



Speaker & Session Booklet



27 January 2026



10:00–17:30 CET



Krakow, Poland

**Rethinking scientific applications for
exascale and emerging architectures**

the Centre of Excellence challenge

Rethinking scientific applications for exascale and emerging architectures

Abstract

The Centres of Excellence challenge

Centres of Excellence (**CoEs**) are initiatives supported by **EuroHPC Joint Undertaking**, aimed at developing a robust software ecosystem for High Performance Computing (HPC) and Quantum Computing (QC) across academia and industry. These efforts are designed to deliver tangible benefits, particularly in addressing scientific and societal challenges.

HPC and quantum computing are **rapidly evolving, disruptive technologies**. Harnessing their full potential demands significant adaptations or even complete rewrite of existing application codes. This presents a considerable challenge, as legacy algorithms must be re-engineered to perform efficiently on next-generation, heterogeneous supercomputing architectures.

This workshop **builds on the success** of “From petascale to exascale and beyond: the Centres of Excellence

challenge” and “Tackling software exascale challenges: the Centres of Excellence in High Performance Computing perspective” held at HiPEAC 2025, continuing the dialogue among Centres of Excellence (CoEs) in HPC and the European Quantum Excellence Centres.

It will bring together researchers from **different CoEs to share their experiences in modernising or developing algorithms for emerging architectures**. Key challenges and potential solutions in porting CoE applications to cutting-edge platforms will be discussed, fostering knowledge exchange within the community.

The event represents a remarkable opportunity to exchange best practices, explore innovative approaches, and strengthening collaboration across CoEs and the broader HPC community. To promote inclusivity, contributions from early-career researchers and scientists from underrepresented groups will be especially encouraged.



Eva Hernández
GEO3BCN-CSIC



ChEESE

About the author

Eva Hernández Plaza is a Computer Engineer at Geosciences Barcelona (GEO3BCN), a centre of the Spanish National Research Council (CSIC), and is currently pursuing a Master's-PhD track in Data Science. Her work focuses on scientific data management, FAIR digital assets, and workflow reproducibility in high-performance computing environments. She has contributed to the ChEESE Centre of Excellence, working on FAIR metadata design, CI/CD automation, repository structuring, and HPC deployment workflows associated with codes such as FALL3D.

She has also participated in the Geo-INQUIRE project, supporting large-scale dataset management, and collaborates with the ESA GET-IT project to implement satellite data assimilation into FALL3D. In addition, she has contributed to the DT-GEO project through the development of the volcanic ash dispersion Digital Twin. Her interests span HPC workflows, scientific data infrastructures, urgent computing, and reproducible research practices.

Data Storage and Availability for Scientific Workflows in HPC Systems

Data Storage and Availability for Scientific Workflows in HPC Systems
High-performance computing (HPC) applications increasingly rely on automated workflows, large simulation datasets, and interoperable software ecosystems. This talk presents the data-centric strategy adopted in the ChEESE-2P Centre of Excellence, where FAIR principles—Findable, Accessible, Interoperable, and Reusable—are embedded by





design to ensure sustainable and machine-actionable digital assets.

We describe the transition from a DOI-centric model in ChEESE-IP to a unified GitLab-based infrastructure featuring structured RDA-aligned metadata, CI/CD pipelines, and automated documentation workflows.

Complementary platforms such as the EOSC-Synergy SQAaaS environment and the Scientific Data Lake strengthen reproducibility, quality assurance, and long-term data stewardship.

The presentation also highlights current efforts in HPC deployment using Spack and EESSI, enabling portable, consistent, and optimized execution of codes like FALL3D across EuroHPC systems.

We describe the transition from a DOI-centric model in ChEESE-IP to a unified GitLab-based infrastructure featuring structured RDA-aligned metadata, CI/CD pipelines, and automated documentation workflows. Complementary platforms such as the EOSC-Synergy SQAaaS environment and the Scientific Data Lake strengthen reproducibility, quality assurance, and long-term data stewardship. The presentation also highlights current efforts in HPC deployment using Spack and EESSI, enabling portable, consistent, and optimized execution of codes like FALL3D across EuroHPC systems.

Together, these advances illustrate a modernized approach to managing data and workflows for scientific computing at scale.



Martin Kronbichler
Ruhr University Bochum



About the author

Martin Kronbichler is a Professor at Ruhr University Bochum, Germany. He holds a PhD degree in scientific computing with specialization in numerical analysis from Uppsala University, Sweden (2012).

His research interests include high-order finite element methods for flow problems with matrix-free implementations, efficient numerical linear algebra, and their parallel and high-performance implementation on emerging exascale hardware using generic numerical software.

As a PI in several exascale projects, he has been contributing to advances in the research field of mathematical software, with many important contributions made to the widely used finite-element library deal.II.

Towards Exascale Finite Element Solvers for the Human Body: A Perspective from the dealii-X CoE

The talk will present insights gained in the first year of the centre of excellence dealii-X, which aims to advance science in the simulation of the human brain, the lungs, the circulatory system, the liver as well as cellular processes.

We develop solvers in this variety of applications by a multi-layered software approach: We provide a generic framework for finite element computations through the open-source library deal.II, a project aimed at providing mathematical tools that assist the rapid development of application codes.



Several activities in our CoE are directed towards improving the infrastructure of the library in terms of coupling algorithms, advanced discretizations, multigrid methods and matrix-free implementations to optimally use compute capabilities in modern parallel computer.

Furthermore, we integrate advanced linear algebra packages for providing fast and robust algorithms.

Research groups specializing in each applications can then concentrate on the application details, such as the biomechanical processes, specific challenges and perspectives of increased simulation accuracy.



Helen Schottenhamml
IFP Energies Nouvelles



About the author

Helen Schottenhamml is a research engineer for scientific computing at IFP Energies Nouvelles in France. After her bachelor's degree in mechanical engineering and her master's degree in computational engineering, she started her PhD in computer science on the efficient simulation of wind farm flows using the lattice Boltzmann method.

This PhD, shared between FAU Erlangen in Germany and IFP Energies Nouvelles in France, deepened her passion for numerical simulations, software engineering, and high-performance computing, which remain part of her daily work. Helen's current research focuses on the simulation of wind farm flows, atmospheric dispersion and porous media.

Large-scale wind farm flow simulations within the Energy-oriented Center of Excellence

The Energy-oriented Center of Excellence (EoCoE) aims to improve the numerical simulation efficiency of applications in the low-carbon energy domain. With its four pillars in Fusion, Materials, Water, and Wind, EoCoE tackles a broad range of numerical solvers and shows the benefits of high-performance computing for the transition towards clean energy.

This presentation will focus on the lighthouse application waLberla-wind within EoCoE's "Wind for Energy" pillar. waLberla-wind was created during EoCoE and enables large-scale high-fidelity simulations of wind turbines in an atmospheric boundary layer using the lattice Boltzmann method.



We will introduce and motivate the underlying numerical and physical models, as well as waLBerla-wind's software design, which provides a portable code base for different hardware architectures. We will conclude the presentation with remarks on computational performance and performance engineering.



Juniper Tyree
University of Helsinki



About the author

Juniper Tyree (orcid:0000-0002-7923-9609) is a doctoral researcher in computational atmospheric sciences at the University of Helsinki, where she teaches a course on Numerical Meteorology. She holds an MEng in Computing from Imperial College London and an MSc in Theoretical and Computational Methods from the University of Helsinki. Her research focuses on developing novel computational methods to support and advance environmental research in areas such as ecology, atmospheric chemistry, and meteorology, while promoting and developing reproducible and open-science tools. In her PhD, which is funded by ESiWACE3, Juniper is working on making lossy compression safe and fearless to use for reducing the data volume produced by weather and climate models. As part of this work, she has been developing compression safeguards and an open-science online laboratory for easily trying out different compression methods. At the EGU26 conference in May 2026, Juniper is convening session ESS12.4 and short course SC2.5 on data compression in earth system sciences. Juniper also enjoys playing board games with friends and cooking with her family.

Exploring Lossy Data Compression in an Online Open Science Laboratory for Climate Science and Meteorology

Interactive code notebooks (e.g. Jupyter) have become popular for sharing and communicating computational experiments, analyses, and visualizations. While sharing the notebooks is easy, running them requires hosting a JupyterLab server and installing all Python and system libraries



required for the notebook. This initial setup cost hinders quickly experimenting with a shared notebook and testing, e.g. a practical example of lossy data compression for oneself.

As part of the EuroHPC ESiWACE, Phase 3, Centre of Excellence (<https://www.esiwace.eu/>), we have been developing an Online Laboratory for Climate Science and Meteorology (<https://lab.climet.eu/>), a JupyterLab instance that runs serverless just within your web browser and comes with many Python packages pre-installed. With the online lab, which builds on the Pyodide and JupyterLite community projects, running and exploring a shared notebook can start within a minute. We use the online laboratory to provide domain scientists with an online compression laboratory, <https://compression.lab.climet.eu>, to reduce the barrier to experimenting with the effect of lossy compression on their own data. The lab also supports URL schemas to preload other third-party notebooks (and repositories) hosted via Git, as Gists, or behind any URL, so that sharing a ready-to-run notebook is as easy as sharing, e.g., <https://lab.climet.eu/v0.3/github/juntyr/climet-lab-demo/v0.3.0/demo.ipynb>. We are also working on quickly turning existing static-documentation example-notebooks into interactive documentation that invites immediate further exploration.

In this presentation, we will showcase the online laboratory and the open science communication services it can provide to the HPC community by live demonstrating its applications in the compression laboratory and others. We also hope to gather feedback on the future direction of its development and collaborations with the broader HPC community.



Andrea Ferretti

CNR-NANO



About the author

Andrea Ferretti is a Research Director at the National Research Council of Italy (CNR), Institute of Nanoscience (CNR-NANO), where he works on electronic-structure simulations in condensed-matter and solid-state physics. He coordinates the High-Performance Computing activity of the institute and serves as its spokesperson. He is also Chair of the Executive Committee of the MaX Centre of Excellence, focused on materials design at the exascale.

Ferretti is a long-time developer and contributor of widely used scientific software for electronic-structure calculations, including YAMBO, Quantum ESPRESSO, and WanT. His research is centred on first-principles simulations of materials, from density-functional theory (DFT) to advanced approaches beyond DFT, including many-body perturbation theory and Koopmans-compliant functionals. His current interests include the electronic and optical properties of organic semiconductors, hybrid organic-inorganic interfaces, and nanoscale transport, with particular attention to correlation effects and theoretical spectroscopy. He is also active in methodological developments, with particular interest in the exploitation of dynamical potential in electronic structure theories.

Materials design towards the exascale: codes, workflows and data

Materials are crucial to science and technology, and connected to major societal challenges ranging from energy and environment, to information





and communication, and manufacturing. Electronic structure methods have become key to materials simulations, allowing scientists to study and design new materials before running actual experiments.

The MaX Centre of Excellence - Materials design at the eXascale - is devoted to enable materials modelling on exascale-class HPC machi

MaX's action focuses on popular open source community codes in the electronic structure field (Quantum ESPRESSO, Yambo, Siesta, Fleur, BigDFT). In this Talk I will discuss the main strategies and targets considered during the optimization and porting of MaX flagship codes, with particular emphasis on the support of GPU accelerated machines. I will also present the case of exascale workflows in materials science, in turn enabling the production of accurate curated data for selected materials properties.



Helena Vela Beltrán

Do IT Now

Multi  **scale**

About the author

Helena Vela Beltran is a Computational Scientist at Do IT Now (Spain), where she leads European projects in high-performance computing, AI/ML environments, and complex software stacks. She specializes in optimizing workflows, tuning HPC applications, and contributing to open-source projects, while also supporting researchers through training, documentation, and user engagement.

Status Update on EESSI: the European Environment for Scientific Software Installations

This talk gives an overview of the latest developments in the European Environment for Scientific Software Installations (EESSI). We'll highlight the newest release of EESSI and the expansion of CPU targets, including updated support for A64FX and RISC-V, alongside the growing catalogue of software projects available through the stack. Recent work has improved the experience of building on top of EESSI and strengthened compatibility with newer CUDA releases, while also looking into ROCm support to broaden GPU coverage.

We will touch on the ongoing effort to make Spack available within EESSI, as well as new integrations with the European Federation Platform and Open OnDemand to simplify adoption across HPC sites, along with the PoC integration with EOSC. The session also introduces the EESSI dashboard and provides a brief look at the community-focused activities launched this year, such as the EESSI happy hours and the expanding





webinar series, among other integrations and developments to the test suite and the bot.

Altogether, this update captures the steady progress of EESSI as it continues to grow into a flexible, multi-architecture software ecosystem for the HPC and research computing community.



Pratik Nayak
Technical University of
Munich



About the author

Pratik Nayak is a postdoctoral researcher within the Chair of Computational Mathematics at the Technical University of Munich. He obtained his PhD from the Infomatics department at Karlsruhe Institute of Technology, Germany.

His research lies at the intersection of numerical linear algebra, HPC and research software engineering. He contributes to many open-source software, and is one of the primary developers of Ginkgo, a high performance numerical linear algebra library.

Preparing legacy codes for large-scale computing: Experiences from different applications

Scientific applications often live much longer than the machines they are built for. As computing architectures evolve—sometimes radically—software that once performed flawlessly can struggle to keep up. The result is a familiar challenge: maintaining decades of scientific insight trapped inside code that no longer runs efficiently, or at all, on modern systems.

In this talk, we will explore how to break this cycle. Using examples of legacy scientific software, we will examine what makes some codes remarkably portable and sustainable, while others become obsolete with each new generation of hardware. We will discuss strategies for designing software that is not only high-performing today, but also flexible enough



to thrive on the heterogeneous, accelerated, and energy-aware architectures of tomorrow.

Ultimately, the goal is to show that portability and maintainability are not trade-offs against performance—they are the foundation for scientific progress that endures across generations of machines.



Luca Pennati

KTH Royal Institute of
Technology



About the author

Luca Pennati holds a Bachelor's degree in Engineering Physics and a Master's degree in Nuclear Engineering, both from Politecnico di Milano, and is currently a PhD student in High Performance Computing at KTH Royal Institute of Technology, under the supervision of Prof. Stefano Markidis and Prof. Ivy Peng.

His research focuses on plasma simulation codes for exascale systems, with an emphasis on algorithmic innovation, performance optimization, GPU acceleration, and performance portability. He is also working on advanced data analysis and data compression techniques, integrating AI and machine learning methods to enhance the efficiency and scalability of plasma simulations.

Physics-Aware Data Compression for Exascale Plasma Simulations

Adapting plasma physics applications to exascale systems requires rethinking not only their solvers, but also how we store, move, and interpret the massive amounts of data they produce. At these scales, traditional I/O and post-processing workflows become major bottlenecks, and the high volume of high-dimensional data demands new strategies. In this context, artificial intelligence and machine learning are emerging as key tools for automated feature detection, reduced-order modeling, and in-situ analysis that can extract scientific insight without overwhelming storage and bandwidth. Within the Plasma-PEPSC



European Centre of Excellence, we explore such approaches for large-scale particle-in-cell (PIC) simulations, for example through in-transit training of deep models and by leveraging machine-learning methods for data compression. In this contribution, we focus on physics-aware compression of velocity distribution functions. Instead of dumping raw particle data, we construct local velocity distributions and approximate them with compact statistical models based on Gaussian Mixture Models (GMMs).

This replaces billions of particles with a much smaller set of parameters that still encode the essential physical information needed to identify beams, bulk flows, and temperature anisotropies.

Our compression pipeline is designed to run in situ and to exploit accelerator resources, overlapping analysis with the core simulation workflow. This allows us to achieve compression factors of up to 104 relative to the original particle data, while keeping information loss at a level comparable to or better than state-of-the-art numerical compressors and preserving a representation that remains directly interpretable for plasma diagnostics. The resulting reduced data products are also well suited as inputs to downstream AI and machine learning techniques for tasks such as pattern recognition and anomaly detection.



Sandra Mendez
(BSC-CNS)

About the author

Sandra Mendez is a Research Associate at the Barcelona Supercomputing Center (BSC) in the Performance Tools group. She holds a PhD in High-Performance Computing (HPC) from the Universitat Autònoma de Barcelona (UAB), where she also teaches in the Department of Computer Architecture and Operating Systems. Her work focuses on the performance analysis and optimisation of parallel applications, with expertise in identifying performance bottlenecks and evaluating hybrid parallel codes on modern HPC systems.

Within the Performance Optimisation and Productivity (POP) Centre of Excellence, she serves as a performance analyst, conducting performance assessments of HPC applications. She is also the main developer of a BSC tool that computes POP efficiency metrics from Paraver traces, enabling more automated and accurate performance evaluations.

Performance Analysis of Hybrid Parallel Applications Using the POP Methodology

The Performance Optimisation and Productivity 3 (POP3) project, running from January 2024 to December 2026, is a European Centre of Excellence (CoE) in High-Performance Computing (HPC). Funded by the European High-Performance Computing Joint Undertaking under Horizon Europe, POP3 focuses on analysing and improving the performance and efficiency of parallel HPC applications.





POP3 is structured around three pillars—services, users, and co-design—and is delivered by a consortium of leading European research institutions and HPC centres with strong expertise in performance analysis, programming models, and application optimisation. Together, the partners provide mature methodologies and tools that enable both academic and industrial users to better understand and enhance the behaviour of their parallel codes.

The core POP service is the Performance Assessment, a first-level analysis designed to evaluate code behaviour and scaling, identify the principal causes of inefficiency, and deliver targeted optimisation recommendations

This talk will present insights gained from assessing parallel applications using the POP methodology. This methodology offers a structured, quantitative workflow for analysing application performance and identifying bottlenecks. It begins with a description of the application setup, code structure, input, and execution environment, followed by identification of the Focus of Analysis (FoA), which is selected based on the application structure and user priorities. The FoA may correspond to the initialization phase, a representative subset of iterations, or specific regions such as I/O or communication phases. The scalability of the FoA is then examined, and a hierarchy of POP efficiency metrics is computed to quantify the impact of different parallelisation factors. These efficiencies range from 0 to 1, with lower values—particularly below 0.8—indicating potential performance issues.



Based on the metric results, a detailed investigation is carried out using POP tools to analyse computation, communication, I/O, or memory-related bottlenecks. The methodology concludes with a set of actionable insights and optimisation suggestions, enabling users to understand the root causes of inefficiency and compare application performance across scales, architectures, and stages of optimisation. This approach is particularly relevant in the transition toward exascale and heterogeneous computing systems, where understanding the behaviour and scalability limits of hybrid parallel applications is essential for guiding application modernisation and ensuring they can fully exploit next-generation HPC architectures.



Ilaria Siloi
(University of Padua)



About the author

Ilaria Siloi is a computational physicist. She received her PhD in 2014, with a thesis on quantum correlations in molecular nanomagnets. Her interest in solid-state physics and real materials led her to the University of North Texas. The opportunity to combine her interest in quantum computing with problems in materials physics brought her to the University of Southern California.

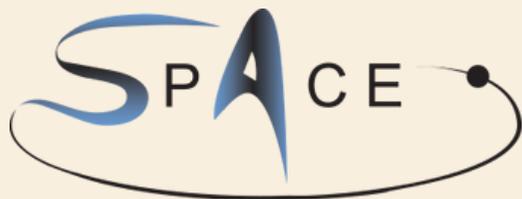
Since February 2022, she has been a staff researcher at the University of Padua, where she works on quantum computer emulation and quantum algorithms for solving optimization problems.

Early Directions on Quantum-Enhanced Applications within QEC4QEA

Quantum computing has reached a level of maturity where the development of concrete quantum-enhanced applications—spanning simulation, optimization, and machine learning—has become a central focus of current research efforts. This talk presents the scope and early activities of QEC4QEA, a newly established European Centre of Excellence focused on enabling quantum-enhanced applications through coordinated access to expertise, software, and infrastructure. QEC4QEA adopts an application-centric approach based on a selected set of quantum-enhanced applications, developed in a hardware-independent manner and supported by compilation and execution layers that enable their integration into hybrid HPC-quantum workflows across diverse



Tine Colman
(ENS Lyon, CNRS)



About the author

Tine Colman is a Belgian postdoctoral researcher at CNRS, based at ENS Lyon (France). She obtained her PhD at the University of Zürich under the supervision of Romain Teyssier, principal developer of the astrophysical simulation code RAMSES. She then held a postdoctoral position at CEA Paris-Saclay, where she worked on the ERC project ECOGAL, which aims to understand star formation in the Milky Way.

Her expertise lies in using high-performance computing and numerical simulations to address astrophysical questions, with a particular focus on star formation. With nearly ten years of experience working with RAMSES, she is currently involved in the SPACE Center of Excellence, where she works on the optimization, modernization and long-term maintenance of the RAMSES code.

The legacy astrophysics code RAMSES on the road to exascale

The Center of Excellence SPACE brings together scientists, computing centers and hardware vendors around several long-established astrophysical simulation codes, with the goal of preparing them and their communities for efficient use of current and upcoming exascale systems. In this talk, I will present ongoing work on RAMSES, one of the codes in the SPACE project. RAMSES is an adaptive mesh refinement (AMR) code written in Fortran and first released almost 25 years ago. It is used for applications ranging from cosmology to planet formation.



My focus has been on improving the time-to-solution on CPU architectures. Using profiling tools, we identified and addressed several memory access and vectorization bottlenecks. To improve scalability, we are introducing intra-node shared memory parallelism with OpenMP on top of the existing MPI implementation. In addition to the optimization work, I will present our efforts to modernize the development workflow with a more robust CI/CD, which is essential for the verification of our developments.

Porting a Fortran AMR code such as RAMSES to GPUs is challenging. I will outline the different approaches that are currently being explored within SPACE and the wider RAMSES community.



The Centres of Excellence

The Centres of Excellence (CoEs) are a core part of CASTIEL2, the European coordination and support action. Under CASTIEL2, these CoEs also collaborate with National Competence Centres (NCCs) to develop advanced HPC applications, share expertise and best practices, and drive innovation. Together, CASTIEL2, the CoEs, and the NCCs promote training, code development and cross-border collaboration to strengthen HPC capacity and foster the future of high-performance computing in Europe. Follow us on our digital channels and stay engaged with the latest developments in HPC.



hpc-portal.eu/



company/eurocc/



[@EuroCC2_EuroCC4SEE](https://www.youtube.com/@EuroCC2_EuroCC4SEE)



[@eurocc.bsky.social](https://www.bsky.social/@eurocc)



[Supercomputing in Europe](https://www.supercomputing.eu/)