Nordic numerical linear algebra meeting

University of Southern Denmark (SDU), Odense

June 17–18, 2024

Nordic numerical linear algebra meeting 2024

Scope

The workshop is devoted to all aspects of applied and numerical linear algebra, spanning theory, analysis, computation and methods to specific research areas in applications as well as mathematical software. The topics include but are not limited to: linear systems, nonlinear systems, iterative methods, direct methods, eigenvalue problems, matrix functions, structured matrices, least squares problems, tensor methods, matrix equations, graph theory and high performance computing. Specific application areas include but are not restricted to: physical sciences, systems & control, data science, machine learning, biology and statistics.

Invited speakers

Erin Claire Carson, Charles University, Czech Republic

Iveta Hnetynkova, Charles University, Czech Republic

Alex Townsend, Cornell University, US

Organizing committee

Philipp Birken (Lund University)

Andrii Dmytryshyn (Örebro University)

Elias Jarlebring (KTH Royal Institute of Technology)

Ralf Zimmermann (University of Southern Denmark)

List of participants

Joe Alexandersen, University of Southern Denmark
Magnus Appel, University of Southern Denmark
Amirhossein Bayat, Southern Denmark university
Philipp Birken, Lund University
Erin Claire Carson, Charles University Prague
Marco Chiarandini, University of Southern Denmark
Sweta Das, Örebro University
Kristian Debrabant, University of Southern Denmark
Andrii Dmytryshyn, Örebro University
Ivo Dravins, Ruhr University Bochum
Miryam Gnazzo, Gran Sasso Science Institute
Iveta Hnetynkova, Charles University Prague
Elias Jarlebring, KTH Royal Institute of Technology
Rasmus Jensen, University of Southern Denmark
Carl Christian Kjelgaard Mikkelsen, Umeå University
Michael Koch, Hamburg University of Technology
Niklas Kotarsky, Lund University
Hao Li, University of Southern Denmark
Gustaf Lorentzon, KTH Royal Institute of Technology
Søren Madsen, Aarhus University
Souad Mohaoui, Örebro University
Jan Martin Nicolaus, University of Potsdam
Lauri Nyman, Aalto University
Vilhelm Peterson Lithell, KTH Royal Institute of Technology
Claudio Pica, University of Southern Denmark
Achim Schroll, University of Southern Denmark
Bonan Sun, Max Planck Institute for Dynamics of Complex Technical Systems
Alex Townsend, Cornell University
Ralf Zimmermann, University of Southern Denmark

Programme

Monday, June 17	
08:30	Registration opens
08:50	Welcome note
09:00-10:00	**Keynote 1: Erin Claire Carson**
10:00-10:30	Coffee break
10:30-11:00	**Bonan Sun**
11:00-11:30	**Souad Mohaoui**
11:30-12:00	**Lauri Nyman**
12:00-12:30	**Rasmus Jensen**
12:30-13:30	Lunch break
13:30-14:30	**Keynote 2: Alex Townsend**
14:30-15:00	**Michael Koch**
15:00-15:30	**Joe Alexandersen**
15:30-15:45	Coffee break
15:45-16:15	**Miryam Gnazzo**
16:15-16:45	**Sweta Das**
19:00	Conference dinner Den Gamle Kro Overgade 23 5000 Odense C

Programme

Tuesday, June 18	
08:30	Registration opens
09:00-10:00	**Keynote 3: Iveta Hnětynková**
10:00-10:30	Coffee break
10:30-11:00	**Ivo Dravins**
11:00-11:30	**Philipp Birken**
11:30-12:00	**Niklas Kotarsky**
12:00-12:30	**Jan Martin Nicolaus**
12:30-13:30	Lunch break
13:30-14:00	**Carl Christian Kjelgaard Mikkelsen**
14:00-14:30	**Elias Jarlebring**
14.30-14:45	Coffee break / Closing

Keynote 1: Balancing Inexactness in Large-Scale Matrix Computations

Carson, Erin Claire
Charles University, Prague

On supercomputers that exist today, achieving even close to the peak performance is incredibly difficult if not impossible for many applications. Techniques designed to improve the performance of matrix computations – making computations less expensive by reorganizing an algorithm, making intentional approximations, and using lower precision – all introduce what we can generally call 'inexactness'. The questions to ask are then:

1. With all these various sources of inexactness involved, does a given algorithm still get close enough to the right answer?
2. Given a user constraint on required accuracy, how can we best exploit and balance different types of inexactness to improve performance?

Studying the combination of different sources of inexactness can thus reveal not only limitations, but also new opportunities for developing algorithms for matrix computations that are both fast and provably accurate. We present a few recent examples of this approach, in which mixed precision computation is combined with other sources of inexactness.

Keynote 2: Rank–revealers must use near–local maximum volume pivoting

Townsend, Alex
Cornell University, Ithaca

Attend any numerical linear algebra conference these days and you will likely find a large fraction of talks about matrix compression, singular value estimation, and column subset selection. In this talk, we will study algorithms called rank–revealers that reveal a matrix's rank structure, which provides theoretical guarantees on the quality of low-rank approximants and condition number bounds on bases for column space and nullspace. We show that the concept of near–local maximum volume pivoting is a crucial and practical pivoting strategy for rank–revealers based on Gaussian elimination (GE) and QR, showing that it is both necessary and sufficient. This insight elevates Gu and Eisenstat's rank-revealing QR as an archetypal rank–revealer, and we design a metric to measure the success of any pivoting strategy. This partially explains the popularity of column-pivoted QR and GE with complete pivoting as rank–revealers.

Keynote 3: Multilinear approximation problems: Variants, solvability analysis and extraction of core data

Hnětynková, Iveta
Charles University, Prague

In many areas, there is a need to solve multidimensional approximation problems $AX \approx B$ of various forms. Such problems suffer from the presence of errors and irrelevant or redundant data in B and/or A. Additional difficulties are related to the fact that the individual observations (columns of B) can be correlated with different subsets of columns of A. Least squares methods (such total or scaled least squares) are widely used here, where generally a minimum norm data correction is searched ensuring existence of the approximate solution X. However, such a correction may not always exist.

Irrelevant and redundant data can be removed by the core reduction introduced in [8] and [9] for single–observation problem. This yields a minimally dimensioned subproblem called the core problem with better solvability properties. Recently, the core reduction was extended to multiple observation problems [4], bilinear, tensor or other highly structured problems [5], [6]. The reduction can be direct (using the SVD or Tucker decomposion) or iterative, based on the block generalization of the Golub–Kahan bidiagonalization [1]. The recent paper [7] gives a unifying universal Krylov subspace approach yielding core data with a band model matrix and orthogonal observation matrix.

In this talk, we summarize various formulations of multilinear approximation problems and analyze their solvability. Then we turn to core data reduction. Direct as well as iterative approach will be explained in details, properties of the resulting minimal core problem will be studied. We give also some notes on difficulties appearing in FPA computations.

References

[1] G. H. Golub, W. Kahan Calculating the singular values and pseudo-inverse of a matrix. SIAM J. Numer. Anal. Ser. B 2 (1965), pp. 205–224.

[2] G. H. Golub, C. F. Van Loan. An analysis of the total least squares problem. SIAM J. Numer. Anal., 17(6) (1980), pp. 883–893.

[3] I. Hnětynkovà, M. Plešinger, Z. Strakoš. The core problem within linear approximation problem $AX \approx B$ with multiple right-hand sides. SIAM J. Matrix Anal. and Appl., 34(3) (2013), pp. 917–931.

[4] I. Hnětynková, M. Plešinger, Z. Strakoš. Band generalization of the Golub-Kahan bidiagonalization, generalized Jacobi matrices, and the core problem. SIAM J. Matrix Anal. and Appl., 36(2) (2015), pp. 417–434.

[5] I. Hnětynková, M. Plešinger, J. Žáková. TLS formulation and core reduction for problems with structured right-hand sides. Linear Algebra and its Appl., 555 (2018) pp. 241–265.

[6] I. Hnětynková, M. Plešinger, J. Žáková. On TLS formulation and core reduction for data fitting with generalized models. Linear Algebra and its Appl., 577 (2019), pp. 1–20.

[7] I. Hnětynková, M. Plešinger, J. Žáková. Krylov Subspace Approach to Core Problems within Multilinear Approximation Problems: A Unifying Framework. SIAM J. on Matrix Anal. and Appl., 44 (2023), pp. 53–79.

[8] C. C. Paige, Z. Strakoš. Scaled total least squares fundamentals. Numerische Mathematik, 91 (2006), pp. 117–146.

[9] C. C. Paige, Z. Strakoš. Core problem in linear algebraic systems. SIAM J. Matrix Anal. and Appl., 27(3) (2006), pp. 861–875.

Approximation to Multivariate Functions in the Extended Tensor Train Format

Strössner, Christoph[a]

Sun, Bonan[b]

Kressner, Daniel[a]

[a] *École Polytechnique Fédérale de Lausanne (EPFL)*
[b] *Max Planck Institute for Dynamics of Complex Technical Systems.*
Work was done when the author was at EPFL

This talk introduces the extended functional tensor train (EFTT) format for compressing and working with multivariate functions on tensor product domains. Our compression algorithm combines tensorized Chebyshev interpolation with a low-rank approximation algorithm that is entirely based on function evaluations. Compared to existing methods based on the functional tensor train format, our approach often reduces the required storage, sometimes considerably, while achieving the same accuracy. In particular, we reduce the number of function evaluations required to achieve a prescribed accuracy by up to over 96% compared to the algorithm from [Gorodetsky, Karaman and Marzouk, Comput. Methods Appl. Mech. Eng., 347 (2019)].

References

[1] Strössner, C., Sun, B. and Kressner, D. Approximation in the extended functional tensor train format. *arXiv preprint*, arXiv:2211.11338, 2022.

Tensor Decomposition for Missing Marker Recovery in Motion Capture Data

Mohaoui, Souad[a]

Dmytryshyn, Andrii[a]

[a] *School of Science and Technology, Örebro University.*

Motion Capture (MoCap) is an essential technology for capturing and analyzing complex human movements. However, MoCap systems frequently encounter the challenge of missing markers due to occlusion or equipment limitations. In this study, we address this challenge using a tensor completion framework. We model the MoCap data as a third-order tensor. We propose a tensor completion approach based on the Canonical Polyadic (CP) tensor decomposition. Experimentally, we test our approach on diverse MoCap sequences from the CMU Motion Capture dataset.

Riemann-Oracle: A general-purpose Riemannian optimizer for matrix nearness problems

Nyman, Lauri[a]

Gnazzo, Miryam[b]

Noferini, Vanni[a]

Poloni, Federico[c]

[a] *Aalto University, Finland*
[b] *Gran Sasso Science Institute, Italy*
[c] *University of Pisa, Italy*

We propose a versatile approach to address a large family of matrix nearness problems. The method is based on splitting a matrix nearness problem into two nested optimization problems, of which the inner one can be solved either exactly or cheaply, while the outer one can be recast as an unconstrained optimization task over a smooth real Riemannian manifold. We also show that the objective function to be minimized on the Riemannian manifold can be discontinuous, thus requiring regularization techniques, and we give conditions for this to happen. Finally, we demonstrate the practical applicability of our method by implementing it for a number of matrix nearness problems that are relevant for applications, including the nearest singular matrix with a given sparsity structure and the nearest singular matrix polynomial.

Optimization on the symplectic Stiefel manifold using second order information

Jensen, Rasmus[a]

Zimmermann, Ralf[a]

[a] *University of Southern Denmark*

Riemannian computing methods have proven useful in various problems in numerical linear algebra. For example, one can obtain the eigenvalue decomposition and the singular value decomposition via minimizing special objective functions on the (product) Stiefel manifold. In this talk, we will discuss recent results in optimization on the symplectic Stiefel manifold. Application scenarios include computing symplectic eigenvectors and eigenvalues, solving the nearest symplectic matrix problem, and computing a proper symplectic decomposition. The latter is an important tool when performing snapshot–based, structure–preserving model order reduction of Hamiltonian systems.

Smaller stencil preconditioners for linear systems in RBF-FD discretizations

Koch, Michael[a]

Le Borne, Sabine[a]

Leinen, Willi[b]

[a] *Hamburg University of Technology*
[b] *Helmut Schmidt University Hamburg*

Radial basis function finite difference (RBF-FD) discretization has recently emerged as an alternative to classical finite difference or finite element discretization of (sys- tems) of partial differential equations. We focus on the construction of preconditioners for the iterative solution of the resulting linear systems of equa- tions. In RBF-FD, a higher discretization accuracy may be obtained by increasing the stencil size. This, however, leads to a less sparse and often also worse conditioned stiffness matrix which are both challenges for subsequent iterative solvers. We pro- pose to construct preconditioners based on stiffness matrices resulting from RBF-FD discretization with smaller stencil sizes compared to the one for the actual system to be solved. In our numerical results, we focus on RBF-FD discretizations based on polyharmonic splines (PHS) with polynomial augmentation. We illustrate the per- formance of smaller stencil preconditioners in the solution of the three-dimensional convection-diffusion equation.

Towards faster topology optimisation of transient problems using a parallel space-time multigrid preconditioner

Alexandersen, Joe[a]

[a] *SDU*

This work presents the first steps towards faster topology optimisation of time-dependent problems with many timesteps.

Topology optimisation is a process for automatically generating design concepts for complex engineering problems using a combination of simulations and optimisation algorithms. It is an iterative process, where the performance of the system is analysed using one or several simulations. The design is then updated until the performance is satisfactory, a demanding process usually requiring around 100-1000 simulations. This means that each simulation must be as fast as possible in order to reduce the overall time-to-solution for the user.

The current state-of-the-art is somewhat limited to steady-state (static) problems, since time-dependent (transient) simulations generally take a very long time due to the inherent serial nature of stepping through time to resolve the system behaviour. Practically, the computational time of time-stepping methods is infeasible in both industrial and academic timescales, because even simple time-dependent problems can take days to optimise. To overcome this limitation, this work considers a space-time formulation of the governing equations and solves the full time-dependent system using an "all-at-once" approach. Multigrid in space-time is applied as a preconditioner to Krylov solvers and combined with parallelisation to reduce the time-to-solution using distributed-memory computing clusters.

In this work, a continuous Galerkin finite element formulation with piecewise-linear basis functions is used to discretise the thermal diffusion equation with varying coefficients. Simple space-time "geometric" coarsening is used in conjunction with a Galerkin-projection geometric multigrid preconditioner. It is shown that a simple, and somewhat naive, implementation of the space-time framework for two-dimensional (spatial) problems leads to a 10% time-to-solution compared to time-stepping, albeit at a cost of 6 times the computational resources (core-hours). The simple setup works very well despite the issue of large contrast in coefficients in space (topology optimisation), even for nonlinear problems with large contrast also in time (phase change).

Backward errors for nonlinear eigenvalue problems

Gnazzo, Miryam[a]

Robol, Leonardo[b]

[a] *Gran Sasso Science Institute, Italy*
[b] *Università di Pisa, Italy*

The backward error analysis is an important part of the perturbation theory and it is particularly useful for the study of the reliability of the numerical methods. Given an approximate eigenpair $(\hat{\lambda}, \hat{v})$ of the nonlinear eigenvalue problem $F(\lambda) = \sum_{i=1}^{k} f_i(\lambda) F_i$, with $F_i \in \mathbb{C}^{n\times n}$, the backward error can be defined as the smallest perturbations ΔF_i such that $(\hat{\lambda}, \hat{v})$ is an exact eigenpair of $\sum_{i=1}^{k} f_i(\lambda)(F_i + \Delta F_i)$. We provide theoretical results about the backward error of a set of approximate eigenpairs $(\hat{\lambda}_1, \hat{v}_1), \ldots, (\hat{\lambda}_p, \hat{v}_p)$, that is finding a set of matrices ΔF_i such that $\sum_{i=1}^{k} f_i(\hat{\lambda}_j)(F_i + \Delta F_i)\hat{v}_j = 0$, for each $j = 1, \ldots, p$. Indeed, small backward errors for separate eigenpairs do not imply small backward errors for a set of approximate eigenpairs. In this talk, we provide inexpensive upper bounds, and a way to accurately compute the backward error by means of direct computations or through Riemannian optimization. We also discuss how the backward error can be determined when the F_i have particular structures (such as symmetry, sparsity, or low-rank), and the perturbations are required to preserve them. For special cases (such as for symmetric coefficients), explicit and inexpensive formulas to compute the δF_i are also given.

References

[1] M. Gnazzo. L. Robol. Backward errors for multiple eigenpairs in structured and unstructured nonlinear eigenvalue problems. *arXiv*, 2024

Minimal degeneration of skew-symmetric matrix polynomials

Das, Sweta[a]

Dmytryshyn, Andrii[a]

[a] *School of Science and Technology, Orebro University*

In this presentation, we describe qualitatively how eigenstructure of skew-symmetric matrix polynomials of odd degree changes under small perturbations in the matrix coefficients. Using strong linearization we prove a necessary and sufficient condition for one orbit of the linearization of a matrix polynomial to be a proper subset of the closure of the orbit of linearization of another polynomial. To achieve this, we introduce a set of rules describing the structure transitions of the canonical blocks of the polynomial's linearization. These rules facilitate the construction of the stratification graphs of linearization of the polynomials. Finally, we state a method that allows us to sketch the entire or a part of the stratification graph of one matrix polynomial's linearization in relation to another matrix polynomial's linearization, provided both polynomials share the same degree but have different dimensions.

Stage-parallel preconditioning for implicit Runge–Kutta methods

Dravins, Ivo[a]

[a] *Ruhr University Bochum, Germany*

Fully implicit Runge–Kutta methods offer the possibility to use high-order-accurate time discretization with desirable stability properties. For general implicit Runge–Kutta methods all stages are coupled leading to a potentially costly and involved solution procedure which has been a major barrier to their widespread use.

We present a stage-parallel block preconditioner for the class of L-stable Radau IIA Runge–Kutta methods. The preconditioner exploits a property of the coefficient matrices to construct a block lower-triangular preconditioner. During the application of the preconditioner, a basis change can be applied to obtain a block-diagonal form, in this way allowing us to decouple the stages when solving for the blocks. In the linear case this basis change can be applied directly – for non-linear equations further approximations are needed to achieve this decoupling.

For the linear case, we discuss the analysis of the preconditioned matrices which are non-symmetric and in tensor form. We give eigenvalue bounds for the two- and three-stage methods under symmetric positive definite assumptions and discuss what can be inferred about the general case.

We illustrate the performance by numerical examples, including also applications to non-linear problems. Finally the parallel behavior is demonstrated, comparing space-parallel but serial in stages against fully stage-parallel implementations on HPC platforms.

Convergence analysis of waveform relaxation for heterogeneous coupled heat equations

Birken, Philipp[a]

[a] *Lund University*

We consider two coupled linear heat equations on separate domains, with possibly discontinuous material coefficients across the interface. This can be used to model conjugate heat transfer, but also the transfer of wind stress between ocean and atmosphere. Our interest is in so called waveform relaxation. There, the subproblems are solved iteratively on a time window, given an approximation of the solution to the other problem. This can be combined with a relaxation step or Quasi-Newton acceleration.

In practice, one uses these methods in a fully discrete setting, where one has discretized in space and time. One is then interested in the convergence behaviour and robustness of these methods, and the dependence on the material parameters, the mesh width and the time steps. Over the years, this has been approached on different levels, meaning fully discrete, semidiscrete, or fully continuous. Thereby, one has obtained both superlinear convergence results, and norm estimates. Currently, these results give an incomplete and seemingly conflicting picture. We review these results, and then fill some of the gaps.

This is joint work with Niklas Kotarsky and Martin J. Gander.

Time adaptive Quasi Newton waveform iterations for thermal interaction

Kotarsky, Niklas[a]

Birken, Philipp[a]

[a] *Lund University*

We consider methods for dynamic coupled problems, or more specifically coupled PDEs that interact through a lower dimensional interface. Examples here are thermal interaction between two different materials or Fluid-structure interaction. To solve these problems we will be using a partitioned approach, where separate solvers are used for the different subdomains resulting in a fixed point iteration. To improve the computational efficiency, different and adaptive time steps in the sub solvers are employed. Using so called waveform iterations, these goals have been achieved for heat transfer problems in [1] and [2]. Quasi- Newton methods have recently been combined with waveform iterations for the case of constant time grids in [3]. In this talk, we discuss the convergence properties of the proposed Quasi- Newton method for waveform iterations in [3] for two coupled heat equations with fixed time steps, where the discrete interface system is reduced to a linear system. We also extend the Quasi-Newton waveform iterations to the time adaptive case, where both of the sub-solvers use an adaptive time stepping scheme. We present the general framework of waveform iterations and their implementation in the open source coupling library PreCICE. Lastly, we show that using a time adaptive solver results in faster run times for a thermal transfer test case.

References

[1] Meisriemel, P. *Adaptive time-integration for goal-oriented and coupled problems*. PhD Thesis, Lund University (2021)

[2] Monge, A. *Partitioned methods for time-dependent thermal fluid-structure interaction*. PhD Thesis, Lund University (2018)

[3] Ruth, B., Uekermann, B., Mehl, M., Birken, P., Monge, A. and Bungartz, H. J. Quasi-Newton waveform iteration for partitioned surface-coupled multiphysics applications, *Int. J. Num. Meth. Engng.* (2021) 122(19):5236–5257

Nonintrusive model order reduction for stochastic differential equations

Nicolaus, Jan Martin[a]

Freitag, Melina[a]

Redmann, Martin [b]

[a] *University of Potsdam*
[b] *Martin-Luther-University Halle-Wittenberg*

This work introduces a non-intrusive model order reduction method provided by extending the operator inference approach to linear and controlled stochastic differential equations with additive noise. The drift and diffusion coefficients of the reduced order model (ROM) are inferred from state observations by solving appropriate least-squares problems. The closeness of the ROM obtained by the presented approach to the intrusive ROM obtained by the proper orthogonal decomposition (POD) method is investigated. Two possible generalisations of the snapshot-based dominant subspace construction to the stochastic case are presented. Numerical experiments are provided to compare the developed approach to POD.

Newton's method revisited: How accurate do we have to be?

Kjelgaard Mikkelsen, Carl Christian[a]

Lopez-Villellas, Lorien [b]

Garcia-Risueno, Pablo[c]

[a] *UmU, Sweden*
[b] *BSC, Spain*
[c] *Independent researcher, Germany*

We analyze the convergence of quasi-Newton methods in exact and finite precision arithmetic using old and new techniques. We derive an upper bound for the stagnation level and we show that any sufficiently exact quasi-Newton method will converge quadratically until stagnation. In the absence of sufficient accuracy, we are likely to retain rapid linear convergence. We confirm our analysis by computing square roots and solving bond constraint equations in the context of molecular dynamics. In particular, we apply both a symmetric variant and Forsgren's variant of the simplified Newton method. This work has implications for the implementation of quasi-Newton methods regardless of the scale of the calculation or the machine.

Linearization of rational eigenvector nonlinearities

Jarlebring, Elias[a]

Upadhyaya, Parikshit[a]

Meerbergen, Karl[b]

Claes, Rob[b]

[a] *KTH Royal Institute of Technology, Sweden*
[b] *KU Leuven, Belgium*

In this talk, we introduce a method to linearize, without approximation, a distinct class of eigenvalue problems featuring eigenvector nonlinearities (NEPv). These nonlinearities are characterized by scalar functions that are represented as a quotient of linear functions of the eigenvector

$$\lambda v = (A_0 + f_1(v)A_1 + \cdots + f_m(v)A_m)v$$

where $f_i(v) = c_i^T v / d_i^T v$. Our approach employs an exact linearization through an associated multiparameter eigenvalue problem (MEP), which encapsulates the precise solutions of the NEPv. By framing MEPs within the context of a generalized eigenvalue problem, we provide a procedure to compute all solutions of the NEPv for moderate sized problems. This framework also paves the way for the development of locally convergent iterative techniques suitable for larger-scale challenges. We propose two numerical strategies that leverage the structured linearization: inverse iteration and residual inverse iteration. We will discuss how exploiting symmetry within the MEP can enhance the reliability and reduce the computational demands of these methods. Our presentation will include numerical examples that confirm our theoretical insights, along with a third example that demonstrates the efficacy of a hybrid approach combining both proposed methods.

References

[1] R. Claes, E. Jarlebring, K. Meerbergen, P. Upadhyaya. Linearizability of eigenvector nonlinearities. *SIAM J. Matrix Anal. Appl.*, 43:764-786, 2022.