

Preconditioning Techniques for Scientific and Industrial Applications

ICMS, Edinburgh, 27–29 May 2026

<https://icms.ac.uk/activities/workshop/precond26/>
<https://sites.google.com/view/precond26/plenary-speakers>

Conference Timetable and Book of Abstracts

Conference Chairs: Esmond Ng (Lawrence Berkeley National Laboratory), Yousef Saad (University of Minnesota), Andy Wathen (Rutherford Appleton Laboratory/University of Oxford)

Local Organisers: John Pearson (University of Edinburgh; chair), Jacek Gondzio (University of Edinburgh), Jennifer Pestana (University of Strathclyde), Alison Ramage (University of Strathclyde), Alex Liu (ICMS, Edinburgh)

Timetable

The conference timetable may be found on the following pages. Changes to the original version of the timetable are coloured in [blue](#). Below are links to the various conference sessions:

[Link to Abstracts for Invited Talks](#)
[Link to Abstracts for Minisymposium M1](#)
[Link to Abstracts for Minisymposium M2](#)
[Link to Abstracts for Minisymposium M3](#)
[Link to Abstracts for Minisymposium M4](#)
[Link to Abstracts for Minisymposium M5](#)
[Link to Abstracts for Minisymposium M6](#)
[Link to Abstracts for Minisymposium M7](#)
[Link to Abstracts for Minisymposium M8](#)
[Link to Abstracts for Minisymposium M9](#)
[Link to Abstracts for Minisymposium M10](#)
[Link to Abstracts for Minisymposium M11](#)
[Link to Abstracts for Minisymposium M12](#)
[Link to Abstracts for Minisymposium M13](#)
[Link to Abstracts for Minisymposium M14](#)
[Link to Abstracts for Minisymposium M15](#)
[Link to Abstracts for Contributed Talks](#)
[Link to Minisymposium Titles and Organisers](#)

Wednesday 27th May

9:05–9:45	Registration	Bayes Centre Ground Floor		
9:45–9:50	Welcome	Bayes Theorem, ICMS Lecture Theatre, and ICMS Seminar Room		
	M1: Bayes Theorem	M2: ICMS Lecture Theatre	M3: ICMS Seminar Room	
9:50–10:10	K Soodhalter M1 Extensions of Admissible and Attainable Convergence Theory for GMRES	C Kruse M2 Learning to Smooth: Enhancing Multigrid Solvers with Data	T Rees M3 Parallel in Time Preconditioning for Solving the Paraxial Wave Equation in Ptychographic Applications	
10:10–10:30	N Spillane M1 Weighted GMRES Accelerated by Preconditioning and Deflation	Y Gong M2 A Learning-Enhanced Nonlinear Preconditioner for Some Hyperelasticity Problems	C Jerez-Hanckes M3 Operator Preconditioning of Multiple Trace Formulations for Maxwell Transmission Problems	
10:30–10:50	J Pestana M1 A GMRES Convergence Bound Based on Pseudospectra for Preconditioned Toeplitz-like Systems	D Parthasarathy M2 Automated Grammar-based Design Of Multigrid Preconditioners With Evolutionary Algorithms	Y Voet M3 Deflation-Based Preconditioning for Immersed Finite Element Methods	
10:50–11:10	P Marchand M1 Deflation and Preconditioning of GMRES for Helmholtz Problems with (Quasi-)Resonances	B Jiang M2 A Polynomial Dimension-Dependence Analysis of Bramble–Pasciak–Xu Preconditioners	E Treister M3 Advances in Multigrid Preconditioning for the Helmholtz Equations	
11:10–11:40	COFFEE/TEA			
	M4: Bayes Theorem	M5: ICMS Lecture Theatre	M3: ICMS Seminar Room	
11:40–12:00	D Fortunato M4 Fast and Accurate Simulation of Close-to-Touching Discs in Stokes Flow	M Freitag M5 The Potential Importance of Re-orthogonalisation for Least Squares Problems	C Stolk M3 Multigrid and Dispersion Correction for Preconditioning Helmholtz Equations	
12:00–12:20	S Leveque M4 Mixed-Precision Parallel-in-Time Solvers for Runge–Kutta Methods	M Meier M5 Backward Stable (Randomized) Preconditioning for Least Squares Problems	P-H Cocquet M3 Comparison of Two Dispersion Correction Methods for Finite-Difference Schemes Applied to the Helmholtz Equation	
12:20–12:40	H Wolles Ljósheim M4 Parallel-in-Time Preconditioning for Time-Dependent Variational Mean Field Games as Generative Models	S Gürol M5 A Spectral Preconditioner for the Conjugate Gradient Method with Iteration Budget	A Tonnoir M3 Dispersion Correction Methods for the Convected Helmholtz Equation	
12:40–14:00	LUNCH			
<i>Chair:</i>	TBC	ICMS Lecture Theatre (streamed to Bayes Theorem)		
14:00–14:40	M Embree	Approximate Preconditioning and Pseudospectral Shattering		
14:40–15:20	J Tabcart	Preconditioning Weak-Constraint 4D Variational Data Assimilation: A Parallelisable Implementation in Firedrake		
15:20–15:50	COFFEE/TEA			

Wednesday 27th May

M6: Bayes Theorem		M5/M2: ICMS Lecture Theatre		M7/CT: ICMS Seminar Room	
15:50–16:10	Y Xi Neural Approximate Inverse Preconditioners	M6	J Papež Error Estimate and Stopping Criteria for (Preconditioned) Least-Squares Problems Solved by CGLS, LSQR, and Their Block Variants	M5	M Deiml The Practical Problems of Quantum Linear System Solvers and a Krylov-based Solution
16:10–16:30	F Brarda Preconditioning Via Spectral Density Driven Graph Neural Networks	M6	J Scott Low Precision Incomplete Factorization Preconditioners for Linear Least-Squares Problems	M5	D Peterseim Multilevel Preconditioning for PDEs on Quantum Computers
16:30–16:50	M Scott Design Criteria for SGD Preconditioners: Local Conditioning, Noise Floors, and Basin Stability	M6	A Lister Mixed Precision Iterative Refinement – Solving Least Squares with More Least Squares	M5	E Fressart Investigating Preconditioning for Quantum Linear System Solvers
16:50–17:10	Z Di Preconditioned Stochastic Gradient Methods for Ptychographic Reconstruction	M6	X-C Cai A Data-Enhanced Coarse Preconditioner for Linear System of Equations	M2	M Baboulin A Quantum-Classical Mixed Precision Algorithm for Solving Linear Systems
17:10–17:30	D Cai A Stochastic View of Preconditioning with Applications to Stochastic Learning	M6	A Klawonn Machine-Learning-Enhanced Domain Decomposition Preconditioners	M2	S Shah Block Jacobi Preconditioning on Analog Hardware
17:30–19:00 WELCOME RECEPTION (soft drinks and snacks; opportunity for discussion between early-career and more senior researchers)					

Thursday 28th May

M8/Invited: Bayes Theorem		M9: ICMS Lecture Theatre		M10: ICMS Seminar Room	
9:10–9:30	H Hénon Multi-fidelity Preconditioner for Ocean Variational Data Assimilation	M8	N Bootland Parallel Preconditioning for Time-Harmonic Waves from Toeplitz Structured Sweeping	M9	S Dasari Extrapolated Multigrid Methods for Elliptic PDEs on Tokamak Geometries
9:30–9:50	I Daužickaitė Randomized Preconditioning for an Ensemble of Data Assimilations	M8	L Gontier A Stochastic BDD strategy for Stochastic Elliptic Problems	M9	F Durastante Preparing deal.II for Exascale with PSCToolkit Preconditioners
9:50–10:10	H Elzayyadi The Role of Rounding Errors on Spectral-Limited Memory Preconditioner	M8	J Weber-Hamacher A Flow-Rate Conserving CNN-Based Domain Decomposition Method for Blood Flow Simulations	M9	C Janna Aggressive Coarsening and Energy Minimization for Lightweight AMG in CFD Problems
10:10–10:30	T Faney Invited Industry Talk Machine Learning for Faster Numerical Simulation: From Nonlinear Solvers to Learned Preconditioning		D Szyld Multiprecision Computations with Schwarz Methods	M9	P Jolivet Robust Algebraic Preconditioning in PETSc with Domain Decomposition Methods
10:30–10:50	T Faney (continued)				A Franceschini A Reverse Augmented Constraint Preconditioner for Industrial Contact Mechanics Applications
10:50–11:20 COFFEE/TEA					
CT: Bayes Theorem		CT: ICMS Lecture Theatre		CT: ICMS Seminar Room	
11:20–11:40	J Nagy Low-Precision Preconditioning for Inverse Problems	CT	S Le Borne Coupled Ordering Techniques for Coupled Partial Differential equations	CT	P Vacek Multigrid Methods with Approximate Coarsest-Level Solvers
11:40–12:00	P Khatun HSS Matrices are Riemannian Manifolds	CT	R Noffke Exploiting Port-Hamiltonian Structures for Preconditioning Saddle Point Problems	CT	L Rooch Optimal Transfer Operators and Convergence Bounds for Nonsymmetric Two-Grid Methods
12:00–12:20	A Meda mdBFGS: A Diagonal BFGS Method with Momentum for Stochastic Optimisation	CT	T Akter Scalable ILU-Based Approximate Selected Inversion with Spectral Corrections for Sparse Systems	CT	D Osei-Kuffuor Robust MGR Preconditioning for Large-Scale Coupled Subsurface Flow and Geomechanics
12:20–12:40	R Turner Basis Preconditioning for Gas Flow Optimization	CT	M Ferronato A Multi-Physics Reduction Preconditioner for a Fully Coupled 3D Dynamic Biot Model	CT	B Sidi Hida Monolithic and Schur-Based Domain Decomposition Strategies for Hybrid High-Order Discretizations
12:40–13:55 LUNCH					

Thursday 28th May

<i>Chair:</i>	TBC	ICMS Lecture Theatre (streamed to Bayes Theorem)			
13:55–14:35	A Kopaničáková	Preconditioning for the Trustworthy Use of Machine Learning in Large-Scale Numerical Simulations			
14:35–15:15	T Mary	In What Precision(s) Should One Precondition?			
15:15–15:45	COFFEE/TEA				
	M11: Bayes Theorem	M12: ICMS Lecture Theatre	M13: ICMS Seminar Room		
15:45–16:05	V Dolean M11 Local Feature Filtering for Scalable and Well-Conditioned Domain-Decomposed Random Feature Methods	M Benzi M12 Scalable Augmented Lagrangian Preconditioning of Fictitious Domain Problems	I Hnětynková M13 Revisiting Flexibly Preconditioned Hybrid LSQR for Efficient Large-Scale Regularization		
16:05–16:25	P Zunino M11 Geometry-Aware Neural Preconditioning for Parametrized PDEs	F Mugnaioni M12 Modified Augmented Lagrangian Preconditioners for Fictitious Domain Formulations of Elliptic Interface Problems	M Sabaté Landman M13 The Role of Inexactness in Krylov Subspace Regularization		
16:25–16:45	P Mycek M11 Accelerating FGMRES for Helmholtz Problems via Learned Nonlinear Preconditioning	G Luhana M12 Truncated Domain Preconditioner for the Implicit-in-Time Immersed Boundary Method	M Donatelli M13 Preconditioning Strategies for a Nested Primal–Dual Method		
16:45–17:05	L Luo M11 Nonlinear Preconditioning Algorithms with Learning Capability	A Martínez Calomardo M12 Eigenvalue Bounds for Symmetric Multiple Saddle-Point Systems Preconditioned with Block-Diagonal Schur Complement matrices	H Al Daas M13 Extended-Krylov-Subspace Methods for Trust-Region and Norm-Regularization Subproblems		
17:05–17:25	R Li M11 AI-Driven Algebraic Multigrid: Graph Neural Networks for Operator Design	L Bergamaschi M12 Block Triangular Preconditioning of Symmetric, Multiple Saddle-Point Matrices: Eigenvalue Bounds with Inexact Schur Complements	E de Sturler M13 Acceleration of Block-Column Based Solvers for Structured Inverse Problems		
17:25–17:35	BREAK				
	Bayes Theorem	ICMS Lecture Theatre			
17:35–18:25	Panel discussion on quantum computing [*]	Panel discussion on machine learning and AI [*]			

[*] The idea of the two panel discussion sessions at the end of the day is to enable an exchange of perspectives on the state-of-the-art and future challenges for preconditioning and numerical linear algebra, in relation to the fields of machine learning, AI, and quantum computing.

Friday 29th May

M14: Bayes Theorem		M12: ICMS Lecture Theatre		M15: ICMS Seminar Room	
9:10–9:30	H Wolkowicz The Omega- and Kappa-Condition Numbers and Optimal Preconditioning	M14	J Liesen Exploiting Non-Standard Inner Products for Saddle Point Problems	M12	C Brezinski History of Extrapolation Methods
9:30–9:50	J Gondzio Challenges in Design of Preconditioners for Interior Point Methods	M14	R Yovel Approximate Commutator Preconditioners for Saddle-Point Systems with Indefinite Leading Blocks	M12	C Brezinski and M Redivo-Zaglia M15
9:50–10:10	A Gevorgyan An Incomplete Factorization Preconditioner for Saddle Point Systems Arising in Interior Point Methods	M14	C Greif A BFBt Preconditioner for Double Saddle-Point Systems	M12	M Redivo-Zaglia Extrapolation Methods for Vector, Matrices and Tensors
10:10–10:30	L-R Santos Randomized Nyström Preconditioned Interior Point-Proximal Method of Multipliers	M14			A Simon Acceleration Techniques for Implicit Adaptive Multistep Coupling in Aerospace Engineering Applications
10:30–10:50	A Montoison Preconditioned Krylov Solvers for KKT systems on GPUs	M14			Y Saad Acceleration Methods for Fixed Point Iterations
10:50–11:15 COFFEE/TEA					
<i>Chair:</i>	TBC		Informatics Forum G.07		
11:15–11:55	Y Nakatsukasa		Adaptive CUR Preconditioning for LSQR		
11:55–12:35	B Southworth		Preconditioning in the Land of Optimization		
12:35–12:45 BREAK					
12:45–13:25	M Udell		Online Scaled Gradient Methods		
13:25–15:00 LUNCH					

Abstracts

▷ Invited Talks

Approximate Preconditioning and Pseudospectral Shattering

Mark Embree, Virginia Tech, USA

In a variety of settings one can establish compelling convergence bounds for GMRES, based on some ideal form of a preconditioner. Often such preconditioners involve infeasible quantities (such as exact inverses of block matrices, or exact invariant subspaces) that can only be approximated in practice. How much do such approximations delay the expected performance of GMRES? We will address this question, describing rigorous bounds on the convergence degradation in terms of the approximation accuracy. We will draw an analogy with the “pseudospectral shattering” phenomenon identified by Nikhil Srivastava and colleagues in their resolution of Brian Davies’s conjecture about “approximate diagonalization”. (This talk describes collaborative work with Ron Morgan, Josef Sifuentes, and others.)

→ Back to timetable

Preconditioning Weak-Constraint 4D Variational Data Assimilation: A Parallelisable Implementation in Firedrake

Jemima Tabcart, Eindhoven University of Technology, The Netherlands

Data assimilation algorithms combine measurement information with numerical models to obtain an updated estimate of the state of a dynamical system. In recent years there has been a significant mathematical effort to develop new solvers and preconditioners to accelerate the solution of variational data assimilation methods. However, testing new approaches on a wide variety of problems has historically been challenging, due to the need to write code for differential equation solvers (including tangent linear and adjoint models) for each new test case.

I will begin by introducing the variational data assimilation problem, and highlighting the importance of preconditioners within practical implementations. I will then discuss our new work which embeds variational data assimilation methods within Firedrake. The new Firedrake-VarDA toolbox allows user to immediately apply variational data assimilation methods to any problem that can be written in Unified Form Language, with automatic parallelisation over observation times, and no need to manually define PDE solvers or adjoints. I will demonstrate the flexibility of our implementation both for solving an array of test problems, and applying different preconditioners and solvers.

→ Back to timetable

Machine Learning for Faster Numerical Simulation: From Nonlinear Solvers to Learned Preconditioning

Thibault Faney, IFP Energies Nouvelles & Casablanca, France

[Dr Faney is an invited speaker for his work in industry.]

Large-scale numerical simulations are often limited by the cost of solving nonlinear and sparse linear systems arising from PDE discretizations. In this talk, I will present three recent works showing how machine learning can accelerate simulation at different levels of the solver stack.

The first contribution addresses difficult well-opening and well-closing events in porous media flow simulation for CO₂ storage. Using Fourier Neural Operators, it predicts a Newton initialization directly in the domain of quadratic convergence, acting as a nonlinear preconditioning strategy that reduces Newton iterations and mitigates time-step reductions.

The second contribution focuses on Poisson problems in CFD. It uses recurrent and implicit GNN-based solvers on unstructured meshes, trained in a physics-informed way to produce accurate approximations.

The third, ongoing contribution investigates GNNs for algebraic preconditioning of sparse linear systems through a hybrid incomplete-factorization strategy combining learned and model-based components.

Together, these works illustrate how machine learning can enhance numerical simulation from nonlinear acceleration to learned preconditioning.

→ Back to timetable

Preconditioning for the Trustworthy Use of Machine Learning in Large-Scale Numerical Simulations

Alena Kopaničáková, Toulouse-INP, France

Recently, scientific machine learning (SciML) has expanded the capabilities of traditional numerical methods by simplifying computational modeling and enabling cost-effective surrogate models. However, despite their growing success, SciML approaches often lack explicit error control, require computationally intensive training, and raise concerns regarding robustness and reliability. Addressing these limitations is a central challenge for the trustworthy deployment of SciML in large-scale numerical simulations. In this talk, we argue that preconditioning provides a unifying principle for bridging classical numerical methods and SciML. We first show how advanced preconditioning strategies, such as multilevel and domain-decomposition methods, can be exploited to accelerate training and improve the accuracy of SciML models. We then demonstrate how SciML models can, in turn, be used to enhance standard numerical preconditioners, leading to hybrid solvers that combine the reliability of classical numerical methods with the efficiency of SciML. The effectiveness of the proposed strategies will be illustrated through a series of numerical experiments, including the training of DeepONets and employing DeepONet-enhanced preconditioners.

→ Back to timetable

In What Precision(s) Should One Precondition?

Theo Mary, Sorbonne Université, CNRS, LIP6, France

This talk will discuss what precision(s) one should use for preconditioning, with an emphasis on finite precision floating-point arithmetic. We will consider both the precision to construct (set up) the preconditioner and the precision to apply it to a vector. While constructing the preconditioner in lower precision is common practice and reasons to do so are well understood, the precision that one should use for applying it is still a matter of debate. In the context of Krylov solvers, we will

compare the diametrically opposed strategies of applying the preconditioner in a lower precision than the rest of the Krylov solver operations, or on the contrary in a higher precision. Then, we will turn to mixed precision preconditioning, where the preconditioner is constructed and/or applied using multiple precisions. We will give examples where mixed precision preconditioning is meaningful and successful, including domain decomposition, incomplete factorizations, and block low-rank solvers.

→ Back to timetable

Adaptive CUR Preconditioning for LSQR

Yuji Nakatsukasa, University of Oxford, UK

We describe a preconditioned LSQR algorithm for large-scale least-squares problems. The preconditioner is built using a CUR approximation whose rank gradually increases as the LSQR iteration progresses. To find a CUR approximation the algorithm uses a single small randomized sketch, which is reused. It does not require the matrix to be tall-skinny, and performs particularly well on sparse or numerically low-rank systems. Time permitting, I will also discuss recent results on the numerical stability of preconditioned Krylov methods for linear systems and least-squares problems, and randomized preconditioned iterative methods for computing the smallest singular triplet of a tall-skinny matrix.

→ Back to timetable

Preconditioning in the Land of Optimization

Ben Southworth, Los Alamos National Laboratory, USA

Modern optimization problems, from PDE-constraints to machine learning (ML), are often high-dimensional, nonconvex, and only partially structured, so robustness and efficiency depend on preconditioning: altering geometry, scales, and dynamics so 1st-order gradient methods see an easier problem. Here I translate classical numerical methods into practical optimization algorithms. First, I introduce a framework for multilevel training of Kolmogorov–Arnold-type Networks (KANs) based on their spline-basis structure. Using a nested hierarchy from knot refinement, analytic transfer operators, and complementary “relaxation”, we obtain coarse-to-fine optimization that accelerates training and improves robustness, including substantial accuracy gains in physics-informed settings. Second, I discuss nonlinear splitting as a preconditioning for gradient dynamics. By decomposing gradients into explicit and implicit components, we arrive at semi-implicit updates that mitigate stiffness, and significantly accelerate PDE-constrained optimization and ML training. Last, I will discuss preconditioning to approximately decorrelate gradient updates of matrix parameters in transformer training, yielding significant improvements in time-to-perplexity when training large language models. Altogether, these examples demonstrate that widely established preconditioning techniques can serve as principled, scalable methods for modern optimization problems.

→ Back to timetable

Online Scaled Gradient Methods

Madeleine Udell, Stanford University, USA

We introduce a framework to accelerate the convergence of gradient-based methods with online learning. This framework, the Online Scaled Gradient Method (OSGM), delivers simple, elegant proofs of convergence, mitigates the challenge of hyperparameter tuning, and yields improved variants of a wide variety of algorithms, from adaptive gradient methods to stochastic optimizers to operator splitting methods like ADMM and PDHG. This talk showcases highlights from the theory and practical applications of OSGM.

→ [Back to timetable](#)

▷ M1: Convergence Analysis and Acceleration of GMRES

Extensions of Admissible and Attainable Convergence Theory for GMRES

Kirk Soodhalter, Trinity College Dublin, Ireland

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 behavior.

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.

→ Back to timetable

Weighted GMRES Accelerated by Preconditioning and Deflation

Nicole Spillane, CNRS, École polytechnique, France

GMRES is a very well established linear solver which can be applied to a large variety of problems. Many, if not all, of us have encountered the problematic situation where GMRES does not converge – or converges very slowly. The purpose of this work is to identify which properties of a linear system govern the efficiency of GMRES. The analysis makes explicit the role of three convergence accelerators:

- Weighting, i.e., changing the inner product,
- Preconditioning, i.e., providing the solver with an approximation of the problem matrix,
- Deflation, i.e., precomputing a well chosen part of the solution.

Based on the analysis, combined choices of weighted inner product, preconditioner and deflation are proposed. Their performance is illustrated numerically.

This is joint work with Daniel B. Szyld (Temple University) and Pierre Matalon (École polytechnique).

→ Back to timetable

A GMRES Convergence Bound Based on Pseudospectra for Preconditioned Toeplitz-like Systems

Jennifer Pestana, University of Strathclyde, UK

The numerical solution of non-self-adjoint PDEs often requires the solution of large systems of linear equations with a nonnormal coefficient matrix. GMRES is often used to solve these systems, but preconditioning is usually required to increase efficiency. While descriptive convergence theory

exists to guide preconditioner choice when the coefficient matrix is normal, for nonnormal problems it can be challenging to accurately describe GMRES convergence.

In this talk we will consider linear systems arising from the numerical solution of PDEs, whose coefficient matrices exhibit (multilevel) Toeplitz or generalized locally Toeplitz (GLT) structure. For these problems the preconditioned coefficient matrices can often be decomposed into the identity plus a low-rank perturbation and a small-norm perturbation, resulting in an eigenvalue clustering near 1 with a few outliers. We present a GMRES bound, based on pseudospectra, that is applicable to such problems and investigate its properties.

This is joint work with Razan Abu-Labdeh.

→ Back to timetable

Deflation and Preconditioning of GMRES for Helmholtz Problems with (Quasi-)Resonances

Pierre Marchand, Inria, France

Finite element discretization methods are well suited for solving Helmholtz problems, because they can handle complex geometries and heterogeneous media. However, the algebraic systems obtained after the discretization are often large, ill-conditioned, non-Hermitian or possibly highly indefinite. Problems with high wavenumbers and/or resonances or quasi-resonances are especially affected by such limitations.

First, a GMRES convergence bound is given, taking into account the nonlinear behavior of the relative residual. We relate the GMRES convergence to the harmonic Ritz values, which are the roots of the minimizing polynomial of the residual. We show that a set of eigenvalues no longer affect convergence when harmonic Ritz values are sufficiently close to these eigenvalues. Using this interpretation tool, we can identify the impact of the small eigenvalues corresponding to the resonances or quasi-resonances on the GMRES convergence.

Then, we use numerical experiments on different configurations to illustrate the convergence difficulties in the context of resonances or quasi-resonances. We investigate the deflation techniques with a suitable choice of (approximated) eigenvectors, adapted to resonance and quasi-resonance benchmark.

→ Back to timetable

▷ M2: Intelligent Preconditioning for Partial Differential Equations

Learning to Smooth: Enhancing Multigrid Solvers with Data

Carola Kruse, Cerfacs, France

Multigrid methods are among the most efficient solvers for elliptic partial differential equations, yet their performance depends critically on the choice of its components, such as the smoothing operators. We propose a hybrid approach to designing data-driven smoothers for multigrid solvers applied to the Poisson equation on structured grids. Exploiting the regularity of the grid, the smoothing operations are defined through convolutional kernels, whose weights can be optimized, or learned.

We investigate the learning of a 9-point smoothing convolutional kernel, possibly varying between the different levels within a V-cycle, enabling fully parallel smoothing operations. Numerical experiments demonstrate that the learned smoothers yield faster convergence in terms of iteration counts compared to classical Gauss-Seidel and optimal weighted Jacobi methods, while keeping the algorithmic multigrid structure.

We discuss the resulting solver's robustness with respect to grid refinement and outline an extension of the approach to unstructured meshes using graph neural networks, opening the door to data-driven multigrid components beyond regular geometries.

This is joint work with Samuel Duarte Dos Santos, Luc Giraud and Paul Mycek.

→ Back to timetable

A Learning-Enhanced Nonlinear Preconditioner for Some Hyperelasticity Problems

Yujie Gong, University of Cologne, Germany

The deformation of the human carotid arterial wall can be modeled by nonlinear hyperelasticity equations whose discretized version can be solved using the inexact Newton method. However, if there is a large atherosclerotic plaque in the artery, the convergence may become very slow. To accelerate the convergence, we propose a novel learning-based nonlinearly preconditioned iterative method that captures and eliminates the subspace associated with the slow convergence with a specially designed targeted quasi-Newton method. Numerical experiments demonstrate that our method significantly outperforms the traditional inexact Newton method by reducing the number of nonlinear iterations and the overall compute time. The parallel scalability is also close to be linear.

→ Back to timetable

Automated Grammar-based Design Of Multigrid Preconditioners With Evolutionary Algorithms

Dinesh Parthasarathy, Friedrich-Alexander-Universität Erlangen-Nürnberg, Germany

Although multigrid is asymptotically optimal for many important PDE problems, its efficiency relies heavily on the careful selection of the individual algorithmic components. In contrast to recent approaches that can optimize certain multigrid components using machine learning techniques, we adopt a complementary strategy, employing evolutionary algorithms to construct efficient multigrid cycles from proven algorithmic building blocks. Here, we will present its application to generate efficient algebraic multigrid methods within the hypre software. The AMG implementation is extended to support flexible cycling, incorporating level-specific smoothing sequences and non-recursive cycling patterns; and these are automatically generated using genetic programming (GP). Numerical experiments demonstrate the potential of these non-standard GP cycles to improve multigrid performance as a preconditioner.

→ Back to timetable

A Polynomial Dimension-Dependence Analysis of Bramble–Pasciak–Xu Preconditioners

Boou Jiang, King Abdullah University of Science and Technology, Saudi Arabia

We investigate the dimension dependence of Bramble–Pasciak–Xu (BPX) preconditioners for high-dimensional partial differential equations and establish that the condition numbers of BPX-preconditioned systems grow only polynomially with the spatial dimension. Our analysis requires a careful derivation of the dimension dependence of several fundamental tools in the theory of finite element methods, including the elliptic regularity, Bramble–Hilbert lemma, trace inequalities, and inverse inequalities.

We further introduce a new quasi-interpolation operator into finite element spaces, a variant of the classical Scott–Zhang interpolation, whose associated constants scale polynomially with the dimension. Building on these ingredients, we prove a multilevel norm equivalence theorem and derive a BPX preconditioner with explicit polynomial bounds on its dimensional dependence. This result has notable implications for emerging quantum computing methodologies: recent studies indicate that polynomial dependence of BPX preconditioners on dimension can yield exponential speedups for quantum-algorithmic approaches over their classical counterparts.

→ Back to timetable

A Data-Enhanced Coarse Preconditioner for Linear System of Equations

Xiao-Chuan Cai, University of Macau, Macau SAR, China

In this talk, we discuss a data-driven enhancement preconditioner for improving the convergence of a preconditioned conjugate gradient method for solving a sequence of symmetric positive definite linear systems of equations. The methodology is designed for the situation that a subset of the systems has been solved and the convergence is considered too slow. In such a situation, data generated from the solved problems are analyzed by an unsupervised learning algorithm, and based on the results of the analysis, one or more enhancement preconditioners are constructed and are additively incorporated into the existing preconditioners to solve other systems in the sequence. Numerical experiments for time-dependent problems demonstrate that the proposed

approach improves the convergence considerably for other systems in the sequence when classical preconditioners are insufficient. This is a joint work with Jing-Yuan Wang.

→ Back to timetable

Machine-Learning-Enhanced Domain Decomposition Preconditioners

Axel Klawonn, University of Cologne, Germany

Domain decomposition methods (DDMs) are robust and parallel efficient iterative solvers for discretized partial differential equations. However, the convergence rate of classic DDMs deteriorates for coefficient distributions with large contrasts. To retain the robustness for such problems, the coarse space of the DDM can be enriched by additional coarse basis functions, often obtained by solving local generalized eigenvalue problems. Within overlapping Schwarz methods, we consider the AGDSW (adaptive generalized Dryja-Smith-Widlund) coarse space to obtain robustness and within nonoverlapping DDMs, related approaches are used in FETI-DP and BDDC methods. However, the computation of the AGDSW, FETI-DP, and BDDC coarse basis functions is computationally expensive due to the solution of many local eigenvalue problems. In this talk, we train a surrogate model based on a deep feedforward regression neural network which directly learns the necessary coarse basis functions. Additionally, we present first results where we also replace the solution of local subdomain problems by surrogate models.

This talk is based on joint work with Martin Lanser and Janine Weber-Hamacher, University of Cologne, Germany.

→ Back to timetable

▷ M3: Advances in Preconditioning Techniques for Wave Propagation Problems

Parallel in Time Preconditioning for Solving the Paraxial Wave Equation in Ptychographic Applications

Tyrone Rees, STFC Rutherford Appleton Laboratory, UK

Ptychography is a lensless imaging technique that allows the high-resolution visualisation of samples. The method works by taking a series of overlapping scans of a specimen and numerically reconstructs an image by assembling the resulting diffraction patterns. Standard ptychography is defined only for thin samples, and robustly applying the technique to thick or heterogeneous media is a challenge.

The state-of-the-art is a multi-slice scheme, which works by advancing the field sequentially slice-by-slice. This works for simple problems, but introduces splitting errors that limit stability as the thickness grows. We present a reformulation of the ptychographic reconstruction in terms of the paraxial wave equation, and will show how this allows the resulting linear solves to use FFT-friendly spectral/shifted preconditioners alongside split variants that better capture coupling in stronger-contrast media.

We will report accuracy–runtime trends relative to conventional multi-slice, and outline paths to higher-order outer integration and adjoint operators for robust thick-sample imaging.

→ Back to timetable

Operator Preconditioning of Multiple Trace Formulations for Maxwell Transmission Problems

Carlos Jerez-Hanckes, Inria Chile, Chile

We consider the three-dimensional time-harmonic Maxwell transmission problem in heterogeneous media and study operator preconditioning strategies for its discretization via the local Multiple Trace Formulation (MTF). In this framework, independent electric and magnetic trace unknowns are assigned to each subdomain, while electromagnetic transmission conditions and Calderón identities are imposed weakly across interfaces. This structure naturally accommodates complex non-overlapping decompositions with junctions and shared edges and is well suited for parallel solvers.

The performance of Maxwell MTFs, however, is strongly affected by the conditioning of the resulting operator, especially as the number of subdomains increases. We investigate block operator preconditioners based on approximate electromagnetic Calderón projectors and On-Surface Radiation Conditions (OSRCs). Numerical experiments demonstrate the robustness of the proposed operator preconditioners with respect to frequency, mesh refinement, and increasing numbers of subdomains, confirming their effectiveness for iterative solution of Maxwell transmission problems.

→ Back to timetable

Deflation-Based Preconditioning for Immersed Finite Element Methods

Yannis Voet, École polytechnique fédérale de Lausanne, Switzerland

Trimming is a ubiquitous operation in computer-aided-design whereby parts of a geometry are merged, intersected or simply discarded. While it grants virtually unlimited flexibility in geometric design, it introduces a plethora of other difficulties when such geometries are used within immersed finite element methods. In particular, small cut elements cause severely ill-conditioned system matrices requiring dedicated penalization, stabilization or preconditioning techniques. In this talk, we explore deflation-based preconditioners for immersed finite element methods in a variety of applications, including wave propagation problems.

→ Back to timetable

Advances in Multigrid Preconditioning for the Helmholtz Equations

Eran Treister, Ben-Gurion University of the Negev, Israel

This talk addresses recent advances in multigrid preconditioners for the Helmholtz equation, with a focus on handling the challenges posed by coarse-grid solution. High-frequency Helmholtz problem requires very fine meshes, and therefore typically solved with a few levels of multigrid. The problem is that even the coarsest grid, in large real-world problems, might be too large for direct solution. We present two complementary approaches to address this problem: the first is designing a special Vanka block-smoother that enables deep V-cycles, allowing the use of extremely coarse grids without sacrificing convergence. The second approach employs a three-level multigrid strategy, implemented in a matrix-free manner, and solves the coarse grid problem using a Lippmann-Schwinger preconditioner. The solution of the Lippmann-Schwinger preconditioner is done by FFT and is therefore blind to sparsity and enables using wide interpolation stencils without degradation. These developments illustrate practical strategies for overcoming multigrid limitations and achieving robust solvers for high-frequency Helmholtz problems.

→ Back to timetable

Multigrid and Dispersion Correction for Preconditioning Helmholtz Equations

Chris Stolk, University of Amsterdam, The Netherlands

In this work we study a two-grid Helmholtz preconditioner, in a setting of regular grid finite elements. In particular we study how the convergence depends on the detailed choice of coarse level discretization. It turns out that (i) there is a strong dependence, and (ii) that the best results are obtained if the coarse level matrix is a discrete Helmholtz operator with accurate dispersion relation, which differs from standard approaches. This is supported by three types of evidence: (1) an analysis in a 1-D toy model relating convergence factors to dispersion errors; (2) results from local Fourier analysis; and (3) numerical results. This for example explains that convergence is much better for higher order finite element discretizations than for low orders (using p-type coarsening). Very good convergence is obtained with optimized finite differences for the coarse level operator.

→ Back to timetable

Comparison of Two Dispersion Correction Methods for Finite-Difference Schemes Applied to the Helmholtz Equation

Pierre-Henri Cocquet, Université de Pau et des Pays de l'Adour, Laboratoire SIAME, France

This talk will present two dispersion correction methods for finite-difference (FD) schemes applied to the Helmholtz equation. The first one is called asymptotic dispersion correction, and is based on the introduction of a shifted wavenumber in the FD stencil that can then be determined in closed form by minimizing the dispersion error. The second technique builds FD stencils from plane-wave interpolation, and we will show how to obtain the interpolating plane-waves that give the least dispersion error. We will then apply both techniques to some specific FD stencils and we will show, through some numerical experiments, that both methods yield very similar results and reduce the relative error compared to standard FD stencils without dispersion correction.

→ Back to timetable

Dispersion Correction Methods for the Convected Helmholtz Equation

Antoine Tonnoir, LMI INSA Rouen Normandie, France

Numerical computations of wavelike problems in time-harmonic regime raises many difficulties, in particular for high frequency problems. It is now well-known that the solution suffers from the pollution effect which imposes to drastically increase the number of points per wavelength (or to use high order methods).

In our work, we are interested in the convected Helmholtz equation which describes the propagation of a sound in a moving fluid. This typically appears in the context of aeronautic. As other wave problems, the resolution of this equation suffers from the same difficulties. Yet, in the case of homogeneous media, there exists a strong connexion between Helmholtz equation and the convected Helmholtz equation (via the Lorenz transformation). This allows to consider usual techniques for the Helmholtz equation and to apply them to the convected one. In particular, we are interested in developing dispersion correction methods for finite difference discretisation of the convected Helmholtz equation. We will present two approaches for performing dispersion correction in that case.

This work is a joint work with Pierre-Henri Cocquet (from of University of Pau) and Martin J. Gander (from the University of Geneva).

→ Back to timetable

▷ M4: Preconditioned Iterative Methods for Discretised PDEs

Fast and Accurate Simulation of Close-to-Touching Discs in Stokes Flow

Daniel Fortunato, Flatiron Institute, USA

Consider the flow of a dense suspension of rigid bodies in a Stokesian fluid. Such flows are difficult to compute numerically due to the presence of close-to-touching interactions between the bodies, which may require a large number of unknowns to discretize, a large number of GMRES iterations to solve, and an extremely small time step. In this talk, we present a fast and accurate boundary integral method that mitigates these challenges without introducing artificial repulsion forces. Through precomputation, compression, and interpolation of the close-to-touching part of the interaction operator, our method—termed “interpolated compressed inverse preconditioning”—efficiently handles close-to-touching interactions down to distances of $1e-10$ with only a coarse discretization of the boundary. Additionally, we present a preconditioner that significantly reduces the number of GMRES iterations required to solve the Stokes mobility problem at each time step by effectively reusing the Krylov subspace from previous time steps. Coupled with high-order, adaptive time-stepping using spectral deferred correction, we are able to take larger time steps, mitigating the temporal stiffness resulting from close-to-touching interactions.

→ Back to timetable

Mixed-Precision Parallel-in-Time Solvers for Runge–Kutta Methods

Santolo Leveque, Charles University, Czech Republic

Time-dependent PDEs arise quite often in many scientific areas, such as mechanics, biology, economics, or chemistry, just to name a few. Of late, researchers have devoted their effort in devising parallel-in-time methods for the numerical solution of time-dependent PDEs, adding a new dimension of parallelism and allowing to speed-up the solution process on modern supercomputers.

In this talk, we present a parallel-in-time preconditioner for the all-at-once linear system arising when employing a Runge–Kutta method in time. We apply the preconditioner in mixed-precision arithmetic. The resulting system is solved iteratively for the numerical solution and for the stages of the method. The proposed preconditioner results in a block-diagonal solve for the stages, and a Schur complement obtained by solving again systems for the stages.

A range of numerical experiments validate the robustness of the preconditioner with respect to the discretization parameters and to the number of stages, showing the speed-up achieved by mixed precision arithmetic.

→ Back to timetable

Parallel-in-Time Preconditioning for Time-Dependent Variational Mean Field Games as Generative Models

Heidi Wolles Ljósheim, University of Edinburgh, UK

We present a parallel-in-time preconditioned iterative framework for time-dependent variational mean field games solved with the Chambolle–Pock primal–dual algorithm. To efficiently handle the ill-conditioned linear systems arising at each iteration, we employ a class of parallel-in-time preconditioners based on temporal diagonalisation via fast Fourier transforms. For structured grids, fast recursive solvers are able to exploit the tensor-product structure of the discrete differential operators, while allowing for more general geometries. Numerical experiments demonstrate the robust solution of very large-scale mean field game systems, including low-viscosity regimes, offering improved efficiency and parallel scalability. Recent work has shown that many generative models can be expressed as time-dependent variational mean field games, providing a stepping stone towards explainable AI. We demonstrate this with a number of popular examples from machine learning.

→ [Back to timetable](#)

▷ M5: Challenges in Preconditioning for Large Scale Least Squares Problems

The Potential Importance of Reorthogonalisation for Least Squares Problems

Melina Freitag, University of Potsdam, Germany

We consider an overdetermined linear least-squares (LS) problem in a large scale setting, where the system matrix is only accessible via a matrix-vector product. Krylov subspace methods are commonly used to solve such problems. Popular choices include CGLS, LSQR, and LSMR, all of which rely on short-term recurrences. The convergence of these methods is highly problem dependent. Preconditioning is often required to accelerate convergence, but constructing effective preconditioners for LS problems is challenging and may still lead to slow convergence. In applications such as data assimilation, only a limited number of iterations can be afforded due to time constraints, motivating alternative acceleration strategies. In this talk, we investigate the use of reorthogonalisation within Krylov methods for LS problems. Although this increases memory usage and per-iteration cost, these overheads may be offset when matrix-vector products with the system matrix dominate the computational expense. We demonstrate the potential gain of using reorthogonalisation in this setting and support our findings by numerical examples. This is joint work with I. Daužickaitė, S. Gürol, A. Lawless, A. Ramage, J. Scott, and J. Tabcart.

→ Back to timetable

Backward Stable (Randomized) Preconditioning for Least Squares Problems

Maïke Meier, University of Groningen, The Netherlands

The finite-precision behavior of short-recursion Krylov subspace methods for least squares problems, such as CGLS and LSQR, is elusive in the best of times. When a preconditioner is included, the observed convergence behavior is even more surprising. In the sketch-and-precondition algorithm for highly-overdetermined least squares problems, a very accurate randomized preconditioner is formed. Despite the quality of this preconditioner, this algorithm is not backward stable. In this talk, we discuss an adaptation to the algorithm, called SPIR, to provably ensure backward stability. Furthermore, we explore more generally how preconditioning changes the finite-precision behavior of Krylov subspace methods for least squares problems.

→ Back to timetable

A Spectral Preconditioner for the Conjugate Gradient Method with Iteration Budget

Selime Gürol, CERFACS, France

In this talk, we address the solution of large symmetric positive-definite linear systems in a matrix-free setting with a limited iteration budget. We focus on the preconditioned conjugate gradient method with spectral preconditioning. Spectral preconditioners map a subset of eigenvalues to a positive value via a scaling parameter, and leave the remaining part of the spectrum unchanged, with the goal of faster convergence. We formulate the design of the spectral preconditioners as a constrained optimization problem. The optimal cluster placement is defined to minimize the error in energy norm at a fixed iteration. We propose three practical strategies for selecting the scaling parameter, hence the cluster position, that incur negligible computational cost. Numerical experiments highlight the importance of cluster placement and demonstrate significant improvements in terms of error in energy norm, particularly during the initial iterations.

→ Back to timetable

Error Estimate and Stopping Criteria for (Preconditioned) Least-Squares Problems Solved by CGLS, LSQR, and Their Block Variants

Jan Papež, Czech Academy of Sciences, Czech Republic

In [Meurant, Papež, Tichý; Numerical Algorithms 88, 2021], we introduced an adaptive estimate for the energy norm of the error in the conjugate gradient (CG) method. This estimate is computationally inexpensive and numerically reliable in finite-precision computations.

In this talk, we present the results from [Papež, Tichý; Numerical Algorithms 97, 2024], where we have extended the estimate to the LSQR and CGLS algorithms for solving least-squares problems involving general rectangular matrices. The estimate is applicable to preconditioned variants of these algorithms and can be seamlessly integrated into existing codes. The extension of the results to block variants of LSQR and CGLS was recently presented in [Meurant, Papež, Tichý; BIT Numerical Mathematics 66, 2026] and will also be discussed.

We emphasize the applicability of the estimates in finite-precision arithmetic, noting that their derivation relies exclusively on local orthogonality, a property generally well preserved in computations. Additionally, the evaluation of the estimates remains inexpensive. Finally, we explore their use in stopping criteria.

→ Back to timetable

Low Precision Incomplete Factorization Preconditioners for Linear Least-Squares Problems

Jennifer Scott, STFC Rutherford Appleton Laboratory & University of Reading, UK

In recent years, hardware developments have led to a surge in interest in exploiting mixed precision arithmetic within numerical linear algebra algorithms to take advantage of potential savings in memory requirements, runtime and energy use, whilst still achieving the requested accuracy. Our interest lies in the robust and efficient solution of large sparse linear least-squares problems. We explore employing mixed precision when solving such problems, focusing on the practicalities of developing robust approaches using general-purpose low precision incomplete Cholesky factorization preconditioners. Using problems from practical applications, we show that half precision

arithmetic can be considered if high accuracy is not required in the solution or the memory for the incomplete factors is very restricted; otherwise, single precision can be used, and double precision accuracy recovered while reducing memory consumption, even for ill-conditioned problems.

→ Back to timetable

Mixed Precision Iterative Refinement – Solving Least Squares with More Least Squares

Andrew Lister, STFC Rutherford Appleton Laboratory, UK

When solving the least squares problem with iterative methods, the dominant cost is from matrix-vector products. Computing the majority of these matrix-vector products in a lower precision may lead to improved performance. In this talk we present a new algorithm for solving the least squares problem using iterative refinement. At each iteration, subproblems are solved using low precision, the least squares solution is updated using higher precision, and the residual is updated in high precision. The subproblems are two least squares problems (one is overdetermined and the other is underdetermined) that are well suited to preconditioned iterative methods where early stopping is possible, as each subproblem only requires an approximate solution. We show conditions on the stability of the algorithm and to what accuracy the subproblems must be solved. Numerical experiments show favourable performance when compared to alternative mixed precision iterative refinement techniques.

→ Back to timetable

▷ M6: Preconditioning Across Scientific Computing and Learning

Neural Approximate Inverse Preconditioners

Yuanzhe Xi, Emory University, USA

In this talk, we present a data-driven framework for constructing efficient approximate inverse preconditioners for elliptic partial differential equations (PDEs) by learning the Green’s function of the underlying operator with neural networks (NNs). Once trained, the NN-approximated Green’s function is directly compressed into either a hierarchical matrix or a sparse matrix—using only the mesh geometry and the network output. This geometric construction achieves nearly linear complexity in both setup and application while preserving the spectral properties essential for effective preconditioning. Numerical experiments on challenging elliptic PDEs demonstrate that the resulting preconditioners consistently yield fast convergence and small iterations.

→ Back to timetable

Preconditioning Via Spectral Density Driven Graph Neural Networks

Francesco Brarda, Emory University, USA

We propose a learning-based framework for sparse approximate inverse preconditioners that promotes spectral clustering for fast Krylov convergence via two core components. First, we introduce a Line-Graph with Virtual Node (LG-VN) architecture that transforms the matrix adjacency graph into its line graph, enabling message passing directly on algebraic entries. A virtual node provides global context at constant communication depth, improving expressivity for long-range interactions. Second, we approximate the spectral density of the preconditioned operator and penalize spectral mass away from unity, encouraging eigenvalue clustering. To stabilize optimization with nonconvex spectral objectives, we adopt a two-stage curriculum: Frobenius warm-up for a robust initial factorization, followed by spectral fine-tuning. For SPD systems, we learn a factorized form; for nonsymmetric ones, we employ a dual-head architecture to predict asymmetric factors, while ensuring inference uses only sparse matrix–vector products. Experiments on diffusion–reaction, Helmholtz, and convection–diffusion problems show consistent reductions in Krylov iteration counts and improved robustness over to classical factorized sparse approximate inverse methods.

→ Back to timetable

Design Criteria for SGD Preconditioners: Local Conditioning, Noise Floors, and Basin Stability

Mitchell Scott, Emory University, USA

Stochastic Gradient Descent (SGD) often slows in the late stage of training due to anisotropic curvature and gradient noise. We analyze preconditioned SGD in the geometry induced by a symmetric positive definite matrix M . Our bounds make explicit how both the convergence rate and the stochastic noise floor depend on M . For nonconvex objectives, we establish a basin-stability

guarantee in a local M -metric neighborhood around a minimizer set: under local smoothness and a local PL condition, we give an explicit lower bound on the probability that the iterates remain in the basin up to a time horizon. This perspective is particularly relevant in Scientific Machine Learning (SciML), where reaching small training losses under stochastic updates is closely tied to physical fidelity, numerical stability, and constraint satisfaction. Our framework covers both diagonal/adaptive and curvature-aware preconditioners and yields a practical criterion: choose M to improve local conditioning while attenuating noise in the M^{-1} -norm. Experiments on a quadratic diagnostic and three SciML benchmarks support the predicted rate–floor behavior.

→ Back to timetable

Preconditioned Stochastic Gradient Methods for Ptychographic Reconstruction

Zichao Wendy Di, Argonne National Laboratory, USA

Ptychographic phase retrieval is a highly ill-conditioned, nonlinear inverse problem in which an unknown object is recovered from overlapping diffraction measurements. In practice, simple projection-based schemes such as PIE and its variants are remarkably effective—especially when updates are applied one frame at a time or with very small mini-batches. In this talk, I reinterpret these algorithms as preconditioned stochastic gradient methods and use a spectral viewpoint to explain how the batch size and sampling strategy can affect PIE convergence.

→ Back to timetable

A Stochastic View of Preconditioning with Applications to Stochastic Learning

Difeng Cai, Southern Methodist University, USA

We introduce a stochastic view of preconditioning and practical estimators based on stochastic Krylov methods for accelerating stochastic learning. The theory reveals connection between preconditioning in matrix computations and variance reduction in stochastic algorithms. The proposed stochastic estimators (termed Preconditioned Truncated Single-Sample (PTSS) estimators) offer reduced variance compared to commonly used unbiased estimators. The new methodology can be used to accelerate various matrix operations such as linear system solutions, log-determinant, as well as parameter-gradient calculations in stochastic learning. Experiments show improved computational efficiency as well as convergence behavior compared to existing methods.

→ Back to timetable

▷ M7: Quantum Computing Meets Numerical Analysis

The Practical Problems of Quantum Linear System Solvers and a Krylov-based Solution

Matthias Deiml, University of Augsburg, Germany

In this talk we consider linear solvers for quantum computers, whose historic development largely ignored classical numerical linear algebra. This led to non-adaptive solvers, similar to the Chebyshev iteration, which have to be hand-tuned to the individual equation in order to be practically useful. We propose to fix this by finding approximations in Krylov subspaces based on estimated spectral information.

→ Back to timetable

Multilevel Preconditioning for PDEs on Quantum Computers

Daniel Peterseim, University of Augsburg, Germany

This talk presents a quantum algorithm for efficiently solving linear systems that arise from the finite element discretization of elliptic partial differential equations. The key ingredient is a multilevel preconditioner of BPX type, which transforms the system into one with uniformly bounded condition number. This makes it suitable for quantum linear system solvers. For any fixed spatial dimension, the algorithm computes linear functionals of the solution with rate-optimal complexity proportional to the inverse of the error tolerance, up to logarithmic terms. The method does not require smooth solutions and can handle rough coefficients, including random inputs. We demonstrate the feasibility of this approach through implementations on quantum circuit simulators and runs on current quantum hardware.

→ Back to timetable

Investigating Preconditioning for Quantum Linear System Solvers

Elise Fressart, École polytechnique & Thales, France

Quantum algorithms can speedup linear system solving under certain conditions. As in the classical framework, the performance of these algorithms is affected by the condition number. Preconditioned quantum linear system solvers have received growing attention. We investigate quantum preconditioning for the Poisson equation based on multilevel and domain decomposition methods. This is joint work with Michel Nowak and Nicole Spillane.

→ Back to timetable

A Quantum-Classical Mixed Precision Algorithm for Solving Linear Systems

Marc Baboulin, Université Paris-Saclay, France

The existing quantum algorithms for solving linear systems require huge quantum resources if we want to achieve an accuracy that is acceptable for HPC applications. We propose a hybrid quantum-classical algorithm that reduces this “quantum” cost by adding iterative refinement in mixed-precision. Our algorithm computes a first solution with a quantum processor using the Quantum Singular Value Transformation, in low precision, and then refines in higher precision with a classical processor until we get a satisfying accuracy. For this solver, we present an error and complexity analysis, a communication scheme, and first experiments using the quantum software stack myQLM.

→ [Back to timetable](#)

▷ M8: Preconditioning Strategies for Large-Scale Data Assimilation

Multi-fidelity Preconditioner for Ocean Variational Data Assimilation

Hélène Hénon, Inria de l'Université Grenoble Alpes, France

Advances in computational power have enabled the use of increasingly precise ocean models, but their high cost makes advanced data assimilation algorithms difficult to use in real-time forecasting systems. This issue motivates the need for data assimilation techniques using cheaper ocean models to reduce the cost of data assimilation without sacrificing accuracy.

This study investigates how multi-fidelity approaches can reduce the computational cost of variational data assimilation algorithms. We consider two models: a high-fidelity model that accurately represents ocean dynamics, but is computationally expensive, and a low-fidelity model that is less accurate but significantly cheaper to run.

On top of the traditional use of the background error covariance matrix B as a preconditioner, we propose to use the low-fidelity model to construct a second-level limited-memory preconditioner. Using a realistic ocean test case, we discuss the influence of the trade-off between computation cost and accuracy of the low fidelity model on the total gain in efficiency. With the right choice of low-fidelity model, this multi-fidelity preconditioning approach can significantly accelerate convergence compared to using B alone.

→ Back to timetable

Randomized Preconditioning for an Ensemble of Data Assimilations

Ieva Daužickaitė, CERFACS, France

An ensemble of data assimilations (EDA) is used for uncertainty quantification in numerical weather prediction. This method requires solving an ensemble of perturbed large-scale symmetric positive definite systems of linear equations in parallel. The solutions can be obtained via the conjugate gradient (CG) method using a limited memory preconditioner (LMP). The LMP is constructed using approximate leading eigenpairs of the coefficient matrix. A tractable approach for attaining the eigenpairs is computing a randomized Nystrom approximation of one of the coefficient matrices. The same LMP is then used to precondition all of the perturbed systems. Computing matrix-vector products with the coefficient matrix in the EDA is highly expensive, and thus a single-pass Nystrom algorithm has to be used. In this talk, we show how the underlying EDA structure enables the construction of a sketching matrix that captures information about the subspace spanned by the leading eigenvectors. This sketching matrix gives eigenvalue approximation of quality comparable to when using one iteration of the power method, but without the need for computing the extra matrix-vector products. Numerical tests confirm the effectiveness of the resulting LMP.

→ Back to timetable

The Role of Rounding Errors on Spectral-Limited Memory Preconditioner

Hisham Elzayyadi, Eindhoven University of Technology, The Netherlands

Rounding errors, often overlooked in iterative solvers like Conjugate Gradient (CG), are introduced and intensified by preconditioning, compromising convergence. We present a full error analysis of the Spectral Limited Memory Preconditioner (Spectral-LMP) within the preconditioned CG (PCG) framework, deriving error bounds. This analysis allows us to strategically select the Spectral-LMP parameter to minimize rounding error, demonstrating how this choice is fundamentally influenced by the distribution of the right-hand side relative to the system's top eigenvector directions. Our findings are validated through numerical experiments motivated by applications in Data Assimilation, showcasing the resulting improvements in robust PCG convergence.

→ Back to timetable

▷ M9: New Results in Domain Decomposition

Parallel Preconditioning for Time-Harmonic Waves from Toeplitz Structured Sweeping

Niall Bootland, STFC Rutherford Appleton Laboratory, UK

We consider the potential for parallelism within an approach to precondition problems involving high-frequency time-harmonic waves. We take inspiration from domain decomposition strategies known as sweeping methods, which have gained notable interest for their ability to yield nearly-linear asymptotic complexity and which can also be favourable for high-frequency problems. While successful approaches exist, such as those based on higher-order interface conditions, perfectly matched layers (PMLs), or complex tracking of wave fronts, they can often be quite involved or tricky to implement and inherently sequential in nature. We introduce here parallelisable versions of simple sweeping techniques for the Helmholtz equation that allow for each subdomain to be solved in parallel with minimal impact on the overall convergence. We further link to similar possibilities in very high-frequency applications, such as thick sample ptychography, using a simplified time-harmonic forward model.

→ Back to timetable

A Stochastic BDD strategy for Stochastic Elliptic Problems

Lucien Gontier, CNRS, France

This work presents a hybrid method combining Balancing Domain Decomposition (BDD) with Polynomial Chaos Expansions (PCE) to accelerate Monte Carlo simulations for an elliptic PDE with a random coefficient. We develop PCE-based local approximations of pseudo-inverses of local Schur complement operators within the BDD framework, creating sample-dependent preconditioners that maintain optimal convergence properties. For the coarse space, we introduce a median-based GenEO approach, avoiding sample-dependent eigenvalue computation while preserving robustness.

Numerical experiments demonstrate the effectiveness and scalability of our method, particularly for problems with small subdomains. The approach reduces the overall computational cost by pre-computing the PCE coefficients and the Median-based GenEO, thus reducing the per-sample overhead cost of our Monte Carlo realization.

Keywords: Balancing Domain decomposition, Stochastic PDE, Monte Carlo method, Polynomial Chaos Expansion, Median-based GenEO.

This is joint work with: Olivier Le Maitre, CNRS, CMAP, Inria; Paul Mycek, CERFACS, Inria; Nicole Spillane, CNRS, CMAP

→ Back to timetable

A Flow-Rate Conserving CNN-Based Domain Decomposition Method for Blood Flow Simulations

Janine Weber-Hamacher, University of Cologne, Germany

In this talk, we consider the prediction of blood flow with non-Newtonian viscosity in stenosed arteries using convolutional neural network (CNN) surrogate models. An alternating Schwarz domain decomposition method is proposed which uses CNN-based subdomain solvers. A universal subdomain solver (USDS) is trained on a single, fixed geometry and then applied for each subdomain solve in the Schwarz method. Results for two-dimensional stenotic arteries of varying shape and length for different inflow conditions are presented and statistically evaluated. One key finding, when using a limited amount of training data, is the need to implement a USDS which preserves some of the physics, as, in our case, flow rate conservation. A physics-aware approach outperforms purely data-driven USDS, delivering improved subdomain solutions and preventing overshooting or undershooting of the global solution during the Schwarz iterations, thereby leading to more reliable convergence. The presented results are based on joint work with S. Klaes (University of Cologne), A. Klawonn (University of Cologne), N. Kubicki (University of Cologne), M. Lanser (University of Cologne), K. Nakajima (University of Tokyo), and T. Shimokawabe (University of Tokyo).

→ Back to timetable

Multiprecision Computations with Schwarz Methods

Daniel Szyld, Temple University, USA

We explore and analyze the use of multiprecision arithmetic for several classes of Schwarz methods and preconditioners, where the approximate solution of the local problems is performed at a lower precision, i.e., with fewer digits of accuracy than in the underlying (double precision) computation. Conditions for the appropriate round-off criteria for the lower precision are presented. It is found experimentally that for the model problems about 5 digits of accuracy are sufficient to achieve the theoretical restrictions, and thus, single precision suffices for the local solves. Several numerical experiments illustrate the obtained results. (Joint work with Michal Outrata, Charles University, Prague)

→ Back to timetable

▷ M10: Advanced Preconditioners for Large-Scale Linear Systems in HPC

Extrapolated Multigrid Methods for Elliptic PDEs on Tokamak Geometries

Sai Aakash Dasari, CERFACS, France

Extrapolated multigrid methods use coarse-grid information to improve both the accuracy of the discrete solution (differential convergence) and the efficiency of the solution process (algebraic convergence). For elliptic finite element systems on polygonal domains, this approach is equivalent to a defect-correction multigrid scheme with a higher-order discretization when using specific quadrature rules.

We extend this methodology to domains with curved boundaries, such as the cross-sections of tokamak reactors. Initial numerical experiments using extrapolation with linear finite elements on a non-nested grid hierarchy show accuracies of the same order as quadratic isoparametric elements.

This approach avoids the implementation difficulties of optimized higher-order isoparametric finite element codes while preserving multigrid efficiency and providing improved solutions on non-polygonal domains.

→ Back to timetable

Preparing deal.II for Exascale with PSCToolkit Preconditioners

Fabio Durastante, Università di Pisa, Italy

The Dealii-X EuroHPC project aims to enable exascale-ready multiphysics simulations by extending the deal.II finite element library with scalable algorithms and software abstractions suited for emerging HPC architectures. A key bottleneck in this context is the solution of large, sparse linear systems, where robust and scalable preconditioning is essential for performance and algorithmic robustness. This talk presents the integration of the PSCToolkit suite of parallel sparse solvers and preconditioners into deal.II's linear algebra framework. PSCToolkit provides high-performance building blocks for sparse matrix factorization and preconditioning, designed to scale efficiently on massively parallel systems. We describe the software design choices enabling interoperability between deal.II's distributed data structures and PSCToolkit's solver interfaces, as well as the extension of deal.II's solver infrastructure to support advanced preconditioning strategies. Performance results from representative Dealii-X applications demonstrate improved scalability and convergence behavior on large-scale systems. This integration represents a significant step toward exascale-capable preconditioning workflows within deal.II.

→ Back to timetable

Aggressive Coarsening and Energy Minimization for Lightweight AMG in CFD Problems

Carlo Janna, Università di Padova, Italy

The solution of large and sparse linear systems of equations remains a major computational bottleneck in computational fluid dynamics. Although Algebraic Multigrid (AMG) methods exhibit excellent weak scalability, they are often burdened by high memory requirements and expensive setup phases. This work presents an aggressive coarsening framework specifically designed for GPUs and implemented within the Chronos library. The proposed approach reduces both memory footprint and setup cost by combining long-distance interpolation strategies with energy minimisation to improve the prolongation. Numerical experiments show that aggressive coarsening leads to lightweight yet robust AMG preconditioners. Large-scale CFD benchmarks demonstrate significant speed-ups and notable memory savings when compared with standard AMG configurations.

→ Back to timetable

Robust Algebraic Preconditioning in PETSc with Domain Decomposition Methods

Pierre Jolivet, Sorbonne Université, CNRS, LIP6, France

Domain decomposition methods provide a powerful algebraic framework to address large-scale linear systems, particularly in parallel and heterogeneous computing environments. In this talk, we present recent advances in algebraic preconditioning within PETSc, with a specific focus on new methods implemented in the PCHPDDM preconditioner.

We discuss how these developments leverage modern domain decomposition techniques to improve robustness with respect to problem parameters and coefficient contrasts, while maintaining strong parallel scalability. Emphasis is placed on the design of coarse spaces, adaptive enrichment strategies, and flexible solver compositions that can be applied in a fully algebraic setting, without requiring intrusive problem-specific modifications. Through numerical experiments and practical examples, we illustrate the performance and robustness of the new PCHPDDM capabilities on challenging non-symmetric and saddle-point problems. The talk aims to provide both an overview of the underlying algorithms and concrete guidance for PETSc users interested in deploying advanced domain decomposition preconditioners in large-scale simulations.

→ Back to timetable

A Reverse Augmented Constraint Preconditioner for Industrial Contact Mechanics Applications

Andrea Franceschini, University of Padova, Italy

Frictional contact is a challenging problem in computational mechanics due to its strong non-linearity and the difficulties it introduces in both nonlinear and linear solvers. When enforced via Lagrange multipliers, contact conditions are satisfied exactly, but the resulting Jacobian is indefinite and requires specialized solution strategies. This talk revisits the Reverse Augmented Constraint Preconditioner (RACP), where the primal Schur complement is computed after augmenting the zero block, in contrast to standard approaches that augment only the structural block. The method effectively handles problems characterized by singular structural blocks, which often arise in contact mechanics, while being significantly cheaper than traditional constraint preconditioning for this class of problems and remaining well suited for efficient HPC implementation.

Numerical results on mid- and large-scale problems demonstrate the efficiency of the approach. In addition, new techniques for automating solver configuration will be presented, moving toward a black-box strategy based on RACP that minimizes user knowledge and parameter tuning.

→ [Back to timetable](#)

▷ M11: Preconditioning for Neural Networks and Neural Networks for Preconditioning

Local Feature Filtering for Scalable and Well-Conditioned Domain-Decomposed Random Feature Methods

Victorita Dolean, Eindhoven University of Technology, The Netherlands

Random Feature Methods (RFMs), including extreme learning machine finite-basis PINNs (ELM-FBPINNs), offer a scalable approach to PDEs by approximating solutions with localized, overlapping random neural bases and training via structured least-squares minimization of PDE residuals. While efficient and parallelizable, the resulting systems are often severely ill-conditioned due to redundancy and overlap-induced correlations. We introduce a block rank-revealing QR (RRQR) filtering and preconditioning strategy applied directly to the structured least-squares problem. Local RRQR factorizations remove redundant basis functions, reduce problem size, and improve conditioning, while enabling a block-sparse, numerically stable right preconditioner. We derive deterministic condition-number bounds with probabilistic refinements. Experiments on multi-scale PDEs in 1D–3D show condition-number reductions up to eleven orders of magnitude, LSQR speedups of 10–1000 \times , and higher accuracy than unpreconditioned and additive Schwarz methods at lower cost.

→ Back to timetable

Geometry-Aware Neural Preconditioning for Parametrized PDEs

Paolo Zunino, Politecnico di Milano, Italy

We present a neural-network-based framework for learning preconditioners for parameter-dependent linear systems arising from PDE discretizations, with a focus on ill-conditioned mixed-dimensional problems. Our approach treats preconditioning as a nonlinear, matrix-free learning task, embedding neural preconditioners directly within Krylov subspace solvers while preserving the full accuracy of the underlying discretization.

We first introduce an unsupervised training strategy based on residual minimization, which enables neural preconditioners to generalize across geometric configurations and mesh resolutions without retraining. We then develop a solver-aware learning paradigm in which the training objective is explicitly aligned with Krylov solver dynamics. By differentiating through the Flexible GMRES algorithm, we construct a geometry-aware loss based on the angles between residuals and Krylov subspaces, directly targeting early-stage convergence.

Numerical results demonstrate improved robustness, convergence rates, and generalization across parameter variations. Overall, this work illustrates how neural networks can bridge numerical linear algebra and scientific machine learning.

→ Back to timetable

Accelerating FGMRES for Helmholtz Problems via Learned Non-linear Preconditioning

Paul Mycek, Cerfacs, France

This work proposes a hybrid numerical approach for solving linear systems arising from the discretization of the two-dimensional parametric Helmholtz equation. A neural network is trained in a self-supervised manner, using a loss function based on the residual norm, to approximate the inverse of the discretized operator. The trained network is employed as a nonlinear preconditioner within the Flexible GMRES (FGMRES) algorithm. Numerical experiments show that, although the network is not sufficiently accurate as a standalone solver, it significantly accelerates FGMRES convergence when used as a nonlinear preconditioner. The neural preconditioner demonstrates robustness and generalization with respect to variations in the velocity field and domain size. Comparisons with classical algebraic preconditioners based on sparsified versions of the discretised operator, which can be made arbitrarily good, are reported and discussed. The approach is PDE-agnostic and can be extended to other parametric PDEs.

→ Back to timetable

Nonlinear Preconditioning Algorithms with Learning Capability

Li Luo, University of Macau, Macau SAR, China

Nonlinear preconditioning algorithms have been applied successfully for solving nonlinearly difficult partial differential equations by identifying and balancing the nonlinearities in the system. One of the challenging tasks when applying the methods is to identify the unbalanced nonlinearities. To address this, we leverage learning-based approaches that identify the bad behavior of a Newton solver from the nonlinear residual and error data collected during the iteration. We use principal component analysis to encode these datasets into low-dimensional latent representations, and propose two strategies to learn an effective mapping between them in this reduced space. This mapping forms the core of a nonlinear preconditioner that produces improved initial guesses. Numerical experiments demonstrate that our learning-based algorithms are more robust and more efficient than existing nonlinear solvers.

→ Back to timetable

AI-Driven Algebraic Multigrid: Graph Neural Networks for Operator Design

Rui Peng Li, Lawrence Livermore National Laboratory, USA

We develop ML and data-driven strategies to accelerate algebraic multigrid (AMG) by learning the key operators. Using graph neural networks (GNNs) and reinforcement learning (RL), we model sparse matrices as graphs and optimize AMG coarsening directly from data. Our results demonstrate that neural operators can reproduce and potentially improve AMG components using LLNL's expertise in scalable PDE solvers and scientific machine learning.

→ Back to timetable

▷ M12: Advances in Preconditioning for (Multiple) Saddle Point Linear Systems

Scalable Augmented Lagrangian Preconditioning of Fictitious Domain Problems

Michele Benzi, Scuola Normale Superiore, Italy

We present preconditioning techniques to solve linear systems of equations with a block two-by-two and three-by-three saddle point structure arising from finite element discretizations of the fictitious domain method with Lagrange multipliers. In particular, we propose augmented Lagrangian-based preconditioners to accelerate the convergence of iterative solvers for such classes of linear systems. We consider two relevant examples to illustrate the performance of these preconditioners when used in conjunction with flexible GMRES: the Poisson and the Stokes fictitious domain problems. A spectral analysis is established for both exact and inexact versions of the preconditioners. We show the effectiveness of the proposed approach and the robustness of our preconditioning strategy through extensive numerical tests in both two and three dimensions obtained using the C++ finite element library Deal.II.

This is joint work with Luca Heltai, Marco Feder, and Federica Mugnaioni.

→ Back to timetable

Modified Augmented Lagrangian Preconditioners for Fictitious Domain Formulations of Elliptic Interface Problems

Federica Mugnaioni, Scuola Normale Superiore, Italy

The efficient numerical solution of partial differential equations modeling interactions across interfaces with complex, possibly moving geometries is crucial in many scientific applications. An effective approach to address such problems is the Fictitious Domain with Distributed Lagrange Multipliers (FD-DLM) formulation, which extends the Fictitious Domain method with boundary-supported Lagrange multipliers used in the co-dimension one setting to the co-dimension zero case. While FD-DLM effectively handles complex geometries, it yields large and often ill-conditioned linear systems. To tackle this issue, we extend the augmented Lagrangian (AL) strategy previously proposed for the co-dimension one case to a co-dimension zero elliptic interface problem, where the equation coefficients may jump across the interface between different subdomains. The method is robust with respect to mesh refinement and possibly large coefficient jumps. We also introduce a computationally cheaper block triangular variant, the modified AL preconditioner, which is easier to implement while preserving low iteration counts. Numerical experiments performed using the C++ finite element library deal.II confirm the robustness and efficiency of the approach.

→ Back to timetable

Truncated Domain Preconditioner for the Implicit-in-Time Immersed Boundary Method

Gautam Luhana, University of British Columbia, Canada

The Immersed Boundary (IB) method is a popular method for simulating fluid-solid interactions. Implicit schemes for the IB method have been around for a while but the real challenge has been to develop a scheme that is computationally competitive with the explicit schemes. Two common approaches for the implicit IB method are to either eliminate the Lagrangian variables and solve the fully Eulerian system or vice versa. The first approach leads to a familiar 2×2 block Saddle-Point system that is difficult to solve efficiently due to the Eulerian structural force term. We propose an alternative approach of solving the full 3×3 block Double Saddle-Point system and present efficient and scalable preconditioners based on recent work on Double Saddle-Point systems. The preconditioners are based on *LDU* block factorizations and we present efficient approximations for the resulting Schur Complements.

→ Back to timetable

Eigenvalue Bounds for Symmetric Multiple Saddle-Point Systems Preconditioned with Block-Diagonal Schur Complement matrices

Ángeles Martínez Calomardo, University of Trieste, Italy

Eigenvalue bounds are developed for symmetric, block tridiagonal multiple saddle-point linear systems preconditioned with block diagonal Schur complement matrices. The results extend those for 3 by 3 block systems by Bradley and Greif (2023) and for 4 by 4 systems by Pearson and Potschka (2024) to systems with an arbitrary number of blocks. Furthermore, these findings generalize the bounds in Sogn and Zulehner (2019), developed for block tridiagonal multiple saddle-point linear systems with an arbitrary number of blocks but with null diagonal blocks. An extension of the bounds is also provided for cases in which the Schur complements are approximated. We show that the eigenvalues of the preconditioned matrices are described in terms of the extremal roots of a sequence of parametric polynomials defined by recurrence, and give a practical procedure to compute them.

Numerical experiments validating our findings will be presented.

→ Back to timetable

Block Triangular Preconditioning of Symmetric, Multiple Saddle-Point Matrices: Eigenvalue Bounds with Inexact Schur Complements

Luca Bergamaschi, University of Padova, Italy

We will first present a class of inexact block triangular (BT) preconditioners for double saddle-point (SP) symmetric systems, arising from the mixed (hybrid) FE discretization of Biot's poroelasticity equations. The preconditioned matrix is shown to have complex eigenvalues lying in a circle of center (1,0) and radius smaller than 1. By contrast, the real eigenvalues are described in terms of the roots of a third-degree polynomial with real coefficients [1]. Some numerical results onto realistic models confirm the quality of the theoretical bounds.

Next, we will develop eigenvalue bounds for symmetric, block tridiagonal multiple SP linear systems, with BT preconditioners, based on approximate Schur complements [2]. We show original

bounds for the truly complex eigenvalues, and how the real ones are characterized by the roots of a sequence of polynomials defined by recurrence.

References:

[1] L. Bergamaschi, M. Ferronato, A. Martinez, BT preconditioners for double SP linear systems arising in the mixed form of poroelasticity equations, *SIAM J. Matrix Anal. Appl.*, 47 (2026), pp. 132-157.

[2] L. Bergamaschi, M. Bergamaschi, J. Pearson, Eigenvalue bounds for preconditioned symmetric multiple SP matrices with BT preconditioners.

→ Back to timetable

Exploiting Non-Standard Inner Products for Saddle Point Problems

Joerg Liesen, TU Berlin, Germany

Saddle point matrices are frequently symmetric, enabling mathematically optimal and numerically efficient iterative solvers like MINRES. Standard preconditioners (e.g., block triangular) often sacrifice symmetry for other seemingly desirable properties like eigenvalue clustering, forcing the use of non-optimal iterative methods or methods without short recurrences. Non-standard inner products are tool to retain symmetry-like properties after preconditioning. The Bramble-Pasciak preconditioner (1988) pioneered this paradigm, leading to numerous variants and extensions over the years. In this talk we will discuss some recent results about non-standard inner products in the context of saddle point problems. Among them are spectral results for certain nonsymmetric saddle point matrices, and properties and simplifications of the non-Hermitian Lanczos algorithm in non-standard inner products. The talk will be based on joint work with Justus Ramme (TU Berlin).

→ Back to timetable

Approximate Commutator Preconditioners for Saddle-Point Systems with Indefinite Leading Blocks

Rachel Yovel, Ben-Gurion University of the Negev, Israel

In this talk we focus on the elastic Helmholtz equation, a challenging system of PDEs whose size and complex physics make it harder to solve than the acoustic Helmholtz equation. Recently, it was observed that the elastic problem can be reformulated as a generalized saddle-point system, which motivates new preconditioning strategies. Exploiting the approximate commutativity of the underlying differential operators, we propose a block-triangular preconditioner with acoustic Helmholtz operators on the diagonal. This allows the elastic problem to be solved using existing acoustic solvers as black boxes.

We prove a sufficient condition for convergence, shedding light on the long-questioned role of the commutator in approximate commutator preconditioners. We demonstrate scalability with respect to both the Poisson ratio and grid size, and show that a block-acoustic multigrid preconditioner reduces computational cost and memory use compared to a recent monolithic multigrid method for highly heterogeneous 2D and 3D media.

If time permits, we may also discuss extensions to other saddle-point systems with indefinite leading blocks.

→ Back to timetable

A BFBt Preconditioner for Double Saddle-Point Systems

Chen Greif, University of British Columbia, Canada

We consider double saddle-point systems. Given the 3-by-3 block structure of the matrix, the associated block-LU decomposition features two Schur complements. A theoretical question we explore is what happens when one of the Schur complements is inverted exactly and the second, nested one, is approximated. Eigenvalue analysis sheds some light on the effect of this type of inexactness on the speed of convergence of minimum residual iterative solvers. We develop a variant of Elman's BFBt method and adapt it to this family of linear systems. Our findings are illustrated on a Marker-and-Cell discretization of the Stokes-Darcy equations.

→ Back to timetable

▷ M13: Preconditioning for Structured Inverse Problems

Revisiting Flexibly Preconditioned Hybrid LSQR for Efficient Large-Scale Regularization

Iveta Hnětynková, Charles University, Czech Republic

Many applications require solving large-scale ill-posed problems contaminated by noise. In this context, Krylov subspace methods are particularly effective, as they can be enhanced with flexible preconditioning to enforce solution properties such as nonnegativity or sparsity. While Flexible LSQR (FLSQR) enables robust hybrid regularization, its underlying Flexible Golub-Kahan process is computationally demanding due to its reliance on two long-term recurrences.

In this talk, we present a novel Fast Flexible Golub-Kahan method that reduces this overhead by employing only one long-term and one short-term recurrence. This leads to the Fast Flexible LSQR (FaFLSQR) algorithm, which maintains the full hybrid regularization capabilities of FLSQR at a computational cost comparable to Flexible CGLS (FCGLS). We discuss the properties of FaFLSQR and show its mathematical equivalence to FCGLS. Finally, we present numerical experiments demonstrating that in floating-point arithmetic, FaFLSQR outperforms both FCGLS and FLSQR in terms of computational efficiency.

→ Back to timetable

The Role of Inexactness in Krylov Subspace Regularization

Malena Sabaté Landman, University of Bath, UK

Linear discrete inverse problems arise in many areas of science and engineering, from medical imaging and geophysics to atmospheric modelling. Their numerical solution often relies on iterative algorithms, particularly Krylov subspace methods, that can efficiently handle large-scale, ill-posed systems. In many practical settings, however, exact computations of matrix–vector products, preconditioners, or right-hand sides are either infeasible or unnecessary, leading to inexact iterations. This talk explores the interplay between inexactness and the regularizing behaviour of Krylov subspace methods for inverse problems. This is joint work with Silvia Gazzola.

→ Back to timetable

Preconditioning Strategies for a Nested Primal–Dual Method

Marco Donatelli, University of Insubria, Italy

Regularization techniques for image reconstruction problems typically consist of a smooth term and a potentially non-smooth convex term. A common approach to solving these problems is to use proximal gradient methods. To accelerate the convergence of these first-order iterative algorithms, extrapolation strategies have been introduced in the literature. We propose a variable metric strategy that can be reinterpreted as a right-preconditioning method. Consequently, we also explore a left-preconditioned version of the same proximal gradient method, along with possible approximation techniques. We prove the convergence of the preconditioned iterations under a suitable choice of the sequence of preconditioners. Numerical results show that preconditioning accelerates convergence while reducing the overall CPU time.

→ Back to timetable

Extended-Krylov-Subspace Methods for Trust-Region and Norm-Regularization Subproblems

Hussam Al Daas, STFC Rutherford Appleton Laboratory, UK

We consider an effective new method for solving trust-region and norm-regularization problems that arise as subproblems in many optimization applications. We show that the solutions to such subproblems lie on a manifold of approximately very low rank as a function of their controlling parameters (trust-region radius or regularization weight). Based on this, we build a basis for this manifold using an efficient extended-Krylov-subspace iteration that involves a single matrix factorization. The problems within the subspace using such a basis may be solved at very low cost using effective high-order root-finding methods. This then provides an alternative to common methods using multiple factorizations or standard Krylov subspaces. We provide numerical results to illustrate the effectiveness of our TREK/NREK approach.

→ Back to timetable

Acceleration of Block-Column Based Solvers for Structured Inverse Problems

Eric de Sturler, Virginia Tech, USA

Big data applications are becoming ever more prominent, and in many applications we need to solve very large linear or nonlinear inverse problems while handling only a relatively small amount of data at a time. We discuss solving very large-scale inverse problems using a block-column-based approach and accelerations of these methods for structured inverse problems.

→ Back to timetable

▷ M14: Preconditioners for Iterative Methods in Optimization

The Omega- and Kappa-Condition Numbers and Optimal Preconditioning

Henry Wolkowicz, University of Waterloo, Canada

Preconditioning is essential in iterative methods for solving linear systems $Ax = b$.

1. We first consider the classical approach that reduces the standard kappa-condition number of A , i.e., the ratio of largest to smallest singular values, a pseudoconvex, nonsmooth function that is inverse invariant. We present new models and algorithms with convergence guarantees that are significantly more efficient and accurate than the current SDP approaches in the literature.

2. Second, we present properties and exploit the structure of the omega condition number of A , i.e., the ratio of the arithmetic and geometric mean of the singular values, a pseudoconvex smooth function that is not inverse invariant. We discuss the issue of inverse invariance and show how we can exploit the structure to get explicit cheap formulae for omega-optimal preconditioners. Several of these are popular heuristics used in the literature, e.g., Jacobi, partial Cholesky with chordal completion.

3. We show empirically that omega-optimal preconditioners are better at clustering eigenvalues around 1, and so also better for iterative methods.

→ Back to timetable

Challenges in Design of Preconditioners for Interior Point Methods

Jacek Gondzio, University of Edinburgh, UK

Due to the use of logarithmic barrier function, interior point methods (IPMs) have to solve very ill-conditioned KKT systems. When problems are too large for direct methods of linear algebra Krylov-subspace methods are usually applied. However, standard preconditioners for saddle point systems are challenged in this case. This talk will address some of the key difficulties in the design and use of preconditioners for such systems.

→ Back to timetable

An Incomplete Factorization Preconditioner for Saddle Point Systems Arising in Interior Point Methods

Artyom Gevorgyan, University of Edinburgh, UK

Our focus will be on solving very large KKT systems that arise in primal–dual interior-point methods using Krylov-subspace iterative solvers. As these systems are indefinite and often severely ill-conditioned, good preconditioning is essential for reliable performance. We will present indefinite factorization-based preconditioners and discuss practical reformulations and regularization strategies that can stabilise the linear algebra and make the iterations more robust. We will demonstrate the proposed techniques on large-scale linear and quadratic programming problems solved using primal–dual interior-point methods.

→ Back to timetable

Randomized Nyström Preconditioned Interior Point-Proximal Method of Multipliers

Luiz-Rafael Santos, Universidade Federal de Santa Catarina, Brazil

In this talk, we present a novel algorithm for convex separable quadratic programming (QP) called Nys-IP-PMM, a regularized interior-point solver that uses low-rank structure to accelerate the solution of the Newton system. The algorithm combines the interior point proximal method of multipliers (IP-PMM) with the randomized Nyström preconditioned conjugate gradient method as the inner linear system solver. Our algorithm is matrix-free: it accesses the input matrices solely through matrix-vector products, as opposed to methods involving matrix factorization. It works particularly well for separable QP instances with dense constraint matrices. We establish convergence of Nys-IP-PMM. Numerical experiments demonstrate its superior performance in terms of wall-clock time compared to previous matrix-free IPM-based approaches.

→ Back to timetable

Preconditioned Krylov Solvers for KKT systems on GPUs

Alexis Montoison, Argonne National Laboratory, USA

Modern interior-point methods (IPMs) for nonlinear optimization rely on solving large and often ill-conditioned Karush-Kuhn-Tucker (KKT) systems. While direct factorizations remain the standard approach, they have clear limitations, especially on GPUs. In this talk, we present HybridKKT.jl, a Julia package for solving KKT systems using a hybrid approach that combines Cholesky and conjugate gradient method. These methods rely on Krylov.jl and CUDSS.jl, and are integrated into MadNLP.jl, enabling new approaches for GPU-accelerated nonlinear optimization.

We further demonstrate the use of GPU-based Krylov solvers in optimization through CompressedSensingIPM.jl, an IPM-based solver for sparse signal recovery formulated as an l_1 -regularized least-squares problem. In this context, matrix-free Krylov methods on GPU are the only viable alternative, as FFT operators are never materialized explicitly as dense matrices. For this specific problem, we show how a GPU-tailored preconditioner can be designed and implemented efficiently.

Finally, we discuss scaled variants of KKT systems for more general optimization frameworks, and outline how these formulations differ from standard ones and their potential impact on GPU.

→ Back to timetable

▷ M15: Acceleration Methods: From Fundamental Issues to Application in Code Coupling

History of Extrapolation Methods

Claude Brezinski, University of Lille, France

When a sequence or a series is slowly converging, it can be transformed into a new one converging faster to the same limit under some assumptions. Since it has been proved that a sequence transformation able to accelerate the convergence of all sequences (or even restricted classes of them) cannot exist, it is useful to design and study several such transformations. They are all based on the idea of extrapolation. It consists in assuming that a certain number of consecutive terms of the sequence are produced by a known function and to extrapolate it, usually at zero or at infinity, which furnishes a guess of its limit. Changing these consecutive terms leads to another guess and, thus, step by step, a new sequence is built.

This talk begins by describing linear extrapolation methods for computing the value of π . Then, the well-known Richardson extrapolation method is presented, as well as Romberg method for extrapolating the trapezoidal rule, which derives from it. The first nonlinear extrapolation method was Aitken process. It was then generalized by and recursively implemented by the algorithm of Wynn.

C. Brezinski, M. Redivo-Zaglia, *Extrapolation and Rational Approximation. The Works of the Main Contributors*, Springer. 2020

→ Back to timetable

Extrapolation Methods for Vector, Matrices and Tensors

Michela Redivo-Zaglia, University of Padua, Italy

One of the purpose of preconditioning a problem to be solved via iterative methods that produce a sequence of elements of a vector space is to reduce the number of iterations for obtaining a good approximation of the exact solution. The aim of this talk is to show that it is possible to apply directly to the original sequence an extrapolation method for accelerating its convergence.

Thus, suppose that the sequence of elements of a general vector space converges slowly to its limit. We can transform it into a new sequence which converges faster to the same limit.

A well known transformation is due to Shanks (1949, 1955), and it can be implemented by the scalar eps-algorithm of Wynn (1956). Wynn (1962) proposed also versions vectors (VEA) and square matrices.

Shanks transformation was generalized in several different ways: in short the MPE, MMPE, RRE, TEA. The last one can be implemented by the simplified topological eps-algorithms (STEA, 2014) that allow to treat sequences of scalar, vectors, matrices and also tensors.

A Matlab toolbox was also written and it is freely available (2017).

Numerical examples showing the effectiveness of these techniques will be presented.

→ Back to timetable

Acceleration Techniques for Implicit Adaptive Multistep Coupling in Aerospace Engineering Applications

Antoine E. Simon, École polytechnique & ONERA, France

Multiphysics simulations often rely on separate codes coupled through a partitioned-in-time numerical scheme, which usually involves low-order and stability issues. A novel high-order time-adaptive multistep coupling scheme has recently been proposed [1] yielding indication on stability limit for explicit scheme and high-order in time, and enabling larger time steps in the implicit variant at the cost of solving nonlinear fixed-point problems. Such an approach offers important computational gain as well as precision [2]. The objective of the contribution is to present how convergence acceleration can lead to even improved efficiency of the implicit multistep coupling approach. Fixed-point acceleration techniques will be discussed, with particular attention to how the structure of the multistep coupling influences their effectiveness. Several benchmarks will be provided, with emphasis on a conjugate fluid-solid heat transfer configuration involving unsteady dynamics, with metrics on iteration counts, achievable time step, and overall efficiency.

[1] A.E. Simon, L. François, M. Massot. *Comptes Rendus. Mécanique* 353 (G1) (2025) 1159-1184

[2] B. Dias, L. François, M. Massot, A.E. Simon. *AIAA SCITECH 2026 Forum* (2026) 2937

→ Back to timetable

Acceleration Methods for Fixed Point Iterations

Yousef Saad, University of Minnesota, USA

A wide range of computational problems may be cast as the determination of the limit of a sequence and extrapolation techniques sought to improve convergence by extrapolating toward this limit, usually by forming linear combinations of recent iterates. This extrapolation framework is often overly restrictive, motivating the development of alternatives that exploit both the iterates and the underlying fixed-point mapping that generates them. Among the most prominent of these fixed-point accelerators is the method introduced by D. Anderson in 1965, now commonly referred to as Anderson acceleration. In parallel, Krylov methods evolved along a different trajectory. Originally developed for linear systems, simple acceleration ideas naturally led to Krylov subspace techniques. Quasi-Newton and Inexact Newton methods may also be interpreted as acceleration strategies. Thus, acceleration methods have emerged from multiple perspectives and were often developed independently despite being founded on identical core principles. This presentation will be a survey of a broad class of acceleration methods, examining their origins, their historical successes, and the many - often subtle - connections among them.

→ Back to timetable

▷ Contributed Talks

Scalable ILU-Based Approximate Selected Inversion with Spectral Corrections for Sparse Systems

Tahamina Akter, TU Braunschweig, Germany

We analyze two parallel algebraic techniques for computing selected entries of the inverse of large, sparse symmetric systems arising in preconditioning-related computations: selected inversion and a factorized approximate inverse based on ILU factorizations. The selected inverse approach exploits an LU/ILU factorization to recover inverse entries on the sparsity pattern induced by the (incomplete) factors. In contrast, the factorized approximate inverse method applies a truncated Neumann series expansion to the (incomplete) LU factors, reducing computational cost at the expense of accuracy. To improve this accuracy–sparsity trade-off, we introduce a low-rank correction together with eigenvector deflation, providing an alternative to tighter drop tolerances and the resulting fill-in. Numerical experiments in both sequential and parallel settings demonstrate the robustness and scalability of the proposed approach for preconditioning-related applications.

→ Back to timetable

A Multi-Physics Reduction Preconditioner for a Fully Coupled 3D Dynamic Biot Model

Massimiliano Ferronato, University of Padova, Italy

An efficient multi-field coupled dynamic model for saturated anisotropic porous materials is developed. The mathematical formulation of the dynamic deformation–diffusion problem is obtained starting from the mixture theory and the definition of the effective stress for anisotropic poro–elasticity, taking into account also the fluid phase compressibility and anisotropic permeability. Numerical solution of the coupled problem is obtained by inf-sup stable Finite Element spaces. A fundamental issue for the computational efficiency of the coupled model is the numerical solution of the resulting large–size and non–symmetric discrete problem. In this work, we develop a fully implicit monolithic solver based on the Bi–Conjugate Gradient Stabilized (Bi-CGStab) algorithm accelerated via an ad–hoc Multi–Physics Reduction (MPR) preconditioning technique. Numerical analyses are performed to test the potential and computational efficiency of the proposed tool. In particular, we focus on 3D wave propagation applications in fully saturated single and multi-layered anisotropic media.

→ Back to timetable

HSS Matrices are Riemannian Manifolds

Pinki Khatun, Università degli Studi di Firenze, Italy

Hierarchically semi-separable (HSS) matrices provide an efficient representation of large-scale structured matrices with low-rank off-diagonal blocks. In this work, we reveal and formalize the intrinsic geometric structure of fixed-rank HSS matrices by modeling their parameter space as a Riemannian quotient manifold. We identify the natural equivalence relation induced by local

unitary gauge transformations in HSS representations and construct the associated total space, as well as the vertical and horizontal tangent spaces. A Riemannian metric invariant under the group action is then defined, yielding a well-posed Riemannian structure on the quotient space of HSS matrices. This geometric formulation enables the application of Riemannian optimization techniques directly on the HSS manifold; for instance, one can use this technology to address the HSS matrix nearness problem. The proposed approach generalizes existing geometric treatments of low-rank matrix manifolds, provides new theoretical insights into hierarchical matrix formats, and establishes a principled foundation for solving efficient structured-matrix approximation problems and developing optimization algorithms for large-scale HSS-structured matrices.

→ Back to timetable

Coupled Ordering Techniques for Coupled Partial Differential equations

Sabine Le Borne, Hamburg University of Technology, Germany

Fluid flow problems can be modeled by the Navier-Stokes equations, a system of partial differential equations that couples the velocity and pressure of the fluid. Their linearization and discretization result in discrete saddle point problems of Oseen type which are typically very large and need to be solved iteratively. Standard (block) preconditioning techniques rely on an approximation of the (inverse) Schur complement. In this talk, we discuss coupled ordering techniques for the pressure and velocity unknowns that facilitate the computation of such Schur complement approximations. In particular, such orderings can be combined with block clustering strategies in the construction of hierarchical matrices and accelerate the construction of hierarchical LU factorizations of the Schur complement. Numerical results illustrate the performance of the resulting saddle point preconditioners.

→ Back to timetable

mdBFGS: A Diagonal BFGS Method with Momentum for Stochastic Optimisation

Andrea Meda, University of Edinburgh, UK

We introduce a diagonal quasi-Newton optimiser called momentum diagonal BFGS (mdBFGS). This optimizer combines BFGS-style curvature updates with exponential moving averages to form a per-parameter preconditioner for large-scale stochastic and non-convex problems. The preconditioner is updated via a damped secant condition evaluated on exponentially averaged curvature pairs, which reduces stochastic noise. A damping mechanism further stabilizes the preconditioner under mini-batch gradients. The method is broadly applicable, from linear regression to large-scale deep neural network training, and offers a simple, scalable alternative to both first-order momentum methods and quasi-Newton schemes.

→ Back to timetable

Low-Precision Preconditioning for Inverse Problems

James Nagy, Emory University, USA

In recent years a substantial amount of work has been done on developing mixed-precision algorithms for linear systems, methods that can exploit capabilities of modern GPU architectures. However, very little work has been done for ill-conditioned problems that arise from large-scale inverse problems. Special considerations, which normally do not arise when solving well-conditioned problems, such as incorporating regularization into the developed methods, need to be considered. In this talk we consider the use of low-precision, regularized preconditioners in flexible Krylov subspace methods.

→ Back to timetable

Exploiting Port-Hamiltonian Structures for Preconditioning Saddle Point Problems

Rene Noffke, University of Wuppertal, Germany

Port-Hamiltonian systems (pHS) provide a powerful framework for modeling and coupling physical environments. Structure-preserving and energy-consistent discretizations often yield large, sparse, and indefinite saddle point problems (SPPs). Standard preconditioning techniques often ignore the underlying physical properties. In this ongoing PhD project, we develop efficient preconditioners that suit the physical problem and exploit structural properties for computational efficiency. As an initial step, we introduce a 1D transmission line and a fully discrete scheme fitting the pHS framework in both continuous and discrete settings. This preserves all energy properties and naturally forms an SPP, introducing the necessary frameworks. Coupled versions of this system will be used as our benchmark. Building on this, we present current work on designing structure-exploiting preconditioners that utilize the block structure while aiming to preserve physical properties. We will show preliminary numerical results using general Krylov subspace methods on this coupled 2D/3D transmission line benchmark, highlighting the potential of structure-aware solvers.

→ Back to timetable

Robust MGR Preconditioning for Large-Scale Coupled Subsurface Flow and Geomechanics

Daniel Osei-Kuffuor, Lawrence Livermore National Laboratory, USA

Solving the massive, ill-conditioned linear systems that arise from coupled multiphysics simulations requires solvers that are both algorithmically resilient and computationally highly efficient. This talk details recent extensions to the Multigrid Reduction (MGR) block preconditioning framework, specifically tailored for challenging subsurface applications including multiphase poromechanics, thermal and isothermal compositional flow, and fractured reservoir environments. By integrating these advanced MGR strategies into the hypre library, we achieve seamless portability and high-performance execution across contemporary CPU and accelerated GPU platforms. We will present extensive numerical experiments conducted within the GEOS multiphysics simulator, illustrating the solver's insensitivity to strong physical couplings and highly heterogeneous media.

These results validate the enhanced MGR method as a robust, highly scalable solver technology capable of efficiently converging real-world problems with billions of unknowns.

This is work with Victor Magri and Nicola Castelletto (LLNL).

→ Back to timetable

Optimal Transfer Operators and Convergence Bounds for Nonsymmetric Two-Grid Methods

Ludwig Roach, Technische Universität Berlin, Germany

Algebraic Multigrid methods have been proven to be effective solvers for large-scale linear algebraic systems with Hermitian positive definite (HPD) system matrix. For such problems the convergence, but for nonsymmetric indefinite systems fewer results exist. Recently, convergence results for more general norms induced by certain HPD matrices were established. We present a theoretical framework for nonsymmetric algebraic two-grid methods for arbitrary inner products which naturally includes the HPD case and all recent results for the nonsymmetric case. We consider two different two-grid error operators, the first one being the generalization of the error operator in the HPD case. The second operator has been studied before and is simpler, but requires additional assumptions to achieve convergence. We explain the differences and similarities of both operators, the necessity of the extra condition and generalize some previous results. We give sharp estimates for the norms of the error propagation matrices and establish optimal Transfer operators for both methods. Finally, we analyze the effect of the number of coarse variables on the convergence speed. This talk is based on a joint work with Reinhard Nabben (TU Berlin).

→ Back to timetable

Block Jacobi Preconditioning on Analog Hardware

Shikhar Shah, Emory University, USA

Analog in-memory computing uses non-volatile resistive memory devices to perform linear algebra operations directly through physical laws such as Ohm's and Kirchoff's laws. Unlike digital architectures, these operations are inherently approximate and stochastic. However, key kernels such as the matrix-vector product can be executed in time independent of the matrix size, making analog hardware attractive for large-scale linear algebra problems, particularly for the solution of large, sparse linear systems.

In this talk, we present theoretical and experimental convergence results for a block Jacobi preconditioner implemented on analog hardware. We compare several construction methods (including the sparse approximate inverse and the Monte Carlo approximate inverse) and evaluate performance on the Laplace and convection-diffusion-reaction equations. Our results show that optimal preconditioner parameters differ substantially from those used on digital hardware. We further extend this approach to a hierarchical two-level preconditioner scalable to very large systems.

→ Back to timetable

Monolithic and Schur-Based Domain Decomposition Strategies for Hybrid High-Order Discretizations

Bahaâ Eddine Sidi Hida, EDF R&D, France

Hybrid High-Order (HHO) discretizations yield block-structured linear systems separating cell and face unknowns. Independently of the polynomial degree, the cell block is block-diagonal by element and can be inverted locally at negligible cost. This structural property motivates solver strategies that explicitly exploit the separation of unknowns.

We investigate two complementary domain decomposition approaches based on the preconditioner HPDDM. The first is a monolithic strategy, where the global HHO system is preconditioned directly. This approach is simple to implement but requires careful parameter choices to achieve robustness and scalability.

The second approach relies on a Schur complement reduction on face unknowns. Leveraging the block-diagonal structure of the cell block, we assemble the exact Schur complement explicitly and use it as the effective operator of the face problem. The reduced system is then preconditioned with HPDDM.

Numerical experiments on 3D elasticity problems illustrate the efficiency of both strategies and include comparisons with state-of-the-art multigrid solvers.

→ Back to timetable

Basis Preconditioning for Gas Flow Optimization

Rowan Turner, University of Edinburgh, UK

We present a basis preconditioner for a PDE-constrained optimization problem, arising from gas transport in hydrogen gas networks. We discretize-then-optimize, which generates a saddle point problem with an indefinite 1-1 block, due to the nonlinearity of the problem. To precondition the problem, we derive a basis for the constraint matrix, which corresponds to solving a problem with fixed boundary conditions/controls. We apply this preconditioner in the context of an interior point method (IPM), on two problems – one which considers the governing isothermal Euler equations along with compressor equations for controlling pressure in the network, and the second which also considers integer valued valve equations for controlling flow. We present results on scalability for time horizons and mesh refinement, both in terms of the iterative solve and the outer IPM iteration count.

→ Back to timetable

Multigrid Methods with Approximate Coarsest-Level Solvers

Petr Vacek, CNRS-IRIT, France

Multigrid methods are frequently used when solving systems of linear equations. They are applied either as standalone solvers or as preconditioners for a Krylov subspace method. They compute an approximate solution by using smoothing on fine levels and solving a system of linear equations on the coarsest level.

The choice and configuration of the coarsest-level solver can significantly affect the overall performance of the multigrid method. The typical approach is to use a direct method based on

LU or Cholesky factorization. There are, however, settings where using approximate solvers, such as (preconditioned) Krylov subspace methods or block low-rank direct approximate solvers leads to better performance. Achieving good performance with these solvers requires finding the right balance: their accuracy must be sufficient not to slow down the overall convergence, while the computational cost remains low.

In this talk we presents new computable stopping criteria for iterative coarsest-level solvers tailored to multigrid methods. Our work builds upon the result in [Vacek et al., SIAM Journal on Scientific Computing (2024)]. We present new theoretical results as well as numerical experiments on GPUs using the Ginkgo library.

→ Back to timetable

Minisymposium Titles and Organisers

M1: Convergence Analysis and Acceleration of GMRES

▷ Organisers: Pierre Marchand, Nicole Spillane

M2: Intelligent Preconditioning for Partial Differential Equations

▷ Organisers: Axel Klawonn, Xiao-Chuan Cai

M3: Advances in Preconditioning Techniques for Wave Propagation Problems

▷ Organisers: Racheli Yovel, Eran Treister

M4: Preconditioned Iterative Methods for Discretised PDEs

▷ Organisers: Jennifer Pestana, Kirk Soodhalter

M5: Challenges in Preconditioning for Large Scale Least Squares Problems

▷ Organisers: Melina Freitag, Jennifer Scott

M6: Preconditioning Across Scientific Computing and Learning

▷ Organisers: Yuanzhe Xi, Daniel Osei-Kuffuor, Rui Peng Li

M7: Quantum Computing Meets Numerical Analysis

▷ Organiser: Nicole Spillane

M8: Preconditioning Strategies for Large-Scale Data Assimilation

▷ Organisers: Selime Gürol, Ieva Daužickaitė

M9: New Results in Domain Decomposition

▷ Organisers: Nicole Spillane, Victorita Dolean

M10: Advanced Preconditioners for Large-Scale Linear Systems in HPC

▷ Organisers: Andrea Franceschini, Carlo Janna, Massimiliano Ferronato

M11: Preconditioning for Neural Networks and Neural Networks for Preconditioning

▷ Organisers: Alena Kopaničáková, Alexander Heinlein

M12: Advances in Preconditioning for (Multiple) Saddle Point Linear Systems

▷ Organiser: Chen Greif

M13: Preconditioning for Structured Inverse Problems

▷ Organisers: Malena Sabaté Landman, James Nagy

M14: Preconditioners for Iterative Methods in Optimization

▷ Organiser: Jacek Gondzio

M15: Acceleration Methods: From Fundamental Issues to Application in Code Coupling

▷ Organisers: Marc Massot, Yousef Saad

→ [Back to timetable](#)