

HPCSE 2026
High Performance Computing
in Science and Engineering

VSb TECHNICAL UNIVERSITY OF OSTRAVA | IT4INNOVATIONS NATIONAL SUPERCOMPUTING CENTER

HPCSE 2026

High Performance Computing in Science and Engineering

Collection of Abstracts

May 18-21, 2026
Hotel Soláň, Czech Republic
Organised by IT4Innovations National Supercomputing Center,
VSb - Technical University of Ostrava

Conference partners



Contributing sponsors



Supporting sponsors



Acknowledgements



Preface

Welcome to the High-Performance Computing in Science and Engineering 2026 conference (HPCSE 2026) organised by IT4Innovations National Supercomputing Center, VSB – Technical University of Ostrava. We aim to bring together international specialists in high-performance computing, applied mathematics, numerical analysis, high-performance data analytics, machine/deep learning, quantum computing, and advanced visualisation, exchange experience and ideas and initiate new research collaborations.

We are honoured to host distinguished experts from leading research institutions and supercomputing centres worldwide. We hope this conference inspires and motivates new students and researchers to join the growing and vital supercomputing community.

Set in the heart of the Beskydy Mountains, Soláň offers a tranquil atmosphere and breathtaking views, creating the perfect backdrop for our event. We thank all contributors and scientific and organising committee members for their tireless efforts to make this conference successful. Your dedication ensures that the HPC community thrives and expands its impact on science and technology. We also acknowledge the partial support of the EuroHPC Joint Undertaking EuroCC3 project. We would also like to sincerely thank our sponsors for their generous support, which helps make HPCSE 2026 possible.

Thank you for joining us. We look forward to engaging in discussions and insightful discoveries during this gathering.

On behalf of the Scientific and Organising Committee of HPCSE 2026,

Prof. Tomáš Kozubek

About IT4Innovations National Supercomputing Center

IT4Innovations National Supercomputing Center at VSB – Technical University of Ostrava is a leading research, development, and innovation centre active in the fields of High-Performance Computing (HPC), Data Analysis (HPDA), Artificial Intelligence (AI), and Quantum Computing (QC) and their application to other scientific fields, industry, and society.

IT4Innovations operates the most powerful supercomputing systems in the Czech Republic, provided to Czech and foreign research teams from academia and industry. Together with the CESNET and CERIT-SC institutions, IT4Innovations constitutes e-INFRA CZ, a strategic research infrastructure of the Czech Republic.

IT4Innovations runs two supercomputers: Karolina (15.7 PFlop/s, installed in the summer of 2021), and Barbora (849 TFlop/s, installed in the autumn of 2019). Apart from the supercomputers operated at IT4Innovations, Czech research communities also have access to the LUMI supercomputer thanks to IT4Innovations' membership in the LUMI consortium. Within its data room, IT4Innovations also operates complementary systems composed of several different hardware platforms. These complementary systems provide users with access to emerging, non-traditional, and highly specialised hardware architectures. In 2025, the first Czech quantum computer, named VLQ, operated by the LUMI-Q consortium, was installed at IT4Innovations.

IT4Innovations is also actively involved in the European network of AI Factories – it is a member of the LUMI AI Factory and coordinates the Czech AI Factory. The primary goal of AI Factories is to support the development of AI startups and the deployment of artificial intelligence in industry, public administration, and society at large. AI Factories aim to achieve this goal by providing expertise in the field of artificial intelligence, supporting the development of new AI models, and providing software tools for the easy use of supercomputers and the processing of large datasets.

IT4Innovations is dedicated to research excellence in High Performance Computing (HPC), Data Analytics (HPDA), Artificial Intelligence (AI) and Quantum Computing (QC). IT4Innovations' core research areas are big data processing and analysis, machine learning, development of parallel scalable algorithms and algorithms for quantum computers and simulators, solving challenging engineering problems, advanced visualisation, virtual reality, modelling for nanotechnology, and material design.

IT4Innovations research activities are executed across 5 laboratories:

- Advanced Data Analysis and Simulations Lab,
- Infrastructure Research Lab,
- Parallel Algorithms Research Lab,

- Modelling for Nanotechnologies Lab,
- Quantum Computing Lab.

IT4Innovations actively cooperates with the industry. Experts from the European Digital Innovation Hub Ostrava (EDIH Ostrava) and the National Competence Centre in HPC are available to assist, especially SMEs and startups, by providing a wide range of services in supercomputing and in the adoption and application of digital technologies.

IT4Innovations contributes significantly to all activities of the EuroHPC Joint Undertaking and is a member of key European infrastructures, initiatives and associations:

BDVA (Big Data Value Association),
 EOSC (European Open Science Cloud),
 ETP4HPC (European Technology Platform in the area of High-Performance Computing),
 EUDAT Collaborative Data Infrastructure,
 iRODS (Integrated Rule-Oriented Data System),
 VI-HPS (Virtual Institute – High Productivity Supercomputing),
 PRACE (Partnership for Advanced Computing in Europe),
 WHPC (Central European Chapter of Women in High Performance Computing),
 QUIC (European Quantum Industry Consortium).

IT4Innovations has long organised and offered a wide array of specialised training events focused on HPC, HPDA, AI, and QC, and on acquiring skills necessary for the effective use of supercomputing infrastructure. It also focuses on developing digital skills. In these areas, IT4Innovations is actively involved in educating students, particularly through study programmes at VSB – Technical University of Ostrava. IT4Innovations is also a member of the EUMaster4HPC consortium, which implements the eponymous pan-European MSc study programme focused on high-performance computing.

About the National Competence Center in HPC Czechia

The National Competence Centre in HPC Czechia (NCC Czechia) is the Czech Republic's reference and single point of contact and coordination for high-performance computing (HPC). It is represented by the IT4Innovations National Supercomputing Center.

On the national level, the mission of the NCC is to analyse, implement, and coordinate all necessary activities and offer comprehensive, high-performance computing services to end-users from industry, public institutions and academia. It includes, for example, expert technology consultancy tailor-made to the needs of interested parties, also guidance through the whole process of the proof of concept of the case, access to supercomputers themselves, assistance with participation in calls for proposals to gain computing time, providing training services at levels from beginners to advanced, in-depth courses and last but not least disseminate information about applications possibilities of supercomputers to the general public.

Since its beginning of operation, the NCC Czechia has achieved a number of success stories and whitepapers with industry and the public sector in a variety of fields, including healthcare, water management, architecture, biotechnology, environment, mechanical engineering, transport, energy industry, security, steel-making industry, agriculture, and more.

More information can be found on the official website www.eurocc-czechia.cz/en



Honouring Prof. Zdeněk Dostál on His 80th Birthday

This year's edition of HPCSE is dedicated to the **80th birthday of Prof. Zdeněk Dostál** in recognition of his outstanding and long-standing contributions to applied mathematics and high-performance computing.

Throughout his distinguished career, Zdeněk Dostál has been closely connected with VSB – Technical University of Ostrava, particularly the Department of Applied Mathematics and the IT4Innovations National Supercomputing Center. He has authored more than 230 scientific publications, served as a visiting professor at several international institutions, and written two monographs presenting his work on optimal algorithms for quadratic programming.

Through his research, international collaborations, and dedication to mentoring young scientists, he has significantly contributed to the development of a strong research community at the intersection of mathematical modelling, algorithm design, and large-scale scientific computing. We are honoured to celebrate this important milestone during HPCSE and to acknowledge Prof. Dostál's lasting impact on generations of students and researchers tackling complex computational challenges.



Scientific and Organising Committees

Scientific Committee

Jaroslav Hron	Charles University, Prague
Jiří Jaroš	Brno University of Technology
Tomáš Kozubek	IT4Innovations, VSB-TUO
Jaroslav Kruis	Czech Technical University in Prague
Dalibor Lukáš	VSB – Technical University of Ostrava
Jan Martinovič	IT4Innovations, VSB-TUO
Tomáš Oberhuber	Czech Technical University in Prague
Jakub Šístek	Institute of Mathematics of the CAS
Petr Tichý	Charles University, Prague
Vít Vondrák	IT4Innovations, VSB-TUO

Organising Committee

Tomáš Kozubek	IT4Innovations, VSB-TUO
Karina Pešatová	IT4Innovations, VSB-TUO
Lucie Rožnovská	IT4Innovations, VSB-TUO
Veronika Pakostová	IT4Innovations, VSB-TUO

List of Speakers

Invited Speakers

Erin Claire Carson	Charles University, CZ
Michael Granitzer	University Passau, DE
Rolf Krause	King Abdullah University of Science and Technology, SAU
Ulrich Langer	Johannes Kepler Universität Linz, AT
Michael McGettrick	University of Galway, IE
Jakub Mielczarek	Jagiellonian University, PL
Matthias Möller	Delf University of Technology, NL
Matthias Obst	University of Gothenburg, SE
Ulrich Ruede	Friedrich-Alexander-Universität Erlangen-Nürnberg, DE
Martin Schulz	Technical University of Munich, DE
Kirk M. Soodhalter	University of Dublin, IE
Olaf Steinbach	Technical University Graz, AT
Daniel Sýkora	Czech Technical University in Prague, CZ

Contributing Sponsors Speakers

Lukáš Vach	M Computers, CZ
------------	-----------------

Regular Speakers

Joao Barbosa	IT4Innovations, VSB-TUO
Thomas Bake	Charles University
Simona Béréšová	Institute of Geonics, Czech Academy of Sciences
Ada Böhm	IT4Innovations, VSB-TUO
Tomas Hammerbauer	Charles University
Jakub Homola	IT4Innovations, VSB-TUO
Jakub Kružík	Institute of Geonics, Czech Academy of Sciences
Dalibor Lukáš	VSB-Technical University of Ostrava
Tomáš Oberhuber	Czech Technical University in Prague
Dominika Regéciová	IT4Innovations, VSB-TUO
Petra Svobodová	IT4Innovations, VSB-TUO
Jakub Šístek	Institute of Mathematics of the Czech Academy of Sciences
Petr Tichý	Faculty of Mathematics and Physics, Charles University
Jiří Tomčala	IT4Innovations, VSB-TUO
Štěpán Vaněček	TU Munich
Filip Vaverka	IT4Innovations, VSB-TUO
Van Binh Henri Vu	IT4Innovations, VSB-TUO

Conference Programme

Monday, 18 May

10:00 - 12:00 [Registration, discussion, refreshment](#)

12:00 - 13:45 [Lunch](#)

13:45 - 14:00 **Tomáš Kozubek:** Opening

14:00 - 14:30 **Vít Vondrák:** News from IT4Innovations and the Czech National Competence Center

14:30 - 16:00 Chairman: T. Kozubek

Ulrich Langer: Goal-oriented adaptive space-time finite element solvers for regularized parabolic p-Laplace problems

Kirk M. Soodhalter: Extensions and applications of admissible and attainable convergence theory for GMRES

16:00 - 16:30 [Coffee break](#)

16:30 - 18:00 Chairman: J. Jaroš

Daniel Sýkora: The Devil in the detAll: Breaking Through GenAI's Quality Ceiling

Dominika Regéciová: Best Practices: The Best Time To Rotate Your Credentials Is Now

Petra Svobodová: Multi-Camera Vision System for Operator Monitoring and Scene Understanding in Isolators

Simona Běrešová: Gradient-Informed Proposals for Surrogate-Assisted Bayesian Sampling

19:00 [Welcome Party](#)

Tuesday, 19 May

07:00 - 09:00 [Breakfast](#)

09:00 - 10:30 Chairman: T. Oberhuber

Ulrich Ruede: On the Parallel Solution of Large Scale Finite Element Problems

Jakub Šístek: Speeding up an unsteady flow simulation by adaptive BDDC and Krylov subspace recycling

Jakub Homola: GPU acceleration of a FETI solver

Tomas Hammerbauer: Scalable two-level Schwarz preconditioners for discontinuous Galerkin of PDEs

10:30 - 11:00 [Coffee break](#)

11:00 - 12:30 Chairman: J. Šístek

Petr Tichý: On the Dubrulle variant of the block CG method

Thomas Bake: Forward and backward error bounds for a mixed precision preconditioned conjugate gradient method

Jakub Kružík: Accelerating MPPG-type Methods Through Preconditioning

Dalibor Lukáš: Operator Preconditioning in Boundary Element Methods on GPUs

Tomáš Oberhuber: TNL: Numerical library for modern parallel architectures

Lukáš Vach, M Computers: Hardware 2026: From Core Wars Intel/AMD to Seagate Storage Revolution

12:30 - 14:00 [Lunch](#)

14:00 - 19:00 [Free afternoon](#)

19:00 - 20:00 [Dinner](#)

20:00 [Poster session](#)

Wednesday, 20 May

07:00 - 09:00 [Breakfast](#)

09:00 - 10:30 Chairman: J. Martinovič

Matthias Obst: Digital Twin applications for dynamic assessments of marine biodiversity

Michael Granitzer: OpenWebSearch.eu - Building an Open Index of the Web on HPC Infrastructure

10:30 - 11:00 [Coffee break](#)

11:00 - 12:30 Chairman: M. Lampart

Matthias Möller: Is the future of HPCSE quantum?

Jakub Mielczarek: How far are we from practical quantum cryptanalysis?

12:30 - 14:00 [Lunch](#)

14:00 - 15:30 Chairman: M. Lampart

Michael McGettrick: Quantum Computing via discrete quantum walks and quantum games

Jiří Tomčala: Quantum Optimization of a Community Energy Network Using a Simulator and a Real Quantum Computer

Van Binh Henri: From Smooth Valleys to Barren Plateaus: The Hidden Geometry of Quantum Optimization

Ivana Miháliková: Using symmetries to prepare ground states on quantum computers

15:30 - 16:00 [Coffee break](#)

16:00 - 18:15 Chairman: D. Lukáš

Rolf Krause: Decomposition Methods for Coupled and Strongly Non-linear Problems in HPC

Olaf Steinbach: Space-time tensor-product finite element methods for parabolic problems

Special session for Zdeněk Dostál

19:00 [Conference Dinner](#)

Thursday, 20 May

07:00 - 09:00 **Breakfast**

09:00 - 10:30 Chairman: P. Tichý

Martin Schulz: From HPC to HPCQC: Software Foundations for Hybrid Classical–Quantum Computing

Erin Claire Carson: Mixed-precision Computing: High Accuracy with Low Precision

10:30 - 11:00 **Coffee break**

11:00 - 12:15 Chairman: L. Říha

Ada Böhm: HyperQueue: Scaling to Millions of Tasks

Roman Iakymchuk: Energy-aware CFD simulations: Sharing Experience from the CEEC project

Filip Vaverka: MaX: Co-design, technology exploitation & energy efficiency

Stepan Vanecek: Understanding Heterogeneous Systems Through Topology-Aware Tools

Joao Barbosa: Supercharging Your HPC Applications: Free Expert Support through EPICURE and POP-3

12:15 - 12:30 **Tomáš Kozubek:** Closing remarks

12:30 - 14:00 **Lunch**

Abstracts – Invited Talks

Monday, 18 May

Ulrich Langer - Goal-oriented adaptive space-time finite element solvers for regularized parabolic p-Laplace problems

We consider goal-oriented adaptive space-time finite-element discretizations of the regularized parabolic p-Laplace problem on completely unstructured simplicial space-time meshes. The adaptivity is driven by the dual-weighted residual (DWR) method since we are interested in an accurate computation of some possibly nonlinear functionals at the solution. Such functionals represent goals in which engineers are often more interested than in the solution itself. The DWR method requires the numerical solution of a linear adjoint problem that provides the sensitivities for the mesh refinement. This can be done by means of the same full space-time finite element discretization as used for the primal non-linear problems. The numerical experiments performed on the HPC cluster Radon1 at Linz demonstrate that this goal-oriented, space-time finite element solver efficiently provides accurate numerical results for different functionals [1]. We also refer to the overview paper [2]. Furthermore, we apply the space-time solver to the KKT system derived from the optimal control of some simplified Carreau flow [3]. References [1] B. Endtmayer, U. Langer, A. Schafelner. Goal-oriented adaptive space-time finite element methods for regularized parabolic p-Laplace problems, *Comput. Math. Appl.* 167 (2025), 286–297. [2] B. Endtmayer, U. Langer, T. Richter, A. Schafelner, T. Wick. A posteriori single- and multi-goal error control and adaptivity for partial differential equation, In F. Chouly, S.P.A. Bordas, R. Becker, and P. Omnes, editors, *Error Control, Adaptive Discretizations, and Applications, Part 2*, volume 59 of *Advances in Applied Mechanics*, pp. 19–108. Elsevier, 2024. [3] S. Beuchler, B. Endtmayer, U. Langer, A. Schafelner, T. Wick. A 5d concept for space-time optimal control problems with application to simplified Carreau flow, *arXiv:2511.09086* (2025).

Kirk M Soodhalter - Extensions and applications of admissible and attainable convergence theory for GMRES

It has long been well understood that the residual convergence behavior of GMRES need not have any connection to the eigenvalue distribution [Greenbaum, Pták, Strakoš; 1996]. The proof of this fact is constructive; pathological matrix/right-hand-side pairs exhibiting this phenomenon are built, demonstrating the result. This work has since been extended to allow for almost-arbitrary assignment of Ritz values, specification of residual convergence behavior for restarted GMRES and block GMRES, etc. This theory has at times been presented from an overly pessimistic point-of-view with regard to our ability to robustly analyze convergence of GMRES. In fact, it has the potential to be a robust framework with which to understand the mechanics of GMRES behaviour. In this talk, we present some recent extensions to this theory

and discuss how one can use it to analyse the convergence of GMRES applied to highly-structured linear systems, e.g., systems with Toeplitz structure.

Daniel Sýkora - The Devil in the detAll: Breaking Through GenAI's Quality Ceiling

Over the past few years, we have witnessed immense progress in the field of generative artificial intelligence. The quality of image and video generation has improved so significantly that outputs often appear indistinguishable from reality. While each new model demo seems more impressive than the last, our overhyped internal beliefs are often shaped by a set of cherry-picked results. In reality, performance frequently falls short of the polished social media advertisements produced by major GenAI players. In this talk, we examine current GenAI capabilities in two interesting domains: the reproduction of artistic styles and the rendering of nuanced facial expressions. We will examine how the infamous 'six-fingered' hallucination problem manifests in these new contexts, discuss the underlying sources of discrepancy, and propose potential strategies to overcome them.

Tuesday, 19 May

Ulrich Ruede - On the Parallel Solution of Large Scale Finite Element Problems

This talk presents research on massively parallel solution algorithms for elliptic PDE. Fast iterative algorithms can compute the solution in asymptotically optimal complexity and also with a surprisingly low operation count for the standard problems. However it remains a challenge to design such algorithms for more complex applications and to implement them on modern large-scale parallel computers. Here we will present the HyTeG multigrid framework that implements a geometric multigrid algorithm on hybrid tetrahedral grids. We demonstrate that it is possible to achieve scalability on up to several hundred thousand processors and then solve linear systems with up to 10^{13} (10 trillion) unknowns in a solution time of a few minutes.

Wednesday, 20 May

Matthias Obst - Digital Twin applications for dynamic assessments of marine biodiversity

The triple planetary crisis that arises from the combined effects of climate change, biodiversity loss, and pollution has profound impacts on ocean ecosystems and human health. Today researchers know the forces driving environmental degradation but have little insight into the complex interplay between human pressures and local ecosystems. Such knowledge however is essential for any serious attempt to reverse environmental destruction. To achieve such insight continuous ecosystem monitoring becomes imperative. Over the last decade, many new monitoring technologies including genomics, hydroacoustics, telemetry, and image-based surveillance have become available to marine biologists. The transformation to such new high throughput methods, however is not trivial as they produce enormous volumes

of data. The European Commission's Digital Twin of the Ocean (DTO) programme addresses these challenges and develops a digital environment that brings together data and analytical services supporting effective monitoring, restoration, and management of marine ecosystems by 2030. This talk will introduce some important concepts behind the DTO for marine biodiversity and highlight the critical role of high-performance computation infrastructures for the DTO. The presentation will explain with several examples how digital science and engineering can contribute to an otherwise highly fragmented research landscape and facilitate complex sensor-to-modelling systems. The talk also emphasises the importance of the human role in connecting HPC infrastructures to new scientific disciplines such as data-driven ecology and how this connection can boost innovation and ecological sustainability.

Michael Granitzer - OpenWebSearch.eu: Building an Open Index of the Web on HPC Infrastructure

Web search has become a critical infrastructure for the digital economy, knowledge access, and increasingly for generative AI systems. Yet today, this infrastructure is largely controlled by a small number of gatekeepers. Their dominance shapes how information is discovered, ranked, and reused, creating biased and economically driven access to the Web. This not only affects users and content providers, but also limits Europe's ability to build sovereign, transparent, and trustworthy AI technologies.

This talk presents the OpenWebSearch.eu initiative, which aims to develop an Open Web Index as a shared European infrastructure for Web search, AI, and future agentic systems. Rather than treating search merely as an application, OpenWebSearch.eu approaches it as a large-scale data and compute infrastructure challenge: crawling, processing, indexing, enriching, and serving Web-scale data requires advanced distributed systems, high-performance computing, scalable storage, and sustainable data pipelines.

A particular focus will be placed on the role of HPC infrastructures in enabling an Open Web Index. We will discuss how large-scale crawling and indexing workflows can benefit from European supercomputing and cloud infrastructures, and how such an index can serve as a foundational resource for generative AI, retrieval-augmented generation, autonomous agents, and other data-intensive AI applications.

Matthias Möller - Is the future of HPCSE quantum?

For decades, high-performance computing in science and engineering has relied on Moore's law to scale computing power. However, this approach is slowly reaching its limits. The question is, whether and for which type of problems novel compute paradigms like quantum computing can provide a more efficient and scalable alternative.

After a brief introduction of the working principle of quantum computers, we will discuss an end-to-end quantum algorithm for solving fluid dynamics problems using the lattice Boltzmann method. I will present encoding schemes that enable an exponential compression

of grid data into quantum bits making it possible to solve fluid flow problems at unparalleled resolutions. I will also discuss building blocks for implementing streaming and collision as quantum circuits and address the challenges of tackling non-linear operations on quantum computers. If time permits, I will address a second use case: quantum physics-informed neural networks (QPINNs) with trainable embeddings. QPINNs are a natural extension of classical PINNs to quantum computers, where the training process is realized through learning the rotation angles of a parametrized quantum circuit. Derivatives can be computed exactly by the so-called parameter-shift rule. The specialty of our approach is the use of classical feed-forward neural networks to learn the mapping of classical data into the rotation angles of the QPINNs' embedding layer. I will present numerical results for Burgers' equation, Stokes flow and linear elasticity.

Jakub Mielczarek - How far are we from practical quantum cryptanalysis?

Despite the growing concern that quantum computing may undermine modern cryptography, the practical severity and timeline of this threat remain uncertain. This talk surveys the current state of quantum cryptanalysis and its implications for symmetric encryption, asymmetric algorithms, hash functions, and authentication protocols. It explains how quantum superposition enables attack models that can provide polynomial or even exponential speedups over classical cryptanalytic techniques. Beyond theoretical advances, the presentation examines the practical limitations of contemporary quantum hardware. Key resource metrics, including quantum volume, are analyzed to evaluate the feasibility of large-scale quantum attacks against cryptographic primitives in current use. Current projections indicate that, under conservative assumptions about hardware progress, practically relevant quantum cryptanalysis is likely to remain infeasible for several decades, although unforeseen scientific and technological breakthroughs could substantially shorten this timeline.

Michael McGettrick - Quantum Computing via discrete quantum walks and quantum games

Our talk will present the intrinsic properties of quantum algorithms that make them different from classical algorithms. Progress in quantum computing is made in two ways:

- 1) Applying (versions of) existing algorithms to new/different problems. Here, new domains in which HPC is important can adopt quantum algorithms, often for parts of the computation (semi-classical). Progress is often made by re-writing a problem using different representation/notation.
- 2) Discovering/constructing fundamentally new algorithms. Progress here has been poor (for decades).

The most widely understood quantum computation model (the gate or circuit model) is computationally equivalent to the Discrete Time Quantum Walk (DTQW) model. We describe basic DTQWs, and more specifically alternating walks and walks with memory, focusing on unusual properties such as ballistic spread and localization. While DTQWs utilize the quantum phenomena of superposition and interference, quantum games utilize entanglement to achieve outcomes that are not achievable classically. We explain how changing a classical

game to a quantum one leads to new strategies enabling better payoff for both parties (This is the basic model behind evolutionary quantum games, AKA quantum agents.)

Rolf Krause - Decomposition Methods for Coupled and Strongly Non-linear Problems in HPC

Large scale problems, as occurring in contact mechanics or fluid structure interaction, pose certain challenges for the solution of the arising large-scale non-linear systems of equations. We will discuss parallel solution methods for this type of non-linear and possibly non-smooth problems and will explain how domain decomposition can be designed in the non-linear case. Particular focus will be on global convergence and control of the non-linear iteration process. In addition to the solution of (coupled) PDEs, we will also consider the training of neural networks, and will show that decomposition approaches can be used also in the context of machine learning.

Olaf Steinbach - Space-time tensor-product finite element methods for parabolic problems

In this talk we will review space-time variational formulations for parabolic evolution equations. In the particular case of space-time tensor-product meshes we can use test and ansatz spaces of different regularity in time. We will provide a detailed stability and error analysis, with a particular focus on initial data with low regularity. Numerical results confirm the theoretical findings. The talk is based on joined work with Richard Löscher (Graz) and Michael Reichelt (Hamburg).

Thursday, 21 May

Martin Schulz - From HPC to HPCQC: Software Foundations for Hybrid Classical–Quantum Computing

The evolution of quantum computing is reshaping the landscape of high-performance computing, where quantum processors are increasingly viewed as accelerators complementing classical architectures. Successfully transitioning from HPC to hybrid HPCQC systems requires software solutions that bridge established HPC workflows with rapidly evolving quantum technologies. Within the Munich Quantum Valley (MQV), we are addressing this challenge through the open-source Munich Quantum Software Stack (MQSS) and the Quantum Device Management Interface (QDMI), a vendor-independent interface for quantum hardware integration. This talk presents our vision and practical approach to building a sustainable software foundation for HPCQC, discusses key integration challenges and lessons learned, and outlines the current status and future potential of hybrid classical–quantum computing infrastructures.

Erin Claire Carson - Mixed-precision Computing: High Accuracy with Low Precision

Mixed-precision algorithms have launched an era in which efficiency and accuracy are no longer mutually exclusive. Rather than rely entirely on high-precision formats like double (64-bit) precision, mixed-precision algorithms apply lower precisions such as single (32-bit) or half (16-bit) precision whenever possible, reserving higher precision only for critical steps. Doing so can drastically reduce memory requirements, improve performance, and lessen energy consumption on modern computer hardware without sacrificing accuracy or stability. In this talk, we discuss the challenges of using low/mixed precision, and present five cases, common in scientific applications, where using mixed precision makes sense.

Abstract – Contributing Sponsors Talk

Tuesday, 19 May

Lukáš Vach, M Computers - Hardware 2026: From Core Wars Intel/AMD to Seagate Storage Revolution

This presentation explores the current landscape of high-performance hardware, highlighting the latest CPU breakthroughs from Intel and AMD. It also analyses shifting price trends in the RAM and SSD markets and introduces cutting-edge HDD innovations from Seagate.

Abstracts – Regular Talks

Monday, 18 May

Dominika Regéciová - Best Practices: The Best Time To Rotate Your Credentials Is Now

In recent weeks, the global research community has witnessed a surge in supply-chain attacks that abuse trusted, often fully automated, software distribution and deployment processes to install backdoors on victim systems. In many cases, these attacks required no user interaction, enabling adversaries to silently exfiltrate sensitive credentials such as SSH private keys and API tokens. The academic sector is particularly exposed to such threats. Its decentralized infrastructure, open collaboration model, and extensive online footprint make research institutions and their members attractive targets for threat actors. This risk is further amplified by rapid advances in large language models (LLMs), which are increasingly improving their capabilities to discover vulnerabilities and zero-day attacks. High-performance computing (HPC) environments have also long been lucrative targets due to their valuable data and computational resources. In the past, there were directed attacks on HPC providers, but nowadays it is simple to target individual researchers and developers to

gain indirect access to these systems by stealing personal credentials. For these reasons, the security awareness is more important than ever before. It is not enough to think before clicking on links, even widely trusted and reputable projects can be compromised and used as vectors for attack. This presentation discusses practical steps that both the research community as a whole and individual users can take to reduce risk and improve resilience against modern supply-chain threats.

Petra Svobodová - Multi-Camera Vision System for Operator Monitoring and Scene Understanding in Isolators

This talk presents ongoing work on a computer vision system for monitoring operator activities in an isolator environment used in gene therapy manufacturing, where manual operations play a key role in process quality and safety. The focus is on visual scene understanding using multiple complementary tasks, including segmentation of laboratory objects, detection and classification of 2D codes (QR and Data Matrix), and hand segmentation for capturing operator interactions. The individual components are developed and evaluated separately using a custom dataset reflecting real operating conditions. We present preliminary results from the individual modules, including the isolator setup, camera configuration, and the design of segmentation and detection modules.

Simona Béréšová - Gradient-Informed Proposals for Surrogate-Assisted Bayesian Sampling

This contribution presents a surrogate-assisted sampling framework for Bayesian inverse problems based on the delayed-acceptance Metropolis-Hastings algorithm. A computationally cheap neural network surrogate model is built on the fly from evaluations of the exact forward solver and shared by several chains running in parallel. The surrogate model allows long proposal subchains at low cost. It also provides approximate gradients that can be used for longer informed moves via Hamiltonian proposals. For priors induced by Gaussian random fields, pCN proposals offer another dimension-robust alternative. The methodology is motivated by geological inverse problems with heterogeneous materials.

Tuesday, 19 May

Jakub Šístek - Speeding up an unsteady flow simulation by adaptive BDDC and Krylov subspace recycling

We deal with accelerating the solution of a sequence of large linear systems solved by preconditioned conjugate gradient method (PCG). The sequence originates from time-stepping within a simulation of an unsteady incompressible flow. We apply a pressure correction scheme and focus on the solution of the Poisson problem for the pressure corrector. Its scalable solution presents the main computational challenge in many applications. The right-hand side of the problem changes in each time step, while the system matrix is constant and symmetric positive definite. The acceleration techniques are studied on a representative problem of flow around a unit sphere. Our baseline approach is based on a parallel solution of each problem in the sequence by a nonoverlapping domain decomposition method. The interface problem is solved by PCG with the three-level BDDC

preconditioner. As a preliminary step, an appropriate stopping criterion for the PCG iterations is chosen. Next, two techniques for accelerating the solution are gradually added to the baseline approach. Deflation is used within PCG with several approaches to Krylov subspace recycling. Finally, we add the adaptive selection of the coarse space within the three-level BDDC method. We present a number of experiments with careful measurements of computational times on the Karolina supercomputer. The combination of the acceleration techniques eventually leads to saving more than 40 % of the computational time. This presentation is based on joint work with Martin Hanek and Jan Papež and the paper [M. Hanek, J. Papež, and J. Šístek. Speeding up an unsteady flow simulation by adaptive BDDC and Krylov subspace recycling. *Computer Methods in Applied Mechanics and Engineering*, 452:118788, 2026].

Jakub Homola - GPU acceleration of a FETI solver

A wide range of engineering problems can be solved using the Finite Element Method (FEM). Finite Element Tearing and Interconnecting (FETI) extends this approach by utilizing domain decomposition techniques, enabling efficient execution on distributed-memory HPC clusters. Nowadays, as the HPC world shifts toward GPU-accelerated supercomputers, we are motivated to extend the solver to leverage the power of GPUs as well. In this talk, we present our approach to GPU acceleration of our FETI solver. A straightforward offloading of computations to the GPU proved insufficient; instead, we adopt a different approach that increases preprocessing costs but reduces iteration time, with the overall benefit depending on the number of iterations performed. We focus on the efficient assembly and application of Schur complement matrices on the GPU and further improve performance by exploiting sparsity of the matrices. We present results across all three major GPU vendors -- Nvidia, AMD and Intel.

Tomas Hammerbauer - Scalable two-level Schwarz preconditioners for discontinuous Galerkin of PDEs

We present and analyze a class of two-level Schwarz preconditioners for discontinuous Galerkin discretizations of partial differential equations. The proposed domain decomposition preconditioners are based on a non-overlapping partition of the computational domain into subdomains. They combine local solvers on the subdomains with global coarse solvers to ensure scalability and efficiency. We show the convergence properties of the preconditioned iterative method, and provide numerical experiments to demonstrate the efficiency and scalability of the proposed approach. Results indicate that our preconditioners significantly reduce the number of iterations required for convergence of the iterative method, making them a promising tool for solving large-scale partial differential equations discretized with discontinuous Galerkin methods.

Petr Tichý - On the Dubrulle variant of the block CG method

In this talk, we analyze a variant of the block conjugate gradient (CG) method introduced by A. Dubrulle, which we refer to as DR-BCG. For the solution of linear systems with a symmetric positive definite matrix and multiple right-hand sides, this variant avoids difficulties caused by possible rank deficiency within the block vectors. We present results showing that DR-BCG is a breakdown-free algorithm, establish its fundamental theoretical properties, and prove global convergence. Numerical experiments demonstrate that DR-BCG consistently outperforms other block CG variants, including deflation-based approaches. Based on these results and experience, we consider DR-BCG to be a natural block generalization of the standard CG algorithm. The DR-BCG algorithm is simple, behaves well numerically, and exploits the full potential of block methods in both theoretical aspects (faster convergence) and practical (use of block operations).

Thomas Bake - Forward and backward error bounds for a mixed precision preconditioned conjugate gradient method

The preconditioned conjugate gradient (PCG) algorithm is one of the most popular algorithms for solving large-scale linear systems $Ax = b$, where A is a symmetric positive definite matrix. Rather than computing residuals directly, it updates the residual vectors recursively. Current analyses of the conjugate gradient (CG) algorithm in finite precision typically assume that the norm of the recursively updated residual goes orders of magnitude below the machine precision, focusing mainly on bounding the residual gap thereafter. This work introduces a framework for the PCG algorithm and provides rigorous proofs that the relative backward and forward errors of the computed results of PCG can reach the levels $O(u)$ and $O(u)\kappa(A)^{1/2}$, respectively, after a sufficient number of iterations without relying on an assumption concerning the norm of the recursively updated residual, where u represents the unit roundoff and $\kappa(A)$ is the condition number of A . Our PCG framework further shows that applying preconditioners in low precision does not compromise the accuracy of the final results, provided that reasonable conditions are satisfied. Our theoretical results are illustrated through a set of numerical experiments.

Jakub Kružík: Accelerating MGP-type Methods Through Preconditioning

MGP-type algorithms are highly successful in solving quadratic programming problems. A natural improvement is to use preconditioning to accelerate convergence. However, this preconditioning is applicable only on the free set—a so-called preconditioning in face. This means the preconditioner must be recomputed every time the approximation of the free set changes. In this talk, we propose an approximate preconditioning in face that computes the preconditioner only once. We characterize the error of the new approach compared to the optimal preconditioning in face and provide a sharp bound on the condition number of the preconditioned operator. Finally, numerical experiments demonstrate that, unlike standard preconditioning in face, our new approach achieves very large speedups.

Dalibor Lukáš: Operator Preconditioning in Boundary Element Methods on GPUs

A less-known advantage of boundary element methods when compared to volume discretization techniques is that the single-layer and hyper-singular operators have opposite-order mapping properties, which allows their mutual operator preconditioning. However, in 3 spatial dimensions the lowest-order discretizations of the operators by discontinuous piecewise constant and continuous piecewise linear functions, respectively, do not match in terms of degrees of freedom. Therefore, a dual mesh is introduced to discretize the single-layer operator. Unfortunately, the assembly is significantly more expensive. In this talk we propose and analyze a novel construction of continuous piecewise polynomial basis functions to discretize the hyper-singular operator, both in 2 and 3 dimensions. We prove that it forms an optimal preconditioner to the original single-layer operator discretized by the piecewise constants. We avoid the dual mesh. The efficiency of our approach is documented by numerical experiments performed on GPUs.

Tomáš Oberhuber: TNL: Numerical library for modern parallel architectures

TNL (www.tnl-project.org) is a collection of building blocks that facilitate the development of efficient numerical solvers and HPC algorithms. It is implemented in C++ using modern programming paradigms in order to provide a flexible and user-friendly interface similar to, for example, the C++ Standard Template Library. TNL provides native support for modern hardware architectures such as multicore CPUs, GPUs, and distributed systems, which can be managed via a unified interface. In our presentation, we will demonstrate the main features of the library together with efficiency of the implemented algorithms and data structures.

Wednesday, 20 May

Jiří Tomčala - Quantum Optimization of a Community Energy Network Using a Simulator and a Real Quantum Computer

In recent years, the integration of local energy sources and local energy consumption into unified community energy networks has gained increasing importance. This development is driven not only by economic considerations but also by environmental concerns. Through the interconnection of a community energy network, local energy sources can be utilized more efficiently, thereby reducing the amount of energy drawn from the public distribution grid. Optimizing these interconnections represents a compelling use case that can be addressed using quantum optimization algorithms. The principal contribution of this work lies in the comparison of quantum optimization results obtained from a simulator and a real quantum computer. This comparison highlights the limitations of current quantum hardware and helps explain why contemporary quantum computers are not yet capable of solving large-scale optimization problems in scenarios where a quantum advantage is theoretically expected.

Van Binh Henri - From Smooth Valleys to Barren Plateaus: The Hidden Geometry of Quantum Optimization

This research investigates the topological properties of cost function manifolds in Variational Quantum Algorithms (VQAs). Rather than benchmarking accuracy, we employ Spectral Hessian Analysis and Lyapunov Exponents to quantify the ruggedness and chaoticity of high-dimensional optimization landscapes. Hessian spectral statistics characterize the local geometric structure of the variational landscape, while Lyapunov exponents capture the stability of optimization trajectories under gradient-based dynamics. Together, they provide complementary information about trainability and dynamical transitions in VQAs. Using state-vector simulations on GPU-accelerated HPC nodes (Karolina IT4I), we aim to identify structural phase transitions—such as the onset of Barren Plateaus or Information Scrambling—that serve as predictive signatures for quantum advantage. This work establishes a quantitative framework for determining problem-to-algorithm suitability based on the intrinsic dynamical properties of the optimization task.

Ivana Miháliková - Using symmetries to prepare ground states on quantum computers

Quantum computers are promising devices for simulating quantum systems. However, the simulation usually starts with preparing an initial quantum state, and this step can strongly affect the whole algorithm. Many physical Hamiltonians have symmetries, which can be used to restrict the search to the physically relevant part of the Hilbert space. In this work, we use this idea for two spin models: a neutrino-inspired model with all-to-all interactions and a small Heisenberg lattice. We start from simple product states and project them into symmetry sectors with fixed total spin J and J_z . For the Heisenberg model, we also use translation and mirror symmetries. Then we improve these states using a variational ansatz based on swap gates, which keeps the state inside the correct symmetry structure. The results show that symmetries are not only important from the theoretical point of view, but they can also strongly improve the practical performance of quantum algorithms. In the neutrino model, we obtain a fidelity of 98.8% with the exact ground state. For the Heisenberg model, the fidelity becomes larger than 98% when all symmetries are included. The symmetry constraints also reduce the effective search space in the Heisenberg case from 4096 states to only 9 relevant states. This suggests that using symmetries in state preparation can make searching for the ground state smoother, faster, and more accurate, especially for near-term quantum computers.

Thursday, 21 May

Ada Böhm - HyperQueue: Scaling to Millions of Tasks

HyperQueue is a high-performance scheduler and meta-scheduler designed to bridge the gap between complex, large-scale scientific workflows and traditional High-Performance Computing (HPC) resource managers like Slurm or PBS. It addresses the limitations of system-level schedulers when handling millions of short-lived, heterogeneous tasks by providing a lightweight, sub-node execution layer that operates as a transparent runtime. By utilizing a central server and a dynamic pool of workers, HyperQueue decouples task management from

the underlying batch system, offering advanced features such as automatic resource allocation, NUMA-aware placement, and fine-grained management of GPUs and memory.

Roman Iakymchuk - Energy-aware CFD simulations: Sharing Experience from the CEEC project

Energy efficiency has emerged as a central challenge for modern high-performance computing (HPC) systems, where escalating computational demands and architectural complexity have led to significant energy footprints. I will present the collective experience of the EuroHPC JU Center of Excellence in Exascale CFD (CEEC) in measuring, analyzing, and optimizing energy consumption across major European HPC systems. We briefly review key methodologies and tools for energy measurement as well as define metrics for reporting results. Through case studies using representative CFD applications (waLBerla, FLEXI/GALÆXI, Neko, and NekRS), we evaluate energy-to-solution and time-to-solution on diverse architectures, including CPU- and GPU-based partitions of LUMI, MareNostrum5, MeluXina, and JUWELS Booster. Our results highlight the advantages of accelerators and mixed-precision techniques for reducing energy consumption while maintaining computational accuracy. Finally, we advocate the need to facilitate energy measurements on HPC systems in order to raise awareness, teach the community, and take actions toward more sustainable exascale computing.

Filip Vaverka - MaX: Co-design, technology exploitation & energy efficiency

The third phase of the MaX Centre of Excellence is designed to provide the European materials simulation community with the necessary tools to adapt to the technological advancements in massively parallel heterogeneous computing systems, enabling them to tackle previously unattainable scientific challenges. In this talk, we explore the behavior of some of the MaX flagship codes (such as YAMBO) when running on advanced hardware platforms with unique feature sets (such as NVIDIA Grace or Intel Sapphire Rapids with HBM). The primary focus of our analysis is energy efficiency and impact of the memory hierarchy on the codes and the associated mini-apps.

Štěpán Vaněček - Understanding Heterogeneous Systems Through Topology-Aware Tools

State-of-the-art HPC and AI systems increasingly rely on heterogeneous and dynamic architectures, while existing tools, applications, and workflows often fail to adapt. This talk presents a coherent toolkit that treats topological information as a first-class input for understanding system and application behaviour. We introduce MT4G, a tool for collecting detailed NVIDIA and AMD GPU topology and attributes using vendor APIs where available and over 50 auto-evaluated microbenchmarks. sys-sage, a library for managing and correlating static and dynamic topological information from multiple sources, includes MT4G, hwloc, and other, especially dynamic. Beyond CPUs and GPUs, sys-sage supports QPUs, with existing integrations of QDMI and IQM readouts. Building on this foundation, xPPV provides an end-to-end workflow for measuring, analysing, and exploiting CPU performance variability across topological levels in large-scale systems.

Joao Barbosa - Supercharging Your HPC Applications: Free Expert Support through EPICURE and POP-3

Getting the most out of Europe's world-class supercomputers is no small feat. This talk introduces two free, EuroHPC JU-funded services that make it easier: EPICURE and POP-3. EPICURE (epicure-hpc.eu) delivers hands-on Level 2 and Level 3 application support — covering code porting, scaling, performance analysis, benchmarking, and optimisation — through a single portal at eurohpcsupport.eu, backed by expert teams across all EuroHPC hosting sites. POP 3 (pop-coe.eu) provides free, independent, and quantitative performance assessments for academic and industrial HPC codes in any domain, helping users identify inefficiencies and chart a clear path to better performance — now with a focus on exascale readiness. Together, they cover the full journey: POP-3 diagnoses, EPICURE fixes. Both are free to eligible EU organisations. Come find out how your project can benefit.

Accepted Conference Contributions Not Presented On Site

Sharapov Mykhaylo - Strategies For Ensuring Functional Stability Of A Complex Sensor Network Based On The Research Of The Dynamics Of The Behavior Of Evolutionary Equations

Wireless sensor networks (WSNs) are a promising solution for continuous monitoring, risk detection, and adaptive response in large-scale and hazardous environments, such as disaster zones. However, their operation is affected by numerous destabilizing factors, requiring approaches to enhance functional resilience. Understanding network behavior under adverse conditions is essential for predicting degradation and preventing critical failures through adaptive strategies. This work applies qualitative differential equation theory to model WSN dynamics under risk. The classical SIR model is used to formalize system behavior, and conditions for stability and asymptotic stability are derived. Strategies to minimize compromised devices are analyzed via phase plane methods. The study establishes how stability depends on failure and recovery rates, and proposes a model to estimate recovery time. A mathematical framework with impulsive control is developed to mitigate risk impact.

List of Posters

No.	Poster Title	Authors
1	Collective and individual oscillations of prominence threads excited by vortex shedding	Sofya Belov ,Petr Jelínek, Marian Karlický, Robert Erdélyi
2	Parallel implementation of slope stability computations: from Matlab to HPC	Michal Béreš, Stanislav Sysala, Simona Bérešová, Tomáš Luber
3	Visual Workflow Builder for HPC	Dhia Bouaziz, Martin Golasowski, Jan Martinovič, Tomáš Martinovič
4	Accelerating Aerosol Hygroscopicity Prediction with Deep Learning and HPC Deployment	Radek Halfar, Billy Braithwaite, Juha Tonttila, Laakso Anton, Kokkola Harri, Sarkar Arup, Antti Vartiainen
5	Accelerating MPGP-type Methods Through Preconditioning	David Horák, Jakub Kružík
6	Multiphysics Finite Element Applications in Simcenter STAR-CCM+	Pavel Hron, Andreas Hauck
7	Visualisation of Massive Black Hole Orbital Decay in Dark Matter Halos	Milan Jaroš, Alice Damiano, Osman Yasal, Petr Strakoš, Lubomír Říha
8	Turbulent plasma flow, its energies, and structures: Velocity vortices, magnetic field cocoons, and plasmoids	Petr Jelínek, Marian Karlický
9	A Novel Clustering-Based Approach for Analyzing Behavior in Nonlinear Dynamical Systems Using Self-Organizing Maps	Syeda Sarwat Kazmi, Tomáš Kozubek, Tomáš Martinovič
10	EUMaster4HPC study programme and related projects	Tomáš Kozubek
11	From Pascal to Blackwell: Benchmarking Multi-GPU k-Wave with a Global FFT Solver	Oliver Kunik, Jiří Jaroš
12	The conjugated gradients and low-rank Lyapunov equation	Jana Lungová, Martin Plešinger
13	Marine video monitoring advances leveraging HPC infrastructure	Louis Roman Fiorina, Ghaith Chaabane, Kalindi Fonda, Tuomas Rossi, Victor Anton, Tomáš Martinovič, Matthias Obst
14	Fluid-Structure interaction of a high-aspect-ratio wing	Ondřej Meca, Tomáš Brzobohatý
15	Adaptive ODE Solving with the Taylor Series Method	Petr Veigend, Václav Šátek, Gabriela Nečasová
16	HPCTRAIN: Building Europe's Future HPC Workforce through Structured Traineeships Abstract	Veronika Pakostová, Karina Pešatová
17	EVITA: Building a European Virtual Training Academy for HPC	Karina Pešatová

18	Resource-Efficient Variational Quantum Classifier	Petr Ptáček, Paulina Lewandowska, Ryszard Kukulski
19	SEANERGENS - Boosting HPC and AI energy efficiency with a European software solution	Lubomír Říha, Ondřej Vysocký
20	AI-Based Shape Optimization of High-Lift Flap Configurations for Custom-Built Aircraft	Tomáš Brzobohatý, Marie Sadowská, E. Vandame
21	Deep Learning-Based Segmentation of AFM Data for Epitaxial Graphene Analysis and SiC Terrace Identification	Khyati Sethia, Petr Strakoš
22	Scalable Simulation and Evaluation of Adaptive Fuzzy Traffic Signal Controller under Multi-Scenario Urban Conditions	Sumaira Shaheen, Jiří Dvorský
23	Enabling HPC-Accelerated Virtual Screening for Drug Discovery through the LEXIS Platform	Paulo Silva, Nestor Pylat, Martin Golasowski, Jan Martinovič, Ludovica Picarone, Daniele Pietrucci, Giovanni Chillemi, Carmine Talarico
24	Accelerating Research in Brain Science and Neurodegenerative Diseases: CLARA Testbed	Kateřina Slaninová, Jan Martinovič, Vít Vondrák
25	Extreme Physics Compression: Adaptive Domain Decomposition for Neural Surrogate Models	Tomáš Brzobohatý, Marcel Smolík
26	Multi-Resolution Grid Reconstruction for Efficient Volume Rendering of Large-Scale Simulation Data	Markéta Faltýnková, Ondřej Meca, Milan Jaroš, Petr Strakoš, Lubomír Říha
27	POP3: The Third Year of Advancing HPC Performance and Productivity	Radim Vavřík, Tomáš Panoc, Ondřej Vysocký, and Lubomír Říha
28	Energy Efficiency Runtime System for GPUs	Daniel Velička, Ondřej Vysocký, Lubomír Říha
29	AI-based surrogate modelling for incremental sheet metal forming	Tomáš Brzobohatý, Erika Viščorová, Peter Scholz, Marta Jarošová
30	DARE: European RISC-V chiplet processors for HPC and AI	Ondřej Vysocký, Osman Yasal, Daniel Velička, Lubomír Říha
31	OPTKIT: from x86 to RISC-V universal performance analysis	Osman Yasal
32	From Redundant Metrics to New Insight: A Cross-Case Analysis of WSS-Derived Hemodynamic Features	Tomáš Brzobohatý, Viktor Gajdošík, Václav Procházka
33	Stokes Problem with Stick-Slip Boundary Conditions Application of the SCD Semismooth* Newton Method	Vladimír Arzt, Petr Beremlijski
34	Quantum-Inspired Ising Framework for Cholera Transmission Dynamics	Umar Shafique, Marek Lampart, Ali Raza
35	Towards Benchmarking the VLQ Quantum Computer	Ivana Miháliková, Paulina Lewandowska, Adam Bílek, Ryszard Kukulski, Petr Ptáček

36	Data-Driven Inefficiency Analysis and AI-Based Predictions for Real-World Industrial Plant Data	Kristina Motyčková, Marcel Smolík
37	Empowering Women in High Performance Computing: Activities of the Central European Chapter of Women In HPC	Lucie Rožnovská, Karina Pešatová
38	Advancing Scientific Job Dispatch: Core Modernization, Cloud Integration, and Visual Workflow Tooling for k-Dispatch	Marta Jaroš, Marcel Feiler, Oleksandr Musiichuk, Silvia Šlachtovská, Jíří Jaroš
39	A QA-Gated Ensemble Framework for robust AI	David Číž, Joao Carias, Monika Weissová, Petra Bobíková

List of Participants

Name Surname	Institution
Peter Arbenz	ETH Zurich
Vladimír Arzt	VSB-TUO/KAM
Thomas Bake	Charles University
Joao Barbosa	IT4Innovations, VSB-TUO
Sofya Belov	University of South Bohemia, Faculty of Science
Petr Beremlijski	VSB-Technical University of Ostrava
Michal Béréš	Institute of Geonics, Czech Academy of Sciences
Simona Béréšová	Institute of Geonics, Czech Academy of Sciences
Ada Böhm	IT4Innovations, VSB-TUO
Dhia Bouaziz	IT4Innovations, VSB-TUO
Tomáš Brzobohatý	IT4Innovations, VSB-TUO
Erin Claire Carson	Charles University MFF
David Číž	IT4Innovations, VSB-TUO
Zdeněk Dostál	VSB-Technical University of Ostrava
Jiří Filipovič	Masaryk University
Michael Granitzer	University of Passau
Radek Halfar	IT4Innovations, VSB-TUO
Tomáš Hammerbauer	Charles University
Jakub Homola	IT4Innovations, VSB-TUO
David Horák	Institute of Geonics, Czech Academy of Sciences
Pavel Hron	Siemens Industry Software
Roman Iakymchuk	Umeå University and Uppsala University
Branislav Jansík	IT4Innovations, VSB-TUO
Marta Jaroš	Brno University of Technology
Milan Jaroš	IT4Innovations, VSB-TUO
Jiří Jaroš	Brno University of Technology
Petr Jelínek	University of South Bohemia, Faculty of Science
Syeda Sarwat Kazmi	IT4Innovations, VSB-TUO

Tomáš Kozubek	IT4Innovations, VSB-TUO
Rolf Krause	KAUST - King Abdullah University of Science and Technology
Michal Krupička	Škoda Auto
Jakub Kružík	Institute of Geonics, Czech Academy of Sciences
Oliver Kunik	Brno University of Technology
Marek Lampart	IT4Innovations, VSB-TUO
Ulrich Langer	Johannes Kepler University Linz
Dalibor Lukáš	VSB-Technical University of Ostrava
Jana Lungová	Technical University of Liberec
Jan Martinovič	IT4Innovations, VSB-TUO
Tomáš Martinovič	IT4Innovations, VSB-TUO
Michael Mc Gettrick	University of Galway
Ondřej Meca	IT4Innovations, VSB-TUO
Jakub Mielczarek	Jagiellonian University
Ivana Miháliková	IT4Innovations, VSB-TUO
Matthias Möller	Delft University of Technology
Kristina Motyčková	IT4Innovations, VSB-TUO
Gabriela Nečasová	Brno University of Technology
Tomáš Oberhuber	Czech Technical University in Prague
Matthias Obst	University of Gothenburg
Veronika Pakostová	IT4Innovations, VSB-TUO
Karina Pešatová	IT4Innovations, VSB-TUO
Martin Plešinger	Technical University of Liberec
Tomáš Posker	IT4Innovations, VSB-TUO
Petr Ptáček	IT4Innovations, VSB-TUO
Dominika Regéciová	IT4Innovations, VSB-TUO
Lubomír Říha	IT4Innovations, VSB-TUO
Lucie Rožnovská	IT4Innovations, VSB-TUO
Ulrich Ruede	TU Ostrava and FAU Erlangen-Nürnberg
Marie Sadowská	IT4Innovations, VSB-TUO

Khyati Sethia	IT4Innovations, VSB-TUO
Umar Shafique	VSB-Technical University of Ostrava
Sumaira Shaheen	VSB-Technical University of Ostrava
Martin Schulz	Technical University of Munich
Paulo Silva	IT4Innovations, VSB-TUO
Pavlna Smolková	IT4Innovations, VSB-TUO
Marcel Smolík	IT4Innovations, VSB-TUO
Kirk M Soodhalter	Trinity College Dublin
Olaf Steinbach	TU Graz
Petr Strakoš	IT4Innovations, VSB-TUO
Petra Svobodova	IT4Innovations, VSB-TUO
Daniel Sýkora	Czech Technical University in Prague
Jakub Šístek	Institute of Mathematics of the Czech Academy of Sciences
Petr Tichý	Faculty of Mathematics and Physics, Charles University
Jiří Tomčala	IT4Innovations, VSB-TUO
Lukáš Vach	M Computers
Štěpán Vaněček	TU Munich
Filip Vaverka	IT4Innovations, VSB-TUO
Radim Vavřík	IT4Innovations, VSB-TUO
Daniel Velička	IT4Innovations, VSB-TUO
Erika Viščorová	IT4Innovations, VSB-TUO
Vít Vondrák	IT4Innovations, VSB-TUO
Van Binh Henri Vu	IT4Innovations, VSB-TUO
Ondřej Vysocký	IT4Innovations, VSB-TUO
Osman Yasal	IT4Innovations, VSB-TUO

