total cost of the project totals 96 Million Euro over six years, which includes electricity, maintenance and personnel.

Written by Sabrina Eisenreich
Leibniz Supercomputing Centre
Contact: sabrina.eisenreich@lrz.de

JURECA Booster starts Operation

In November 2017, the Jülich Supercomputing Centre (JSC) opened user access to the JURECA Booster system. As announced previously [1], the system was installed by Intel, Dell, and partners in late summer and autumn 2017. The Booster consists of 1,640 computes nodes with one 68-core Intel Xeon Phi "Knights Landing" 7250-F, 96 GiB main memory and additional 16 GiB high-bandwidth, multi-channel DRAM (MCDRAM) memory. The compute nodes are interconnected by a 100 Gb/s Intel Omni-Path Architecture interconnect with full fat tree topology. The Booster is tightly integrated with the architecturally complementary JURECA Cluster, an 1,884-node cluster based on multi-core CPUs and Mellanox InfiniBand interconnect, through 198 bridge nodes. Cluster and Booster are jointly administrated using ParTec's ParaStation Cluster Suite and offer users a homogeneous software layer across the two underlying heterogeneous hardware modules.

In the first months following the start of operation, JSC and Intel have been working to mature the system and address I/O constraints affecting particularly data-intensive applications. For a number of users, however, the system has already been an important and productive workhorse.

The system first showed its capabilities when the Cluster and Booster jointly entered the

Fig. 1: The JURECA supercomputer at JSC consisting of the Cluster module (left), delivered in 2015 by T-Platforms, and the Booster module (right), delivered by Intel with Dell in 2017. © Forschungszentrum Jülich GmbH

Top500 list in November, 2017 at spot 29 with a Linpack performance of 3.78 PFLOP/s, boosting the system from spot 80 on the June 2017 list. This success proved the capabilities of the Cluster-Booster communication mechanism in the ParaStation software layer as the enabler of high-performance communication across the two architecturally diverse high-speed interconnects in the Cluster and Booster. In addition, it also shows that JSC's modular supercomputing approach—primarily designed to accelerate complicated simulation codes with their multiple different scalability characteristics—can also benefit applications such as those resembling performance and scalability characteristics of the Linpack benchmark that are not the primary design target, albeit require additional tuning efforts.

The project partners are jointly working to further mature the Cluster-Booster software stack starting from the current prototype implementation and to enable the larger JURECA user community to concurrently use Cluster and Booster in their simulation projects.

References

[1] **D. Krause, Jülich:**
Supercomputing Centre starts deployment of a Booster for JURECA, Innovatives Supercomputing in Deutschland 15 (2017).

Written by Dorian Krause
Jülich Supercomputing Centre (JSC)
Contact: d.krause@fz-juelich.de

Sociopolitical Advisory Board to Help Guide HLRS Vision

Fig. 1.: Bringing together professionals in media, medicine, education, art, design, and architecture, HLRS hosted its inaugural meeting for its sociopolitical advisory board. Copyright: Julian Rettig.

Everyday at supercomputing centers like the High-Performance Computing Center Stuttgart (HLRS), simulations provide unique contributions for basic and applied science. From predicting weather and discovering new medicines, to enabling the more efficient design of cars and traffic flow, simulation indirectly benefits people in their everyday lives in many ways.

At the same time, however, the use of simulation is often driven by scientific or economic interests. Moreover, knowledge about the value of simulations or even their existence is often limited to a group of experts in a particular field. This reality can lead to limitations in vision for other ways in which simulation could be applied for the benefit of society.

"There is no question that computer simulations can provide added value, but they must not be a privilege of elites," says HLRS director Michael Resch. Together with Andreas Kaminski, head of the Department of Philosophy of Computer Simulation at HLRS, he recently convened a new sociopolitical advisory board; the group held its inaugural meeting at HLRS on April 9. "We want to work together to ensure that in the future a wider cross-section of society benefits from the use of simulation," Resch explains, describing his vision for the board.

Improving engagement with society

The sociopolitical advisory board includes 13 members who come from different areas of society, ensuring representation of a wide diversity of socially relevant topics. Among the disciplines represented are nursing, architecture, design, art, education, and journalism.

Many board members recognize the potential of simulation to make complex events understandable. "Simulation can make social phenomena more accessible to decision-makers and affected people," says sociopolitical advisory board Chairman Prof. Ortwin Renn, Scientific Director at the Institute for Advanced Sustainability Studies in Potsdam. He points out, "Against the background of this progress, however, it must be ensured that ethical values are not violated and that social preferences are sufficiently taken into account."

The meeting started with a presentation of research projects currently underway at HLRS that have everyday social relevance. This led to a discussion among the advisory board participants on how simulation could be leveraged to support progress in their fields of expertise. Discussions centered around specific topics such as urban development, the emergence of social inequality, and populism. But also questions regarding the methodology of simulations arose: could it be simulated how Facebook would look if there were no filtering algorithm? Could a simulation provide information on desired realities, showing how it would look like to live in them?

Such thematic and methodological ideas will in the future be structured, united, and refined. Discussions will continue in order to develop a pilot proposal for a possible research project, given the availability of a suitable funding program. "We have received more input than we hoped for", Kaminski states. "As a result, the next meeting will likely take place before the end of year in order to continue working on the pilot."

Adding to existing HLRS efforts in social responsibility

With the advisory board HLRS will seek to discover new areas in which computer simulation can make a social contribution. This new effort will enhance other activities on social responsibility that HLRS has already undertaken. So far, for instance, a sustainability group at the HLRS has been tackling issues of energy and water use in supercomputing, while a project called "Simulated Worlds" gave pupils a chance to expand their digital skills by completing a one-year science project focused on traffic simulation.

Written by Lena Bühler
High Performance Computing Centre Stuttgart (HLRS)
Contact: buehler@hlrs.de

Human Brain Project Successfully Enters the Next Phase

Human Brain Project

The Human Brain Project (HBP) [1] is one of the first two projects funded by the European Union's Future and Emerging Technologies (FET) Flagship Initiative [2]. With over a hundred partner institutions from more than 20 countries, HBP is building a European research infrastructure for the neuroscience community. It was launched in October 2013 and successfully entered its third project phase in April, 2018.

Gaining a comprehensive understanding of the brain is one of the grand scientific challenges of our time. Several countries have initiated longer-term brain research initiatives in recent years, and HBP was among the first. While other projects focus on certain neuroscientific research aspects, the HBP takes the unique and sustainable approach to build and operate a research infrastructure (RI) co-designed by the neuroscience community. This RI will enable scientists to gather, organize and disseminate data, to build and simulate multi-scale models of the brain, and to develop brain-inspired computing, data analytics, robotics software and technology. In December, 2017, the HBP and four other international projects signed the Canberra Declaration to create an International Brain Initiative (IBI) and join forces in neuroscience research [3].

The HBP RI is organized in six information and communication technology (ICT) platforms that were first released at the end of the project's ramp-up phase in March, 2016 [4]. In the last two years, these platforms significantly advanced and were more closely integrated. The High-Performance Analytics and Computing (HPAC) Platform—coordinated by the Jülich Supercomputing Centre (JSC) at Forschungszentrum

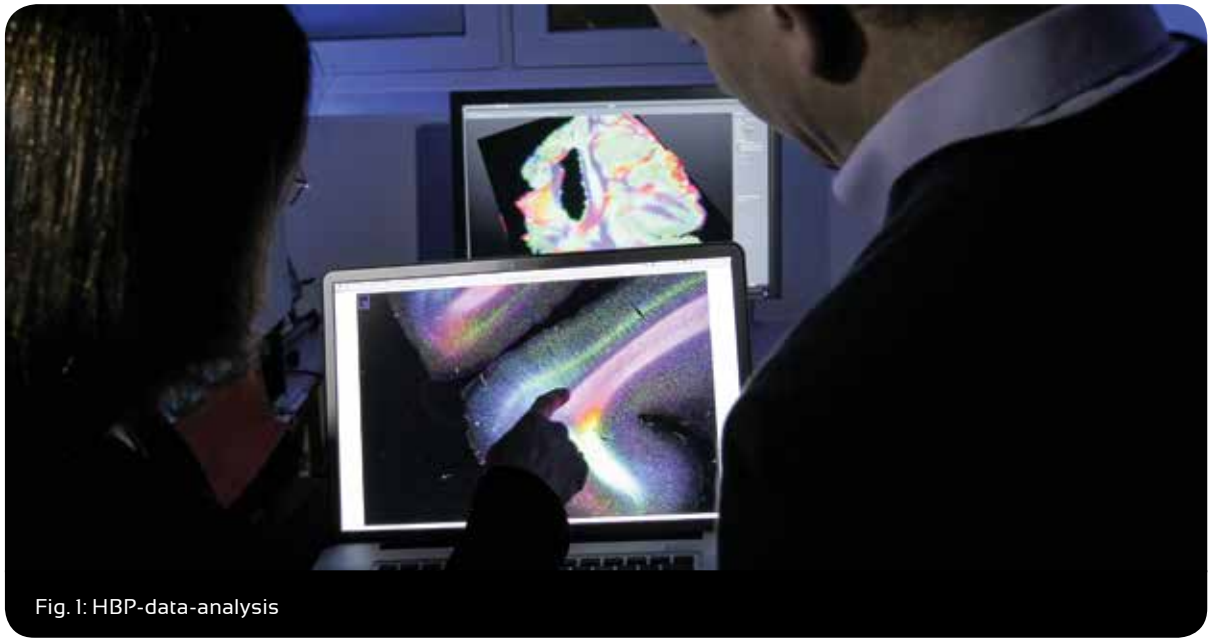

Fig. 1: HBP-data-analysis

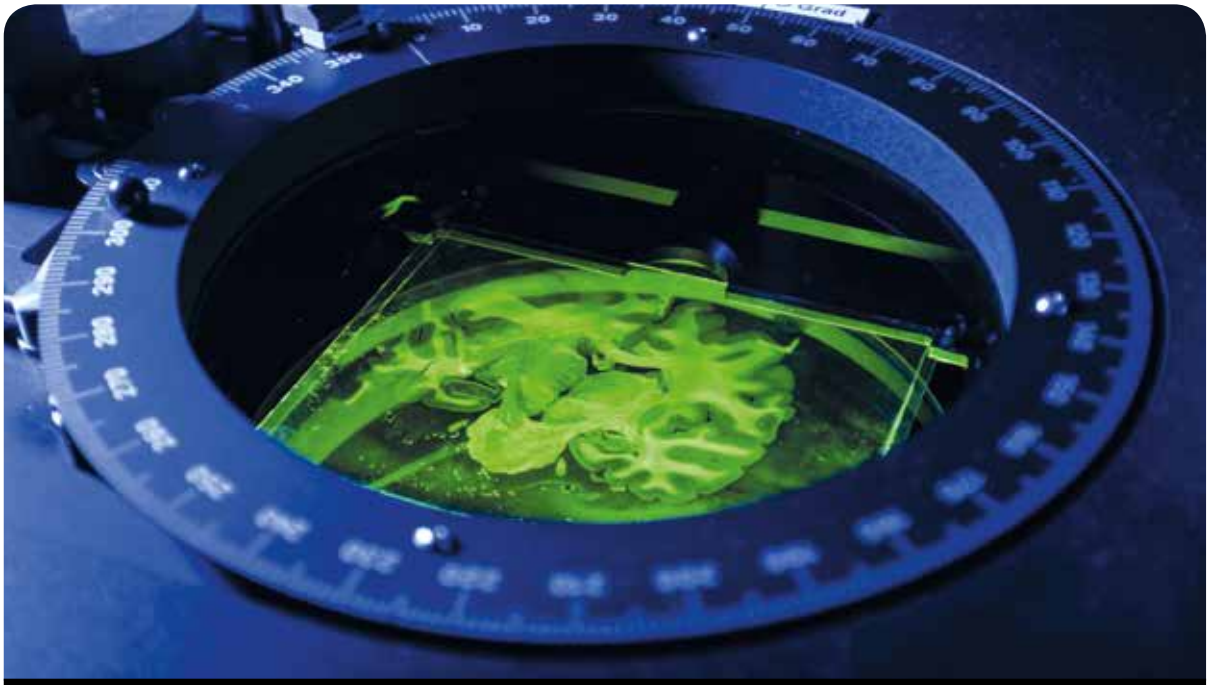

Fig. 2: HBP-data-gathering

Jülich and the Swiss National Supercomputing Centre in Lugano (CSCS)—plays a crucial role in the HBP RI. The mission of the HPAC Platform is to provide the basic data and computing infrastructure that enables scientists to store their data, integrate it into models, use it in simulations, as well as to analyze and visualize it. The participating centers BSC, Cineca, CSCS, JSC and, starting with the third phase of HBP, the TGCC of CEA in France are jointly developing the infrastructure in co-design with early adopters from the HBP community. Including CEA, tier-0 HPC systems of all five PRACE hosting members are now integrated into the HPAC Platform and available to HBP scientists for their research. The five centers are also working closely together to develop the federated Fenix infrastructure, which receives funding through the recently started ICEI project [5]. While primarily driven by HBP use cases, the scope of Fenix goes beyond neuroscience, as the infrastructure should also benefit other research areas with similar requirements, such as materials science.

One of the main goals for the recently started HBP phase will be to bring the six ICT platforms together in a single platform, the HBP Joint Platform [6]. The strength of the HBP approach lies in the support of complex, interactive workflows including modelling, simulation and data analysis. These workflows require the combined use of several of the individual platforms and—importantly—their seamless integration,

to provide a user experience that allows researchers to focus on scientific problem solving rather than technical details. To this end, the HBP Collaboratory [7] will be transformed into a convenient and reliable user portal for the HBP Joint Platform. In addition, the HBP is about to establish the HBP High-Level Support Team (HLST), a group of active, experienced scientists and technology experts that will support users in the implementation of complex workflows, with the welcome effect that parallel development will be avoided and synergies used. The HLST will also help with applications for resource allocations.

The HBP Joint Platform will build on the base infrastructure provided by Fenix. The HPAC Platform will be migrated to operate on top of this federated infrastructure and take care of the community-specific, platform-level interfaces to the generic Fenix services and resources. Already in the previous project phase some of the other HBP ICT Platforms migrated parts of their service portfolios to the HPAC infrastructure. These efforts will be intensified in the next two years of the HBP Flagship endeavor in order to enable neuroscience research that will bring us closer to a comprehensive understanding of the human brain.

References

[1] **Human Brain Project:**
https://www.humanbrainproject.eu

[2] **FET Flagships:**
http://ec.europa.eu/programmes/horizon2020/en/h2020-section/fet-flagships

[3] **Canberra declaration:**
https://www.humanbrainproject.eu/en/follow-hbp/news/worlds-brain-initiatives-move-forward-together/

[4] **Human Brain Project:**
Towards a European infrastructure for brain research, http://inside.hlrs.de/editions/16autumn.html#human-brain-project-towards-a-european-infrastructure-for-brain-research (Reference to an article in a previous InSiDE edition)

[5] **Fenix Consortium to build a Federated Data and Computing Infrastructure for the Human Brain Project**
(Reference to another article in the same InSiDE edition)

[6] **HBP Joint Platform:**
https://www.humanbrainproject.eu/en/hbp-platforms/hbp-joint-platform/

[7] **HBP Collaboratory (registration required):**
https://collab.humanbrainproject.eu/

Written by Anna Lührs and Boris Orth
Jülich Supercomputing Centre
Contact: a.luehrs@fz-juelich.de
b.orth@fz-juelich.de

Leibniz Supercomputing Centre Named as A Bavarian Big Data Competence Centre

Fig. 1: Panel discussion at the kick-off event "Big Data: Driving Science and Economy in Bavaria". From left to right: Prof. Dr. Dieter Kranzlmüller (LRZ), Dr. Stefan Roskos (ONELOGIC), Prof. Dr. Thomas Seidl (LMU), Prof. Dr. Fabian Theiss (TUM), and Dr René Wies (BMW).

Big Data: Driving Science and Economy in Bavaria. This was the motto of a kick-off event for naming the LRZ as A Bavarian Big Data Competence Centre (BBDCC) on February 22, 2018.

High-ranking officials from the Bavarian Ministry of Education and Research, including State Secretary Georg Eisenreich, as well as representatives from industry and Munich-area universities came to the LRZ to attend a multifaceted programme that included ample time for networking and discussion.

In his keynote address, Dr. René Wies from the R&D department at BMW Group IT, shared countless examples of how big data, machine learning and artificial intelligence have changed information technology (IT) as well as auto manufacturer BMW's manufacturing and business processes.

A panel discussion followed the keynote, which featured Dr. Wies, Prof. Dr. Thomas Seidl, Chair for Data Mining at LMU, Prof. Dr. Fabian Theiss, Ordinarius for Bioinformatics at TU Munich and leader of the Helmhotz Institute for Computational Biology, Dr. Stefan Roskos, Managing Director at ONE LOGIC and LRZ's Chairman of the Board of Directors, Prof. Dr. Dieter Kranzlmüller. The panellists engaged with the audience, which led to a dialogue on topics ranging from which technologies need to be mastered in order to successfully turn big data into smart data and what is needed to more effectively train data scientists and AI experts.

The LRZ used the opportunity at the event to introduce itself as a Bavarian Big Data Competence Centre. In his introductory talk, Prof. Kranzlmüller outlined the parameters of the new centre. Building on its strong foundation as an IT competence centre, LRZ will continue to increase its support for users from science and industry (in the pre-competitive area)—as well as potentially small- and medium-sized enterprises (SMEs)—by helping to solve IT- and big-data-related problems . The aim is to enable partners to turn big data into smart data in order to generate new knowledge and added value to the businesses, and LRZ serving as a „one-stop shop" for all matters relating to big data. The BBDCC is completely integrated into the Bavarian state government's DIGITAL STRATEGIE, and LRZ's efforts ideally complement the activities of the Zentrum Digitalisierung.Bayern (ZD.B)—an organization dedicated to strengthening networking between science and industry—as well as the development activities of the Fraunhofer Institutes.

One focus of the BBDCC will be on the various aspects of managing data, especially as it relates to the storage, processing, and provision of open data, which should be based on the FAIR principle. It will help users to make their publishable scientific data discoverable, accessible, interoperable and reusable. Another focus is to process confidential data securely (data security and privacy).

With respect to infrastructure, the Munich Scientific Network (Münchner Wissenschaftsnetz, or MWN), operated by LRZ, offers the ideal infrastructure to transfer data quickly and securely between users and LRZ resources, making it available on the required hardware resources within LRZ's compute environment.

In terms of hardware, LRZ provides resources especially suited for processing and analysing large amounts of data, such as the LRZ Data Science Storage or special hardware such as the NVIDIA DGX-1 for deep learning applications. These processing and analysis offerings will continue to be expanded. In addition, LRZ wants to scale machine learning from desktop solutions to the supercomputing level in the future.

In addition to these focus areas, the BBDCC will also focus on training and education. In this arena, LRZ contributes its many years of experience and already offers courses for beginners and intensive workshops for advanced users on topics ranging from the introduction to machine learning tools to the use of the programming language R for statistics and data analysis.

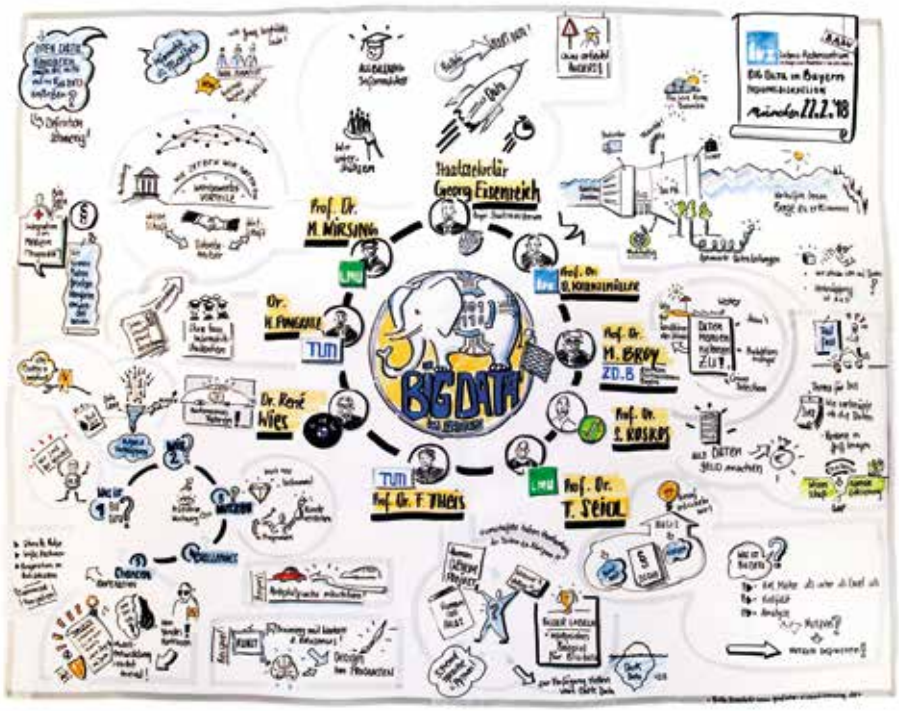

Fig. 2: Graphic recording of the whole kick-off event by Britta Krondorf.

For more information on the LRZ as Bavarian Big Data Competence Centre please visit: www.bigdata.lrz.de

Written by Sabrina Eisenreich
Leibniz Supercomputing Centre
Contact: sabrina.eisenreich@lrz.de

Designing for Energy Efficiency in HPC Centers

High-performance computing (HPC) has become an essential tool for investigating many kinds of problems in research and technology development. But the opportunities that it offers come at a cost. Operating a supercomputer can burn through the energy needed to power a small city, requires large cooling systems to prevent electronic equipment from overheating, and relies on literally tons of electronic hardware whose creation and disposal have sizable environmental impacts.

All of these facts also make HPC systems expensive to run, meaning their operators have a vested interest in designing them to be as efficient as possible. Currently, however, each computing center must individually find its own way to make its supercomputing more sustainable, not just in terms of its economic costs, but also with respect to its environmental and social implications.

In an effort to promote discussion connecting these issues within this broader context of sustainability, HLRS organized and hosted its first Energy Efficiency Workshop for Sustainable High-Performance Computing on October 25–26. The event brought together representatives from supercomputing facilities in the Gauss Centre for Supercomputing (GCS) and other academic institutions. The three GCS centres held presentations that focused on recent innovations in cooling systems and heat reuse, which currently offer some of the greatest opportunities for increasing efficiency in computing centers.

HLRS Deputy Director for Infrastructure Norbert Conrad and sustainability team member Ursula Paul described HLRS's sustainability strategy. They discussed plans for further improving the ecological, economic, and social aspects of sustainability in the years to come and for gaining Eco-Management and Audit Scheme (EMAS) and ISO 50001 certifications, two demanding standards for environmental sustainability and energy management, respectively.

Willi Homberg from the Jülich Supercomputing Centre (JSC) gave an overview of how JSC's cooling concept has evolved over time, from Freon-cooled supercomputers in the 1980s to its current-day hybrid cooling approach. In addition, Homberg gave his perspective on modern cooling concepts such as direct water cooling and the prospect of more efficiency coming from future technologies, including an approach currently under investigation at JSC—immersion cooling, a process of completely submerging computer processors in a high-tech fluid to immediately catch heat radiation, in turn minimizing the amount of addition cooling that would be needed.

He also indicated that, as part of Forschungzentrum Jülich's campus development plan, JSC has plans to capture waste heat from its supercomputer to heat the surrounding buildings.

Leibniz Supercomputing Centre (LRZ) staff member Torsten Wilde discussed the adsorption cooling system at LRZ—the only one of its kind in a Top500 supercomputer. He briefed

HLRS brought together members of the German supercomputing community to focus on ways to improve energy efficiency and sustainability in HPC centers.

workshop participants what tests his team has undertaken to further optimize energy efficiency in the system. He also described a current collaborative project between LRZ and its academic and industrial partners to develop software that uses machine learning to analyse data collected from sensors in ways that would make it possible to predict and manage energy consumption in real time.

In addition to the GCS presenters, researchers representing the Gauss Allianz, the Center for Information Services and High-Performance Computing at TU Dresden, the Steinbuch Centre for Computing at the Karlsruhe Institute of Technology, and the University of Stuttgart's Institute for Energy Economics and Rational Energy Use also presented. The group agreed that bringing together these different perspectives on energy efficient supercomputing was valuable and the representatives from the partner organizations will be discussing ways to build on the success of the first workshop in the coming years.

Written by Christopher Williams
High Performance Computing Centre Stuttgart (HLRS)
Contact: williams@hlrs.de

Navigating A Post-Moore's-Law World

In the 1960s, researcher and businessman Gordon Moore stated that the number of transistors on a computer chip would double every 1–2 years for the coming decades, leading to a near exponential growth of computing power.

Until recently, this concept, dubbed Moore's Law, has served as the rough roadmap for building faster, more powerful supercomputers. Recently, however, progress has begun to lag behind the pace specified by Moore's Law because of physical limits. In addition, intra-node communication as well as data transport and storage have become the next frontiers in helping high-performance computing (HPC) users achieve new heights in scientific research.

At the 17th annual Workshop on Scalable Global File Systems, which took place on April 16–18 at the High-Performance Computing Center Stuttgart (HLRS), experts came together to address next-generation computing challenges related to data management. HLRS and partner Höchleistungsrechner für Wissenschaft und Wirtschaft (hww, or in English, High-Performance Computers for Science and Business) hosted the event.

Taking "Post-Moore I/O?" as this year's theme, the workshop brought together experts from academia, computing centres, and hardware and software suppliers to discuss solutions to the challenges of helping users effectively move, store, and manage their datasets.

As the HPC community approaches exascale computing—a thousand-fold increase over a current-generation petaflop machine, capable of at least one quadrillion calculations per second—computer scientists from industry and academia must develop new methods to send data to and from compute nodes, on and off HPC platforms, and to and from users' home institutions.

"This year's theme was a reflection on the development path for high-performance I/O," said Dr. Thomas Bönisch, group leader for the HLRS Project User Management group and one of the primary oraganizers of the workshop. "Will devices and technologies used in and for high-performance input/output (I/O) soon face similar issues like silicon based parts and if yes, do we have alternatives?"

Storage technology, like many cultural trends, sees an ebb and flow in the prevalence of various storage methods. Older technologies get replaced by newer innovations, but are often repurposed or rethought later, once again increasing in popularity.

For example, before disks were a ubiquitous part of computing, tapes were a common method for storing information. Even as disks and virtual storage become more and more common for commercial storage technology, tapes have remained a valuable storage tool for high volumes of data and remain an important part of an HPC centre's long-term storage portfolio.

Another "throwback" technology, flash memory, is becoming an increasingly integral part of an HPC centre's data storage options and was discussed heavily throughout the conference.

Flash memory was developed in the 1980s as an alternative to hard-disk-drive (HDD) memory. The format allows data to be saved and secured without a continuous power source—USB sticks, for example, use this format—and is the principle behind solid-state disk drives (meaning there are no moving components).

While flash primarily served as external storage for personal computing in the last couple of decades, HPC administrators see a scaled-up version of flash's energy efficiency and speed as a possible solution to the increasing demand for large volumes of data processing. Some believe it could also speed up the cumbersome (I/O) process used in moving data on and off HPC systems.

During the workshop, experts from DDN, Seagate, Cray, IBM, Megware, ThinkParQ, Dell EMC, Intel, and Adva presented various visions for the future mix storage between hard-disk drives, tape storage, flash memory, persistent memory, and, especially, the file system software to manage storage devices.

Toward the end of the workshop program, Bönisch summarized some of the key challenges that came up during the conference, and offered up some questions from the perspective of an HPC centre's future storage-related needs.

Near the end of the 17th annual Workshop on Scalable Global File Systems, attendees participated in an open discussion about the future of storage technologies.

Bönisch asked the attendees to consider whether power would be one of the primary limiting factor for I/O performance, whether scaling out a centre's storage capabilities was the only valid development path, whether software and I/O libraries would be able to keep pace with hardware improvements in the field, and the role persistent memory will play in future systems, among other questions. This led to a discussion between representatives of the different stakeholders from hardware vendors, software developers, computer center representatives and high-end system administrators.

Written by Eric Gedenk

Gauss Centre for Supercomputing (GCS), based at High Performance Computing Centre Stuttgart (HLRS)

Contact: gedenk@hlrs.de
eric.gedenk@gauss-centre.eu

Fenix Consortium to build a Federated Data and Computing Infrastructure for the Human Brain Project and other communities

The European supercomputing centers BSC (Spain), TGCC/CEA (France), Cineca (Italy), CSCS (Switzerland) and JSC (Germany) have joined forces to design and build a federated data and computing infrastructure. The realization of the so-called Fenix research infrastructure, which in the beginning will be used primarily by the Human Brain Project (HBP) [1], started officially in January 2018 with the launch of the "Interactive Computing E-Infrastructure" (ICEI) project. The ICEI project is co-funded by the European Commission through a Specific Grant Agreement (SGA) under the umbrella of the HBP Framework Partnership Agreement (FPA) [2].

The five partners in the Fenix consortium—all of which are also members of the HBP's High Performance Analytics and Computing(HPAC) Platform subproject and hosting members of PRACE—plan to deliver a generic set of e-infrastructure services which will serve the HBP and other communities as a basis for the development and operation of community-specific platforms. To this end, the design and the implementation of the Fenix infrastructure are driven by the needs of the HBP as well as other scientific communities with similar requirements (e.g., materials science) [3, 4].

The key services provided by Fenix will encompass interactive computing services (supporting direct user interaction with data-intensive distributed workloads), scalable computing services (supporting scalable workloads), and a federated data infrastructure (including federated archival data repositories optimized for capacity, reliability and availability, and site-local active data repositories close to compute or visualization resources). A distinctive feature of the Fenix infrastructure is that data storage and scalable computing resources will be in close proximity to each other and tightly integrated. In this manner, Fenix will enable HBP use cases such as GUI-based interaction with large-scale neural network simulations, massive data processing and analysis for the HBP Human Brain Atlas, the validation of simulation results obtained on neuromorphic hardware, and the implementation of the HBP Neurorobotics Platform. Virtual Machine (VM) services provided as part of the Fenix infrastructure are hosting the HBP Collaboratory, the central gateway to all HBP Platforms, and other HBP platform services.

The ICEI project will realize key elements of the planned infrastructure through procurements of R&D services and equipment, which will be coordinated between the partners. In preparation of these procurements, a public information event was held in Barcelona on March 15 in order to inform interested suppliers about the expectations and plans of the ICEI partners, and to gather their feedback. Other significant components of the infrastructure will be realized through in-kind contributions from the participating supercomputing centers.

Beside the technical infrastructure, the Fenix consortium will also establish a suitable governance structure and define mechanisms for the allocation of Fenix resources. Stakeholders,

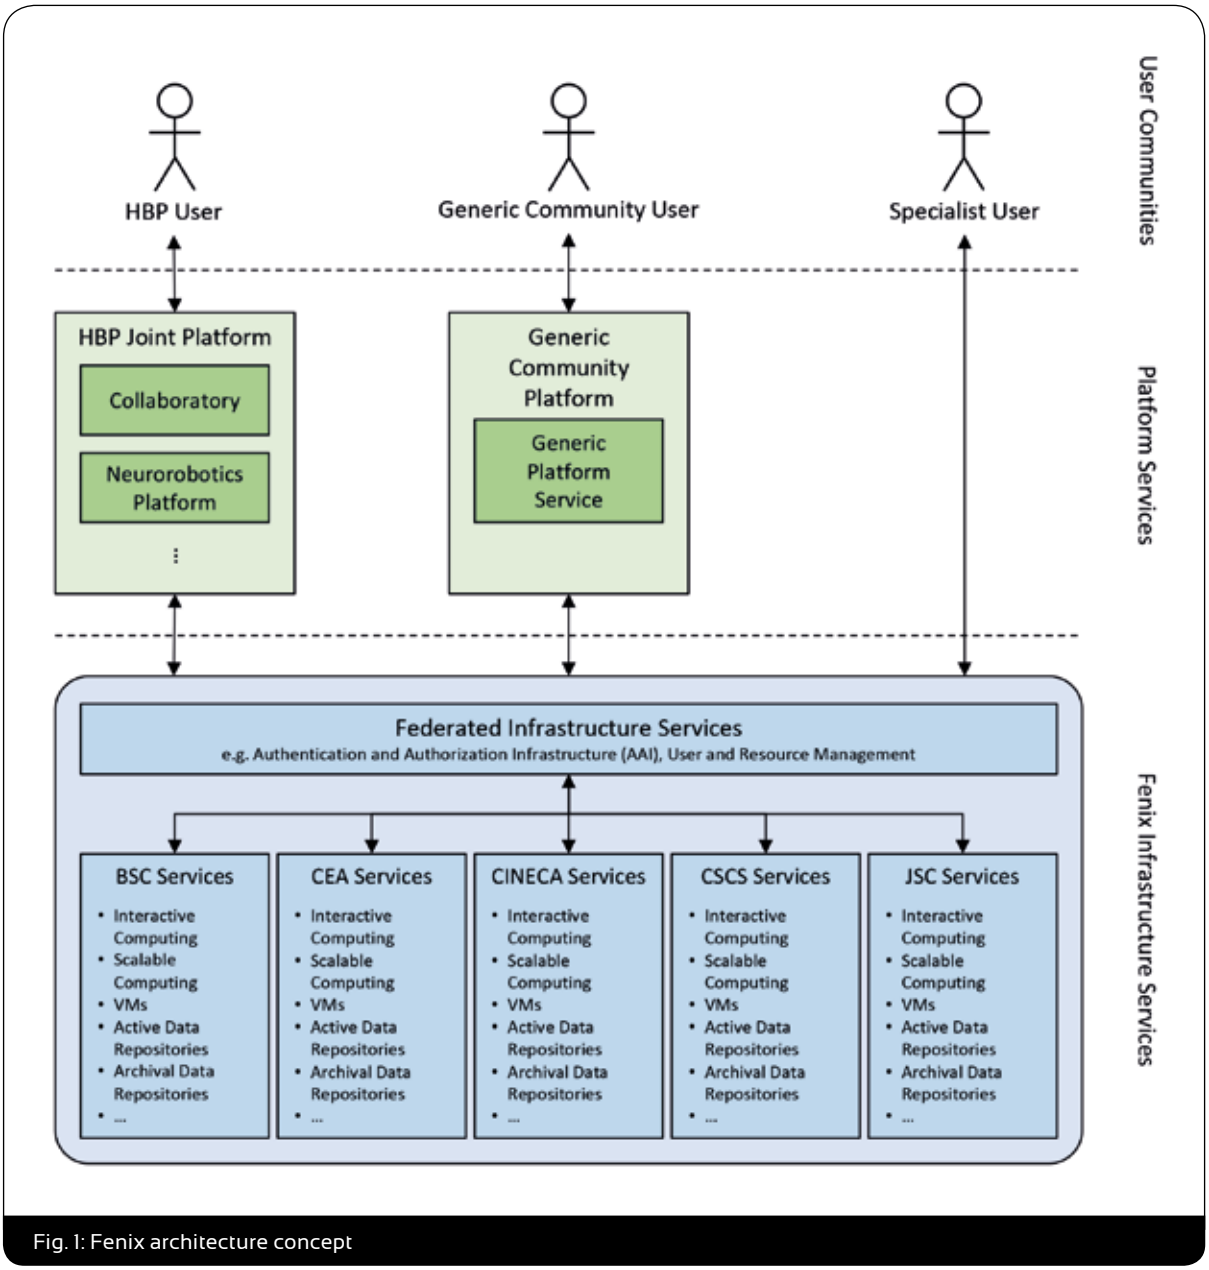

Fig. 1: Fenix architecture concept

like the Fenix partners and the HBP, who contribute substantial financial or other resources to the Fenix infrastructure, will be represented in the supreme body of the Fenix governance structure, the Fenix Council. Among other things, the Fenix Council is responsible for defining and updating the requirements that guide the infrastructure co-design. While the Fenix Council will

Fig. 2 ICEI Public Information Event on 15 March 2018 in Barcelona

also determine the general principles for the allocation of Fenix resources within stakeholder communities (which must be based on peer review), each community will be responsible for the actual distribution of their share within that community. Overall, of the Fenix resources funded through ICEI, a total share of 25% will be made available to the HBP, and 15% to European researchers at large (through PRACE). Fenix is open for other communities who want to contribute resources to and use the infrastructure.

According to the ICEI project schedule, the deployment of the first Fenix infrastructure components at the majority of the participating centers and the first demonstration of the key services are expected to start towards the end of 2019. All infrastructure services are planned to be operational in early 2021.

References

[1] Human Brain Project, https://www.humanbrainproject.eu

[2] Framework Partnership Agreement – Signature (press release), https://www.humanbrainproject.eu/en_GB/fpa-signature

[3] Human Brain Project: Towards a European infrastructure for brain research, http://inside.hlrs.de/editions/16autumn.html#human-brain-project-towards-a-european-infrastructure-for-brain-research (Reference to an article in a previous InSiDE edition)

[4] Human Brain Project successfully entered the next Phase (Reference: see page 17)

Written by Boris Orth
Jülich Supercomputing Centre (JSC)
Contact: b.orth@fz-juelich.de

Dirk Pleiter
Jülich Supercomputing Centre (JSC)
Contact: d.pleiter@fz-juelich.de

GPU Hackathon Dresden 2018

The Dresden/Jülich Hackathon—taking place March 5–March 9, 2018 at TU Dresden—as the first in a series of GPU Hackathons planned for 2018. The event was locally organized by ZIH Dresden and Helmholtz Zentrum Dresden Rossendorf (HZDR). More than 50 people from different scientific fields attended, with the goal of accelerating codes using graphics processing units (GPUs).

Teams from all over Europe applied to participate in the hackathon. Eight teams were admitted and participated in an intense week-long event of optimization and conversion about scientific programs. Mentors from academia and industry, all experts in parallel programming, assisted with programming strategies, tools, and consulting on techniques and tools for concurrent programming. Teams could quickly evaluate new ideas while sitting together in one large room and, if necessary, discuss their thoughts beyond the scope of their team with experts from other groups.

Team Ptycho_Imaging, for example, attended with an x-ray imaging application written in Matlab that approaches its performance limits quickly. During the hackathon, together with their two mentors, they converted the code base to Python and wrote an interface to C++, from where they call the GPU-accelerated solvers offered by PETSc. During the course of the week, they went from a serial, prototype-like application working on a laptop, to a HPC-level program making use of the massively parallel performance offered by GPUs.

The team LeMonADErs from the Leibniz Institute of Polymer Research Dresden was one of several teams, using GPUs for accelerating numerical simulations. Their molecular dynamics application was already ported to the GPU architecture

All participants and mentors that joined the GPU Hackathon. © Guido Juckeland

prior to the event and optimized thoroughly. During the Hackathon, the team investigated further, low-level optimizations and strengthened their underlying physical simulation by integrating various external libraries. Not all of the evaluated strategies were successful, but testing during the concentrated atmosphere made the team's investigation efficient. The team valued the scientific discourse during the event.

Team McStas-McXtrace optimized a set of physics applications for tracing neutron and x-ray scattering experiments. Their legacy codebase proved hard to accelerate, so the team focused on creating a code generator for a specific function which then contained GPU-accelerated code. They achieved a three-fold speedup for the first GPU version of the application. The main benefit of the hackathon for the team was the ability to jumpstart the modernization of its application with a GPU-parallel backend in mind, allowing them to gain first insights into the ecosystem.

Guido Juckeland, the main organizer, moderated the hackathon. Every day adhered to a certain schedule: Teams meet in the morning and work until lunch, after which every team presents their current status, their successes and failures, and their future plans in a short, two-minute stand-up presentation. Afterwards it is back to hacking away until the evening. Guiding goals were set for each day: At the beginning of the week, it was stock-taking of the current state of the apps. Plans were formulated, first easy-to-attain, then harder ones. Finally, advanced topics were tackled, like asynchronous execution of functions of the program. At the end of the week, all teams achieved their self-set goals. For many of them it was an achieved net-speedup of their application, now running on GPUs. Others were able to extend the scope of their applications and libraries. All learned valuable software engineering methods along the way and could work together closely with HPC researchers and scientists.

The next GPU hackathon in Europe is at CSCS in Switzerland, October 1–5, 2018. See the official website at *https://www.olcf.ornl.gov/calendar/2018-gpu-hackathons/* for upcoming registration deadlines.

References

[1] http://www.dresden-concept.de/home/news-detailseite/article/marathon-fuer-hacker-internationaler-grafikprozessoren-hackathon-findet-in-dresden-statt.html

[2] https://gcoe-dresden.de/the-hack-is-on-in-2018/

[3] https://gcoe-dresden.de/meet-the-gpuhackdd-2018-mentors/

[4] https://twitter.com/search?q=%23gpuhackdd

Written by Andreas Herten
Jülich Supercomputing Centre (JSC)
Contact: a.herten@fz-juelich.de

Dirk Pleiter
Jülich Supercomputing Centre (JSC)
Contact: d.pleiter@fz-juelich.de

Guido Juckeland
Helmholtz Zentrum Dresden Rossendorf
Contact: g.juckeland@hzdr.de

The Promise and Perils of Simulation in the Social Sciences

Although computer simulation has provided tools for the social sciences for more than 50 years, it remains a method awaiting its breakthrough discovery. In part this is because—in contrast with physical systems—it is extremely difficult to comprehensively represent phenomena like human behavior, social interactions, and historical change in computer models.

The success of simulation across many other scientific domains has nevertheless invited growing interest in using computers in social research, particularly for supporting political decision-making. But troubling questions about simulation's abilities remain: Are today's computational social sciences mature enough to make such predictions reliably? Would it make a difference if we gained more and better data, more computing power, or better algorithms? What questions can we expect simulation to answer in the future? And what problems will always be too complex for simulation to solve?

Recently, members of the HLRS Department of Philosophy of Science and Technology of Computer Simulation, under the direction and initiative of HLRS Director Prof. Michael Resch, launched a new collaboration with investigators in the Friedrich-Alexander Universität Erlangen-Nürnberg (FAU) Institute for Sociology that will focus on the intersection of simulation and the social sciences. On March 9–10, 2018, the group held its inaugural symposium in Erlangen to begin identifying social science problems that simulation could help address and developing a better understanding of simulation's capabilities

Participants in the symposium, "Simulation in the Social Sciences and the Sociology of Simulation." Photo: HLRS

and limits. The meeting was the first in a series of annual symposia that will be held alternately in Stuttgart and Erlangen in the coming years.

Titled "Simulation in the Social Sciences and the Sociology of Simulation," the symposium brought together two key communities of social scientists—those who use simulation to do social research and those who study how computational social research is interpreted and used in society. The event also featured talks by HLRS simulation scientists who discussed scientific uses of simulation and its relevance for society.

The gathering revealed that in both "toy models"—used, for example, in game theory—and „microscale" models based on large collections of empirical data, simulation will always face limitations, as input conditions must be dramatically simplified. For example, models can not reflect the emergence and effects of things like public opinion, laws, and different cultural backgrounds. In addition, no model can predict the

consequences of unforeseen events like wars, mass migrations, or sudden political realignments like Brexit, which have major impacts on societies.

"Are social scientists condemned to only developing toy models?" asked Prof. Dr. Nicole J. Saam, a co-organizer of the symposium, in a follow-up interview. "This is an important question that we as social scientists need to discuss. Some people think it should be possible to simulate pretty much anything and come to us with completely unrealistic expectations. We need to understand the foundations of simulation better, particularly what kinds of simulation models can't be developed. This will help us to learn what models we can indeed develop, what characteristics these models have, and what kinds of knowledge they can produce."

The symposium also considered the social environment of simulation. Whenever simulations are developed—whether in the social sciences or physical sciences—the context in which they originate is often obscure to those who did not participate. In the natural sciences, for example, simulations often arise out of collaborations among large groups of investigators. For social scientists, understanding the social processes through which such projects grow also presents interesting questions.

"Considering the size of the research groups and the number of different disciplines involved in the development of technologically and mathematically complex computer simulations, it is understandable that it would be difficult even for simulation scientists to understand them completely," says Dr. Andreas Kaminski, leader of the HLRS Department of Philosophy of Science and Technology of Computer Simulation. Philosophers of simulation refer to a condition called „epistemic opacity," the fact that not only nature but also the scientific method itself can appear opaque.

"Experts in the field of simulation sciences are beginning to become aware of this problem," Kaminski explains. "Now it has also become very important that those who want to use computer simulations from a social science perspective or in the political field understand the scientific method behind simulation, as well as its limitations."

The Erlangen workshop was one step toward this goal. By bringing together critical perspectives from the social sciences, philosophy of science, and simulation science, the group aims to chart a path forward for simulation in social research. The next workshop will take place at HLRS in spring, 2019, with the location of the meeting alternating between Stuttgart and Erlangen in the following years.

Written by Christopher Williams
High Performance Computing Centre Stuttgart (HLRS)
Contact: williams@hlrs.de

HLRS Strengthens Collaborations in Asia

In spring, 2018, senior members of the High-Performance Computing Center Stuttgart (HLRS) travelled to Asia to meet with leaders of several prominent computing centers. The tour offered opportunities to exchange ideas about the state of the art in high-performance computing (HPC)and to discuss opportunities for collaboration.

In Daejon, South Korea, HLRS Director Michael Resch met with Choi Heeyoon, president of the Korean Institute of Science and Technology Information (KISTI). At the meeting—which followed up on the renewal of a memorandum of understanding signed by the two institutions in 2017—conversation identified collaboration opportunities that the two organizations could pursue in a variety of fields including HPC, data analytics, and machine learning.

One specific topic the institutes plan to pursue together involves developing new tools that could be used in digital product design and fabrication. Researchers at KISTI have been developing a technology for these purposes called HEMOS-Fluid. HEMOS (which stands for High-Performance Engineering Modeling & Simulation) is a software platform designed to support fluid analysis, particularly in the large-scale post-processing of simulation data for both large data displays and desktop monitors.

At the meeting, the institutes agreed to explore how HEMOS-Fluid might be combined with COVISE (Collaborative Visualization and Simulation Environment), a technology developed at HLRS. COVISE is an extendable, distributed software environment that seamlessly integrates simulation results, post-processing, and visualization.

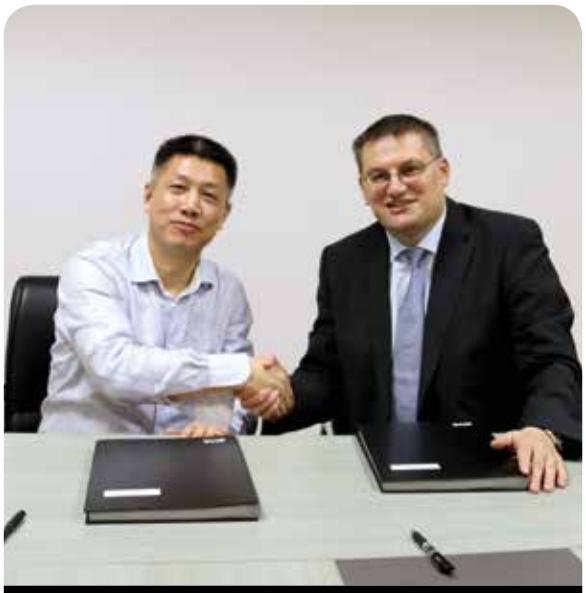

Supercomputing Center of Wuhan University Director Dengyi Zhang and HLRS Director Michael Resch sign the new collaboration agreement.
© HLRS

Ultimately, Dr. Heeyoon and Dr. Resch anticipate that combining these two technologies could offer a powerful new approach for visualizing simulation data. Pooling their institutes' expertise and strengthening mutual collaboration on this topic, they also stated, could ultimately benefit the competitiveness of small- and medium-scale enterprises (SMEs) that use simulation in their product design and development.

Later in the trip, Dr. Resch traveled to Wuhan, China, where he signed a collaboration

KISTI President Dr. Choi Heeyoon (speaking, left) meets with Dr. Michael Resch and visitors from HLRS in Daejeon, South Korea. © KISTI

agreement with Supercomputing Center of Wuhan University Director Dengyi Zhang. The agreement pledges collaboration for the next three years.

HLRS and the Supercomputing Center at Wuhan University plan to exchange scientists and to focus on key research topics in HPC. Both sides will also share their experiences in installing large-scale computing systems, particularly because both Wuhan and Stuttgart aim to develop exascale systems.

In addition to its collaborations with KISTI and Wuhan University, HLRS has made formal agreements of collaboration with other HPC centers in Asia, including the Shanghai Supercomputing Center, the National Center for High Performance Computing (Taiwan), and the Cyberscience Center, Tohuku University (Sendai, Japan). HLRS also maintains agreements with many other HPC centers in Europe, Russia, and the United States aimed at promoting the exchange of HPC knowledge and expertise.

Written by Christopher Williams
High Performance Computing Centre Stuttgart (HLRS)
Contact: williams@hlrs.de

Cosmologists Create Largest Simulation of Galaxy Formation, Break Their Own Record

Humans have long tried to explain how stars came to light up the night sky. The wide array of theories throughout history have one common (and correct) governing principle that astrophysicists still use to this day: by understanding the stars and their origins, we learn more about where we come from.

However, the vastness of our galaxy—let alone our entire universe—means experiments to understand our origins are expensive, difficult, and time consuming. In fact, experiments are impossible for studying certain aspects of astrophysics, meaning that in order to gain greater insight into how galaxies formed, researchers rely on supercomputing.

In an attempt to develop a more complete picture of galaxy formation, researchers from the Heidelberg Institute for Theoretical Studies, the Max-Planck Institutes for Astrophysics and for Astronomy, the Massachusetts Institute of Technology, Harvard University, and the Center for Computational Astrophysics in New York have turned to supercomputing resources at the High-Performance Computing Center Stuttgart (HLRS. The resulting simulation will help to verify and expand on existing experimental knowledge about the universe's early stages.

Recently, the team expanded on its 2015 record-breaking "Illustris" simulation—the largest-ever hydrological simulation of galaxy formation. Hydrodynamic simulations allow researchers to accurately simulate the movement of gas. Stars form from cosmic gas, and stars' light provides astrophysicists and cosmologists with important information for understanding how the universe works.

The researchers improved on the scope and accuracy of their simulation, naming this phase of the project, "Illustris, The Next Generation," or "IllustrisTNG." The team released its first round of findings across three journal articles appearing in the *Monthly Notices of the Royal Astronomical Society* and are preparing several more for publication.

Magnetic modelling

Just as humanity cannot envision exactly how the universe came to be, a computer simulation

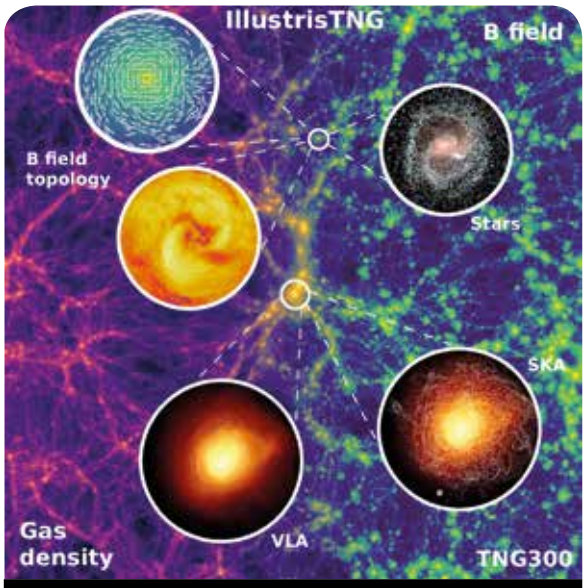

Gas density (left) and magnetic field strength (right) across the TNG300 box, centered on the most massive galaxy cluster. Zoomed panels show the magnetic field orientation and stellar light (top) and xray and radio emission from massive cluster (bottom). © Illustris team

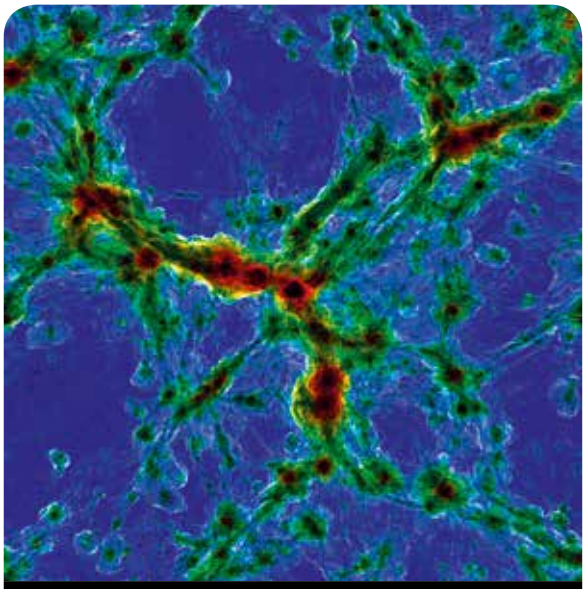

Composite which combines gas temperature (as the color) and shock mach number (as the brightness). Red indicates 10 million Kelvin gas at the centers of massive galaxy clusters, while bright structures show diffuse gas from the intergalactic medium shock heating at the boundary between cosmic voids and filaments. © Illustris team

cannot recreate the birth of the universe in a literal sense. Instead, researchers feed equations and other starting conditions—observations coming from satellite arrays and other sources—into a gigantic computational cube representing a large swath of the universe and then use numerical methods to set this "universe in a box" in motion.

For many aspects of the simulation, researchers can start their calculations at a fundamental, or ab initio, level with no need for preconceived input data, but processes that are less understood—such as star formation and the growth of supermassive black holes—need to be informed by observation and by making assumptions that can simplify the deluge of calculations.

As computational power and know-how have increased, so too has the ability to simulate larger areas of space and increasingly intricate and complex phenomena related to galaxy formation. With IllustrisTNG, the team simulated 3 different universe "slices" at different resolutions. The largest was 300 megaparsecs across, or roughly 1 billion light years. The team used 24,000 cores on Hazel Hen over the span of 35 million core hours.

In one of IllustrisTNG's major advances, the researchers reworked the simulation to include a more precise accounting for magnetic fields, improving the simulation's accuracy. "Magnetic fields are interesting for a variety of reasons," said Prof. Dr. Volker Springel, Professor and researcher at the Heidelberg Institute for Theoretical Studies and principal investigator on the project. "The magnetic pressure exerted on cosmic gas can occasionally be equal to thermal (temperature) pressure, meaning that if you neglect this, you will miss these effects and ultimatelwwwy compromise your results."

While developing IllustrisTNG, the team also made a surprising advance in understanding black hole physics. Based on observational knowledge, the researchers knew that supermassive black holes propel cosmic gases with a lot of energy while also "blowing" this gas away from galaxy clusters. This helps to "shut off" star formation in the biggest galaxies and thus imposes a limit on the maximum size they can reach.

In the previous Illustris simulation, researchers noticed that while black holes go through this energy transfer process, they would not shut off the star formation completely. By revising the black holes' physics in the simulation, the team saw much better agreement between the data and observation, giving researchers greater confidence that their simulation corresponds to reality.

A long-standing alliance

The team has been using GCS resources since 2015 and been running the IllustrisTNG simulation on HLRS resources since March 2016. Considering that IllustrisTNG's dataset is both larger and more accurate than the original, the researchers are confident their data will be used far and wide while they apply for more time to continue refining the simulation. The original Illustris data release garnered 2,000 registered users and resulted in more than 130 publications.

During that time, the researchers have relied on GCS support staff to help with several low-level issues related to their code, specifically related to memory crashes and file system issues. Team members Drs. Dylan Nelson and Rainer Weinberger also both benefitted from attending 2016 and 2017 machine-level scaling workshops at HLRS. The team's long-standing collaboration with HLRS has resulted in winning 2016 and 2017 Golden Spike awards, which are given to outstanding user projects during HLRS' annual Results and Review Workshop.

Nelson pointed out that while current-generation supercomputers have enabled simulations that have largely overcome most fundamental issues related to massive-scale cosmological modelling, there are still opportunities for improvement.

"Increased memory and processing resources in next-generation systems will allow us to simulate large volumes of the universe with higher resolution," Nelson said. "Large volumes are important for cosmology, understanding the large-scale structure of the universe, and making firm predictions for the next generation of large observational projects. High resolution is important for improving our physical models of the processes going on inside of individual galaxies in our simulation."

This research used Gauss Centre for Supercomputing resources based at the High-Performance Computing Center Stuttgart.

Written by Eric Gedenk

Gauss Centre for Supercomputing (GCS), based at High Performance Computing Centre Stuttgart (HLRS)

Contact: gedenk@hlrs.de
eric.gedenk@gauss-centre.eu

Researchers Create Largest, Longest Multiphysics Earthquake Simulation to Date

Just before 8:00 a.m. local time on December 26, 2004, people in southeast Asia were starting their days when the third strongest recorded earthquake in history ripped a 1,500-kilometer tear in the ocean floor off the coast of the Indonesian island of Sumatra.

The earthquake lasted between 8 and 10 minutes (one of the longest ever recorded), and lifted the ocean floor several meters, creating a tsunami with 30-meter waves that devastated whole communities. The event caused nearly 200,000 deaths across 15 countries, and released as much energy above and below ground as multiple centuries of US energy usage.

The Sumatra-Andaman Earthquake, as it is called, was as surprising as it was violent. Despite major advancements in earthquake monitoring and warning systems over the last 50 years, earth scientists were unable to predict it because relatively little data exists about such large-scale seismological events. Researchers have a wealth of information related to semi-regular, lower-to-medium-strength earthquakes, but disasters such as the Sumatra-Andaman—events that only happen every couple hundred years—are too rare to create reliable data sets.

In order to more fully understand these events, and hopefully provide better prediction and mitigation methods, a team of researchers from the Ludwig-Maximilians-Universität Munich (LMU) and Technical University of Munich (TUM) is using supercomputing resources at the Leibniz Supercomputing Centre (LRZ) to better understand these rare, extremely dangerous seismic phenomena.

"Our general motivation is to better understand the entire process of why some earthquakes and resulting tsunamis are so much bigger than others," said TUM Professor Dr. Michael Bader. "Sometimes we see relatively small tsunamis when earthquakes are large, or surprisingly large tsunamis connected with relatively small earthquakes. Simulation is one of the tools to get insight into these events."

The team strives for "coupled" simulations of both earthquakes and subsequent tsunamis. It recently completed its largest earthquake simulation yet. Using the SuperMUC supercomputer at LRZ, the team was able to simulate 1,500 kilometers of non-linear fracture mechanics—the earthquake source—coupled to seismic waves traveling up to India and Thailand over a little more than 8 minutes of the Sumatra-Andaman earthquake. Through several in-house computational innovations, the team achieved a 13-fold improvement in time to solution. In recognition of this achievement, the project won the best paper award at SC17, one of the world's premier supercomputing conferences, held this year on November 12–17 in Denver, Colorado.

Megathrust earthquakes, massive scale simulations

Earthquakes happen as rock below Earth's surface breaks suddenly, often as a result of the slow movement of tectonic plates.

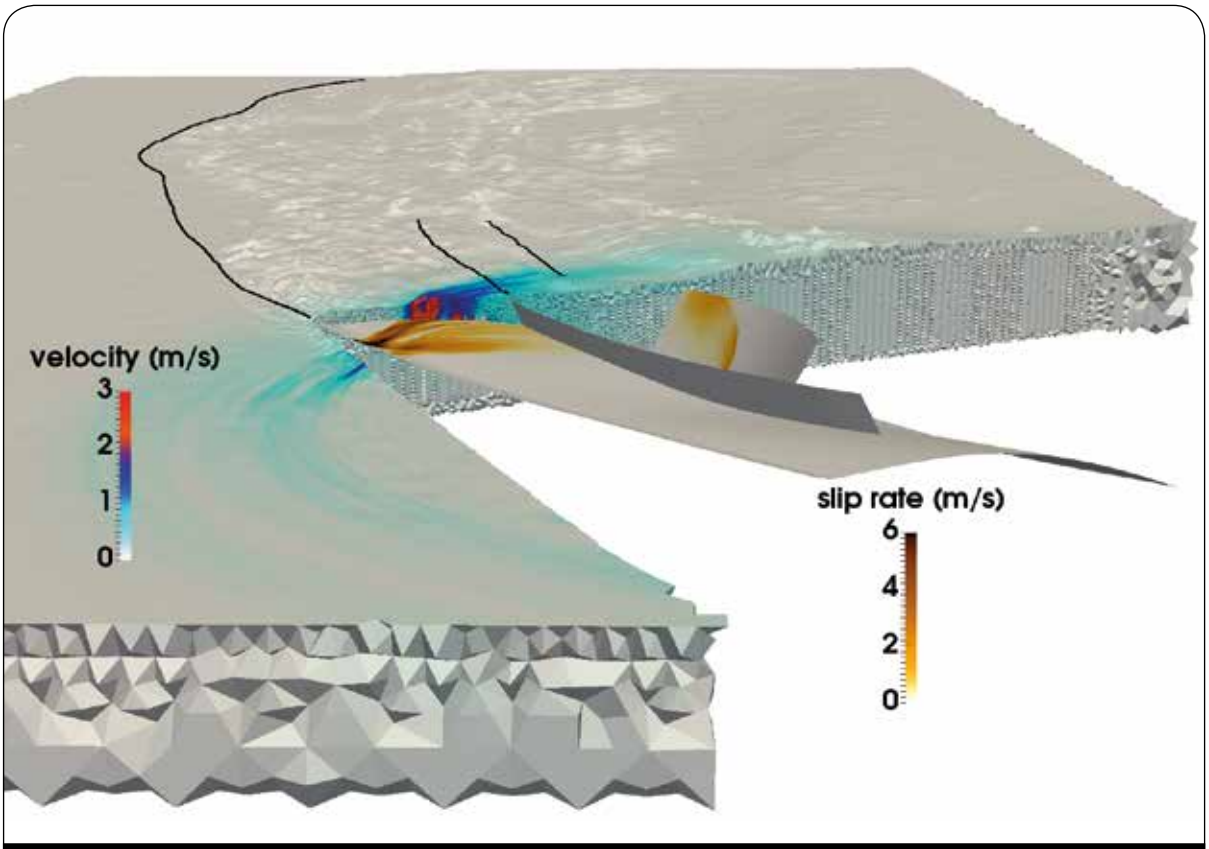

Rupture propagation and the resulting seismic wave field during 2004 Sumatra-Andaman earthquake.
Copyright: C. Uphoff, S.Rettenberger, M. Bader, Technical University of Munich. E. Madden, T. Ulrich, S. Wollherr, A. Gabriel, Ludwigs-Maximilians-Universität.

One rough predictor of an ocean-based earthquake's ability to unleash a large tsunami is whether plates are grinding against one another or colliding head-on. If two or more plates collide, one plate will often force the other below it. Regions where this process occurs are called subduction zones and can host very large, shallowly dipping faults—so called "megathrusts." Energy release across such huge zones of weakness tends to create violent tsunamis, as the ocean floor rises a significant amount, temporarily displacing large amounts of water.

Until recently, though, researchers doing computational geophysics had great difficulties simulating subduction earthquakes at the necessary level of detail and accuracy. Large-scale earthquake simulations are difficult generally, but subduction events are even more complex.

"Modeling earthquakes is a multiscale problem in both space and time," said Dr. Alice Gabriel, the lead researcher from the LMU side of the team. "Reality is complex, meaning that incorporating the observed complexity of earthquake

sources invariably involves the use of numerical methods, highly efficient simulation software, and, of course, high-performance computing (HPC). Only by exploiting HPC can we create models that can both resolve the dynamic stress release and ruptures happening with an earthquake while also simulating seafloor displacement over thousands of kilometers."

When researchers simulate an earthquake, they use a computational grid to divide the simulation into many small pieces. They then compute specific equations for various aspects of the simulation, such as generated seismic shaking or ocean floor displacement, among others, over "time steps," or simulation snapshots over time that help put it in motion, much like a flip book.

The finer the grid, the more accurate the simulation, but the more computationally demanding it becomes. In addition, the more complex the geometry of the earthquake, the more complex the grid becomes, further complicating the computation. To simulate subduction earthquakes, computational scientists have to create a large grid that can also accurately represent the very shallow angles at which the two continental plates meet. This requires the grid cells around the subduction area to be extra small, and often slim in shape.

Unlike continental earthquakes, which have been better documented through computation and observation, subduction events often happen deep in the ocean, meaning that it is much more difficult to constrain a simulation by ground shaking observations and detailed, reliable data from direct observation and laboratory experiments.

Furthermore, computing a coupled, large-scale earthquake-tsunami simulation requires using data from a wide variety of sources. Researchers must take into account the seafloor shape, the shape and strength of the plate boundary ruptured by the earthquake and the material behaviour of Earth's crust at each level, among other aspects. The team has spent the last several years developing methods to more efficiently integrate these disparate data sources into a consistent model.

To reduce the enormous computing time, the team exploited a method called "local time stepping." In areas where the simulations require much more spatial detail, researchers also must "slow down" the simulation by performing more time steps in these areas. Other sections that require less detail may execute much bigger and thus far fewer time steps.

If the team had to run its entire simulation at a uniform small time step, it would have required roughly 3 million individual iterations. However, only few cells of the computational grid required this time step size. Major parts could be computed with much larger time steps, some requiring only 3000 time steps. This reduced the computational demand significantly and led to much of the team's 13-fold speedup. This advancement also led to the team's simulation being the

largest, longest first-principles simulation of an earthquake of this type.

Forward motion

Due to its close collaboration with LRZ staff, the team had opportunities to use the entire SuperMUC machine for its simulations. Bader indicated that these extremely large-scale runs are invaluable for the team to gain deeper insights in its research. "There is a big difference if you run on a quarter of a machine or a full machine, as that last factor of 4 often reveals the critical bottlenecks," he said.

The team's ability to take full advantage of current-generation supercomputing resources has it excited about the future. It's not necessarily important that next-generation machines offer the opportunity for the LMU-TUM researchers to run "larger" simulations—current simulations can effectively simulate a large enough geographic area. Rather, the team is excited about the opportunity to modify the input data and run many more iterations during a set amount of computing time.

"We have been doing one individual simulation, trying to accurately guess the starting configuration, such as the initial stresses and forces, but all of these are still uncertain," Bader said. "So we would like to run our simulation with many different settings to see how slight changes in the fault system or other factors would impact the study. These would be larger parameter studies, which is another layer of performance that a computer would need to provide."

Gabriel also mentioned that next-generation machines will hopefully be able to simulate urgent, real-time scenarios that can help predict hazards as they relate to likely aftershock regions. The team is excited to see the next-generation architectures at LRZ and the other GCS centres, the High-Performance Computing Center Stuttgart and the Jülich Supercomputing Centre.

In Bader's view, the team's recent work not only represents its largest-scale simulation to date, but also the increasingly strong collaboration between the domain scientists and computational scientists in the group. "This paper has a strong seismology component and a strong HPC component," he said. "This is really a 50-50 paper for us. Our collaboration has been going nicely, and it is because it isn't about getting ours or theirs. Both groups profit, and this is really nice joint work."

This research used Gauss Centre for Supercomputing resources based at the Leibniz Supercomputing Centre.

Written by Eric Gedenk

Gauss Centre for Supercomputing (GCS), based at High Performance Computing Centre Stuttgart (HLRS)

Contact: gedenk@hlrs.de
eric.gedenk@gauss-centre.eu

Researchers Compute Their Way to the Centre of the Earth

In order to more fully comprehend the complexities of Earth's interior, humanity has to dig deep—literally. To date, scientists have been able to bore a little over 12 kilometres deep, or about half the average depth of the Earth's crust.

Why would researchers need to peer into deeper depths? Both to better understand how the earth formed and how the interior might have an effect on our life on the surface of the Earth today, such as by the magnitude and reversals of the Earth's magnetic field.

However, experiments investigating materials at conditions deep in the Earth are challenging, meaning that to continue gaining insights into these phenomena, experimentalists must turn to modeling and simulation to support and complement their efforts.

To that end, researchers at the University of Cologne's Institute for Geology and Mineralogy have turned to computing resources at the Jülich Supercomputing Centre (JSC) to help better comprehend how materials behave in the extreme conditions below the surface of the earth.

The team, led by University of Cologne's Prof. Dr. Sandro Jahn and Dr. Clemens Prescher, has been using JSC's JUQUEEN supercomputer to simulate the structure of melts by studying silicate glasses as a model system for melts under ultra-high pressures. The team recently published its initial findings in the Proceedings of the National Academy of Sciences.

"Understanding properties of silicate melts and glasses at ultra-high pressure is crucial to understand how the Earth has formed in its infancy, where impacts of large asteroids led to a completely molten Earth," said Prescher. "In fact, all of the internal layered structure we know today was formed in such events."

It's a glass

When most people think of the word glass, they think of windows or bottles. Glass, however, is a term describing a wide range of non-crystal solids. Atoms in a solid can organize themselves in a variety of ways, and materials considered glasses have some of the more "chaotic" atomic structures possible in solids.

A glass can also be seen as a frozen melt. Thus by understanding the properties of glasses at ultra-high pressures, researchers can gain insights into the melts' properties in the deep Earth's interior, providing a clearer view into the physical processes which made the Earth and might be still occurring today.

Using a variety of geophysics measurements and laboratory experiments, researchers are capable of gaining some degree of insight into material properties under certain pressure conditions without actually being able to make direct observations. Enter supercomputing. As computing power has gotten stronger, geophysics researchers are able to complement and expand their studies of these inner-Earth processes through the use of numerical models.

In the case of the University of Cologne researchers, they wanted to get a more detailed insight into the structure of the silicate glass than their experimental efforts were able to provide. The team utilized ab initio calculations of atoms' electronic structures and put these calculations in motion using molecular dynamics simulations. Ab initio calculations mean that researchers start with no assumptions in their mathematical models, making a simulation more computationally expensive but also more accurate.

Due to having many calculations for each atom's structure and computationally demanding molecular dynamics calculations, the team keeps its simulations relatively small in scale—the team's largest runs typically have between 200-250 atoms in the simulation. This size allows the team to run simulations under a variety of different pressure and temperature combinations, ultimately allowing it to calculate a small but representative sample of material interactions under a variety of conditions.

To test its model and lay the foundation for modeling increasingly complex material interactions, the team decided to simulate silicon dioxide (SiO_2), a common, well-studied material, most well-known as the compound that forms quartz.

Among silicate materials, SiO_2 is a good candidate on which to base computational models—researchers already understand how its atomic structure patterns and material properties change under a variety of pressure conditions.

The team chose to focus on a relatively simple, well-known material in order to expand the range of pressure it could simulate and attempt to validate the model with experimental data. Using JUQUEEN, the team was able to extend its investigation well beyond the experimentally achieved 172 Gigapascals, corresponding to 1.72 million times the Earth's atmospheric pressure, or roughly the amount of pressure the Eiffel Tower would apply by pressing down on the tip of a person's finger.

The researchers also found that at high pressures, oxygen atoms are much more compressible than silicon atoms. The varying size ratio between the two leads to hugely different glass structures of SiO_2 at low and at high pressures.

Digging Deeper

By validating its model, the team feels confident that it can move on to more complex materials and interactions. Specifically, the team hopes to expand its investigations deeper into the realm of melts. Think of lava as a melt—molten rock erupts from below the earth's surface, rapidly cools when it reaches the surface, and may form obsidian, a glassy rock.

In order to do more advanced simulations of melts, the team would like to be able to expand its simulations to account for a wider range of chemical processes as well as expand the number of atoms in a typical run.

As JSC and the other two Gauss Centre for Supercomputing (GCS) facilities—the

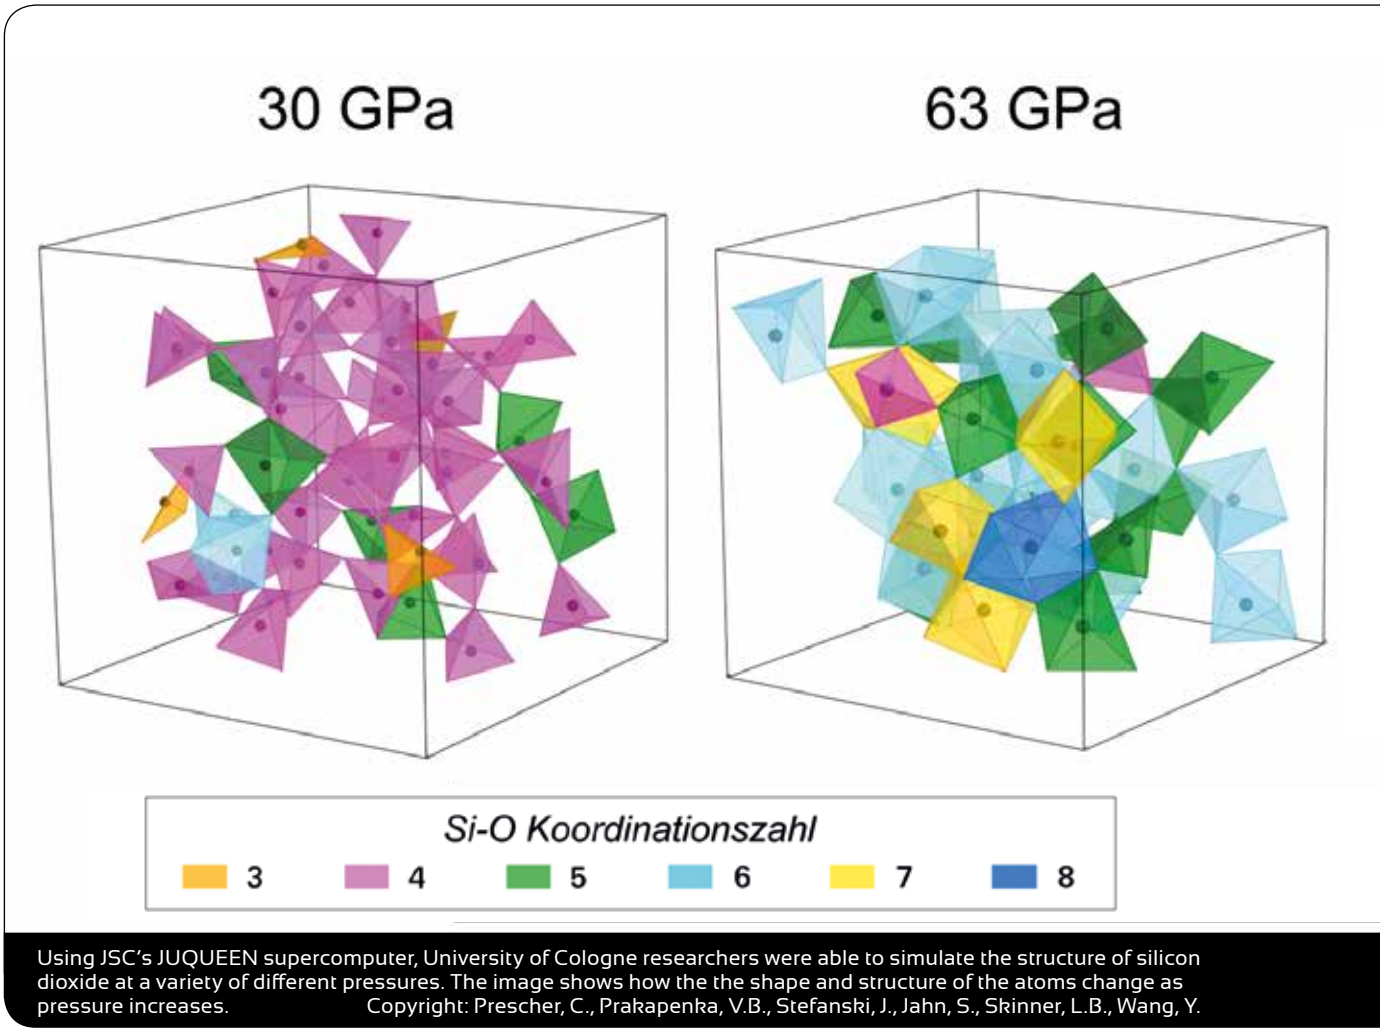

Using JSC's JUQUEEN supercomputer, University of Cologne researchers were able to simulate the structure of silicon dioxide at a variety of different pressures. The image shows how the the shape and structure of the atoms change as pressure increases. Copyright: Prescher, C., Prakapenka, V.B., Stefanski, J., Jahn, S., Skinner, L.B., Wang, Y.

High-Performance Computing Center Stuttgart and the Leibniz Supercomputing Centre in Garching—install next-generation supercomputers, the team is confident that they will be able to gain even greater insight into the wide range of complex material interactions happening many kilometres below the surface.

"A faster machine will enable us to simulate more complex melts and glasses, which is crucial to go from model systems, such as SiO2 glass in this study, to the real-world compositions we expect in the Earth's interior," Prescher said.

Prescher also noted that JSC support staff helped the team work more efficiently by assisting with implementing the team's code.

This type of support represents GCS' plans for the future. With the promise and opportunity

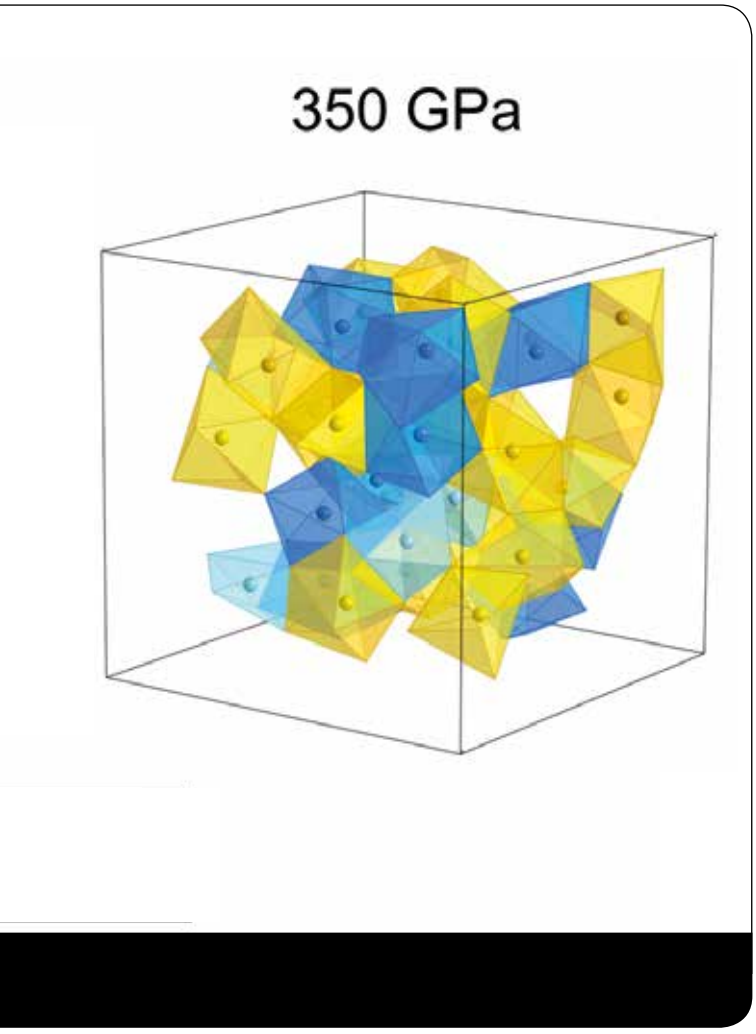

connected to next-generation computing architectures, GCS centre leadership realizes that closer collaboration with users and application co-design will be a key component for ensuring researchers can efficiently solve bigger, more complex scientific problems.

Whether studying deep in space among the stars or deep below the surface of the Earth, the collaboration between supercomputing centres and researchers will play an increasingly important role in solving the world's toughest scientific challenges.

This research used Gauss Centre for Supercomputing resources based at the Jülich Supercomputing Centre.

Written by Eric Gedenk

Gauss Centre for Supercomputing (GCS), based at High Performance Computing Centre Stuttgart (HLRS)

Contact: gedenk@hlrs.de
eric.gedenk@gauss-centre.eu

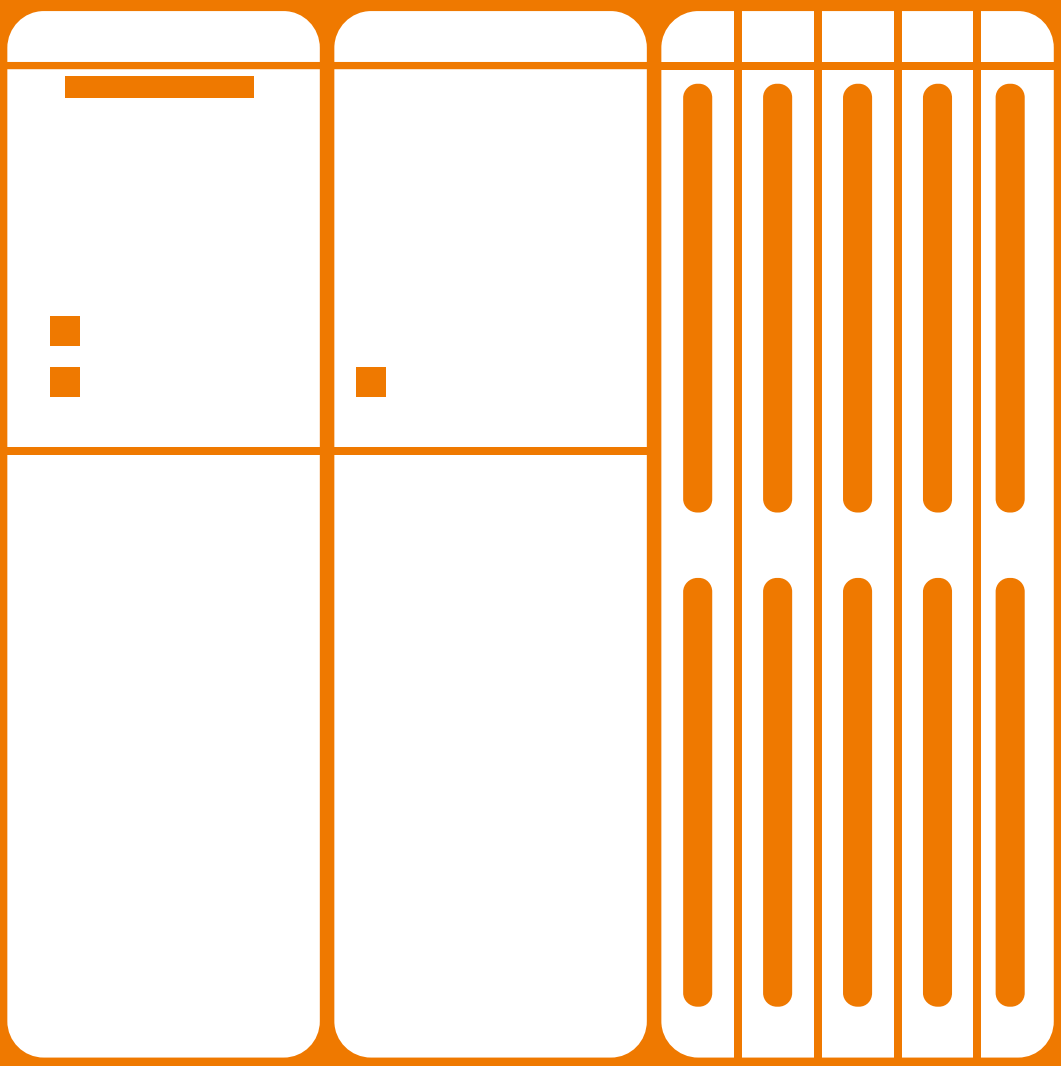

Applications

In this section you will find the most noteworthy applications of external users on GCS computers.

Optimization and HPC-Applications of the Flow Solver FLOWer

With the continual increase in supercomputing power, more advanced numerical methods become feasible and numerical investigations in research and development of aircraft can focus on demanding aerodynamic phenomena previously too computationally intensive. The helicopter and aeroacoustics group within the Institute of Aerodynamics and Gas Dynamics (IAG) at the University of Stuttgart specializes in high-fidelity, unsteady Reynolds-averaged Navier–Stokes equation simulations of rotors and helicopters, where the flows are transsonic, highly unsteady, and three-dimensional and complex flow phenomena like blade-vortex interaction or flow separation occur (see Fig. 1).

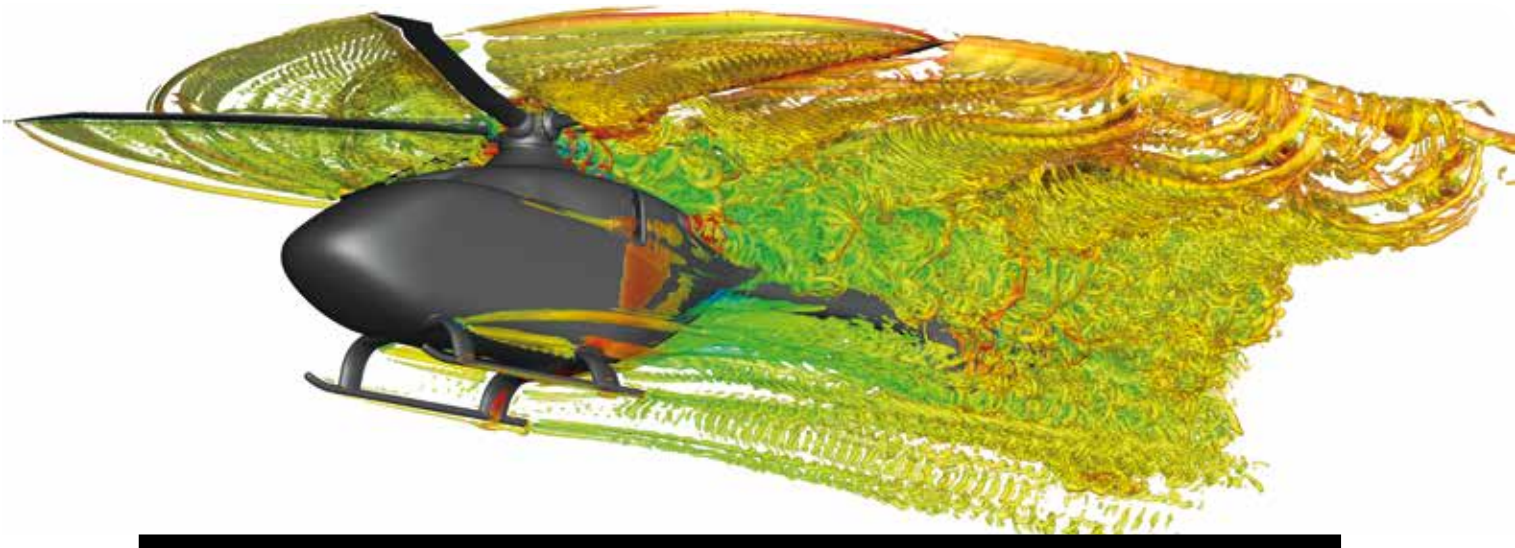

Fig. 1: Visualization of the complex flow field of a modern helicopter.

The "workhorse" code within the group is the block-structured finite volume flow solver FLOWer [1]. Originally developed by the German Aerospace Center (DLR), the code has been significantly enhanced with a fifth-order WENO scheme, hybrid mesh capability, a new coupling library for fluid–structure interactions, and optimized for high-performance computing by IAG. Most recently, it has been expanded by state-of-the-art detached-eddy simulation (DES) methods in order to conduct hybrid RANS/LES simulations [2].

Optimization of Code Parallelization

The scalability of the CFD code FLOWer was recently significantly improved by the implementation of a node-to-node (n2n) communication instead of using core-to-core (c2c) communication [3], allowing simulations with several hundred million grid cells. By using a shared memory to centrally collect the halo data from each single process on a node and a subsequent dispatch of the data as a bundled packet, a significant reduction of the number of messages could be achieved. By sending large data

packages instead of many small messages, a more efficient use of the available bandwidth on Hazel Hen is achieved.

The structured code FLOWer allows the parallelization of the computational effort to run among several cores by dividing the computational domain into blocks. This division is to be carried out in FLOWer by the user with the aim of creating blocks with a preferably equal amount of cells per block. The previous parallelization used a distribution of the blocks to individual computing units with the aim to achieve optimal load balance by equalizing the workload. For this purpose, one to several blocks may also be placed on one core. By providing more fine-grained block units and, in turn, degrees of freedom for the parallelization, an ideal load balance can be better approximated. However, extensive decomposition becomes partially inefficient by, for example, loop initialization overhead. In the previous approach, the communication paths and expense of the halo cells have not been taken properly into account as additional workload. Therefore, a graph partitioning has been introduced, which aims at an optimization of the communication paths for the reduction of the MPI data stream. To carry out the graph partitioning, the widely used METIS library [4] was integrated into the code. A graph is set up whose vertices correspond to blocks, and edges to the communication paths to their neighbor blocks. The vertices and edges are weighted with the block cells and the number of halo cells.

Figure 2 shows the communication paths between individual nodes for an application-oriented computation with 50 nodes. The previous parallelization (Fig. 2a) and the new implementation using the described graph partitioning (Fig 2b) are both illustrated. Due to the stronger weighing of communication in graph

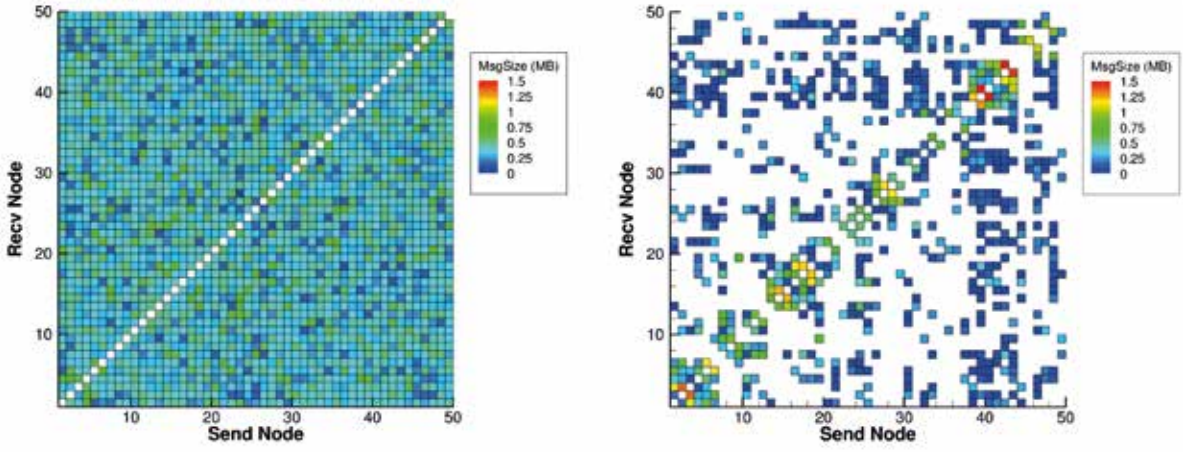

Fig. 2: Comparison of MPI communication pattern using different parallelization schmees in FLOWer.
a) Load balance parallelization b) Graph partitioning

partitioning, there are mainly partitions consisting of spatially adjacent blocks. The chronological assignment of these partitions to the computing nodes results in a concentration of communication paths near the diagonal. Communication with more remote nodes results from the subsequent redistribution of the graph partitioning to reduce the load imbalance. However, these messages are much smaller in their data size. Compared to the old communication pattern, the graph partitioning results in a reduction of the number of communication messages and consolidates communication toward larger messages. In this test case, the overall size of the MPI data stream is reduced by factor 18.3 and the message amount by factor 3.5, with a maximal load imbalance of 2.3 %.

The advances in the optimization of the FLOWer code with regard to its scalability are shown in the Figs. 3a and 3b. A strong and weak scaling depict the benefit of the implementations compared to the previous versions. For weak scaling, the new parallelization scheme (n2n + graph) yields a nearly constant improvement of a 20 % run time speed-up. The scaling shows a similar behavior as the previous parallelization, however, with a slightly reduced slope above 50,000 cores, indicating more efficient MPI communication. In the case of strong scaling, a comparable improvement could be achieved. Overall, very satisfactory scaling results are achieved using the new graph parallelization method.

HPC-Application: Simulation of Static Stall on a Finite Wing

Flow separation plays an important role in the field of aviation. If the angle of attack of a wing or of a rotor blade increases beyond a critical maximum angle, the adverse pressure gradient becomes too strong and eventually, the near-wall flow even reverts its direction. If this happens, the flow detaches from the aerodynamic surface and highly unsteady, three-dimensional and chaotic vortical structures evolve. This

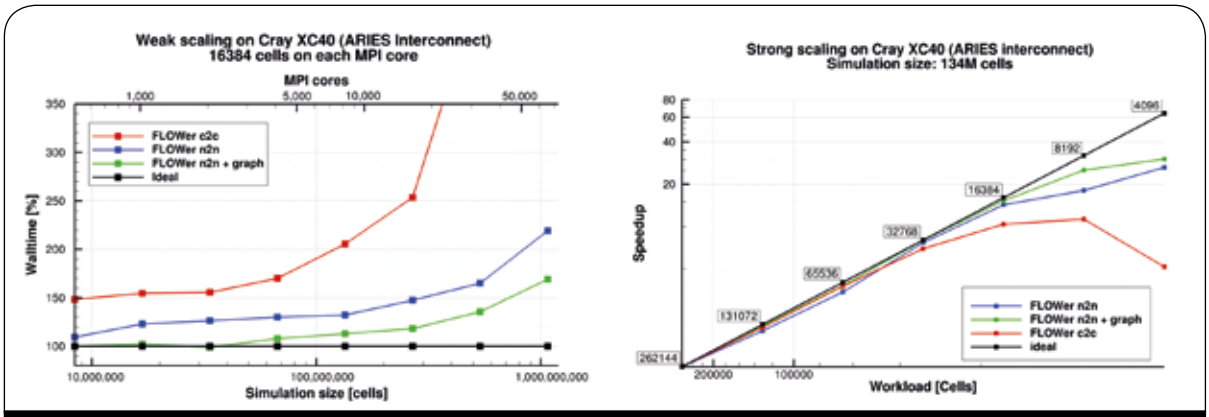

Fig. 3: Scaling studies of the flow solver FLOWer on the Cray XC40 (HazelHen).
a) Weak scaling b) Strong scaling

process is called stall and leads to a more or less sudden reduction in lift and increase in drag and thus limits the flight envelope of an aircraft. The chaotic nature of separated flow presents a great challenge to CFD, as fine spatial and temporal resolutions and low numerical dissipation are essential. Furthermore, the fundamental statistical averaging approach of a URANS simulation makes it inherently impossible to resolve highly unsteady, time-dependent vortical structures.

Therefore, a delayed-detached-eddy simulation (DDES) was carried out to investigate static stall on a finite wing. The wingspan of 1.62 m and the setup is based on an experiment carried out at DLR Göttingen [10]. In total, 150 million grid cells were used, split into 6,800 blocks and computed on 6,000 Hazel Hen cores with a physical time-step size of about 5 microseconds. As Fig. 4 shows, the URANS computation exhibits large-scale vortical structures evolving from the free shear layer in a periodic pattern and the in-plane streamlines indicate rather two-dimensional flow. The vorticity contours of the DDES flow field show several smaller vortical structures distributed over the entire upper side of the wing. Including the streamlines, the flow appears highly chaotic and much more three-dimensional, which is closer to reality.

HPC-Application: Simulation of a Model Contra-Rotation Open Rotor (CROR)

CRORs are a fixed wing aircraft engine concept with two large contra-rotating fans operating at transsonic Mach numbers developed in the 1980s. Recent research concerning eco-efficient engines to meet future fuel saving aims has once again brought the concept to the forefront. The benefit of CRORs lies in their low fuel consumption and the subsequent reduced pollutant emission compared to conventional jet engines enabled by the high bypass ratio. One of the main drawbacks is the high noise emission caused by the contributions of single rotor noise and interaction noise. As a huge part of research on CRORs is undertaken by the industry, recent experimental data is not publicly available. A

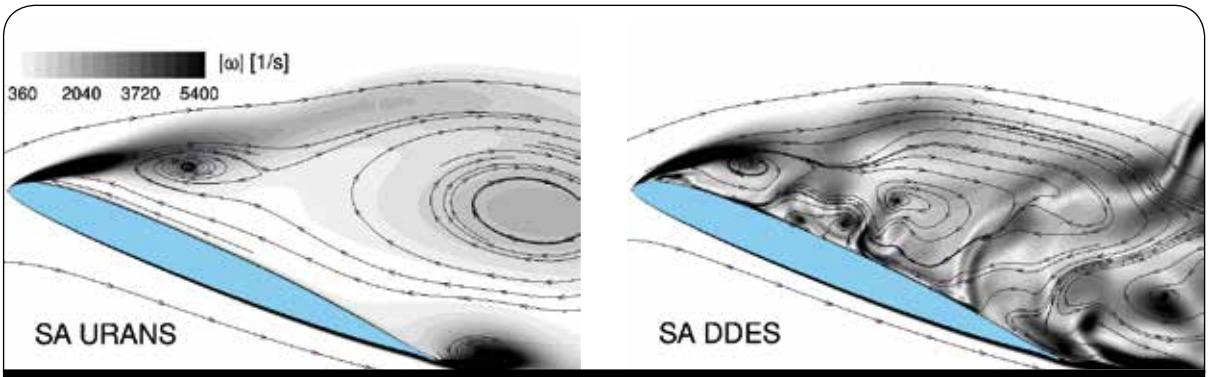

Fig. 4: Comparison of instantaneous vorticity and in-plane streamlines between URANS (4a) and DDES (4b) computations in case of massive flow separation.

test rig featuring a scale model CROR was built at IAG in order to close this gap. The first CFD simulations of the model CROR (see Fig. 5) with different inflow velocities and rotational speeds are presented here to serve as a baseline for validation and further enhancement of the numerical process chain in this application.

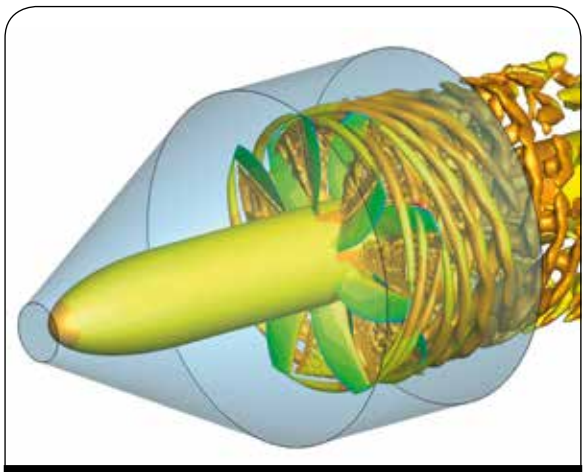

Fig. 5: CROR setup with acoustic hull surface and vortex visualization.

The simulation required 52 grid cells distributed on 1,700 Hazel Hen cores, enabling the computation of six rotor revolutions in 24 hours at a 1-degree time steps. The acoustic evaluation was carried out using the in-house Ffowcs Williams-Hawkings solver ACCO [5]. Figures 6a and 6b show the polar noise characteristics of the CROR model obtained with ACCO at a distance of 10 meters from the center of the CROR. The sound pressure level LP is shown as a function of the polar angle, which is the angle between the connecting line from the CROR's center to the observer and the rotational axis (0 degrees is upstream, 90 degrees is in the rotor plane and 180 degrees is downstream). The sound field is symmetric around the rotational axis so it is fully described by one polar from 0 to 180 degrees.

Overall noise levels (Fig. 6a) increase with rotational speed, as expected. They are also slightly elevated for the zero inflow cases compared to the cases with inflow. This can be attributed to the higher loading in the zero inflow case which causes stronger pressure fluctuations, wakes and tip vortices and thus more single rotor and interaction noise. The directivity patterns are similar for all operating conditions: Noise radiates both upstream and downstream, in both cases slightly inclined in relation to the rotational axis, with the downstream emission louder than the upstream lobe.This behaviour is typical for CRORs in takeoff conditions and is caused by the prevalence of interaction noise, which is typically radiated in oblique directions.

As can be seen from Fig. 6b, the sound field is dominated by the first interaction frequency (which is the sum of the front and rear rotor's blade passing frequencies BPF1 and BPF2). The 8,400-rpm zero inflow case is shown as an example, but the behaviour is similar for all examined operating conditions. Single rotor noise (aft rotor blade passing frequency BPF2 shown exemplarily) is mainly radiated in the rotor plane. For the zero inflow case, minor BPF2 lobes can be observed around the rotational axis while radiation in the rotor plane is approximately equal to the 20 m/s case. This phenomenon can be explained by higher blade

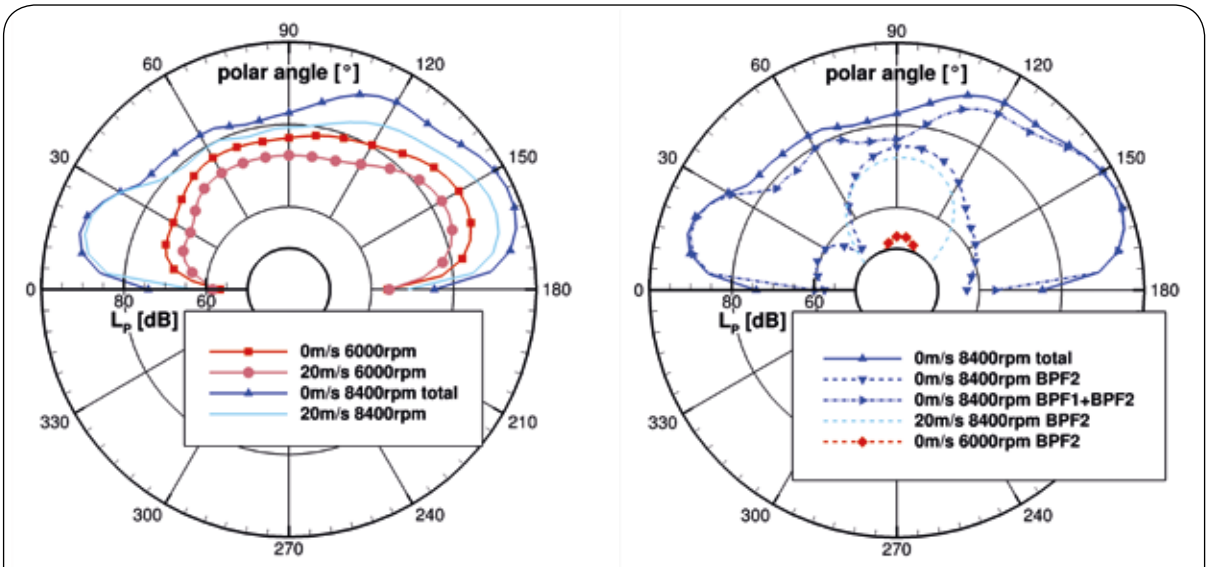

Fig. 6: Polar noise directivity at a distance of 10 m from the CROR's center.
a) Overall sound pressure level b) Selected tones

loading causing more loading noise. As stated, this loading noise is radiated normally to the blade surface (in upstream and downstream directions), while thickness noise remains constant. Single rotor noise gains importance for higher rotational speed but does not contribute significantly to the overall sound pressure level at 6,000 rpm. The observations agree well with expectations and results from the literature and the simulations are therefore believed to form a reliable basis for further investigations on integrated CRORs.

Written by Johannes Letzgus, Lukas Dürrwächter, Ulrich Schäferlein, Ulrich Schäferlein, Manuel Keßler, Ewald Krämer,

University of Stuttgart, Institute of Aerodynamics and Gas Dynamics (IAG)

letzgus@iag.uni-stuttgart.de

References

[1] Raddatz, Jochen and Fassbender, Jens K:
Block Structured Navier-Stokes Solver FLOWer in: MEGAFLOW-Numerical Flow Simulation for Aircraft Design, 2005, pp.27 – 44, Springer.

[2] Weihing, P. and Letzgus,J. and Bangga,G. and Lutz,T. and Krämer, E.:
Hybrid RANS/LES Capabilities of the Flow Solver FLOWer - Application to Flow around Wind Turbines in: 6th Symposium on Hybrid RANS-LES Methods, Strasbourg, 2016.

[3] Kranzinger, Patrick P and Kowarsch, Ulrich and Schuff, Matthias and Keßler, Manuel and Krämer, Ewald:
Advances in parallelization and high-fidelity simulation of helicopter phenomena, in: High Performance Computing in Science and Engineering 15, 2016, pp. 479 – 494, Springer.

[4] Karypis, George and Kumar, Vipin:
METIS - unstructured graph partitioning and sparse matrix ordering system, version 2.0, 1995.

[5] Keßler, M. and Wagner, S.:
Source-Time Dominant Aeroacoustics, in: Computers & Fluids Vol. 33, 2004, pp. 791 – 800.

Automated Code Generation for Maximizing Performance of Detailed Chemistry Calculations in OpenFOAM

Numerical simulations have proven to be an effective tool to aid the development and optimization of modern combustion devices. Due to the complex interplay of combustion chemistry and turbulent flow, combustion processes present a challenging task for simulations. In the past, many combustion models have been developed which approximate the underlying chemical reactions and turbulent flow. The most accurate insights into the flames, however, are gained from simulating the flame and flow directly without the use of modelling assumptions. These direct numerical simulations (DNS) require resolving the flames' smallest structures and capturing the full combustion chemistry, which includes a wide variety of chemical reactions involving multiple chemical species, and billions of computational cells for highly turbulent flows. Because of this, DNS of turbulent flames require enormous computational resources and are only possible with the help of supercomputers. In this work, an optimization strategy for highly efficient computation of complex chemical reaction rates for DNS of turbulent combustion has been implemented in the widespread open-source code OpenFOAM. The developed code is used to simulate a turbulent burner of laboratory scale on 28,800 CPU cores on the Hazel Hen supercomputer at HLRS. The optimized solver is approximately 50 percent faster than its old implementation. The simulation results show a quantitatively good agreement with measured data, validating the accuracy and efficiency of the new implementation.

Code Development

The DNS code [1] for the detailed simulation of combustion is developed at the Engler-Bunte-Institute (EBI) and Steinbuch Centre for Computing (SCC) at the Karlsruhe Institute of Technology. It uses the capabilities of the open-source computational fluid dynamics library OpenFOAM [2] in order to simulate the flow of chemically reacting gases and is able to handle complex geometries via unstructured grids. The solver code is also coupled to the open-source thermo-chemical library Cantera [3] for computing detailed thermodynamical and molecular transport properties. Since its development, the code has been used for the direct numerical simulation of different type of flames and has been validated with experimental data [4, 5].

As mentioned, detailed simulations of turbulent flames require large computational resources. Because combustion processes are governed by a large number of chemical species that interact through many different chemical reactions, one of the simulations' performance bottlenecks is computing chemical reaction rates. In order to overcome this challenge, an optimization technique has been developed that–without a loss of accuracy–can reduce the time to compute the reaction rates by 50 percent [6] compared to the previous solver version, where Cantera was used for computing the reaction rates. Usually, before a flame is simulated, a set of chemical species and chemical reactions are selected for describing the combustion process. For example, the well-known GRI 3.0 [7] reaction mechanism describes the combustion

of methane in air with 53 chemical species and 325 elementary reactions. The reaction mechanism is provided as a text file with parameters for all participating species and chemical reactions and is used as input in the simulation. In the new approach [6], a converter tool has been developed which reads a reaction mechanism file and automatically generates source code for computing the reaction rates. This approach has several advantages: instead of using a general implementation, which can use arbitrary reaction mechanisms, the generated code only contains the instructions to compute reaction rates for one specific reaction mechanism and can therefore be highly optimized by the compiler. When the generated code is compiled, it can directly be used in the simulation due to OpenFOAM's runtime selection mechanism.

In the next section, an application of the optimized code to a realistic burner is presented.

Simulation of a Turbulent Flame

In order to enhance stability of combustion processes and reduce pollutant emissions, modern combustion devices tend to use partially premixed combustion concepts, where fuel and oxidizer are neither perfectly mixed ("premixed") nor perfectly separated ("non-premixed"). This presents a challenging task for the numerical modelling, as available combustion models are designed mostly for either premixed or non-premixed flames. To gain an in-depth understanding of the mechanism of flame stabilization and to provide a comprehensive database for the development of advanced modelling concepts, a highly resolved numerical simulation for the Sandia/Sydney burner [8] has been performed.

The simulation has been split into two parts. In the first simulation, the mixing of methane and air is simulated before the flame is ignited. The computational grid consists of 150 million computational cells. Figure 1a shows a 2D cut through the 3D domain of the mixing process of methane (red) and air (blue) in a pipe leading to the burner nozzle exit. When the gas flow reaches the nozzle on the right border of the figure, the mixing process is not completed, resulting in an inhomogeneous distribution of fuel and air. Due to the relatively high gas velocities of about 60 m/s, the flow is highly turbulent. In Fig. 1b, the turbulent flow structures aiding the mixing in the pipe are visualized by the vorticity field.

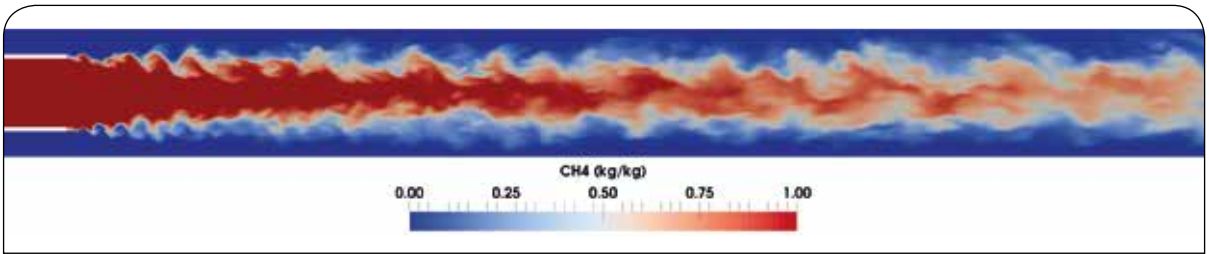

Fig. 1a: Mixing of methane (red) and air (blue) in a pipe leading to the burner nozzle exit.

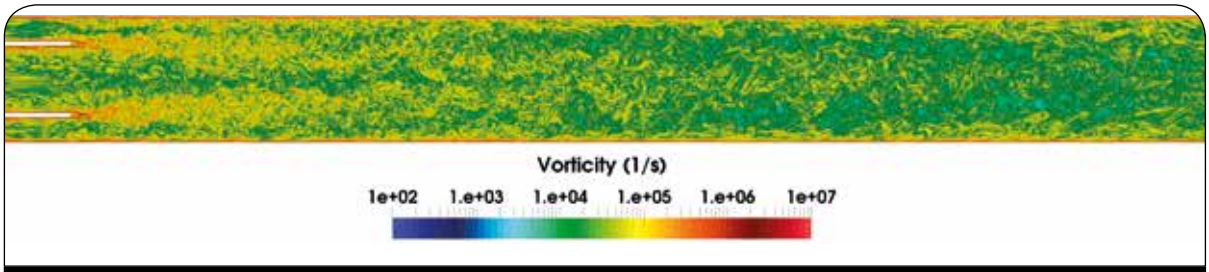

Fig. 1b: Turbulent flow structures visualized by the vorticity field.

The results from the previous simulation of the inhomogeneous mixing of methane and air serve as inflow boundary conditions for the subsequent simulation of the turbulent combustion downstream of the nozzle, which also consists of 150 million cells. In this simulation, the previously partially mixed fuel and oxidizer gas are ignited. Because this simulation considers the full combustion mechanism, it requires more resources than the first. However, due to the aforementioned performance optimization strategy, it was possible to run the simulation on 28,800 CPU cores on the Hazel Hen cluster at the High-Performance Computing Center Stuttgart (HLRS). In total, the simulation required about 15 million core hours and produced 15 TB of data for 3D transient results of flow/chemical variables, which will be analysed in future work. Figure 2a shows a 2D cutting plane of the temperature field through the flame. The inner region on the left enclosed by the high temperature region indicates the cold, unburnt methane-air mixture from the previous simulation, which ignites by mixing with hot gas from a pilot flow. The flame surface becomes increasingly wrinkled due to the interaction of the flame with large-scale vortices developed in the shear layer. The simulation results are for example used to analyse the regions of high heat production, as depicted in Fig. 2b by the heat release rate $\dot{Q}$.

The simulation results show a quantitatively good agreement with the experimental data.

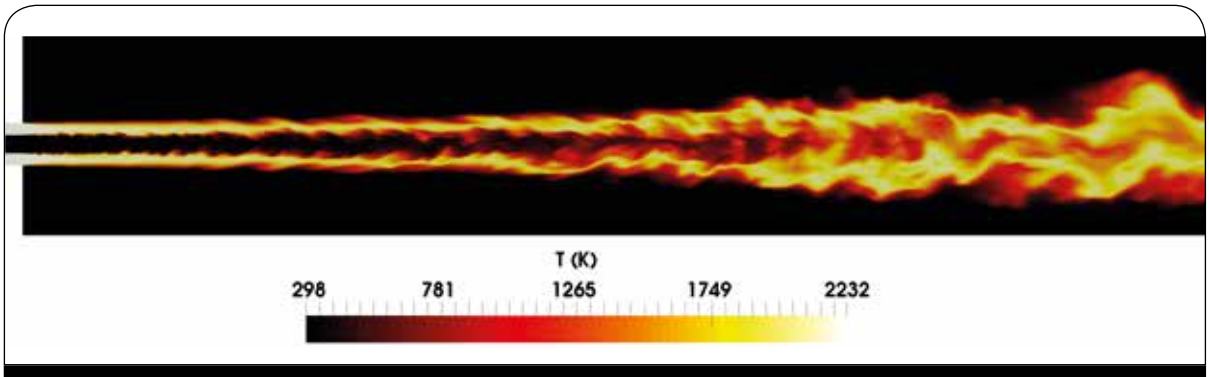

Fig. 2a: The flame visualized by its temperature field.

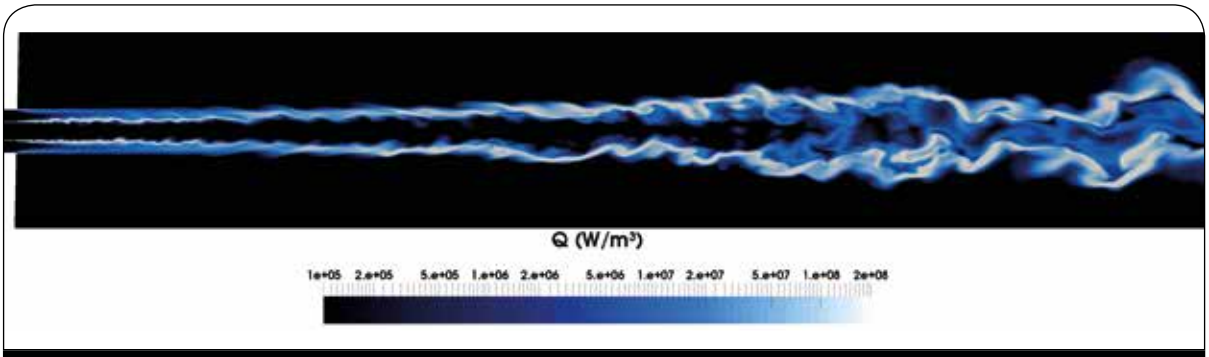

Fig. 2b: Regions with high heat release rates.

In Fig. 3, radial profiles of time averaged (mean) and root mean squared (RMS) values of two solution variables (temperature and mixture fraction Z) are compared at a fixed axial position showing that the simulation results lie well within the experimental uncertainties.

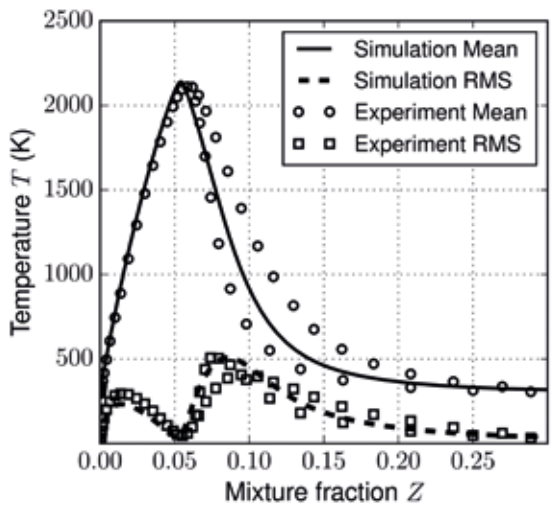

Fig.3: Comparison of mean and RMS values of radial temperature and mixture fraction profiles with experimental data.

Due to its highly resolved nature, the simulation enables studying the flame in great detail, including regions and quantities that are not accessible in experiments. For example, Fig. 4 shows the complex flow structures in the cold inner regions as iso-surfaces of vorticity and how they interact with the flame, depicted as a 2D temperature cutting plane. The simulation results will be compiled into a database and made publicly available. This database will help to understand the fundamental mechanisms of such partially premixed flame concepts and to validate new combustion models. The visualization for this simulation has been done on the ForHLR II cluster at SCC at the Karlsruhe Institute of Technology (KIT).

Acknowledgements

This work was performed on the national supercomputer Cray XC40 Hazel Hen at the High Performance Computing Center Stuttgart (HLRS) and on the computational resource ForHLR II at SCC (KIT) funded by the Ministry of Science, Research and the Arts Baden-Württemberg and DFG (“Deutsche Forschungsgemeinschaft”).

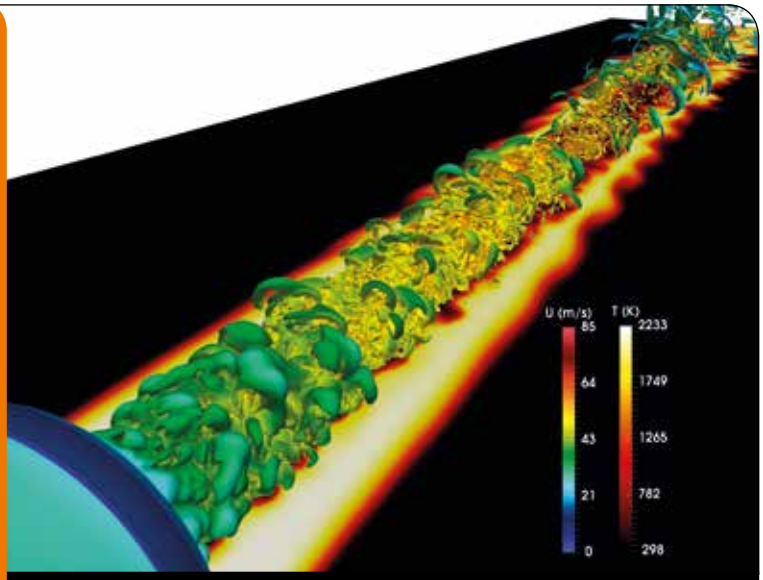

Fig. 4: Turbulent flow structures in the flame visualized by the iso-surface of vorticity, colored by the gas velocity.

References

[1] **F. Zhang, H. Bonart, T. Zirwes, P. Habisreuther, H. Bockhorn, N. Zarzalis:**
Direct Numerical Simulation of Chemically Reacting Flows with the Public Domain Code OpenFOAM. In: Nagel, W.E., Kröner, D.H., Resch, M.M. (eds.) High Performance Computing in Science and Engineering '14, Springer, Berlin Heidelberg (2015), pp. 221–236.

[2] **OpenCFD, OpenFOAM:**
The Open Source CFD Toolbox. User Guide Version 1.4, OpenCFD Limited. Reading UK, Apr. 2007.

[3] **D.G. Goodwin, H.K. Moffat, R.L. Speth. Cantera:**
An Object-oriented Software Toolkit for Chemical Kinetics, Thermodynamics, and Transport Processes. Version 2.3.0b (2016), http://www.cantera.org

[4] **T. Zirwes, F. Zhang, T. Häber, H. Bockhorn:**
Ignition of combustible mixtures by hot particles at varying relative speeds. Combustion Science and Technology (2018), pp. 1-18.

[5] **F. Zhang, T. Zirwes, P. Habisreuther, H. Bockhorn:**
Effect of unsteady stretching on the flame local dynamics. Combustion and Flame 175 (2017), pp. 170-179.

[6] **T. Zirwes, F. Zhang, J. Denev, P. Habisreuther, H. Bockhorn:**
Automated Code Generation for Maximizing Performance of Detailed Chemistry Calculations in OpenFOAM. In: High Performance Computing in Science and Engineering '17, W. Nagel, D. Kröner, and M. Resch, Eds. Springer, 2017, pp. 189–204.

[7] **G. Smith, D. Golden, M. Frenklach, N. Moriarty, B. Eiteneer, M. Goldenberg, C. Bowman, R. Hanson, S. Song, W. Gardiner, V. Lissianski, Z. Qi:**
Gri 3.0 reaction mechanism.

[8] **R. Barlow, S. Meares, G. Magnotti, H. Cutcher, A. Masri:**
Local extinction and near-field structure in piloted turbulent CH4/air jet flames with inhomogeneous inlets. Combustion and Flame (2015), vol. 162, no. 10, pp. 3516–3540.

Written by Thorsten Zirwes[1,2], Feichi Zhang[2], Peter Habisreuther[2], Jordan A. Denev[1], Henning Bockhorn[2]

[1]Steinbuch Centre for Computing, Karlsruhe Institute of Technology

[2]Engler-Bunte-Institute, Combustion Technology, Karlsruhe Institute of Technology

Contact: thorsten.zirwes@kit.edu

Simulating the Emergence of Cosmic Structures

Galaxy formation in a dark universe

Observed galaxies range in mass from a few thousand to a few trillion times the mass of the Sun, encompass physical sizes from a fraction to tens of kilo-parsecs, and span a variety of morphologies. Galaxies can reside in diverse environments—in isolation, or as members of rich groups and clusters. They are self-gravitating systems of stars and gas embedded in a halo of dark matter, and their distribution throughout space traces a cosmic web defined by filaments, nodes, sheets, and voids of matter. The highly clustered, large-scale structure of the Universe today, at mega-parsec and giga-parsec scales, arose from more than 13 billion years of evolution, starting from a nearly homogeneous distribution of matter in the early universe.

In order to gain a theoretical, ab initio understanding of the structure formation process and the role different physical processes play in shaping galaxies, simulations that account for the multi-scale physics involved are the tool of choice. As a starting point, such calculations rely on initial conditions that are known and well constrained by observational data from Cosmic Microwave Background radiation, now measured to exquisite precision by the Planck

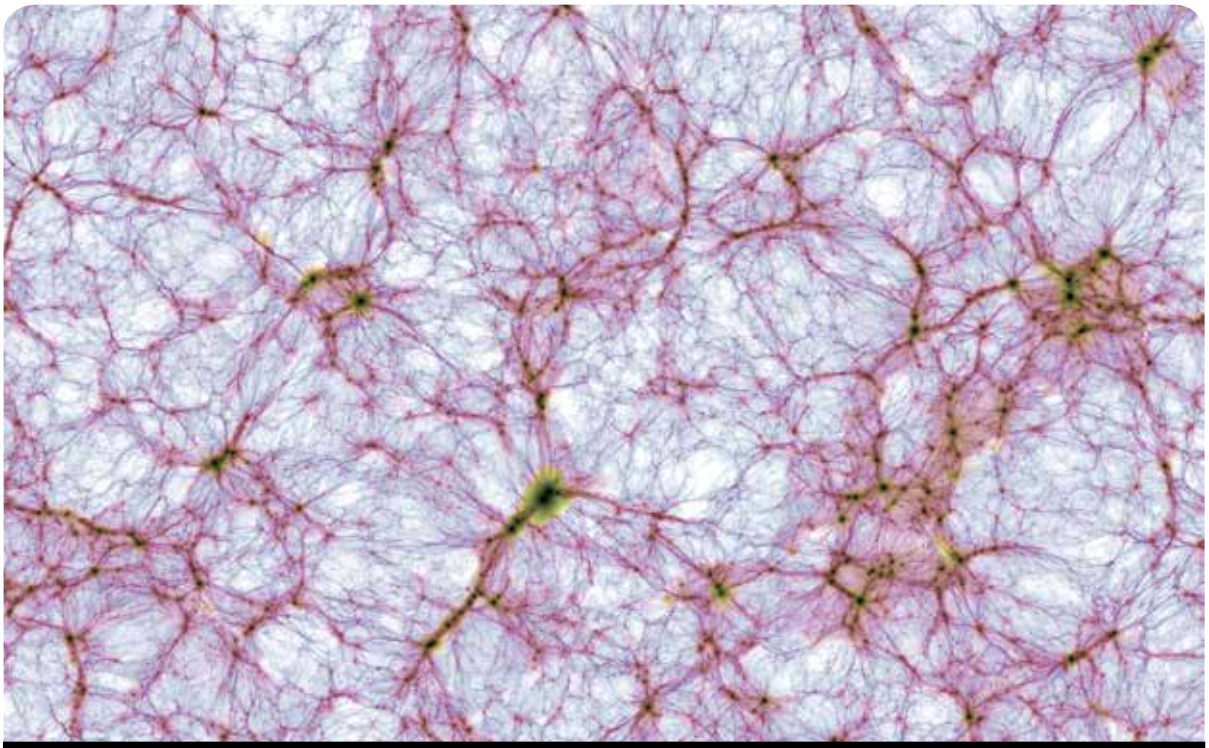

Fig. 1: Thin slice through the cosmic large-scale structure in the largest simulation of the IllustrisTNG project, containing more than 30 billion resolution elements. Shown is a projection of the baryonic density field, where image brightness indicates the projected mass density and color hue visualizes the mean projected gas temperature. The displayed region extends by about 1.2 billion lightyears from left to right.

satellite. The calculations then need to accurately compute the dominant physical force, namely gravity, acting upon all matter, within the accelerating expansion of the Universe. Furthermore, (magneto) hydrodynamical processes for modeling the evolution of the gaseous component of the universe need to be followed. Finally, one must account for all other relevant astrophysical processes: from the atomic level interactions that govern radiative cooling of a metal-enriched gas, to the formation of stars and supermassive black holes, with their subsequent expulsion of mass, metals, and "feedback" energy, which can impact scales even larger than entire galactic halos.

In 2014, the Illustris team—named after its flagship simulation—performed one of the first high-resolution, very large hydrodynamic simulations of galaxy formation. This simulation used a novel and innovative numerical approach, employing a moving, unstructured mesh as realized in the team's AREPO code [1]. One of Illustris' major achievements was its ability to follow the small-scale evolution of gas and stars within a representative portion of the Universe, yielding a population of thousands of well-resolved elliptical and spiral galaxies. For the first time, the simulation reproduced the observed morphological mix of galaxies and its dependence on stellar mass. However, some

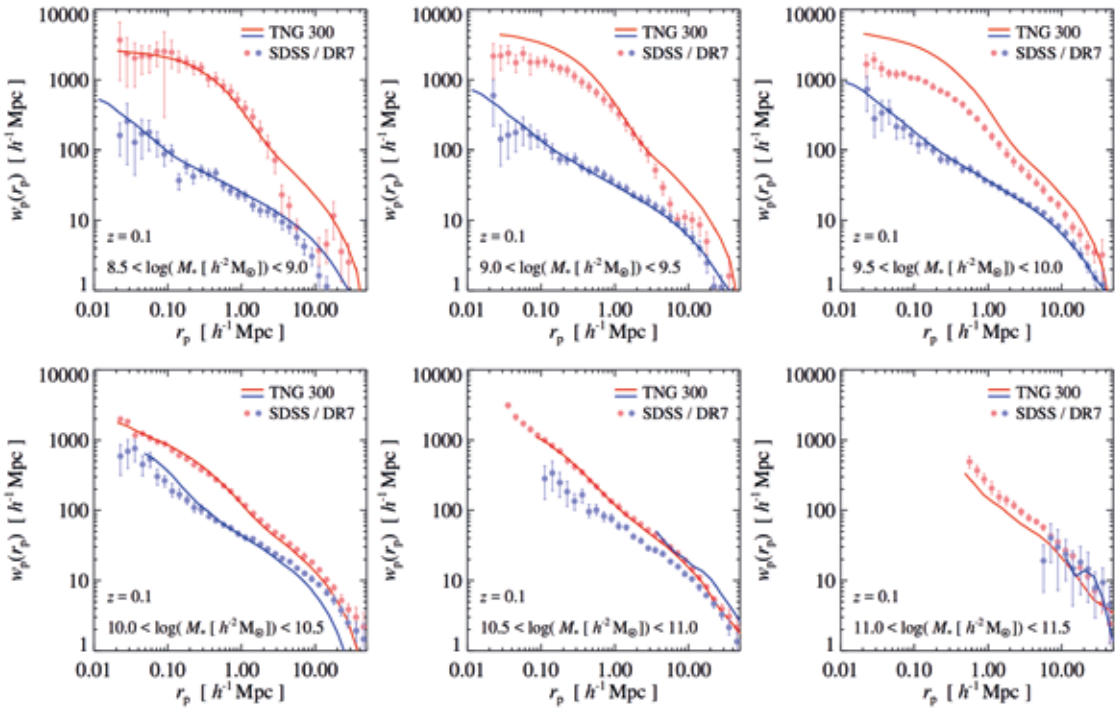

Fig. 2: Comparison of the projected two-point galaxy correlation functions of TNG300 at redshift z = 0.1 to data of the Sloan Digital Sky Survey (symbols), split by g-r color into red and blue galaxies, in six different stellar mass ranges.

aspects of Illustris were found to be in tension with a number of observable constraints. These important discrepancies could be traced to the Illustris feedback model from active galactic nuclei, pointing to deficits in the physics model for supermassive black hole accretion and its associated energy release.

New physics in IllustrisTNG

To address these discrepancies, the Illustris team used the last three years to undertake a campaign of new model development, while also improving upon other lingering problems in the predictions of the original Illustris simulation. This effort led to a new kinetic wind model for black hole feedback physics [2], the incorporation of magnetic field amplification in our code, and new techniques for tracking the production of heavy elements such as magnesium and europium in different types of supernova explosions, among numerous other advances in the implemented physics model of galaxy formation [3].

This substantially improved physical and numerical model has then been applied by the team in "The Next-Generation Illustris" (IllustrisTNG) simulations. This ambitious simulation program actually consists of three complementary cosmological simulations, with the flagship runs TNG300 (consisting of a large volume simulation with a periodic boxsize of 300 Mpc), TNG100 (intermediate volume, boxsize 110 Mpc), and TNG50 (small volume, boxsize 50 Mpc). The calculations were carried out on the Hazel Hen supercomputer at the High-Performance Computing Center Stuttgart, with CPU-time awarded by the Gauss Centre for Supercomputing.

The highest resolution implementation of the TNG300 simulation contains 2 × 25003 total resolution elements (31.25 billion), half in the form of dark matter particles and half as Voronoi gas cells. In addition, 25003 Monte Carlo tracer particles for tracking gas flows were included. The mean baryonic gas mass resolution is 1.1 × 107 M⊙, while the DM particle mass is 5.9 × 107 M⊙. The Plummer equivalent gravitational softening of the collisionless component (dark matter and stars) is 2.95 co-moving kilo-parsecs until redshift z = 1, after which it is fixed to its physical value of 1.48 kilo-parsecs at that epoch. The gravitational softening of gas cells is adaptive, with a minimum of about 370 co-moving parsecs. To assess resolution convergence, the TNG300 simulation (as well as the other flagship runs) has been supplemented by other lower resolution simulations, spaced by factors of 8 in mass resolution. To isolate the effects of the galaxy formation model onto the underlying dark matter distribution, each full baryonic physics run has also been repeated with the same initial condition phases, but including only dark matter and gravity.

Large-scale structure results

Figure 1 gives an overview of the cosmic large-scale structure as seen in the TNG300 simulation. Interestingly, hydrodynamical simulations of galaxy formation have now reached sufficient volume to make precision predictions for clustering on cosmologically relevant scales. For

example, we have used our new IllustrisTNG simulations to study the non-linear correlation functions and power spectra of baryons, dark matter, galaxies and haloes over an exceptionally large range of scales [4]. We found that baryonic effects increase the clustering of dark matter on small scales, and damp the total matter power spectrum on scales up to 3 Mpc. We have also found that the two-point correlation function of the simulated galaxies agrees well with observational data from the Sloan Digital Sky Survey (see Figure 2), both as a function of stellar mass and when split according to galaxy color, a remarkable result that for the first time has been achieved with cosmological hydrodynamical simulations.

Given this agreement, the TNG simulations can make valuable theoretical predictions for the so-called clustering bias of different galaxy samples, which has important cosmological implications. For example, we have detected significant scale dependencies in the bias of different observational tracers of large-scale structure, extending well into the range of the baryonic acoustic oscillations that can be used to reconstruct the expansion history of the universe and thus constrain dark energy. The resulting nominal shifts of the acoustic peaks can be around 5 percent in spatial scale, corresponding to a substantial mistuning of this cosmic ruler. But fortunately, it appears that this is correctable with sufficient amounts of galaxy survey data, thanks to the smooth variation of the bias with spatial scale.

Magnetic fields and metal distribution in clusters of galaxies

On somewhat smaller scales, IllustrisTNG makes interesting novel predictions as well. As an example, the team can consider galaxy clusters, large assemblies of more than a thousand galaxies bound together by a deep gravitational potential well due to a common background dark matter halo and a hot atmosphere of intracluster gas. The largest of these clusters weigh in excess of 1,015 solar masses, yet less than 2 percent of this mass are actually bound in stars. The plasma of the intracluster gas contains most of the cluster baryons, and through its X-ray emission, one can study the thermodynamic properties and the metal enrichment of this component.

The IllustrisTNG simulations make detailed predictions for how galaxies synthesize heavy elements in stars and blow them out of galaxies through supernova-driven winds, stirring and enriching the intra-cluster gas [5]. As an example, Figure 3 shows maps of the metal abundance in the most massive cluster of TNG300 at redshift z=0. Maps of the element abundances are given for iron (upper left), iron over all metals (upper right), silicon over iron (lower left), and oxygen over iron (lower right). The circles in the maps show a radius enclosing a mean overdensity of 500 times the critical density of the universe, a common characteristic radius used in X-ray astronomy. The extent and depth of each map is three times this radius. Note that the element ratios are spatially not uniform, but show small-scale variations. The details of those

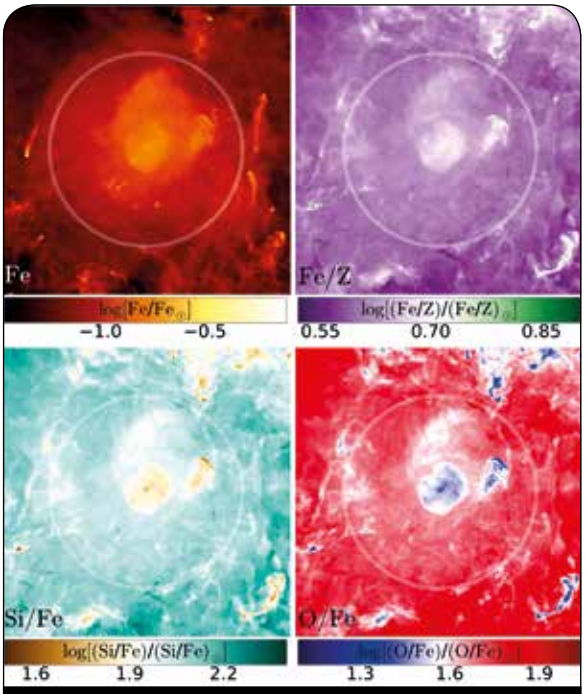

Fig. 3: Maps of the metal abundance and relative abundance ratios in the most massive cluster of galaxies in TNG300 at the present epoch.

patterns are caused by the different enrichment channels of the various elements. Interestingly, some detailed features of the ratio distributions are quite pronounced, but less visible in the iron distribution.

The analysis [5] has shown that the predicted radial metallicity profiles agree well with observations, and the enrichment history is consistent with observational data going beyond z ~ 1, showing nearly no metallicity evolution. In fact, we could demonstrate that more than 80 percent of the metals in the intra-cluster medium have been accreted from the proto-cluster environment, which has been enriched to about 0.1 solar already at z ~ 2. We conclude that the intra-cluster metal distribution is uniform among our simulated cluster sample, nearly time-invariant in the outskirts for more than 10 Gyr, and forms through a universal enrichment history.

Another interesting aspect of the recent work deals with the team using IllustrisTNG to track magnetic fields for the first time. Magnetic fields are prevalent in galaxies, and their build-up is closely linked to structure formation. We find that structure formation amplifies the tiny initial seed fields to the values observed in low-redshift galaxies (1–10 μG). The magnetic field topology is closely connected to galaxy morphology such that irregular fields are hosted by early-type galaxies, while large-scale, ordered fields are present in disc galaxies. Using a simple model for the energy distribution of relativistic electrons, the team can also predict the diffuse radio emission of galaxy clusters [6], and make observational forecasts for forthcoming radio telescopes such as the Square Kilometer Array (SKA).

An example is given in Figure 4 that shows an X-ray map (color coded) overlayed with synchrotron emission contours for the most massive halo of the TNG300 simulation taken from the z = 0.2 snapshot. The panel is 3.5 Mpc on a side and in projection depth, whereas the inset zooms onto the central region and is 800 kpc across. Note how the X-ray and radio morphologies match one another, indicative of the fact that the gas and large-scale magnetic field morphologies are similar. However,

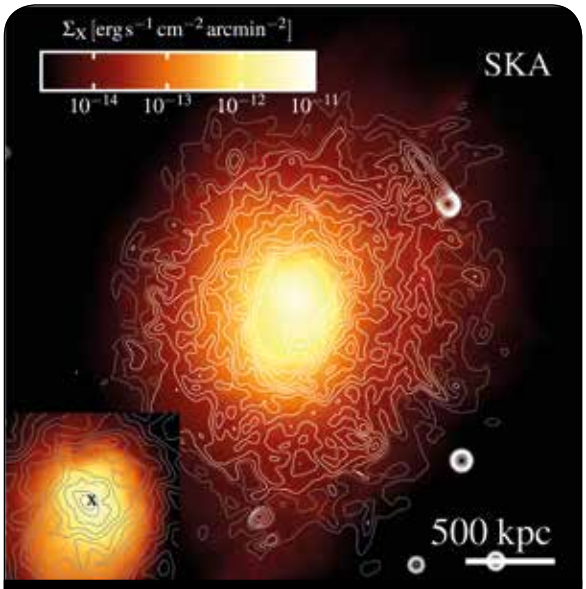

Fig. 4: X-ray emission map overlayed with radio synchrotron emission contours for a massive cluster halo at redshift z = 0.2. The radio maps are smoothed with a Gaussian kernel representative of the typical beam size and sensitivity of future SKA observations of radio haloes in galaxy clusters.

at early times such as this, there are also compact regions of radio emission without any significant X-ray counterpart originating from substructures. In general, these simulated clusters show extended radio emission, whose detectability correlates with their virial mass. We also reproduce the observed scaling relations between total radio power and X-ray emission, cluster mass, and the Sunyaev-Zel'dovich decrement.

Outlook

The results discussed here only scratch the surface of the rich science that can be addressed with the IllustrisTNG simulations. Indeed, already 16 scientific articles have been published by the Illustris team and collaborators on different aspects of the calculations, and much additional work is currently under way (see http://www.tng-project.org/ for an overview). The TNG simulations produced more than 500 terabytes of simulation data, forming an extraordinarily rich resource for astrophysical research. Its full analysis will keep us and all participating scientists busy for many years to come, and promises to deliver many exciting results about diverse astrophysical processes. The team also plans to make all the data publicly available by the end of 2018, which we expect to further strengthen the utility of the project for the international astronomical research community.

In our view, IllustrisTNG is therefore a good example for the enormous scientific possibilities of modern supercomputers. Indeed, computational astrophysics has become an indispensable part of astronomy, allowing complex theories to be tested against observational data. In the future, we want to concentrate on reducing the number of assumptions needed in calculations such as IllustrisTNG, which still need to invoke sweeping simplifications and are severely limited in the dynamic range available for addressing the multi-scale physics determining cosmic structure formation. It is clear that computer simulations together with observational progress hold the key to eventually understand galaxy formation in full, which at the moment is still a distant goal.

Acknowledgements

We gratefully acknowledge support by the High-Performance Computing Center Stuttgart,

and computer time through project GCS-ILLU on Hazel Hen. We also acknowledge financial support through subproject EXAMAG of the Priority Programme 1648 SPPEXA of the German Science Foundation, and through the European Research Council through ERC-StG grant EXAGAL-308037, and we would like to thank the Klaus Tschira Foundation.

References

[1] **Springel V., E pur si muove:**
Galilean-invariant cosmological hydrodynamical simulations on a moving mesh, 2010, Monthly Notices of the Royal Astronomical Society, 401, 791

[2] **Weinberger R. Springel V., Hernquist L., Pillepich A., Marinacci F., Pakmor R., Nelson D., Genel S., Vogelsberger M., Naiman J., Torrey P., 2017:**
Simulating galaxy formation with black hole driven thermal and kinetic feedback, Monthly Notices of the Royal Astronomical Society, 465, 3291

[3] **Pillepich A., Springel V., Nelson D., Genel S., Naiman J., Pakmor R., Hernquist L., Torrey P., Vogelsberger M., Weinberger R., Marinacci F., 2018:**
Simulating Galaxy Formation with the IllustrisTNG Model, Monthly Notices of the Royal Astronomical Society, 473, 4077

[4] **Springel V., Pakmor R., Pillepich A., Weinberger R., Nelson D., Hernquist L., Vogelsberger M., Genel S., Torrey P., Marinacci F., Naiman J., 2018:**
First results from the IllustrisTNG simulations: matter and galaxy clustering, Monthly Notices of the Royal Astronomical Society, 475, 676

[5] **Vogelsberger M., Marinacci F., Torrey P., Genel S., Springel V., Weinberger R., Pakmor R., Hernquist L., Naiman J., Pillepich A., Nelson D., 2018:**
The uniformity and time-invariance of the intra-cluster metal distribution in galaxy clusters from the IllustrisTNG simulations, Monthly Notices of the Royal Astronomical Society, 474, 2073

[6] **Marinacci F., Vogelsberger M., Pakmor R., Torrey P., Springel V., Hernquist L., Nelson D., Weinberger R., Pillepich A., Naiman J., Genel S., 2018:**
First results from the IllustrisTNG simulations: radio haloes and magnetic fields, Monthly Notices of the Royal Astronomical Society, submitted, arXiv:1707.03396

Written by Volker Springel, Rainer Weinberger, Rüdiger Pakmor
Heidelberg Institute for Theoretical Studies and Heidelberg University, Germany

Annalisa Pillepich
Max-Planck Institute for Astronomy, Heidelberg, Germany

Dylan Nelson
Max-Planck-Institute for Astrophysics, Garching, Germany

Mark Vogelsberger, Federico Marinacci, Paul Torrey
Massachusetts Institute of Technology, Cambridge, USA

Lars Hernquist, Jill Naiman
Harvard University, Cambridge, USA

Shy Genel
Center for Computational Astrophysics, Flatiron Institute, New York, USA

Contact: sgenel@flatironinstitute.org

Data Assimilation with the Integrated Terrestrial Systems Modeling Platform, TerrSysMP-PDAF

Data assimilation is a statistical methodology to correct numerical model predictions by measurements. The calculated correction is derived from an optimal weighting of simulated model states on one hand and measurements on the other. The Centre for High-Performance Scientific Computing in Terrestrial Systems (HPSC-TerrSys) coupled the Parallel Data Assimilation Framework (PDAF) to the Terrestrial Systems Modelling Platform (TerrSysMP), which can model the terrestrial system from the bedrock to the lower stratosphere in an integrated fashion and forecasts the full state of the terrestrial system.

Method

The Terrestrial Systems Modelling Platform (TerrSysMP) (Shrestha et al., 2014) is a model that simulates terrestrial system states and fluxes from the bedrock to the lower stratosphere, including aquifers, soils, vegetation, surface water and weather in an integrated fashion. It is one of the few available model systems in the world that represents the water and energy cycles in a fully coupled fashion from groundwater-to-atmosphere (G2A). TerrSysMP consists of three component models: (i) COSMO, an atmospheric circulation model developed by the German Weather Service (Baldauf et al., 2011); (ii) the Community Land Model (CLM), a land surface model developed by the National Center for Atmospheric Research (NCAR) (Oleson et al., 2013), and (iii) the subsurface model ParFlow (Kollet and Maxwell, 2006).

COSMO is based on the primitive thermo-hydrodynamical equations describing compressible flow in a moist atmosphere. As a limited-area model, COSMO requires lateral boundary conditions from a driving model by nesting COSMO into a global scale atmospheric model. ParFlow solves the three-dimensional Richard's-Equation (a non-linear parabolic partial differential equation, or PDE) for variably saturated porous media flow in the mixed form using a finite difference approach in space and a backward Euler scheme in time, which is used for overland flow coupling as well. The overland flow module (a hyperbolic PDE) uses kinematic wave approximation of the shallow water equation and is solved simultaneously with the subsurface equations using the Newton-Krylov method. The three component models are coupled via OASIS3-MCT (Ocean-Atmosphere-Sea-Ice-Soil coupler–Model Coupling Toolkit) (Valcke et al., 2013). Fig. 1 schematically illustrates the model set-up.

TerrSysMP shows excellent parallel scalability on the Jülich supercomputing infrastructure, allowing for model runs at a high spatial resolution over large spatial domains. The Centre for High-Performance Scientific Computing in Terrestrial Systems (HPSC-TerrSys) addresses a series of research topics applying TerrSysMP related to (i) continental-scale hydrological predictions including groundwater, soil moisture, evapotranspiration and discharge; (ii) near real-time prediction of states of agricultural systems including near real-time control of agricultural operations in the context of sustainable bio-economy; (iii) the role of groundwater on land surface and atmospheric processes

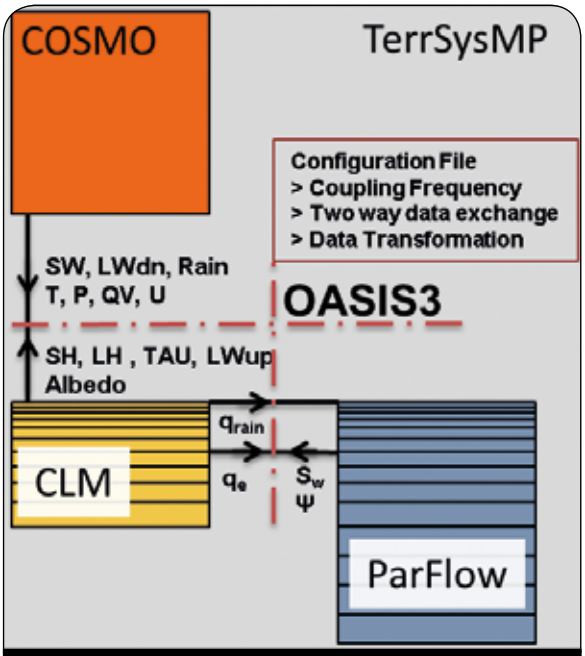

Fig. 1: Schematic of TerrSysMP and the coupling between the three component models. COSMO provides input to CLM in the form of incoming shortwave and longwave radiation, precipitation, air temperature, pressure, humidity and wind speed. CLM provides input to COSMO by outgoing shortwave and longwave radiation, sensible and latent heat flux and land surface temperature. CLM provides input to ParFlow in the form of net precipitation and evapotranspiration and ParFlow calculates saturation and matric potential which are used by CLM to calculate the land surface water and energy balance.

including weather prediction; (iv) human water use (e.g., groundwater pumping, irrigation) and its impact on the hydrological cycle and climate; (v) high-resolution, convection-permitting regional climate modelling.

Predictions with TerrSysMP are affected by different sources of uncertainty, like uncertain initial conditions, uncertain parameters (especially for soil and vegetation properties) and also uncertainty with respect to small-scale processes which are strongly parameterized (for example cloud microphysics and vegetation functioning). Short and medium-term predictions with TerrSysMP can be improved with data assimilation combining uncertain predictions and observations to provide the best possible characterization of terrestrial system states and fluxes. Long-term predictions can also be improved with this approach, if data assimilation is combined with parameter estimation. A widely used approach to assimilate data in numerical models is using ensemble-based Kalman filter methods requiring a large number of model runs. The model runs have, among other variations, different initial conditions, and possibly also different model parameters and boundary conditions. These differences result in an ensemble of different simulations results. The simulation spread should be such that the uncertainty is adequately covered.

HPSC-TerrSys has therefore coupled the Parallel Data Assimilation Framework (Nerger and Hiller, 2013) to TerrSysMP, and the framework TerrSysMP-PDAF shows also an excellent scalability for the use of up to 50,000 processors (Kurtz et al., 2016). This is one of the few data assimilation frameworks available world-wide for integrated terrestrial system modeling.

Regional soil moisture network

One of the research questions addressed recently concerns the value and assimilation of neutron count intensity measurements by cosmic ray probes. Neutron count intensity shows an inverse relation to soil moisture content

and an interesting property of these measurements is that the diameter of the measurement footprint is roughly 600m— much larger than the typical in situ soil moisture measurement (Zreda et al., 2012; Köhli et al., 2015). The footprint is much smaller than most of the satellite-based soil moisture measurements, but satellite measurements have the disadvantage that they are limited to the very few top centimeters of soil, have poor accuracy over densely vegetated areas, and show larger temporal gaps. A number of numerical experiments was carried out and soil moisture data from eight cosmic-ray stations were assimilated into the TerrSysMP-model of the Rur catchment on a daily basis for the period April-December 2013. Cross-validation experiments were performed in which each of the stations was left out from the assimilation for one particular simulation. This procedure was used to interrogate the impacts of the assimilation on the catchment-wide characterization of soil moisture. It was found that the root mean square error (RMSE) between measured and simulated soil moisture content (at verification locations not included in the assimilation) decreased on average over the experiments from roughly 0.09 cm3/cm3 to around 0.05 cm3/cm3. This is a significant error reduction, considering that the density of the cosmic ray probe network is less than 1 per 200 km2, and illustrates the potential of cosmic ray probe networks to improve the characterization of the land surface state. Fig. 2 illustrates that the regional distribution of soil moisture content is strongly affected by data assimilation.

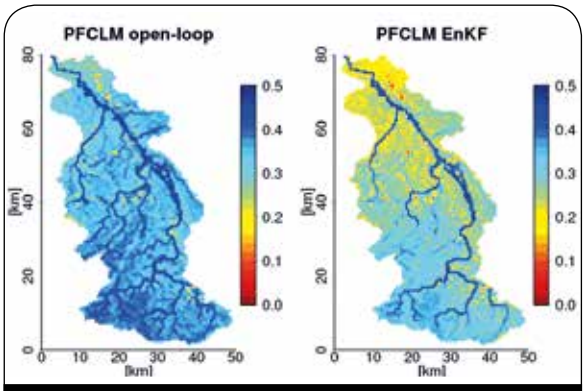

Fig.2.: Simulated soil moisture content in 2013 without data assimilation (left) and with the data assimilation of soil moisture data from a network of cosmic ray probes (right)

Groundwater level data

Another research question HPSC-TerrSys addressed is the value of groundwater level data to improve the characterization of root zone soil moisture content. Root zone soil moisture content is a key variable in terrestrial models, because this is the available water for plant growth and transpiration, and exerts strong control on water fluxes from the land to the atmosphere. Remotely sensed soil moisture only provides limited information on root zone soil moisture content. Groundwater level data could potentially provide additional important information about root zone soil moisture content and large measurement networks are world-wide available (but the databases difficult to access). After developing a methodology to assimilate groundwater level data in TerrSysMP and updating soil moisture content with groundwater level data also under very dry (and therefore very non-linear) conditions, HPSC-TerrSys investigated under which conditions groundwater level data can improve root zone soil moisture

characterization. Simulations were performed for four different soil types combined with five different plant functional types and five different climates. Figure 3 summarizes results from these experiments. In the boxplots, red and green lines smaller than 1.0, indicate that groundwater level data helped to estimate root zone soil moisture content. Groundwater level data provide more information if the groundwater level is between 1m and 5m below the land surface, and for deeper rooting vegetation and loam soils.

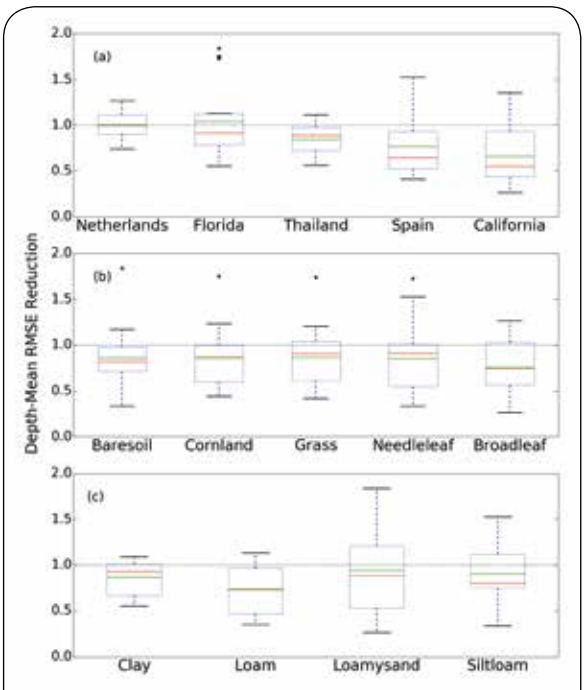

Fig. 3: Boxplots of the root zone averaged ratio of RMSE for simulation experiments where both surface soil moisture and groundwater levels are assimilated and simulation experiments where only surface soil moisture data are assimilated. Values below 1.0 (black dot line) indicate that groundwater level assimilation improved root zone soil moisture characterization. Plotted are results for all 100 cases in terms of climate conditions (a), plant functional types (b) and soil texture types (c). In the boxplots the red line represents median value and the green line mean value.

Continental scale assimilation

More recently, data assimilation with TerrSysMP has also been performed at the continental scale. The soil moisture product CCI from the European Space Agency, a composite constructed from different satellite products (e.g., Dorigo et al., 2017), was assimilated for the period 2000–2006 in CLM, using the TerrSysMP-PDAF framework. The high-resolution reanalysis COSMO-REA6 from the Hans-Ertel Centre for Weather Research (HErZ) was used as model forcing (Simmer et al., 2016). The CCI soil moisture was only assimilated for 100 spatial locations scattered over Europe, and soil moisture was updated at a 3 by 3 km grid over the EURO-CORDEX domain. The assimilation results were evaluated by CCI soil moisture data which were not assimilated and a gridded runoff product (Gudmundsson and Seneviratne, 2016). Results indicate that soil moisture characterization improves especially for the summer and autumn seasons with a bias reduction of 60% compared to simulations without data assimilation. Figure 4 illustrates these results. The modelled discharge is for runs with data assimilation closer to the gridded runoff product than for the simulations without data assimilation. These results show the potential of combining terrestrial systems modelling and data assimilation for improving hydrological predictions at the continental scale.

The data assimilation capacity of TerrSysMP-PDAF will be extended in the coming years for more data types, and further development is foreseen in order to perform data assimilation

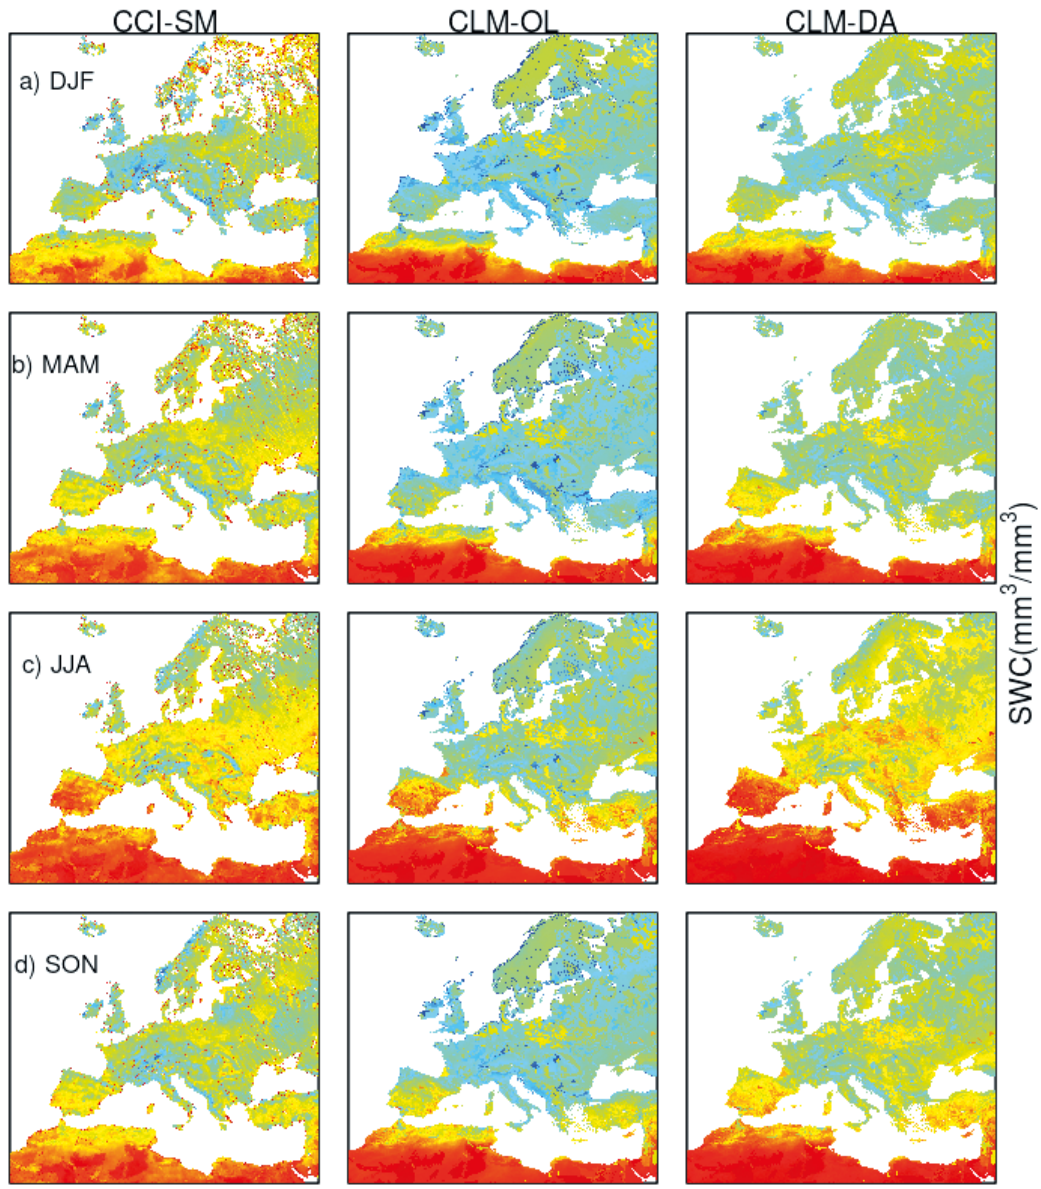

Fig. 4. Soil moisture content (cm3/cm3) averaged for the four different meteorological seasons, according the CCI composite satellite product (left), the simulation runs without data assimilation (middle) and the runs with data assimilation (right).

with the fully coupled model that includes all component models, and applications at the continental scale. The simulations are centered towards answering the main scientific questions listed earlier in this article.

References

[1] **Baldauf, M., Seifert, A., Förstner, J., Majewski, D. , Raschendorfer, M., Reinhardt, T.:**

Operational Convective-Scale Numerical Weather Prediction with the COSMO Model: Description and Sensitivities, Mon. Wea. Rev., 139, 3887–3905, 2011.

[2] **Dorigo, W., Wagner, W., Albergel, C., Albrecht, F., Balsamo, G., Brocca, L., Chung, D., Ertl, M., Forkel, M., and Gruber, A.:**

ESA CCI Soil Moisture for improved Earth system understanding: state-of-the art and future directions, Remote Sens. Environ., doi: 10.1016/j.rse.2017.07.001, 2017.

[3] **Gudmundsson, L., Seneviratne, S. I.:**

Observational gridded runoff estimates for Europe (E-RUN version 1.0), Earth Sys. Sci. Data, 1-1, 2016.

[4] **Köhli, M., Schrön, M., Zreda, M., Schmidt, U., Dietrich, P., Zacharias, S.:**

Footprint characteristics revised for field-scale soil moisture monitoring with cosmic-ray neutrons, Water Resour. Res., 51, 5772–5790, 2015.

[5] **Kollet, S. J., Maxwell, R. M.:**

Integrated surface–groundwater flow modeling: a free-surface overland flow boundary condition in a parallel groundwater flow model, Advances in Water Resources, 29(7), 945-958, 2006.

[6] **Kurtz, W., He, G., Kollet, S. J., Maxwell, R. M., Vereecken, H., Hendricks Franssen, H.J.:**

Terrsysmp-pdaf (version 1.0): a modular high-performance data assimilation framework for an integrated land surface-subsurface model, Geoscientific Model Development Discussions, 8, 9617-9668, 2016.

[7] **Nerger, L., Hiller, W.:**

Software for ensemble-based data assimilation systems—implementation strategies and scalability, Computers & Geosciences, 55(5), 110-118, 2013.

[8] **Oleson, K. W., Niu, G., Yang, Z., Lawrence, D. M., Thornton, P. E., Lawrence, P. J., et al.:**

Improvements to the community land model and their impact on the hydrological cycle, Journal of Geophysical Research Biogeosciences, 113(G1), G01021, 2015.

[9] **Shrestha, P., Sulis, M., Masbou, M., Kollet, S., Simmer, C.:**

A scale-consistent terrestrial systems modeling platform based on COSMO, CLM, and ParFlow, Mon. Weather Rev., 142, 3466-3483, 2014.

[10] **Simmer, C., Adrian, G., Jones, S., Wirth, V., Göber, M., Hohenegger, C., Janjic, T., Keller, J., Ohlwein, C., and Seifert, A.: HErZ:**

The German Hans-Ertel Centre for Weather Research, B. Am. Meteorol. Soc., 97, 1057-1068, 2016.

[11] **Valcke, S.:**

The oasis3 coupler: a European climate modelling community software, Geoscientific Model Development Discussions, 6(2), 373-388, 2013.

[12] **Zreda, M., Shuttleworth, W. J., Xeng, X., Zweck, C., Desilets, D., Franz, T. E., Rosolem, R., Ferre, P.A.:**

COSMOS: The COsmic-ray Soil Moisture Observing System, Hydrology and Earth System Sciences, 16: 4079-4099, 2012.

Written by Harrie-Jan Hendricks Franssen, Wolfgang Kurtz, Mukund Pondkule, Hongjuan Zhang, Bibi Naz, Harry Vereecken, Klaus Goergen, Stefan Kollet.

Jülich Supercomputing Centre (JSC)

Contact: h.hendricks-franssen@fz-juelich.de

Affiliations:

Agrosphere (IBG-3), Forschungszentrum Jülich GmbH

Centre for High Performance Scientific Computing in Terrestrial Systems (HPSC-TerrSys), Geoverbund ABC/J, Jülich, Germany

Rapidly Rotating Convection in Geophysical Systems

Buoyancy driven flows, also called convective flows, are ubiquitous in geophysical and astrophysical objects. Deep inside the Earth, the geomagnetic field is generated by the inductive effect of turbulent convective motions in its outer, liquid iron core. Convection also occurs in giant planets like Jupiter, and may be responsible for the strong zonal winds dominating its visual appearance. Because of the large spatial extent and the rapid rotation of these objects, the flows are typically highly turbulent, but nevertheless they are strongly constrained by Coriolis forces. As this situation is characteristic for many geophysical and astrophysical objects, the goal of this project is to develop a deep physical understanding of turbulent, rotationally constrained convection.

While numerical simulations of the geophysical systems described above have become a common research tool, all current models lack the resolution to represent the key physical processes involved in a realistic fashion. This is only partly due to the usual problems associated with resolving small-scale turbulence. Rapid rotation represents a more fundamental obstacle because it requires extremely small length and time scales that characterize even the primary convective instabilities driving the flow. The common way to deal with this problem is to reduce the strength of rotation in the simulations by many orders of magnitude as compared to realistic values—far enough to guarantee that the leading-order dynamics occur entirely within the limited range of explicitly resolvable scales. This procedure is problematic in a number of ways. Clearly, relating such simulations to natural systems requires certainty that both share a common dynamical regime, such that the simulation results can safely be extrapolated over many orders of magnitude to planetary rotation rates. The scaling laws used in such extrapolations are currently largely empirical and lack a solid theoretical foundation. Finally, the turbulence level realizable in the simulations is fairly limited because increasing the flow amplitude at moderate rotation rates quickly leads to an unrealistic loss of rotational influence on the turbulent eddies.

The problems described above reflect a general lack of physical understanding of rapidly rotating convective systems. The goal of this project is therefore to complement models that are tailored to reproduce geophysical observations as closely as possible with research efforts focusing on the basic physical processes involved. Conceptually simple, plane layer models are studied because (i) they allow for extreme parameter values and (ii) the simple geometry facilitates theoretical advances. Furthermore, (iii) laboratory studies of rotating convection are usually confined to a similar setup.

So far, reliable data that can be used to test theoretical ideas or to guide theoretical work is exceedingly sparse. Experiments on rapidly rotating convective flows have proven to be difficult to conduct. While complex, large-scale facilities are currently built that will hopefully improve the situation in the future, simulations also provide a powerful tool to explore

previously uncharted territory and remain the only way to reveal the complex three-dimensional flow structure. Such simulations also allow for an inclusion of magnetic fields and strong compressibility effects, which are hard to study in the laboratory, but are clearly relevant in natural systems.

Model and Methods

Our simulation code supports a variable background density in order to allow for compression or expansion of fluid particles as they travel between different depth levels. This effect can be strong in geophysical systems, as the mass of the overlying material typically causes a large pressure increase with depth. In Jupiter for example, density is expected to vary by four orders of magnitude from the 1 bar level to the bottom of the outer convection zone. The so-called anelastic approximation is used to account for these effects in our model. Different thermal diffusion processes, ranging from the molecular Fourier law to parameterized turbulent entropy diffusion, are supported. Simulations using the classical Boussinesq approximations can also be performed to facilitate comparisons with laboratory experiments. The code can solve the full magneto-hydrodynamic equations to compute flow-induced magnetic fields and to study their effect on the fluid flow.

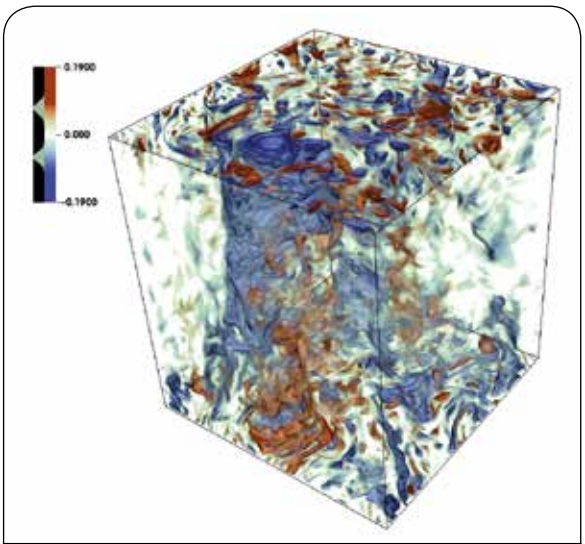

Fig. 1: Temperature anomaly in rapidly rotating convection (Ekman number E=10-7, Rayleigh number Ra=3.2x10[11], Prandtl number Pr=1, resolution 1152x1152x1280, no-slip boundaries). The fluid is heated from below and rapidly rotates about a vertical axis.

Technically, the momentum density and the magnetic field are represented by poloidal and toroidal potentials, while the energy equation is formulated in terms of either entropy or temperature deviations from an adiabatic background. The code supports a mixed pseudo-spectral / high-order finite difference discretization or a fully spectral method based on an integral preconditioned Chebyshev approach, depending on the application scenario. Time stepping is done by an AB/BDF IMEX method with second- or third-order accuracy. Memory requirements are proportional to the number of grid points, with the number of floating point operations per time step exhibiting this proportionality up to a small logarithmic factor. The code shows good scaling on up to 262,144 cores for 3.6 billion grid points on JUQUEEN. Parallel NetCDF is used for I/O.

Results

Computing time provided on JUQUEEN has allowed us to run extensive simulations in

previously unexplored regions of parameter space. In a first series of computations, background density stratifications and magnetic fields were neglected in order to allow detailed comparisons with experiments and existing theory. Figure 1 shows a visualization of the temperature field in one simulation out of a systematic parameter study. The flow is turbulent, yet strongly influenced by Coriolis forces down to small scales.

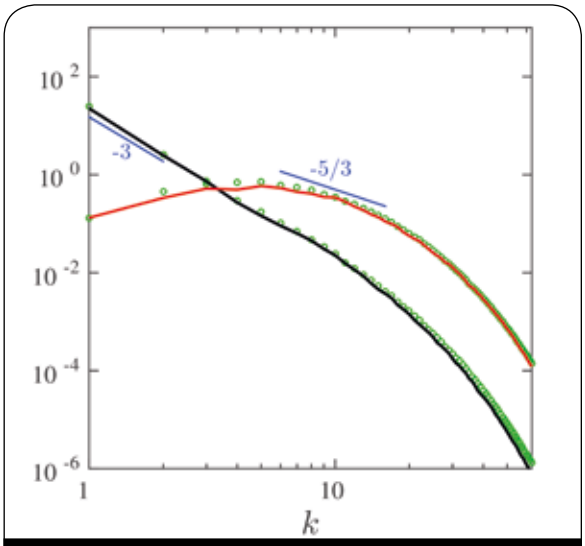

Fig. 2: Barotropic (black) and barocilic (red) horizontal kinetic energy spectra as predicted from asymptotic theory, with k denoting the horizontal wave number. The green dots represent a DNS performed on JUQUEEN. Details are given in [3], from where the figure is reproduced with permission. (Ekman number E=10-7, Rayleigh number Ra=1.94 x 10^{11}, Prandtl number Pr=1).

Simulations like the one illustrated in figure 1 are sufficiently far within the rapidly rotating regime to allow for detailed comparisons with existing theoretical models based on an asymptotic reduction of the governing equations. Our simulations have shown that in contrast to expectations, this theory only holds for very idealized mechanical boundary conditions [1]. For realistic, laboratory-style boundary conditions, no quantitative agreement between the theory and simulations could be found. The problem has been traced to active Ekman boundary layers which, although covering less than one thousand of the layer depth, have a leading-order influence on the flow field [1]. Guided by simulations on JUQUEEN, the theory has been revised [2], and the improved version reproduces the simulation results very well [3,4]. An example is shown in figure 2, where power spectra obtained from direct numerical simulations are compared with results obtained using the refined asymptotic model equations. Excellent agreement is evident in this figure.

A systematic parameter study, covering varying flow amplitudes and fluid properties, has demonstrated that the theory works well over large parameter ranges. The confirmation of the validity of the asymptotic model equations is a key result of our work, as these can be used to investigate the system behavior for parameter values many orders of magnitude beyond the capabilities of DNS [3].

The simulations provide further insight into the nature of turbulence in rapidly rotating convection. In particular, a strong, non-local inverse cascade of barotropic kinetic energy has been found [1,3]. If the boundaries do not support lateral shear stresses, this results in the formation of a coherent, depth-independent dipole condensate [1]. For rigid no-slip boundaries, Ekman friction prevents the formation of long-lived,

coherent large-scale vortices [3]. Instead, the upscale kinetic energy transport results in erratic and intermittent large-scale structures which frequently break up, a process visible in figure 1. The presence of upscale transport in turbulent rotating convection is an important result in the geophysical context, and one of the key predictions to test in future laboratory experiments.

Simulations involving a background density stratifications have also been performed, and we have been able to demonstrate that the interplay of rotation with the compression and expansion of fluid parcels can give rise to the formation of multiple jets and to potential vorticity layering. Furthermore, the scaling of the jet thickness has been shown to be compatible with simple theoretical arguments [5]. The underlying effect may be involved in driving the zonal winds observed on Jupiter and other giant planets. As a by-product of this work, we have also been able to systematically test different modeling approaches for sub-sonic, compressible convective flows frequently used in geo- and astrophysics [6].

Currently, the topology and dynamical influence of a flow-induced magnetic field is studied in detail. Future simulations will also shed light on the important dynamical regime where the rotational influence on the small flow scales is lost, while the large-scale dynamics is still dominated by Coriolis forces. This parameter regime remains largely unexplored to date, but certainly has geophysical relevance.

References

[1] **Stellmach, S., Lischper, M., Julien, K., Vasil,G., Cheng, J., Ribeiro, A., King, E., Aurnou, J.M.:**
Approaching the asymptotic regime of rapidly rotating convection: Boundary layers versus interior dynamics. Physical Review Letters 113 (25), 254501 doi: 10.1103/Phys- RevLett.113.254501 (2014)

[2] **Julien K., Aurnou J., Calkins M., Knobloch E., Marti P., Stellmach S., Vasil G.:**
A nonlinear model for rotationally constrained convection with Ekman pumping. Journal of Fluid Mechanics, 798, pp. 50-87. doi: 10.1017/jfm.2016.225 (2016)

[3] **Plumley M., Julien K., Marti P., Stellmach S.:**
The effects of Ekman pumping on quasi-geostrophic Rayleigh-Bénard convection. Journal of Fluid Mechanics, 803, pp. 51-71. doi: 10.1017/jfm.2016.452 (2016)

[4] **Plumley M., Julien K., Marti P., Stellmach S.:**
Sensitivity of rapidly rotating Rayleigh-Bénard convection to Ekman pumping. Phys. Rev. Fluids 2, 094801, doi: 10.1103/PhysRevFluids.2.094801 (2017)

[5] **Verhoeven J., Stellmach S. The compressional beta effect:**
A source of zonal winds in planets?. Icarus, 237, pp. 143-158. doi: 10.1016/j.icarus.2014.04.019 (2014)

[6] **Verhoeven J., Wiesehöfer T., Stellmach S.:**
Anelastic Versus Fully Compressible Turbulent Rayleigh-Bénard Convectio. Astrophysical Journal, 805 (1). doi: 10.1088/0004-637X/805/1/62 (2015)

Written by Stephan Stellmach
Institut für Geophysik, Westfälische Wilhelms-Universität Münster
Contact: stellma@earth.uni-muenster.de

Ulrich Hansen
Institut für Geophysik, Westfälische Wilhelms-Universität Münster
Contact: hansen@earth.uni-muenster.de

Cavitation Erosion in Injection Systems

Developments in direct diesel injection systems increase rail pressures to more than 2,500 bar. This trend aims at enhancing jet break-up and mixing to improve combustion and thermo-hydrodynamic effects, such as cavitation, which occur when the liquid evaporates locally. The collapse of such vapor structures causes strong shock waves. When bubbles collapse near a solid wall, high-velocity liquid jets directed toward material surfaces are created. Imposed structure loads can lead to material erosion that may be so strong that it leads to severe performance degradation or device failure. On the other hand, these loads are used to clean nozzle holes and throttles from surface deposits, and can promote jet break-up. Furthermore, two-phase flows can be used to maintain choked nozzle conditions and a constant mass flow rate.

Understanding the flow phenomena inside an injection system is necessary to quantify the effects of turbulence and cavitation, and their influence on jet and spray characteristics. Small dimensions, high operating pressures and short timescales make the instrumentation of fuel nozzles with experimental equipment challenging. Computational fluid dynamics (CFD) can provide time-resolved information on flow structures in arbitrary small geometries. Numerical simulations thus have become an important tool in the design process of injection systems.

The present research project focuses on the prediction of cavitation erosion in fuel injection systems using a CFD approach. This includes wave dynamics, interaction of cavitation and turbulence as well as flow transients due to moving geometries. In our project, we use large eddy-simulation (LES) to understand the flow dynamics. Our simulations run on SuperMUC in the massively parallelized numerical framework INCA [1].

Numerical Method

With LES, the smallest turbulent flow scales are not resolved on the computational grid. Effects of these scales thus must be modelled. We employ an implicit LES approach based on the adaptive local deconvolution method (ALDM) [2]. To consider two-phase effects, we apply the homogenous-mixture cavitation model. The actual vapor-liquid interface of cavitation structures is not reconstructed in this barotropic model. Surface tension thus is neglected. Recently, we have extended the single-fluid, two-phase model by a component of non-condensable gas. Complex (moving) bodies are considered by a conservative cut-cell method [3].

Results

Interaction of turbulence and cavitation

We have performed wall-resolved LES of the flow through a generic throttle to validate our model in the context of turbulent, cavitating nozzle flows [4]. At pressure difference 300 to 115 bar, we observe periodic formation of vapor in the detached shear layer at the throttle inlet (left column of Fig. 1). When the pressure difference is increased (300 to 55 bar, right column of Fig. 1), a stable vapor sheet develops at the throttle inlet. In the throttle center, cavitation in large, stable vortices is observed.

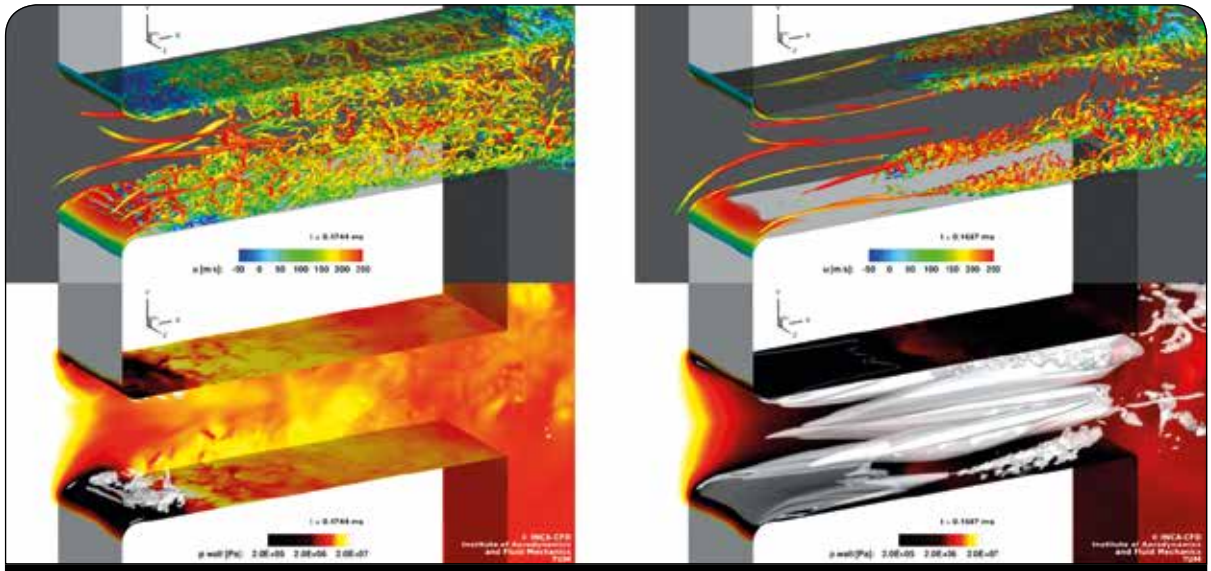

Fig. 1: Instantaneous turbulence (top) and cavitation structures (bottom) in a generic throttle valve. Pressure difference 300 to 115 bar (left) and 300 to 55 bar (right).

We find that while turbulence and vortex dynamics play a dominant role at low pressure differences, the formation of a stable sheet cavity and cavitating vortices suppress turbulence when the differences between pressures is high.

Break-up of cavitating liquid jets

Recently, we have studied the break-up of cavitating liquid jets in a free gas phase [5]. The setup resembles a generic, scaled-up automotive fuel injector and consists of a cavitating water jet emanating from a rectangular nozzle injected into air. We investigated several operating conditions that lead to different cavitation characteristics.

In a supercavitating state, see right column of Fig. 2, we observe a stronger break-up of the liquid jet than in case of developing cavitation, see left column of Fig. 2.

From an analysis of the transient data we have identified three main mechanisms that lead to distortions of the jet surface and, ultimately, to a widening and break-up of the jet. First, turbulent fluctuations, which are induced by collapse events in the proximity of the exit plane of the nozzle, add to the momentum in wall-normal direction.

Second, low pressure vapor regions near the nozzle exit and the gas filled plenum form a pressure gradient, which enables entrainment of gas from the outlet region into the nozzle. When the gas is being ejected again, the water is accelerated toward the side walls and creates large- scale bulges of liquid. Third, collapse events of cavitation structures inside the jet near the liquid-gas interface induce high velocity liquid jets directed towards the interface.

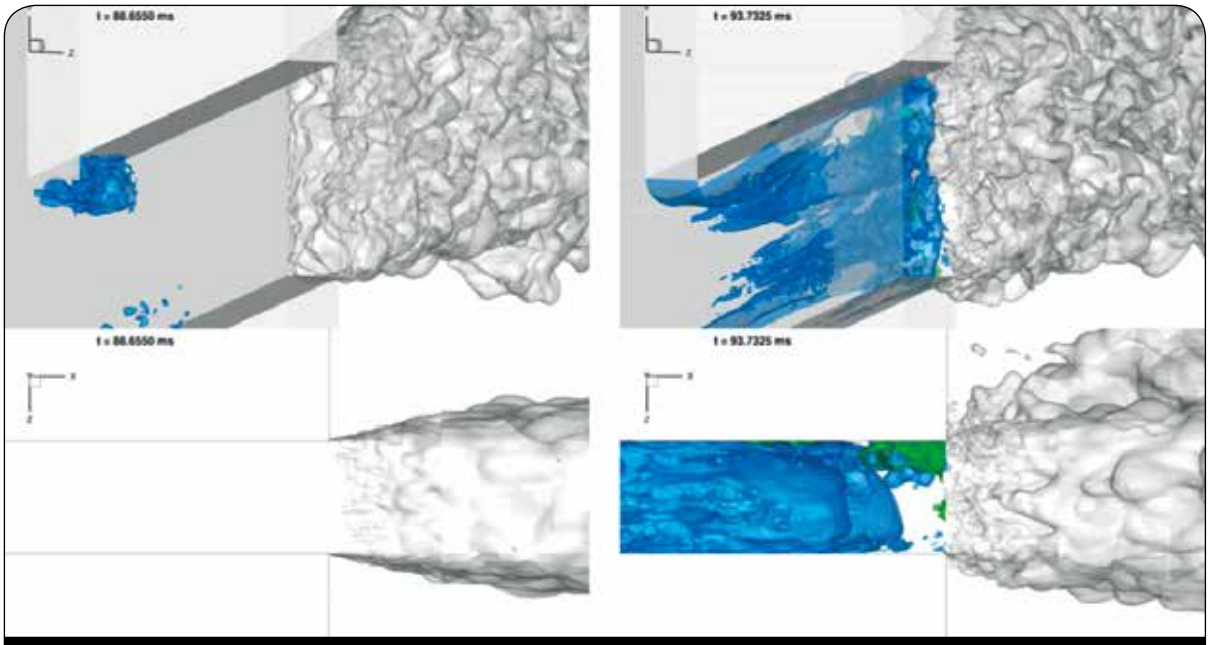

Fig. 2: Cavitation structures (blue) and jet surface (grey) for developing cavitation (left) and supercavitation (right): perspective view (top) and top view (bottom).

9-hole Diesel injector with moving needle

To study our models in realistic environments, we investigate the turbulent multiphase flow inside a 9-hole common rail diesel injector during a full injection cycle of ISO 4113 diesel fuel at a pressure of 1,500 bar into air. Our simulation includes a prescribed needle movement. The nozzle holes have a mean diameter of 150µm.

The analysis of the turbulent flow field reveals that the opening and closing phase are dominated by small-scale turbulence, while in the main injection phase, large vortical structures are formed in the volume upstream of the needle seat, and reach into the nozzle holes, see Fig. 3. In each hole, several of these structures are present at the same time. During and after the closing phase, cavitation structures are detected in the nozzle holes and in the sac hole region, see Fig. 4, and cause violent collapse events. Subsequently, the collapse of the sac hole cavity

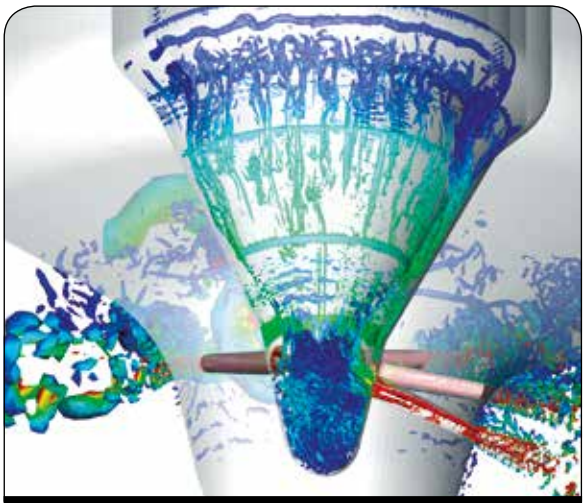

Fig. 3: Coherent vortical structures during the main injection phase colored by velocity in the half-domain.

and rebound effects cause a large number of strong events near the lowest point of the sac hole. These events during this phase are considered to be most likely to cause surface erosion inside the device during operation.

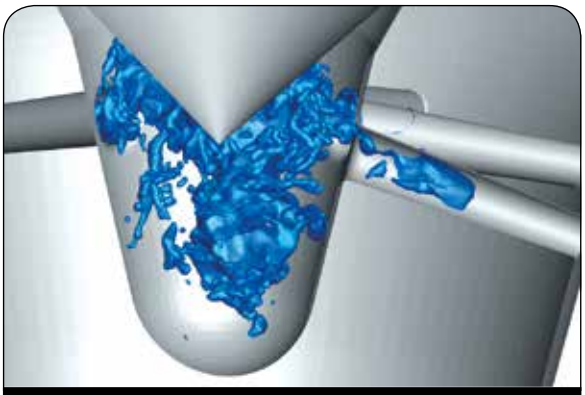

Fig. 4: Cavitation structures inside the nozzle holes and sac hole shortly after full closing of the injector needle.

On-going Research and Outlook

Our studies helped us better understand the dynamics of cavitating, turbulent fluid flows in injection systems. High-performance computing is a necessary tool to address the requirements that the investigations of turbulent, cavitating flows impose on spatial and temporal resolution. Due to the high speed of sound, the time step size usually is on the order of less than a nano-second, while near-wall turbulence requires a high grid resolution and thus causes a large number of cells. Our simulations usually run on up to 5,600 cores to compute physical time scales on the order of micro-seconds. Current research topics include the development of improved gas models with degassing and solution of gas in liquid, the analysis of fluid structure interaction in the context of cavitation, the modeling of quantitative erosion prediction in cavitating flow environments, and acceleration methods for CFD codes for the simulation of cavitating flows.

References and Links

[1] http://www.inca-cfd.com
http://www.aer.mw.tum.de

[2] Hickel, S., Egerer, C. P., & Larsson, J. (2014):
Subgrid-scale modeling for implicit large eddy simulation of compressible flows and shock-turbulence interaction. Physics of Fluids, 26(10), 106101. doi:10.1063/1.4898641

[3] Örley, F., Pasquariello, V., Hickel, S., & Adams, N. A. (2015):
Cut-element based immersed boundary method for moving geometries in compressible liquid flows with cavitation. Journal of Computational Physics, 283(C),1–22. doi:10.1016/j.jcp.2014.11.028

[4] Egerer, C. P., Hickel, S., Schmidt, S. J., Adams, N. A., 2013:
Large-eddy simulation of turbulent cavitating flow in a micro channel. Physics of Fluids, 26(8), 085102. doi:10.1063/1.4891325.2

[5] Örley, F., Trummler, T., Hickel, S., Mihatsch, M. S., Schmidt, S. J., Adams, N. A., (2015):
Large-eddy simulation of cavitating nozzle flow and primary jet break-up. Physics of Fluids, 27(8), 086101. doi:10.1063/1.4928701

Written by Theresa Trummler, Felix Örley, Christian Egerer, Steffen Schmidt

Institute of Aerodynamics and Fluid Mechanics, Prof. N. A. Adams, TUM

Contact: theresa.trummler@aer.mw.tum.de

Secondary Circulations at an Isolated Semi-Arid Forest

Studying the impacts of afforestation in semi-arid regions is important for mitigation of anthropogenic climate change and desertification. However, lowering the albedo by artificially planting trees in semi-arid regions, combined with a large incoming radiation, can lead to a substantial increase in available energy. To guarantee the survival of the ecosystem, this energy load at the surface has to be removed, mainly by turbulent transport to the atmosphere. Due to limited water availability in semi-arid regions, the common pathway of evaporative cooling can often not be applied. However, an alternative cooling mechanism for semi-arid ecosystems was discovered from investigations of the isolated, semi-arid pine forest Yatir [2], located at the northern edge of the Negev Desert in Israel (Fig. 1). This cooling mechanism is mainly realized by an enhanced sensible heat flux above the forest. However, the heterogeneous nature of the isolated Yatir forest can also influence energy transport by generating secondary circulations between the forest and the surrounding shrubland. These circulations affect the incoming circulations and influence the atmospheric boundary layer. They are driven by mechanical effects (roughness of forest canopy) or by buoyancy (due to albedo differences), and can cause effects on weather and climate at the regional scale.

Within the "Climate Feedbacks and Benefits of Semi-Arid Forests (CLIFF)" project [1], we study the influence of these secondary circulations on the surface-atmosphere exchange of energy at the Yatir forest by means of large-eddy simulations (LES). We performed three simulations, varying the background wind speed (U) to study the effect on the location, extension and strength of secondary circulations and the mechanisms triggering them. Furthermore, we study the effect of these circulations on the air temperature within the forest.

Results and Methods

We use the LES model PALM [3] from the Leibniz University Hannover, which is available under the Gnu General Public License(GPL-3). The time-stepping is done with a third-order Runge-Kutta scheme, the advection scheme with Wicker-Skamarock. The pressure solver uses a fast Fourier transform. PALM has implicit filtering of the subgrid-scales, for which it employs 1.5-order turbulence closure with a prognostic equation for the turbulent kinetic energy. At the lowest grid cell, the Monin-Obukhov similarity theory is applied. We prescribe an inversion layer at the top of the domain. Our lateral boundary conditions are periodic and the surface fluxes fixed, as we model a constant incoming net radiation (conditions at noon). The domain of PALM is split into vertical columns, which communicate by MPI calls.

We used grid dimensions of 5.0 m in the main wind direction, 7.5 m in the crosswind direction and 2.5 m in the vertical. We performed six simulations in total, a weakly convective scenario (WC, U = 5.7 m s-1), a mildly convective scenario (MC, U = 2.8 m s-1), a strongly convective scenario (SC, U = 0 m s-1) and

Fig. 1: Satellite image of Yatir forest (Israel).

three corresponding preliminary simulations. These preliminary simulations are used for turbulent spin-up and run for 2 hours of simulated time, while the main simulations run for 1 hour. The time steps were selected dynamically to maintain a Courant-Friedrichs-Levy number of 0.9, which yielded time steps of approximately 0.23 seconds. The number of grid points and the used core hours are shown in Table 1.

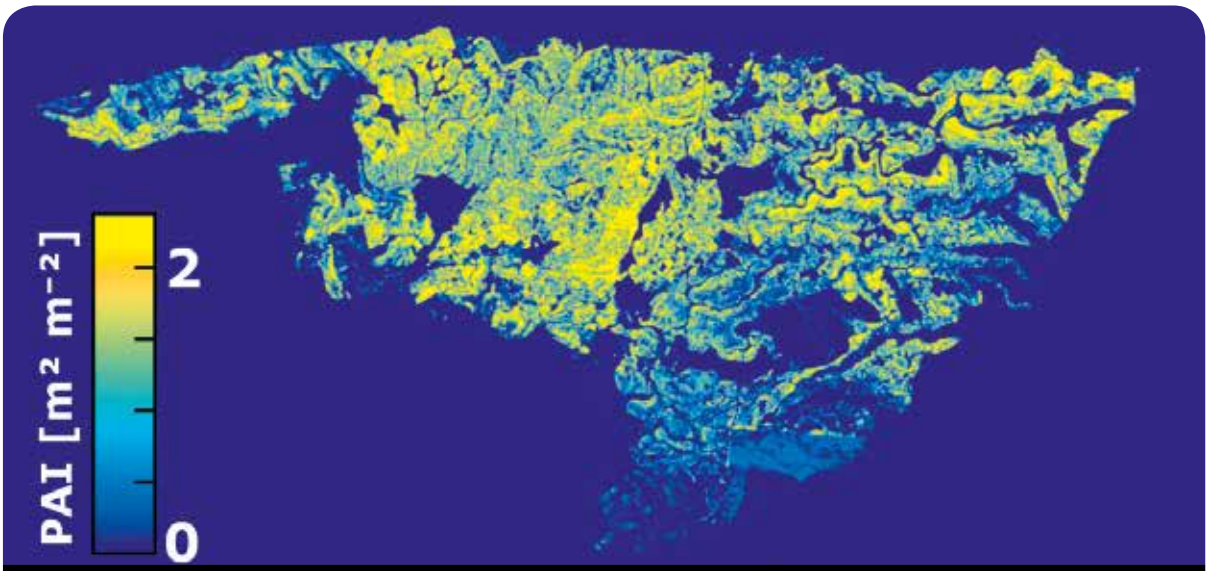

Fig. 2: PAI map of Yatir forest.

	# Grid point	Core hours (106 cpu h)
WC pre+main	6144 x 2048 x 1024	0.34 + 0.53
MC pre+main	6144 x 2048 x 1024	0.30 + 0.66
SC pre+main	4096 x 2816 x 1024	0.26 + 0.65

Table 1: Grid points and core hours for the six LES performed. We used 8192 cores of Phase 1.

The output of the six LES consists of 260 files (400 GB) which are stored on PROJECT. As the main simulations were initialized with the prognostic variables of the preliminary simulations at every grid point, restart data of 30k files (15TB) also had to be stored. For this task we used the PROJECT and SCRATCH file systems.

To initialize the LES, we used a two dimensional map of the plant area index (PAI) of the forest canopy (~tree density), which was derived from satellite data alongside lidar measurements. The derived PAI map is illustrated in Figure 2. Furthermore, the forest is rotated by 45° counter-clockwise, such that the background wind is located in positive x direction (Figure 3). The Coriolis force in the model was adjusted accordingly.

From these LES simulations, we found that secondary circulations emerge in all three cases (updrafts depicted by positive w values in Figure 3), however the location and the extension of these circulations change between all three

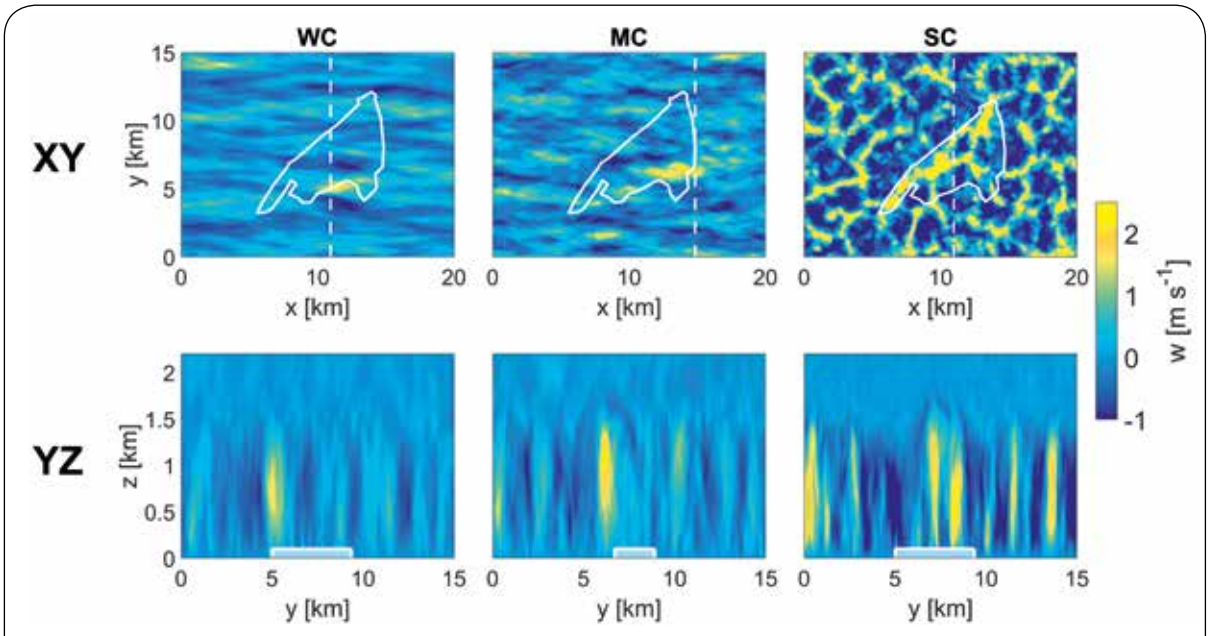

Fig. 3: Cross sections of the horizontal (first row) and vertical cross-wind (second row) of the 30 min mean of the vertical velocity component w, for the three cases WC (first column), MC (second column), and SC (third column). The location of the forest is depicted by the white solid lines in the first row and by white boxes in the second row. The location of the vertical cross section with respect to the forest is depicted by the white dashed lines in the first row.

cases. This is explained by the different mechanisms triggering those simulations. For WC and MC, the secondary circulations are mainly generated by mechanical effects, as they appear in the lee of the densest part of the forest (Figure 2). However, for SC, the secondary circulations are mainly produced by buoyancy, which is indicated by the central location of the updrafts above the forest. The strong buoyancy in the SC case also leads to a fundamental difference in the appearing structures throughout the simulation domain. While in WC and MC roll like structures appear (stripes in Figure 2), for SC hexagonal cells emerge. In addition, the strength of the secondary circulations and the altitudes these circulations reach also differ, as the strength of the updrafts increases from WC to SC.

Ongoing Research / Outlook

We used Phase 1 for our entire project, as the queuing time on Phase 2 became too long when asking for the required 8,192 cores. The main issues we faced was running PALM at the specified number of grid cells and to manage a stable output of the NetCDF files.

In the next phase of this research project, we want to quantify the surface atmosphere exchange of energy between a heterogeneous surface like Yatir forest, and the atmosphere, by means of aerodynamic resistance. As there are only parametrizations for homogeneous surfaces available right now [4], we developed a heterogeneous extension of those parametrizations. To test these extensions for several different surface heterogeneities, we will run a couple of LES simulations of similar dimensions as the ones already performed.

References and Links

[1] www.imk-ifu.kit.edu/projects_2244.php

[2] E Rotenberg, D Yakir, 2010:
Science 22,327(5964), 451-4. DOI: 10.1126/science.1179998.

[3] www.palm.muk.uni-hannover.de/trac

[4] T Banerjee, F De Roo and M Mauder 2017:
Hydrol. Earth Syst. Sci., 21, 2987–3000, DOI:10.5194/hess-2017-4

Written by Frederik De Roo

Institute of Meteorology and Climate Research, Atmospheric Environmental Research, Karlsruhe Institute of Technology

Contact: frederik.deroo@kit.edu

Konstantin Kröniger

Weizmann Institute of Science, Faculty of Chemistry, Department of Earth and Planetary Sciences

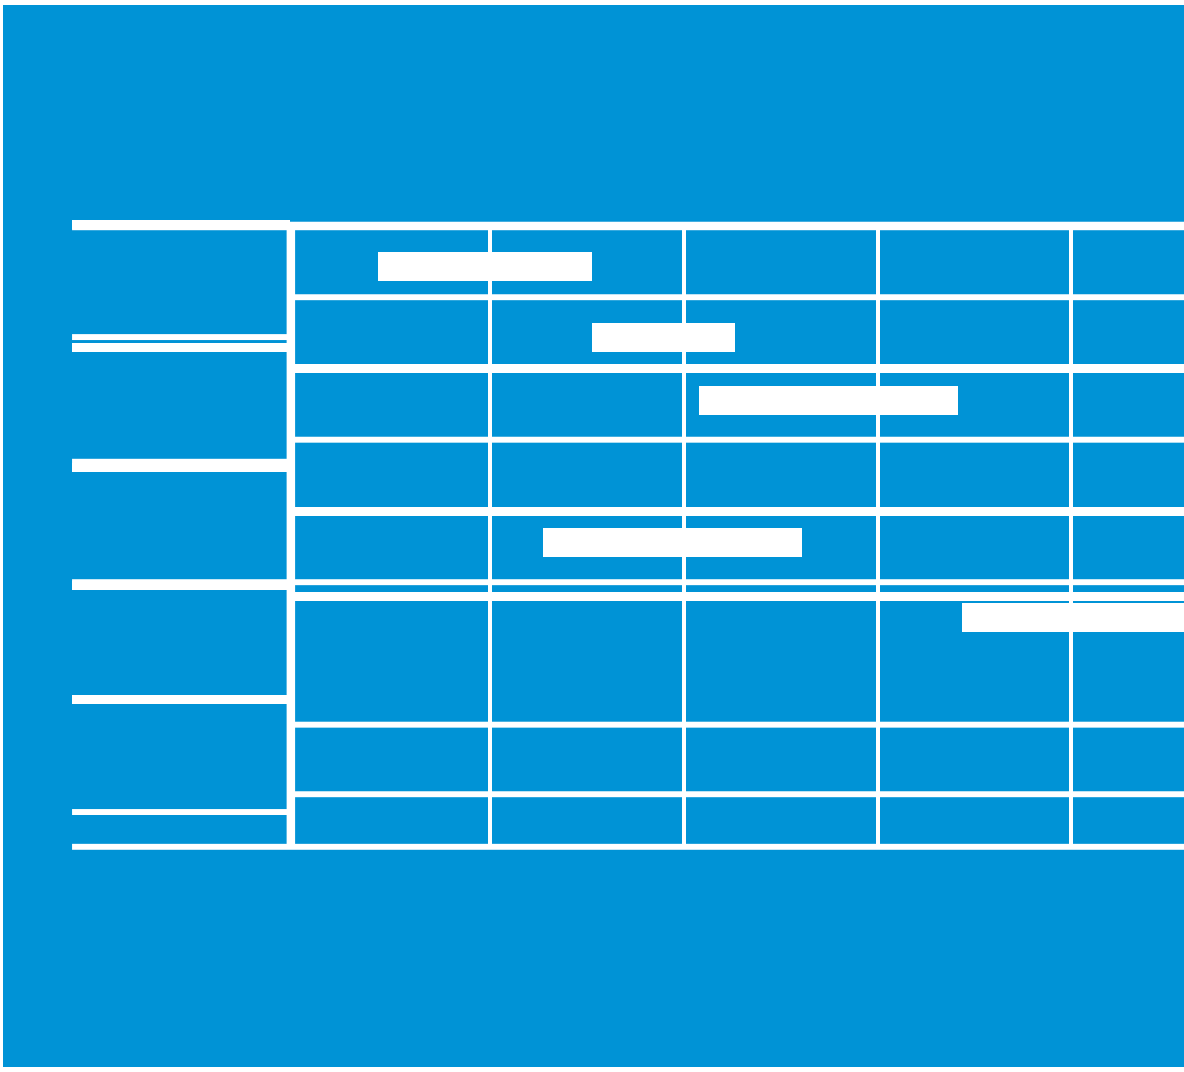

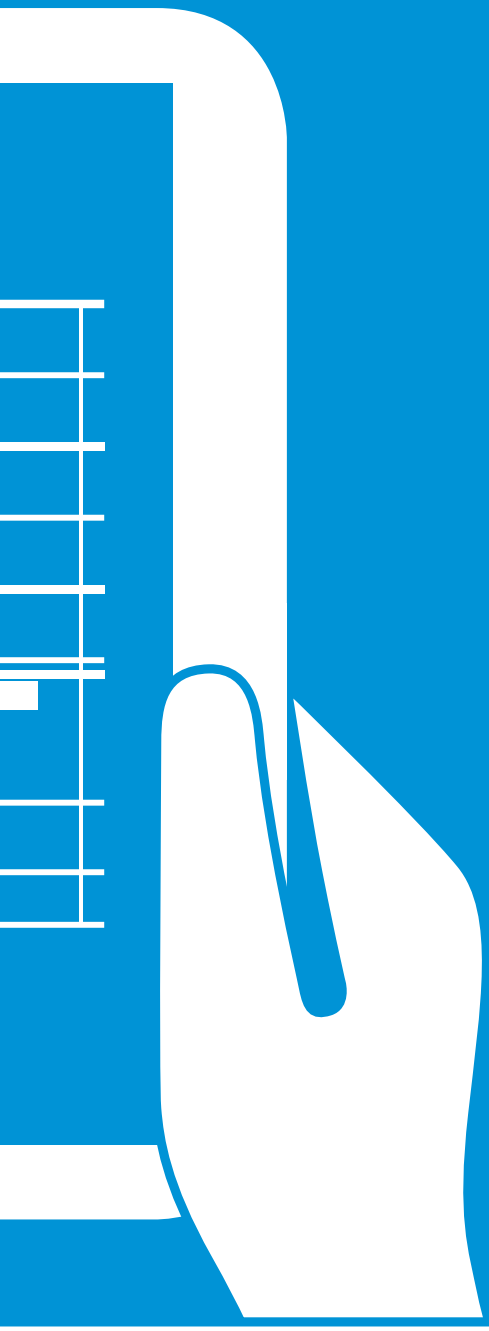

Projects

In this section you will find information about current internal (research) projects of GCS.

ORPHEUS–Fire Safety in the Underground

In January, 2018, the three-year JSC-coordinated research project ORPHEUS ended. ORPHEUS' aim was the investigation of new experimental and numerical methods for fire safety in underground stations. A consortium of 13 partners from academia, industry, and fire brigades contributed to achieve these goals [1].

One of the most prominent aspects of the project were real-scale fire experiments. These experiments were carried out in the Berlin metro station 'Osloer Straße' during off-hours, i.e. between 1 am and 3 am. The experiments were mainly used for underground climate investigations and model validation. The project partner from Aachen (Institut für Industrieaerodynamik Aachen) set up propane burners to physically simulate the thermal and flow conditions, that may be found during the initial phase of a fire (Figure 1). The resulting flows and air temperatures were measured by colleagues from the Ruhr-University Bochum and the smoke detector manufacturer Hekatron Vertriebs GmbH. With over 600 temperature, concentration and flow velocity sensors, they were able to map the full station. Additionally, artificial smoke was added to the burner plumes to visualise the smoke spread and stratification

Fig. 1: Real-scale fire experiment in the metro station 'Osloer Strasse' in Berlin. Propane burners simulate the initial phase of a fire, while the full station is monitored with temperature and concentration sensors. © Lukas Arnold

inside the station. After the measurements, Berlin's fire brigade used the smoke-filled station to critically evaluate their new firefighting tactics in metro stations (Figure 2).

Fig. 2: Berlin's fire brigade trains new tactics and tools in a smoke-filled station. © Ralf Eisenbach

Besides the real-scale fire experiments, the project partner ROM Technik from Hamburg designed novel smoke extraction systems. These concepts were successfully evaluated in small scale experiments, where the physical model represented the full metro station in a 1:15 scale (Figure 3). In order to compare the experimental data to simulations, all experimental studies were accompanied by numerical simulations using FDS (Fire Dynamics Simulator) and Ansys CFX. Whereby most of the production level computations were carried out by the Bundesanstalt für Materialforschung und -prüfung.

JSC's scientific contribution to the ORPHEUS project was the development of new numerical methods. Four JSC PhD students contributed to the research activities in the fields of pedestrian and smoke dynamics. The intersection of these two fields was the focus of Benjamin Schröder's dissertation [2]. He investigated the analysis of life safety in a complex underground infrastructure. His work includes the definition of a large scenario ensemble that covers multiple fire locations and design fires. Based on this ensemble, a spatio-temporal analysis was carried out in order to evaluate the criticality of each scenario. For the first time, this approach allows to identify critical locations in a complex building (Figure 4).

Other JSC work packages covered the development of new numerical approaches to simulate the smoke spread. Here, two paths were investigated: adaptive mesh refinement and GPU-based real-time concepts. The GPU-based implementation aims at smoke spread

Fig. 3: Inside view of the 1:15 scaled model of the metro station 'Osloer Strasse'. Small-scale smoke sources are used to model the smoke spread. © Lukas Arnold

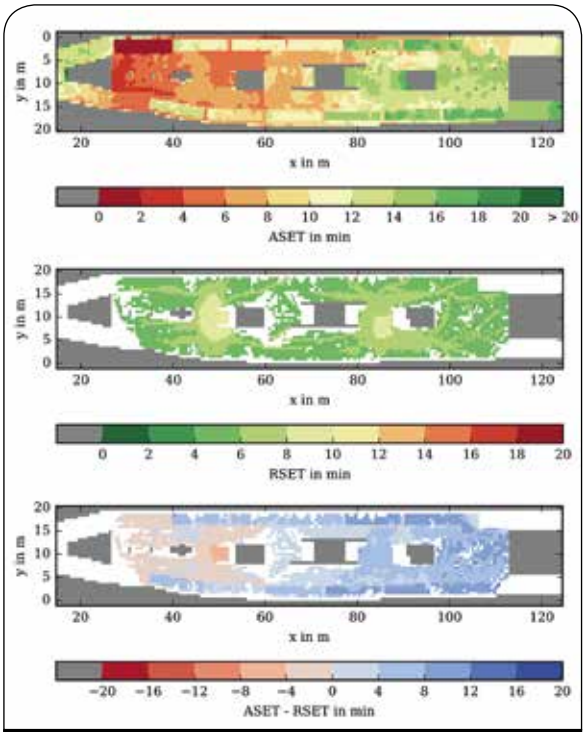

Fig. 4: Identification of critical locations in a complex building. Shown is the lowest level of the metro station 'Osloer Strasse' for a single scenario with a fire seat located at the top right corner of the platform. The top figure shows the computed available time for a safe evacuation at all platform positions. This time is compared to the required time (middle figure), which is based on evacuation simulations. The combination of both results in the difference map is shown at the bottom. Locations with negative time differences indicate critical locations with respect to life safety.

simulations in real-time, while using a portable programming approach [3]. Both activities lead to first proof-of-concept implementations and their development will be further continued in follow-up projects.

In ORPHEUS, new fire safety engineering methods and tools were applied, evaluated and developed in the context of underground stations. However, many of these concepts can be

Fig. 5: The ORPHEUS consortium. © Alexander Belt

transferred to other complex infrastructures like airports or stadiums. In total, the project members published more than 20 contributions in national and international conferences and journals, where an overview of the achieved results of all project partners was given in a public event held in January, 2018 in Berlin [4].

References

[1] **ORPHEUS project webpage:**
http://www.orpheus-projekt.de

[2] **Benjamin Schröder:**
Multivariate Methods for Life Safety Analysis in Case of Fire. Dissertation at the Bergische Universität Wuppertal 2017, http://elpub.bib.uni-wuppertal.de/servlets/DerivateServlet/Derivate-7013/dd1611.pdf

[3] **Anne Küsters, Sandra Wienke and Lukas Arnold:**
Performance Portability Analysis for Real-Time Simulations of Smoke Propagation Using OpenACC. Lecture Notes in Computer Science 2017. http://dx.doi.org/10.1007/978-3-319-67630-2_35

[4] **ORPHEUS Abschlusstreffen Januar 2018:**
http://www.orpheus-projekt.de/SharedDocs/Meldungen/ORPHEUS/DE/020_abschlusstreffen.html

Written by Lukas Arnold
Civil Safety and Traffic
Jülich Supercomputing Centre (JSC)
Contact: l.arnold@fz-juelich.de

EU H2020 CoE Performance Optimization and Productivity (POP) Successfully Finished

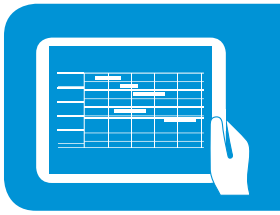

From October 2015 to March 2018, the EU Horizon 2020 Center of Excellence (CoE) for Performance Optimisation and Productivity (POP) provided performance optimisation and productivity services for academic and industrial codes. Europe's leading high-performance computing experts helped application developers get a precise understanding of their respective applications' and systems' behaviour. Both established codes and codes which had never undergone any analysis or performance tuning profited from POP services, which used latest state-of-the-art tools to detect and locate bottlenecks in applications, suggested possible code improvements, and even helped with proof-of-concept experiments for customer codes on their own platforms.

Today's complexity of high-performance computer systems and codes makes it increasingly difficult to get applications running fast and efficiently on the latest hardware. Often expert knowledge and a good amount of experience is needed to figure out the most productive direction for code refactoring. Many domain experts in in industry and academia use computer simulations, but lack this knowledge.

As a result, their codes are often far away from using the hardware efficiently, using much more compute time than needed. This lack of optimization can waste energy, require superfluously oversized and expensive hardware, or just miss research potential, as their codes can only handle smaller or less complex problems in the available amount of compute time.

To overcome this situation, the POP CoE brought users a service that tightly couples two disciplines crucial for the efficient use of parallel computers in the future: First, powerful performance analysis tools, methodologies, and expertise needed to precisely understand and gain real insight into the actual application and system behaviour as well as a deep understanding of programming models and best-practice guidance needed to express algorithms in the most flexible, maintainable and portable way, while still being able to maximise the performance achieved.

POP Services

The POP CoE team comprised six partner organisations with experts in high-performance computing with long-standing experience in performance tools and tuning as well as researchers in the field of programming models and programming practices. All partners have a research and development background and proven commitment to applying their know-how to real academic and industrial use cases. The POP CoE provided three kind of service levels to its customers – depending on their background, knowledge, and demands:

?: Application Performance Audit

This was the primary service of the POP CoE and the starting point for any further work. Applications using this service were analyzed by POP experts after an initial discussion with respect to their best practices and provided a first impression of the code status. Within the Performance Audit, the customer's code

performance issues could be identified at the location a customer would normally run his or her code. It also served as a starting point for further analysis or initial code refactoring. The duration for a Performance Audit averaged around two months and a successful Performance Audit may be seen as a code quality certificate in HPC.

!: Application Performance Plan

The Performance Plan service followed the Performance Audit when the customer needed more detailed knowledge where and how to address specific issues in the code. POP experts developed together with the customer a plan how and with which tools to analyse the issues under investigation. The POP experts then analysed the code in detail and gave quantified advice to overcome the problems that could be fixed by the customer. The duration for a Performance Plan is very problem-specific, but always included a closer look into the source code.

√: Proof-of-Concept

When requested, proof-of-concept studies were performed. This included experiments and mock-up tests for customer codes. The details of the proof-of-concept study were decided in very close collaboration with the customer and could include kernel extraction from the application, parallelisation, or mini-apps experiments to show effects of the POP experts' proposed optimisations. As this very complex task goes into deep detail, proof-of-concept work sometimes required about six months.

Besides the above three key services, the POP CoE also provided a variety of training activities in the field of performance analysis and optimisation to improve basic high-performance programming knowledge and increase the awareness of performance issues and potentials in general.

Codes Analyzed

In its 30 months of operation, POP has undertaken over 150 assessments of codes drawn from a wide range of scientific domains covering astronomy, chemistry, Earth science, energy, engineering, health, mathematics and physics.

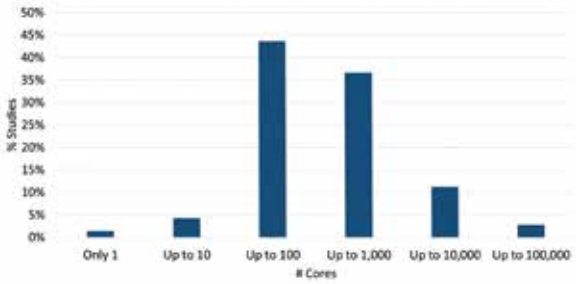

Although the bulk (80%) of POP studies have looked at codes that run best on more than 10 but less than 1,000 cores, POP also performed assessments on larger scale codes over 100,000 cores.

Roughly half of the assessments originated from academic institutions, a quarter from research or government laboratories, and a quarter had an industrial background.

Languages and Parallelism

As one might expect, Fortran codes dominate, with over half the studies (82 of 151) written either entirely in Fortran or Fortran combined with C

and/or C++. C++ seems more prevalent than C, but C is more likely than C++ to be combined with Fortran:

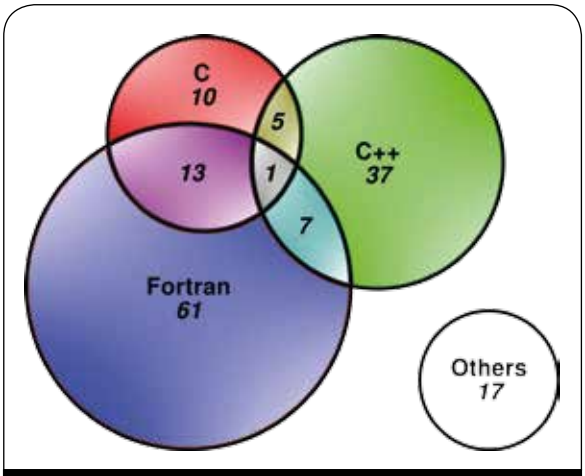

Fig.2: Programming languages used by POP customer applications. The areas of the circles in the Venn diagram are proportional to the number of studies they contain.

The "Others" category makes up about 10% of studies. Of these, 13 involved Python, either stand-alone (3) or in conjunction with one or more compiled language, and the remainder were a combined C/Fortran/Octave code, a Java code, a Matlab code, and a Perl code. The fact that 10% of codes that POP assessed are written in languages other than C/C++/Fortran demonstrate that it's important to have tools and methodologies capable of handling a wide range of languages and not just the usual suspects.

A similar situation arises looking at the types of parallelism used by the codes studied by POP. MPI is the most common form of parallelisation, with nearly 80% of codes using either pure MPI or MPI+OpenMP. Over a third of codes POP has assessed are hybrid combining MPI+OpenMP.

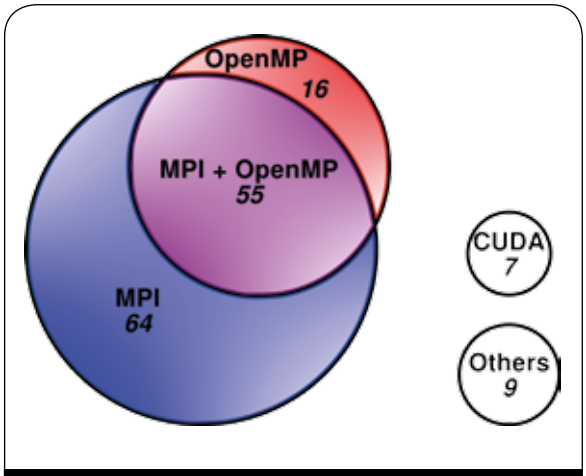

Fig.3: Parallelization paradigm used by POP customer applications. The areas of the circles in the Venn diagram are proportional to the number of studies they contain.

7 codes could also use CUDA and there is again an „Others" category, but unlike the corresponding category for the languages every member of this set is unique. Examples of the other types of parallelism we encountered include Intel Threading Building Blocks (TBB), C++ threading and Coarray Fortran.

Causes of Low Efficiencies

The POP efficiency metrics [1] provide a methodology for characterising the performance of parallel codes and for providing insight into where the most pressing problems lie. Looking across all the POP studies we can use the metrics to discern whether there are any overarching performance trends.

Of the analysed codes, 66% had a Parallel Efficiency less than 80%, meaning that they typically required improvement to run efficiently in parallel. Indeed, 22% of codes had Parallel Efficiency below 50%, which means that less than half of their runtime is dedicated to computation. Note that analysis generally omits initialisation and finalisation, so in practice their efficiency is even worse.

Looking at the actual numbers reported in the studies, we find that Load Balance Efficiency is often either very good or very bad. This suggests that load balance is something users must ensure is done correctly or else can have significant impact on efficiency, particularly when scaling to larger numbers of cores.

We can also use the hierarchical nature of the metrics to look at the common underlying causes of low efficiencies. Low Communication Efficiency is mostly caused by data transfer (high volume of data or high number of communications) rather than serialisation of communication. Low Computation Efficiency is often caused by poor instruction scalability rather than reduced instructions-per-cycle IPC values; when strong scaling, growth in the total number of instructions executed often corresponds to undesirable code replication.

POP staff can look further into the categories of problems to see if there is a link between programming approach and the types of problems. Although there was no obvious correlation between language and inefficiency (e.g. we couldn't conclude things like "C programmers were more likely to write badly load-balanced code"), there was an interesting distinction to be drawn based on the type of parallelism the code employed.

For each study we recorded the main cause of inefficiency that was identified (i.e. load balance, computation or communication) and looked at how this varied across the three main types of parallelism:

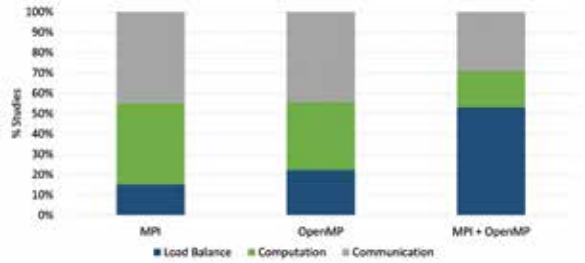

From the graph, we can see that studies of hybrid codes were much more likely to report problems with load balance than studies of pure MPI or pure OpenMP codes. This is perhaps understandable—when writing hybrid code, users need to take into account both how the work is divided across MPI ranks and also how it is split up between threads of a process.

Results

The POP project was very successful: more than 90% of the customers were either very satisfied or satisfied with the service they received. More than half of the customers of a Performance Audit requested a follow-up service. More than 2/3 of the customers with Performance Plans indicated they plan to use POP services again.

The proof-of-concept studies, where POP experts worked together with the developers on improving their codes, were especially successful: in many cases, they were able to demonstrate a doubling of the performance and/or scalability of the codes investigated. In some cases, a six- to ten-fold improvement could be implemented. A more detailed description of these accomplishments can be found on the POP blog under the tag "success stories" [2].

In the case of Performance Plans, where developers improved their codes themselves based on the recommendations given by the POP experts, they reported a 25% performance or scalability improvement on average, in some cases even 50% to 70%, allowing them to treat larger problems or better exploit new architectures. Customers also reported that, in most cases, only a few days' effort was necessary to perform this work; the remainder required either a few weeks' or a few months' effort.

The return-on-investment (ROI) in these cases are enormous, as the following two examples demonstrate: In the first case, where an application running on the UK national academic supercomputer (Archer) was first analyzed and then in a proof-of-concept study, had a 72% improvement in time-to-solution could be implemented. The saving in compute time for a typical run of this code was €15.58, which resulted in a yearly saving of €56,000 for this specific customer based on their monthly usage data. In the other example, the customer reported that the costs for implementing the recommendations of the POP experts were €2,000 and resulted in a 62% performance improvement. By this, €12,400 of the customer's annual €20,000 operating cost could be saved, resulting in a ROI of 620%.

Impact

The POP CoE, with its service and training activities, had a wide impact within all areas of research and industry, making it a real transversal activity:

• It provided access to computing application expertise that enables researchers and industry to be more productive, leading to scientific and industrial excellence.

• It improved competitiveness for the centre's customers by generating a tangible return-on-investment (ROI) in terms of savings, elimination of waste, errors, and delays by making their applications leaner and issue-free.

• As the centre represents European world-class expertise in this area, its deployment strengthened Europe's leading position in

the development and use of applications that address societal challenges or are important for industrial applications through better code performance and better code maintenance and availability.

• The centre's services included training on the use of computational methods and optimisation of applications; especially successful were webinars [3].

It is planned to continue the POP CoE project for additional three years (2019 to 2021). A proposal was prepared and submitted by the partners for the call INFRAEDI-02-2018: HPC PPP – Centres of Excellence on HPC.

Partners

Barcelona Supercomputing Centre (BSC), High-Performance Computing Center Stuttgart of the University of Stuttgart (HLRS), Jülich Supercomputing Centre (JSC), Numerical Algorithms Group (NAG), Rheinisch-Westfälische Technische Hochschule Aachen (RWTH), and TERATEC.

Timeframe

October 2015 – March 2018

POP Coordination

Prof. Jesus Labarta, Judit Gimenez Barcelona Supercomputing Center (BSC)

Email: pop@bsc.es

Web: https://www.pop-coe.eu

Acknowledgements

This project was supported by the European Commission under H2020 Grant Agreement No. 676553

References

[1] https://pop-coe.eu/node/69

[2] https://pop-coe.eu/blog/tags/success-stories

[3] https://pop-coe.eu/blog/tags/webinar

Written by Bernd Mohr, Brian J. N. Wylie
Jülich Supercomputing Centre (JSC)
Contact: b.mohr@fz-juelich.de, b.wylie@fz-juelich.de

José Gracia, Christoph Niethammer
University of Stuttgart (HLRS)
Contact: gracia@hlrs.de, niethammer@hlrs.de

TELEMAC - A hydrodynamic solver for HPC?

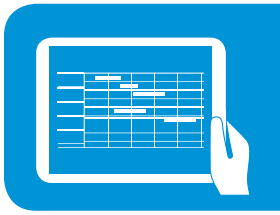

The open-source software TELEMAC-MASCARET is a powerful tool for river engineering applications, such as accurately representing flood inundations, long-term developments of rivers, as well as investigating ecological flow conditions. The software is MPI-parallelized and tested on diverse high-performance computing systems [1]. For the research project "Wasser-Zukunft Bayern", the potential and suitability of the software on the SuperMUC is evaluated.

The TELEMAC-MASCARET system consists of different modules that can be directly coupled depending on a user's requirements [2]. Here, the two-dimensional shallow-water solver TELEMAC2D is applied using the distributed version 7.2. The software is designed on open-source applications where the standard compilation is GCC and OpenMPI. However, the hardware at LRZ is optimized for Intel so we installed the software with the Intel compiler (v.17) and IntelMPI.

Performance and scalability test cases

We developed three real simulation cases varying in length and the number of elements (0.5, 5 and 20 million elements, respectively), but with identical flow characteristics. Based on these cases, software scalability was analyzed for strong scaling. For each case, simulations are performed with increasing number of nodes, specifically cores (1 node = 28 cores). The reference time is set as the time that is needed for simulating one node. The software shows good scaling effects until 32 nodes, despite the 20 Mio. case where memory issues prevented simulation during domain decomposition.

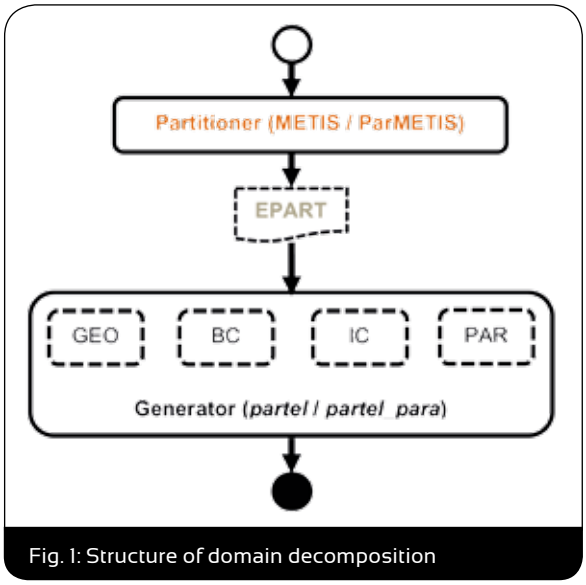

Fig. 1: Structure of domain decomposition

Code analysis

The domain decomposition for TELEMAC is done with a graph-decomposer named METIS (5.1.0). This follows the workflow sketched in Figure 1. The output of METIS is the decomposed graph and is stored in a variable named EPART. With the result of METIS, each element is assigned to a subdomain. This information will be used in the later stages of the decomposition to process the remaining information. This step is completed by a TELEMAC internal function named partel that generates the necessary files for the simulations. Through analyzing the code in detail we were able to identify the problems related to the memory usage. The code contains some variables, whose dimensions depend on the size of elements and number of partitions (e.g. CUT_P and GELEGL). Furthermore, in the source code of TELEMAC some fragments of a parallel domain decomposition exist. Here ParMETIS is used for graph

decomposition followed by a function named partel_para. The resulting structure is similar to the one sketched in Figure 1, however, this code is not working properly, as some required information is not generated correctly (e.g. the PAR-file). This is mandatory for the information exchange between the subdomains during the simulations. To overcome those problems, we developed a workaround. The basic idea of the modifications is to use ParMETIS for the graph decomposition, and then use a simple I/O-interface to pass the correctly produced variables to the serial partel function.

Results

With the developed code, we were able to perform the domain decomposition and simulation runs up to 60 nodes for the two smaller domains. Unfortunately, for a 20 Mio. element case there was the reoccurrence of memory problems, which might be a result of nested loops. In figures 2, 3 and 4, the strong scaling of the software is shown separately for the domain decomposition (split), the simulation (run) and the final merging (merge) as well as the total time. In the figures, the linear speed-up for the simulation until approximately 2,000 elements/core can be seen quite clearly. Furthermore, the drop in the split performance is a result of the I/O-interface, which is relatively slow but working. Here one must consider that the simulations should run several days and simulate multiple years, thus the duration for domain decomposition, which lasts only a few minutes, does not affect the overall performance significantly.

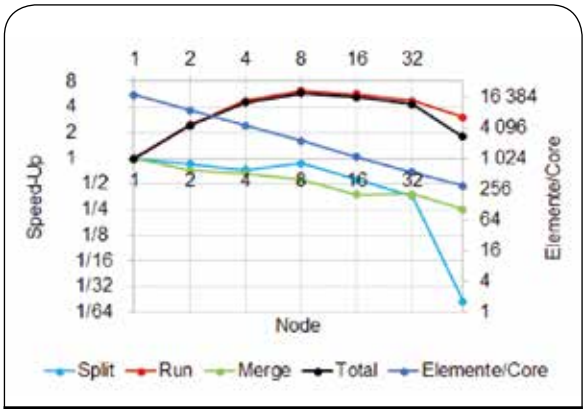

Fig. 2: Speed-Up for a 0.5 Mio. element case

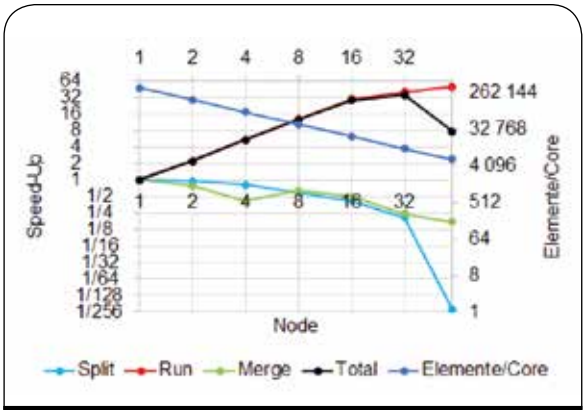

Fig. 3: Speed-Up for a 5.0 Mio. element case

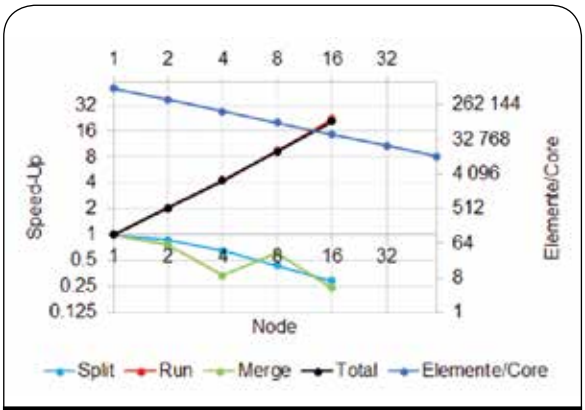

Fig. 4: Speed-Up for a 20 Mio. element case

Unfortunately, we were not able to perform the simulation for the largest case, Figure 4, so the performance chart stops at 16 nodes.

We will drive our development by implementing a fully parallelized domain decomposition without the I/O-interface and reducing the memory requirements for the allocation of the mentioned variables.

Acknowledgements

We gratefully acknowledge the support of two application laboratories at the Leibnitz Supercomputing Centre, namely the CFDLab and the GeoLab for the collaboration. Furthermore, we would like to thank the KONWIHR initiative (Bavarian Competence Network for Technical and Scientific HPC) for funding this study.

References

[1] **Moulinec, C., et al., 2011:**
TELEMAC: An efficient hydrodynamics suite for massively parallel architectures. Computers & Fluids, 51(1), 30-34. https://doi.org/10.1016/j.compfluid.2011.07.003

[2] **Hervouet, J.M., 2007:**
Hydrodynamics of Free Surface Flows: Modelling with the finite element method. John Wiley and Sons, Chichester, England. http://dx.doi.org/10.1002/9780470319628

Written by Markus Reisenbüchler, Ngoc Dzung Nguyen, Tobias Liepert, Minh Duc Bui, Peter Rutschmann
Leibniz Supercomputing Centre (LRZ)
Contact: markus.reisenbuechler@tum.de

SMITH: Smart Medical Information Technology for Healthcare

The SMITH project aims at making better use of the wealth of patient data in order to validate the relevance of novel diagnostic and therapeutic methods in systematic trials, ultimately hoping to directly improve care. One goal is to optimize diagnosis and therapy in intensive care unit (ICU) patients with lung failure. University Hospital RWTH Aachen, RWTH Aachen University, Bayer AG and the Jülich Supercomputing Centre are collaborating in the clinical use case Algorithmic Surveillance of ICU patients (ASIC).

The SMITH project is one of four consortia funded by the German Medical Informatics Initiative (MI-I) including partners from universities, university hospitals, research institutions and IT companies. The benefits for patient care that drives the SMITH project include the availability of actionable, data-driven decision support standardized across healthcare organizations such as the seven German hospitals that are part of the SMITH consortium. The overarching goals are to establish data integration centers (DICs) at each SMITH partner university hospital and to implement use cases, which demonstrate the usefulness of the SMITH approach.

SMITH Use Cases

The methodological use case Phenotype Pipeline (PheP) develops algorithms for annotations and analyses of patient-related phenotypes according to classification rules or structured-data-based statistical models. The word 'phenotype' can be interpreted in a very general sense—in this context, it refers to a set of attributes that can be attached to an individual. Unstructured textual data will be subject to natural language processing (NLP) to permit integration into the phenotyping algorithms. The goal of this use case is creating annotations (aka tags) for a broad spectrum of attributes linked to a specific patients' patterns of care.

The first clinical use case, Algorithmic Surveillance of ICU Patients (ASIC), focuses on ICU patients with acute respiratory distress syndrome (ARDS). A model-based decision-support system will advise medical professionals about mechanical ventilation and is planned to support the daily work within selected university hospital ICUs. The second clinical use case, HELP, develops a hospital-wide electronic medical record-based computerized decision support system to improve outcomes of patients with bloodstream infections. Both ASIC and HELP will use elements of the methodological use case PheP. The clinical benefit of the use cases ASIC and HELP will be demonstrated in a change of care clinical trial based on a step wedge design in hospitals. This article focusses on the ASIC use case.

Clinical Use Case Algorithmic Surveillance of ICU Patients (ASIC)

The demand for intensive care medicine will increase over the next 10—15 years. This can be explained by the epidemiological challenges Germany faces. There is a need for improvement and it is especially urgent in patients

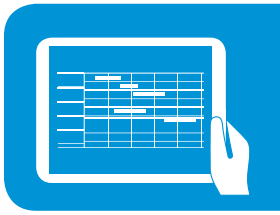

suffering from acute respiratory distress syndrome (ARDS), the focus of the ASIC use case. Incidence of ARDS worldwide remains high, with 10.4% of total ICU admissions and 23.4% of all patients requiring mechanical ventilation [2]. The goal of the ASIC use case is to perform a continuous analysis of data obtained from the hospital patient data management system (PDMS) in order to enable model-based 'algorithmic surveillance' of the state of critically ill patients. The ASIC use case, therefore, develops the following two key elements in the form of a model and a decision support system.

Virtual Patient Model

In order to predict individual disease progression of ICU patients, ASIC will utilize pattern recognition technologies as well as established mechanistic systems medicine models, complemented by machine learning and high-performance computing (HPC), both integrated in a hybrid virtual patient (VP) model [3]. This model will enable individual prognoses to support therapy decisions, clinical trials, and training of future clinicians. Training the VP model requires HPC and will take advantage of Markov Chains Monte Carlo (MCMC) methods and deep learning algorithms [4, 5]. HPC computing time estimations for one given VP model is only roughly 2 sec (1 core) but given the uncertainty, 106 runs are required for the MCMC approach, leading to 1 patient in roughly 1,000 core hours. Of course the SMITH project aims to compute at least 1,000 patients' data in many iterations in order to improve modeling. For the foreseen clinical practice with new patients on a daily basis, however, a model reduction is planned by mapping the model onto a deep learning network. Another computationally intensive part of the VP approach is the deep learning network training process. This, in turn, requires HPC for training but also for validation and model selection in order to scan the full parameter space of such a deep learning network.

Diagnostic Expert Advisor

Beside the above described hybrid modeling approach, the ASIC use case will also contribute to the development of an online, rule-based computerized decision-support system (CDSS). This system will extensively use a wide variety of DIC services and the Phenotype Rules-Engine that applies phenotyping rules developed in the PheP use case. The goal of CDSS in clinical practice is to accelerate correct and sound diagnosis and increase guideline compliance regarding ventilation. The rules that drive the CDSS design are realized by explicit decision trees, complex models, mechanistic or machine learning-based, or combinations of both. The main component of this CDSS system is the Diagnostic Expert Advisor (DEA) and all elements of the ASIC system will be accessible via an ASIC app. This app will provide the interface between the DIC, the surveillance algorithms outlined above and the medical professionals at the bedside. It enables the functionality of automatically informing physicians if a priori specified limits are exceeded. This alert function will be enhanced by the latest action-based algorithms, preventing medical professionals from moot alerts (and, thus, alert fatigue).

References

[1] A. Winter et al.:
Smart Medical Information Technology for Healthcare (SMITH) – Data Integration based on Interoperability Standards, in Methods, 2018, to be published

[2] K. Henry:
A targeted real-time early warning score (TREWScore) for septic shock, in Sci Transl Med 2015, 7(299):299ra122.

[3] O. Wolkenhauer et al.:
Enabling multiscale modeling in systems medicine, In Genome Med 2014; 6(3):21

[4] M. Krauss et al.:
Bayesian Population Physiologically-Based Pharmacokinetic (PBPK) Approach for a Physiologically Realistic Characterization of Interindividual Variability in Clinically Relevant Populations, Doi: https://doi.org/10.1371/journal.pone.0139423

[5] M. Krauss et al.:
Translational learning from clinical studies predicts drug pharmacokinetics across patient populations, in Systems Biology and Applications (3), 2016

Written by Prof. Dr. med. Gernot Marx
RWTH Aachen University Hospital, Department of Intensive Care and Intermediate Care; Aachen
Contact: nreimer@ukaachen.de

Prof. Dr. Ing. Morris Riedel
Forschungszentrum Juelich – Juelich Supercomputing Centre
Contact: m.riedel@fz-juelich.de

Prof. Dr. Andreas Schuppert
RWTH Aachen University; Institute for Computational Biomedicine II, Aachen
Contact: schuppert@aices.rwth-aachen.de

M.Sc. Oliver Maaßen
RWTH Aachen University Hospital, Department of Intensive Care and Intermediate Care; Aachen
Contact: omaassen@ukaachen.de

Bringing HPC into the Cloud

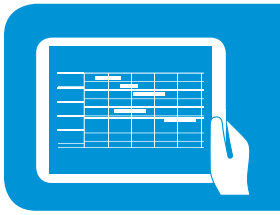

As partners in a multinational EU research project, HLRS investigators developed software that facilitates the integration of two powerful technologies.

Oftentimes researchers or companies need to solve problems that require high-performance computing but don't have access to the computing systems they need. One option is to calculate on physical high-performance computing (HPC) systems like the one at HLRS, although such systems are not always configured to accommodate all applications (for example, because they don't support specific programming languages or operating systems).

Another option that is becoming increasingly attractive is to be able to perform such calculations in the cloud. Cloud-based services offer the benefits of computing availability and cost efficiency, and have more options for customizing resources to an individual user's needs. However, they also show shortcomings in performance when doing the kinds of compute- and data-intensive tasks that HPC demands.

Funded with roughly €6 million under the auspices of the European Union's Horizon 2020 research program, an international collaborative project called MIKELANGELO developed a software stack that enables big data and HPC applications to be carried out efficiently in the cloud. As a partner in the MIKELANGELO team, HLRS contributed to the project's success by creating a job management system that makes this possible.

Improving resource allocation and I/O

Bringing together HPC and the cloud requires virtual machines (VMs). VMs enable users to reproduce their local computing tools and architectures remotely, while at the same time taking advantage of large CPU and memory amounts that they do not have on their own.

When users run computing jobs in virtual HPC environments, virtual resource administrators must define the memory, CPUs, software, operating system, and other resources and system requirements that they need. This enables users to utilize a virtual HPC environment—including all necessary application components or even operating systems—just as they might a traditional HPC environment.

As part of MIKELANGELO, HLRS developed a software tool for flexible resource allocation in the cloud that was essential to the project's success. The software, called vTorque, is designed not only to schedule traditional batch jobs but also to manage entire virtual infrastructures. With vTorque users can control an entire application remotely from their own computers, without being restricted to one specific operating system or software.

Another challenge that virtualization faces results from the process of data input and output (I/O). This challenge is even more problematic when using HPC applications. Because of the large amounts of data they need to process, they demand vast resources for data processing, management, and storage.

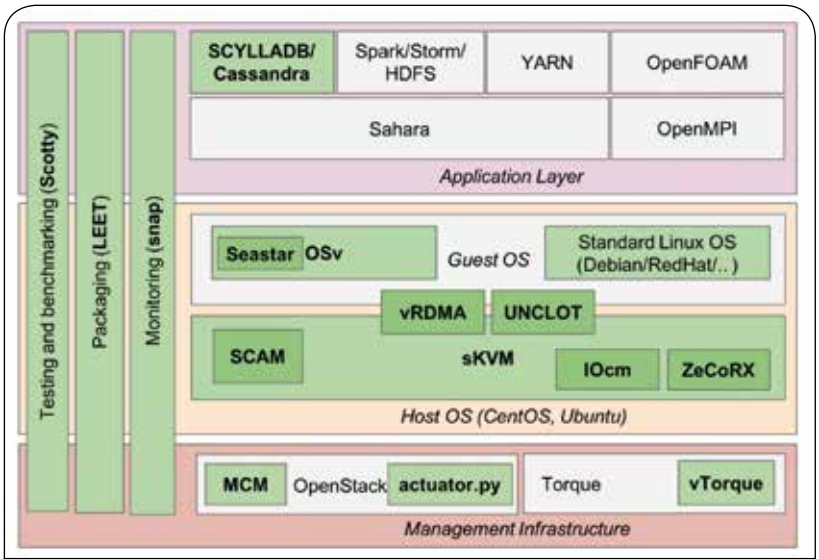

MIKELANGELO final architecture.

HLRS's MIKELANGELO partners developed software components to address such I/O challenges—for example, by accelerating virtual I/O and establishing high-performance interconnect networks. HLRS staff successfully integrated these software components into vTorque, demonstrating its functionality on both cloud-based and HPC-based system environments.

When they compared an HPC application's performance in both environments HLRS investigators found that it only ran a maximum of 3% more slowly on a VM using vTorque than it did running on a physical HPC system, an impressive accomplishment.

A roadmap for the future

MIKELANGELO was recognized in an annual publication of the Horizon 2020 program as one of the top success stories in EU-funded German information and communication technologies research. The project also showed great impact on the ETP4HPC Strategic Research Agenda (SRA) of the European Commission. The SRA outlines a roadmap for achieving exascale capabilities within the European HPC ecosystem, serving as a foundation for future projects and for assessment of the entire Horizon 2020 program. The SRA's latest update in November 2017 showed great interest in conducting additional research on topics related to MIKELANGELO. Virtualization techniques in particular have been identified as technical research priorities for HPC system architecture and components and are on the list of trends to be monitored by the community.

MIKELANGELO was coordinated by Slovenian R&D company XLAB, with partners from Germany, Israel, Ireland, and Slovenia.

Written by Lena Bühler

Integrated Services for the European Open Science Cloud

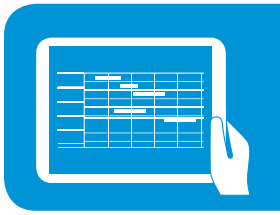

The EC-funded project EOSC-hub project started on January 1, 2018, bringing together an extensive group of national and international service providers to create the hub: a central contact point for European researchers and innovators to discover, access, use and reuse a broad spectrum of resources for advanced data-driven research.

The consortium of 100 partners from more than 50 countries will develop the vision of the hub as the integration and management system of the future European Open Science Cloud. The EOSC-hub project mobilizes providers from the EGI Federation, EUDAT collaborative data infrastructure (CDI), INDIGO-DataCloud and other major European research infrastructures to deliver a common catalogue of research data, services and software for research.

For researchers, this will mean a broader access to services supporting their scientific discovery and collaboration across disciplinary and geographical boundaries.

EOSC-hub will significantly reduce the fragmentation of IT facilities and digital tools in Europe. By bringing together a broad range of services from general and domain-specific research digital infrastructures under a common integration and operation layer, the EOSC-hub will foster new modes of working for collaborative research to deliver trusted services.

EOSC-hub will collaborate closely with OpenAIRE Advance, a project set up to support the open access and open data mandates in Europe.

The Jülich Supercomputing Centre (JSC) is one of the funding members of the EUDAT CDI. JSC operates B2ACCESS, the federated cross-infrastructure authorisation and authentication framework for user identification and community-defined access control enforcement in the EUDAT CDI, which allows EUDAT users to authenticate themselves using a variety of credentials. JSCs second major contribution to the EUDAT CDI is the operation of the B2DROP service, which is a secure and trusted cloud storage to store and exchange data, stipulating how, with whom, and for how long and which allows automatic desktop synchronization of files. In the EOSC-hub project, JSC will integrate B2ACCESS and B2DROP with services from the EGI Federation and INDIGO-DataCloud and support communities in EOSC Competence Centres and thematic service providers in exploiting these services.

Reference

Website: http://eosc-hub.eu

Written by Daniel Mallmann
Jülich Supercomputing Centre
Contact: d.mallmann@fz-juelich.de

MYX: MUST Correctness Checking for YML and XMP Programs

Exascale systems are expected to consist of tens of thousands of compute nodes complemented by specialized accelerators, resulting in heterogeneous system architectures on multiple levels. Such architectures challenge programmers to write multi-level parallel programs, which means employing multiple different paradigms to address each individual level of the system's parallelism. The levels range from inter-node parallelism addressed by distributed memory programming, to inner-node parallelism addressed by shared memory programming exploiting multi-core processors and acceleration units, to vector-style programming targeting corresponding hardware units.

Regardless of the programming model used, parallel programming will always be more challenging than programming for sequential code, as parallel programming introduces a new dimension of possible mistakes to make. Parallel applications can exhibit a wide range of errors, ranging from rather simple mistakes such as invalid arguments used in communication routines to complex issues such as deadlocks or race conditions that only occur when a large amount of parallelism is employed. The use of correctness tools to detect and report such errors and to avoid the occurrence of such errors in production runs can greatly extend the productivity of HPC. There are different approaches to check for correctness. Static source code analysis is very limited in scope, as it can hardly deal with indirect addressing, pointers, and similar constructs. A model-checking approach is, in principle, more powerful. However, model checking parallel programs in general tend to suffer from the state explosion problem. For partitioned global address space (PGAS) programs, this problem is even worse, as these programs allow concurrent and/or parallel execution as well as memory sharing. This restricts the use of model checking only to programs with limited functionality or limited use of the features of the parallelization paradigm. Runtime error detection is the most practical approach for correctness checking tools, since it can be deployed transparently and avoids the potentially exponential analysis time of model checking an entire application.

Exascale Programming Paradigms in MYX

A long-term challenge is to evolve existing programming models, and develop new ones to better support application development on exascale machines. For different domains and different abstraction levels, various programming models have gained momentum. While there is ongoing research on how to make the currently predominant HPC programming model—namely MPI+X— scale well in such systems, the emerging and more high-level PGAS programming models have shown to deliver high productivity for researchers using certain types of codes. With the multi-level programming paradigm FP3C, programmers can express high-level parallelism in the YML workflow language and employ parallel components written in the XcalableMP paradigm.

XcalableMP [1], or XMP for short, is a directive-based language extension which allows

users to develop parallel programs for large-scale systems in a highly productive manner and to tune the performance by having minimal amounts of simple notation. It is a PGAS language specified by Japan's PC Cluster Consortium for high-level programming and the main research vehicle for Japan's post-petascale programming model research targeting exascale. XMP employs a SPMD approach and provides both a global-view model and a local-view model—the first targeting node-level parallelism similar to OpenMP, and the latter offering a PGAS programming model.

The YvetteML workflow language [2], YML, is used to describe the parallelism of an application at a very high level. YML provides a compiler to translate the YvetteML notation into MPI- or XMP-parallel programs, and a just-in-time scheduler to manage the execution of parallel programs. It hides the rather low-level communication details from the programmer, particularly when coupling complex applications.

MUST Correctness Analysis for XMP

By exploiting the MPI profiling interface to record the messages between processes, the MUST [3] runtime correctness checking tool currently can detect a wide range of issues in MPI. It does so by collecting program information and aggregating this in a tree-based overlay network (TBON) capable of running different types of analysis, all at program runtime in a very scalable manner.

Runtime tools often need knowledge coming from the parallel runtime system. MPI, with the rich set of API functions, provides with its profiling interface an interception layer that allows tools to easily intercept MPI API calls. OpenMP provides the parallel functionality in the form of compiler directives. To gather information about the parallel program and finally analyze the program, a tool needs help by either the OpenMP compiler or the OpenMP runtime system. OpenMP just recently added a tools interface that provides the necessary information for runtime tools.

XMP is very similar to OpenMP in the way that parallelism is also expressed by directives. Therefore, it makes sense to define the tools interface for XMP (XMPT) similarly to the OpenMP tools interface. A tool can register callback functions invoked by the parallel runtime at certain defined events, such as the begin and end of a loop. Based on these events, the runtime tool can apply analyses, in the case of MUST, for example, data race analysis. Fig. 1

```
#pragma xmp nodes p(8)
int main(int argc, char** argv){
  xmp_init(&argc, &argv);
#pragma xmp task on p(5:8)
  { do_some_stuff();
#pragma xmp task on p(1)
    {
      printf("PASS\n");
    }
  do_some_other_stuff(); }
  return 0;
}
```

Fig. 01: The XMP code example testcase will never print PASS. The outer task region makes only nodes 5-8 active. The inner task region tries to make only node 1 active, but the node never reaches the code. This program therefore violates against the XMP specification and a runtime tool can detect this semantic error. The color map on the right shows the active nodes for each line.

shows a simple incorrect XMP program. This code also shows a main advantage of XMP over MPI: The programmer can express parallel semantics by using the xmp pragma directives. For an MPI program with similar semantics and result, one would replace the task regions with if-blocks to filter the MPI ranks. In the MPI case, a runtime tool would not understand the issue, and the MPI program would not even violate against the MPI standard.

In a first step, we identified how a programmer might violate the constraints provided by XMP. This analysis also decided whether such a violation should be analyzed at compile time or runtime. In the next step, we defined and implemented XMPT in a way that provides all necessary information to apply the necessary runtime analysis. For a broader applicability of the interface, we developed XMPT in dialogue with the ScoreP team. In the current phase, we are implementing MUST analyses on base of the knowledge provided by the XMPT interface.

MUST Correctness Analysis with YML

The YvetteML workflow language provides a mean to describe high level work and data flow. Each YML task might be an independent application, which is wrapped by some YML initialization and communication code. The high-level Yvette description can be analyzed statically by the Yvette compiler. For the programmer it is also important that the individual YML task executes correct. Therefore, we provide a workflow which allows to apply MUST on a per-task basis as shown in Fig. 2.

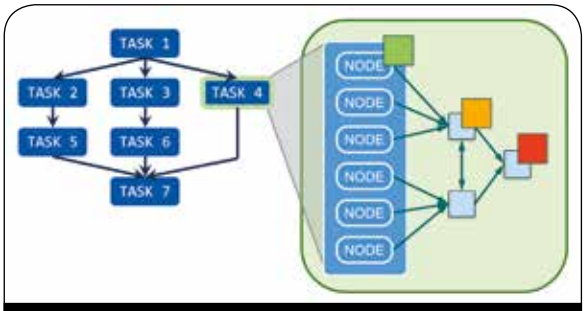

Fig. 02: The YML scheduler starts a specific YML task under the control of MUST. Therefore, the scheduler launches additional processes which are used by the tool for distributed and global analysis. The programmer can specify on a per-task base which parts of the execution should be analyzed by the tool.

References

[1] http://www.xcalablemp.org/

[2] http://yml.prism.uvsq.fr/

[3] https://www.itc.rwth-aachen.de/MUST

Written by Joachim Protze

RWTH Aachen

Contact: protze@itc.rwth-aachen.de

MoeWE Project Opens Supercomputing-Akademie

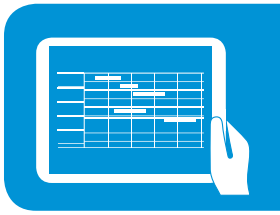

In April 2018, the project MoeWE (Modular training as experts in high-performance computing) opened its Supercomputing-Akademie – a new continuing education program for professionals – with the first module "parallel programming".

Based on HLRS' interactions with industrial partners, there is a clear need for expertise in IT areas like simulation, modeling, and parallel programming. The four cooperation partners, HLRS, SICOS BW, and the Universities of Ulm and Freiburg are helping fill this void with the project MoeWE and founded the Supercomputing-Akademie.

The program uses a "blended learning" approach, which combines classroom teaching and online learning. In order to help IT professionals become HPC-experts, the MoeWE concept consists of several modules:

- Parallel Programming
- Simulation
- Visualization
- Optimization
- Cluster, Cloud, and High-Performance-Computing
- Business Issues, Ecology, and Economy
- Data Management
- Introduction to IT.

Since November 2016, the project team has been working on the modular and occupational training concept. Among the key themes, the MoeWE team hopes to facilitate knowledge transfer between HLRS experts and industry professionals. For the module "parallel programming," the project team was able to recruit several HLRS experts to serve as lecturers in order to help shape the module. The training program focuses on code development on parallel systems, architecture and programming models such as MPI and OpenMP.

In April 2018, the Supercomputing-Akademie started with 17 participants from academia and industry. The group continues its studies online over a period of 15 weeks, with interactive learning scripts, animations, programming exercises, and explanatory videos. Lecturers field questions during virtual meetings. In July 2018, the module will end with a test, with those completing the program receiving a certificate confirming participants' knowledge in parallel programming.

The MoeWE team is looking forward to providing further modules for the Supercomputing-Akademie, such as simulation and optimization.

The project is funded by the European Social Fund (ESF) and the Ministry of Science, Research and the Arts of the state of Baden-Württemberg.

Written by Hanna Skubski
High Performance Computing Centre Stuttgart (HLRS)
Contact: skubski@hlrs.de

Centers / Syste

ems / Trainings

In this section you will find an overview about the upcoming training program and information about the members of GCS.

The Leibniz Supercomputing Centre of the Bavarian Academy of Sciences and Humanities (Leibniz-Rechenzentrum, LRZ) provides comprehensive services to scientific and academic communities by:

- Giving general IT services to more than 100,000 university customers in Munich and for the Bavarian Academy of Sciences
- Running and managing the powerful communication infrastructure of the Munich Scientific Network (MWN)
- Acting as a competence centre for data communication networks
- Being a centre for large-scale archiving and backup
- Providing high-performance computing resources, training and support on the local, regional, national and international level.

Research in HPC is carried out in collaboration with the distributed, statewide Competence Network for Technical and Scientific High Performance Computing in Bavaria (KONWIHR).

Picture of the Petascale system SuperMUC at the Leibniz Supercomputing Centre.

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 Phase 1" IBM System x	IBM iDataPlex 9216 nodes, 147456 cores, 288 TByte, FDR10	3,185	Capability Computing	
	205 nodes, 8,200 cores Westmere EX 52 TByte, QDR	78	Capability Computing	
	32 accelerated nodes Knights Corner 76 GByte, FDR14	100	Prototype System	
"SuperMUC Phase 2" Lenovo Nextscale	3072 nodes, 86016 cores, Haswell EP 197 TByte, FDR 14 IB	3,580	Capability computing	German universities and research institutes, PRACE (Tier-0 System)
"CooLMUC2" Lenovo Nextscale	252 nodes, 7,056 cores Haswell EP 16.1 TByte, FDR 14 IB	270	Capability computing	Bavarian Universities (Tier-2)
"CoolMUC3" Megware Slide SX	148 nodes, 9472 cores, Knights Landing, 14.2 TByte, Omnipath	383	Capability Computing	Bavarian Universities (Tier-2)
Compute Cloud Linux-Cluster	200 nodes, 2700 cores	18	Capability Computing	Bavarian Universities, LCG Grid

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

H L R S

First German National Center

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.

View of the HLRS Cray XC40 "Hazel Hen"

Höchstleistungsrechenzentrum Stuttgart (HLRS), Universität 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

Bundling competencies

In order to bundle service resources in the state of Baden-Württemberg HLRS has teamed up with the Steinbuch Center 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 centers 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.

Compute servers currently operated by HLRS

System	Size	Peak Performance (TFlop/s)	Purpose	User Community
Cray XC40 "Hazel Hen"	7,712 nodes 185,088 cores 1 PB memory	7,420	Capability Computing	European (PRACE) and German Research Organizations and Industry
NEC Cluster (Laki, Laki2) heterogenous compunting platform of 2 independent clusters	911 nodes 9,988 cores 23 TB memory	170	Laki: 120,5 TFlops Laki2: 47,2 TFlops	German Universities, Research Institutes and Industry
NEC SX-ACE	64 nodes 256 cores 4 TB memory	16	Vector Computing	German Universities, Research Institutes and Industry

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

JÜLICH Forschungszentrum | JÜLICH SUPERCOMPUTING CENTRE

he 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..

The Cluster module of JSC's Modular Supercomputer "JUWELS" (artistic rendering).

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 cooperation with neighbouring universities.

Implementation of strategic support infrastructures including community-oriented simulation laboratories and cross-sectional teams, e.g. on mathematical methods and algorithms and parallel performance tools, enabling the effective usage of the supercomputer resources.

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
Atos BullSequana X1000 Cluster "JUWELS"	10 cells, 2,559 nodes 122,448 cores Intel Skylake 192 graphics processors (NVIDIA V100) 273 Tbyte memory	12,001	Capability Computing	European (PRACE) and German Universities and Research Institutes
T-Platforms Cluster + Intel/Dell Booster "JURECA"	Cluster: 1,884 nodes 45,216 cores Intel Haswell 150 graphics processors (NVIDIA K80) 281 TByte memory	2,245	Capacity and Capability Computing	German Universities, Research Institutes and Industry
	Booster: 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

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

Timetable
High Performance Computing
Courses and Tutorials

Course / Workshop Title	Location	Date
Concepts of GASPI and Interoperability with other communication APIs (PRACE)	Stuttgart	Jul 2-3, 2018
Introduction to UPC and Co-Array Fortran (PRACE)	Stuttgart	Jul 5-6, 2018
Advanced C++, Focus on Software Engineering	Stuttgart	Jul 10-13, 2018
Introduction to parallel programming with MPI and OpenMP	Jülich	Aug 13-17, 2018
Parallel Programming with MPI / OpenMP	Zürich	Aug 20-23, 2018
Iterative Linear Solvers and Parallelization	Garching	Sep 10-14, 2018
Introduction to Computational Fluid Dynamics	Stuttgart	Sep 10-14, 2018
Advanced Fortran Topics (PRACE)	Garching	Sep 17-21, 2018
CFD with OpenFOAM®	Siegen	Sep 24-28, 2018
Introduction to Cluster Filesystems	Stuttgart	Sep 25, 2018
Porting code from Matlab to Python	Jülich	Oct 8-9, 2018
C Language for Beginners	Garching	Oct 9-11, 2018
Parallel Programming Workshop (MPI, OpenMP and advanced topics) (PRACE)	Stuttgart	Oct 15-19, 2018
Advanced C++ with Focus on Software Engineering	Garching	Oct 23-25, 2018
Scientific Visualization	Stuttgart	Oct 25-26, 2018
Introduction to GPU programming using OpenACC	Jülich	Oct 29-30, 2018
Cray XC40 Workshop on Scaling and Node-Level Performance	Stuttgart	Nov 5-9, 2018 (tbc)
C++ Language for Beginners	Garching	Nov 6-9, 2018
Software Development in Science	Jülich	Nov 19-20, 2018
Advanced C++ with Focus on Software Engineering	Stuttgart	Nov 19-22, 2018
Introduction to the Programming and Usage of the Supercomputer Resources at Jülich	Jülich	Nov 22-23, 2018
Advanced Parallel Programming with MPI and OpenMP	Jülich	Nov 26-28, 2018
Fortran for Scientific Computing	Stuttgart	Dec 3-7, 2018 (tbc)
Parallel Programming with HPX (PRACE)	Stuttgart	(tba)

Visit InSiDE Online for Details

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

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

Get the Latest News about Innovative High-Performance Computing in Germany

Researchers from the Heidelberg Institute for Theoretical Studies, Max Planck Institutes for Astrophysics and Astronomy, Harvard University, and the Center for Computational Astrophysics in New York have used supercomputing to develop a more complete picture of galaxy formation. The composite image on the cover combines the gas temperature (as the color) and shock mach number (as the brightness). Red indicates gas at 10 million degrees Kelvin at the centre of massive galaxy clusters, while the bright structures show diffuse gas from the intergalactic medium shock heating at the boundary between cosmic voids and filaments.

GCS
Gauss Centre for Supercomputing