METT VII — 7th Workshop on Matrix Equations and Tensor Techniques

Pisa, 13-14 Feb 2017

Program

Monday, February 13

8:30 Registration

Session I (chair: B. Meini)

9:00 Opening

- 9:15 P. Kürschner: Large generalized Sylvester equations with additional low-rank terms
- **9:40** G. Mele: Krylov methods for generalized Sylvester equations with low rank commutative coefficients
- 10:05 E. Ringh: Sylvester-based preconditioner for the waveguide eigenvalue problem

10:30 Coffee break

Session II (chair: F. Poloni)

- 11:00 V. Kazeev: Improved convergence bounds for the tensor-structured multilevel approximation of polynomial and piecewise-analytic functions
- 11:25 L. Perisa: Recompression of Hadamard Products of Tensors in Tucker Format
- 11:50 S. Kraemer: Consistent ALS in the TT-Format for Rank-Adaptive Tensor Completion

12:15 Free time (lunch)

Session III (chair: H. Fassbender)

- **14:30** B. Baran: Differential Riccati Equations for the Feedback Control of Complex Flows with Moving Interfaces
- 14:55 A. Koskela: Differential Riccati Equations: a matrix exponential approach
- 15:20 I. Gosea: Model reduction of quadratic-bilinear systems in the Loewner framework

Session IV (chair: V. Simoncini)

15:55 Poster blitz

16:15 Coffee break + Poster Session

20:00 Conference dinner

List of posters in the poster session

- 1. A. Bernardi: Tensor decomposition and homotopy continuation
- 2. G. Heidel: Second Order Riemannian Methods for Low-Rank Tensor Completion
- 3. J. Heiland: POD Reduced Tensor Bases in Optimal Control of PDEs
- 4. M. Hund: Parametrized Sylvester Equations in Model Order Reduction
- 5. D. Kressner: Low-rank approximation for tensor structured Markov chains
- 6. I. Kuzmanović: Parameter dependent generalized Sylvester and T-Sylvester equations
- 7. C. Löbbert: Parallel Tensor Arithmetic for iterative solvers
- 8. S. Massei: Solving quadratic matrix equations with semi-infinite quasi-Toeplitz coefficients
- 9. D. Palitta: Large-scale matrix equations with banded right-hand side
- 10. L. Robol: On the decay of the off-diagonal singular values in cyclic reduction
- 11. J. Roman: Low-rank solution of linear matrix equations with SLEPc
- 12. D. Sun: A flexible variant of block GMRES method for solving linear systems with multiple shifts and multiple right-hand sides
- 13. M. Sutti: Multiple shooting for the Stiefel geodesic distance

Tuesday, February 14

Session V (chair: P. Benner)

- **9:15** M. Köhler: Efficient Implementation BLAS Level-3 solvers for Sylvester-type Matrix Equations
- 9:40 F. Poloni: When is a system of (star)-Sylvester equations well posed?
- 10:05 N. Mastronardi: The revisited implicit Q-Theorem and the fast computation of roots of polynomials
- 10:30 Coffee break

Session VI (chair: D. Bini)

- 11:00 P. Boito: Efficient Solution of Parameter Dependent Quasiseparable Systems and Computation of Meromorphic Matrix Functions
- 11:25 M. Fasi: A Schur algorithm for rational matrix equations

11:50 R. Luce: Fast Toeplitz Matrix Exponential

12:15 Free time (lunch)

Session VII (chair: L. Grasedyck)

- 14:30 I. Domanov: On the largest multilinear singular values of higher-order tensors
- 14:55 N. Vannieuwenhoven: A Riemannian optimization method for computing tensor rank decompositions
- 15:20 B. Meini: Perron-based algorithms for the multilinear pagerank

15:55 Closing

Useful Information

The conference will take place in the *aula magna* of the Department of Mathematics of the University of Pisa, Largo B. Pontecorvo 5, in the northeastern part of the city centre. The Eduroam wi-fi network is available in the department.

The conference dinner will take place in Ristorante La Clessidra, Via del Castelletto 26/30, within walking distance from the conference venue. A map is enclosed below.



List of Abstracts

Talks

B. Baran: Differential Riccati Equations for the Feedback Control of Complex Flows with Moving Interfaces Problems featuring moving interfaces appear in many applications. They can model solidification and melting of pure materials, crystal growth and other multi-phase problems. The control of the moving interface allows to, for example, influence production processes and, thus, the product material quality.

We consider the two-phase Stefan problem that models a solid and a liquid phase separated by the moving interface. In the liquid phase, the heat distribution is characterized by a convectiondiffusion equation. The fluid flow in the liquid phase is described by the Navier–Stokes equations which introduces a differential algebraic structure to the system. The interface movement is coupled with the temperature through the Stefan condition, which adds additional algebraic constraints. Our formulation uses a sharp interface representation and we define a quadratic tracking-type cost functional as a target of a control input.

We compute an open loop optimal control for the Stefan problem using an adjoint system. For a feedback representation, we linearize the system about the trajectory defined by the open loop control. This results in a linear-quadratic regulator problem, for which we formulate the differential Riccati equation with time varying coefficients. This Riccati equation defines the corresponding feedback gain.

In this talk, we show how the differential Riccati equation is derived. Further, we present the feedback formulation that takes into account the structure and the differential algebraic components of the problem. Also, we discuss how the complexities that come, for example, with mesh movements, can be handled in a feedback setting.

Authors: Björn Baran (Max Planck Institute for Dynamics of Complex Technical Systems), Peter Benner (Max Planck Institute for Dynamics of Complex Technical Systems), Jan Heiland (Max Planck Institute for Dynamics of Complex Technical Systems), Jens Saak (Max Planck Institute for Dynamics of Complex Technical Systems)

P. Boito: Efficient Solution of Parameter Dependent Quasiseparable Systems and Computation of Meromorphic Matrix Functions In this work we focus on the solution of shifted quasiseparable systems and of more general parameter dependent matrix equations with quasiseparable representations. We propose an efficient algorithm exploiting the invariance of the quasiseparable structure under diagonal shifting and inversion. This algorithm is applied to compute various functions of matrices. Numerical experiments show the effectiveness of the approach.

This is joint work with Yuli Eidelman and Luca Gemignani. Authors: Paola Boito (XLIM-Université de Limoges, LIP-ENS Lyon)

I. Domanov: On the largest multilinear singular values of higher-order tensors Let σ_n denote the largest mode-*n* multilinear singular value of an $I_1 \times \cdots \times I_N$ tensor \mathcal{T} . We prove that

 $\sigma_1^2 + \dots + \sigma_{n-1}^2 + \sigma_{n+1}^2 + \dots + \sigma_N^2 \le (N-2) \|\mathcal{T}\|^2 + \sigma_n^2, \quad n = 1, \dots, N.$

We also show that at least for third-order cubic tensors the inverse problem always has a solution. Namely, for each σ_1 , σ_2 and σ_3 that satisfy

$$\sigma_1^2 + \sigma_2^2 \le \|\mathcal{T}\|^2 + \sigma_3^2, \quad \sigma_1^2 + \sigma_3^2 \le \|\mathcal{T}\|^2 + \sigma_2^2, \quad \sigma_2^2 + \sigma_3^2 \le \|\mathcal{T}\|^2 + \sigma_1^2,$$

and the trivial inequalities $\sigma_1 \geq \frac{1}{\sqrt{n}} \|\mathcal{T}\|$, $\sigma_2 \geq \frac{1}{\sqrt{n}} \|\mathcal{T}\|$, $\sigma_3 \geq \frac{1}{\sqrt{n}} \|\mathcal{T}\|$, there always exists an $n \times n \times n$ tensor whose largest multilinear singular values are equal to σ_1 , σ_2 and σ_3 . We also show that if the equality $\sigma_1^2 + \sigma_2^2 = \|\mathcal{T}\|^2 + \sigma_3^2$ holds, then \mathcal{T} is necessarily equal to a sum of multilinear rank- $(L_1, 1, L_1)$ and multilinear rank- $(1, L_2, L_2)$ tensors and we give a complete description of all its multilinear singular values.

We establish a connection with honeycombs and eigenvalues of the sum of two Hermitian matrices. This seems to give at least a partial explanation of why results on the joint distribution of ML singular values are scarce.

Authors: Lieven De Lathauwer (KU Leuven), Ignat Domanov (KU Leuven)

M. Fasi: A Schur algorithm for rational matrix equations We consider rational matrix equations of the form $p(X) q(X)^{-1} = A$, where A is a complex square matrix and p and q are polynomials. It is easy to see that any solution X also satisfies p(X) = A q(X). When taking into account only primary solutions, the other implication is also true, and thus we can focus our attention on the latter simpler equation.

We develop a novel Schur method for the computation of primary solutions of the aforementioned equations, which generalises the algorithm of Smith [SIAM. J. Matrix Anal. & Appl., 24 (2003), pp. 971–989] for the computation of primary pth roots, and has a similar computational cost. We show that the algorithm can be implemented using only real arithmetic if A is real, and discuss the existence of real solutions.

Authors: Massimiliano Fasi (School of Mathematics - The University of Manchester), Bruno Iannazzo (Università di Perugia)

I. Gosea: Model reduction of quadratic-bilinear systems in the Loewner framework We introduce a data-driven model order reduction (MOR) approach which can be viewed as the generalization of the Loewner framework for quadratic-bilinear (QB) control systems. Reduction of such systems have been treated extensively in recent years since, for certain types of nonlinear systems, one can always find an equivalent QB model without performing any approximation.

The first step is to give an appropriate definition of generalized higher order transfer functions (which are multi-variate rational functions) for the class of QB systems. Proceed by constructing system matrices using only measured/computed data (samples of the new transfer functions). The advantage is that the approach is data-driven since one would only need data points to construct a reduced order QB system.

Then, provide appropriate generalizations of the Loewner matrices which can be directly obtained by solving generalized Sylvester equations with quadratic terms. We are concerned with optimally solving these equations, which turns out it can be achieved via tensor calculus techniques. We avoid explicit computations of Kronecker products between large scale matrices by using various properties of tensor matricizations.

We illustrate the practical applicability of the proposed method by means of several numerical experiments resulting from semi-discretized nonlinear partial differential equations (such as Burgers' and Chafee-Infante) as well as the FitzHugh-Nagumo system. Compare the results of our method against the results obtained when applying some recent proposed extensions of balanced truncation and IRKA for QB systems.

Authors: Athanasios C. Antoulas (*Rice University Houston*), Ion Victor Gosea (*Jacobs University Bremen*)

V. Kazeev: Improved convergence bounds for the tensor-structured multilevel approximation of polynomial and piecewise-analytic functions Tensor decompositions, separating variables in high-dimensional arrays, are mostly seen as an efficient tool for lifting the so-called «curse of dimensionality» in high-dimensional settings. In such cases, polynomial approximations are often used in the «physical» and, possibly, parametric spaces to translate a continuous problem into a tensor problem. Such approximations are however known to converge slowly when the solution has low regularity or when the regularity constants are prohibitively large, which is the case for singular and highly oscillatory solutions. This poses a challenge even in low dimensions. Many discretizations have been developed for such problems, such as the adaptive hp-FEM and the heterogeneous multiscale method.

When applied to resolve the multiscale structure of the data, a tensor decomposition may be expected to realize the efficient representation of solutions while retaining a standard, problemnonspecific discretization scheme based on tensor-product grids and low-order approximation. Recently, the so-called quantized-tensor-train (QTT) decomposition has been shown to achieve exponential convergence (with respect to the number of effective degrees of freedom) in certain problems with singular and highly oscillatory solutions. These convergence results build, in particular, upon the tensor-rank bound for the polynomial functions by Grasedyck.

This talk presents refined bounds for the QTT approximation of polynomial and piecewiseanalytic functions, based on the subspace interpretation of the QTT decomposition. For a piecewiseanalytic function, a characterization of the ranks and accuracy of its QTT approximations in terms of its holomorphy is given and a quasi-optimal adaptive algorithm for QTT approximation with certified accuracy in function norms is proposed.

Authors: Vladimir Kazeev (University of Geneva)

M. Köhler: Efficient Implementation BLAS Level-3 solvers for Sylvester-type Matrix Equations The solution of dense Sylvester equations has been considered a solved problem after the seminal paper by Bartels and Stewart. Variants of their algorithm have increased the efficiency by rearranging the operations, especially for the generalized equations. None of the direct approaches goes beyond a BLAS level-2 formulation, except of the Generalized Schur method of Kågström and Poromaa from 1994. On the other hand, the RECSY family of methods works on top of a recursive subdivision into smaller problems. RECSY can thus accelerate the solution process on many processors. Still, it has a rather irregular memory access pattern, especially when the subdivision is not equally sized.

In our contribution, we present a BLAS level-3 type tiling approach for the solution of (generalized) Sylvester equations, including the (generalized) Lyapunov and Stein equations, after the reduction of the coefficients to Schur form. The methods thus accelerate the triangular solution phase, compared to the BLAS level-2 methods, by the use of different efficiency improving implementation techniques. On the one hand, the proposed BLAS level-3 solvers act as an outer iteration that exploits the existing solvers for solving the arising block problems on the single tiles. On the other hand, we optimize the existing BLAS level-2 solvers with respect to the capabilities of modern computer architectures. Beside this classical BLAS level-3/level-2 type tiling approaches, we combine the newly evolved methods with features from the latest OpenMP 4/4.5 standard to get parallel implementation which works independent of the parallel features provided by the BLAS library.

When the RECSY solver is used as the inner method, the outer iteration contributes a largely regularized memory access pattern. Therefore also here we get a notable acceleration with respect to the standalone solver.

For a single equation the runtime is, however, dominated by the reduction of the coefficient matrices to Hessenberg/Schur form. Therefore, the overall performance gain is limited. However, the gain can be increased, e.g, when the equations with the same coefficient needs to be solve for several right hand sides. This is for example the case if autonomous differential matrix equations are solved or the condition number of the operator should be estimated.

Authors: Martin Köhler (*Max Planck Institute for Dynamics of Complex Technical Systems*), Jens Saak (*Max Planck Institute for Dynamics of Complex Technical Systems*)

A. Koskela: Differential Riccati Equations: a matrix exponential approach The numerical treatment of linear quadratic regulator, linear quadratic Gaussian design and stochastic control problems of certain type require solving Riccati equations. In the finite time horizon case, the differential Riccati equation arises.

We consider symmetric differential Riccati equations (DREs) of the form

$$X'(t) = Q + X(t)A + A^{T}X(t) - X(t)SX(t), \quad X(t_{0}) = X_{0},$$

where $t \in [0, T]$, and $Q, A, S \in \mathbb{R}^{n \times n}$. Moreover, $Q^T = Q$ and $S^T = S$.

This is equivalent to solving the system of linear first-order matrix differential equations

$$\frac{d}{dt} \begin{bmatrix} U(t) \\ V(t) \end{bmatrix} = \begin{bmatrix} -A & S \\ Q & A^T \end{bmatrix} \begin{bmatrix} U(t) \\ V(t) \end{bmatrix}, \quad \begin{bmatrix} U(0) \\ V(0) \end{bmatrix} = \begin{bmatrix} U_0 \\ V_0 \end{bmatrix}$$

where $U(t), V(t) \in \mathbb{R}^{n \times n}, V_0 U_0^{-1} = X_0$ for some $U_0 \in \mathbb{R}^{n \times n}$ invertible and $V_0 \in \mathbb{R}^{n \times n}$.

The first-order system above is Hamiltonian, i.e., for the coefficient matrix

$$H = \begin{bmatrix} -A & S \\ Q & A^T \end{bmatrix}$$

it holds that $(JH)^T = JH$, where

$$J = \begin{bmatrix} 0 & I \\ -I & 0 \end{bmatrix}.$$

Therefore also a certain energy functional is invariant along the solution.

We propose efficient methods for the large-scale DRE which are based on approximating the solution, i.e., the product $\exp(tH)W_0$, where $W_0 = \begin{bmatrix} U_0 \\ V_0 \end{bmatrix}$, in the Krylov subspace

$$\mathcal{K}_m(H, W_0) = \text{span}\{W_0, HW_0, \dots, H^{m-1}W_0\}.$$

The methods we consider also preserve the Hamiltonian structure of the original equation.

In the low-rank case $Q = BB^T$, $S = CC^T$, $B, C \in \mathbb{R}^{n \times k}$, $k \ll n$, one particular Krylov subspace approach turns out to be very efficient. This method is fully structure preserving and is based on making suitable projections using a basis of the subspace

$$\mathcal{K}_m(A,C) \cup \mathcal{K}_m(A^T,B) \cup \mathcal{K}_m(A,U_0) \cup \mathcal{K}_m(A^T,V_0)$$

Authors: Antti Koskela (KTH Royal Institute of Technology), Hermann Mena (Yachay Tech)

S. Kraemer: Consistent ALS in the TT-Format for Rank-Adaptive Tensor Completion Low rank tensor completion is a hard task if either the given data is noisy or the underlying data

model is not accurate. In both cases some sort of regularization is required in order to solve the highly ill-posed inverse problem. In this article we focus on the calibration of the data model.

For alternating optimization, we observe that neither thresholding nor existing rank adaption methods enable a continuous transition between manifolds of different ranks. We denote this flaw as *inconsistency with truncation*. As a consequence of this flaw, arbitrarily small changes in the singular values of an iterate can have arbitrarily large influence on the further reconstruction. We therefore introduce a singular value based regularization to the standard alternating least squares (ALS), which is motivated by averaging in micro-steps. We prove its *consistency* and derive a natural semi-implicit rank adaption strategy. We further prove that the standard ALS microsteps are only continuous on manifolds of fixed ranks, and only around points that have what we define as *internal tensor restricted isometry property iTRIP*. Finally, we provide numerical examples that show improvements of the reconstruction quality up to orders of magnitude in the new C(onsistent)ALS compared to standard ALS.

Authors: Lars Grasedyck (IGPM, RWTH Aachen), Sebastian Kraemer (IGPM, RWTH Aachen)

P. Kürschner: Large generalized Sylvester equations with additional low-rank terms We consider multi-term Sylvester equations $A_1X + XA_2 + \sum_{i=1}^{L} N_iXM_i + BC^T = 0$, where the matrices B, C and N_i, M_i in the additional L summands are of very small rank. For small-sized equations of this form, a direct solution approach is presented in [1] and is related to a formal generalization of the Sherman-Morrison-Woodbury formula. Our ultimate goal is to make this approach also efficiently applicable in the large-scale case. State-of-the-art low-rank methods for standard two-term Sylvester equations, e.g. rational Krylov or low-rank ADI methods, are one important cornerstone for achieving this. The proposed numerical method will be tested and compared to other existing algorithms for large-scale multi-term Sylvester and Lyapunov equations, e.g., splitting based iteration schemes [1,2] and low-rank matrix valued Krylov methods [1,3]. If time permits, special cases with more than two full rank terms will be discussed.

References:

[1] T. Damm: Direct methods and ADI-preconditioned Krylov subspace methods for generalized Lyapunov equations, Numer. Lin. Alg. Appl., 15, pp. 853-871, 2008.

[2] S. D. Shank, V. Simoncini, D. B. Szyld: Efficient low-rank solutions of generalized Lyapunov equations, Numer. Math., 134(2), pp. 327-342, 2016.

[3] D. Kressner, C. Tobler: Krylov subspace methods for linear systems with tensor product structure, SIAM J. Matrix. Anal. and Appl., 31(4), pp. 1688-1714, 2010.

Authors: Patrick Kürschner (Max Planck Institute for Dynamics of Complex Technical Systems)

R. Luce: Fast Toeplitz Matrix Exponential The computation of the matrix exponential is a ubiquitous operation in numerical mathematics, and for a general, unstructured $n \times n$ matrix it can be computed in $O(n^3)$ operations. An interesting problem arises if the input matrix is a Toeplitz matrix, for example as the result of discretizing integral equations with a time invariant kernel. In this case it is not obvious how to take advantage of the Toeplitz structure, as the exponential of a Toeplitz matrix is, in general, not a Toeplitz matrix itself. The main contribution of this work is an algorithm of quadratic complexity for the computation of the Toeplitz matrix exponential. It is based on the scaling and squaring framework, and connects classical results from rational approximation theory to matrices of low displacement rank. As an example, the developed methods are applied to Merton's jump-diffusion model for option pricing.

Authors: Daniel Kressner (EPFL), Robert Luce (EPF Lausanne)

N. Mastronardi: The revisited implicit Q-Theorem and the fast computation of roots of polynomials In this talk we revisit the Implicit-Q Theorem and analyze the the problem of performing a QR-step on an unreduced Hessenberg H matrix when we know an "exact" eigenvalue λ_0 of H. Under exact arithmetic, this eigenvalue will appear on diagonal of the transformed Hessenberg matrix \tilde{H} and will be decoupled from the remaining part of the Hessenberg matrix, thus resulting in a deflation. But it is well known that in finite precision arithmetic the so-called perfect shift could get blurred and the eigenvalue λ_0 can not be deflated and/or is perturbed significantly. In this talk we develop a new strategy for computing such a QR step so that the deflation is indeed successful. The method is based on the preliminary computation of the corresponding eigenvector \mathbf{x} such that the residual $(H - \lambda_0 I)\mathbf{x}$ is sufficiently small. The eigenvector is then transformed to a unit vector by a sequence of Givens transformations, which are also performed on the Hessenberg matrix. Such a QR step is the basic ingredient of the QR method to compute the Schur form, and hence the eigenvalues of an arbitrary matrix.

In particular, the new technique will be applied to compute the eigenvalues of companion matrices in a fast way.

Authors: Nicola Mastronardi (Istituto per le Applicazioni del Calcolo, Consiglio Nazionale delle Ricerche), Paul Van Dooren (Catholic University of Louvain, Department of Mathematical Engineering, Avenue Georges Lemaitre 4, B-1348 Louvain-la-Neuve, Belgium)

B. Meini: Perron-based algorithms for the multilinear pagerank We consider the equation

$$x = \alpha P(x \otimes x) + (1 - \alpha)v$$

where $\alpha \in [0, 1]$, $x, v \in \mathbb{R}^n_{\geq 0}$ (the vector x being the unknown), $P \in \mathbb{R}^{n \times n^2}_{\geq 0}$, with v and P stochastic (in a suitable sense). This problem arises in the computation of *multilinear pagerank*, an approximation to the invariant measure of a second-order Markov chain considered, for instance, in [Gleich, Lim, Yu, SIMAX 2015]. We study the equation and the properties of its nonnegative solutions drawing a parallel with the very similar setting of *quadratic vector equations* considered

in [P., LAA 2013]. The study reveals the role of the minimal nonnegative solution, which is not usually considered because often it is not stochastic and has no immediate applicative interest.

We adapt to this problem the iterative algorithms based on computing at each step a Perron eigenvector, suggested in [Meini, Poloni, SIMAX 2011] and [Bini, Meini, Poloni, NLAA 2011]. **Authors:** Beatrice Meini (University of Pisa), Federico Poloni (Dipartimento di Informatica, Università di Pisa)

G. Mele: Krylov methods for generalized Sylvester equations with low rank commutative coefficients We consider the following generalized Sylvester equation

$$AX + XB^T + \sum_{i=1}^m N_i XM_i^T = C$$

and the special case of the generalized Lyapunov equation

$$AX + XA^T + \sum_{i=1}^m N_i X N_i^T = C.$$

These matrix equations arise in applications such as solving discretized PDEs and control theory. In control theory, the solution to a generalized Lyapunov equation represents the Gramian associated to a bilinear system. The singular values of the Gramian give quantifiable information on reachability and observability of the states. This it is used to perform model order reduction of the bilinear system based on balanced truncation [1].

Under the assumption that the matrices N_i and C have low rank, by using a matrix equation version of the Sherman-Morrison-Woodbury formula, it is possible to show that the solution of the generalized Lyapunov equation has a fast decay of the singular values and there are efficient algorithms for computing a low rank approximation of the solution [2].

To our knowledge, the general case of the generalized Sylvester equation with no assumption on the rank of the matrices M_i, N_i and C, is not well understood. We denote by $\mathcal{L}(X) := AX + XB$ the Sylvester part, and by $\Pi(X) := \sum_{i=1}^{m} N_i X M_i^T$ the rest. Under the assumption that $\rho(\mathcal{L}^{-1}\Pi) < 1$, we show that the solution to the Sylvester equation can be expressed by means of a Neumann series expansion and a strong singular value decay is expected. Moreover, we propose a new Krylov method for computing a low rank approximation of the solution to the generalized Sylvester equation under the condition that the commutators $[A, N_i] = AN_i - N_iA, [B, M_i] = BM_i - M_iB$ and C have low rank. The effectiveness of the proposed method is illustrated by solving generalized Sylvester and Lyapunov equations associated to a discretization of a PDE and to a MIMO (multiple input multiple output) problem arising in control theory.

[1] T. Damm, Lyapunov equations, energy functionals, and model order reduction of bilinear and stochastic systems. Control and Optimization 49.2 (2011): 686-711.

[2] P. Benner, T. Breiten, Low rank methods for a class of generalized Lyapunov equations and related issues. Numerische Mathematik 124.3 (2013): 441-470.

Authors: Elias Jarlebring (*KTH Royal Institute of Technology*), Giampaolo Mele (*KTH Royal Institute of Technology*), Davide Palitta (*Dipartimento di Matematica, Universita' di Bologna*), Emil Ringh (*KTH Royal Institute of Technology*)

L. Perisa: Recompression of Hadamard Products of Tensors in Tucker Format The Hadamard product features prominently in tensor-based algorithms in scientific computing and data analysis. Due to its tendency to significantly increase ranks, the Hadamard product can represent a major computational obstacle in algorithms based on low-rank tensor representations. It is therefore of interest to develop recompression techniques that mitigate the effects of this rank increase. In this work, we investigate such techniques for the case of the Tucker format, which is well suited for tensors of low order and small to moderate multilinear ranks. Fast algorithms are attained by combining iterative methods, such as the Lanczos method and randomized algorithms, with fast matrix-vector products that exploit the structure of Hadamard products. The resulting complexity reduction is particularly relevant for tensors featuring large mode sizes I and small to

moderate multilinear ranks R. To implement our algorithms, we have created a new Julia library for tensors in Tucker format.

Authors: Daniel Kressner (EPFL), Lana Perisa (University of Split)

F. Poloni: When is a system of (star)-Sylvester equations well posed? We study systems of Sylvester, \top -Sylvester, or *-Sylvester equations, i.e., systems of equations of the form

$$\begin{cases} A_1 X_{i_1}^{s_1} B_1 + C_1 X_{j_1}^{t_1} D_1 = E_1 \\ A_2 X_{i_2}^{s_2} B_2 + C_2 X_{j_2}^{t_2} D_2 = E_2 \\ \vdots \\ A_r X_{i_r}^{s_r} B_r + C_r X_{j_r}^{t_r} D_r = E_r, \end{cases}$$

where the unknowns X_1, X_2, \ldots, X_r and the coefficients A_k, B_k, C_k, D_k, E_k (for $k = 1, 2, \ldots, r$) are $n \times n$ matrices, the indices i_k and j_k are in $1, 2, \ldots, r$, and each of the symbols s_k, t_k is either $1, \top$ (complex transpose) or * (conjugate transpose).

We are interested in determining for which values of A_k, B_k, C_k, D_k (k = 1, 2, ..., r) the system has a unique solution for each choice of the right-hand sides $E_k, k = 1, 2, ..., r$ (unique solvability for each right-hand side, or *well-posedness*). This is equivalent to determining when the linear system obtained after vectorization has an invertible matrix.

We make several transformations to reduce the problem to a simpler form, and then we give a criterion to check well-posedness by working directly on $n \times n$ matrices, without Kronecker products and vectorization. This procedure generalizes analogous criteria that exist for single equations and simpler cases, which are based on the spectral properties of suitable matrices and pencils.

Our strategy is constructive, and can be used to find a $O(rn^3)$ solution algorithm to solve these systems.

Authors: Fernando De Terán (Universidad Carlos III de Madrid), Bruno Iannazzo (Università di Perugia), Federico Poloni (Dipartimento di Informatica, Università di Pisa), Leonardo Robol (KU Leuven)

E. Ringh: Sylvester-based preconditioner for the waveguide eigenvalue problem We consider a nonlinear eigenvalue problem (NEP) arising from absorbing boundary conditions in the study of a partial differential equation (PDE) describing a waveguide. We propose a new computational approach for this large-scale NEP based on residual inverse iteration with preconditioned iterative solves. Similar to many preconditioned iterative methods for discretized PDEs, this approach requires the construction of an accurate and efficient preconditioner. For the waveguide eigenvalue problem, the associated linear system can be formulated as a generalized Sylvester equation $AX + XB + A_1XB_1 + A_2XB_2 + K \circ X = C$, where \circ denotes the Hadamard product. The equation is approximated by a low-rank corrected Sylvester equation, which we use as a preconditioner. The action of the preconditioner is efficiently computed by using the matrix equation version of the Sherman-Morrison-Woodbury (SMW) formula. We show how the preconditioner can be integrated into residual inverse iteration, and complexity moved between a precomputational-and iterative-phase. The results are illustrated by applying the method to large-scale problems.

Authors: Elias Jarlebring (*KTH Royal Institute of Technology*), Johan Karlsson (*KTH Royal Institute of Technology*), Giampaolo Mele (*KTH Royal Institute of Technology*), Emil Ringh (*KTH Royal Institute of Technology*)

N. Vannieuwenhoven: A Riemannian optimization method for computing tensor rank decompositions Riemannian optimization is the natural framework for optimization problems on smooth Riemannian manifolds. In the context of tensor decompositions such methods have already been considered for the Tucker decomposition by Kressner, Steinlechner and Vandereycken (2014) and for the Hierarchical Tucker and Tensor Trains decompositions by Uschmajew and Vandereycken (2013), Lubich et al. (2013), and Da Silva and Herrmann (2016). Hitherto no Riemannian optimization method exists for computing a tensor rank decomposition. The reason for this is that the collection of real tensors of fixed rank is not a smooth manifold but rather a semialgebraic set. Such sets admit a singular locus where no local diffeomorphism to a Euclidean space exists. Current research for alleviating this problem in the context of tensor rank decompositions for matrices appears targeted at optimizing instead over a desingularization of the semi-algebraic set.

In this presentation, we explore a much simpler approach: the singular locus of the semialgebraic set can be avoided by modifying the cost function so that we optimize instead over the Cartesian product of several smooth Segre varieties. We show that the rank-1 ST-HOSVD approximation (Vannieuwenhoven et al., 2012) is retraction operator for the Segre variety that can be computed efficiently. Then, a Riemannian Gauss-Newton (RGN) method with dogleg trust region is proposed, which employs the aforementioned retraction. It is shown that the radius and rate of convergence of this RGN method are intimately related to the condition number proposed by Breiding and Vannieuwenhoven (2016) of the corresponding tensor rank decomposition problem.

This is joint work with Paul Breiding (TU Berlin). **Authors:** Paul Breiding (*TU Berlin*), Nick Vannieuwenhoven (*KU Leuven*)

Posters

A. Bernardi: Tensor decomposition and homotopy continuation In this talk I will present a joint work with N.S. Daleo, J.D. Hauenstein and B. Mourrain where we we develop computational methods for computing ranks and border ranks of tensors along with decompositions.

Authors: Alessandra Bernardi (University of Trento), Noah Daleo (Worcester State University), Jonathan Hauenstein (Notre Dame), Bernardi Mourrain (INRIA)

G. Heidel: Second Order Riemannian Methods for Low-Rank Tensor Completion The goal of tensor completion is to fill in missing entries of a partially known tensor (possibly inculding some noise) under a low-rank constraint. This may be formulated as a least-squares problem. The set of tensors of a given multilinear rank is known to admit a Riemannian manifold structure, thus methods of Riemannian optimization are applicable.

In our work, we derive the Riemannian Hessian of an objective function on the low-rank tensor manifolds and discuss the convergence properties of second order methods for the tensor completion problem, both theoretically and numerically. We compare our approach to Riemannian tensor completion methods from recent literature. Our examples include the recovery of multidimensional images, approximation of multivariate functions and recovery of partially missing data from survey statistics.

Authors: Gennadij Heidel (Trier University), Volker Schulz (Trier University)

J. Heiland: POD Reduced Tensor Bases in Optimal Control of PDEs We consider partial differential equations (PDE) on a space-time cylinder $\Omega \times (0,T)$, where Ω is a domain in \mathbb{R}^n and where T > 0. We seek to approximate the solution $x \in L^2(0,T;L^2(\Omega))$ to the PDE in the finite-dimensional space-time tensor space $S \cdot \mathcal{Y}$, where S is a finite-dimensional subspace of $L^2(0,T)$ and \mathcal{Y} is a finite-dimensional subspace of $L^2(\Omega)$. The naive space-time Galerkin projection of the PDE onto $S \cdot \mathcal{Y}$ would lead to a (possibly nonlinear) equation system of dimension equal to the dimension of $S \cdot \mathcal{Y}$ which is prohibitively large for generic ansatz spaces. In this talk, we show how, based on measurements of the solution x, one can find low-dimensional subspaces of S and \mathcal{Y} that optimally parametrize the solution and that give rise to a space-time Galerkin scheme of lower dimension. By means of a Burgers' equation, we show how this approach, which is a generalization of the well-know POD method for model reduction, can be used for efficiently controlling nonlinear PDEs.

Authors: Peter Benner (Max Planck Institute for Dynamics of Complex Technical Systems), Jan Heiland (Max Planck Institute for Dynamics of Complex Technical Systems)

M. Hund: Parametrized Sylvester Equations in Model Order Reduction Sylvester equations of the form

$$AVE + EVA = BL$$

are a key player in model order reduction (MOR) of linear input-output systems as

$$E\dot{x}(t) = Ax(t) + Bu(t),$$
$$y(t) = Cx(t),$$

with a regular matrix $E \in \mathbb{R}^{n \times n}$, as well as matrices $A \in \mathbb{R}^{n \times n}$, $B \in \mathbb{R}^{n \times p}$, $C \in \mathbb{R}^{q \times n}$ and $\tilde{A}, \tilde{E} \in \mathbb{R}^{m \times m}$, where $m \ll n$.

These arise in the context of the \mathcal{H}_2 optimality conditions of Wilson, the variation of the orthogonal polynomial time domain MOR framework of Jiang and Chen, as well as in classical moment matching.

There are different ways parametrized Sylvester equations may arise in MOR. On the one hand, in the time domain approach, one ends up with the above mentioned Sylvester equation, where \tilde{A} is the identity $I_m \in \mathbb{R}^{m \times m}$. The solution of this Sylvester equation spans the same Krylov subspace as moment matching with the generalized eigenvalues of (I_m, \tilde{E}) as expansion points. Now, when using parametrized families of Jacobi polynomials, a Sylvester equation

$$AV\tilde{E}(\mu) + EV = BL$$

arises, that is parametrized with $\mu \in \mathbb{R}^2$ in the small matrix $\tilde{E}(\mu)$.

On the other hand, considering parametrized dynamical systems

$$E(\mu)\dot{x}(t) = A(\mu)x(t) + B(\mu)u(t)$$

$$y(t) = C(\mu)x(t),$$

with parameter $\mu \in \mathbb{R}^d$, provides a Sylvester equation that is now parametrized in its large, but sparse matrices

$$A(\mu)V\tilde{E} + E(\mu)V\tilde{A} = B(\mu)L.$$

Our final goal is to solve the \mathcal{H}_2 optimal MOR problem via a fixed point iteration on the Wilson conditions along the lines of the two sided iterative approximation (TSIA) algorithm. This could in turn serve as an alternative parametrized \mathcal{H}_2 optimal MOR algorithm to the 'parametric IRKA''. In our contribution, we will present first results towards this goal.

Authors: Manuela Hund (Max Planck Institute for Dynamics of Complex Technical Systems), Jens Saak (Max Planck Institute for Dynamics of Complex Technical Systems)

D. Kressner: Low-rank approximation for tensor structured Markov chains Stochastic models that describe interacting processes, such as stochastic automata networks, feature a dimensionality that grows exponentially with the number of processes. This state space explosion severely impairs the use of standard methods for the numerical analysis of such Markov chains. In this talk, we discuss the approximation of solutions by matrix product states or, equivalently, by tensor train decompositions. Our approach significantly extends existing approaches based on product form solutions and can, in principle, attain arbitrarily high accuracy. Numerical experiments demonstrate that the newly proposed algorithms are particularly well suited to deal with pairwise neighbor interactions. For more general models, the growth of the involved ranks and tensor sizes makes it necessary to more carefully consider the underlying algorithms and incorporate preconditioners. We propose a combination of multigrid preconditioners with low-rank tensor approximation for this purpose. This talk is based on joint work with Francisco Macedo, Matthias Bolten, Karsten Kahl, and Sonja Sokolović.

Authors: Daniel Kressner (*EPFL*)

I. Kuzmanović: Parameter dependent generalized Sylvester and *T*-Sylvester equations In this talk Sherman—-Morrison—-Woodbury—type formula for the solution of generalized Sylvester equation of the form

$$(A_0 + \mu U_1 V_1) X E_1 + E_2 X (B_0 + \mu U_2 V_2) = E$$

as well as for generalized T-Sylvester equation of the form

$$(A_0 + \mu U_1 V_1) X E_1 + E_2 X^T (B_0 + \mu U_2 V_2) = E,$$

will be presented, where U_1 , U_2 , V_1 and V_2 are low-rank matrices and $\mu \in \mathbb{R}$ is parameter. These formulas can be used to obtain algorithms for efficient calculation of the solution of structured generalized Sylvester and *T*-Sylvester equations for many values of parameter μ as well as for optimization of the solution with respect to μ .

Authors: Ivana Kuzmanović (Department of Mathematics, University of Osijek), Ninoslav Truhar (Department of Mathematics, University of Osijek)

C. Löbbert: Parallel Tensor Arithmetic for iterative solvers Solutions of high-dimensional parameter-dependent problems can be approximated as tensors in the data-sparse Hierarchical Tucker Format, if the parameter dependencies fulfill some low rank property. We developed parallel algorithms which perform standard arithmetic operations like the addition or the dot product between tensors directly in the data-sparse Hierarchical Tucker Format. Furthermore an operator can be applied to the data-sparse tensor. The parallel algorithms result in a time to solution which grows like log(d), where d is the dimension of the tensor (i.e. number of parameters). We present results on using parallel tensor arithmetic for iterative methods in order to solve parameter dependent problems directly in the data-sparse format.

Authors: Lars Grasedyck (RWTH Aachen), Christian Löbbert (RWTH Aachen)

S. Massei: Solving quadratic matrix equations with semi-infinite quasi-Toeplitz coefficients We consider the problem of solving matrix equations of the kind $A_{-1} + A_0X + A_1X^2 = 0$ where A_{-1}, A_0, A_1 are semi-infinite Quasi-Toeplitz (QT) matrices. A matrix A is QT if it can be written in the form T(a) + E, where $T(a) = (t_{i,j})_{i,j\in\mathbb{Z}^+}$ is the semi-infinite Toeplitz matrix associated with the symbol $a(z) = \sum_{j=-\infty}^{\infty} a_j z^j$ such that $\sum_{j=-\infty}^{\infty} |ja_j| < \infty$ and $E = (e_{i,j})_{i,j\in\mathbb{Z}^+}$ is such that $\sum_{i,j=1}^{+\infty} |e_{i,j}|$ is finite. This problem is encountered in the solution of certain QBD processes like the tandem Jackson queue [2].

Relying on the fact that the class of QT matrices is a Banach algebra [1] we prove that the solution X of interest of this quadratic equation belongs to the class QT as well, and analyze algorithms for its computations. In particular, we represent QT matrices with a finite number of parameters, consider algorithms based on fixed point iterations and on cyclic reduction, and analyze their convergence properties. We discuss on possible generalizations of this approach to multilevel structures. Numerical experiments are presented and discussed. Our approach is general and can be applied to other problems, like the computation of matrix functions of quasi Toeplitz matrices.

[1] Bini, D. A., Massei, S., Meini, B. (2016). Semi-Infinite Quasi-Toeplitz Matrices with Applications to QBD Stochastic Processes. arXiv preprint arXiv:1611.06337.

[2] Motyer, A. J., Taylor, P. G. (2006). Decay rates for quasi-birth-and-death processes with countably many phases and tridiagonal block generators. Advances in applied probability, 522-544. **Authors:** Dario Bini (*University of Pisa*), Stefano Massei (*Scuola Normale Superiore*), Beatrice Meini (*University of Pisa*)

D. Palitta: Large-scale matrix equations with banded right-hand side We are interested in the numerical solution of the large-scale Lyapunov equation

$$AX + XA^T = B, (1)$$

where the $n \times n$ matrices A, B are both large, sparse and banded with bandwidth β_A, β_B respectively. We suppose that A is symmetric and positive definite. If A is well-conditioned, the entries of the solution X to (1) decay in absolute value as their indeces move away from the sparsity pattern of B. This result relies on recent decay bounds for the entries of $(A \otimes I + I \otimes A)^{-1}$ where the Kronecker structure can be exploited. We show how such property can be employed to determine a memory-saving appproximation to X.

More precisely, the matrix-oriented CG method can be employed in the numerical solution of (1). We show that all the iterates of the solution process are banded matrices with a bandwidth

that linearly depends on the number of CG iterations. Therefore, for well-conditioned problems the memory requirements and computational costs of the algorithm remain of the order of n. Our description will be accompanied by numerical experiments to illustrate the performance of the proposed strategy.

Authors: Davide Palitta (Dipartimento di Matematica, Universita' di Bologna), Valeria Simoncini (Dipertimento di Matematica, Universita' di Bologna)

L. Robol: On the decay of the off-diagonal singular values in cyclic reduction It was recently observed that the singular values of the off-diagonal blocks of the matrix sequences generated by the Cyclic Reduction algorithm decay exponentially. This property can be used to solve, with a higher efficiency, certain quadratic matrix equations encountered in the analysis of queuing models.

We provide a sharp theoretical bound to the basis of this exponential decay together with a tool for its estimation based on a rational interpolation problem. Applications to solving $n \times n$ block tridiagonal block Toeplitz systems with $n \times n$ semiseparable blocks and certain generalized Sylvester equations in $\mathcal{O}(n^2 \log n)$ arithmetic operations are shown.

Authors: Dario Bini (Università di Pisa), Stefano Massei (Scuola Normale Superiore), Leonardo Robol (KU Leuven)

J. Roman: Low-rank solution of linear matrix equations with SLEPc We present preliminary developments in SLEPc to compute the low-rank solution of linear matrix equations. A new module LME has been added to SLEPc, aiming at solving a matrix equation written generically as

$$AXE + DXB = C.$$

As in the rest of SLEPc, only projection methods will be considered, so we restrict to cases where a low-rank solution exists, being C and X expressed as the outer product of tall-skinny matrices. Also, the focus is on large-scale problems and hence the solvers have parallel capabilities.

We have started with a Krylov solver for the continuous-time Lyapunov equation, $AX + XA^T = -C_1C_1^T$. Later on, solvers based on extended Krylov or rational Krylov subspaces will be added, as well as support for other variants of the equation. Our Krylov solver implementation includes an Eiermann-Ernst-type restart as proposed by Kressner [2008 IEEE Symp. CACSD].

Since SLEPc is mainly concerned with eigenvalue problems, one of our goals is to use the Lyapunov solver in the context of Lyapunov inverse iteration (see, e.g., [Ellman and Wu, SIMAX 2013]) to compute rightmost eigenvalues of matrices. We will show preliminary results for this use case.

Authors: Jose E. Roman (Universitat Politecnica de Valencia)

D. Sun: A flexible variant of block GMRES method for solving linear systems with multiple shifts and multiple right-hand sides We consider iterative solutions of linear systems with multiple right-hand sides, in which the coefficient matrices differ from each other by a scalar multiple of the identity,

$$(A - \sigma_i I)X_i = B, \text{ with } i = 1, 2, \dots, L,$$
(1)

where $A - \sigma_i I$ are square nonsingular matrices of large dimension $n, B = [b^{(1)}, b^{(2)}, \ldots, b^{(p)}]$ is a full rank matrix consisting of $p \ll n$ given right-hand sides $b^{(j)}(j = 1, \ldots, p)$, and $\sigma_i(i = 1, \ldots, L)$ are called shifts. Sequences of shifted linear systems with multiple right-hand sides arise, e.g., in lattice quantum chromo dynamics applications, PageRank problems and image restorations.

We propose a new flexible shifted block GMRES method (FDBGMRES-Sh) that solves the whole sequence of linear systems simultaneously, and accomodates for the use of variable preconditioning at every iteration step. The new algorithm detects the convergence of each linear system in the sequence, and exhibits the attractive numerical properties of both the block GMRES method with initial deflation [1] and the shifted block GMRES method for solving multiple linear systems [2].

In our talk we present the main lines of development of the FDBGMRES-Sh method, describing its theoretical properties, and we report on numerical experiments on a set of representative matrix problems to show its numerical behavior against other state-of-the art solvers, with and without preconditioner.

[1] H. Calandra, S. Gratton, J. Langou, X. Pinel, and X. Vasseur. Flexible variants of block restarted GMRES methods with application to geophysics, *SIAM J. Sci. Comput.*, 34(2):A714-A736, 2012.

[2] G. Wu, Y.-C. Wang, and X.-Q. Jin. A preconditioned and shifted GMRES algorithm for the pagerank problem with multiple damping factors, *SIAM J. Sci. Comput.*, 34(5):A2558-A2575, 2012.

Authors: Bruno Carpentieri (School of Science and Technology, Nottingham Trent University), Ting-Zhu Huang (School of Mathematical Sciences/Institute of Computational Science, University of Electronic Science and Technology of China), Yan-Fei Jing (School of Mathematical Sciences/Institute of Computational Science, University of Electronic Science and Technology of China), Dong-Lin Sun (Johann Bernoulli Institute for Mathematics and Computer Science, University of Groningen)

M. Sutti: Multiple shooting for the Stiefel geodesic distance Several applications in optimization, image and signal processing deal with data that belong to the Stiefel manifold St(n, p), that is, the set of $n \times p$ matrices with orthonormal columns. Some applications, for example, the computation of the Karcher mean, require evaluating the geodesic distance between two arbitrary points on St(n, p). This can be done by explicitly constructing the geodesic connecting these two points. We propose a new computational framework based on the multiple shooting method for computing the geodesic distance between two points on St(n, p). After generating a starting guess by the Leapfrog algorithm from L. Noakes, we observe that our Newton method converges quadratically to the global solution, also in those cases where existing algorithms fail or converge very slowly. Numerical experiments demonstrate the computational efficiency of the proposed framework.

Authors: Marco Sutti (University of Geneva), Bart Vandereycken (University of Geneva)