


Precision Diagnostics

How 3D Modelling Could Improve Ear, Nose, and Throat Surgery



Smart Scale Strategy

Germany plans for its high-performance computing future

Making Supercomputing Accessible

Forschungszentrum Jülich celebrates 30 years supporting advanced research

Publishers



Prof. Dr. D. Kranzlmüller
Director, Leibniz Supercomputing
Centre



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



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

Staff Writers



Lena Bühler
Contact: buehler@hlrs.de

Lena Bühler is a communications manager and responsible for project communication at the High-Performance Computing Center Stuttgart.



Eric Gedenk
Contact: gedenk@hlrs.de

Eric Gedenk is a science writer for the Gauss Centre for Supercomputing. He is based at the High-Performance Computing Center Stuttgart.



Christopher Williams
Contact: williams@hlrs.de

Christopher Williams is a scientific editor at the High-Performance Computing Center Stuttgart.

© HLRS 2017

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

Managing Editor
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

Welcome to this new issue of InSiDE, the journal on Innovative Supercomputing in Germany published by the Gauss Centre for Supercomputing (GCS). This time InSiDE is full of celebrations and events, the most important of which being the announcement of a German smart scale strategy by Prof. Lukas from the Federal Ministry of Science and Education (BMBF). The announcement was made during ISC 2017 in Frankfurt. Here we cover not only his announcement but also its consequences for GCS.

As always, GCS showed its research and technology at ISC in a common booth showcasing all three centers. This year was special, however, in that GCS also celebrated its 10th anniversary. Together with friends and colleagues from all over the world we heard a short speech and enjoyed a long evening discussing the history of GCS and future trends in HPC.

Another highlight of the last six months was the celebration of 30 years of the John von Neumann-Institute for Computing (NIC) and High-Performance Computing Center (HLRZ) at Jülich. With a tradition of much longer than 30 years, the HLRZ was the first step toward a nationally accessible supercomputer in Germany. In this way Jülich was preparing the path to the national supercomputing center strategy currently implemented by GCS.

This year also saw the opening of a new training building at the High-Performance Computing Center Stuttgart (HLRS) in July. The facility

will enable HLRS to better meet the requirements of its international training activities.

This issue of InSiDE presents a variety of applications using GCS supercomputers over recent months. Their variety shows that GCS systems are able to cover a wide spectrum of disciplines, making supercomputing available for all fields that rely on compute performance to break new barriers and make new discoveries. Exciting results are presented, for example, for the simulation of nasal cavity flows, which could enable better understanding of the medical implications of surgical treatment.

The Projects section shows how much GCS and its centers are enmeshed in the European and German scientific community when it comes to research in the field of high-performance computing. We find a variety of project reports that range from basic research to applied research and industrial development.

As usual, this issue also includes information about events in supercomputing in Germany over the last months and gives an overview of workshops in the field. Readers are invited to participate in these workshops, which are part of the PRACE Advanced Training Center and hence open to all European researchers.

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

1 Editorial

News Events

6 German Federal Ministry of Research and Education Commits to Investments in "Smart Scale"

7 Germany's Strategy for "Smart Exascale"

9 GCS Celebrates a Decade of Computing Excellence with Strong Showing at ISC17

12 Festive Colloquium Marks the 30th Anniversary of the HLRZ and NIC

14 HLRS Celebrates Opening of New HPC Training Building

18 State Minister Announces New State HPC Strategy, Praises HLRS as Key Component

20 On the Heels of an HPC Innovation Excellence Award, "Maya the Bee" Keeps Flying

22 HLRS Researchers Support Design of MULTI Elevator System

27 GPU Hackathon in Jülich

29 Summer School on Fire Dynamics Modelling 2017

31 Jülich Supercomputing Centre starts deployment of a Booster for JURECA

34 Two New Research Groups Established at the John von Neumann Institute for Computing in Jülich

36 Quantum Annealing & Its Applications for Simulation in Science & Industry, ISC 2017

40 The Virtual Institute—High-Productivity Supercomputing Celebrates its 10th Anniversary

42 Hazel Hen's Millionth Compute Job

46 Summer School Focuses on Links between Computer Simulation and the Social Sciences

48 FAU Students Win Highest Linpack Award at ISC17's Student Cluster Competition

51 Changing Gear for Accelerating Deep Learning: First-Year Operation Experience with DGX-1

55 Intel MIC Programming & HPC for Natural Hazard Assessment and Disaster Mitigation Workshops @ LRZ

58 New PATC Course: HPC Code Optimisation Workshop @ LRZ

59 OpenHPC

63 Scaling Workshop 2017: Emergent Applications

Applications


- 66 FS3D – A DNS Code for Multiphase Flows
- 72 Distribution Amplitudes for η and η' Mesons
- 77 Performance Optimization of a Multi-resolution Compressible Flow Solver
- 82 Performance Evaluation of a Parallel HDF5 Implementation to Improve the Scalability of the CFD Software Package MGLET
- 86 ECHO-3DHPC: Relativistic Magnetized Disks Accreting onto Black Holes
- 89 Environmental Computing at LRZ
- 116 HLRS Researchers Do the Math for Future Energy Savings
- 118 Fortissimo 2 Uses HPC and Machine Learning to Improve Self-Driving Vehicle Safety
- 120 Fostering Pan-European Collaborative Research: 3rd Edition of the “HPC-Europa” Project Launched
- 122 HLRS Scholarships Open Up “Simulated Worlds” to Students
- 124 NOMAD: The Novel Materials Discovery Laboratory
- 128 The ClimEx project: Digging into Natural Climate Variability and Extreme Events

Projects

- 94 PPI4HPC: European Joint Procurement of Supercomputers Launched
- 96 DEEP-EST: A Modular Supercomputer for HPC and High Performance Data Analytics
- 101 Helmholtz Analytics Framework
- 106 Rhinodiagnost: Morphological and Functional Precision Diagnostics of Nasal Cavities
- 110 DASH – Distributed Data Structures in a Global Address Space

Centers Systems Trainings

- 136 Leibniz Supercomputing Centre (LRZ)
- 138 Höchstleistungsrechenzentrum Stuttgart (HLRS)
- 140 Jülich Supercomputing Centre (JSC)
- 142 Timetable HPC Courses and Tutorials
- 143 Visit InSiDE Online for Details



News/Events

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

German Federal Ministry of Research and Education Commits to Investments in “Smart Scale”

MinDr. Prof. Dr. Wolf-Dieter Lukas, Head of the Key Technologies Unit at the German Federal Ministry of Research and Education (BMBF), lauded the Gauss Centre of Supercomputing (GCS) as one of Germany’s great scientific success stories, and indicated that the federal government would be investing deeper in high-performance computing (HPC).

Speaking at ISC17, the 32nd annual international conference on HPC, Lukas reaffirmed the German government’s commitment to funding HPC, and indicated that were GCS not to have been created 10 years ago, there would be a significant need to create it today. “GCS is a good example that shows investment in research pays off,” he said.

In describing the next decade of funding for GCS, Lukas emphasized that German supercomputing would be focused on “smart scale” along its path toward exascale computing—a thousand-fold increase in computing power over current-generation petascale machines, capable of at least 1 quadrillion calculations per second.

“GCS is about smart scale; it isn’t only about computers, but computing,” he said. GCS’s smart exascale efforts are funded through the BMBF’s smart scale initiative.

In addition to new supercomputers at each of the three GCS centres, GCS also plans to invest heavily in furthering education and training programs.



MinDr. Prof. Dr. Wolf-Dieter Lukas, Head of the Key Technologies Unit at the German Federal Ministry of Research and Education announces funding for the next decade of education.

© Prof. Dieter Kranzlmüller, LRZ

While Lukas acknowledged the need to develop exascale computing resources in Germany, he indicated that the government wanted to fund initiatives that would enable researchers to make the best possible use of supercomputers. He also emphasized that GCS will continue to support German and European researchers through excellence-based, competitive grants for access to diverse supercomputing architectures at the three GCS centres. Lukas emphasized that supporting GCS made a major impact in supporting German and European researchers’ HPC needs. In general, the BMBF aims to increase research investment from 3% to 3.5% of German GDP by 2025.

Written by Eric Gedenk

Germany's Strategy for "Smart Exascale"



On June 19, MinDir. Prof. Dr. Wolf-Dieter Lukas, Head of the Key Technologies Unit at the German Federal Ministry for Education and Research (BMBF), announced Germany's smart scale strategy in HPC for the coming years. With this announcement, he set the agenda for GCS and its partners for their ongoing discussions and future planning.

The vision of the Gauss Centre for Supercomputing (GCS), founded in 2007, was always to unite Germany's premier supercomputing facilities toward one goal—delivering competitive, world-class high-performance computing (HPC) resources to users so they could solve some of the world's most challenging problems. Over the last decade, the three GCS centres were successful in achieving this goal by consistently putting systems in the top 10 of the top500 list. At the same time they integrated their expertise, resources, and manpower to become Europe's premier HPC centre.

GCS users in engineering have helped design safer cars and more efficient power plants. Climate scientists have gone from simulating past weather events to using their models to predict the Earth's climate moving forward. Physicists have made gigantic strides in understanding the universe's most basic building blocks. Biologists have helped develop more useful medications and moved toward personalizing knee replacement and hip replacement surgeries. These are just a few of the research areas our resources have supported.

Today the HPC ecosystem, always rapidly evolving, is entering an entirely new and disruptive phase. Computer scientists see computing speed constrained not only by the amount of processors running in parallel and how much power is available to a facility, but also how quickly processors can talk to one another, how well users can create codes to take full advantage of new, increasingly complex systems, and, perhaps most importantly, how to keep the environmental footprint of supercomputing sustainable long into the future. The HPC community has to find creative ways to solve these problems if we are to get to exascale computing—a thousand-fold increase over current-generation petascale machines.

When GCS received its newest round of federal and state funding in 2017—securing our place as the European HPC leader for another decade—GCS decided to dedicate itself to not just building newer, faster machines, but also taking a leading role in solving these challenges at the frontline of supercomputing, strengthening our role as an end-to-end service provider.

Based on GCS's "smart exascale," approach, GCS is going to continue to increase the computing power at the three GCS centres with new machines for the next two generations. GCS will make sure that our users are equipped to take full advantage of each next-generation machine and the new challenges that go along with them.

To that end, GCS centres will be hiring 30 new scientific staff members to support the mission.

The centres are redesigning their user support models, bringing in scientific domain experts to support users on a scientific level. At the same time, HPC experts continue to work on continually improving hardware and software support while working with vendors to develop new technologies.

GCS will foster—and incentivize—deeper teamwork between numerical methods experts, computer scientists, and computational domain scientists, ultimately nurturing much closer teamwork in the scientific community. If we see more of the GCS centre scientific staff members as co-authors on users' scientific papers, we will know we are closer achieving the results we are hoping for as it relates to collaboration. Recently, the three centres have held exascale workshops and founded the Hi-Q club to foster these kinds of collaborations, and researchers are already seeing major speedups in their codes' performance.

In order to spread HPC knowledge and to help domain scientists to better use scalable systems, GCS is further expanding its robust training programs. By doing so, HPC experts and relative newcomers to HPC alike can learn how supercomputing can increase productivity in their research lab, multinational enterprise, or small business. Due to the increasing diversity of computing architectures, GCS will also train users on how to develop applications in a flexible, portable way so they are prepared to take advantage of the wealth of smart exascale HPC architectures.



Dr. Claus Axel Müller is the managing director for the Gauss Centre for Supercomputing.

While the challenges of reaching exascale are real, we at GCS feel that we will deliver not only exascale systems, but also stable, environmentally sustainable systems that the users understand and make the best use of from day one. By delivering on these promises, GCS will turn Prof. Lukas' announcement into a productive and successful strategy while simultaneously helping both German research and German industry stay competitive in a rapidly changing world.

Written by Dr. Claus Axel Müller

Contact: claus-axel.mueller@gauss-centre.eu

GCS Celebrates a Decade of Computing Excellence with Strong Showing at ISC17



Five years ago, the three GCS centres began attending the ISC High Performance conference (ISC) under the GCS banner.

At this year's conference—which ran June 18–22 in Frankfurt, Germany—GCS staff provided a wide range of interactive exhibits, including videos about ongoing GCS research, augmented and virtual reality presentations, and demonstrations of cutting-edge monitoring and visualization programs.

Using a motorcycle, HLRS staff demonstrated an augmented reality air flow simulation. JSC staff showcased a live demo of its self-developed system monitoring tool, Llview. Staff at LRZ presented a visualization of convection currents within the Earth's mantle using a head-mounted display. All the centres also presented videos and gave talks highlighting simulations and visualizations done at the three centres, the centres' commitment to energy efficiency, and how the three centres support users' applications.

Beyond the exhibition and GCS centres' staffs' regular active participation in panel discussions, birds of a feather sessions, tutorials, and session chair roles they regularly participate in during ISC, GCS employees also came together to celebrate 10 years of leading-edge German supercomputing.


During the conference's opening session on Monday, June 19, MinDir. Prof. Dr. Wolf-Dieter Lukas, Head of Key Technologies at the German

Federal Ministry of Research and Education (BMBF), announced the German strategy for the next decade of high-performance computing, and noted that GCS was an essential organization in accomplishing the German "smart scale" mission. The BMBF sees its "smart scale" strategy as increasing investment in not only computing power, but also the software advances and training needed to make the best use possible of next-generation machines.

"GCS is a good example that shows investment in research pays off," he said. GCS's primary funding agencies, the BMBF and the science ministries of the states of Baden-Württemberg, Bavaria, and North Rhein-Westphalia, acknowledged GCS's last decade of success by funding it to the tune of nearly €500 million for the next 9 years.

GCS Managing Director Dr. Claus-Axel Müller and GCS Board of Directors Chair Prof. Michael M. Resch also joined Lukas on stage and presented the Gauss Award, given to the conference's best technical paper, to a team from Boston University (BU) and Sandia National Laboratories in the United States. The award, accepted by team member and BU doctoral candidate Emre Ates, was given for the team's paper, "Diagnosing Performance Variations in HPC Applications Using Machine Learning."

On Tuesday, Ates presented the team's paper during a 30-minute session, describing how the team built a framework using machine learning that can help diagnose performance anomalies



On Tuesday, GCS staff invited press, users, and friends to celebrate 10 years of GCS.

for applications running on supercomputing resources, ultimately helping researchers improve code performance.

Toward the end of Tuesday's exhibition, GCS staff invited press, users, friends of the GCS centres, and European partners to toast 10 years of GCS and hear Director Resch speak about the next 10 years. The catered event was well-attended, and the food also supported one

of GCS's major ISC-related investments—the undergraduate team it sponsored in the Student Cluster Competition.

Student Cluster Competitions take place at the world's largest supercomputing conferences in Europe, North America, and Asia, and pit undergraduate teams against one another in a fast-paced, multi-day event that involves building a cluster computer and using it to run a variety



Leadership from GCS and BMBF present the Gauss Award at the opening session at ISC17. From Left: GCS Managing Director Claus-Axel Müller, BMBF Head of Key Technologies MinDir. Prof. Dr. Wolf-Dieter Lukas, award recipient and Boston University doctoral candidate Emre Ates, and GCS Board of Directors chair Michael Resch. © Philip Loeper for ISC

of challenging applications while keeping to a very strict power limit. Teams are graded on speed through the high performance Linpack (HPL) benchmark, accuracy, energy efficiency, and application knowledge. At the end of the conferences, teams are given awards in various categories.

This year, GCS sponsored a team from Friedrich-Alexander-University Erlangen-Nürnberg (FAU). The 6-student team, FAU Boyzz, comprised of computational engineering, computer science, and medical engineering students, won the HPL award by achieving peak performance of 37.05 teraflops (1 teraflop equals 1 trillion

floating point operations per second)—nearly tripling the result from last year's ISC and setting a new record for HPL in the competition.

"The Gauss Centre for Supercomputing, who by definition is highly interested in drawing young talent's attention toward High Performance Computing, is always open to support up and coming HPC talent, also in the framework of these kinds of events," explains Claus Axel Müller, Managing Director of the GCS. "We are well aware of the financial constraints students are facing when trying to participate in international competition, especially if travel

and related expenses are involved. Thus we are happy to be of help, and we would like to sincerely congratulate the FAU Boyzz for their great achievements at ISC."

At this year's ISC, GCS also released a new brochure in both English and German, *Computing for a Better Tomorrow*.

Written by Eric Gedenk

Festive Colloquium Marks the 30th Anniversary of the HLRZ and NIC

The scientific community has historically struggled with a large variety of numerical problems in fields such as fluid dynamics, quantum mechanics, and many more. The advent of supercomputers in the 1970s led to major steps forward in solving these problems. However, these resources were not readily available. This changed in 1987, when the Höchstleistungsrechenzentrum (HLRZ)–the forerunner of today’s John von Neumann Institute for Computing (NIC)–was founded at the Forschungszentrum Jülich (FZJ). This made supercomputing far more accessible to a broad community. Thirty years later, FZJ can celebrate NIC’s 30th anniversary by commemorating the centre’s contribution to the field of computational science.

It all began a few years earlier, when Prof. Friedel Hoßfeld, who was heading the Zentralinstitut für Angewandte Mathematik (ZAM)

in Jülich at the time, began to campaign decisively for the use of the then newly emerging supercomputers to be used in natural sciences and technology, taking an early lead and acting as a visionary in this new field. He argued that next to the proven methods of theoretical models and experimental verification, the new tool of simulation should advance sufficiently to become a third source of scientific knowledge and insight. Naturally, he always insisted on the latest and most powerful computers available for a given period. As such, the era of supercomputing in Jülich began in 1983 with the installation and use of the vector processor Cray X-MP22, which was the world’s fastest computer at the time.

From the mid-80s onwards, Prof. Hoßfeld also began to engage himself increasingly in science management. He came to be a member of the commissions of the scientific council, which developed recommendations for the development of high-performance computing (HPC) in Germany. An important step was the founding of the HLRZ in 1987 as a joint venture between the FZJ, the Deutsche Elektronen-Synchrotron (DESY) and the Gesellschaft für Mathematik und Datenverarbeitung (GMD). The HLRZ was the first institution in Germany to offer supercomputer capacity together with the ZAM as well as consulting and support for the use of the machines on a national level. After the GMD left the venture, the FZJ and DESY confirmed their commitment and founded the NIC as a successor to the HLRZ in 1998. Later, the Gesellschaft für Schwerionenforschung (GSI) joined this



Prof. Binder (right), the chairman of the NIC Scientific Council congratulates Prof. Hoßfeld.



Prof. Lippert, the head of the JSC, gives a presentation of the current state of today's supercomputing.

was a PhD student of Prof. Hoßfeld and JSC staff member in the 90s, looked back at the origins of parallel computing. Prof. Thomas Lippert, the head of JSC, discussed the current state of today's supercomputing, and the evolving possibilities in neural networks and deep learning. Looking to the future, Prof. Hans

newly formed institute. The original ZAM went on to become today's Jülich Supercomputing Centre (JSC), while Prof. Hoßfeld became synonymous with HPC in Germany, most notably in Jülich itself. Prof. Hoßfeld retired in 2002.

De Raedt from the University of Groningen gave an outlook to the coming revolution of quantum computing and the associated challenges and opportunities.

On September 1, 2017, JSC hosted HLRZ's 30th anniversary celebration—together with Prof. Hoßfeld's 80th birthday—with a festive colloquium. The event provided a welcome opportunity to look back on some remarkable achievements and highlights made possible by simulation on high-performance computers in recent years. After a warm welcome by Prof. Sebastian Schmidt, member of the board of directors of FZJ, Prof. Kurt Binder, the chairman of the NIC Scientific Council, presented the answers to complex questions arising in the field of soft matter physics—a research area in which supercomputing has provided insights. Prof. Wolfgang Nagel from TU Dresden, who

The staff and users of today's JSC and NIC would like to sincerely thank Prof. Hoßfeld for laying the foundations of these institutions and for all his ground-breaking and pioneering work. We look forward to the coming innovations with excitement and will always fondly remember the first steps into the then new field of HPC.

Written by Alexander Trautmann

NIC Coordination Office at the Jülich Supercomputing Centre (JSC)

Contact: coordination-office@fz-juelich.de

HLRS Celebrates Opening of New HPC Training Building



The new HLRS training center offers a modern, state-of-the-art learning environment.

© Simon Sommer

High-ranking representatives from the state of Baden-Württemberg, the city of Stuttgart, the University of Stuttgart, and invited guests gathered to mark the beginning of an important new chapter in HLRS's development.

In addition to conducting research and providing high-performance computing (HPC) services to academic and industrial users, the High-Performance Computing Center Stuttgart (HLRS) is a European leader in training scientists, programmers, and system developers in the use of HPC. As simulation and data analytics become increasingly essential tools across many fields of research and technology, the need is greater than ever for people trained with such skills.

To address this demand, HLRS recently undertook a major expansion, opening a new building that is making it possible to provide HPC training to more visitors than ever. Funded for a sum of €6.8 million by HLRS, the approximately 1000 square-meter facility opened its doors on March 7. The building offers a state-of-the-art lecture hall, smaller seminar rooms, and space to expand HLRS staff. It will enable HLRS to expand its continuing professional training and support other educational symposia for students and the general public.

At an inauguration ceremony on July 14, 2017, Annette Ipach-Öhmann (Director, Baden-Württemberg State Office of Real Estate and Construction), Gisela Splett (State Secretary,



Students from the Stuttgart Academy of Music and Performing Arts opened the inauguration ceremony of the new HLRS training center.
© Zimmermann Visuelle Kommunikation, Stuttgart

Baden-Württemberg Ministry of Finance), Hans Reiter (Baden-Württemberg Ministry of Science, Research, and Art), Isabel Fezer (Minister for Youth and Education, City of Stuttgart), and Simone Rehm (Chief Information Officer, University of Stuttgart) delivered a series of speeches celebrating this milestone, spotlighting the important role that the new training center will play in the future of HPC research and education for Stuttgart, Baden-Württemberg, Germany, and the world.

Following a musical greeting performed by students from the Stuttgart Academy of Music and Performing Arts, Annette Ipach-Öhmann began the event by pointing out that the new HLRS training facility is the final pillar in the institute's growth, which began in 2005 with the opening of its main building and went through

expansions in 2011 and 2012. In her welcoming remarks she greeted the other speakers and recognized many key individuals who collaborated on the planning and realization of the new facility, including government representatives, administrators and staff at the University of Stuttgart, as well as the architecture firm Wenzel + Wenzel, Pfefferkorn Engineers, and artistic director Harald F. Müller.

Gisela Splett remarked that although the „manicured understatement“ of the area surrounding HLRS and the University of Stuttgart's Vaihingen campus may be deceiving, the supercomputer that HLRS maintains — presently ranked 17 on the Top500 list — is a global “lighthouse” for advanced science and research. The opening of the new HLRS training facility, she suggested, constitutes an important step in raising

the visibility of HLRS and will enhance its global reach. "This will be a meeting point for a network that will play an important role in future generations for the scientific and industrial competitiveness of Baden-Württemberg and Europe," she said. Commenting on the elegance of the building itself, Splett remarked, "I'm impressed by how peaceful the atmosphere is... It creates an environment that is appropriate to the high standards of what happens here." Splett also pledged that supporting high-performance computing will remain a priority for Baden-Württemberg in the future.

Hans Reiter emphasized the significance of the new training center in the context of digital@BW, the digitalization strategy for research and education that is currently being pursued by the majority political coalition in Baden-Württemberg. Considering the crucial role that data analytics, simulation, and visualization play in key industries and research projects across the region, Reiter argued, "we need a first-class IT infrastructure for high-performance computing so that science and business can remain internationally competitive in the future." For this reason, the state is not only promoting the development of supercomputing infrastructure but is also intensifying its support for HPC users, both in academia and in small to mid-sized businesses. The opening of the new HLRS training center, he suggested, will "both qualitatively and quantitatively improve training opportunities in high-performance computing," and will be further evidence of the University of Stuttgart's place as the leading scientific center

for simulation technology in Germany and in Europe.

Isabel Fezer argued that the opening of the new HLRS training center will benefit not just researchers and industry in the region, but the City of Stuttgart as a whole. "The spirit that has grown and is nurtured here permeates the whole city, improves its reputation, and in the end benefits our children and youth," she enthused. Fezer also celebrated the international networks that develop when professionals from around the world visit institutions like HLRS for training and education. Offering advanced training in topics like HPC, she stated, should spread awareness globally of Stuttgart as a center for advanced research and technology.

Simone Rehm saw in the opening of the new HLRS training center the fulfillment of an important component of its mission as a part of the University of Stuttgart. "You all know the saying 'Do good and talk about it,'" she mused. "Today I'd like to modify the saying and suggest we should 'Do good and teach about it.'" She argued that HLRS should continually seek to disseminate the knowledge that emerges from its research, and that the new training center offers exciting new opportunities to do so.

Professor Michael Resch, Director of HLRS, closed the event by thanking the state, federal, city, and university agencies that contributed to the growth of HLRS over the past several years. He also recognized the support that the University of Stuttgart contributed to the building, the



Gisela Splett (State Secretary, Baden-Württemberg Ministry of Finance) emphasized the importance of high-performance computing for the state's future.
© Zimmermann Visuelle Kommunikation, Stuttgart

university departments involved in managing its physical construction, and the HLRS employees whose efforts enable the institute to conduct its work. He also specifically highlighted the unique and critical contributions of Erika Fischer, who worked on behalf of HLRS to manage many facets of the construction of the new building and to ensure that it met the organization's needs.

Resch highlighted some features of the new building before adjourning the formal program. In a festive mood, the approximately 100 visitors then received tours of HLRS and tested the new building's reception area, exchanging impressions and ideas over light refreshments.

Written by Christopher Williams

State Minister Announces New State HPC Strategy, Praises HLRS as Key Component

During a press conference at HLRS, the Baden-Württemberg Minister for Science, Research, and the Arts announced €500 million to support HPC efforts as part of the state wide digitalization strategy.

At an August 24th visit to the High-Performance Computing Center Stuttgart (HLRS), Theresia Bauer, Baden-Württemberg Minister for Science, Research, and the Arts, pointed to HLRS as a key component in the state's digitalization strategy and pledged €500 million toward further advancing the state's high-performance computing (HPC) capabilities over the next decade.

"Digitalization permeates education, research, and technology usage in almost all areas of study," Bauer said. "Digitalization not only promotes research, it highlights the clear benefits technology has on society and industry."

Bauer continued, noting that the success of the state's digitalization strategy did not just depend on infrastructure—building bigger, faster machines—but also on the cooperation of academia and industry and the ability to train next-generation HPC experts. She pointed to HLRS's partnerships with industrial users as an example of the center's commitment to not only building some of the world's fastest



HLRS Director Michael Resch gives Minister Theresia Bauer a tour of HLRS' computer room after an Aug. 24 press conference.

© Eric Gedenk



supercomputers, but also training researchers from a variety of disciplines in science and engineering how to use these tools.

The HLRS visit offered reporters a chance to speak with both Bauer and HLRS Director Michael Resch about the center's future plans and how this funding will impact both HLRS's and Baden-Württemberg's competitive position as a hotbed of research and innovation. The visit focused primarily on three lesser-known research areas where supercomputing has accelerated innovation or helped solve problems—industry and design, participation in democracy, and security.

Resch contextualized how HLRS's goals ultimately help researchers solve some of the hardest challenges facing humanity. "The overall goal of HLRS is to bring together HPC technologies and the expertise of scientists to solve not only problems in science but also problems that affect the lives of everybody in fields like mobility, energy supply, health, and the environment," he said.

Resch described the statewide HPC infrastructure and explained how companies, from small enterprises to multinational corporations, can apply for time on HLRS resources. Afterwards, he invited the minister and press to tour the HLRS CAVE visualization environment, where the HLRS visualization team presented immersive 3D visualizations of industrial design, city planning models for the Baden-Württemberg city of Herrenberg, and forensics research being

done in cooperation with the State Office of Criminal Investigation (LKA) Baden-Württemberg, as well as a tour of the HLRS computer room.

Written by Eric Gedenh

On the Heels of an HPC Innovation Excellence Award, "Maya the Bee" Keeps Flying



Maya the Bee, made with the help of HLRS resources.

©2017 Studio 100 Media. Studio B Animation.

HLRS' successful collaboration with the studio M.A.R.K. 13 continues with production of "Maya the Bee—The Honey Games"

At this year's ISC conference, Hyperion Research presented its annual High-Performance Computing (HPC) Innovation Excellence Award. The award highlights small- and medium-sized businesses' successful use of HPC.

M.A.R.K. 13 and Studio 100 Media took home an award in "HPC User Innovation" for using HPC to calculate about 115,000 computer-generated stereo images (CGI) for the animated film *Die Biene Maja* (*Maya the Bee*) on supercomputing resources at the High Performance Computing

Center Stuttgart (HLRS). Each frame was calculated from the perspectives of both the left and right eye, a prohibitively large computation using traditional computers.

M.A.R.K. 13 began its collaboration in 2013 with HLRS through SICOS BW GmbH, a company founded by the University of Stuttgart and the Karlsruhe Institute of Technology and based at the High Performance Computing Center Stuttgart (HLRS).

SICOS BW helps small- and medium-sized businesses connect with HPC experts to level the competitive playing field with large, multinational companies, training companies' staffs



how to better integrate computation, simulation, and smart data processes into their respective research areas or product development.

By gaining access to HLRS supercomputing resources, companies like M.A.R.K. 13 drastically reduce their production timelines, and are able to use supercomputing resources in a cost-effective way that would be impossible outside the context of a collaborative agreement.

By working with increasingly diverse companies, SICOS BW also benefits by having a stronger portfolio to attract new users across a wider range of industries.

“To be able to support a company like M.A.R.K. 13 is what our work is all about,” said Dr. Andreas Wierse, CEO of SICOS BW. “At the same time, we grow to better understand the difficulties that they face, which ultimately helps us strengthen the support that HLRS can provide to such innovative companies.”

Learning from the different challenges during the production of the first movie, HLRS decided to further engage in the field of media. Since 2015, the Media Solution Center BW has collaborated on projects at the intersection of media, art and HPC.

The partnership looks to continue its success, and recently announced that M.A.R.K. 13 would continue working with HLRS as they produce *Maya the Bee: The Honey Games*. The film is still in its early stages of production and has been

funded with €750,000 from the MFG Film Fund Baden-Württemberg (Medien- und Filmgesellschaft Baden-Württemberg).

Maya the Bee 2: The Honey Games will explore another story of Maya and her friends where Maya is forced to compete in “the honey games” for a chance to save her hive. A teaser trailer for the movie has been released.

Written by Eric Gedenh

HLRS Researchers Support Design of MULTI Elevator System

Using virtual reality and numerical simulation, visualization experts at the High-Performance Computing Center Stuttgart make important contributions to the development of a groundbreaking technology.

In June 2017, the engineering company thyssenkrupp Elevator AG began operating the world's first elevator capable of moving horizontally as well as vertically. Called the MULTI, the new concept mounts elevator cabins on rails instead of

suspending them on cables, offering increased flexibility of movement and exciting opportunities for architects to begin rethinking how large buildings and building complexes are designed. The first fully operational prototype is now running in a specially built tower located in Rottweil, Germany.

Working behind the scenes since 2014, the High-Performance Computing Center Stuttgart (HLRS) played an important role in the MULTI's



A functional, three-dimensional virtual model of the MULTI system at HLRS.

© HLRS



development. Researchers in the HLRS Visualization Department collaborated with thyssenkrupp Elevator engineers and construction managers at Ed. Züblin AG to conduct simulations that tested key features of the new system before it was built. Using numerical simulation as well as virtual reality tools, HLRS made it possible for the engineers to spot design flaws early, assess mechanical configurations and physical stresses that could affect the MULTI's operation, and investigate the experience that its users could expect. The tools that HLRS developed will also support the design of future MULTI installations.

"Combining virtual building models and elevator models makes a lot of sense," says Uwe Wössner, head of the HLRS Visualization Department and leader of HLRS's participation in the project. "I see a great opportunity here."

Virtual Reality used to test design before prototyping

Early in the development process thyssenkrupp Elevator realized that it was important to understand how users would interact with the MULTI system and how this experience would differ from more conventional elevators.

To help address such questions, the HLRS visualization team used data from thyssenkrupp Elevator to create a virtual reality simulation of the elevator system and tower. Once displayed in the CAVE, an immersive three-dimensional virtual reality facility located at HLRS, engineers and architects could interact with the model, moving

through it to gain a sense of how a user might experience the actual elevator and observing its highly complex mechanics in action. The simulation helped the developers identify features in the design that either caused usability problems or that could be improved upon, such as collisions between machine parts in motion that would have been much harder to detect in computer aided design (CAD) software alone.

Working with Züblin, HLRS also integrated its virtual reality simulation into the firm's building information modeling (BIM) strategy. In this case BIM started with a CAD model of the building and elevator system as well as functional data about all of the parts and materials that go into their construction. This made it possible to simulate, for example, whether the elevator would operate properly, stopping at the correct pickup points and moving most efficiently through the structure.

The BIM model also enabled the construction managers to plan the build itself—for example, determining how large components such as motor blocks should be brought into the building during construction and how they should be rotated into position.

During multiple iterations in the development of the elevator and tower models, representatives from thyssenkrupp Elevator and Züblin periodically visited the CAVE to explore the virtual reality model. HLRS also supported the thyssenkrupp Elevator engineers in building a simpler version of the VR facility in their home office


so that they could use virtual modeling as they made improvements during their day-to-day design activities.

Simulating airflow in new elevator shaft geometries

When designing a new elevator, one other factor that developers must consider is how air moves through an elevator shaft. As cabins move passengers from place to place, the air that they pass through has to go somewhere, causing turbulence that can affect the elevator's operation. Two cars passing one another in a shaft, for example, can cause noise and vibrations that

disturb passengers' comfort, generate stresses on machinery, and increase energy consumption.

In addition, planners must also consider what happens when a moving cabin compresses air as it approaches the end of an elevator shaft. Typically, shafts have empty rooms at their ends to accommodate this pressure, but engineers must always optimize the size of this space and the holes through which the air passes to balance the system's technical requirements against the importance of using space and materials most efficiently. In a flexible and adaptable system



Visualization of turbulence surrounding two elevator cars as they pass through a MULTI air shaft.

© HLRS



such as the one that MULTI offers, projecting these kinds of effects on airflow becomes even more challenging.


“Because the whole concept was so new,” says Thomas Obst, an HLRS researcher who investigated these problems, “thyssenkrupp Elevator wanted to know how simulation could help them. We discovered that it could continually inform changes in the system that would lead to improvements.”

Even in conventional vertical elevator systems, airflow dynamics are different in every building. In developing simulations of the airflow in the MULTI system installed in Rottweil, therefore, an additional goal for the team was to develop a method in which new shaft geometries could quickly be imported and simulated.

“Normally every time you lay out a new elevator geometry you need to start again at the beginning and redo a lot of calculations,” says Wössner. “With the simulation approach we used, it’s just a matter of importing the new geometry digitally and running a couple of scripts. It makes it much quicker and thus less expensive than other methods to test new ideas.”

Simulating airflow around large buildings to improve stability

Also important to the construction of the prototype in Rottweil was predicting the effects of airflow around the exterior of the building. Because no one lives or works in offices there, the structure is much thinner than a conventional high-rise, meaning that it is also more susceptible to swaying due to wind. To counteract this force,



HLRS investigators simulated airflow around a model of the Rottweil tower to assist in the design of stabilization measures.
© HLRS

the building incorporates counterweights that dampen the wind's effects.

During the design process, engineers created a small-scale model of the building and tested it in a wind tunnel in order to determine the optimal size and configuration of the dampening system. HLRS then performed numerical simulations and showed that in reality the structure was actually likely to behave differently than what was predicted in the laboratory tests. In follow-up investigations, engineers confirmed that HLRS's simulations were probably more realistic than what was measured in the wind tunnel, though the more conservative experimental model was still valid.

"Once MULTI begins being installed in real buildings," Wössner says, "builders will want to base their plans on accurate models of reality rather than conservative estimates. When they optimize the system, they will need to find the sweet spot where the tolerances properly balance safety needs with the practical need to control costs. The precise values that come from computational modeling will be valuable in achieving this."

Future applications

At the same time that the Rottweil tower began operation, thyssenkrupp Elevator also announced that it had found a customer for MULTI to be installed in a new development being planned in Berlin. The company hopes that this will be the first of many projects that exploit their system's unique capabilities. And

as additional installations accumulate, simulation will have an important ongoing role to play.

"In the future it's going to be a huge advantage to be able to virtually install a MULTI system in a specific building during the development phase," Wössner points out. "This will make it possible for the client to see, long before construction begins, how it operates and to try out different combinations to see what options best meet his or her needs. Does a new building need two shafts, for example, or would three make more sense? You wouldn't be able to answer questions like this by looking at one specific building because the needs will be different every time."

When building physical prototypes becomes prohibitively time-consuming and expensive, simulation and visualization offer powerful tools that can save time and costs, and prevent construction delays. As the MULTI system becomes more widely adopted, HLRS's contributions to its development will continue to be important for its future success.

Written by Christopher Williams



From March 6-10, JSC hosted the first GPU Hackathon of 2017. In this series of events, people come together for a full week to enable scientific applications for GPUs, optimize the performance and parallelise to many GPUs. Intensive mentoring allows application developers to make significant progress on using this promising exascale technology efficiently.

GPU hackathons are a series of events organized by Fernanda Foertter from Oak Ridge National Laboratory [1]. They are hosted by different sites in Europe and the USA. During five days, teams of 3-5 application developers are mentored full-time by two experts. The event is organized such that participants can fully concentrate on their applications. Many of the

experts come from relevant vendors as NVIDIA and PGI as well as from supercomputing centers. The participants thus can expect to have access to most advanced hardware architectures like the nodes of JSC's JURECA cluster that are accelerated by K80 GPUs and the even more advanced OpenPOWER cluster JURON with its NVLink attached P100 GPUs.

All teams made best use of the available time. After spending nine hours at JSC, many continued hacking after returning to their hotels in Jülich. Every day, each team had to present their status and report on challenges during a scrum session. The slides have been made publicly available [2]. At the end of these sessions, new tasks were assigned to the teams. The ability



All participants and mentors that joined the GPU Hackathon at Jülich Supercomputing Centre.

to flexibly provide additional training sessions depending on the needs of the participants are another important part of the concept.

Teams that want to join a GPU hackathon have to submit an application. The applications are reviewed by an international panel. For the Hackathon in Jülich, more good applications were submitted than could be accepted. After making additional efforts in recruiting more mentors, ten teams coming from all over the world were accepted. The applications covered a broad range of science including brain research, lattice QCD, materials science and fluid dynamics. While some teams came with an already mature GPU application and used the event for more in-depth tuning, other teams came without any prior GPU knowledge and worked on their very first steps into realm of GPUs.

When the Hackathon concluded on Friday with final presentations of all teams, everyone found the time well-spent and praised the intense working atmosphere. The closeness to the experts from science (CSCS, JSC, HZDR/MPI-CBG, RWTH) and industry (IBM, NVIDIA, PGI) was held in high regard. By Friday afternoon, over 1,000 jobs were submitted to JURECA and JURON. Four more Hackathons took place throughout 2017. We recommend all interested developers to watch for the announcement of more hackathons to come in 2018.

References

- [1] **Hackatons:**
<https://www.olcf.ornl.gov/training-event/2017-gpu-hackathons/>.
- [2] **INDICO:**
<http://indico-jsc.fz-juelich.de/e/Eurohack17>

Written by Dirk Pleiter

Jülich Supercomputing Centre (JSC)

Contact: d.pleiter@fz-juelich.de

Andreas Herten

Jülich Supercomputing Centre (JSC)

Contact: a.herten@fz-juelich.de

Guido Juckeland

Helmholtz Zentrum Dresden Rossendorf

Contact: g.juckeland@hzdr.de

Summer School on Fire Dynamics Modelling 2017



In the beginning of August 2017, JSC's "Civil Safety and Traffic" division organised a one-week summer school about fire dynamics modelling. Over the last decades, fire modelling became very popular in fire safety engineering and science. As the models evolve, they become more complex, making it harder to understand the underlying principles as well as their application limits. This summer school was intended to educate students and researchers on the fundamental theory and algorithms of fire modelling. The theoretical part was accompanied by practical exercises – mostly with the popular Fire Dynamics Simulator (FDS) – with focus on the discussed models. Besides the presentation

of models, scientific pre- and post-processing, as well as validation methods, were part of the agenda.

This school not only targeted learning more about the underlying numerical models in common fire simulation software, but also enabled participants to get in touch with model developers one-on-one. Since special emphasis was put on scientific work, the participants were mainly PhD students and PostDocs.

The topics covered included an introduction to computational fluid dynamics, turbulence, combustion, thermal radiation, and pyrolysis



Summer school participants, lecturers and local organisation team.

© Ralf-Uwe Limbach, FZJ



Simo Hostikka (Aalto University) teaching the transport of thermal radiation in fires.

modelling. A short introduction to Python allowed the students to learn how to post-process simulation data. The sections were presented by seven lecturers: Simo Hostikka (Aalto University), Bjarne Husted (Lund University), Susanne Kilian (hhp Berlin), Randall McDermott (NIST), Kevin McGrattan (NIST), and Lukas Arnold (JSC).

The 30 participants came mainly from Europe (8 from Germany, 4 from United Kingdom and Poland, 3 from Finland, Italy and Czech Republic, as well as participants from Denmark, Sweden, Hungary, Slovenia and Australia).

The evaluation of the school showed that the participants were satisfied with the organisation and contents. Based on that outcome, the lecturers decided to repeat this event in 2019.



Simulation of a pool fire.

Written by Lukas Arnold

Jülich Supercomputing Centre (JSC)

Contact: l.arnold@fz-juelich.de

Jülich Supercomputing Centre starts deployment of a Booster for JURECA



Since its installation in autumn 2015, the JURECA ("Jülich Research on Exascale Cluster Architectures") system at the Jülich Supercomputing Centre (JSC) has been available as a versatile scientific tool for a broad user community. Now, two years after the production start, an upgrade of the system in autumn 2017 will extend JURECA's reach to new use cases and enable performance and efficiency improvements of current ones. This new "Booster" extension module, utilizing energy-efficient many-core processors, will augment the existing "Cluster" component, based on multi-core processor technology, turning JURECA into the first "Cluster-Booster" production system of its kind.

The "Cluster-Booster" architecture was pioneered and successfully implemented at prototype-level in the EU-funded DEEP and DEEP-ER projects [1], in which JSC has been actively engaged since 2011. It enables users to dynamically utilize capacity and capability computing architectures in one application and optimally leverage the individual strengths of these designs for the execution of sub-portions of, even tightly coupled, workloads. Lowly-scalable application logic can be executed on the Cluster module whereas highly-scalable floating-point intense portions can utilize the Booster module for improved performance and higher energy efficiency.



Fig. 1: The JURECA Cluster module at Jülich Supercomputing Centre.

The JURECA system currently consists of an 1,872-node compute cluster based on Intel "Haswell" E5-2680 v3 processors, NVidia K80 GPU accelerators and a Mellanox 100 Gb/s InfiniBand EDR (Extended Data Rate) interconnect [2]. The system was delivered by the company T-Platforms in 2015 and provides a peak performance of 2.2 PFlops/s. The new Booster module will add 1,640 more compute nodes to JURECA and increase the peak performance by five PFlops/s. Each compute node is equipped with a 68-core Intel Xeon Phi "Knights Landing" 7250-F processor and offers 96 GiB DDR4 main memory connected via six memory lanes and additional 16 GiB of high-bandwidth MCDRAM memory. As indicated by the "-F" suffix, the utilized processor model has an on-package Intel Omni-Path Architecture (OPA) interface which connects the node to the 100 Gb/s OPA network organized in a three-level full-fat tree topology. The Booster, just as the Cluster module, will

connect to JSC's central IBM Spectrum Scale-based JUST ("Jülich Storage") cluster. The storage connection, realized through 26 OPA-Ethernet router nodes, is designed to deliver an I/O bandwidth of up to 200 GB/s. In addition, 198 bridge nodes are deployed as part of the Booster installation. Each bridge node features one 100 Gb/s InfiniBand EDR HCA and one 100 Gb/s OPA HFI, in order to enable a tight coupling of the two modules' high-speed networks. The Booster is installed in 33 racks directly adjacent to the JURECA cluster module in JSC's main machine hall. JSC and Intel Corporation co-designed the system for highest energy efficiency and application scalability. Intel delivers the system with its partner Dell, utilizing Dell's C6320 server design (see Figure 2). The group of partners is joined by the software vendor ParTec, whose ParaStation software is one of the core enablers of the Cluster-Booster architecture. The Cluster and Booster module of JURECA will



Fig. 2: Example of a Dell C6320P server system. The model utilized in the JURECA Booster slightly deviates from the shown version due to the utilized processor type. Copyright: Dell Technologies.



Fig. 3: The JURECA Booster module at the Jülich Supercomputing Centre. The Cluster module is visible at the left border of the photograph.

be operated as a single system with a homogeneous global software stack.

As part of the deployment, the partners engage in a cooperative research effort to develop the necessary high-speed bridging technologies that enables high-bandwidth, low-latency MPI communication between Cluster and Booster compute nodes through the bridge nodes. The development will be steered by a number of real-world use cases, such as earth systems modeling and in-situ visualization.

The compute time on the Booster system will be made available primarily to scientists at Forschungszentrum Jülich and RWTH Aachen University. During a two-year interim period, all admissible researchers at German universities can request computing time by answering the calls of the John von Neumann Institute for Computing (NIC) until the second phase of the JUQUEEN successor system has been fully deployed.

The realization of the Cluster-Booster architecture in the JURECA system marks a significant evolution of JSC's dual architecture strategy as it brings "general purpose" and highly-scalable computing resources closer together. With the replacement of the JUQUEEN system in 2018, JSC intends to take the next step in its architecture roadmap and, in phases, deploy a Tier-0/1 "Modular Supercomputer" that tightly integrates multiple, partially specialized, modules under a global homogeneous software layer.

References

- [1]: **DEEP and DEEP-ER projects:**
<http://www.deep-projects.eu>
- [2]: **Jülich Supercomputing Centre, JURECA:**
General-purpose supercomputer at Jülich Supercomputing Centre. Journal of large-scale research facilities, Volume 2, A62.
<http://dx.doi.org/10.17815/jlsrf-2-121>, 2016

Written by Dorian Krause

Jülich Supercomputing Centre (JSC)

Contact: d.krause@fz-juelich.de

Two New Research Groups Established at the John von Neumann Institute for Computing in Jülich

The three founding partners of the John von Neumann Institute for Computing (NIC)—namely the Forschungszentrum Jülich (FZJ), Deutsches Elektronen-Synchrotron DESY Zeuthen and GSI Helmholtzzentrum für Schwerionenforschung Darmstadt—support supercomputer-oriented research and development through research groups dedicated to selected fields of physics and other natural sciences. Recently, two new groups have formed for the field of computational biology, but place their respective emphases on different aspects.

The NIC research group, “Computational Biophysical Chemistry” began work at FZJ at the end of April, 2017. The group is headed by Holger Gohlke from Heinrich Heine University Düsseldorf. Prof. Gohlke obtained his diploma in chemistry from the Technische Universität Darmstadt and his PhD from the Philipps-Universität Marburg. He subsequently conducted postdoctoral research at The Scripps Research Institute, La Jolla, USA. After appointments as an assistant professor in Frankfurt and a professor in Kiel, he moved to Düsseldorf in 2009. Prof. Gohlke was awarded the “Innovationspreis in Medizinischer und Pharmazeutischer Chemie” (innovation award for medicinal and pharmaceutical chemistry) by the German Chemical Society (GDCh) and the German Pharmaceutical Society (DPhG), the Hansch Award of the Cheminformatics and QSAR Society, and the Novartis Chemistry Lectureship. His current research focuses on the understanding, prediction, and modulation of interactions involving biomolecules and supramolecules from a computational perspective. Prof. Gohlke’s group

develops and applies techniques founded in structural bioinformatics, computational biology, and computational biophysics. In line with these research interests, the group is excited about the possibility of a dual affiliation with the Jülich Supercomputing Centre (JSC) and the Institute of Complex Systems/Structural Biochemistry (ICS-6). This will pave the way to bridging the supercomputing capabilities of the JSC with the structural biochemistry capabilities of the ICS-6 in order to address complex questions regarding the structure, dynamics, and function of biomolecules and supramolecules.



The other new group, “Computational Structural Biology” is led by Alexander Schug. It became operational in September, 2017 and is affiliated with the JSC. Dr. Schug studied physics at the University of Dortmund and obtained his PhD 2005 at the Forschungszentrum Karlsruhe and the University of Dortmund. Afterwards, he has worked as Postdoctoral Scholar in Kobe (Japan) and San Diego (US) before becoming an Assistant Professor in Chemistry (Umeå, Sweden). In 2011, he



returned to Germany to head a research group at the Karlsruhe Institute of Technology (KIT). Dr. Schug has received multiple awards, including a FIZ Chemie Berlin Preis from the GdCH and a Google Faculty Research Award 2016. His general research interests include theoretical biophysics, biomolecular simulations, and high-performance computing. In this role, his group leverages the incredible, constantly growing capabilities of HPC by integrating data from multiple sources in simulations to gain new insight about topics ranging from biomolecular structure and dynamics at atomic resolution to understanding neural cell tissue growth and differentiation. As understanding these properties is key to understanding biological function, this work promises to provide significant new insight with impact in fields ranging from basic molecular biology to pharmacological and medical research.

Further information about these new and all other NIC research groups can be found on our web site [1].

Reference

- [1] **John von Neumann Institut:**
http://www.john-von-neumann-institut.de/nic/EN/Research/ResearchGroups/_node.html

Written by Alexander Trautmann

NIC Coordination Office at the Jülich Supercomputing Centre (JSC)

Contact: coordination-office@fz-juelich.de

Quantum Annealing & Its Applications for Simulation in Science & Industry, ISC 2017

At ISC 2017, the international supercomputing conference held in Frankfurt am Main June 18–22, Prof. Dr. Kristel Michielsen from the Jülich Supercomputing Centre hosted the special conference session “Quantum Annealing & Its Applications for Simulation in Science & Industry”. The goal of the session was to introduce the

general principles of quantum annealing and quantum annealer hardware to the global HPC community and to discuss the challenges of using quantum annealing to find solutions to real-world problems in science and industry. These topics were addressed in four presentations:

- An Introduction to Quantum Annealing
Prof. Dr. Kristel Michielsen, Jülich Supercomputing Centre (ISC)
Summary and YouTube video: <http://primeurmagazine.com/live/LV-PL-06-17-36.html>
- Qubits, Couplers & Quantum Computing in 2017
Dr. Denny Dahl, D-Wave Systems
Summary and YouTube video: <http://primeurmagazine.com/weekly/AE-PR-08-17-7.html>
- Quantum Annealing for Aerospace Planning Problems
Dr. Tobias Stollenwerk, Deutsches Zentrum für Luft- und Raumfahrt (DLR)
Summary and YouTube video: <http://primeurmagazine.com/weekly/AE-PR-08-17-27.html>
- Maximizing Traffic Flow Using the D-Wave Quantum Annealer
Dr. Christian Seidel, Volkswagen (VW)
Summary and YouTube video: <http://primeurmagazine.com/live/LV-PL-06-17-37.html>



Session speakers, from left: Christian Seidel, Denny Dahl, and Tobias Stollenwerk. Right: session host Kristel Michielsen.



Quantum annealing and discrete optimization

New computing technologies, like quantum annealing, open up new opportunities for solving challenging problems including, among others, complex optimization problems. Optimization challenges are omnipresent in scientific research and industrial applications. They emerge in planning of production processes, drug-target interaction prediction, cancer radiation treatment scheduling, flight and train scheduling, vehicle routing, and trading. Optimization is also playing an increasingly important role in computer vision, image processing, data mining and machine learning.

The task in many of these optimization challenges is to find the best solution among a finite set of feasible solutions. In mathematics, optimization deals with the problem of finding numerically minima of a cost function, while in physics it is formulated as finding the minimum energy state of a physical system described by a Hamiltonian, or energy function. Quantum annealing is a new technique, exploiting quantum fluctuations, for solving those optimization problems that can be mapped to a quadratic unconstrained binary optimization problem (QUBO). A QUBO can be mapped onto an Ising Hamiltonian and the simplest physical realizations of quantum annealers are those described by an Ising Hamiltonian in a transverse field, inducing the quantum fluctuations. Many challenging optimization problems playing a role in scientific research and in industrial applications naturally occur as or can be mapped by clever modeling strategies onto QUBOs.




D-Wave Systems

Founded in 1999, D-Wave Systems is the first company to commercialize quantum annealers, manufactured as integrated circuits of superconducting qubits which can be described by the Ising model in a transverse field. The currently available D-Wave 2000Q™ systems have more than 2000 qubits (fabrication defects and/or cooling issues render some of the 2048 qubits inoperable) and 5600 couplers connecting the qubits for information exchange. The D-Wave 2000Q™ niobium quantum processor, a complex superconducting integrated circuit with 128,000 Josephson junctions, is cooled to less than 15 mK and is isolated from its surroundings by shielding it from external magnetic fields, vibrations and external radiofrequency fields of any form. The power consumption of a D-Wave 2000Q™ system is less than 25 kW,

most of which is used by the refrigeration system and the front-end servers.

Roughly speaking, programming a D-Wave machine for optimization consists of three steps: (i) encode the problem of interest as an instance of a QUBO; (ii) map the QUBO instance on the D-Wave Chimera graph architecture connecting

a qubit with at most six other qubits, which in the worst case requires a quadratic increase in the number of qubits; (iii) specify all qubit coupling values and single qubit weights (the local fields) and perform the quantum annealing, a continuous time (natural) evolution of the quantum system, on the D-Wave device. The solution is not guaranteed to be optimal. Typically a





user performs thousands of annealing runs for the problem instance to obtain a distribution of solutions corresponding to states with different energy.

The potential of quantum annealing for some applications in science and industry

The exploration of quantum annealing's potential for solving some real-world problems on D-Wave Systems' hardware is a challenge that nowadays is taken up not only in the US, but also in Europe. For these exploratory endeavors, it is essential that users from science and industry have easy access to this new computing technology at an early stage. As explained by Kristel Michielsen, JSC aims to establish a Quantum Computer User Facility hosting a D-Wave quantum annealer and various other quantum computing systems. The development of applications in the field of quantum computers by research groups in science and industry in Germany and the rest of Europe will largely profit from opportunities of being able to access the various available technologies.

In his presentation, Denny Dahl from D-Wave Systems focused on the possible benefits of new annealing controls, introduced with the latest-generation system, that allow the user to have more control over the annealing process. These controls help improve performance in finding solutions to certain problems or simulating particular quantum systems. As sample problems, he considered prime factorization and the simulation of a three-dimensional Ising model in a transverse field.

Tobias Stollenwerk from DLR reported on a research project pertaining to an aerospace planning problem, which he performed in close collaboration with researchers from NASA Ames. There are about 1,000 transatlantic flights per day. In order to fit more flights in the limited airspace, one considers wind-optimal or fuel saving trajectories which might lead to conflicts (airplane collisions). To solve the deconflicting problem with minimum flight delays, it was first formulated as a QUBO and then solved on a D-Wave machine. Problem instances with up to 64 flights and 261 conflicts were solved.

Christian Seidel from VW showed how to maximize traffic flow using the D-Wave quantum annealer. For this project the VW team used a public data set containing data of 10,000 taxi's driving in Beijing during one week. They restricted the traffic flow maximization problem, in which the travel time for each car has to be smaller than the one in the un-optimized traffic flow, using a trajectory from Beijing city centre to the airport for 418 cars thereby allowing each car to take three possible routes. They formulated this constrained optimization problem as a QUBO and solved it on the D-Wave machine.

Written by Prof. Kristel Michielsen

Jülich Supercomputing Centre (JSC)

Contact: k.michielsen@fz-juelich.de

The Virtual Institute– High-Productivity Supercomputing Celebrates its 10th Anniversary

The perpetual focus on hardware performance as a primary success metric in high-performance computing (HPC) often diverts attention from the role of people in the process of producing application output. But it is ultimately this output and the rate at which it can be delivered, in other words the productivity of HPC, which justifies the huge investments in this technology. However, the time needed to come up with a specific result or the “time to solution,” which it is often called, depends on many factors, including the speed and quality of software development. This is one of the solution steps where people play a major role. Obviously, their productivity can be enhanced with tools such as debuggers and performance profilers, which help find and eliminate errors or diagnose and improve performance.

Ten years ago, the Virtual Institute–High-Productivity Supercomputing (VI-HPS) was created with exactly with this goal in mind. Application developers should be able to focus on accomplishing their research objectives instead of having to spend major portions of their time solving software-related problems. With initial funding from the Helmholtz Association, the umbrella organization of the major national research laboratories in Germany, the institute was founded on the initiative of Forschungszentrum Jülich together with RWTH Aachen University, Technische Universität Dresden, and the University of Tennessee. Today, the institute encompasses twelve member organizations from five countries, including all three members of the Gauss Centre for Supercomputing.

Since then, the members of the institute have developed powerful programming tools, in particular for the purpose of analyzing HPC application correctness and performance, which are today used across the globe. Major emphasis was given to the definition of common interfaces and exchange formats between these tools to improve the interoperability between them and lower their development cost. A series of international tuning workshops and tutorials taught hundreds of application developers how to use them. At the multi-day VI-HPS Tuning Workshops, attendees are introduced to the tool suite, learn how to handle the tools effectively, and are guided by experts when applying the tools to their own codes. 25 tuning workshops have been organized at 19 different organizations in 9 countries all over the world. Numerous tools tutorials at HPC conferences and seasonal schools have been presented, especially at ISC in Germany and at SC in the US, where in some years five or more tutorials were given by VI-HPS members. Finally, the institute organized numerous academic workshops to foster the HPC tools community and offer especially young researchers a forum to present novel program analysis methods, namely the two workshop series Productivity and Performance (PROPER) at the Euro-Par conference from 2008 to 2014, and Extreme-Scale Programming Tools (ESPT) at SC since 2012.

One June 23, 2017, the institute celebrated its 10th anniversary at a workshop held in Seeheim, Germany. Anshu Dubey from Argonne National Laboratory, one of the keynote speakers,



explained that in HPC usually all parts of the software are under research, an important difference to software development in many other areas, leading to an economy of incentives where pure development is often not appropriately rewarded. In his historical review, Felix Wolf from TU Darmstadt, the spokesman of VI-HPS, looked back on important milestones such as the bylaws introduced to cover the rapid expansion of VI-HPS taking place a few years ago. In another keynote, Satoshi Matsuoka from the Tokyo Institute of Technology / AIST, Japan highlighted the recent advances in artificial intelligence and Big Data analytics as well as the

challenges this poses for the design of future HPC systems. Finally, all members of VI-HPS presented their latest productivity-related research and outlined their future strategies.

Reference

Workshop website:

www.vi-hps.org/symposia/10th-anniversary/

Written by Bernd Mohr¹ and Felix Wolf²

¹Jülich Supercomputing Centre (JSC)

²Technische Universität Darmstadt

Contact: wolf@cs.tu-darmstadt.de

Hazel Hen's Millionth Compute Job



A 3D visualization of the data set investigated in Hazel Hen's millionth job, displayed in the CAVE virtual reality environment at HLR5.

© HLR5

The HLR5 supercomputer recently reached a milestone, crossing into seven digits in the number of compute jobs it has executed. The millionth job is an example of the essential role that high-performance computing is playing in many research fields, including fluid dynamics.

Traditional laboratory experimentation has been and continues to be indispensable to the advance of scientific knowledge. Across a wide range of fields, however, researchers are being compelled to ask questions about phenomena

that are so complex, so remote, or that are found at such small or large scales that direct observation just isn't practical anymore. In these cases, simulation, modeling, data visualization, and other newer computational approaches offer a path forward. The millionth job on Hazel Hen was one such case.

Leading the research behind the millionth job was Professor Bernhard Weigand, Director of the Institute of Aerospace Thermodynamics at the University of Stuttgart. His laboratory



studies multiphase flows, a common phenomenon across nature in which materials in different states or phases (gases, liquids, and solids) are simultaneously present and physically interact. In meteorology, for instance, raindrops, dew, and fog constitute multiphase flows, as does the exchange of gases between the oceans and the atmosphere. Such phenomena also occur in our daily lives, such as when water bounces off our skin in the shower or when we inhale nasal sprays to control the symptoms of a cold.

In engineering, multiphase flows can also be extremely important. Perhaps their most familiar application is in the design of fuel injection systems in automobiles, gas turbines, and rockets. Other examples include the spreading of fertilizers for farming or the use of spray drying in the production of pharmaceuticals and foods.


In all of these cases, understanding how multiphase flows behave in detail could both enhance our ability to study the natural world and improve the design of more effective and more efficient products. But because of the enormous numbers of droplets that are involved in multiphase flows and the extremely small scale at which they interact, our ability to gain precise knowledge about them purely through observation has been limited.

For this reason, Weigand turned to HLRS and its Hazel Hen high-performance computer to simulate multiphase flows computationally. His work and that of his colleagues has led to a variety of insights with wide-ranging practical relevance.

Supercomputing simulates droplet dynamics

Professor Weigand and his group are primarily interested in basic multiphase flows involving droplets, such as those that fall as rain from the sky. In the past Weigand and his group investigated topics related to the dynamics of cloud formation, for example, gaining insights into what happens when droplets in the atmosphere collide; these findings were subsequently used by other scientists to develop better weather forecast models. Weigand is also speaker of the Collaborative Research Council SFB-TRR 75 (a research project funded by the Deutsche Forschungsgemeinschaft (DFG) that also includes investigators at TU Darmstadt and DLR Lampoldshausen). In this capacity his team has been investigating the fluid dynamics of super-cooled water droplets in extreme situations, such as when ice crystals develop in clouds. This problem is important for precipitation forecasting (for example, hail) and also in air travel, as ice formation on airplane wings can negatively affect flight stability and decrease fuel efficiency.

To study the dynamics of droplets' physical behavior, Weigand and his group use a mathematical approach called direct numerical simulation (DNS). Over many years he and members of his lab have been building DNS methods into an in-house software program called FS3D (Free Surface 3D), which they use to model droplet dynamics. FS3D can, for example, precisely simulate what happens when a water droplet falls onto a liquid film and forms a "crown," taking a new shape and breaking apart into smaller droplets.



High-performance computing (HPC) is absolutely essential to the success of FS3D because the software requires an extremely high “gate resolution.” Like the frame rate in a video or movie camera, the program must represent the complex collisions, adhesions, and breaking apart of droplets and molecules at extremely small scales of space and time. FS3D can simulate such interactions in 2 billion “cells” at once, each of which represents a volume of less than 7 cubic micrometers, tracking how the composition of every cell changes over time.

Achieving such a high resolution generates massively large datasets, and it is only by using a supercomputer as powerful as HLRS’s Hazel Hen that these simulations can be run quickly enough to be of any practical use. Moreover, during simulations, HPC architectures can rapidly and reliably save enormous collections of data that are output from one round of calculations and efficiently convert them into inputs for the next. In this way, simulation becomes

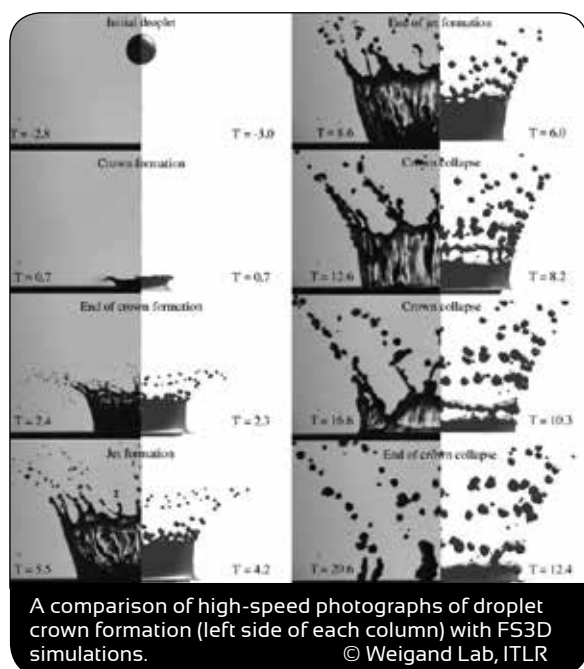
an iterative process, leading to better and better models of complex phenomena, such as the multiphase flows the Weigand Lab is investigating.

Having so much power at your disposal presents some unique challenges, though. In order to take full advantage of the opportunities that supercomputers offer, software behind algorithms like FS3D must be written specifically for the parallel computing architecture of high-performance computing systems. Programming in this way requires special expertise, and as FS3D has developed, staff members at HLRS and at Cray, the company that built Hazel Hen, have helped the Weigand Lab to optimize it for HPC.

“It’s not really practical for us to have HPC experts in our lab, and so staff at HLRS and Cray have been very supportive in helping us to run FS3D effectively on Hazel Hen,” says Dr. Weigand. “Their knowledge and advice have been very important to the success of our recent studies.”

The millionth job: visualizing how non-newtonian fluids break apart in jets

The millionth job on Hazel Hen was not focused on atmospheric water, but instead on multiphase flows in non-Newtonian fluids. Such fluids—which include materials like paint, toothpaste, or blood—do not behave in ways that Newton’s laws of viscosity would predict; instead, their fluid dynamic properties follow other rules that are not as thoroughly understood.



More specifically, Weigand's team wanted to use computational simulations to gain a better understanding of how non-Newtonian jets break up when injected into a gaseous atmosphere. This question is important because droplet sizes and the increase in a fluid's surface area as it becomes atomized can be important factors in optimizing the efficiency of a process—such as in the application of aerosolized paint to a car body.

The researchers simulated the injection of aqueous solutions of the polymers Praestol2500® and Praestol2540® through different pressure nozzles into air. When used in water treatment, the viscosity of these polymers decreases due to shear strain. The fluid properties for this case were approximated by flow curves obtained from experiments by colleagues at the University of Graz.

Running FS3D on Hazel Hen, the Weigand team performed a variety of "virtual" experiments on the supercomputer to investigate specific features of these flows, gaining a much more precise picture of how the solutions disperse. For example, they modeled jet breakup after

injection and how factors such as flow velocity and the shape of the nozzle changed the fluids' viscous properties. (This work was undertaken under the auspices of DFG-funded priority program SPP 1423-Process Spray. Speaker: Prof. Udo Fritsching, University of Bremen).

The millionth job run on Hazel Hen was one of several post-processing visualizations the team undertook in cooperation with VISUS (University of Stuttgart Visualisation Research Centre) to investigate the development of a liquid mass over time. In this series of studies, they generated extremely fine-grained visualizations of changes in the shape of the flow passing through the jet, identified differences in the loss of flow cohesion under different conditions, and discovered changes in surface area as the flow becomes atomized, among other characteristics. This led to insights about similarities and differences between Newtonian and non-Newtonian flows, and about how nozzle shape affects flow properties.

In the future, such information could enable engineers to improve the efficiency of their nozzle designs. In this sense, the millionth compute job on Hazel Hen was just one page in a long and continuing scientific story. Nevertheless, it embodies the unique kinds of research that HLRS makes possible everyday.

Written by Christopher Williams

Summer School Focuses on Links between Computer Simulation and the Social Sciences

A five-day seminar hosted by the High-Performance Computing Center Stuttgart (HLRS) provided philosophers, social scientists, and historians a detailed understanding of the capabilities and challenges of simulation.

On September 25-29, 2017, students and lecturers with backgrounds in the humanities, social sciences, and computational sciences gathered at the High-Performance Computing Center Stuttgart (HLRS) to explore such issues in the Summer School on Computer Simulation Methods. The event was a joint effort organized by Michael Resch (Director, HLRS) and Andreas Kaminski (Head, HLRS Department for Philosophy of Science & Technology of Computer Simulation at HLRS), in collaboration with Giuseppe Primiero (Middlesex University London) and Viola Schiaffonati (Politecnico Milano).

The goal of the Summer School was to increase cross-disciplinary understanding between computational scientists on the one side and researchers in the humanities and social sciences on the other. Such bridge building is important because each has different interests and fundamental questions when it comes to thinking about computational science methods. Whereas computer scientists are interested in methodological and technical issues involved in simulation technologies, for example, researchers in the humanities and social sciences explore issues that operate on the level of self-observation. Such approaches can support critique of the extent to which simulations can be trusted in their representation of real-life phenomena.

Exploring key questions in simulation

In order to give the participants a framework for exploring such questions, the summer school was divided into several main topics. Led by HLRS investigators Uwe Küster and Ralf Schneider, the first day provided an introduction to some foundations of high-performance computing, numerical methods, and the differences between simulation and experimentation. In the afternoon, Schiaffonati discussed relationships between experiment and simulation in the sciences, presenting a nuanced look at different types of simulations and their uses in research.

The second day was dedicated to the role of mathematics in simulation. For centuries mathematics has been an important framework for describing the world and predicting what will happen in the future. As Küster, Kaminski, and Johannes Lenhard (Universität Bielefeld) suggested, however, several fundamental problems related to the nature of mathematics and its implementation in software raise questions about the extent to which we can always trust it. Tuesday afternoon also featured a lecture looking specifically at the use of computer simulations to understand biological systems in the cell.

In a day focusing on social applications of simulation, empirical sociologist Nicole Saam (Friedrich-Alexander-University Erlangen-Nürnberg), presented a two-part lecture on the challenges facing social science simulation. "Computer simulations in the social sciences," Saam said, "are mainly done either as theory-free micro-simulations or on the basis of rational choice



theory, although these approaches play a rather marginal role in sociology. What is missing are computer simulations on (different) theoretical fundamentals.” Saam discussed several paradigms in sociology that are still lacking in suitable applications in computer simulation and explored the reasons for this with the participants. Primiero followed her lectures with a talk using an epistemological approach to think about case studies in agent-based simulation.

Friday focused on the use of visualization in simulation and data analysis. Following a tour of the CAVE, a three-dimensional virtual reality environment at HLRS, Visualization Department Head Uwe Wössner discussed the opportunities and challenges related to the use of virtual reality in research and technology development. The lectures concluded with a talk by Angelo Verneulen (Delft University of Technology), a scientist and artist who is developing new concepts for starship development.

The start of a new collaboration among HLRS, Middlesex University, and Politecnico Milano

The decision to hold the first edition of this Summer School at HLRS in Stuttgart the result of its unique mandate. “The Department of Philosophy of Science & Technology of Computer

Simulation,” Kaminski explained, “manages to close the gap between social science researchers interested in computation and their research subject. Philosophers and social scientists are able to study numeric scientific practice and the methodology of simulation at the same place where expertise in these fields is gathered.”

Considering the warm reception that Summer School received from its attendees, the organizers plan to deepen their collaboration and organize future events focusing on the intersection of philosophy, social science, and computer science. “We have considered the possibility of organizing a workshop in Milano end of next year,” Viola Schiaffonati revealed.

Participant profiles

Ramon Alvarado, Philosophy of sciences student, University of Kansas:

“What I liked most about this Summer School was the exchange on heavy epistemological questions, like what we are doing with computers, and why we are doing it the way we are.”

Marta Conti, Student of Physics and Philosophy, University of Barcelona:

“Having studied differential equations and simulation myself as a physicist, I was fascinated by the concept of questioning scientific practice, particularly how under the same conditions simulated outcomes can differ.”

Written by Lena Bühler

FAU Students Win Highest Linpack Award at ISC17's Student Cluster Competition

GCS-sponsored team FAU Boyzz, six students at Friedrich-Alexander-Universität Erlangen-Nürnberg, Germany (FAU), walked away with a highly coveted championship title from the Student Cluster Competition (SCC), held in the framework of the International Supercomputing Conference 2017 (ISC). Team FAU Boyzz, made up of bachelors students studying computational engineering, computer science, and medical engineering, captured the trophy for the hotly competed SCC High Performance Linpack (HPL) benchmark challenge. The amazing HPL score of 37,05 Teraflops (1 Teraflop = 1 trillion floating point operations per second), delivered on the students' self-assembled Hewlett Packard Enterprise (HPE) cluster featuring 12 NVIDIA P100 GPUs, marks a new all-time high in the history of ISC's SCC. The score almost triples the result of the previous year's SCC Linpack high mark achieved at ISC.

The HPL benchmark traditionally enjoys special attention among the challenges the student teams face in the course of a gruelling, ambitious three-day competition. The event is an integral part of the annually recurring ISC, the international counterpart of SC, the worlds largest HPC conference, held in the United States.

"This competition is quite fun and quite challenging," said Jannis Wolf, team captain of FAU Boyzz. "We have been preparing for this for a year and we've met people that we otherwise never would have—our team had different disciplines coming together."

During the contest, teams of undergraduate students are exposed to a variety of application codes and asked to solve a variety of problems related to HPC. In addition to application performance, teams are judged on their clusters' energy efficiency and power consumption, application performance and accuracy, and interviews by subject matter experts assessing their knowledge of their systems and applications.

"One of the best parts is the practical knowledge that comes from this process," said team member Lukas Maron. Indeed, the teams are given real-world applications and work closely with mentors who are already active in the HPC community. This type of experience is invaluable for students' future career prospects and also for exposing them to possible new avenues to explore.

"I think this is a great opportunity for students to get a feeling for what it is like at an HPC conference, to deal with a wide variety of applications, and to get to be able to design a cluster from scratch," said FAU researcher and team mentor Alexander Ditter. "Of course, it would not be possible for us to participate in these kinds of friendly competitions were there no support from the research community as well as the industry. Thus I would like to express big thanks to our sponsors GCS and SPPEXA who helped us financially, and to our hardware sponsors HPE and NVIDIA. We hope our success made them proud."



ISC High Performance

The HPC Event.

JUNE 18 - 22, 2017
FRANKFURT, GERMANY



Team FAU Boyzz of Friedrich-Alexander-Universität Erlangen-Nürnberg, proud winners of the LINPACK benchmark challenge at ISC17's SCC. From left: Phillip Suffa, team captain Jannis Wolf, Benedikt Oehlrich, Lukas Maron, Fabian Fleischer, Egon Araujo. © GCS

The complete list of teams participating in the ISC Student Cluster Competition is:

- Centre for High Performance Computing (South Africa)
- Nanyang Technological University (Singapore)
- EPCC University of Edinburgh (UK)
- Friedrich-Alexander University Erlangen–Nürnberg (Germany)
- University of Hamburg (Germany)
- National Energy Research Scientific Computing Center (USA)
- Universitat Politècnica De Catalunya Barcelona Tech (Spain)
- Purdue and Northeastern University (USA)
- The Boston Green Team (Boston University, Harvard University, Massachusetts Institute of Technology, University of Massachusetts Boston) (USA)
- Beihang University (China)
- Tsinghua University (China)

“The Gauss Centre for Supercomputing, which by definition is highly interested in drawing young talent’s attention toward High Performance Computing, is always open to support up and coming HPC talent, also in the framework of these kinds of events,” explains Claus Axel Müller, Managing Director of the GCS. “We are well aware of the financial constraints students face when trying to participate in international competitions, especially if travel and related expenses are involved. Thus, we are happy to be of help and we would like to sincerely congratulate the FAU Boyzz for their great achievements at ISC.”

Written by Regina Weigand

GCS Public Relations

Contact: r.weigand@gauss-centre.eu

Changing Gear for Accelerating Deep Learning: First-Year Operation Experience with DGX-1



The rise of GPU for general purpose computing has become one of the most important innovations in computational technology. The current phenomenal advancement and adaptation of deep learning technology in many scientific and engineering disciplines won't be possible without GPU computing. Since the beginning of 2017, the Leibniz Supercomputing Centre of the Bavarian Academy of Sciences and Humanities has deployed several GPU systems, including a DGX-1 and OpenStack cloud-based GPU virtual servers (with Tesla P100). Among many typical deep-learning-related research areas, our users tested the scalability of deep learning on DGX-1, trained recurrent neural networks to optimize dynamical decoupling for quantum memory, and performed numerical simulations of fluid motion, utilizing the multiple NVLink P100 GPUs on DGX-1. These research activities demonstrate that GPU-based computational platforms, such as DGX-1, are valuable computational assets of the Bavarian academic computational infrastructure.

Scaling CNN training on the DGX-1

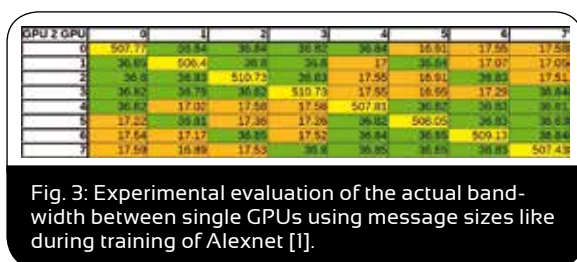

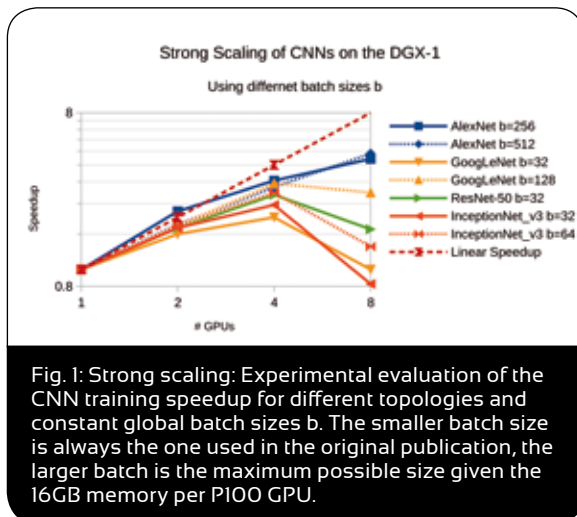
The training of deep neural networks (DNN) is a very compute- and data-intensive task. Modern network topologies [3,4] require several exaFLOPS until convergence of the model. Even training on a GPU still requires several days of training time. Using a multi-GPU system could ease this problem. However, parallel DNN training is a strongly communication bound problem [5]. In this study, we investigate if the NVLINK interconnect,

with its theoretical bandwidth of up to 50 GB/s, is sufficient to allow scalable parallel training.

We used four popular convolutional neural network (CNN) topologies to perform our experiments: AlexNet [1], GoogLeNet [2], ResNet [3] and InceptionNet [4]. The software stack was built on NVIDIA-Caffe v0.16, Cuda 8, and cuDNN 6. We used the data-parallel training algorithm for multi-GPU systems [5], which is provided by the Caffe framework.

Figure 1 shows the results for a strong scaling of the CNN training. Notably, the parallelization appears to be efficient up to four GPUs, but drops significantly when scaling to eight GPUs. This might be caused by the NVLINK interconnection topology of the DGX-1 (shown in Fig 3), where the GPUs are split into two fully connected groups of four. However, looking at the results for AlexNet (which has the largest communication load) shows that the maximum possible batch size is actually the problem. As shown in [5], data-parallel splitting of smaller batch sizes causes inefficient matrix operations at the worker level.


Large batch sizes can be preserved by a weak scaling approach, shown in figure 2. Using the maximum global batch size leads to better scaling performance. However, it should be noticed that increasing the batch size usually leads to reduced generalization abilities of the trained model [5].



Shifting gears: gear-train simulations on the DGX-1 using nanoFluidX

Besides common utilization of GPUs on DGX-1 for machine learning (deep learning), GPUs can be used for numerical simulations of fluid motion. One of the GPU-based CFD codes on the market is the nanoFluidX (nFX) code based on the smoothed particle hydrodynamics (SPH) method, developed by FluidDyna GmbH.

nFX is primarily used for simulations of gear- and power-train components in the automotive industry, allowing quick execution of transient, multiphase simulations in complex moving geometries that would otherwise be prohibitively computationally expensive or impossible to do with conventional finite-volume methods.



The SPH method is based on an algorithm that is perfectly suited for parallelization, as it involves a large number of simple computations repeated over regions that are spatially independent. This allows for easy distribution of tasks over threads and efficiently harnesses the power of the GPUs.



Performance and scaling of the nFX code on DGX-1 are shown in Figs. 4 and 5. The chosen test case for scaling and performance tests is a single gear immersed in an oil sump. The case contains 8,624,385 particles, which at maximum number of GPUs results in approximately 1 million particles per GPU device. Each case ran for exactly 1000 steps, resulting in a minimum run time of 37.78 seconds and maximum of 2 minutes, 54 seconds.

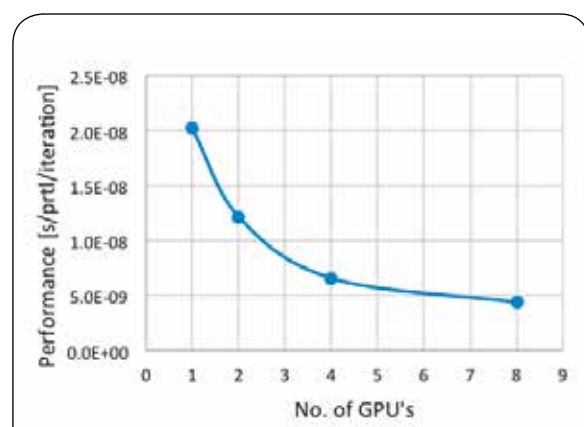


Fig. 5: nFX code performance [s/particle/iteration]. It is noticeable that the scaling drops off as the number of particles per GPU decreases. This is common behaviour under suboptimal load of the cards, as communication becomes more prominent while the GPU memory is underutilized.

It has been noted that scaling on GPUs is heavily influenced by the relative load put on each card. In reality, this transfers to the issue of having an upper limit on the acceleration of the simulation for a limited size of the case. As a counter-example, one can imagine having a case with 100 million particles, and scaling would likely be almost ideal in the range 1-10 GPUs, but would likely drop off to about 80% at 100 GPUs.




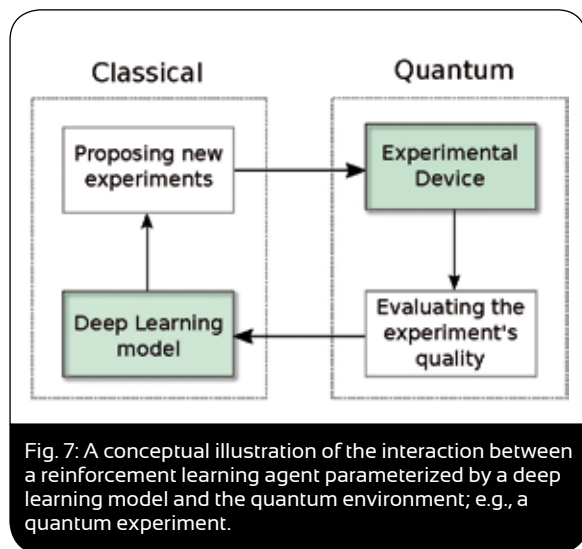
Fig. 6: Strong scaling efficiency, calculated as 1 minus the ratio of the single GPU performance to the current number of GPUs. We see that the performance tops off at around 80% efficiency, corresponding to the relative performance drop from Fig. 2.

Deep learning models for simulating quantum experiments

This work aims at developing deep learning models to automatically optimize degrees of freedom and predict results of quantum physics experiments. In order for these algorithms to be broadly applicable and be compatible with quantum mechanical particularities—i.e., measurements influence the results—we take a black-box perspective and, for instance, do not assume the error measure representing the experiment's result to be differentiable.

August and Ni have recently introduced an algorithm [6] for the optimization of protocols for quantum memory. The algorithm is based on long short-term memory (LSTM) recurrent neural networks that have been successfully applied in the fields of natural language processing and machine translation. Tackling this

problem from a different perspective, August has now casted it to a reinforcement learning setting where the agent's policy is again represented as an LSTM.



References

- [1] **AlexNet: Krizhevsky, Alex, Ilya Sutskever, and Geoffrey E. Hinton:**
"Imagenet classification with deep convolutional neural networks." Advances in neural information processing systems. 2012.
- [2] **GoogLeNet: Szegedy, Christian, et al.:**
"Going deeper with convolutions." Proceedings of the IEEE conference on computer vision and pattern recognition. 2015.
- [3] **ResNet: He, Kaiming, et al.:**
"Deep residual learning for image recognition." Proceedings of the IEEE conference on computer vision and pattern recognition. 2016.
- [4] **InceptionNet: Szegedy, Christian, et al.:**
"Rethinking the inception architecture for computer vision." Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. 2016.
- [5] **Keuper, Janis, and Franz-Josef Preundt:**
"Distributed training of deep neural networks: theoretical and practical limits of parallel scalability." Machine Learning in HPC Environments (MLHPC), Workshop on. IEEE, 2016.
- [6] **August, Moritz and Ni, Xiaotong:**
"Using recurrent neural networks to optimize dynamical decoupling for quantum memory." Phys. Rev. A 95, 012335

Written by Yu Wang

The Leibniz Supercomputing Centre of the Bavarian Academy of Sciences and Humanities

Janis Keuper

Fraunhofer-Institut für Techno- und Wirtschaftsmathematik ITWM

Milos Stanic

Fluidyna GmbH

Moritz August

Technical University of Munich

Contact: Yu.Wang@lrz.de

Intel MIC Programming & HPC for Natural Hazard Assessment and Disaster Mitigation Workshops @ LRZ



IT4Innovations
national
supercomputing
center



For the fourth time, the Czech-Bavarian Competence Centre for Supercomputing Applications (CzeBaCCA) organised a technical Intel Many Integrated Core (MIC) programming workshop, combined with a scientific workshop about HPC simulations in the field of environmental sciences. The Czech-Bavarian Competence Centre was established in 2016 by the Leibniz Supercomputing Centre (LRZ), the Department of Informatics at the Technical University of Munich (TUM), and the IT4Innovations National Supercomputing Centre in the Czech

Republic to foster Czech-German collaboration in high-performance computing. One of the main objectives of the Competence Centre is to organise a series of Intel Xeon Phi-specific technical workshops in concert with scientific symposia on topics like optimisation of simulation codes in environmental science.

The successful series of workshops began in February 2016 with an introductory Intel MIC programming workshop concentrating on the Salomon supercomputer at IT4I, currently the



Fig. 1: Participants of the "Intel MIC programming workshop" at LRZ, June 26 – 28, 2017. The main lecturers in the first row (Dr. M. Allalen (LRZ), Dr. V. Weinberg (LRZ), Dr.-Ing. Michael Klemm (Intel), Dr.-Ing. Jan Eitzinger (RRZE), f.l.t.r.) are holding Intel Xeon Phi co-processors. © A. Podo (LRZ)

largest European Intel Knights Corner (KNC)-based system, combined with a symposium on "SeisMIC–Seismic Simulation on Current and Future Supercomputers" at IT4Innovations (see inSiDE Vol. 14 No. 1 p. 76ff, 2016). In June 2016 the series continued at LRZ with extended Intel MIC programming workshops on LRZ's SuperMIC system or on Salomon that was combined with a scientific workshop on "High-Performance Computing for Water Related Hazards" (see inSiDE Vol. 14 No. 2 p. 25ff, 2016) and a scientific symposium on "High-Performance Computing in Atmosphere Modelling and Air Related Environmental Hazards" (see inSiDE Vol. 15 No. 1 p. 48ff, 2017) in February 2017 at IT4Innovations.

The fourth edition of this workshop series took place at LRZ June 26–30, 2017. The three-day Intel MIC programming workshop was organised as a PRACE Advanced Training Centre (PATC) event. It covered a wide range of topics, from the description of the Intel Xeon Phi co-/processors' hardware, through information about the basic programming models and information about vectorisation and MCDRAM usage, up to tools and strategies for analysing and improving applications' performance. The workshop mainly concentrated on techniques relevant for Intel Knights Landing (KNL)-based many-core systems. During a public plenary session on Wednesday afternoon (joint session with the scientific workshop) eight invited speakers from IPCC@LRZ, IPCC@TUM, IPCC@IT4Innovations, Intel, RRZE, the University of Regensburg, IPP, and MPCDF talked about Intel Xeon Phi experiences and best-practice

recommendations for KNL-based systems. Hands-on sessions were done on the KNC-based system SuperMIC and two KNL test systems at LRZ. The workshop attracted over 35 international participants.

The Intel MIC programming workshop was followed by a two-day symposium on "HPC for natural hazard assessment and disaster mitigation." Presenters from the Technical University of Munich, the Ludwig-Maximilians-University Munich, the University of Augsburg, the Munich Reinsurance Company, and IT4Innovations addressed topics such as simulation of geological or meteorological hazards, floods, tsunamis, earthquakes, dangerous terrain motion, diseases, and other hazards to society. A special focus was on demands and desired features of (future) simulation software, parallelisation for current and novel HPC platforms, as well as scalable simulation workflows on supercomputing environments. The next Intel MIC programming workshops will take place at LRZ in June 2018 and will concentrate on the new KNL cluster, CoolMUC3, at LRZ.

Acknowledgements

The Czech-Bavarian Competence Centre for Supercomputing Applications is funded by the Federal Ministry of Education and Research. The Intel MIC programming workshop was financially also supported by the PRACE-4IP and PRACE-5IP projects funded by the European Commission's Horizon 2020 research and innovation programme (2014-2020) under grant agreements 653838 and 730913.



References

- [1] <https://www.lrz.de/forschung/projekte/forschung-hpc/CzeBaCCA/>
- [2] <http://www.training.prace-ri.eu/>
- [3] <https://www.lrz.de/services/compute/courses/>
- [4] <http://training.it4i.cz/kurzy-it4innovations/>

Written by Volker Weinberg, Momme Allalen, Arndt Bode, Anton Frank, Dieter Kranzlmüller, Megi Sharikadze

Leibniz Supercomputing Centre (LRZ), Germany

Ondřej Jakl, Branislav Janský, Martin Palkovic, Vít Vondrák

IT4Innovations, Czech Republic

Michael Bader

Technical University of Munich (TUM), Germany

Contact: Volker.Weinberg@lrz.de

New PATC Course: HPC Code Optimisation Workshop @ LRZ



As code optimisation techniques are getting more and more important in HPC, LRZ @ GCS—one of the six European PRACE Advanced Training Centres (PATCs)—has extended its curriculum by adding a new PATC course, “HPC code optimisation workshop,” which took place at LRZ on May 4, 2017 for the first time.

The workshop was organised as a compact course and focused on code improvement and exploration of the latest Intel processor features, particularly the vector units. During the optimisation process, attendees learned how to enable vectorisation using simple pragmas and more effective techniques, like changing the data layout and alignment. The process was guided by hints from the Intel compiler reports, and by using the Intel Advisor tool. The outline of the workshop included basics of modern computer architectures, optimisation process, vectorisation, and the Intel tools. An N-body code was used to support the described optimisation solutions with practical examples. Through a sequence of simple, guided examples of code modernisation, the attendees developed awareness on features of multi- and many-core architectures, which are crucial for writing modern, portable, and efficient applications. The exercises were done on the SuperMIC system at LRZ.

The lectures were given by Dr. Fabio Baruffa and Dr. Luigi Iapichino, who are both members of the Intel Parallel Computing Center (IPCC) established at LRZ in 2014. Within this framework, the team at LRZ is active in the performance optimisation of the Gadget code, a widely used community code for computational astrophysics, on multi- and many-core computer architectures. The experiences pertaining to the optimisation work for the IPCC were passed to the participants of the workshop as best practice recommendations.



Fig. 1: Participants of the new PATC course: “HPC code optimisation workshop” at LRZ, May 4, 2017.

© A. Podo (LRZ)

Due to the great success of the workshop, it will be repeated in May 2018 at LRZ.

Written by Volker Weinberg

Leibniz Supercomputing Centre (LRZ), Germany

Contact: Volker.Weinberg@lrz.de



Launched initially in November 2015 and formalized as a collaborative Linux Foundation project in June 2016, OpenHPC is a community driven project currently comprising over 25 member organizations with representation from academia, research labs, and industry. To date, the OpenHPC software stack aggregates over 60 components, ranging from tools for bare-metal provisioning, administration, and resource management to end-user development libraries that span a range of scientific/numerical uses. OpenHPC adopts a familiar repository delivery model with HPC-centric packaging in mind, and provides customizable recipes for installing and configuring reference designs of compute clusters. OpenHPC is intended both to make available current best practices and provide a framework for delivery of future innovation in cluster computing system software. The OpenHPC software stack is intended to work seamlessly on both Intel x86 and ARMv8-A architectures.

Motivation

Many HPC sites spend considerable effort aggregating a large suite of open-source components to provide a capable HPC environment for their users. This is frequently motivated by the necessity to build and deploy HPC focused packages that are either absent or outdated in popular Linux distributions. Further, local packaging or customization typically tries to give software versioning access to users (e.g., via environment modules or similar equivalent). With this background motivation in mind, combined with a desire to minimize duplication and

share best practices across sites, the OpenHPC community project was formed with the following mission and vision principles:

Mission

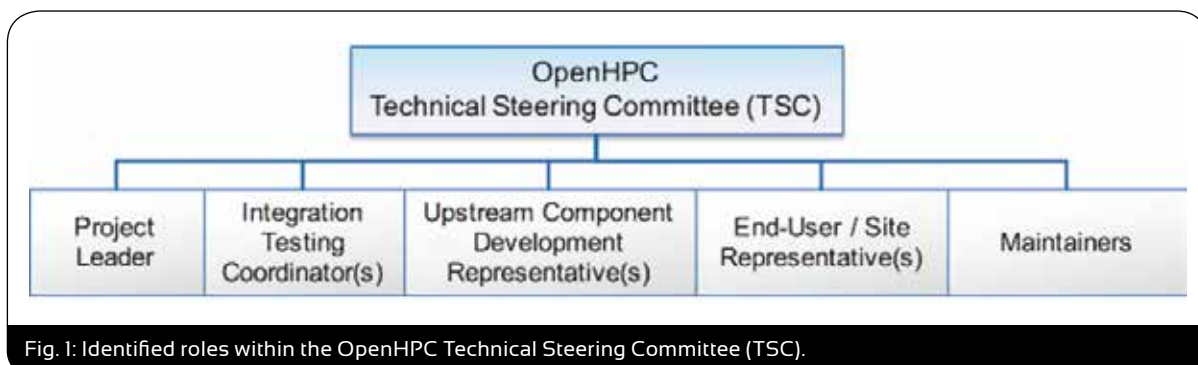
To provide a reference collection of open-source HPC software components and best practices, lowering barriers to deployment, advancement, and use of modern HPC methods and tools.

Vision

OpenHPC components and best practices will enable and accelerate innovation and discoveries by broadening access to state-of-the-art, open-source HPC methods and tools in a consistent environment, supported by a collaborative, worldwide community of HPC users, developers, researchers, administrators, and vendors.

Governance & community

Under the auspices of the Linux Foundation, OpenHPC has established a two-pronged governance structure consisting of a governing board and a technical steering committee (TSC). The governing board is responsible for budgetary oversight, intellectual property policies, marketing, and long-term road map guidance. The TSC drives the technical aspects of the project, including stack architecture, software component selection, builds and releases, and day-to-day project maintenance. Individual roles within the TSC are highlighted in Figure 1. These include common roles like maintainers and testing coordinators, but also include unique HPC roles designed to ensure influence and capture points of view from two



key constituents. In particular, the component development representative(s) are included to represent the upstream development communities for software projects that might be integrated with the OpenHPC packaging collection. In contrast, the end user/site representative(s) are downstream recipients of OpenHPC integration efforts and serve the interests of administrators and users of HPC systems that might leverage OpenHPC collateral. At present, there are nearly 20 community volunteers serving on the TSC with representation from academia, industry, and government R&D laboratories.

Build infrastructure

To provide the public package repositories, OpenHPC utilizes a set of standalone resources running the Open Build Service (OBS). OBS is an open-source distribution development platform written in Perl that provides a transparent infrastructure for the development of Linux distributions and is the underlying build system for openSUSE. The public OBS instance for OpenHPC is available at <https://build.openhpc.community>.

While OpenHPC does not, by itself, provide a complete Linux distribution, it does share many of the same packaging requirements and targets a delivery mechanism that adopts Linux sysadmin familiarity. OBS aids in this process by driving simultaneous builds for multiple OS distributions (e.g. CentOS, SLES and Red-hat), multiple target architectures (e.g. x86 64 and aarch64), and by performing dependency analysis among components, triggering downstream builds as necessary based on upstream changes.

Each build is carried out in a chroot or KVM environment for repeatability, and OBS manages publication of the resulting builds into package repositories compatible with yum and zypper. Both binary and source RPMs are made available as part of this process. The primary inputs for OBS are the instructions necessary to build a particular package, typically housed in an RPM .spec file. These .spec files are version controlled in the community GitHub repository and are templated in a way to have a single input drive multiple compiler/MPI family combinations.



Integration testing

To facilitate validation of the OpenHPC distribution as a whole, we have devised a stand-alone integration test infrastructure. In order to exercise the entire scope of the distribution, we first provision a cluster from bare-metal using installation scripts provided as part of the OpenHPC documentation. Once the cluster is up and running, we launch a suite of tests targeting the functionality of each component. These tests are generally pulled from component source distributions and aim to ensure development toolchains are functioning correctly and to ensure jobs perform under the resource manager. The intent is not to replicate a particular component's own validation tests, but rather to ensure all of OpenHPC is functionally integrated. The testing framework is publicly available in the OpenHPC GitHub repository. A Jenkins continuous integration server manages a set of physical servers in our test infrastructure. Jenkins periodically kickstarts a cluster master node using out-of-the-box base OS repositories, and this master is then customized according to the OpenHPC install guide. The LATEX source for the install guide contains markup that is used to generate a bash script containing each command necessary to provision and configure the cluster and install OpenHPC components. Jenkins executes this script, then launches the component test suite.

The component test suite relies on a custom autotools-based framework. Individual runs of the test suite are customizable using familiar autoconf syntax, and make check does what

one might expect. The framework also allows us to build and test multiple binaries of a particular component for each permutation of compiler toolchain and MPI runtime if applicable. We utilize the Bash Automated Testing System (BATS) framework to run tests on the cluster and report results back to Jenkins.

As the test suite has grown over time to accommodate a growing set of integrated components, the current test harness has both short and long configuration options. The short mode enables only a subset of tests in order to keep the total runtime to approximately 10 minutes or less for more frequent execution in our CI environment. For the most recent OpenHPC release, the long mode with all relevant tests enabled requires approximate 90 minutes to complete approximately 1900 individually logged tests.

Papers, presentations & tutorials

Papers:

- Schulz, K., Baird, C.R., Brayford, D., et al.: Cluster computing with OpenHPC. In: Supercomputing HPC Systems Professionals (2016)

Presentations:

- ISC 2017 - BoF Session; OpenHPC Birds of a Feather Session; David Brayford (LRZ), Chulho Kim (Lenovo), Karl W. Schulz (Intel) and Thomas Sterling (Indiana University)
- SC16 Birds of a Feather; OpenHPC Birds of a Feather Session; Karl W. Schulz (Intel) and David Brayford (LRZ)

- MVAPICH User Group Meeting 2016; OpenHPC Overview; Karl W. Schulz (Intel)
- ISC 2016 - BoF Session; OpenHPC Birds of a Feather Session; Karl W. Schulz (Intel)
- FOSDEM 2016; OpenHPC: Community Building Blocks for HPC Systems; Karl W. Schulz (Intel)

Tutorials:

- PEARC17; Getting Started with OpenHPC; Karl W. Schulz (Intel), Reese Baird (Intel), Eric Van Hensbergen (ARM), Derek Simmel (PSC), and Nirmala Sundararajan (DELL)

How to get involved

OpenHPC encourages participation from across the HPC community and draws on a global membership from over two dozen commercial companies, government, and academic organizations. The community <http://www.openhpc.community/>

- Mailing lists available to post questions, report issues, or stay aware of community announcements:

- openhpc-announce
<https://groups.io/g/openhpc-announce>

This is intended to be a low-volume list that is used to announce new versions and important updates. Only members from the core development team can post to this list.

- openhpc-users
<https://groups.io/g/openhpc-users>

This list is for general questions and discussions by end users and is appropriate for subscribers to post questions, issues, suggestions related to the use of OpenHPC packages and documentation.

- openhpc-devel
<https://groups.io/g/openhpc-devel>

This list is intended for primary use by developers contributing to OpenHPC.

- Contribute code to the github repository
<https://github.com/openhpc/ohpc>
- Become a member
info@openhpc.community

Written by David Brayford

Leibniz Supercomputing Centre (LRZ), Germany

Contact: David.Brayford@lrz.de

Scaling Workshop 2017: Emergent Applications



In recent years, LRZ had regularly conducted extreme scaling workshops with the goal of scaling applications to the full SuperMUC Phase I system at 147,000 cores. This year, we focused programs that showed decent scaling within one island (512 nodes) and aimed to be run on multiple islands. Six teams applied and were invited to the “Scaling Workshop 2017: Emergent Applications” in May 2017, namely

- MPAS (D. Heinzeller, KIT), weather forecasting
- VLASOV6D (K. Reuter et al., TUM/IPP/ MPCDF), plasma physics
- ECHO (M. Bugli, MPA), simulation of accretion disks
- BFPS (M. Wilczek et al., MPDS), turbulence
- MGLET (Y. Sakai et al., TUM), CFD
- TERRA-NEO (S. Bauer et al., LMU/TUM), geophysics

During the four-day workshop, participants were supported by the LRZ application support group and experts from IBM, Lenovo, Intel, and Allinea. A special reservation on SuperMUC allowed for fast testing of code modifications on multiple islands instead of the regular 20 nodes of the test queue.

Five projects succeeded in scaling up to eight islands of SuperMUC, which comprises half of the SuperMUC Phase I system. A special highlight was the awarding of the “Leibniz Scaling Award 2017” to Matteo Bugli from the Max-Planck-Institute for Astrophysics by Prof. Kranzlmüller for his progress during the scaling work-

shop. Dr. Bugli’s ECHO project of deals with the simulation of accretion disks around neutron stars or black holes in the framework of relativistic magnetohydrodynamics. He used the software darshan to optimize the I/O and enhance the program’s performance by 18% and showed excellent scaling behaviour. A detailed report is found in this issue of InSiDE on page 86.




Prof. Kranzlmüller awards the Leibniz Scaling Award 2017 to Dr. Bugli.

**Written by Ferdinand Jamitzky,
Nicolay Hammer**

Leibniz Supercomputing Centre (LRZ), Germany

Contact: Ferdinand.Jamitzky@lrz.de



Applications

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

FS3D – A DNS Code for Multiphase Flows

The subject of multiphase flows encompasses many processes in nature and a broad range of engineering applications, such as weather forecasting, fuel injection, sprays, and spreading of substances in agriculture. To investigate these processes the Institute of Aerospace Thermodynamics (ITLR) uses the direct numerical simulation (DNS) in-house code Free Surface 3D (FS3D). The code is continuously optimized and expanded with new features and has been in use for more than 20 years.

The program FS3D was specially developed to compute the incompressible Navier-Stokes equations as well as the energy equation, with free surfaces. Complex phenomena demanding strong computational effort can be simulated because the code works on massive parallel architectures. Due to DNS, and thus resolving the smallest temporal and spatial scales, no turbulence modeling is needed. In the last years a vast number of investigations were performed with FS3D: for instance, phase transitions like freezing and evaporation, basic drop and bubble dynamics processes, droplet impacts on a thin film ("splashing"), and primary jet breakup, as well as spray simulations, studies involving multiple components, wave breaking processes, and many more.


Method

The flow field is computed by solving the conservation equations of mass, momentum, and energy in a one-field formulation on a Cartesian grid using finite volumes. The different fluids and phases are treated as a single fluid with

variable thermophysical properties that change across the interface. Based on the used Volume-of-Fluid (VOF) method additional indicator variables f_i are used to identify different phases. The VOF variables are defined as

$$f_i(x, t) = \begin{cases} 0 & \text{in the continuous phase,} \\]0, 1[& \text{in interfacial cells,} \\ 1 & \text{in the disperse phase} \end{cases}$$

and represent the different phases liquid ($i=1$), vapour ($i=2$) and solid ($i=3$). To ensure a successful advection of the VOF variable, a sharp interface, as well as its exact position, is required. This is done using the piecewise linear interface reconstruction (PLIC) method, which reconstructs a plane on a geometrical basis and, therefore, can determine the liquid and gaseous fluxes across the cell faces. The advection can be achieved with second-order accuracy by using two different methods [1]. For the computation of the surface tension several models are implemented in FS3D; for instance, the conservative continuous surface stress model (CSS), the continuum surface force model (CSF) or a balanced force approach (CSFb), which allows a significant reduction of parasitic currents. Due to the volume conservation in incompressible flow Poisson's equation of pressure needs to be solved, which is achieved by using a multigrid solver. In order to perform simulations with high spatial resolutions, FS3D is fully parallelized using MPI and OpenMP. This makes it possible to perform simulations with more than a billion cells on the supercomputer Cray-XC40 at HLRS. Some



applications of FS3D and results are presented in the following.

Applications

Freezing

Supercooled water droplets exist in liquid form at temperatures below the freezing point. They are present in atmospheric clouds at high altitude and are important for phenomena like rain, snow, and hail. The understanding of the freezing process, its parametrization, and the link to a macrophysical system such as a whole cloud is essential for the development of meteorological models.

The diameter of a typical supercooled droplet, as it exists in clouds, is on the order of 100 μm whereas the ice nucleus is in the nanometer range. This large difference in the scales requires a fine resolution of the computational grid. To capture the complex anisotropic structures that develop as the supercooled droplet solidifies, an anisotropic surface energy density is considered at the solid-liquid boundary using the Gibbs-Thomson equation. The energy equation is solved implicitly in a two-field formulation in order to remove the severe timestep constraints of solidification processes. The density of both ice and water are considered equal. This is a reasonable assumption and greatly simplifies the problem at hand. A typical setup consists of a computational grid with $512 \times 512 \times 512$ cells where the initial nucleus is resolved by roughly 20 cells. A visualization of a hexagonally growing ice particle embedded in a supercooled water droplet is shown in Fig. 1.



Fig. 1: Visualization of a hexagonally growing ice particle embedded in a supercooled water droplet.

Evaporation of supercooled water droplets

Not only freezing processes but also the evaporation of supercooled water droplets need to be understood for the improvement of meteorological models. In the presented study the evaporation rate, depending on the relative humidity of the ambient air, is in the focus of numerical investigations with FS3D.

Several simulations of levitated supercooled water droplets are performed at different constant ambient temperatures and varying relative humidities Φ , with one example shown in Fig. 2. The evaporation rate β is determined and compared to experimental measurements [4]. The setup consists of an inflow boundary on the left side, an outflow boundary on the right side, and free slip conditions on all lateral boundaries. The grid resolution is $512 \times 256 \times 256$ cells and the diameter of the spherical droplet is resolved by approximately 26 cells.




Fig. 2: Simulation of an evaporating supercooled water droplet with FS3D

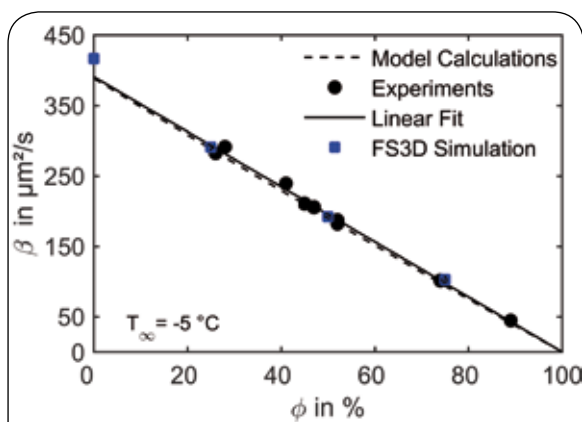


Fig. 3: Measured evaporation rates β at $T_{\infty}=268,15$ K.

The resulting dependency of the evaporation rate on the relative humidity is depicted in Fig. 3, for an ambient temperature of $T_{\infty}=268,15$ K. The numerical results agree very well with experimental data. This shows that FS3D is capable of simulating the evaporation of supercooled water droplets and therefore can help to improve models for weather forecast. For example, future numerical simulations of the evaporation of several supercooled water droplets and their

interaction could be investigated, a goal that is currently not feasible experimentally.

Non-newtonian jet break up

Liquid jet break up is a process in which a fluid stream is injected into a surrounding medium and disintegrates into many smaller droplets. It appears in many technical applications; for instance, fuel injection in combustion gas turbines, water jets for firefighting, spray painting, spray drying, or ink jet printing. In some of these cases an additional level of complexity is introduced if the injected liquids are non-Newtonian; i.e., they have a shear dependent viscosity. Due to the complex physical processes, which happen on very small scales in space and time, it is hard to capture jet break up by experimental methods in great detail. For this reason it is a major subject for numerical investigations, and therefore, for investigations with FS3D.

We are simulating the injection of aqueous solutions of the polymer Praestol into ambient air. The shear-thinning behavior is incorporated by using the Carreau-Yasuda model. The largest simulations are done on a $2304 \times 768 \times 768$ grid, using over 1.3 billion cells, where the cells in the main jet region have an edge length of $4 \cdot 10^{-5}$ m. The simulated real time is in the order of 10 ms.

We investigate the influence of different destabilizing parameters on the jet (see Fig. 4), such as the Reynolds number, the velocity profile at the nozzle or the concentration of the injected solutions (and therefore the severity of the non-Newtonian properties). We analyze the





Fig. 4: Visualization of a jet break up simulated with FS3D

influence of these parameters on the jet break up behavior, quantified by the liquid surface area, the surface waves disturbing the jet surface and the droplet size distribution [2]. We

then investigate the three-dimensional simulation data, such as the velocity field or the internal viscosity distribution, in detail to explain the differences in jet behavior (see Fig. 5).

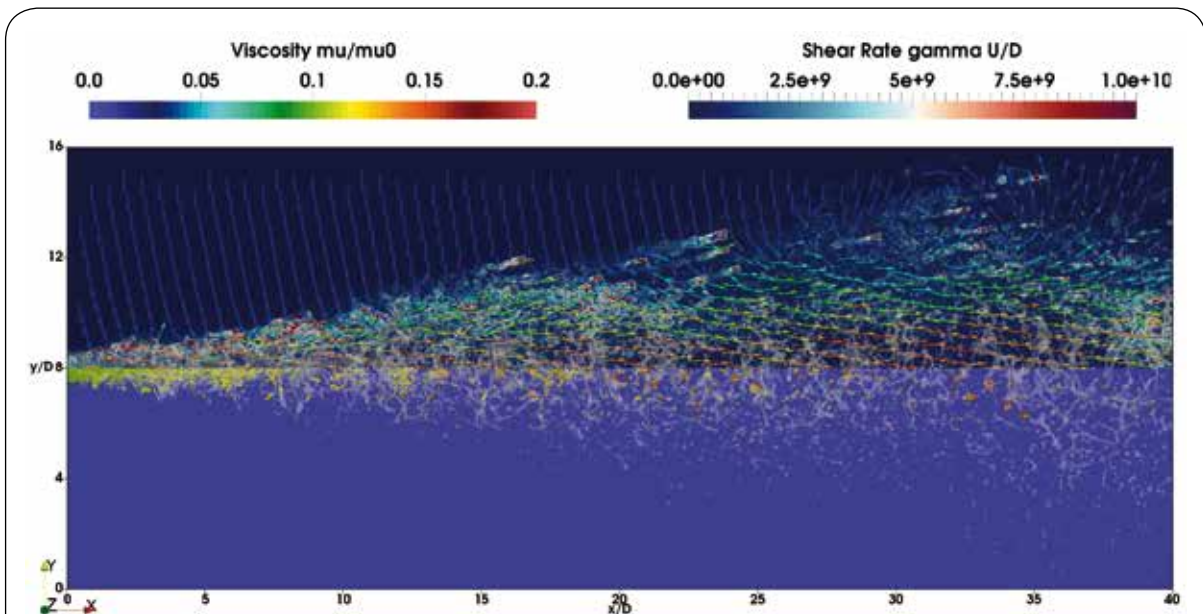


Fig. 5: Visualization of a transparent jet. In the background we show a slice through the centerline displaying the viscosity distribution on the lower half and the shear rate as well as the velocity vector on the upper half.

Wave breaking

The interaction between an airflow and a water surface influences many environmental processes. This is particularly important for the formation and amplification of hurricanes. Water waves, wave breaking processes, and entrained water droplets play a crucial role in the momentum, energy, and mass transfer in the atmospheric boundary layer.

In order to simulate a wind wave from scratch a quiescent water layer with a flat surface and an air layer with a constant velocity field are initialized. The computational domain, corresponding to one wavelength of $\lambda=5$ cm, has a resolution of $512 \times 256 \times 1024$ cells. Every simulation is performed on the Cray-XC40 at HLRS with at least several thousand processors. Due to transition, the air interacts with the water surface and a wind wave develops, shown in Fig. 6. In the first step the occurring parasitic capillary waves on the frontside of the wind wave are evaluated. Wave steepnesses and the different wave lengths of all parasitic capillary waves offer detailed insights into energy dissipation mechanisms, which could not be gained from experiments. In a second step the wind is enhanced by applying a wind stress boundary condition at the top of the computational domain. This leads to the growth of the wave amplitude and finally to wave breaking. Not only phenomenological comparison of this process with experiments, but also information about temporal evolution of the wave energy, structures in the water layer, or dynamics of vortices are remarkable results of these simulations. For future investigations




Fig. 6: Simulation of a gravity-capillary wind wave with FS3D. The water surface is visualized in the front and the turbulent velocity field of the air layer on the left and rear boundaries of the computational domain.


of wind waves and, for example, droplet entrainment from the water surface higher velocities, higher resolutions, and therefore, higher computational power will be needed. Such simulations requiring more than one billion cells makes the use of supercomputers indispensable.

Droplet splashing

If a liquid droplet impacts on a thin wall film, the resulting phenomena can be very complex. Impact velocity, droplet size and wall film thickness have a large influence on the shape and morphology of the observed crown. If the conditions are such that secondary droplets are ejected, this phenomenon is called splashing.

The splashing process is highly unsteady and its appearance is dominated by occurring instabilities that have a wide range of different scales. However, only a limited amount of properties are accessible through experiments. For example, thickness of the crown wall and velocity profiles are difficult to obtain experimentally.

Currently, we are able to perform simulations with up to one billion cells. A rendering of an exemplary simulation is shown in Fig. 7. In order to capture splashing processes on the smallest scale, a very high resolution is required. Therefore, often only a quarter of the physical domain is simulated by applying symmetry boundary conditions.



When the droplet and the wall film consist of two different liquids, additional phenomena occur that cannot be explained anymore with single-component splashing theories. One reason for this is that not only the properties of the liquids themselves but also their ratio matters.

Due to this, a multi-component module is implemented in FS3D, which captures the concentration distribution of each component within the liquid phase. This makes it possible to evaluate, for example, composition of the secondary droplets. One technical application for which this is important is the interaction of fuel droplets with the lubricating oil film on the cylinder in a diesel engine. This interaction occurs during the regeneration of the particle filter and leads to both a dilution of the engine oil wall film and to higher pollutant emissions. Here, a better understanding of two-component splashing dynamics can be a great advantage in order to minimize both engine emissions and lubrication losses.




Fig. 7: Visualization of a splashing droplet.

Acknowledgements

The FS3D team gratefully acknowledges support by the High Performance Computing Center Stuttgart over all the years. In addition we kindly acknowledge the financial support by the Deutsche Forschungsgemeinschaft (DFG) in the projects SFB-TRR75, WE2549/35-1, and SimTech.

References

- [1] **Eisenschmidt, K., Ertl, M., Goma, H., Kieffer-Roth, C., Meister, C., Rauschenberger, P., Reitzle, M., Schlottke, K., Weigand, B.:**
Direct numerical simulations for multiphase flows: An overview of the multiphase code FS3D, *Applied Mathematics and Computation*, 272, pp. 508-517, 2016.
- [2] **Ertl, M., Weigand, B.:**
Analysis methods for direct numerical simulations of primary breakup of shear-thinning liquid jets. *Atomization and Sprays* 27(4), 303–317, 2017.
- [3] **Reitzle, M., Kieffer-Roth, C., Garche, H., Weigand, B.:**
A volume-of-fluid method for three-dimensional hexagonal solidification processes, *J. Comput. Phys.* 339: 356-369, 2017.
- [4] **Ruberto, S., Reuttsch, J., Roth, N., Weigand, B.:**
A systematic experimental study on the evaporation rate of supercooled water droplets at subzero temperatures and varying relative humidity, *Exp Fluids*, 58:55, 2017.

**Written by Moritz Ertl, Jonas Kaufmann,
Martin Reitzle, Jonathan Reuttsch, Karin
Schlottke, Bernhard Weigand**

Institute of Aerospace Thermodynamics (ITLR),
University of Stuttgart


Contact: karin.schlottke@itlr.uni-stuttgart.de

Distribution Amplitudes for η and η' Mesons

Theoretical background

The development of Quantum Field Theory (QFT) was without a doubt one of the greatest cultural achievements of the 20th century. Within its proven range of applicability, its predictions will stay valid as long as the universe exists within quantified error ranges given for each specific quantity calculated. In light of this far-reaching perspective, great effort is invested to improve the understanding of every detail. QCD, which describes quarks, gluons, and their interactions and thus, most properties of the proton and neutron, is a mature theory which nevertheless still holds many fascinating puzzles. Therefore, present-day research addresses often quite intricate questions difficult to explain in general terms. Unfortunately, this is also the case here. Highly advanced theory is needed to really explain the underlying theoretical concepts and the relevance of the specific calculations performed, which can, therefore, only be sketched in the following: QCD, like all QFTs realized in nature, is a gauge theory, or a theory whose experimentally verifiable predictions are unchanged if, for example, all quark wave functions are modified by matrix-valued phase factors which can differ for all space-time points. In fact, nearly all properties of QCD can be derived unambiguously solely from this property and Poincaré symmetry (the symmetry associated with the special theory of relativity). The matrix properties of these phase factors are completely specified within a classification of group theory which was already completed in the 19th century. Within this classification its invariance properties with respect to the “color”

of quarks are named SU(3)-symmetry. SU(3) has SU(2) subgroups and SU(2) is isomorphic to the group of spatial rotations in three space dimensions (this is why spin and orbital angular momentum have very similar properties). This implies the existence of infinitely many distinct QCD vacuum states which differ by the number of times all SU(2) values occur at spatial infinity when all spatial directions are covered once. Mathematically, these different “homotopy classes” are characterized by a topological quantum number which is also equivalent to the local topological charge density, see Fig.1, integrated over the whole lattice. While all of this might sound pretty abstract and academic, it can actually have very far reaching practical consequences. Still reflecting the bafflement these facts created about 50 years ago, these effects are called “anomalies.” In this specific case one speaks of the “axial anomaly.” By now anomalies are completely understood mathematically. In a nutshell one can say that symmetries of a classical theory can be violated when the theory is quantized leading typically to additional, often surprising, consistency conditions. After the complete theoretical understanding of these features was achieved, anomalies became, actually, one of the most powerful tools of QFT. The requirement that fundamental symmetries of the classical theory have to be preserved, implies, for example, that only complete families of fermions—e.g. consisting in case of the first particle family of the electron, the electron neutrino and three variants of the up and down quarks—can exist. In a similar manner the absence of unacceptable anomaly-induced



effects requires supersymmetric string theories to exist in 1 time and 9 space dimensions. In a way, these are modern physicists versions of Kant's synthetic a priori judgments: Mathematical consistency alone implies certain fundamental structures of physics. The properties of the η , η' meson system are affected by one of these anomalies in a non-catastrophic—i.e. acceptable—manner and are thus perfectly suited to test our understanding of the sketched involved properties. A final level of complication is added by the fact that the mass eigenstates η and η' are quantum mechanical superpositions of the "flavor" singlet and octet states of which only the singlet state is affected by the anomaly. Thus, one of the tasks is to determine the mixing coefficients more precisely.





Fig.1: Topological charge density (red: positive, blue: negative) for one of our field configurations from one of our quantum field ensembles. The typical length scale of these structures is 0.5 fm. The lattice spacing used for this configuration is 0.085 fm, i.e. the structures are clearly resolved.

The numerical approach

Unfortunately, the fundamental concepts of lattice QCD are also mathematically highly non-trivial.

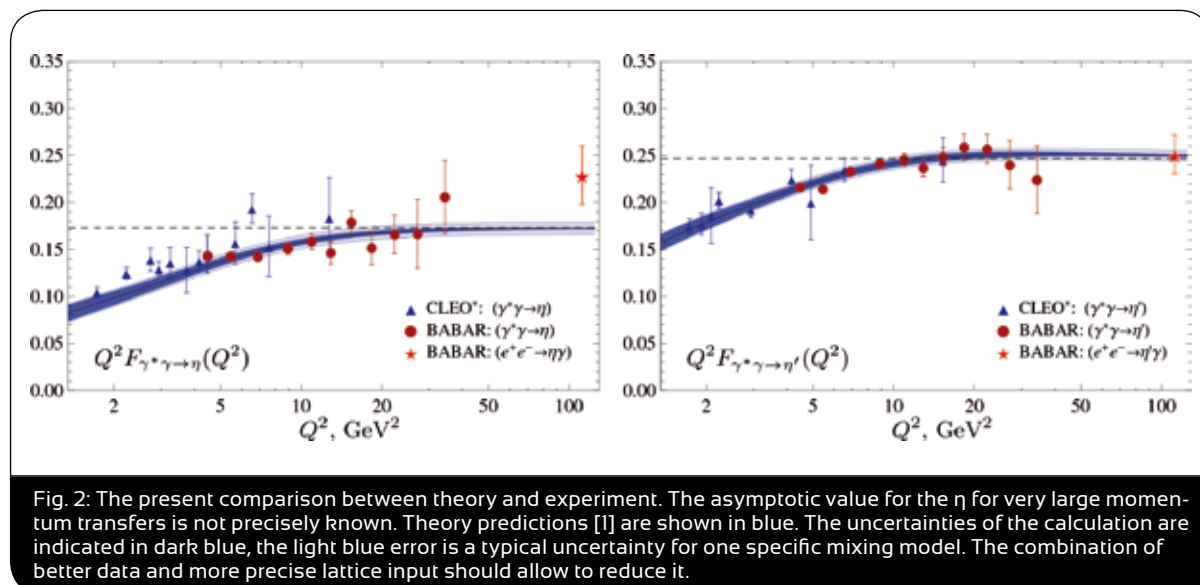

Analysis, the mathematical discipline, allows for the analytic continuation of functions of real variables to functions of complex variables and back. In lattice QCD the whole formulation of QFT is analytically continued from real time to imaginary (in the sense of square root of -1) time. Because QFT is mathematically exact this is possible just as for other functions. Somewhat surprisingly, this mathematical operation maps QFT onto thermodynamics such that problems of quantum field theory become solvable by stochastic algorithms which are perfectly suited for numerical implementation. Because the number of degrees of freedom is proportional to the number of space-time points, to do so, the space time continuum is substituted by a finite lattice of space-time points, such that the quantities to be evaluated are extremely high but finite dimensional integrals, which are computed with Monte Carlo techniques. This gave the method its name. In the end all results have to be extrapolated to the continuum; i.e., to vanishing lattice spacing. To guarantee ergodicity when sampling the states with different topological quantum number, the "topological autocorrelation time" (i.e., the number of Monte Carlo updates needed before another topological sector gets probed) must be much smaller than the total simulation time. Unfortunately, in previous simulations using the standard periodic boundary conditions one has observed a diverging topological autocorrelation time when the lattice spacing is reduced, precluding a controlled continuum extrapolation. As a remedy the CLS collaboration, to which we belong, has started large scale simulations with open—i.e.

not periodic—boundary conditions which allows topological charge to leave or enter the simulation volume and thus solves the sketched problem. The price to pay is that simulated regions close to the open boundaries are strongly affected by lattice artefacts such that the fiducial volume is reduced and the computational cost increases accordingly, typically by roughly 30% for the presently used simulation volumes. However, as topology is crucial for the investigated properties, this overhead is well justified. With the sketched techniques ergodic ensembles of field configurations are generated on which the quantities of interest are then calculated. To do so reliably, many additional steps are necessary which will not be explained except for one: In the continuum, quantum fluctuations lead to divergences which have to be “renormalized” to get physical results. On any discretized lattice the renormalization factors differ from their continuum values by finite conversion factors. These factors also have to be determined numerically.



Distribution amplitudes

All experiments involving hadrons (i.e. bound states of quarks and gluons) are parameterized by a large assortment of functions, each of which isolates some properties of its extremely complicated many particle wave function. The latter ones are chosen specifically for the type of interactions which are studied experimentally. Collision experiments in which all produced particles are detected, so called “exclusive” reactions, are typically parameterised by Distribution Amplitudes. The production of η or η' in electron-positron collisions is the theoretically best understood exclusive reaction and should thus be perfectly suited to determine these DAs. Very substantial experimental efforts were undertaken to do so, especially by the BaBar experiment at the Stanford Linear Accelerator Center. Unfortunately, the result is somewhat inconclusive, showing a 2σ deviation for η production at large momentum transfer Q (see Fig. 2), where the agreement between theory and experiment should be perfect. Here, the reaction probability is parameterised by a function F which is defined in such a way that for large Q values the experimental data points should be independent of Q , which might or might not be the case. Clarifying the situation was one of the motivations for building a 40 times more intensive collider in Japan and upgrading the Belle experiment there. The task of lattice QCD is to produce predictions with a comparable precision, such that both taken together will allow for a much more precise determination of η/η' mixing and the effects caused by the axial anomaly. This is what we are providing.



Up to now we have only analyzed a small fraction of our data. Fig. 3 shows one of the calculated lattice correlators, which are the primary simulation output directly related to the DAs. The different correlators differ by quark type (light, i.e. up or down, quarks and strange quarks). In contrast to earlier work [2] we can avoid any reference to chiral perturbation theory or other effective theories or models, which should reduce the systematic uncertainties. Note that the data points are strongly correlated; i.e., all curves can shift collectively within the size of the error bars. Our final precision should be substantially better. Then a combined fit and extrapolation of all lattice data—for all ensembles—will provide the DAs we are interested in. Together with the expected much improved experimental data this should finally test how the axial anomaly affects the structure of the η and η' mesons. Additional information can be obtained from analyzing decays of, for example, D_s mesons into η and η'

mesons, see [3]. Let us add that even the DA of the most common hadrons like the proton, neutron, or pion are not well-known. This is primarily due to the fact that the investigation of hard exclusive reactions is experimentally harder, such that precision experiments only became feasible with the extremely high-intensity colliders build in the last decades. Their experimental and theoretical exploration can, therefore, be expected to be a most active field in future. We expect that the methods we optimize as part of this project will thus find a wide range of applications in the future.

References

- [1] S.S. Agaev, V.M. Braun, N. Offen, F.A. Porkert and A. Schäfer: "Transition form factors $\gamma^* \gamma \rightarrow \eta$ and $\gamma^* \gamma \rightarrow \eta'$ in QCD" Physical Review D 90 (2014) 074019 doi:10.1103/PhysRevD.90.074019 [arXiv:1409.4311 [hep-ph]].

[2] C. Michael, K. Ottnad and C. Urbach
[ETM Collaboration]:
"η and η' mixing from lattice QCD"
Physical Review Letters 111 (2013) 181602
doi:10.1103/PhysRevLett.111.181602
[arXiv:1310.1207 [hep-lat]].

[3] G.S. Bali, S. Collins, S. Dürr and I. Kanamori:
"D_s → η, η' semileptonic decay form factors with
disconnected quark loop contributions",
Physical Review D 91 (2015) 014503
doi:10.1103/PhysRevD.91.014503
[arXiv:1406.5449 [hep-lat]].





Fig.3: Just one out of very many lattice results showing the correlators calculated on the lattice. After combining all of them for an extrapolation to the physical masses, infinite volume and vanishing lattice constant, the result will provide the DAs we are interested in.

Performance Optimization of a Multiresolution Compressible Flow Solver



Currently, biotechnological and biomedical procedures such as lithotripsy or histotripsy are used successfully in therapy. In these methods, compressible multiphase flow mechanisms, such as shock-bubble interactions are utilized. However, the underlying physics of the processes involved are not fully understood. To get deeper insights into these processes, numerical simulations are a favorable tool. In recent years, powerful numerical methods which allow for accurately simulating discontinuous, compressible multiphase flows have been developed. The immense numerical cost of these methods, however, limits the range of applications. To simulate three-dimensional problems, modern high-performance computing (HPC) systems are required and need to be utilized efficiently in order to obtain results within reasonable times. The sophisticated simulation environment "ALIYAH," developed at the Chair of Aerodynamics and Fluid Mechanics, combines advanced numerical methods—including Weighted Essentially Non-Oscillatory (WENO) stencils and sharp-interface treatment

(Level-Set) in a Multiresolution Finite-Volume Framework with Total-Variation-Diminishing (TVD) Runge-Kutte (RK) time integration—to solve the Euler equations for compressible multiphase problems.

Exemplarily, the simulation result of a collapsing gas bubble near a deformable gelatin interface is shown in Figure 1. This configuration mimics the dynamics of an ultrasound-induced gas bubble near soft tissue as model for *in vivo* cavitation effects. The bubble collapse is asymmetrical and induces a liquid jet towards the gelatin that eventually ruptures this material. The detailed understanding of such phenomena is the overall scope of our research.

The baseline version of ALIYAH runs a block-based MR algorithm as described in [5]. The code is shared-memory parallelized using Intel Threading Building Blocks (TBB). The performance crucial (parallelizable) loops are distributed among the threads using the TBB affinity partitioner. Thus, the load is dynamically




Fig.1: Bubble collapse near deformable Gelatin interface: Interface visualization from simulation with ALIYAH.

re-evaluated every time the algorithm reaches a certain function.


Much of the computational cost in the considered simulation comes from the modeling of the interface between fluids. In our approach the interface is modeled by a conservation ensuring scalar level set function [1], and the interactions across the interfaces need to be considered; this is done with an acoustic Riemann solver which includes a model for surface tension [3]. For the non-resolvable structures—i.e., droplets, bubbles, or filaments with diameters close to the cell size of the finite volume mesh—scale separation of [4] is used.

Performance and scalability test cases

The simulation tests were performed for two cases: A small generic case (“synthetic case”), which executes all methods described in the previous section but with a coarse resolution of only 4096 cells, and the second case (“restart case”), which is a real-application case with a high resolution in all three spatial dimensions. Due to its long run time, only one timestep of this case is analyzed.

The restart case scenario uses an axis-symmetric model, to simulate cylindrical channel geometries in a Cartesian grid. The simulation is conducted with a quarter-model of the full problem; i.e., the Y- and Z-planes are cut into halves with imposed symmetry conditions. Since a full simulation’s runtime is too large to be profiled, the measurements are obtained for just one timestep on the coarsest level. To still


capture a relevant and representative timestep, the simulation is advanced until time $t_s = 3.16\mu s$ without profiling the code. The corresponding physical state of the bubble break-up is shown in Figure 1.



Code analysis

We conduct our analysis and optimization on a dual-socket Intel Xeon E5-2697 v3 (codenamed Haswell). Computational results are presented for an Intel Haswell system at 28 cores. The processor has 2.6 GHz frequency, 32 KB/256 KB L1/L2 caches and 2.3 GB RAM per core.

With the baseline version of the code the two test-cases—restart case and synthetic case, described above—were simulated in a wall clock time of 589 seconds and 666 seconds, respectively.



To find promising starting points for code optimization, a node-level analysis is performed using the Intel VTune Amplifier. To reduce the amount of collected information the Amplifier analysis as well as all subsequent optimization runs are performed using eight threads. The hotspot analysis for the restart case is presented in Figure 2 and for the small synthetic case in Figure 3.

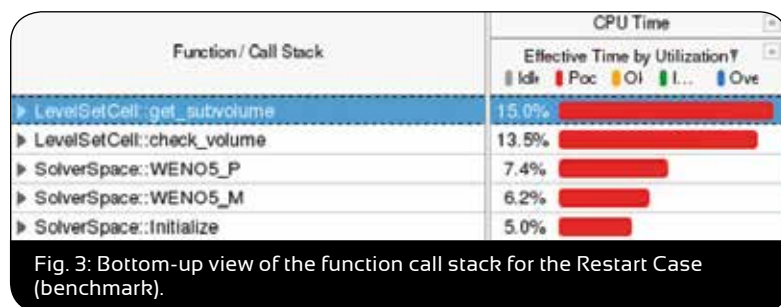


Fig. 3: Bottom-up view of the function call stack for the Restart Case (benchmark).

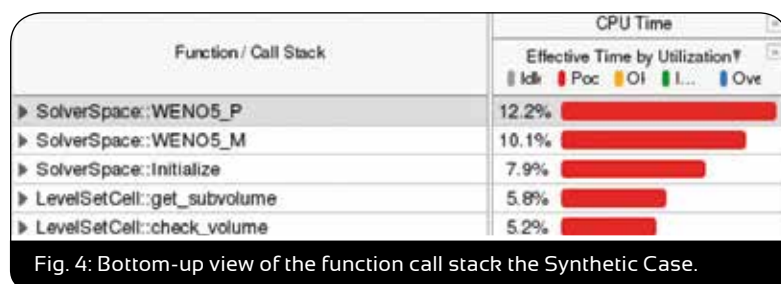


Fig. 4: Bottom-up view of the function call stack the Synthetic Case.

One can clearly identify the functions `get_subvolume`, `check_volume`, and `WENO5_*` as the hotspots. The optimization of `WENO5_*` requires only small reorganization of the corresponding source code. In contrast to the WENO methods, the time spent in the `get_subvolume` function does not increase linearly with the problem size (c.f., relative time spent for the small synthetic case and the larger restart case).

Hence, a focus is laid on the non-straight-forward optimization of the `get_subvolume` and `check_volume` functions.

An essential ingredient to utilize HPC architectures efficiently is the usage of single instruction multiple data (SIMD) instructions in the computationally intensive parts of the code. SIMD instructions allow processing of multiple pieces of data in a single step, speeding up throughput for many tasks. Compilers can auto-vectorize loops that are considered safe for vectorization. In the case of the here-used Intel compiler version 16.0, this happens at default for optimization levels `-O2` or higher.

To analyze the auto-vectorized code the Intel Advisor XE tool is used. The analysis revealed the functions listed in Figure 4 to be the most time consuming non-vectorized ones. In the figure, "self time" represents the time spent in a particular program unit and "total time" includes "self time" of the function itself and "self time" of all functions that were called from within this function.

As seen, the function `get_subvolume`, which is called recursively from the function `get_volume`, is the most time-consuming non-vectorized function. In contrast to compilers assumption, the examination of `get_subvolume`'s source code reveals no crucial dependency problems.

Function Call Sites and Loops	Vector Issues	Self Time	Total Time
[loop in LevelSetCell::get_subvolume at level_set_cell.cpp:602]	1 Assumed dependency prese...	231.177s	3810.722s
[loop in TBB Dispatch Loop at custom_scheduler.h:413]		228.349s	231.398s
[loop in Meshlevel::Update_tags(Meshlevel&)::lambda(tbb::bl...	1 Assumed dependency prese...	126.556s	126.556s
[loop in SolverSpace::Initialize at solver_space.cpp:182]	1 Opportunity for outer loop ve...	125.308s	125.308s
[loop in BlockAndPackageTags::UpdatePackageInterfacePro...	1 Assumed dependency prese...	68.673s	68.673s
[loop in BlockAndPackageTags::UpdatePackageInterfacePro...	1 Assumed dependency prese...	60.704s	60.704s
[loop in LevelSetCell::get_subvolume at level_set_cell.cpp:601]	1 Assumed dependency prese...	45.179s	3855.900s
[loop in LevelSetCell::get_subvolume at level_set_cell.cpp:600]	1 Assumed dependency prese...	39.763s	3895.664s
[loop in LevelSetNode::GetIncrementForReinitialization at lev...	3 Assumed dependency prese...	38.644s	38.644s

Fig. 5: Survey analysis of the vectorization in the baseline version of ALIYAH.


Results

Since it is a recursive call automatic vectorization or OpenMP-SIMD, annotations cannot be applied directly to the body of the function `get_subvolume`. Moreover, due to the presence of the relatively large amount of nested loops with small trip counters the declaration of `get_subvolume` as “vectorizable” is not an optimal strategy in this case. On Haswell, SIMD instructions process four elements (double precision) at once. This means loops with a trip counter of two underutilize the vector registers by a factor of two. It appears OpenMP-SIMD is not able to collapse the two nested loops and apply vectorization automatically. As auto-vectorization fails even with the usage of OpenMP pragmas we follow the more aggressive approach, described below.

The function `get_subvolume` performs temporary subdivisions of the cubic grid cells based on linear interpolation to approximate the volume one phase occupies. Due to the recursive call with a local stopping criterion the data flow in each local volume evaluation

is complex. To apply SIMD vectorization, we combine linear interpolation on several elements into one call. This is profitable since the operation on two neighbor grid points is the same, albeit with different data from the vector. We program vectorized loops directly using Intel AVX instructions.

The explicit SIMD vectorization with intrinsics allows us to reduce the number of micro-operations from 185 for the baseline version down to 88. The block throughput is also reduced from 48 cycles to 24 cycles. The total time spent in the `get_subvolume` function is reduced by a factor of 0.7, which means a gain in performance of 40%. CPU time of the two functions `get_subvolume` and `check_volume` after optimization is reduced by a factor of 0.5 compared to the baseline version. Moreover, the wallclock time of the AVX version is reduced to 531 sec and 558 sec for the restart case and the synthetic case, respectively. For the whole simulation this corresponds to a speedup of 11% for the restart case and 19% for the synthetic cases, correspondingly.



Acknowledgment

The authors gratefully acknowledge the Kompetenznetzwerk für wissenschaftliches Höchstleistungsrechnen in Bayern for the KONWIHR-III funding. S. Adami and N.A. Adams gratefully acknowledge the funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement No 667483).

References

- [1] **X. Y. Hu, B. C. Khoo, N. A. Adams, and F. L. Huang:**
"A conservative interface method for compressible flows," *J. Comput. Phys.*, vol. 219, no. 2, pp. 553–578, Dec. 2006.
- [2] **R. P. Fedkiw, T. D. Aslam, B. Merriman, and S. Osher,**
"A Non-oscillatory Eulerian Approach to Interfaces in Multimaterial Flows (The Ghost Fluid Method)," *J. Comput. Phys.*, vol. 152, pp. 457–492, 1999.
- [3] **R. Saurel, S. Gavriluk, and F. Renaud:**
"A multiphase model with internal degrees of freedom: application to shock–bubble interaction," *J. Fluid Mech.*, vol. 495, pp. 283–321, 2003.
- [4] **J. Luo, X. Y. Hu, and N. A. Adams:**
"Efficient formulation of scale separation for multi-scale modeling of interfacial flows," *J. Comput. Phys.*, vol. 308, pp. 411–420, Mar. 2016.
- [5] **L. H. Han, X. Y. Hu, and N. A. Adams:**
"Adaptive multi-resolution method for compressible multi-phase flows with sharp interface model and pyramid data structure," *J. Comput. Phys.*, vol. 262, pp. 131–152, Apr. 2014.

**Written by Nils Hoppe¹, Igor Pasichnyk²,
Stefan Adami¹, Momme Allalen³, and
Nikolaus A. Adams¹**

¹ Lehrstuhl für Aerodynamik und Strömungsmechanik,
Technische Universität München,
Boltzmannstraße 15, 85748 Garching

² IBM Deutschland GmbH,
Boltzmannstraße 1, 85748 Garching

³ Leibniz-Rechenzentrum der Bayerischen Akademie
der Wissenschaften,
Boltzmannstraße 1, 85748 Garching

Contact: momme.allalen@lrz.de

Performance Evaluation of a Parallel HDF5 Implementation to Improve the Scalability of the CFD Software Package MGLET

This paper presents a performance evaluation for an implementation in parallel HDF5 inside the MGLET code.


The computational fluid dynamics (CFD) code “MGLET” is designed to precisely and efficiently simulate complex flow phenomena within an arbitrarily shaped flow domain. MGLET is capable of performing direct numerical simulation (DNS) as well as large eddy simulation (LES) of complex turbulent flows. It employs a finite-volume method to solve the incompressible Navier–Stokes equations for the primitive variables (i.e. three velocity components and pressure), adopting a Cartesian grid with staggered arrangement of the variables. The time integration is realised by an explicit third-order low-storage Runge–Kutta scheme. The pressure computation is decoupled from the velocity computation by the fractional time-stepping, or Chorin’s projection method. Consequently, an elliptic Poisson equation has to be solved for each Runge–Kutta sub-step.

The current version of MGLET utilises a parallel adaptation of Gauss–Seidel solver as well as Stone’s Implicit Procedure (SIP) within the multigrid framework, both as the smoother during the intermediate steps and the solver at the coarsest level. Such separate usage is justified by the fact that the former is very effective in eliminating low-frequency error predominant over the successive coarsening stage of multigrid algorithms, whereas the latter can be used to solve the Poisson problem at the coarsest level with a broad spectrum of residual error.

Geometrically complex surfaces, or arbitrarily curved, can be represented by an immersed boundary method (IBM). MGLET offers several sub-grid scale models for LES simulation, such as Smagorinsky’s model, two versions of the dynamic formulations and the WALE model. MGLET is written in FORTRAN and the parallelisation strategy is based on Message Passing Interface (MPI).

The code is currently being used by several research groups: At the Chair of Hydromechanics of the Technical University of Munich, for instance, turbulent flow through complex geometries, flow in porous media, and fibre suspensions in fluid media have been investigated using MGLET. The groups of Prof. Helge Andersson and Prof. Bjørnar Pettersen (both NTNU Trondheim) use the code to predict and analyse bluff-body flows primarily using DNS and IBM. At the Institute for Atmospheric Physics (DLR Oberpfaffenhofen), aircraft wake vortices are investigated, including their interaction with atmospheric boundary layers and ground effects. These applications demonstrate MGLET’s excellent numerical efficiency and adaptability to the diverse hydrodynamic problems.

Continuous improvement of its parallel scalability has been, and will remain, critically important for the MGLET development programme, as it allows us to simulate ever-more realistic and engineering-relevant turbulent flows at an adequate resolution of motion. For example, there is a trend towards higher Reynolds numbers, more



complex flow configurations and the inclusion of micro-structural effects such as particles or fibres. The simulation so far that has used the largest number of degrees of freedom is the one simulating a fully turbulent channel flow of a fibre suspension, realised by approximately 2.1 million cells, 66 million Lagrangian particles and 100 fibres [1]. This simulation is the only one currently published that used a full micro-mechanical model for the fibres' orientation distribution function without closure. More recently, we simulated turbulent flow around a wall-mounted cylinder with a Reynolds number up to 78,000 (the results are partially published in [2]), where the utilised number of cells was increased up to approximately 1.9 billion.

In recent years, MGLET has undergone a series of major parallel performance improvements, mainly through the revision of the MPI communication patterns, and the satisfactory scalability of the current version of the code had been confirmed up to approximately 7,200 CPU cores, which is roughly equivalent to the number of cores in one island of the SuperMUC Phase 1 Thin Node at Leibniz Supercomputing Centre (LRZ). As MGLET's parallel scalability improves significantly, however, its I/O performance became progressively the main performance bottleneck, stemming from the fact that the current implementation is entirely serial: the master MPI rank is solely responsible for collecting/distributing data from/to the other ranks. We




Fig. 1: Implemented HDF5 file structure design to store physical field values and grid information. Circles represent HDF5 groups and squares represent datasets. Datasets "Runinfo" and "Gridinfo" store the global header information, whereas other datasets (e.g. U, V, W etc.) are the simulated physical data. Only representative datasets are shown in this figure.

decided to resort to the parallel HDF5 I/O library to overcome the performance bottleneck.

Consequently, we will discuss the implementation details and the results of the performance evaluation and scalability analysis of the new parallel I/O module as the main focus of this paper. Before proceeding any further, it is important to note that this work is funded by KONWIHR (Bavarian Competence Network for Technical and Scientific HPC), which is gratefully acknowledged by the authors.

Parallel I/O implementation using HDF5


The implementation of a new parallel I/O module has been divided into two parts: 1) The I/O related to instantaneous and time-averaged field data; and 2) the data related to immersed boundary (geometry) information. In this contribution, we exclusively discuss our work related to the first part.

Figure 1 shows the file structure that was adopted in the current implementation, where each circle and square represents a HDF5 group and dataset respectively. In this design, the master process writes the global header information to the output file, whereas individual processes write the physical data that are local to their memory in a collective manner.

Experimental evaluation

The new implementation was evaluated in SuperMUC Phase 1. A series of I/O weak-scaling tests was conducted, and showed a consistent factor of 5 speed-up in comparison to

the original serial I/O implementation. Figure 2 shows the data transfer rate for such test case with 512000 cells per MPI process. Despite the significant improvement, however, we observed a noticeable drop in the I/O performance utilising more than 1 island (i.e. 8192 cores).



In order to identify the cause of such performance degradation, an I/O profiling analysis was conducted using the scalable HPC I/O characterisation tool, Darshan. By analysing the request size for collective operations, we noticed that the operations at the POSIX level were done in size of 512 KiB, which is unfavourably small by considering the large I/O operation overhead present in any large-scale parallel file systems. To circumvent this behaviour, we explicitly instructed the I/O library to exploit the collective buffering techniques through the ROMIO hints.





Fig. 3: Data transfer rate at MPI-IO and POSIX-IO level for new HDF5 implementation with enabled collective buffering technique.

Figure 3 shows the results from the same weak-scaling tests as before, but with the collective buffering technique being enabled. First of all, notice that the data transfer rate improved significantly by the modification: the peak performance increased from ≈ 1.2 GiB/sec to ≈ 4.2 GiB/sec at the POSIX level, while it increased from ≈ 1.1 to ≈ 2.2 GiB/sec at the MPI level. Second of all, the gap between the POSIX and the MPI level was widened. Finally, the improved version still suffers from a performance drop observed between 2 (16,384 cores) and 4 islands (32,768 cores). A further analysis showed that this phenomenon can be related to the metadata operations, which are currently performed by the master rank only. This is a known limitation of version 1.8.X of the HDF5 library. Currently, we are testing the newest version 1.10, which allows us to perform the metadata operations collectively in parallel.

Conclusion

The implementation of the new parallel I/O module in the CFD software MGLET was discussed, and the results from the initial performance evaluations were presented. Upon the evaluation, an I/O performance degradation was detected when the employed number of MPI processes exceeded around 8,000. To identify the cause of the performance drop, an in-depth analysis was performed by using Darshan, and it was found that enabling the collective buffering technique boosts the performance drastically, even more than factor of 2 improvement. Consequently, the peak scaling limit of the new I/O module was shifted towards higher limit, somewhere between 16,000 and 33,000

MPI processes, and further analysis is in progress to push the scaling limit even further.

References

- [1] **A. Moosaie and M. Manhart:**
Direct Monte Carlo simulation of turbulent drag reduction by rigid fibers in a channel flow. *Acta Mechanica*, 224(10):2385–2413, 2013.
- [2] **W. Schanderl and M. Manhart:**
Reliability of wall shear stress estimations of the flow around a wall-mounted cylinder. *Computers and Fluids* 128:16–29, 2016.

**Written by Y. Sakai¹, S. Mendez²,
M. Allalen², M. Manhart¹**

¹Chair of Hydromechanics, Technical University of Munich, Munich, Germany

²Leibniz Supercomputing Centre, Munich, Germany
Contact: momme.allalen@lrz.de

ECHO-3DHPC: Relativistic Magnetized Disks Accreting onto Black Holes

Accretion of magnetized hot plasma onto compact objects is one of the most efficient mechanisms in the Universe in producing high-energy radiation coming from sources such as active galactic nuclei (AGNs), X-ray binaries (XRBs) and gamma-ray bursts (GRBs), to name a few. Numerical simulations of accretion disks are therefore of paramount importance in modeling such systems, as they enable the detailed study of accretion flows and their complex structure.

However, numerical calculations are subject to serious constraints based on the required resolution (and hence computational cost). The presence of magnetic fields, which are a fundamental ingredient in current models of accretion disks, can play an especially strong role in setting characteristic length-scales much smaller than the global size of a typical astrophysical flow orbiting around a black hole.

In order to afford multiple high-resolution simulations of relativistic magnetized accretion disks orbiting around black holes, in the last three years we established collaborations with different HPC centres, including the Max Planck Computation and Data Facility (MPCDF) and, in particular, the team of experts of the AstroLab group at the Leibniz Supercomputing Center (LRZ). Here we present the main achievements and results coming from these interdisciplinary efforts.

The code

Our calculations are carried using an updated version of the ECHO (Eulerian Conservative High Order) code [5], which implements a grid-based



finite-differences shock-capturing scheme to integrate the general relativistic magnetohydrodynamics (GRMHD) equations. The GRMHD approximation is widely used in the study of relativistic magnetized plasma, and the code's versatility allows for the numerical study of a wide range of astrophysical problems, such as pulsar wind nebulae [6], neutron star's magnetospheres [7] and magnetic reconnection [4].

The updated version includes a particular prescription for Ohm's law, i.e. the relation that defines the property of the electric field in a conducting fluid, in order to take into account the turbulent dissipation and amplification of magnetic fields that can naturally occur in various astrophysical sites [1, 2].

Parallelization and I/O

The main improvement to the original version of ECHO has been achieved in the parallelization scheme, which was extended from the original one-dimensional MPI decomposition to a multi-dimensional one. This allows one, for any given problem size, to make use of a larger number of cores, since it results in a larger ratio between the local domain volume and the volume of data that needs to be communicated to neighbor processes. The runtime of a typical three-dimensional simulation can therefore be reduced even by a factor of 100.

A proof of the parallel efficiency of this strategy can be evinced from Fig. 1, as the code shows an extremely good strong and weak scaling up to 8 islands (i.e. 65536 cores) on SuperMUC Phase 1.




Moreover, the use of the Intel Profile-guided Optimizations (PGO) compiling options led to an additional speed-up of about 18%.

Another important feature is the use of the MPI-HDF5 standard, which allows for a parallel management of the I/O and hence significantly cuts the computational cost of writing the output files.

Results and perspectives

Despite the wide range of astrophysical problems investigated using the ECHO code, it is in the context of relativistic accretion disks that the first three-dimensional simulations were conducted. By exploiting the vast improvements in the code's parallelization schemes, we were able to conduct a study on the stability of three-dimensional magnetized tori and investigate the development of global non-axisymmetric modes [3]. In the hydrodynamic case, thick accretion disks are prone to develop the so-called Papaloizou-Pringle instability (PPI, see left panel of Fig. 2, which leads to the formation of a smooth large-scale overdensity and a characteristic $m=1$ mode. However,



adding a weak toroidal magnetic field triggers the growth of the magnetorotational instability (MRI), which drives MHD turbulence and prevents the onset of the PPI (right panel of Fig. 2). This result holds as far as the resolution of the simulation is high enough to capture the dynamics of the small-scale fluctuations in the plasma: when the numerical dissipation increases, the PPI can still experience significant growth and not be fully suppressed by a less-effective MRI. An excess of numerical diffusion can hence lead to qualitatively different results, proving how crucial it is to conduct these numerical experiments with an adequate resolution.

In the near future the code will undergo an additional optimization through the implementation of a hybrid OpenMP-MPI scheme, that will allow for a better exploitation of the modern Many Integrated Core architectures (MIC) that supercomputing centres as LRZ currently offer. This new version will be employed in investigating the role of magnetic dissipation in shaping the disk's structure and affecting the efficiency of the MRI, leading to a deeper understanding of the fundamental physical processes underlying accretion onto astrophysical compact objects.

References

- [1] **N. Bucciantini and L. Del Zanna:**
"A fully covariant mean-field dynamo closure for numerical 3+1 resistive GRMHD".
MNRAS , 428:71-85, 2013.
- [2] **M. Bugli, L. Del Zanna, and N. Bucciantini:**
"Dynamo action in thick discs around Kerr black holes: high-order resistive GRMHD simulations".
MNRAS , 440:L41-L45, 2014.
- [3] **M. Bugli, J. Guilet, E. Mueller, L. Del Zanna, N. Bucciantini, and P. J. Montero:**
"Papaloizou-Pringle instability suppression by the magnetorotational instability in relativistic accretion discs". ArXiv e-prints, 2017.
- [4] **L. Del Zanna, E. Papini, S. Landi, M. Bugli, and N. Bucciantini:**
"Fast reconnection in relativistic plasmas: the magnetohydrodynamics tearing instability revisited".
MNRAS, 460:3753-3765, 2016.
- [5] **L. Del Zanna, O. Zanotti, N. Bucciantini, and P. Londrillo:**
"ECHO: a Eulerian conservative high-order scheme for general relativistic magnetohydrodynamics and magnetodynamics". A&A, 473:11-30, 2007.
- [6] **B. Olmi, L. Del Zanna, E. Amato, and N. Bucciantini:**
"Constraints on particle acceleration sites in the crab nebula from relativistic magnetohydrodynamic simulations". MNRAS , 449:3149-3159, 2015.
- [7] **A. G. Pili, N. Bucciantini, and L. Del Zanna:**
"Axisymmetric equilibrium models for magnetized neutron stars in general relativity under the conformally flat condition". MNRAS , 439:3541-3563, 2014.



A success story on the close collaboration between domain scientists and IT experts

Environmental ecosystems are among the most complex research topics for scientists: Not only because of the fundamental physical laws, but also because of the convoluted interactions of basically everything that surrounds us. Consequently, no environmental ecosystem can be understood by itself. A deep understanding of the environment must revolve around complex coupled systems. To this end, modern environmental scientists need to collect vast amounts of data, process this data efficiently, then develop appropriate models to test their hypotheses and predict future developments. In support of this endeavour, the LRZ has started its Environmental Computing initiative. In close collaboration with domain scientists, the LRZ supports this research field with modern IT resources and develops new information systems that will eventually benefit scientists from many other domains.

The fundamental physical laws are well understood in the context of environmental sciences. But a concrete description of an environmental ecosystem is difficult and complex. It is important to understand that different environmental systems interact in different ways. This requires multi-physics, multi-scale and multi-model workflows. Scientists are developing procedures to describe such systems numerically with the goal to understand the environment, including natural hazards and risks. The commonly used models are developed by domain scientists for other researchers in their field. These models often require detailed configurations and setups, which poses a huge

challenge: Natural disasters strike quickly with often little advance warning. To empower decision makers responsible for environmental protection and the mitigation authorities responsible for environmental protection, hazard mitigation or disaster response, they need access to fast, reliable and actionable data. Hence, the models developed by environmental scientists in recent years need to be put into operational services. This requires a very close collaboration between scientists, authorities, and IT service providers. In this context, we—the Leibniz Supercomputing Centre (LRZ)—started the Environmental Computing initiative. Our goal is to learn from scientists and authorities, support their IT needs, jointly develop services, and foster the knowledge transfer from academia to the authorities.

Within the last year, this effort has led to several joint research collaborations. A recurring theme among all of these projects is a lively partnership between our IT specialists and the domain scientists: We plan, discuss and realize research projects together on equal grounds. As an example, our IT experts regularly work with domain scientists onsite at their institutions to form one coherent team. Conversely, domain scientists have the possibility to regularly work at the LRZ to have direct access to our experts during critical phases such as the last steps of their code optimizations for our HPC systems. This close interaction as a team helps the domain scientists to make best use of our modern IT infrastructures. At the same time, our experts benefit from getting a better



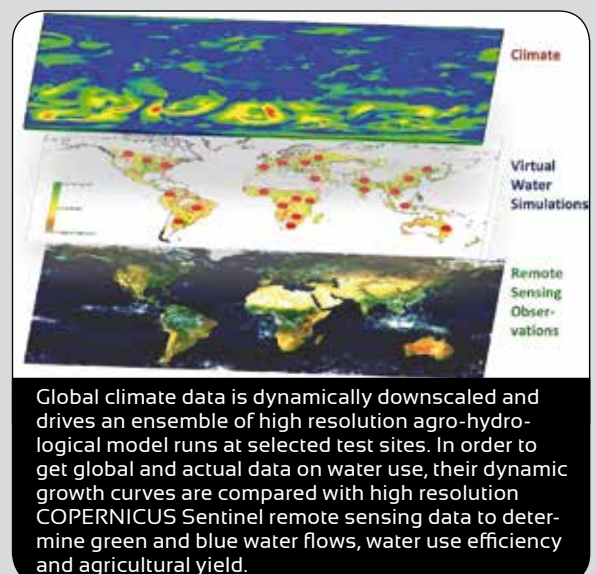
understanding of the needs of the domain scientists, their research questions, and their computational and data-related challenges.


On the technical side, we support researchers from environmental sciences with our key competences: high-performance computing and big data. Environmental systems are increasingly monitored with sensors, cameras and remote sensing techniques such as the Copernicus satellite missions of the European Space Agency. These huge datasets are rich in information about our environment, but the sources need to be made accessible through modern sharing and analysis systems. For understanding and prediction purposes, these systems also need to be modelled with a resolution that corresponds to the resolution of the data, which requires large-scale simulations on modern high-performance systems. Our tight collaboration with domain scientists has resulted in innovations such as a data centre project for the knowledge exchange in the atmospheric sciences, a technical backend for a pollen monitoring system, several projects revolving around hydrological disasters and extreme events such as floods and droughts, and a workflow engine for seismological studies, among many others.

The current and upcoming joint research projects within the environmental context confirm that our approach pays off: Personal consulting and a partnership of equals leads to fruitful collaborations and thus enables successful research.

ViWA

The project ViWA (Virtual Water Values) explores ways of monitoring global water consumption. Its primary goals are to determine the total volume of water required for food production on a global scale, and develop incentives to encourage the sustainable use of water. The new monitoring systems will focus on determining the amount of 'virtual water' contained in various agricultural products, i.e. the water consumed during their production. This will allow researchers to estimate sustainability in our current patterns of water use. In order to do so, an interdisciplinary research team lead by Prof. Wolfram Mauser, Chair of Hydrology and Remote Sensing at the LMU Munich, combines data from remote-sensing satellites with climate and weather information. The LRZ supports the domain scientists with the efficient use of high-performance computers to analyse and model their data. Additionally, LRZ will work with international stakeholders to develop an e-infrastructure that best enables the subsequent use of the collected research data.





Water use efficiency as a function of agricultural yield for three important global crops. Increases in yield are functionally connected with higher water use efficiency (adapted from Zwart and Bastiannsen 2004).

Project duration: 2017 – 2020

Website: <http://viwa.geographie-muenchen.de/>

Funding Agency: German Federal Ministry of Education and Research

Grant: € 3,6 Mio

Project Partners:

- Ludwig-Maximilians-Universität München
- Leibniz-Rechenzentrum
- Helmholtz-Zentrum für Umweltforschung UFZ
- Universität Hannover
- Institut für Weltwirtschaft (IfW)
- Climate Service Center (GERICS)
- VISTA Geoscience Remote Sensing GmbH

ClimEx

The ClimEx project seeks to investigate the occurrence of extreme meteorological events such as floods and droughts on the hydrology in Bavaria and Québec under the influence of climate change. The innovative approaches proposed by the domain scientists require the use of considerable computing power together with the expertise of professional data processing and innovative data management, and LRZ and LMU Munich contribute their expert knowledge.

The Canadian partners share their methodological expertise in performing accessible high-resolution dynamic climate projections. ClimEx further strengthens the international collaboration between Bavaria and Québec as research facilities, and universities and public water agencies intensify their cooperation approaches. For a detailed presentation of the ClimEx project see page 130.

Project duration: 2015 - 2019


Website: <http://www.climex-project.org/>

Funding Agency: Bavarian State Ministry of the Environment and Consumer Protection

Grant: €720.000

Project Partners:

- LMU München
- Bayerisches Landesamt für Umwelt
- Ouranos - Climate Scenarios and Service Group
- Centre d'Expertise hydrique du Québec (CEHQ)
- École de Technologie Supérieure (ETS)
- Montreal (PQ)




Visualization of the storm in May 1999 that led to the Pentecost Floodings in Bavaria.

Written by Jens Weismüller, Sabrina Eisenreich, and Natalie Vogel.

Leibniz Supercomputing Centre (LRZ), Germany

Contact: Jens.Weismueller@lrz.de

[illegible]



Projects

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

PPI4HPC: European Joint Procurement of Supercomputers Launched

Several European supercomputing centers have started a joint effort to buy their next generation of supercomputers within a Public Procurement of Innovative solutions (PPI). This PPI is co-funded by the European Commission and a fraction of the supercomputing resources will be made available to European scientists through PRACE.


Five partners from four different countries agreed to coordinate their procurement activities to facilitate a joint procurement: BSC from Spain, CEA and GENCI from France, CINECA

from Italy and JSC from Germany. The procedure is organized in four different lots such that next-generation supercomputing systems can be realized in each country.

A PPI is a new funding instrument, which the European Commission (EC) introduced within the H2020 framework. It aims to promote innovation through public procurement by providing financial incentives for the public sector, acting as a springboard for innovative products and services. It requires public procurers that face challenges that require solutions which



Sergi Girona (BSC) provides the welcome address during the Open Dialogue Event in Brussels on September 6, 2017.



are almost on the market but are not yet available at scale. This scenario is well-known for leading supercomputing centers.

The PPI4HPC project implements a PPI with several goals in mind. First of all, it wants to foster science and engineering applications in Europe by providing more computing resources. Furthermore, it wants to promote research and development on HPC architectures and technologies in Europe by promoting a strong relationship between the procurers and the suppliers for large-scale testing, tuning and maturation. Finally, it aims to create a greater emphasis and more impact on common topics of innovation in the area of HPC. This should lead to solutions designed according to the needs of scientists and engineers in Europe.

In preparation of the joint tender documentation, the project, which officially started in April 2017, already performed market consultations. In this context, an Open Dialogue Event has been organized in Brussels on September 6 (see Fig. 1). More in-depth technical discussions happened during meetings between the group of procurers and individual vendors. Fifteen one-to-one meetings with major HPC companies including various SMEs took place on 28-29 September in Milan and on 4-6 October in Barcelona. They were open to any interested supplier providing HPC solutions. A joint contract notice is planned to be published in April 2018. Thereafter, in each of the countries, competitive dialogues will take place, resulting in the award of one contract per lot. Systems are

planned to be installed in the time frame 2019-2020.

The PPI4HPC is taking steps toward stronger coordination of different European supercomputing centres' activities as they relate to the path toward exascale. The co-funding by the EC will allow for a significant enhancement of the planned pre-exascale HPC infrastructure from 2018 on.

Acknowledgement

The PPI4HPC project is partially funded by the European Union H2020 Program under grant agreement no. 754271.

Reference

[1] <http://www.ppi4hpc.eu>

Written by Dirk Pleiter

Jülich Supercomputing Centre (JSC)

Contact: d.pleiter@fz-juelich.de

Dorian Krause

Jülich Supercomputing Centre (JSC)

Contact: d.krause@fz-juelich.de

DEEP-EST: A Modular Supercomputer for HPC and High Performance Data Analytics

How does one cover the needs of both HPC and HPDA (high performance data analytics) applications? Which hardware and software technologies are needed? And how should these technologies be combined so that very different kinds of applications are able to efficiently exploit them? These are the questions that the recently started EU-funded project DEEP-EST addresses with the Modular Supercomputing architecture.

Scientists and engineers run large simulations on supercomputers to describe and understand problems too complex to be reproduced experimentally. The codes that they use for this purpose, the kind of data they generate and analyse,

and the algorithms they employ are very diverse. As a consequence, some applications run better (faster, more cost- and more energy-efficient) on certain supercomputers and some run better on others.

The better the hardware fits the applications (and vice-versa), the more results can be achieved in the lifetime of a supercomputer. But finding the best match between hardware technology and the application portfolio of HPC centres is getting harder. Computational science and engineering keep advancing and increasingly address ever-more complex problems. To solve these problems, research teams frequently combine multiple algorithms, or even completely





Fig. 1: DEEP-EST collaboration at the kick-off meeting in Jülich, July 13th.



different codes, that reproduce different aspects of the given topic. Furthermore, new user communities of HPC systems are emerging, bringing new requirements. This is the case for large-scale data analytics or big data applications: They require huge amounts of computing power to process the data deluge they are dealing with. Both complex HPC workflows and HPDA applications increase the variety of requirements that need to be properly addressed by a supercomputer centre when choosing its production systems. These challenges add to additional constraints related to the total cost of the machine, its power consumption, the maintenance and operational efforts, and the programmability of the system.

The modular supercomputing architecture

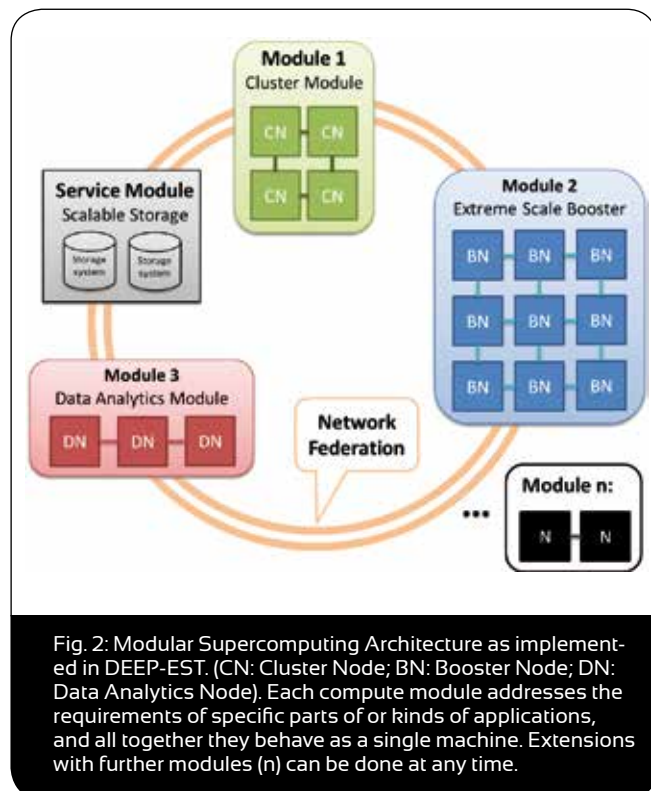
Creating a modular supercomputer that best fits the requirements of these diverse, increasingly complex, and newly emerging applications is the aim of DEEP-EST, an EU project launched on July 1, 2017 (see Fig. 1). It is the third member of the DEEP Projects family, and builds upon the results of its predecessors DEEP[1] and DEEP-ER[2], which ran from December 2011 to March 2017.

DEEP and DEEP-ER established the Cluster-Booster concept, which is the first incarnation of a more general idea to be realised in DEEP-EST: the **Modular Supercomputing Architecture**. This innovative architecture creates a unique HPC system by coupling various compute modules according to the

building-block principle. Each module is tailored to the needs of a specific group of applications, and all modules together behave as a single machine. This is guaranteed by connecting them through a high-speed network and, most importantly, operating them with a uniform system software and programming environment. In this way, one application can be distributed over several modules, running each part of its code onto the best suited hardware.

The hardware prototype

The DEEP-EST prototype (see Fig. 2) to be installed in summer 2019, will contain the following main components:



- Cluster Module: to run codes (or parts of them) requiring high single-thread performance
- Extreme Scale Booster: for the highly-scalable parts of the applications
- Data Analytics Module: supporting HPDA requirements

The three mentioned compute modules will be connected with each other through a “Network Federation” to efficiently bridge between the (potentially different) network technologies of the various modules. Attached to the “Network Federation,” two innovative memory technologies will be included:

- Network Attached Memory: providing a large-size memory pool globally accessible to all nodes
- Global Collective Engine: a processing element at the network to accelerate MPI collective operations

In addition to the three abovementioned compute modules, a service module will provide the prototype with the required scalable storage.

One important aspect to be considered in the design and construction of the DEEP-EST prototype is energy efficiency. It will influence the choice of the specific components and how they are integrated and cooled. An advanced monitoring

infrastructure will be included to precisely quantify the power consumption of the most important components of the machine, and modelling tools will be applied to predict the consumption of a large scale system built under the same principles.

The software stack

The DEEP-EST system software, and in particular its specially adapted resource manager and scheduler, enable running concurrently a mix of diverse applications, best exploiting the resources of a modular supercomputer. In a way, the scheduler and resource manager act similar to a Tetris player, arranging the differently shaped codes onto the hardware so, that no holes (i.e. empty/idle resources) are left between them (see Fig. 3). When an application

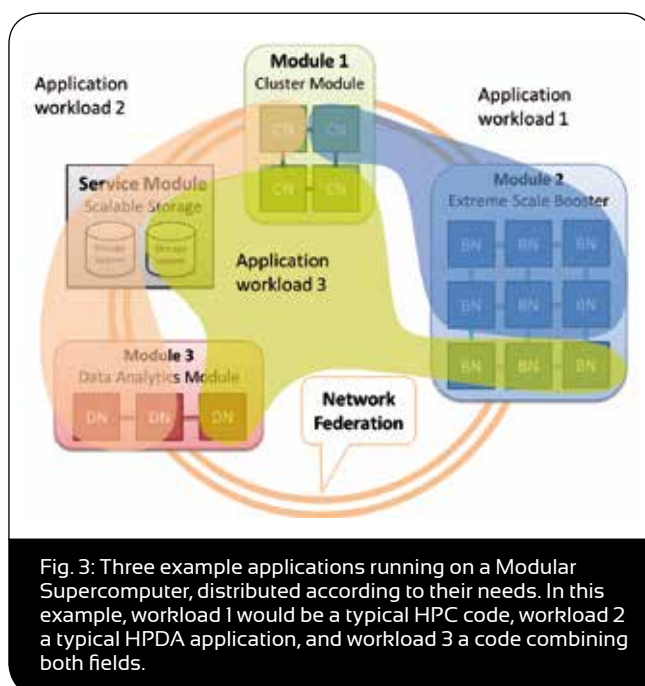



Fig. 3: Three example applications running on a Modular Supercomputer, distributed according to their needs. In this example, workload 1 would be a typical HPC code, workload 2 a typical HPDA application, and workload 3 a code combining both fields.



finishes using some nodes, these are immediately freed and assigned to others. This reservation and release of resources can be done also dynamically, what is particularly interesting when the workloads have different kinds of resource requirements along their runtime.

In DEEP-EST, the particularities and complexity of the underlying hardware are hidden from the users, which face the same kind of programming environment (based on MPI and OpenMP) that exists in most HPC systems. The key components of the programming model used in DEEP-EST have been in fact developed already DEEP. Employing ParaStation MPI and the programming model OmpSs, users mark the parts of the applications to run on each compute module and let the runtime take care of the code-offload and data communication between modules. Further resiliency capabilities were later developed in DEEP-ER. In DEEP-EST, ParaStation MPI and OmpSs will be, when needed, adapted to support the newly introduced Data Analytics Module and combined with the programming tools required by HPDA codes.

The DEEP-EST software stack is completed with compilers, the file system software (BeeGFS), I/O libraries (SIONlib), and tools for application performance analysis (Extrae/Paraver), benchmarking (JUBE) and modelling (Dimemas).

Co-design applications

The full DEEP-EST system (both its hardware and software components) is developed in co-design with a group of six scientific

applications from diverse fields. They come from neuroscience, molecular dynamics, radio astronomy, space weather, earth sciences and high-energy physics. The codes have been chosen to cover a wide spectrum of application fields with significantly different needs, and include traditional HPC codes (e.g. GROMACS), HPDA applications (e.g. HPDBSCAN), and very data intensive codes (e.g. the SKA and the CMS data analysis pipelines).

The requirements of all of these codes will shape the design of the hardware modules and their software stack. Once the prototype is installed and the software is in operation, the application codes will run on the platform, demonstrating the advantages that the Modular Supercomputing Architecture provides to real scientific codes.

Project numbers and GCS contribution

The DEEP-EST project will run for three years, from July 2017 to June 2020. It was selected under call FETHPC-01-2016 ("Co-design of HPC systems and applications") and receives a total EU funding of almost €15 million from the H2020 program. The consortium, led by JSC, includes LRZ within its 16 partners comprising computing centres, research institutions, industrial companies, and universities.

LRZ leads the energy efficiency tasks and the public relations and dissemination activities. It also chairs the project's Innovation Council (IC): a management body responsible to identify innovation opportunities outside the project.

Beyond the management and coordination of the project, JSC leads the application work package and the user-support activities. It will also contribute to benchmarking, and I/O tasks. Furthermore, in collaboration with partners Barcelona Supercomputing Centre and Intel, JSC will adapt the SLURM scheduler to the needs of a modular supercomputer. Last but not least, JSC drives the overall technical definition of the hardware and software designs in the DEEP-EST project as the leader of the Design and Development Group (DDG).

Acknowledgements

The research leading to these results has received funding from the European Community's Horizon 2020 (H2020) Funding Programme under Grant Agreement n° 754304 (Project "DEEP-EST").


References

- [1] **Suarez, E., Eicker, N., Gürich, W.:**
"Dynamical Exascale Entry Platform: the DEEP Project", inSiDE Vol. 9 No.2, Autumn 2011, http://inside.hlr.de/htm/Edition_02_11/article_12.html
- [2] **Suarez, E. and Eicker, N:**
"Going DEEP-ER to Exascale", inSiDE Vol. 9 No.2, Spring 2014, http://inside.hlr.de/htm/Edition_02_11/article_12.html
- [3] **www.deep-projects.eu**

Written by Estela Suarez

Jülich Supercomputing Centre (JSC)

Contact: e.suarez@fz-juelich.de



The Helmholtz Analytics Framework is a data science pilot project funded by the Helmholtz Association. Six Helmholtz centers will pursue a systematic development of domain-specific data analysis techniques in a co-design approach between domain scientists and information experts in order to strengthen the development of data sciences in the Helmholtz Association. Data analytics methods will be applied to challenging applications from a variety of scientific fields in order to demonstrate their potential in leading to scientific breakthroughs and new knowledge. In addition, the exchange of methods among the scientific areas will lead to their generalization.

The Helmholtz Analytics Framework (HAF) is complementary to the Helmholtz Data Federation (HDF) in that the developed libraries will be made available there first. The three-year project

starts in October 2017 and receives funding of close to €3 million.

Scientific Big Data Analytics (SBDA) has become a major instrument of modern research for tackling scientific problems of highest data and computational complexity. SBDA deals with data retrieval, assimilation, integration, processing and federation on an unprecedented scale, made possible through leading-edge high-performance computing and data management technologies.

It is crucial that systematic development of domain-specific data analytics techniques will be carried out as a co-design activity between domain and infrastructure scientists. This happens within a set of highly demanding use cases spanning six Helmholtz centers—DESY, DKFZ, DLR, FZJ, HMGU, and KIT—spanning

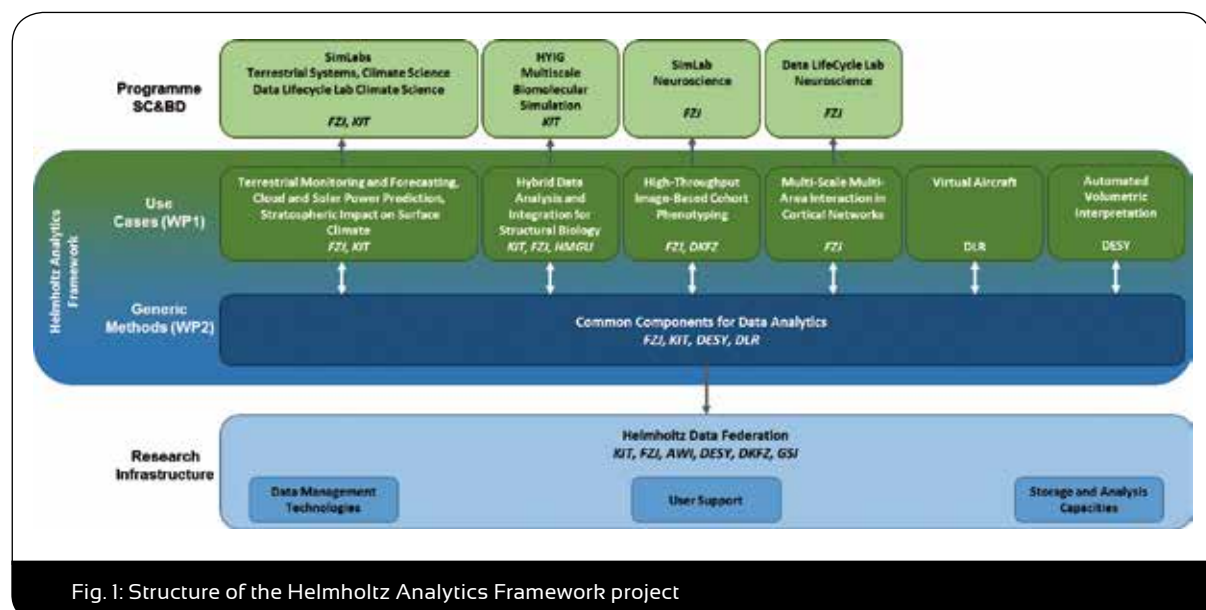


Fig. 1: Structure of the Helmholtz Analytics Framework project

five scientific domains: earth system modeling, structural biology, aeronautics and aerospace, medical imaging, and neuroscience. The exchange of techniques between the use cases will lead to generalizations and standardization to be made available to yet other fields and users.

The HAF will boost the development of the HDF, which is designed to be the hardware and support backbone for the entire Helmholtz Association and will address the dramatically increasing demands of science and engineering for transforming data into knowledge.

Thus, we start an exciting culture for future systematic developments of the HAF on top of the HDF.

Approach

The research strategy of the project is based on co-designing domain-specific data analytics techniques by domain scientists, together with data and computer scientists, evolving data analytics methods, developing the infrastructure, the HDF, with basic software systems and suitable interfaces to the application software. These activities are coherently derived from properly defined “use cases.” The use cases are chosen such that they themselves target, in a complementary manner, scientific challenges with an important societal impact and a high potential for breakthroughs in their respective domains. Through this interdisciplinary cooperation, the HDF investments will be leveraged towards a full-system solution. It is an important


goal of the project to translate specific methods developed within given use cases into generic tools and services. In a first step, they are made available to other use cases raising synergy within the project. Later, the methods will become beneficial to other fields.

Eight use cases from five scientific domains are participating in this project. The domains are Earth System Modeling, Structural Biology, Neuroscience, Aeronautics and Aerospace, and Research with Photons.

Earth system modeling

In the use case Terrestrial Monitoring and Forecasting, forecasts and projections of the terrestrial water and energy cycles constitute a scientific grand challenge due to the complexities involved and the socioeconomic relevance. Prominent examples include forecasts of weather, extreme events (floods, low-flows, droughts), water resources and long-term climate projections emerging as one of the major pillars in Earth system discovery including climate change research. Major SBDA methods encompass ensemble data assimilation technologies and genetic algorithms.

The proper forecasting of clouds in the use case Cloud and Solar Power Prediction is important for the short-term predictions of photovoltaic power, photo-chemically impaired air quality, and precipitation. The transfer of this spaceborne information in prognostic models, to result in a demonstrated beneficial effect on cloud evolution and prediction capabilities, is an



unresolved issue. Major SBDA methods to be applied are supervised learning as well as parallel and scalable classification algorithms.

Recent model developments in meteorology allow more seamless approaches to modeling weather and climate in a unified framework for the use case Stratospheric Impact on Surface Climate. An application of these advances is a hind-cast assessment of well-observed winter seasons in the northern hemisphere. Each of these (retrospective) forecasts will consist of an ensemble of realizations subsequently compared to the “real world” in order to find the most realistic ensemble member and put the real development into context with the ensemble statistics. Simulation runs will produce a large volume of 5 dimensional data, requiring fast processing for building up successive analysis layers for individual winters and for comparing all available winters in a climatological context.

Structural biology

The use case Hybrid Data Analysis and Integration for Structural Biology deals with the determination of structural ensembles of biomolecular complexes required to understand their biological functions. Single experimental techniques cannot describe the complex conformational space and temporal dynamics, and thus the integration of many complementary data with advanced computational modeling is essential. The project vision is to develop the concepts and methods needed to integrate experimental data from NMR spectroscopy, single-particle cryo-electron microscopy, and

co-evolution analysis of genetic sequences with molecular dynamics simulations. The required computational tools such as Bayesian modeling, enhanced sampling techniques, multidimensional statistical inference, feature extraction, and pattern recognition, will be developed within the algorithmic and technological framework of the HDF.

Neuroscience

Advanced medical research, like understanding of the brain or personalized medicine, are facing the challenge to understand the correlation and effect model between environmental or genetic influence and the observed resulting phenotypes (e.g. morphological structures, function, variability) in healthy or pathologic tissue. The use case High-Throughput Image-Based Cohort Phenotyping will involve neuroimaging as piloting image domain to establish time-efficient parallel processing on HPC clusters as well as highly robust but flexible processing pipelines, efficient data mining techniques, uncertainty management, sophisticated machine learning and inference approaches. Such analyses are not only of high value for systems neuroscience and medical science, but also could be generalized for other disciplines searching for causalities between image-based observations and underlying mechanisms.

The use case Multi-Scale Multi-Area Interaction in Cortical Networks employs parallelized data mining strategies paired with statistical Monte-Carlo approaches to evaluate signatures of correlated activity hidden in the

high-dimensional ensemble dynamics recorded simultaneously from visual and motor brain areas in order to link neuronal interactions to behavior. There are two challenges to be tackled by this use case. Multi-dimensional correlation analysis methods of activity due to the combinatorial complexity, strong undersampling of the system, and non-stationarities that prohibit the use of analytic statistical tests lead to increased computational demands. In addition, the heterogeneity and complex structure of the various data streams, including rich metadata, require suitable informatics tools and protocols for the acquisition of metadata and provenance tracking of the analysis workflows.

Aeronautics and aerospace

The use case Virtual Aircraft employs reduced order models that extract relevant information from a limited set of large-scale high-fidelity simulations through elaborate result analysis methods to provide an attractive approach to reduce numerical complexity and computational cost while providing accurate answers. Data classification methods are of interest to gain more physical insight, e.g., to identify (aerodynamic) nonlinearities and to track how they evolve over the design space and flight envelope. The Virtual Aircraft use case will lead to SBDA techniques from other fields of research being evaluated for extracting a comprehensive digital description of an aircraft from a parallel workflow based on high-fidelity numerical simulations. The Virtual Aircraft use case will also contribute a wide range of methods for data fusion, surrogate and reduced-order modeling


to the generic methods that can be applied to the use cases of other partners. The software and SBDA methods to set up the Virtual Aircraft can be developed in such a generic fashion that it will be possible to adapt them to other fields of research that deal with product virtualization.

Research with photons

SBDA techniques can be used for Automated Volumetric Interpretation of time-resolved imaging of materials and biological specimen to provide deep insight into dynamics of bacterial cells, composite materials, or living organisms, among others. Experiments are coming from X-ray imaging at synchrotrons or free-electron lasers. The quality of automated segmentation and interpretation algorithms will strongly increase with the amount of available data combined with SBDA techniques to harvest and mine prior information from similar experiments across facilities and disciplines. To maximize the sample size, we aim to exploit the vast amount of imaging data available in the Helmholtz Data Centers as well as the PaNdata collaboration, which includes almost all European Photon and Neutron sources, and also collaborations with various other light sources, particularly in the USA. The interpretation of 3D-data by volumetric segmentation and interpretation can greatly benefit from SBDA by harvesting and mining prior information from similar experiments across facilities and disciplines.

Work plan

The project has a duration of 36 months. During the initial phase, we will determine common



methods and respective tools among the use cases. An initial set of common method areas, including stochastics, image analysis, supervised and unsupervised learning, has already been identified. During the second phase, the methods for mutual use in the participating use cases will be generalized and the tools will be adapted and rolled out on the HDF. It is expected that this will lead to cross-fertilization in the use of common methods. In the final phase, the common methods and tools will be made available for a wider audience. Care will be taken to make tools available not only among participating scientific domains, but also generically. This will include appropriate documentation of the methods as well as the tools that implement them and their installation and usage on the HDF.

Acknowledgement

The project is funded by the Helmholtz Association Initiative and Networking Fund under project number ZT-I-0003.

Written by Björn Hagemeier

Jülich Supercomputing Centre (JSC)

Contact: b.hagemeier@fz-juelich.de

Daniel Mallmann

Jülich Supercomputing Centre (JSC)

Contact: d.mallmann@fz-juelich.de

Achim Streit

Karlsruhe Institute of Technology, Steinbuch Centre for Computing

Contact: achim.streit@kit.edu

Rhinodiagnost: Morphological and Functional Precision Diagnostics of Nasal Cavities

In this project, globally recognized research centers and market-leading medical technology companies are working on coordinated morphological and functional diagnostics for ear, nose and throat (ENT) physicians. Services are organized as a fast-working network in which important new decision aids, such as 3D models and flow simulation results, are made available to ENT specialists.

The nose is one of the most important organs of the human body and its functions are essential for the comfort of the individual patient. It is responsible for olfaction, supports degustation, filters the air of harmful particles, and tempers and moisturizes the inhaled air to create optimal

conditions for the lung. Diseases and discomfort in the nasal cavity as they occur, for example, in chronic rhinosinusitis, nasal septal deviation, after surgery, or in polyp diseases, often lead to a reduction in one or more of these functionalities. Such a reduction frequently results in a limitation of the respiratory capacity, the formation of inflammatory foci in the nasal cavity, and lung diseases. A meaningful rhinological diagnosis is therefore key in evaluating the effectiveness of patient-specific nasal functionalities, taking into account the respective pathology.

The diagnostic quality is currently primarily based on the quality of the training of the practicing physician and his or her experience




Fig. 1: Inspiratory flow in the human nasal cavity. The streamlines are colored by the velocity magnitude. The inset on the right shows a magnification of the boundary-refined computational mesh. © Institute of Aerodynamics, RWTH Aachen University.

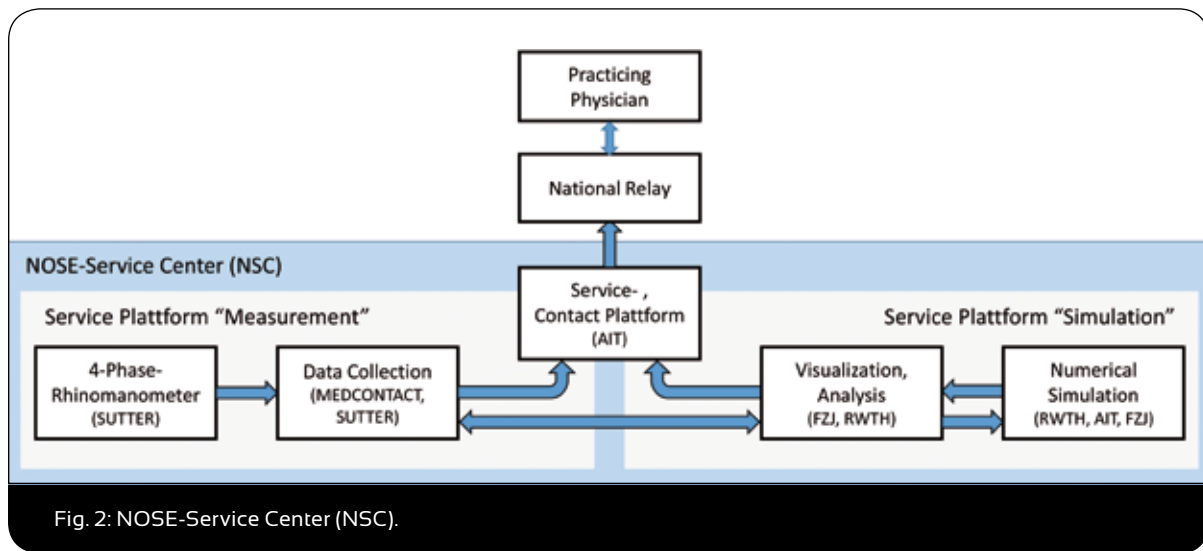



Fig. 2: NOSE-Service Center (NSC).

in the treatment of specific clinical pictures. The according functional diagnostics employ methods of medical imaging, such as computer tomography (CT) or magnetic resonance tomography (MRT), to enable a well-founded diagnosis. Unfortunately, such analyses do not include any information on the respiratory comfort of a patient defined by the fluid mechanical properties of respiration.

Project goals

Current developments in the field of computational fluid dynamics (CFD) and high-performance computing (HPC) allow for patient-specific prediction of the flow in a human nasal cavity by means of numerical simulations [1, 2] (see Fig. 1), enabling identification of anatomical locations of pathologies. In addition, advanced rhinomanometry methods [3, 4] allow medical professionals to determine respiratory resistance in order to provide extended information on the patient's respiratory capacity. Hence,

results from CFD and rhinomanometry can be used to a-priori determine optimal surgery strategies for an individual patient in order to increase surgery success rates and to adapt treatment therapies.

Unfortunately, such methods have not made their way into everyday clinical practice due to their complexity and costs. In order to improve this situation, the implementation of a NOSE Service Center (NSC) is to be prepared within this project, offering extended possibilities of functional diagnostics, and providing a network of service points. Fig. 2 shows schematically the structure and interaction chain within the NSC.

Project consortium and tasks

To reach the project goals, two German medical device companies, namely Sutter Medizintechnik GmbH (SUTTER) and Med Contact GmbH (MEDCONTACT), jointly proceed with the Austrian partner Angewandte


Informationstechnik Forschungsgesellschaft mbH (AIT), and the two research facilities Institute of Aerodynamics (AIA), RWTH Aachen University, and Jülich Supercomputing Centre (JSC), Forschungszentrum Jülich, to implement the NSC.

In more detail, the partners will perform the following tasks to reach the overall project goals:

- **SUTTER** has developed the 4-phase rhinomanometer (4PR, 4RHINO), which was declared a standard in functional diagnostics in November 2016. SUTTER carries out optical examinations in an in-vivo nasal model and performs in-vitro analyses using 4PR to validate the numerical methods. In addition, the influence and the physical properties of the nasal valve, which can produce an airway collapse during accelerated flow, will be investigated by means of elastography methods.
- **MEDCONTACT** expands the 4PR with wireless data transmission functions for automated data collection. NSC compatibility will be ensured in cooperation with SUTTER and the 4PR will be clinically tested and introduced to the market.
- **AIT** sets up a service and contact platform, which is to serve as an interface between the practicing physician and the service platforms behind it. Additionally, AIT evaluates an established CFD method in terms of cost, efficiency and accuracy.
- **AIA** evaluates high-fidelity CFD methods in terms of cost, efficiency and accuracy. Furthermore, in-situ computational steering will be implemented to allow for online modification of the geometry at simulation run time and for an up-to-date fluid mechanical interpretation of the geometrical changes. Therefore, automatic analysis tools for expert analysts as well as tools retrieving key information relevant for direct clinical use will be implemented.
- **JSC** develops software components making the analysis of the simulation data accessible to the physician interactively and purposefully on modern HPC systems. Beyond that, the possibility of using virtual operations with direct updating and analysis of the flow parameters are demonstrated in close cooperation with AIA.

Project funding

Rhinodiagnost is funded as a ZIM (Zentrales Innovationsprogramm Mittelstand) project by the Federal Ministry for Economic Affairs and Energy (BMWi) in Germany. The Austrian partner is funded by COIN (Cooperation and Innovation), Federal Ministry of Science, Research and Economy (BMWFW). The project runs under the auspices of IraSME (International research activities by SMEs).



Project coordination and contact:

AIT – Angewandte Informationstechnik
Forschungsgesellschaft mbH
Klosterwiesgasse 32/I
8010 Graz Austria
Phone: +43 - 316 - 8353590
Email: info@rhinodiagnost.eu
Web: www.rhinodiagnost.eu

References

- [1] **A. Lintermann, M. Meinke, W. Schröder:**
Investigations of Nasal Cavity Flows based on a Lattice-Boltzmann Method, in: M. Resch, X. Wang, W. Bez, E. Focht, H. Kobayashi, S. Roller (Eds.), High Perform. Comput. Vector Syst. 2011, Springer Berlin Heidelberg, Berlin, Heidelberg, 2012, pp. 143–158. doi:10.1007/978-3-642-22244-3.
- [2] **A. Lintermann, M. Meinke, W. Schröder:**
Fluid mechanics based classification of the respiratory efficiency of several nasal cavities, Comput. Biol. Med. 43 (11) (2013) 1833–1852. doi:10.1016/j.compbiomed.2013.09.003.
- [3] **K. Vogt, A. A. Jalowayski:**
4 - Phase-Rhinomanometry, Basics and Practice 2010, Rhinology Supplement 21.
- [4] **K. Vogt, K.-D. Wernecke, H. Behrbohm, W. Gubisch, M. Argale:**
Four-phase rhinomanometry: a multicentric retrospective analysis of 36,563 clinical measurements, Eur. Arch. Oto-Rhino-Laryngology 273 (5) (2016) 1185–1198. doi:10.1007/s00405-015-3723-5.

Written by Andreas Lintermann

Institute of Aerodynamics and Chair of Fluid Mechanics,
RWTH Aachen University, Germany

Contact: A.Lintermann@aia.rwth-aachen.de

Jens Henrik Göbbert

Jülich Supercomputing Centre (JSC),
Forschungszentrum Jülich, Germany

Contact: j.goebbert@fz-juelich.de

Klaus Vogt

Sutter Medizintechnik GmbH, Germany

Contact: rhinovogt@t-online.de

Walter Koch

AIT - Angewandte Informationstechnik Forschungs-
gesellschaft mbH, Austria

Contact: kochw@ait.co.at

Alexander Hetzel

Med Contact GmbH, Germany

Contact: hetzel@reger-med.de

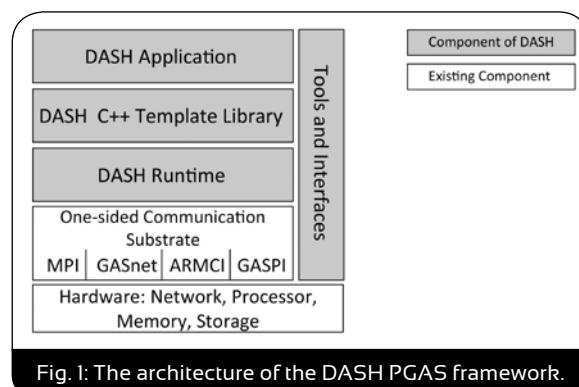
DASH – Distributed Data Structures in a Global Address Space

In 1980, an Intel 8088/87 CPU could perform a floating point operation in around 10 microseconds. In the same amount of time, a whopping 20 bytes of data could be fetched from memory [1]. The systems of this era were largely limited by their ability to perform arithmetic operations—how things have changed since then! Today, the performance most applications can extract from a system is not limited by the raw CPU computing power in terms of floating point operations per second (FLOPs), but by the time it takes to feed the data into the processors' compute units.

As the CPU manufacturing process evolves into smaller feature sizes, data transfers will account for an ever increasing fraction of the time and energy budget of a CPU and this situation is expected to get significantly worse in the future [2]. At the same time, HPC systems are getting bigger in terms of core and node counts and data locality must be taken into account both horizontally (between shared memory domains) and vertically (within a shared memory domain) [3].


Data thus needs to have a place at the center of attention for HPC application developers. This is where the DASH project aims to make contributions by providing distributed data structures and parallel algorithms similar in spirit to what is provided by the STL (standard template library). The goal of DASH is to make working with data easier and more productive for HPC developers.

The overall structure of the DASH project is shown in Fig. 1. DASH makes use of existing



one-sided communication substrates and is based upon a runtime system called DART. This runtime system provides central abstractions such as global memory allocation and addressing and offers one-sided communication operations (puts and gets). The execution model adopted by DART follows SPMD (single program, multiple data) semantics, where the individual participants are called units. Units can be hierarchically grouped into so-called teams to address the growing size of systems and their increasingly multi-leveled organization. To enable seamless integration with the large number of existing MPI applications, DART is based on the MPI-3 RMA (remote memory access) interface [5].

The DASH project started in 2013 with the first funding phase of SPPEXA and continued into the second phase with a consortium of four German partners. The project is led by LMU Munich, where also the bulk of the C++ library development is situated. The partners in HLRS Stuttgart, IHR Stuttgart, and TU Dresden contribute expertise in runtime system development, application engagement and tools integration.



Data structures with global-view and local-view semantics

DASH offers data structures that mimic the interface and behavior of STL containers. The most basic data structure available in DASH is a fixed-size, one-dimensional array. `dash::Array<int> arr(1000)` declares such an array of one thousand integers, where the data to store the individual elements is contributed by all units that execute the program. The array, `arr`, is said to have global-view access semantics, since each process has a unified (global) view of the container – `arr.size()` returns the same (global) size and `arr[42]` refers to the same element on each unit. Global-view access is very convenient when working with dynamically changing and irregular data structures, but it comes with an overhead in terms of access cost, since each access instance might need a locality check and a network transfer to retrieve or store remote data.

To avoid such access overheads entirely when working with known local data, DASH also supports a local-view access mode, using the `“.local”` accessor object. For example `arr.local.size()` returns the number of data elements available locally, `arr.local[0]` returns the first element stored locally, and so on. Using this form of local-view access has large performance benefits and allows for a straightforward realization of the owner computers parallel computation model.

Besides the basic one-dimensional fixed-size array, DASH also supports multidimensional

arrays (`dash::NArray`) where support for slicing in arbitrary dimensions is included. Dynamic (growing/shrinking) data structures are under development, including dynamic lists and hash-maps.

Data distribution patterns

In large-scale HPC systems, parallelism implies distribution, meaning several compute nodes are interconnected by some form of high-speed interconnect network. The way in which the data that is being operated on is distributed among these nodes can have an important influence on program performance. Since DASH offers distributed data structures that can span multiple interconnected compute nodes, it also has to provide flexible ways in which to specify data distributions.

In DASH this is achieved by specifying a so-called pattern. The pattern determines how data elements are distributed among a set of units and how the iteration sequence is determined. In one dimension, DASH offers the usual choice of a blocked, cyclic, and block-cyclic distribution. In multiple dimensions, these specifications can be combined in each dimension. Additionally, a tiled distribution is supported where small contiguous blocks of the iteration space are specified.

Figure 2 shows several examples for DASH data distribution patterns in two dimensions. The colors correspond to the processes (units), the iteration sequence is additionally visualized for unit 0 using connected dots.




Fig. 2: Examples of DASH Data Distribution Patterns.

DASH supports both row-major as well as column-major storage order.

Productivity through algorithms


Distributed data structures are convenient, but ultimately developers are interested in performing computations on the data in an efficient and productive way. In addition to element-wise and bulk data access feeding into existing code, DASH also offers interoperability with sequential STL algorithms and provides a set of parallel algorithms modeled after their STL counterpart.

As an example, standard algorithms such as `std::sort` or `std::fill` can be used in conjunction with the global-view and the local-view mode for a DASH container.

`std::sort(arr.local.begin(), arr.local.end())` sorts the local portion of the distributed array and

`std::min_element(arr.begin(), arr.end())` finds the smallest element in the whole array. In the latter example, the minimum is found without exploiting the available parallelism, since the STL algorithm cannot be aware of the data distribution and available compute units. When the algorithm provided by DASH `dash::min_element(arr.begin(), arr.end())` is used instead, all units collaboratively find the global minimum by first finding their local minima and then collectively determining the global minimum.

The usage of DASH algorithms can enhance programmer productivity significantly. Instead of the classic imperative programming style commonly used in C/C++ or Fortran MPI codes, the usage of algorithms provides a more declarative style that is both more compact and invariant under changes of the underlying data distribution.



```
1 #include <iostream>
2 #include <libdash.h>
3
4 int main(int argc, char *argv[]) {
5     dash::init(&argc, &argv);
6
7     // 2D integer matrix with 10 rows, 8 cols
8     // default distribution is blocked by rows
9     dash::NArray<int, 2> mat(10, 8);
10
11     for (int i=0; i<mat.local.extent(0); ++i) {
12         for (int j=0; j<mat.local.extent(1); ++j) {
13             mat.local(i, j) = 10*dash::myid()+i+j;
14         }
15     }
16
17     dash::barrier();
18
19     auto max = dash::max_element(mat.begin(), mat.end());
20
21     if (dash::myid() == 0) {
22         print2d(mat);
23         cout << "Max is " << (int)(*max) << endl;
24     }
25
26     dash::finalize();
27 }
```

Fig. 3: A basic example DASH application.

Fig. 3 shows a basic complete DASH program using a 2D array (matrix) data structure. The data type (int) and the dimension (2) are compile-time template parameters, the extents in each dimension are set at runtime. In the example, a 10 x 8 matrix is allocated and distributed over all units (since no team is specified explicitly). No specific data distribution pattern is requested, so the default distribution by block of rows over all units is used. When run with four units, each unit gets $\text{ceil}(10/4)$ matrix rows, except for the last unit, which receives only one row.

Lines 10 to 15 in Fig. 3 show data access using the local matrix view by using the proxy object `mat.local`. All accesses are performed using local indices (i.e., `mat.local(1,2)` refers to the element stored locally at position (1,2)) and no communication operation is performed. The barrier in line 17 ensures that all units have initialized their local part of the data structure before the `max_element()` algorithm is used to find the maximum value of the whole matrix. This is done by specifying the global range that encompasses all matrix element (`mat.begin()` to `mat.end()`). In the library implementation of `max_element()`, each unit determines the locally stored part of the global range and performs the search for the maximum there. Afterwards a reduction operation is performed to find the global maximum. The return value of `max_element()` is a global reference for the location of the global maximum. In lines 21 to 24, unit 0 first prints the whole matrix (the code for `print2d()` is not shown) and then outputs the maximum by dereferencing the global reference `max`.

Compile and Run:

```
$> mpicc -L ... -ldash -o
      example example.cc
$> mpirun -n 4 ./example
```

Output:


```
0 1 2 3 4 5 6 7
1 2 3 4 5 6 7 8
2 3 4 5 6 7 8 9
10 11 12 13 14 15 16 17
11 12 13 14 15 16 17 18
12 13 14 15 16 17 18 19
20 21 22 23 24 25 26 27
21 22 23 24 25 26 27 28
22 23 24 25 26 27 28 29
30 31 32 33 34 35 36 37
Max is 37
```

Fig. 4: Commands to compile and run the DASH application and the output produced by the program.

Fig. 4 shows the output produced by this application and how to compile and run the program. Since DASH is implemented on top of MPI, the usual platform-specific mechanisms for compiling and running MPI programs are used. The output shown is from a run with four units (MPI processes), hence the first set of three rows are initialized to 0...9, the second set of three rows to 10...19, and so on.

Memory spaces and locality information

To address the increasing complexity of super-computer systems in terms of their memory



organization and hardware topology, work is currently underway in DASH to offer constructs for productive programming constructs dealing with novel hardware features such as non-volatile and high-bandwidth memory. These additional storage options will be represented as dedicated memory spaces and the automatic management and promotion of parts of a data structure to these separate memory space will be available in DASH. Additionally, a locality information system is under development which supports an application-centric query and exploitation of the available hardware topology on specific machines.

Conclusion

DASH is a C++ template library that offers distributed data structures with flexible data partitioning schemes and a set of parallel algorithms. Stand-alone applications can be written using these features but DASH also allows for integration into existing MPI codes. In this scenario, individual data structures can be ported to DASH, incrementally moving from existing two-sided communication operations to the one-sided operations available in DASH.

DASH is available as open source software under a BSD clause and is maintained on GitHub (<https://github.com/dash-project/dash>). Additional information on the project, including tutorial material, can be found on the projects webpage at <http://www.dash-project.org/>. More information can also be found in a recent overview paper [4].

References

- [1] **McCalpin, J. D.:**
A survey of memory bandwidth and machine balance in current high performance computers, IEEE TCCA Newsletter 19, 25, 1995.
- [2] **Ang, J. A., et al.:**
Abstract machine models and proxy architectures for exascale computing. Hardware-Software Co-Design for High Performance Computing (Co-HPC), IEEE, 2014.
- [3] **Unat, D., et al.:**
Trends in data locality abstractions for HPC systems, IEEE Transactions on Parallel and Distributed Systems, 2017.
- [4] **Führlinger, K., Fuchs T., and Kowalewski R.: DASH:**
a C++ PGAS library for distributed data structures and parallel algorithms, High Performance Computing and Communications; IEEE 14th International Conference on Smart City; IEEE 2nd International Conference on Data Science and Systems (HPCC/SmartCity/DSS); 18th International Conference on, IEEE, 2016.
- [5] **Zhou, H., et al.:**
DART-MPI: an MPI-based implementation of a PGAS runtime system, Proceedings of the 8th International Conference on Partitioned Global Address Space Programming Models, ACM, 2014.

Written by The DASH Project Team:

**Dr. Karl Fűrlinger, Tobias Fuchs, MSc.,
Roger Kowalewski, MSc.**

Ludwig-Maximilians-Universität München

Contact: fuerling@nm.ifi.lmu.de

Dr. Andreas Knüpfer, Denis Hünich, MSc.

Technische Universität Dresden

Contact: andreas.knuepfer@tu-dresden.de

Dr. José Gracia, Joseph Schuchart, MSc.

Dr. Colin Glass, Daniel Bonilla, MSc.

Höchstleistungsrechenzentrum Stuttgart

Contact: gracia@hlrs.de

HLRS Researchers Do the Math for Future Energy Savings

In terms of data and energy reduction, finding the right balance is key

Computational fluid dynamics (CFD) simulations are both energy and time intensive. However, for future industrial applications, CFD simulations are crucial for reducing costs, time-to-market, and increasing product innovation. Within the scope of the EU research project 'ExaFLOW' and DFG (Deutsche Forschungsgemeinschaft)—funded project 'Exasolvers 2', HLRS researchers Björn Dick and Dr. Jing Zhang are rising to the challenge of developing energy efficient algorithms and easing the I/O bottleneck.

Modern high performance computers (HPC) have seen a steady growth in computation speed for the last 10 years and now head towards exascale computing performance—a thousand-fold speedup over current petascale machines. However, data transfer rates are currently not able to keep up with this rapid hardware development. More specifically, despite the theoretical capability of producing and processing high amounts of data quickly, the overall performance is often-times restricted by how fast a system can transfer and store the computed data, not to mention the high energy demand for handling big datasets. Nevertheless, high-performance CFD simulations are increasingly gaining importance for industrial implementation, as they allow for lower costs during the research and development process and offer amplified optimization possibilities. For a business, this means better products reaching the market faster. To this end, HLRS—together with seven more partners from Europe—participates in the 'ExaFLOW'



project with the aim of addressing key algorithmic challenges in CFD to enable simulation at exascale, guided by a number of use cases of industrial relevance, and to provide open-source pilot implementations. The role of HLRS in this project is to work out solutions for two crucial exascale bottlenecks: data volume and energy reduction.

HLRS contributes with energy efficient algorithms and I/O optimization

Principal investigator Jing Zhang is convinced of the importance of 'ExaFLOW'. "CFD simulations are among the most important application areas in HPC and crucial for future progress," she said. "But in order to reach our development goals, the data volume is the main influencing variable



we need to play with." Within the ExaFLOW project, her main task is to apply the data analysis strategy of singular value decomposition (SVD) on the dataset and research the effect on data volume. SVD is used as a method to decompose a single matrix into several matrices with varying characteristics and subsequently reducing the dimensionality of the data. However, this



data-reduction approach would lead to a slight loss of accuracy, so the challenge is to find the right balance between data reduction and loss of accuracy in the I/O process. When it comes to the energy efficiency of applications running in the 'ExaFLOW' project, studies show that moderately changing the CPU clock frequency of nodes—the frequency at which nodes are actively calculating—can result in significant savings in energy-to-solution without excessively increased time to solution. In contrast, if CPU clock frequency is highly reduced, time to solution is significantly extended. So researchers are not just trying to improve efficiency, but

rather find the right balance between energy to solution and time to solution. Synergetic effects within the HLRS research staff contributes heavily to research on future exascale challenges. With Björn Dick working on algorithms and implementations with low power consumption within the Exasolvers 2 project, knowledge and skills concerning exascale environments have been exchanged.

Written by Lena Bühler, Michael Gienger

Höchstleistungsrechenzentrum Stuttgart

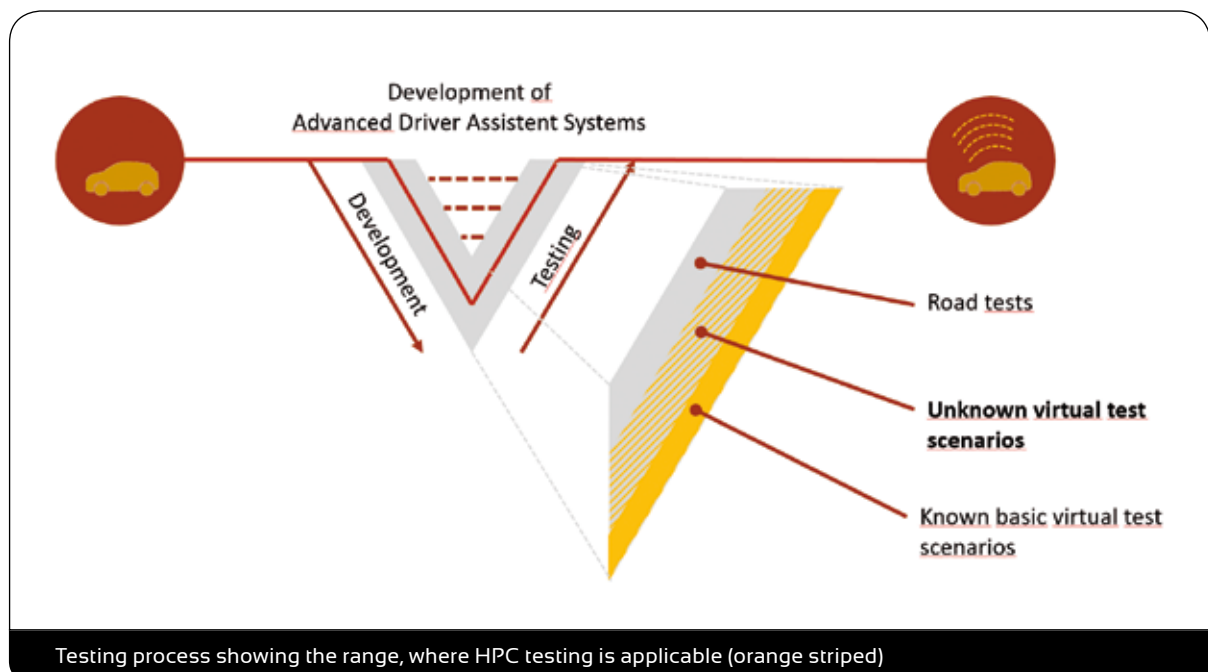
Contact: gienger@hlrs.de

Fortissimo 2 Uses HPC and Machine Learning to Improve Self-Driving Vehicle Safety

An experiment conducted by Stuttgart-based Spicetech GmbH, with the collaboration of the High-Performance Computing Center Stuttgart (HLRS) and the Slovenian research group XLAB, uses a combination of high-performance computing (HPC) and machine learning in order to improve the safety and decrease the cost of automated driving. The experiment is part of the latest call for the EU-project Fortissimo 2 new application experiments.


Despite being one of the most anticipated innovations in the automotive industry, self-driving vehicles are associated with risks. As a result,

of critical events contributing to the security system. On top of this testing being the most cost-intensive and time-consuming part of automotive development, the system will still react exclusively according to the previously encountered test scenarios. "Keeping lane on a perfect street is a simple task for autonomous cars", said Dr. Alexander Thieß, experiment leader and founder of the engineering and consulting enterprise Spicetech GmbH. "This leaves a vast amount of rarely appearing, but yet possible safety-relevant scenarios uncovered. Reacting properly to the unexpected requires a combination of machine learning and HPC."



these vehicles require maximum reliability, leading to a vast testing process. Traditional testing approaches maximize the amount of driven kilometers in order to obtain a high frequency

The process starts with rendering complex, high-resolution scenarios, which are detected by different sensor systems, such as camera, laser, or radar. The output information defines



whether the system reacted correctly to the scenario within a single test run. The challenge is to quickly identify as many source of failure as possible in the vast variety of surrounding conditions, such as weather, the surface of the road, and objects on the road. Therefore, the researchers use HLRS' Hazel Hen supercomputer which could run the test set of one trillion object and street scenarios in less than a day. If e.g. a street sign could not be detected because of hard rain, software engineers can promptly work on refining this specific error source.

To improve speed and detail-level of test space exploration, Spicetech incorporates already existing and successfully applied machine-learning algorithms. These algorithms help define a measure of potential test failure then they share the obtained information multiple times per minute among the used supercomputing nodes. By that, strength and spectrum of

the set of scenarios can be improved by orders of magnitude.

Spicetech GmbH conducts the experiment within the scope of EU-funded research project Fortissimo 2, with the technical collaboration from the HLRS and the research group XLAB, and with the associated partners Porsche, Valeo und Pipistrel. The aim is to significantly improve both financial- and security-related aspects of self-driving vehicles by developing a virtual HPC testing framework and making it available as a web application for car and Advanced Driver Assistance Systems (ADAS) manufacturers.

Written by Lena Bühler

Höchstleistungsrechenzentrum Stuttgart, (HLRS)

Contact: Michael Gienger, gienger@hlrs.de

Fostering Pan-European Collaborative Research: 3rd Edition of the “HPC-Europa” Project Launched

In early May 2017 researchers from all over Europe came together to kick off the 3rd edition of the EC-funded activity “Transnational Access Programme for a Pan-European Network of HPC Research Infrastructures and Laboratories for scientific computing” project, abbreviated as HPC-Europa 3 (<http://www.hpc-europa.eu>). The High-Performance Computing Center in Stuttgart (HLRS) appreciates this continuation of a well-established framework, especially as this project fits into the center’s strategy to support both, junior researchers and Small and Medium Enterprises (SMEs).

With the first initiative dating back to 2002, the HPC-Europa project has proved a successful European Union (EU) instrument for supporting Pan-European collaboration on HPC-related topics. Similar to its predecessors, HPC-Europa3 offers grants to young researchers for an international research exchange to one of the ten participating HPC centres. Young scientists from the EU and beyond who deal with HPC applications, tools, benchmarks or similar topics are entitled to apply for travel grants. HLRS, as one of the participating HPC centres, is planning to provide up to 150 young scientists with access to the systems to the tune of 4.2 Million compute hours. Altogether, the participating centres are planning to host 1,220 junior researchers and provide over 100 Million compute hours.


HPC-Europa supports the international cooperation of researchers in HPC

The goal of HPC-Europa3 is to foster Pan-European and international research collaboration.

Researchers are offered grants to visit an EU-based partner organization which in turn provides compute time at one of the corresponding HPC centres. For example, any researcher who already has or wishes to establish a partnership with one of the research organizations in Germany can apply for a grant and get several thousand compute hours on the HLRS systems. German researchers can also benefit from HPC-Europa3 by visiting one of the other 9 HPC centres in other countries (<http://www.hpc-europa.eu/partners>).

HPC resources and services for SMEs and junior researchers

HPC-Europa 3 aims to foster cooperation in the research community, but also facilitate small and medium enterprises’ (SMEs’) access to HPC resources. For this purpose, HLRS will be the first of three HPC centres conducting a so-called “HPC awareness workshop” in October 2018, followed by the supercomputing centre at the University of Edinburgh (EPCC) and Italian computing centre CINECA in a 6-month intervals. Acting as a mentor, HLRS staff will be inviting representatives from SMEs to Stuttgart to share information on how HPC can be beneficial for their competitive positions. As a centre with strong technical and engineering expertise, HLRS aims to build up long-term technology transfer actions with participating SMEs thus enabling them to effectively use HPC technologies and services. HLRS director Prof. Michael M. Resch emphasizes the importance of supporting both junior researchers and SMEs. “The scope of HPC has amplified many times over in



the last decade, not only in terms of research, but also in terms of commercial exploitation,” Resch explains. “We now need to support junior researchers in order to forward academic achievements in Europe. As for SMEs, they represent the economic backbone, especially in Germany, and must be included to maintain their competitive positions.”.

Recognizing the importance of SMEs for the EU research and innovation horizon, HPC-Europa3—for the first time in its history—will allow young researchers from universities to visit partnering SMEs, which are acting as hosts for them. The inverse—young SME staff visiting academic organizations to foster innovative work and proof-of-concepts for SMEs—is also possible and is supported by HPC-Europa3.

HPC-Europa 3 funded by European Commission with over 9 million euro

In order to increase synergy and collaboration with these various interest groups, access to the participating European HPC centres and travel grants will be provided by the participating centres free of charge and via a single application, easing the administrative burden. The Transnational Access Programme is funded with 9.2 million Euro by the EU within the scope of the H2020 Programme. The program is planned to run 48 Months starting in May 2017.

HPC-Europa3 calls

The first HPC-Europa3 call for applications, which was successfully closed in September 2017, had 69 applications submitted (14 of

them indicated HLRS as a centre and a German research organization as a host). All applications are evaluated by the technical and scientific boards, created with the participation of the leading experts in the HPC field in order to ensure a high quality and a competitiveness of the projects to be conducted under the HPC-Europa3 umbrella. The typical application includes brief information about the applicant, an abstract of the proposed research project, technical details of the HPC requirements as well as a short information on the host(s) for visiting.

The calls happen on a regular (usually quarterly) basis. The deadlines of the upcoming calls are:

- Call #2 - 16 November 2017
- Call #3 - 28 February 2018
- Call #4 - 17 May 2018
- Call #5 - 06 September 2018
- Call #6 - 15 November 2018

HPC-Europa3 support

HPC-Europa3 strives to immediately respond to any emerging requirements or support requests of the applicants and hosts (both from SMEs and academia). For any requests, a variety of contacting possibilities are available, ranging from e-mail to twitter.

For any questions related to the German site, please contact the HPC-Europa3 hosting team at: hpc-europa@listserv.uni-stuttgart.de

Written by Lena Bühler, Alexey Cheptsov

Höchstleistungsrechenzentrum Stuttgart (HLRS)

Contact: Alexey Cheptsov, cheptsov@hlrs.de

HLRS Scholarships Open Up “Simulated Worlds” to Students



On Friday July 7, six scholarship holders accepted certificates for successfully carrying out research projects in the field of simulation. The recipients were funded with €1,000 within the context of “Simulated Worlds,” a project led by the High-Performance Computing Center Stuttgart (HLRS) that aims to raise pupils’ and teachers’ awareness of simulation and the technical skills it involves.

Safe driving, accurate weather forecasting, and resource-friendly production—few know that without simulation, certain accomplishments of the modern age would not exist, including some ubiquitous components of everyday life. The research project Simulated Worlds, funded by the Baden-Württemberg Ministry of Science, Research, and Art, aims to sensitize students to the importance of simulations and their applications, and to enhance their interest in coding by bringing the topic into the classroom.

Simulation covers broad range of topics

The first call for HLRS-funded scholarships in the school year 2016/2017 constitutes a milestone for Simulated Worlds. Beginning in 2011, elements such as study trips, training courses, and course material formed the foundation of the project. These are now being enhanced by actively involving a selected group of 10th- and 11th-grade junior scientists in scientific work. Their efforts to familiarize themselves with subjects such as medical technology, urban planning, and philosophy has been especially pleasing for project leader Jörg Hilpert: “When students get to work on such a wide range of


topics, they get to see far more ways in which simulations are applied in both scientific and societal ways,” Hilpert says. “They also get to see how the technical skills that are required in simulation can be used in fields one wouldn’t necessarily expect.”

Promising trainees enrich science

Six pupils received awards on July 7 for their work on three projects. Focusing on blood flow through the human heart, Jana-Pascal Bode, Cara Buchholz, and Jakob Steimle analyzed the flow conditions and volume flows of four predetermined sections of the multiple parts of the heart’s main artery—the aortic arch and the descending aorta—as well as cranial branches. The team used open source-software called ParaView to visualize and analyze the underlying MRI dataset.

Alexander Kharitonov and Marius Manz conducted traffic simulations of a traffic hub in Herrenberg, a small city located near Stuttgart. They collected three-dimensional data using laser scanning, generated a road network, and integrated the results with another set of three-dimensional measurement data provided by the state of Baden-Württemberg. They visualized these components using the Virtual Reality Modeling Language (VRML) and presented their findings to attendees of the awards ceremony in the CAVE, a virtual reality environment at HLRS.

Kira Fischer approached the topic of simulation on a philosophical level, raising the question



HLRS Director Michael Resch (right) with participants in the Simulated Worlds program. (Photo: HLRS)

of the veracity of results gained by computer simulation. Developing her argument required exploring and combining insights from the study of societal, technical, and mathematical aspects of computer simulation.

Cooperation between schools and universities

In comments near the end of the event, HLRS Director Prof. Michael Resch praised the high standard and ambitiousness of the students' projects. "I have seen some lectures at professional conferences that were not nearly as impressive as some of these projects," he said as he evaluated the scholars' final presentations.

The six pupils study at the Königin-Charlotte-Gymnasium in Stuttgart-Möhringen, the Schelztor-Gymnasium in Esslingen a.N., the

Friedrich-Schiller-Gymnasium in Marbach a.N., and the Königin-Katharina-Stift in Stuttgart. Technical support was provided by four HLRS-employees—Dr. Ralf Schneider and Alena Wackerbarth supervised the computer model in medical technology, Myriam Guedey supported the team of traffic visualizers, and Dr. Andreas Kaminski served as an expert in the field of technological philosophy.

In addition to HLRS, the Steinbuch Centre for Computing (SCC) in Karlsruhe and the Center for Interdisciplinary Risk and Innovation Studies (ZIRIUS) are also involved in the Simulated Worlds project.

Written by Christopher Williams

NOMAD: The Novel Materials Discovery Laboratory




Fig. 1: NOMAD and internal organization.

The Novel Materials Discovery (NOMAD) Laboratory [1] contains a large repository of materials simulations (repository). NOMAD accepts datasets of all commonly used codes in the field. The data is processed into a code-independent representation (archive), which can then be explored within a materials-oriented view (encyclopedia). Big data analytics allow us to find low-computational-cost descriptors for specific properties and to classify materials. Advanced graphics are used to enable effective data exploration and to create dissemination materials, while HPC provides the supporting infrastructure.

NOMAD is a European Center of Excellence that began in November 2015, with roots in the earlier NOMAD Repository. NOMAD performed its 18-month review in Brussels on June 16, 2017, where the latest advances and success stories in the different work packages were showcased. Very positive feedback was given by the evaluators.

NOMAD asks all researchers in the chemistry and material science field to consider uploading their calculations to the repository. A 10-year storage guarantee and open access sharing possibilities (including DOI support) make NOMAD the largest repository for input and output files of computational materials science codes.

Virtual Reality viewer

Within NOMAD, LRZ is in charge of providing a virtual reality viewer optimized for materials science datasets. The viewer can be used either in combination with the rest of the NOMAD infrastructure, or in a stand-alone fashion [6]. The following dataset types are supported: crystal





Fig. 2: Crystal structure of Nb₈As₄ (Google cardboard), silver Fermi surface and molecular dynamics simulation of pyridine (HTC Vive).



structures, Fermi surfaces, molecular dynamics and electron density calculations (figures 2 and 3). In particular, the virtual reality system has been especially well received to study electron-hole interactions, a.k.a. excitons (figure 4).




Fig. 3: Exploration of a CO₂@CaO electron density simulation (423 timesteps, relative densities) during the NOMAD Data Workshop celebrated at the LRZ in April 2017.

The software has been designed in a modular architecture to allow the use of the multiple SDKs powering virtual reality hardware. In particular, we support the LRZ CAVE-like* environment, HTC Vive (OpenVR SDK), Samsung GearVR (Oculus Mobile SDK) and Google Cardboard (GVR SDK). The viewer was showcased during the 1st NOMAD Data Workshop, celebrated at the LRZ on 25-27 April 2017 (figure 3). The software is also available to interested participants in the LRZ BioLab Summer of Simulation 2017.

Videos created using the NOMAD Virtual Reality viewer




Fig. 4: Excitons in Pyridine@ZnO (above) and in graphene-h-BN heterostructure (below).

The user interaction with the material viewer can be used for teaching or outreach purposes by creating videos of the experience. Figure 4 contains two examples of such movies. The stereoscopic movies visualize excitons in Pyridine@ZnO [2], and in a graphene-hexagonal boron nitride heterostructure [3] (figure 4). The former one was shown with great success in the Berlin Long Night of Research on 24 June 2017.

Videos created using the NOMAD Virtual Reality pipeline (360° stereoscopic)


The pipeline to prepare the datasets for the viewer can also be used to prepare panoramic, stereoscopic movies for outreach purposes. In particular, 3-minute movies were created describing CO₂ adsorption on CaO [4] and excitons in LiF [5] (figure 5). The first video was partially rendered using SuperMUC at the LRZ.


Acknowledgements

The project received funding from the European Union's Horizon 2020 research and innovation program under grant agreement no. 676580 with The Novel Materials Discovery (NOMAD) Laboratory, a European Center of Excellence.

Olga Turkina provided the Pyridine@ZnO dataset and Wahib Aggoune provided the graphene-BN heterostructure. The CO₂@CaO dataset was provided by Sergei Levchenko, and the Ag Fermi surface was provided by Artur Garcia. Raison Dsouza provided the pyridine simulation. Andris Gulans recorded the videos shown in figure 4.

*CAVETM is a trademark of the University of Illinois Board of Trustees. We use the term CAVE to denote the both the original system at Illinois and the multitude of variants developed by multiple organizations.





References

- [1] <https://nomad-coe.eu/>
- [2] [Charge-transfer excitons at organic-inorganic interface](#)
<https://www.youtube.com/watch?v=2cOmQp6RYXA>
- [3] [Exciton in a graphene/BN heterostructure](#)
<https://www.youtube.com/watch?v=tQrAPuFpFh8>
- [4] [A 360° movie by NOMAD:](#)
Conversion of CO₂ into Fuels and Other Useful Chemicals https://youtu.be/zHIS_8PwYYs
- [5] [A 360° movie by NOMAD:](#)
An Exciton in Lithium Fluoride - Where is the electron? <https://youtu.be/XPPDeePlcoM>
- [6] [García-Hernández, R. J., Kranzlmüller, D.:](#)
Virtual Reality toolset for Material Science: NOMAD VR tools. 4th International Conference on Augmented Reality, Virtual Reality and Computer Graphics. Lecture Notes on Computer Science, no 10324, Part I, pp 309-319. Springer, 2017.

Written by Rubén Jesús García-Hernández

Leibniz-Rechenzentrum.

Contact: garcia@lrz.de

The ClimEx project: Digging into Natural Climate Variability and Extreme Events

Scientific context

Climate models are the basic tools used to support scientific knowledge about climate change. Numerical simulations of past and future climates are routinely produced by research groups around the world, which run a variety of models driven with several emission scenarios of human-induced greenhouse gases and aerosols. The variety of such climate model results are then employed to assess the extent of our uncertainty on the state of the future climate. Similarly, large ensembles of realizations using a single model but with different sets of initial conditions allow for sampling another source of uncertainty—the natural climate variability, which is a direct consequence of the chaotic nature of the climate system. Natural climate variability adds noise to the simulated climate change signal, and is also closely related to the occurrence of extreme events (e.g. floods, droughts, heat waves). Here, a large ensemble of high-resolution climate change projections was produced over domains covering north-eastern North America and Europe. This dataset is unprecedented in terms of ensemble size (50 realizations) and resolution (12km) and will serve as a tool to implement robust adaptation strategies to climate change impacts that may induce damage to several societal sectors.

The ClimEx project [1] is the result of more than a decade of collaboration between Bavaria and Québec. It investigates the effect of climate change on natural variability and extreme events with a particular focus on hydrology. In order to better understand how climate

variability and meteorological extremes translate into basin-scale hydrological extremes, a complex modelling chain is being implemented. Here we present the outcome of the climate simulation phase of the project that consisted in the production of 50 Regional Climate Model (RCM) simulations run over two domains. This throughput-computing exercise was conducted during year 2016 and 2017 on the SuperMUC supercomputer at the Leibniz Supercomputing Centre (LRZ), requiring a total of 88 million core-hours of resources.

Hydro-climate modelling chain

The ClimEx modelling framework involves three layers, as different spatial and time scales need to be modelled, from global climate changes to basin-scale hydrological impacts. First, a Global Climate Model (GCM) simulates the climate over the entire Earth's surface with typical grid-space resolutions ranging between 150 and 450 km. The GCM's coarse-resolution outputs can then be used as input forcing for boundary conditions of an RCM. An RCM concentrates computational resources over a smaller region, thus allowing the model to reach spatial resolutions of the order of ten kilometers. As the third layer, a hydrological model uses the high-resolution meteorological variables from the RCM simulation and runs simulations over one particular basin in resolutions of tens to hundreds of meters.

The current setup involves 50 realizations of this 3-layer modelling cascade run in parallel. The Canadian Earth System large ensemble

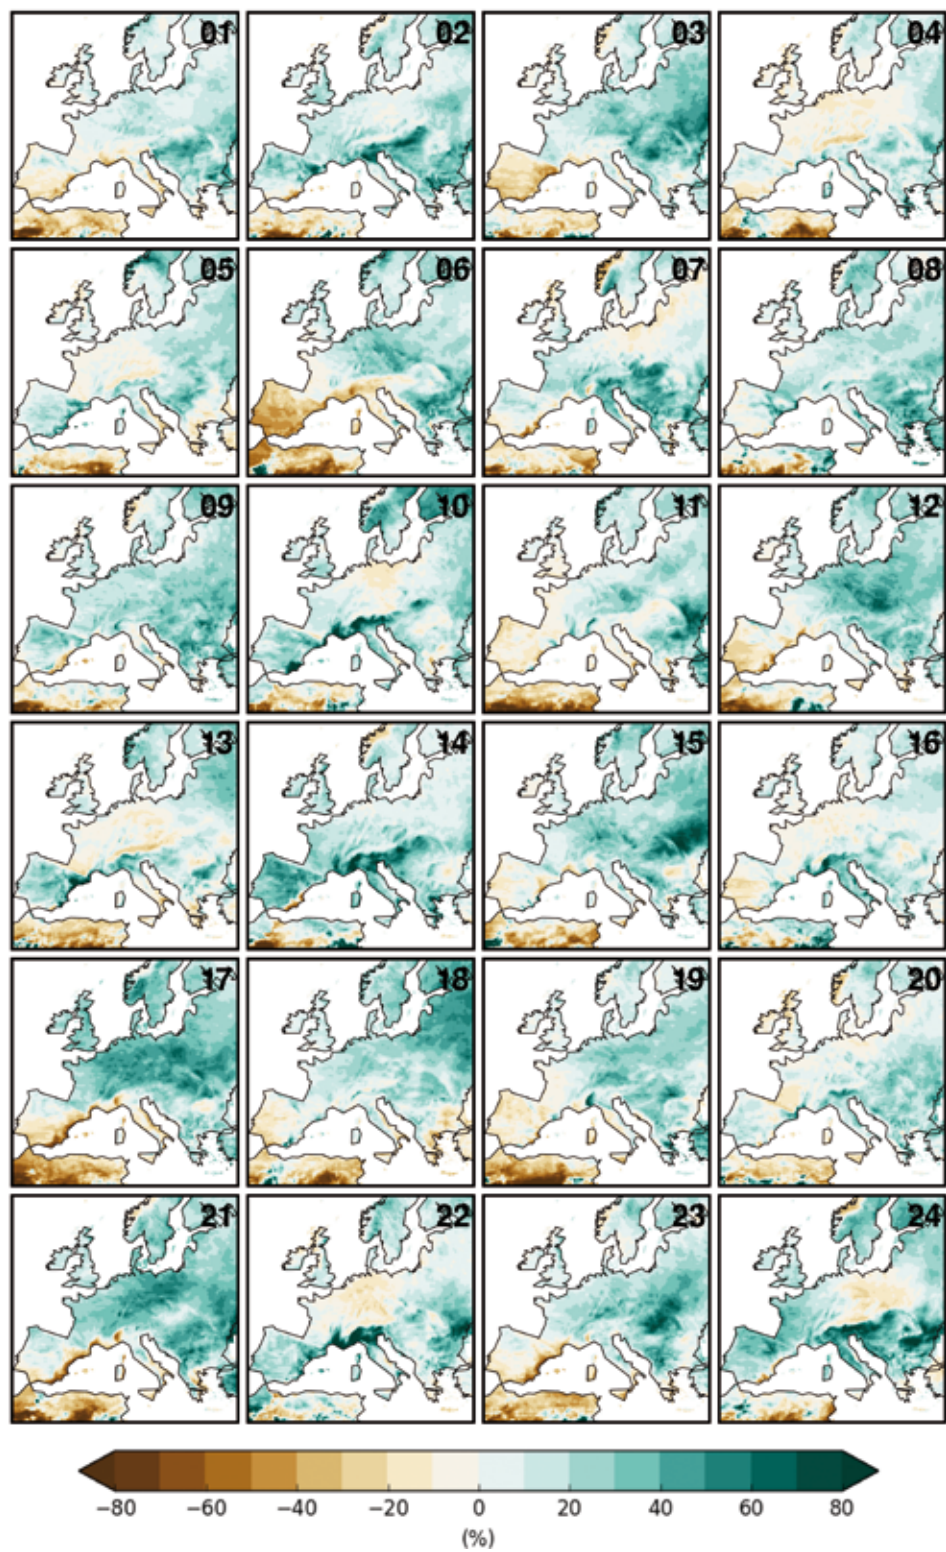



Fig. 1: Climate-change projections of the January mean precipitation over Europe (2040-2060 vs. 1980-2000).

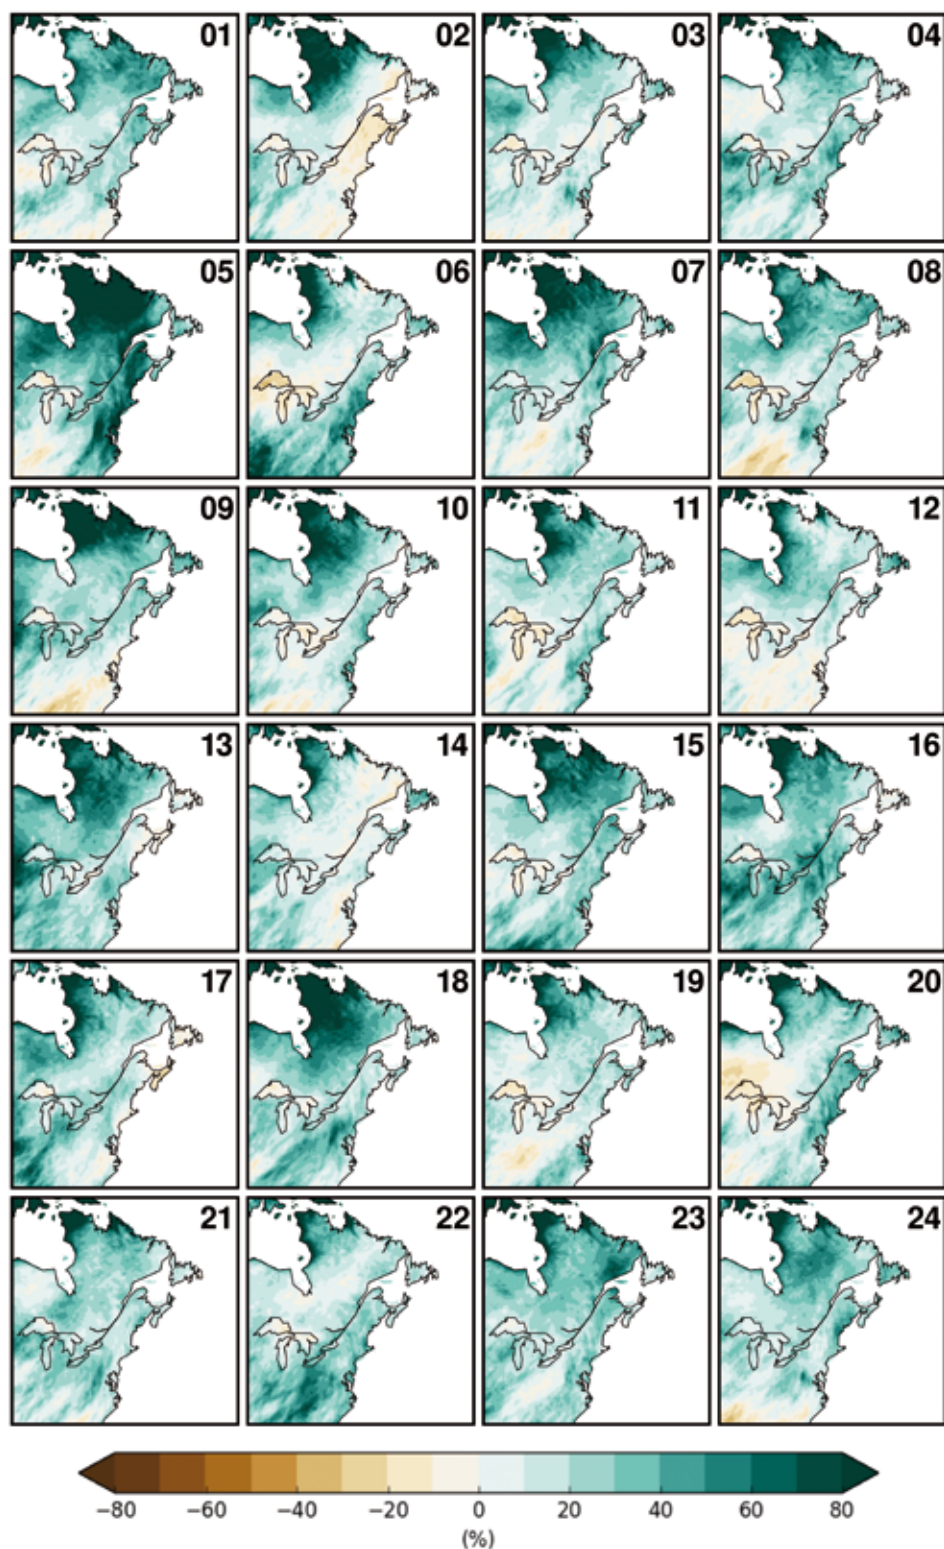



Fig. 2: Climate-change projections of the January mean precipitation over north-eastern North America (2040-2060 vs. 1980-2000).



(CanESM2-LE; [2]) consists of 50 realizations from the same GCM at the relatively coarse resolution of 2.8° ($\sim 310\text{km}$). These realizations were generated after introducing slight random perturbations in the initial conditions of the model. Given the non-linear nature of the climate system, this procedure is widely used to trigger internal variability of climate models, which can be quantified as the spread within the ensemble. All 50 realizations were run using the same human-induced greenhouse gases and aerosols emission pathways (also known as RCP 8.5) as well as natural forcing like aerosol emissions from volcanoes and the modulating incoming solar radiation. CanESM2 is developed at the Canadian Centre for Climate Modelling and Analysis of Environment and Climate Change Canada (ECCC). In the climate production phase described below, the 50 CanESM2 realizations were dynamically downscaled using the Canadian Regional Climate Model version 5 (CRCM5; [3]) at 0.11° ($\sim 12\text{km}$) resolution over two domains in order to cover Bavaria and Québec (see Figures 1 and 2). CRCM5 is developed by Université du Québec à Montréal (UQAM) in collaboration with ECCC. The upcoming phase of the ClimEx project focuses on hydrology, where all simulations over both domains will serve as driving data for hydrological models that will be run over different basins of interest in Bavaria and Québec.

Production of the CRCM5 large ensemble (CRCM5-LE)

Climate models allow to numerically resolve in time the governing equations of the climate

system over a gridded spatial domain. Such models are expensive to run in terms of computational resources because high resolution and long simulation periods are generally required for climate change impact assessments. Here, the CRCM5 was run over two domains using a grid of 380×380 points (i.e. the integration domain). An analysis domain of 280×280 is finally extracted to prevent boundary effects which are well known in the regional climate modelling community (e.g [4] and [5]). The CRCM5-LE thus consists in 50 numerical simulations per domain of the day-to-day meteorology covering the period from 1950 to 2100. The size of the final dataset is about 0.5 petabytes and includes around 50 meteorological variables. The choice of archived variables and time resolution (e.g. hourly for precipitation) was defined in collaboration with project partners and was based on a balance between disk space and priorities for future projects.

Before going into massive production, the workflow, including simulation code and job farming, was optimized for a minimal core-hour consumption and high throughput on SuperMUC. The best compromise was found when running multiple instances of CRCM5 in parallel, each utilizing 128 cores, and an aimed average total utilization of 800 SuperMUC nodes in parallel. The CRCM5-LE was produced in the scope the 14th Gauss Centre Call for Large-Scale Projects, where 88 million core hours were granted on SuperMUC at the Leibniz Supercomputing Centre (LRZ). These resources were spent during a massive production phase and successfully

completed at the end of March 2017. A small portion of the CPU-budget was dedicated to data management and post-processing resulting in a dataset based on standardized data formats (NetCDF convention) including metadata for reproducibility of the scientific workflow.

Results

In Figures 1 and 2, climate change projections of the January mean precipitation are shown over Europe and northeastern North America respectively. For simplicity, only 24 realizations out of 50 are shown for each domain. The climate-change signal is expressed in percentages and represents the relative change in January mean precipitation in the middle of the 21st century (from 2040 to 2060) compared with a recent-past reference period (from 1980 to 2000).


Recalling that simulations from the ensemble differ solely by slight random perturbations in their initial conditions and that the exact same external forcing (GHGA) was prescribed in every case, these figures allow us to appreciate the magnitude of the natural variability existing in the climate system. The ensemble spread at different geographical locations may represent a wide range of outcomes that are permitted by the chaotic behaviour of the climate system. For instance, January mean precipitation in Spain shows a 40% decrease for realization 6 while a 40% increase appears in realization 22 (Figure 1). A similar situation appears in the southern part of the North American domain for realizations 8 and 22 (Figure 2). Features with alternating

sign between individual realizations are considered uncertain, while other features persist consistently throughout the ensemble, and may therefore be considered robust. Good examples are the precipitation decrease in northern Africa (Figure 1) or the precipitation increase in northern Québec (Figure 2), which are detected in all simulations.

These results highlight the importance of performing ensembles of several realizations to assess the robustness of estimated climate-change patterns. However, it is worth noting that these results are specific to the combination between CanESM2 and CRCM5, but many other GCMs and RCMs could be considered as well. Therefore, one caveat of this framework is that it does not address the epistemic uncertainty of climate-change projections, but the aleatory uncertainty associated with the regional CanESM2/CRCM5 climate system is assessed with a degree of robustness that was unprecedented until now.

Perspectives of the ClimEx project

First results of the CRCM-LE were presented during the 1st ClimEx Symposium that took place on June 20th-21st, 2017 at the Ludwig-Maximilian University of Munich. This meeting brought together climate scientists, hydrologists and other impact modellers, as well as decision makers to discuss most recent findings on the dynamics of hydrometeorological extreme events related to climate change. In this context, good contacts were established with other researchers who engage in the analyses of large



scale single model ensembles and it was agreed upon to exchange data and information on this joint research topic. An official announcement was made that the ClimEx dataset will become publicly available to the community during 2018, following a thorough quality control phase and preliminary analyses by the ClimEx project team and close partners.

The project group is currently working on the refined calibration of the hydrological models which are to be driven with processed CRCM-LE data in the case studies in Bavaria and Québec to assess the dynamics of hydrological extremes under conditions of climate change. It is intended that the analysis of hydro-meteorological extremes in the context of water resources is only the first step in a sequence of scientific projects to explore the full capacity of this unique dataset. Potential application cases are obvious in agriculture and forestry, but also in the health or energy sector.

References

- [1] www.climex-project.org
- [2] **Fyfe, J. C. and Coauthors, 2017:**
Large near-term projected snowpack loss over the western united states. *Nature Communications*, 8, 14996, doi:10.1038/ncomms14996.
<https://doi.org/10.1038/ncomms14996>
- [3] **Šeparović, L., A. Alexandru, R. Laprise, A. Martynov, L. Sushama, K. Winger, K. Tete, and M. Valin, 2013:**
Present climate and climate change over north america as simulated by the fifth-generation canadian regional climate model. *Clim Dyn*, 41, 3167–3201, doi:10.1007/s00382-013-1737-5.
<http://dx.doi.org/10.1007/s00382-013-1737-5>
- [4] **Leduc, M., and R. Laprise, 2009:**
Regional climate model sensitivity domain size. *Clim. Dyn.*, 32, 833–854.
- [5] **Matte, D., R. Laprise, J. M. Thériault, and P. Lucas-Picher, 2016:**
Spatial spin-up of fine scales in a regional climate model simulation driven by low-resolution boundary conditions. *Climate Dynamics*, nil, nil, doi:10.1007/s00382-016-3358-2.
<http://dx.doi.org/10.1007/s00382-016-3358-2>

**Written by Martin Leduc, Anne Frigon,
Gilbert Brietzke, Ralf Ludwig, Jens Weis-
müller, Michel Giguère**

Ludwig-Maximilians-Universität München

Contact: Prof. Dr. Ralf Ludwig, Faculty of Geosciences,
Department of Geography, r.ludwig@lmu.de



Centers / Systems

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, and by
- 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. Kranzlmüller

Boltzmannstraße 1, 85748 Garching near Munich
Germany

Phone +49-89-358-31-80 00

kranzlmue@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

HLRS

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 computing 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



The Jülich Supercomputing Centre (JSC) at Forschungszentrum Jülich is committed to enabling scientists and engineers 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 of the highest performance class for projects in science, research and industry in the fields of modeling and computer simulation including their methods. The selection of the projects is performed by international peer-review procedures implemented by the John von Neumann Institute for Computing (NIC), GCS, and PRACE.



JSC's supercomputer "JUQUEEN", an IBM Blue Gene/Q system.



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

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

laboratories and cross-sectional teams, e.g. on mathematical methods and algorithms and parallel performance tools, enabling the effective usage of the supercomputer resources.

Compute servers currently operated by JSC

System	Size	Peak Performance (TFlop/s)	Purpose	User Community
IBM Blue Gene/Q "JUQUEEN"	28 racks 28,672 nodes 458,752 cores IBM PowerPC® A2 448 Tbyte memory	5,872	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 Title	Location	Date
Scientific Visualization	Stuttgart	Nov 6-7, 2017
Software Development in Science	Jülich	Nov 20, 2017
Advanced C++, Focus on Software Engineering	Stuttgart	Nov 20-23, 2017
Vectorisation and Portable Programming Using OpenCL	Jülich	Nov 21-22, 2017
Introduction to the Programming and Usage of the Supercomputer Resources at Jülich	Jülich	Nov 23-24, 2017
Advanced Parallel Programming with MPI and OpenMP	Jülich	Nov 27-29, 2017
Fortran for Scientific Computing	Stuttgart	Nov 27-Dec 1, 2017
Node-Level Performance Engineering (PATC course)	Garching	Nov 30-Dec 1, 2017
Molecular Modeling with Schrödinger-Suite	Garching	Dec 5-7, 2017
Parallel and Scalable Machine Learning (PATC course)	Jülich	Jan 15-17, 2018
Introduction to Hybrid Programming in HPC (PATC course)	Garching	Jan 18, 2018
Intel KNL Many-Core - Usage and Profiling	Jülich	Feb 2018 (tba)
Parallel Programming (MPI, OpenMP) and Tools	Dresden	Feb 12-16, 2018
Programming with Fortran	Garching	Feb 14-16, 2018
Introduction to Computational Fluid Dynamics	Siegen	Feb 19-23, 2018
NIC Symposium 2018	Jülich	Feb 22-23, 2018
CFD with OpenFOAM®	Stuttgart	Mar 5-9, 2018
OpenMP and OpenACC GPU Directives for Parallel Accelerated Supercomputers (PATC course)	Stuttgart	Mar 12-13, 2018 (tbc)
Parallel I/O and Portable Data Formats (PATC course)	Jülich	Mar 12-14, 2018
Parallel Programming of High Performance Systems	Erlangen	Mar 12-16, 2018
Advanced C++, Focus on Software Engineering	Stuttgart	Mar 14-16, 2018
Parallel Programming with HPX	Stuttgart	Mar 15-16, 2018 (tbc)
Introduction to Parallel Programming with MPI and OpenMP	Jülich	Mar 19-22, 2018
Iterative Linear Solvers and Parallelization	Stuttgart	Mar 19-23, 2018
Advanced Topics in High Performance Computing (PATC course)	Garching	Mar 26-29, 2018
Introduction to Parallel In-Situ Visualization with VisIt	Jülich	Apr 2018 (tba)
Usage of VTK for Scientific-technical Visualization	Jülich	Apr 2018 (tba)
Fortran for Scientific Computing (PATC course)	Stuttgart	Apr 9-13, 2018
From zero to hero: Understanding and Fixing Intra-node Performance Bottlenecks	Jülich	Apr 11-12, 2018
Cray XC40 Workshop on Scaling and Node-Level Performance	Stuttgart	Apr 23-26, 2018
VI-HPS Tuning Workshop (PATC course)	Garching	Apr 23-27, 2018
Scientific Visualization	Stuttgart	May 7-8, 2018
Introduction to the Programming and Usage of the Supercomputer Resources at Jülich	Jülich	May 28-29, 2018
Intel MIC / Knights Landing Programming Workshop (PATC)	Garching	Jun 2018 (tba)
High-performance Scientific Computing in C++ (PATC course)	Jülich	Jun 11-13, 2018
Node-Level Performance Engineering (PATC course)	Stuttgart	Jun 14-15, 2018 (tbc)
High-performance Computing with Python (PATC course)	Jülich	Jun 18-19, 2018
Introduction to Hybrid Programming in HPC	Stuttgart	Jun 19, 2018
Cluster Workshop	Stuttgart	Jun 20-21, 2018
Efficient Parallel Programm. with GASPI (PATC course)	Stuttgart	Jul 2-3, 2018
Introduction to UPC and Co-Array Fortran (PATC course)	Stuttgart	Jul 5-6, 2018
Introduction to Cluster Filesystems	Stuttgart	Jul 2018 (tba)



You will also find the detailed course list online at
<http://inside.hlr.de/courses.html>
or by scanning the QR-Code with your mobile device.

Complete and updated list of all GCS courses:

http://www.gauss-centre.eu/gauss-centre/EN/Training/Training/_node.html

Further training courses at members' sites:


<http://www.fz-juelich.de/ias/jsc/courses>


<http://www.lrz.de/services/compute/courses/>

<http://www.hlr.de/de/training/>



© Forschungszentrum Jülich





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

Scientists from the Institute of Aerodynamics (AIA), RWTH Aachen University, and the Jülich Supercomputing Centre (JSC), Forschungszentrum Jülich, aim at personalizing medical treatment of ear, nose, and throat patients together with industrial partners within the IraSME project Rhinodiagnost. Large-scale flow computations on HPC systems help them to better understand respiratory diseases and to find optimized surgery strategies. The image on the front matter depicts the results of such a simulation for an inspiration phase. The flow is visualized by streamlines, colored by the velocity magnitude. The streamlines bend around the complex anatomical structures in the nasal cavity and the flow is accelerated in locally converging regions. Furthermore, the wall-shear stress is mapped to the tissue surface. The insets shows the computational boundary-refined mesh that is necessary to discretize the governing equations of fluid flows and to run the simulation on the HPC infrastructure at JSC.