



***LRZ newsletter, October 2022: Readings about supercomputing and Quantum***

NEWS

[“We are planning a series of quantum systems with superconducting technologies”](#)

[Controlling the power demand in supercomputing](#)

[Simulate the smallest vortices of turbulence](#)

[Let there be light](#)

[Accelerate Supercomputing](#)

WORKSHOPS & EVENTS

USED THINGS FOR FURTHER SERVICE

JOB OFFERS

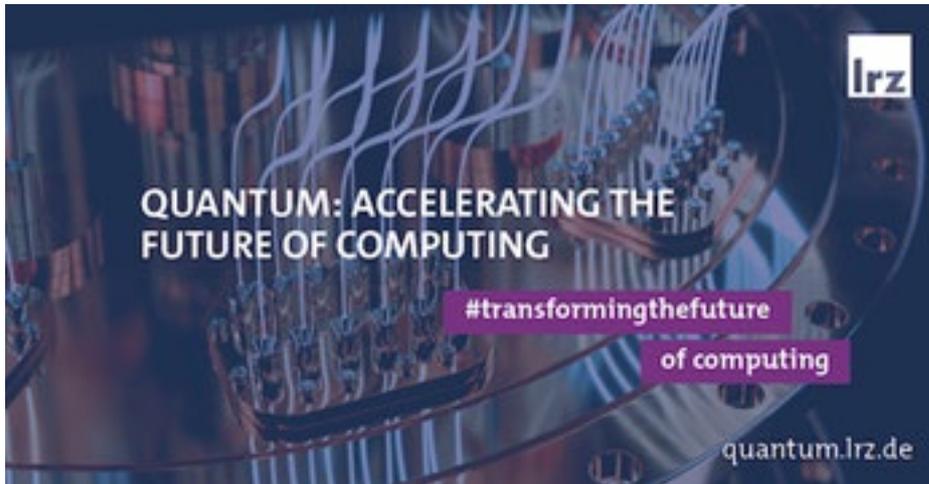
MORE TO READ

INFORMATIONS & IMPRINT

---

## NEWS

### "We are planning a series of quantum systems with superconducting technologies"



The Leibniz Supercomputing Centre (LRZ) will become a hosting site for **European quantum computing**. This was decided by the Governing Board of the European High Performance Computing Joint Undertaking (EuroHPC JU) in early October 2022.

The LRZ-idea for the project "European Quantum Computing for Exascale-HPC", or Euro-Q-Exa for short, is: **to integrate quantum processors into supercomputing** and make the new computer technology more controllable and accessible to users in this way. "We are planning a series of quantum systems with superconducting technology that will eventually comprise up to 100 qubits," says **Prof. Dr. Martin Schulz**, Member of the Boards of Directors at LRZ. "We will also integrate a quantum processing unit directly into the planned next supercomputer for the LRZ." Euro-Q-Exa is for the time being the last research project in the field of quantum computing that the LRZ has been able to win for itself. "Locating the quantum computing system also helps us to further develop services for users and is another important milestone on the way to a next generation of accelerated high performance computing or HPC systems," notes Laura Schulz, head of the Quantum Computing and Technologies department. The new plans give reason to ask about the state of quantum computing at the LRZ and its integration into supercomputing: **The entire interview online**.

---

### Controlling the power demand in supercomputing



Efficiency is in demand: not only because of the current energy crisis is the power demand of a data centre one of its most important key figures. Technology can fail if computers are permanently driven to peak performance; moreover, delivery quantities are limited. If negotiated limits are exceeded, electricity becomes significantly more expensive. For its energy management, the Leibniz

Supercomputing Centre (LRZ) relies on the **open-source software Energy Awareness Runtime (EAR)**, and this is the focus of a study that Carla Guillen and Carmen Navarrete, both computer scientists with doctorates at the LRZ,

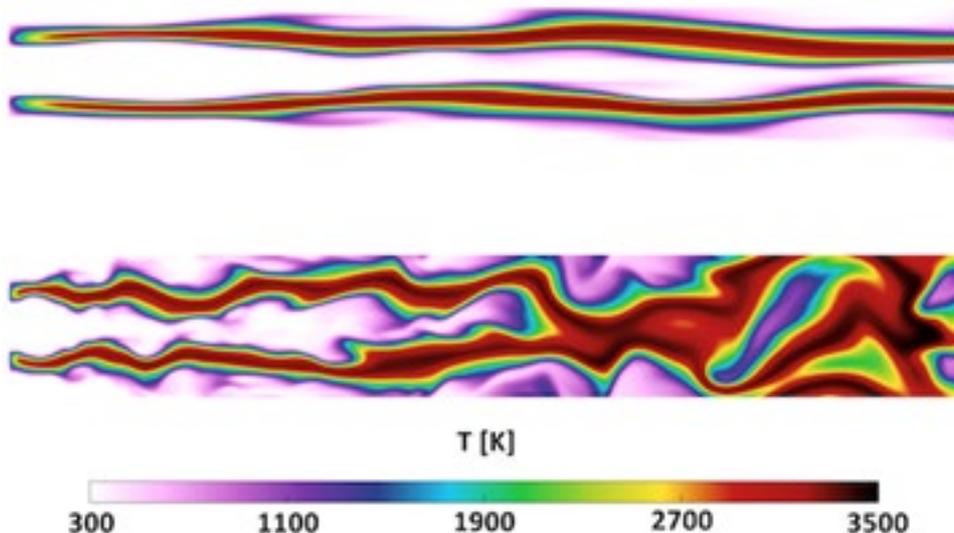
conducted with their colleagues Julita Cobalan and Lluís Alonso from the Barcelona Supercomputing Center (BSC), and which they presented at a workshop on energy efficiency.

The paper "**Soft Powercap at SuperMUC-NG with EAR**" describes how hard and soft powercap limits can be realised in the operation of supercomputers and even controlled down to the application level: "We analysed and optimized EAR's powercap management system for this purpose," the team explains. "SuperMUC-NG is also subject to a power limit, but this can be exceeded for a short time as long as the power is below the limit on average." The team tested how **EAR** addresses which levels of a system - processors, nodes, clusters - for such soft limits and distributes the power in the system on an island of SuperMUC-NG, consisting of 792 nodes, each with two processors and 48 cores, and a memory capacity of 96 gigabytes. Key findings: To maintain a hard, fixed performance limit, EAR only lowers the clock frequency of the processors. Now, limit values can be entered to control a soft limit - for example, from monitoring data, i.e. information on the utilisation of the system. In this way, EAR controls the power demand within a corridor so that the limit is not exceeded. EAR controls upper and lower power limits with its hard power capping algorithms, such that the power demand remains in a corridor. The programme not only controls the processors, but also the energy supply at node and even application level: "At the node level, we do not necessarily have to pay attention to an even and constant distribution of power," the paper says. "At the node, EAR controls the power consumption below the power limit, here the power demand can be dynamically adjusted to the requirements of an application." The working group presented its experiences with EAR during the **13th International Green and Sustainable Computing Conference** at the end of October.

---

### Simulate the smallest vortices of turbulence

The random flows in liquids and gases, the turbulences, still pose many questions for the natural sciences. At **Chair of Astronautics at the Technical University of Munich (TUM)**, a research team led by engineers **Andrej Sternin** and **Daniel Martinez** has used Direct Numerical Simulations (DNS) to calculate turbulence in space propulsion systems (the graphic at the end of text shows a Methane-oxygen diffusion flame at low (top) and high (bottom) Reynolds number).. They succeeded - with the support of SuperMUC-NG and the specialists of the **Computational X Support** of the Leibniz Computing Centre (LRZ). Along the way, new methods were developed with which the smallest vortices of turbulence in particular can be represented. Normally, these can only be determined for short periods of time, but with "quasi-DNS" simulations, it was possible in this case to reduce the computational effort and, for the first time, also model the mini-vortices and their influence on the turbulence: "This was no regular customer service," Sternin and Martinez praised the work of the LRZ specialists and especially Martin Ohlerich. "Without his help, it would not have been possible to carry out our simulations on this scale." More on this productive teamwork for space research can be found at the **Gauss Centre for Supercomputing (GCS)** and at **HPC Wire**



---

### Let there be light

Ancient Pompeii shone very brightly: Oil lamps, candelabras, lamp sculptures or torches stood and hung everywhere in the submerged city beneath Mount Vesuvius. The interplay of these different light sources was

quite deliberate. Light was used to modulate geometrical design on walls at festivals and on streets, to create artistic spaces and to evoke a special atmosphere. This effect of light can be relived at the Staatliche Antikensammlungen (State Collections of Classical Antiquities) in Munich from **8 November 2022 to 2 April 2023**: The exhibition "[New Light from Pompeii](#)" not only presents nearly 130 antique lamps and lights - including loans from the National Archaeological Museum of Naples, rediscoveries from depots and replicas - but also uses animations and simulations to make light and its effect in ancient Pompeii understandable and comprehensible. Contemporary light installations by the [Munich design-team Ingo Maurer](#) complete the exhibition and show how the ancient knowledge of light has evolved. The exhibition is part of a research project at Ludwig-Maximilians-Universität (LMU) led by archaeologist [Prof. Dr. Ruth Bielfeldt](#), and the [Centre for Visualisation and Virtual Reality](#) at Leibniz-Rechenzentrum supported the researchers in their simulations of the lighting effects..

---

## Accelerating Supercomputing

13 November 2022 marks the start of the [Supercomputing 2022 or SC22](#) trade fair and conference in Dallas, where exascale and quantum computing will take centre stage. For six days, the high-performance computing or HPC community will exchange ideas, either on-site or online, and discuss how supercomputing can be



further accelerated and energy efficiency improved. In addition to graphic processing units (GPU), field programmable gate arrays (FPGA) and other accelerators, many hopes rest on the [new quantum processors \(QPU\)](#). The Leibniz Supercomputing Centre (LRZ) is involved in [eleven research projects that focus on the integration of QPUs](#) in supercomputers, the development of control and monitoring options as well as a programming environment for this new technology and the security questions it poses.

Representatives of the [Munich Quantum Valley \(MQV\)](#) and project partner [IQM Quantum Computers](#) will be present at booth no. #1035. The fledging company is contributing its first quantum processors and experience to the planned integration. LRZ specialists will report in workshops and lectures on how QPUs can be integrated into HPC, what the first computing schemes might look like and how hybrid HPC systems can work. Further topics include the processing and analysis of operational data in HPC to reduce energy requirements; how supercomputing can improve disaster prevention;

and last but not least, how and where to find IT and quantum experts that the LRZ and its partners need to develop quantum computing further. Some events will be broadcast online, the time difference between Munich and Dallas is 7 hours (Munich 5 pm = Dallas 10 am).

Lectures and workshops involving LRZ experts

- **November 13, starting from 8.30 a.m.:** "[Latest Advances in Scalable Algorithms for Large-Scale Heterogeneous Systems](#)" ([ScalAH'22](#)): This workshop deals with the computational schemes that hybrid or heterogeneous supercomputing architectures require when different types of processors operate in them. Prof. Dr. Dieter Kranzlmüller will co-host this event with colleagues from other computing centres
- **Starting also at 8.30 the 3rd workshop on quantum computing software.** Burak Mete from the LRZ Quantum Computing and Technology team will present his paper on "Predicting the Optimizability for Workflow Decisions".

- **11:10 to 11:30:** Dr. Tobias Guggemos will show what a [compiler for universal, photonic quantum computers](#) might look like. The lecture is the result of a [joint research project](#) of the University of Vienna, the Ludwig Maximilian University and the LRZ ().
- **November 14, 9:30 a.m. - 10:00 a.m.:** At the [7th Workshop on Extreme Scale Programming Models and Middleware](#), Prof. Dr. Martin Schulz will talk about the Message Passing Interface (MPI) and its limitations in systems that combine different computing resources. The MPI sessions could provide more opportunities
- **November 15, 12:15 - 13:15:** [The CO2 footprint of supercomputing](#) - in six lectures, computer scientists will show how high the CO2 emissions of supercomputing are and how they can be reduced through clever management and tools. You will hear also solutions from LRZ, Prof. Dr. Dieter Kranzlmüller has aligned this session with colleagues.
- **17:15 to 18:45:** [Performance comparison of HPC architectures](#): In order to better assess computers and their performance, data centres rely on specially developed benchmarks. In five short presentations, representatives of international data centres will present the benchmarks they work with. Dr. Josef Weidendorfer, Head of Future Computing, will show which benchmarks the LRZ works with. The aim of the event: to question existing comparative values and to improve them in view of new, diverse architectures.
- **17:15 - 18:45:** The Partnership for Advanced Computing in Europe (PRACE) is dedicated to [European HPC](#) in many short presentations. Dr. Volker Weinberg, coordinator of the LRZ training programme, and Dr. Michele Martone will focus on their own codes and how they can be implemented on supercomputers.
- **November 16, 12.15-13.15:** The [Message Passing Interface \(MPI\)](#) is one of the most important program schemas in HPC. The users are already thinking about further developments of this open source program. So where is the standardization round MPI 5.0 heading? Prof. Dr. Martin Schulz will lead this discussion round.
- **17:15 - 18:45:** [Introducing water cooling](#). Water cools computers more sustainably than air. Water cooling systems help save electricity, especially when they use warm water. Dr. Michael Ott, who works on energy efficiency at the LRZ, shows how SuperMUC-NG achieves best PUR values.
- **17:15 - 18:45:** How can parallelism be increased and improved in HPC systems that combine diverse technologies? In the panel discussion ["Enabling I/O and Computation Malleability in HPC"](#) Prof. Schulz and international colleagues seek answers and solutions.
- **November 17, 12:15-13:45:** [Analysis of operating data](#). When it comes to saving power in HPC, computer scientists and developers are increasingly focusing on operating data. With their help, applications can be better planned and software better tailored to the needs of a computer. LRZ staff member Dr. Michael Ott will present practical experiences and tools for monitoring and evaluating data.
- **November 18, 8:30 - 10:00:** [Calculating with the Unexpected](#) - Disaster Prevention with the Help of Supercomputing. Supercomputers calculate the models that researchers build for earthquakes, climate and environmental developments. These have become more and more precise and detailed in recent years - and now form the basis for forecasting tools and for planning protective measures. Prof. Dr. Dieter Kranzlmüller, Director of the LRZ, presents some samples of the LRZ's work.

---

## WORKSHOPS & EVENTS

### ANSYS Fluent for Computational Fluid Dynamics

Consisting of 13 lectures and about 6 practical exercises, this seminar, which runs from **October 27 to December 8, 2022**, introduces the use of the fluid dynamics software package ANSYS Fluent. This supports computations and simulations of computational fluid dynamics in particular. Participants will learn about typical CFD workflows for ANSYS and application examples. In addition, they learn how ANSYS works and is integrated on the LRZ Linux cluster. For this they will receive the Linux Primer. Participants should know the basics and numerical methods of fluid mechanics and have completed first calculations. [Information & Registration](#)

### Introduction to the Classiq Platform

Manually defining a connection between qubits and gates or matching prefabricated quantum blocks works well for processors with 5 or 15 qubits. But as quantum processing units (QPU) grow, where 100 or even 1000 qubits come together, this is no longer effective. The Classiq development platform shows alternatives: In a

tutorial on November 7, 2022, quantum specialists can learn to develop software and algorithms for larger quantum computers. On the agenda are the most important functional and algorithmic blocks and how to combine them for software, also to consider hardware constraints, optimize circuits for properties such as depth, number of qubits, precision and more. Methods for pure quantum and for hybrid algorithms will be discussed. Participants should be familiar with gate-level design tools (such as Qiskit) and with the operational concept of quantum computing algorithms (e.g. Grover, VQE, QAOA). [Information and registration](#)

### OneAPI-Training

The oneAPI initiative ([www.oneapi.com](http://www.oneapi.com)) is a cross-industry, open, standards-based programming model for different vendors' CPU and accelerator architectures. It enables supercomputing researchers to accelerate application performance and get more productivity out of computers. oneAPI and, more importantly, the toolkits Intel has released for it, will be the focus of a three-day course at **November 8 to-10, 2022**. Participants will learn how to use and exploit the programming standard for their projects. [Information and registration](#)

### Improving and accelerating supercomputing

Arrived in the exascale era: The latest supercomputers from the U.S. and Japan can already perform a trillion calculations or FLOP per second, and their performance and new exascale systems are the focus of SC2022 from **November 13 to 18, 2022**, the international conference on high-performance computing (HPC), networks, storage and data analysis, which is being held this year in Dallas. Of course, it's all about innovative supercomputing technology, but above all it's about how supercomputers can compute and work even faster with the help of quantum computing and artificial intelligence (AI) methods. Speed is also a form of energy efficiency in computing, and ways to achieve more economical cooling will also be discussed at SC2022. [Information und registration](#)

### Developing algorithms with Fortran

Software design and object-oriented programming with Fortran: At the four-day online workshop from **November 21 to 24, 2022**, participants can deepen and professionalize their knowledge of Fortran. They will learn good examples of robust code, as well as the tricks of combining Fortran. programs with applications in C languages, plus how to optimize their algorithms for parallel computer systems. [Information and registration](#)

### OpenFOAM for Supercomputing

OpenFOAM is an open source C++ framework of solvers and tools for numerically solving partial differential equations using the finite volume method. It is embedded in the Linux Bash script workflows and interfaces with the Paraview analysis tool for pre- and post-processing and debugging. Using openFOAM is the focus of this tutorial course from **November 23 to December 1, 2022**. After an introductory day of self-study, the course mixes self-study with two additional online tutorial sessions where participants can clarify comprehension questions and deepen their use of openFOAM. [Information and registration.](#)

### Faster Computing with CUDA and C/C++

The CUDA platform can be used to accelerate and optimize CPU-only applications for execution on GPU. In a workshop on November 28, 2022, participants will learn how to use CUDA as well as the C and C++ programming schemes and how iterative development works in steps. [Information and registration](#)

### Working with OpenMP

Since its introduction, the OpenMP programming model has proven to be the driving force behind parallel programming for shared memory architectures. These are gaining relevance with the proliferation of multicore processors: Good reason to take a closer look at OpenMP and programming in the three-day course from **November 29 to December 1, 2022**. The instructors will provide theory and background in lectures, and participants will practice what they have learned in practical tasks. [Registration and information](#)

### Dealing with processors and nodes

Even application developers who are familiar with OpenMP and MPI often don't know exactly how much performance their code can achieve. Parallelism is not everything for good computing performance. That's why this three-day online course from **December 5-7, 2022**, hosted by PRACE, provides the necessary understanding of how software and hardware interact and work together. Participants will learn to address the core, socket, and node levels with their codes already, and architectural features and the bottlenecks of processors and compute nodes will also be discussed. [Registration and information](#)

## Introduction to LRZ systems and flow simulations

Simulating gases, liquids, flows: These are tasks that are carried out particularly frequently on SuperMUC-NG and the LRZ's High Performance Computers (HPC). On **December 14, 2022**, Master's students and doctoral candidates can get an idea of the possibilities, get to know the Linux cluster, the LRZ supercomputer and software and algorithms for computational fluid mechanics or computational fluid dynamics. And of course it's also about the workflows - how to access and use the resources and plan or organise your own work. [Information and registration](#)

## Visualizing Research with VisIt and OSPRay Studio

The LRZ offers researchers VisIt and OSPRay Studio for the visualization of research data. Remote access enables distributed renderings on multiple compute nodes and thus large visualizations on SuperMUC-NG & Co. The workshop on December 15, 2022 shows different techniques of visualization with the programs VisIt and OSPRay and presents the necessary workflow from data import to image or movie production. In addition, the instructors will demonstrate the basics of visualizations and show excellent case studies. [Information and registration](#)

## Programming for parallel HPC systems

How can scientists develop programs and algorithms for SuperMUC-NG, CoolMUC and other parallel high performance computers and implement their applications on the systems? This three-day online course, March 7-9, 2023, introduces the basics of high performance computing (HPC) and provides practical tricks as well as information on funding programs for HPC work. [Information and registration](#)

---

## USED THINGS FOR FURTHER SERVICE

The LRZ is always getting rid of used hardware and furniture - a constantly updated list of things we want to give away can be found [online](#). Here you can also read where to direct your interest. The equipment and furniture are free of charge for institutes, chairs and other research institutions.

---

## JOB OFFERS

You will find an international and diverse team in Garching, which is constantly growing. If you don't find a suitable job profile below, please visit the [career page](#) of the Leibniz Supercomputing Centre or send an [unsolicited application](#). We are LRZ - and curious about you!

[IT specialist](#) for Multi -Factor-Authentication

[Consultant and networker](#) for managed security services

[IComputer scientist](#) for AI and data driven applications

[Software engineer](#) HPC for the development of a quantum software stack

[DevOps engineer](#) for storage solutions

[Senior Systemengineer](#) for storage services

[IT specialist Systemintegration](#) for client-Mmagement Mac and mobile devices

[Technician](#) for maintenance of communication nets

[Scientific employee](#) für Simulationen in derGeodynamik

[Scientific employee](#) for system administration

[Scientific employee / IT-Engineer](#) for Windows-Server und collaborations

[Scientific employee](#) for the software governance and Trust & Identity Development

[CRM-Manager](#) in part time (20 hrs)

[Student assistant](#) PR und Content

[Student assistant](#) for ITSM-Development

[Student assistant](#) for the service desk

---

## MORE TO READ

Here you will find links to latest information from the german-european supercomputing community and our cooperation partners

- The [newsletter](#) of the Bavarian Academy for Science and Humanities
  - [Publications](#) of the Gauss Centre for Supercomputing (GCS): GCS-News und Inside
  - [Infoletters](#) of the Gauß-Alliance
  - Publications of PRACE: [PRACE Digest, Jahresbericht](#)
- 

## INFORMATIONS & IMPRINT

- The LRZ Newsletter is published in German and English. You can find the latest and former editions on the [LRZ-Website](#).
- You have problems displaying the newsletter? Please send a description of your problem to <pr-team\_AT\_lrz.de>. Thank You!
- You can subscribe or unsubscribe the LRZ-Newsletter via our [website](#).
- Current information about the LRZ and about courses and events can also be found on [Twitter](#) and [LinkedIn](#).

### Editor:

Leibniz-Supercomputing Centre (LRZ) of the Bavarian Academy of Science and Humanities  
Boltzmannstraße 1  
D-85748 Garching  
Phone: +49-89-35831-8000  
Telefax: +49-89-35831-9700  
E-Mail: pr-Team\_AT\_lrz.de; newsletter\_AT\_lrz.de  
Twitter: [LRZ DE](#); LinkedIn: [Leibniz-Rechenzentrum](#)

### Editorial Office: PR-Team

E-Mail: pr-team\_AT\_lrz.de; newsletter\_AT\_lrz.de  
Photos: Chair of Astronautics/ TUM, V. Hohenegger/LRZ, Wikimedia; Anastasia Zhenina/Unsplash