The f(A)bulous workshop on matrix functions and exponential integrators

Monday, 25 September 2023 - Wednesday, 27 September 2023

Max Planck Institute for Dynamics of Complex Technical Systems

Book of Abstracts

Book of abstracts

Contents

Invited talks / 20

Limited memory polynomial Krylov subspace methods for matrix functions

Author: Andreas Frommer[1]

[1] *University of Wuppertal*

Corresponding Author: frommer@math.uni-wuppertal.de

When efficient linear solvers for shifted systems are unavailable, polynomial Krylov subspace methods are often the only methods of choice to compute $f(A)b$, the action of a matrix function on a vector. For less well conditioned problems the number of required Arnoldi steps may then become so large that storing the Arnoldi vectors exceeds the available memory and that orthogonalization costs become unbearable.

This talk will deal with different ways to overcome these shortcomings. We will first review restart procedures, in particular for matrix Stieltjes functions and its recent generalization [A. Frommer, K. Kahl, M. Schweitzer, and M. Tsolakis, *Krylov subspace restarting for matrix Laplace transforms*, SIAM J. Matrix Anal. Appl. 44, 693-717 (2023)] to matrix Laplace transforms. We then continue with the sketching approach developed in [S. Güttel and M. Schweitzer, *Randomized sketching for Krylov approximations of large-scale matrix functions*, SIAM J. Matrix Anal. Appl. 44, 1073-1095 (2023)]. Finally, we will show that polynomial preconditioning can also be used to keep recursions short and memory requirements low in the case of the matrix square root and inverse square root. This is joint work with Gustavo Ramirez, Marcel Schweitzer and Manuel Tsolakis.

Contributed talks / 5

Krylov Subspace Recycling For Matrix Functions

Author: Liam Burke[1]

Co-authors: Andreas Frommer [2]; Gustavo Ramirez-Hidalgo [3]; Kirk Soodhalter [1]

[1] *Trinity College Dublin*

[2] *University of Wuppertal*

[3] *Wuppertal University*

Corresponding Author: burkel8@tcd.ie

We discuss a new augmented Krylov subspace method which allows for the efficient evaluation of a sequence of matrix function applications on a set of vectors using Krylov subspace recycling. If selected appropriately, the recycling subspace can be used to accelerate the convergence of each problem in the sequence, leading to an overall reduction in the computational overhead required to evaluate the full sequence of function applications, in comparison to standard Krylov subspace methods. We present results of numerical experiments demonstrating the effectiveness of the method using examples from practical applications such as Quantum Chromodynamics.

Contributed talks / 9

Adaptive rational Krylov methods for exponential Runge–Kutta integrators

Author: Kai Bergermann[1]

Co-author: Martin Stoll [1]

[1] *TU Chemnitz*

Corresponding Authors: kai.bergermann@math.tu-chemnitz.de, martin.stoll@math.tu-chemnitz.de

We consider the solution of large stiff systems of ordinary differential equations with explicit exponential Runge–Kutta integrators. These problems arise from semi-discretized semi-linear parabolic partial differential equations on continuous domains or on inherently discrete graph domains. A series of results reduces the requirement of computing linear combinations of φ-functions in exponential integrators to the approximation of the action of a smaller number of matrix exponentials on certain vectors. State-of-the-art computational methods use polynomial Krylov subspaces of adaptive size for this task. They have the drawback that the required Krylov subspace iteration numbers to obtain a desired tolerance increase drastically with the spectral radius of the discrete linear differential operator, e.g., the problem size. We present an approach that leverages rational Krylov subspace methods promising superior approximation qualities. We prove a novel a-posteriori error estimate of rational Krylov approximations to the action of the matrix exponential on vectors for single time points, which allows for an adaptive approach similar to existing polynomial Krylov techniques. We discuss pole selection and the efficient solution of the arising sequences of shifted linear systems by direct and preconditioned iterative solvers. Numerical experiments show that our method outperforms the state of the art for sufficiently large spectral radii of the discrete linear differential operators. The key to this are approximately constant rational Krylov iteration numbers, which enable a near-linear scaling of the runtime with respect to the problem size.

Contributed talks / 11

Error bounds for the approximation of matrix functions with rational Krylov methods

Author: Igor Simunec[1]

[1] *Scuola Normale Superiore di Pisa*

Corresponding Author: igor.simunec@sns.it

In this talk we present some error bounds for the approximation of matrix-vector products $f(A)b$ and quadratic forms $b^T f(A)b$ with a matrix function $f(A)$, for a Hermitian matrix A and a vector b by means a rational Krylov subspace method. The error bounds are obtained by exploiting properties of rational Arnoldi decompositions and the Cauchy integral formula to link the matrix function error to the residuals of shifted linear systems. This theory leads to both a priori and a posteriori bounds for the error, and it generalizes the bounds derived in [T. Chen, A. Greenbaum, C. Musco, and C. Musco, SIMAX, 2022] for the Lanczos method to the rational Krylov case. The accuracy of the bounds is demonstrated with several numerical experiments.

Invited talks / 17

f(A)bulous networks: an overview of matrix functions in network science

Author: Francesca Arrigo[1]

[1] *University of Strathclyde*

Corresponding Author: francesca.arrigo@strath.ac.uk

Though seemingly they belong to two different worlds, matrix functions and network science have some degree of overlap thanks to a very simple fact; powers of the adjacency matrix count traversals

in the underlying network. This concept in turn allows for the definition of centrality measures in terms of entries (or sums thereof) of functions of the adjacency matrix.

In this talk, after reviewing basic definitions, we will give an overview of popular walk-based centrality measures in networks, emphasizing the role of matrix functions and of expressions of the form $f(A)b$ and $c^T f(A)b$. We will further discuss nonbacktracking walk-based centralities and describe challenges and open problems.

Focus groups / 27

Introduction to focus group topics

Author: Kathryn Lund[1]

Co-authors: Marcel Schweitzer [2]; Stéphane Gaudreault [3]

[1] *Max Planck Institute for Dynamics of Complex Technical Systems*
[2] *University of Wuppertal*
[3] *Environment and Climate Change Canada*

Corresponding Authors: stephane.gaudreault@ec.gc.ca, marcel@uni-wuppertal.de, lund@mpi-magdeburg.mpg.de

We will provide a brief overview of the following topics for the focus groups:
1) Transfering knowledge among f(A)b, exponential integration, and related problems
2) High-performance and energy-aware computing for matrix functions and exponential integration
3) Building a f(A)bulous benchmark collection and comparison framework

Invited talks / 21

The joys and challenges of exponential integration

Author: Mayya Tokman[1]

[1] *University of California, Merced*

Corresponding Author: mtokman@ucmerced.edu

Exponential integration has taken a more prominent role in scientific computing over the past decades. Exponential schemes offer computational savings for many problems involving large stiff systems of differential equations. Careful design of a practical exponential scheme is crucial, however, to ensure that the resulting method is efficient for a particular equation. In particular, to construct an efficient exponential scheme it is important to take into account what algorithms are used to approximate products of the exponential-like matrix functions of the problem's Jacobian or its approximations and vectors. We will discuss several ways to build classes of exponential integrators that take advantage of the structure of the numerical linear algebra algorithms employed within the time integration method. We will present fully exponential, partitioned and hybrid implicit-exponential schemes and illustrate performance of these methods using several application.

Invited talks / 7

The $\varphi(A)$bulous matri$\otimes$ functions in exponential integrators for problems with Kronecker structure

Author: Marco Caliari[1]

[1] *University of Verona*

Corresponding Author: marco.caliari@univr.it

In this talk, we consider two efficient ways to approximate actions of φ-functions for matrices A with a d-dimensional Kronecker sum structure, that is $A = A_d \oplus \cdots \oplus A_1$. The first one is based on the approximation of the integral representation of the φ-functions by Gaussian quadrature formulas combined with a scaling and squaring technique. The resulting algorithm evaluates the required exponential actions using Tucker operators, which are realized in a μ-mode fashion by exploiting the high performance level 3 BLAS. The whole procedure involves exponentials of the *small sized* matrices A_μ, does not require forming the *large sized* matrix A, and allows for the computation of linear combinations of actions of φ-functions, too. The second one is based on a direction splitting and it is suitable for approximating actions of φ-functions for well-established exponential integrators. The desired actions are realized in an efficient way again in a μ-mode fashion by using few Tucker operators. The two approaches have been successfully tested on Riccati differential equations and on 2D and 3D semidiscretized partial differential equations, with various exponential integrators, showing consistent speedups with respect to state-of-the-art techniques to approximate (linear combinations of) actions of φ-functions.

Contributed talks / 8

Domain-based explicit-exponential methods for flow problems

Author: Philipp Birken[1]

Co-authors: Sigrun Ortleb [2]; Veronika Straub [2]; Andreas Meister [2]

[1] *Lund University*
[2] *University of Kassel*

Corresponding Author: philipp.birken@na.lu.se

In the computation of wall bounded flows, resolving the boundary layer requires a very fine resolution. The CFL condition then makes the use of explicit time integration schemes infeasible. However, these may be used in parts of the domain where the mesh is coarse, and using an implicit method only on the remainder. This gives rise to domain based IMEX methods. In this talk, we consider the combination of explicit Runge-Kutta (ERK) methods and exponential integrators, called EXPEX.

A first result is that such a combination of an ERK method with a standard EPIRK method can not have more than order one. However, methods of higher order can be obtained by making use of the (s)EPIRK methods of Rainwater and Tokman. Care has to be taken to define the splitting in a way such that conservation of the flow quantities is not lost.

These methods are then applied to Discontinuous Galerkin discretizations of the compressible Euler and Navier-Stokes equations. Numerical Results show the competitiveness with existing IMEX methods.

Contributed talks / 6

Efficient numerical solution of spatially variable coefficients ADR equations using exponential integrators

Author: Fabio Cassini[1]

Co-authors: Alexander Ostermann [2]; Lukas Einkemmer [2]; Marco Caliari [1]

[1] *University of Verona*

[2] *University of Innsbruck*

Corresponding Author: fabio.cassini@univr.it

Due to the importance of simulation in various fields of science and engineering, devising efficient numerical methods for solving high-dimensional evolutionary Partial Differential Equations is of considerable interest.

In this talk, we present an efficient technique to employ exponential integrators for solving evolutionary Advection-Diffusion-Reaction equations with *spatially variable* coefficients, i.e., in form

$$\left[\begin{array}{l} \partial_t u(t,x) = \nabla \cdot (\lambda(x)\nabla u(t,x)) + r(t,x,u(t,x)), \quad t \in [0,T], \quad x \in \Omega \subset \mathbb{R}^d \\ u(0,x) = u_0(x) \end{array}\right]$$

coupled with appropriate boundary conditions. The approach is based on the extraction from the original PDE of a constant coefficient diffusion part, which is determined by a linear stability analysis of the chosen temporal scheme. After semidiscretization in space, the arising stiff system of Ordinary Differential Equations can then be numerically integrated efficiently by employing, for instance, Fast Fourier Transform or Kronecker/tensor μ-mode based techniques.

In this context, we also consider two new exponential integrators of Lawson type (of first and second order), which appear to have better unconditional stability bounds compared to other well-established exponential integrators.

The validity and effectiveness of the approach is highlighted by presenting numerical experiments in one space dimension and by showing performance results on Advection-Diffusion-Reaction equations in two and three space dimensions, which exhibit a neat advantage compared to state-of-the-art techniques.

Contributed talks / 14

Projected exponential methods for stiff dynamical low-rank approximation problems

Author: Benjamin Carrel[1]

Co-author: Bart Vandereycken [1]

[1] *University of Geneva*

Corresponding Author: benjamin.carrel@unige.ch

The numerical integration of stiff equations is a challenging problem that needs to be approached by specialized numerical methods. Exponential integrators [1] form a popular class of such methods since they are provably robust to stiffness and have been successfully applied to a variety of problems. The dynamical low-rank approximation [2] is a recent technique for solving high-dimensional differential equations by means of low-rank approximations. However, the domain is lacking numerical methods for stiff equations since existing methods [3, 4, 5] are either not robust-to-stiffness or have unreasonably large hidden constants.

In this talk, we focus on solving large-scale stiff matrix differential equations with a Sylvester-like structure,

$$\dot{X}(t) = AX(t) + X(t)B + G(t, X(t)), \quad X_0 = X(0),$$

that admit good low-rank approximations. We propose two new methods that have good convergence properties, small memory footprint and that are fast to compute. The theoretical analysis shows that the new methods have order one and two, respectively. We also propose a practical implementation based on Krylov techniques.

The approximation error is analyzed, leading to a priori error bounds and, therefore, a mean for

choosing the size of the Krylov space. Numerical experiments are performed on several examples, confirming the theory and showing good speedup in comparison to existing techniques.

[1] Hochbruck, M., & Ostermann, A. (2010). Exponential integrators. Acta Numerica, 19, 209-286.
[2] Koch, O., & Lubich, C. (2007). Dynamical low-rank approximation. SIAM Journal on Matrix Analysis and Applications, 29(2), 434-454.
[3] Ceruti, G., & Lubich, C. (2022). An unconventional robust integrator for dynamical low-rank approximation. BIT Numerical Mathematics, 62(1), 23-44.
[4] Lubich, C., & Oseledets, I. V. (2014). A projector-splitting integrator for dynamical low-rank approximation. BIT Numerical Mathematics, 54(1), 171-188.
[5] Ostermann, A., Piazzola, C., & Walach, H. (2019). Convergence of a low-rank Lie–Trotter splitting for stiff matrix differential equations. SIAM Journal on Numerical Analysis, 57(4), 1947-1966.

Invited talks / 10

Understanding tensor and tensor decompositions with application to the chemical master equation

Author: Roger Sidje[1]

[1] *University of Alabama*

Corresponding Author: roger.b.sidje@ua.edu

Tensors are multidimensional arrays that can play a key role in the representation of big data. Decompositions of higher-order tensors have applications in biochemistry, signal processing, data mining, neuroscience, and elsewhere. Building on earlier decompositions (such as canonical/parallel factor (CANDECOMP/PARAFAC), Tucker or its variants), recent research efforts have been devoted to the tensor-train (TT) or quantized tensor train (QTT) to optimize the storage in some large high-order tensors that arise naturally in different scientific fields. Computations with the decomposed tensors, however, raise the sort of challenges reminiscent of the vast complications that arose when new algorithms had to be designed to tackle large sparse matrices to emulate the tasks that were assumed for small dense matrices. But those challenges also point to the prospect of promising research to offer new insights into solving classic problems. We will consider the chemical master equation as an illustrative application that involves an extremely large matrix exponential function that would be infeasible to handle without recasting through tensor-trains decompositions.

Invited talks / 19

$f(A)b$ulous ladder networks: From "rigorous" computation of $b^* \exp(tA)b$ to learning $f(A)b$ with unknown A

Author: Vladimir Druskin[1]

[1] *Worcester Polytechnic Institute*

It is well known that the Lanczos tridiagonal matrix can be transformed to an equivalent finite-difference scheme, with the coefficients obtained from the Stieltjes continued fraction. We show the usefulness of such representation on two seemingly unrelated problems.

The first one is "rigorous" computation of the exponential matrix moments $b^* \exp -tAb$. Here we use the finite-difference representation of the Lanczos tridiagonal matrix to compute sharp upper and lower a-posteriori bounds on the solution. These bounds converge strictly monotonically and even remain valid under computer roundoff, when the Lanczos algorithm loses orthogonality, hence the name "rigorous".

The second problem is learning or data-driven computation of $f(A)b$, where A is a s.p.d differential operator with unknown coefficients and f being $\exp(-tA)$, $\cos(\sqrt{A})$, or $(A + sI)^{-1}b$. As the data we use the corresponding SISO transfer function $b^*f(A)b$ or the matrix moments. Such problems are paramount in remote sensing and other "noninvasive problems", e.g., radar imaging and medical applications, where measurements are not available in the interior problem. We first compute the finite-difference representation using the data-driven Lanczos algorithm, and then compute the state solution $f(A)b$ via the corresponding finite-difference problem. Time permitting, we show generalization to multidimensional PDE operators in MIMO and SAR (Synthetic Aperture Radar) setups.

Contributors to different stages of this long-term project were Jorn Zimmerling, Mikhail Zaslavsky, Valeria Simoncini, Rob Remis, Shari Moskow, Alexander Mamonov, Leonid Knizhnerman, David Ingerman, Fernando Guevara Vasquez, Elena Cherkaeva, Liliana Borcea, and Justin Baker.

Invited talks / 15

Computational graphs for matrix functions

Author: Massimiliano Fasi[1]

Co-authors: Elias Jarlebring [2]; Emil Ringh [3]

[1] *Durham University*

[2] *KTH Royal Institute of Technology*

[3] *Ericsson Research*

Corresponding Author: massimiliano.fasi@gmail.com

Numerical methods for evaluating a function f at an $n \times n$ matrix A can be based on a variety of different approaches, but for a large class of algorithms the matrix $f(A)$ is approximated by using only three operations:

1. $Z \leftarrow c_X X + c_Y Y$, linear combination of matrices),

2. $Z \leftarrow X \cdot Y$, (matrix multiplication),

3. $Z \leftarrow X^{-1}Y$, (solution of a linear system with n right-hand sides),

where X, Y, and Z are $n \times n$ matrices and c_X and c_Y are scalars.

Algorithms that combine only these three basic building blocks are particularly attractive, as they correspond to functions that are easy to work with: if an expression for the scalar function g features only linear combinations, multiplications, and inversions, and g is defined on the spectrum of A, then a formula for $g(A)$ can be obtained by replacing all occurrences of z in the formula for $g(z)$ with A.

Rephrasing these methods as directed acyclic graphs (DAGs) is a particularly effective approach to study existing techniques, improve them, and eventually derive new ones. The accuracy of these matrix techniques can be characterized by the accuracy of their scalar counterparts, thus designing algorithms for matrix functions can be regarded as a scalar-valued optimization problem. The derivatives needed during the optimization can be calculated automatically by exploiting the structure of the DAG, in a fashion analogous to backpropagation.

GraphMatFun.jl is a Julia package that offers the means to generate and manipulate computational graphs, optimize their coefficients, and generate Julia, MATLAB, and C code to evaluate them efficiently at a matrix argument. The software also provides tools to estimate the accuracy of a graph-based algorithm and thus obtain numerically reliable methods. For the exponential, for example, using a particular form (degree-optimal) of polynomials produces implementations that in many cases are cheaper, in terms of computational cost, than the Padé-based techniques typically used in mathematical software.

Invited talks / 18

Estimating the von Neumann entropy of large matrices

Author: Michele Benzi[1]

[1] *Scuola Normale Superiore di Pisa*

The problem of approximating the von Neumann entropy of a symmetric positive semidefinite matrix A, defined as $\mathrm{tr}(f(A))$ where $f(x) = -x \log x$, is considered. After discussing some useful properties of this matrix function, approximation methods based on randomized trace estimation and probing techniques used in conjunction with polynomial and rational Krylov methods will be described. Bounds and heuristics used in the implementation of the algorithms will be discussed. The performance of the methods will be assessed using test problems arising in Network Science. This is joint work with Michele Rinelli and Igor Simunec.

Reference: M. Benzi, M. Rinelli and I. Simunec, *Computation of the von Neumann entropy of large matrices via trace estimators and rational Krylov methods*, to appear in Numerische Mathematik.

Contributed talks / 12

Analysis of stochastic probing methods for estimating the trace of functions of sparse symmetric matrices

Author: Michele Rinelli[1]

Co-authors: Andreas Frommer [2]; Marcel Schweitzer [2]

[1] *Scuola Normale Superiore di Pisa*
[2] *University of Wuppertal*

Corresponding Author: michele.rinelli@sns.it

We consider the problem of estimating the trace of a matrix function $f(A)$. In certain situations, in particular if $f(A)$ cannot be well approximated by a low-rank matrix, combining probing methods [2] based on graph colorings with stochastic trace estimation techniques can yield very accurate approximations. So far, such methods have not been thoroughly analyzed, though, but were rather used as efficient heuristics by practitioners 1. We perform a detailed analysis of stochastic probing methods and, in particular, expose conditions under which the expected approximation error in the stochastic probing method scales more favorably with the dimension of the matrix than the error in non-stochastic probing. Extending previous heuristic results 1, we also characterize situations in which using just one stochastic vector is always better than the deterministic probing method. Several numerical experiments illustrate our theory and compare with existing methods.

1 E. Aune, D. P. Simpson, and J. Eidsvik. Parameter estimation in high dimensional Gaussian distributions. *Stat. Comput.*, 24:247–263, 2014.

[2] A. Frommer, C. Schimmel, and M. Schweitzer. Analysis of probing techniques for sparse approximation and trace estimation of decaying matrix functions. *SIAM J. Matrix Anal. Appl.*, 42(3):1290–1318, 2021.

Contributed talks / 4

Randomized low-rank approximation of monotone matrix functions

Author: David Persson[1]

Co-author: Daniel Kressner [1]

[1] *École polytechnique fédérale de Lausanne*

Corresponding Author: david.persson@epfl.ch

This work is concerned with computing low-rank approximations of a matrix function $f(A)$ for a large symmetric positive semi-definite matrix A, a task that arises in, e.g., statistical learning and inverse problems. The application of popular randomized methods, such as the randomized singular value decomposition or the Nyström approximation, to $f(A)$ requires multiplying $f(A)$ with a few random vectors. A significant disadvantage of such an approach, matrix-vector products with $f(A)$ are considerably more expensive than matrix-vector products with A, even when carried out only approximately via, e.g., the Lanczos method. In this work, we present and analyze funNyström, a simple and inexpensive method that constructs a low-rank approximation of $f(A)$ directly from a Nyström approximation of A, completely bypassing the need for matrix-vector products with $f(A)$. It is sensible to use funNyström whenever f is monotone and satisfies $f(0) = 0$. Under the stronger assumption that f is operator monotone, which includes the matrix square root $A^{1/2}$ and the matrix logarithm $\log(I + A)$, we derive probabilistic bounds for the error in the Frobenius, nuclear, and operator norms. These bounds confirm the numerical observation that funNyström tends to return an approximation that compares well with the best low-rank approximation of $f(A)$. Furthermore, compared to existing methods, funNyström requires significantly fewer matrix-vector products with A to obtain a low-rank approximation of $f(A)$, without sacrificing accuracy or reliability. Our method is also of interest when estimating quantities associated with $f(A)$, such as the trace or the diagonal entries of $f(A)$. In particular, we propose and analyze funNyström++, a combination of funNyström with the recently developed Hutch++ method for trace estimation.

Contributed talks / 29

Sensitivity of matrix function based network communicability measures

Author: Marcel Schweitzer[1]

[1] *Bergische Universität Wuppertal*

Corresponding Author: marcel@uni-wuppertal.de

When analyzing complex networks, an important task is the identification of nodes which play a leading role for the overall communicability of the network. In the context of modifying networks (or making them robust against targeted attacks or outages), it is also relevant to know how sensitive the network's communicability is to changes in certain nodes or edges.

Recently, the concept of *total network sensitivity* was introduced in [O. De la Cruz Cabrera, J. Jin, S. Noschese, L. Reichel, *Communication in complex networks*, Appl. Numer. Math., 172, pp. 186-205, 2022], which allows to measure how sensitive the total communicability of a network is to the addition or removal of certain edges. One shortcoming of this concept is that sensitivities are extremely costly to compute when using a straight-forward approach (orders of magnitude more expensive than the corresponding communicability measures).

To overcome this problem, we combine Krylov methods for computing Fréchet derivatives with a maximum element estimator for implicitly given matrices, which results in a computational procedure for estimating network sensitivity at a cost that is essentially *linear in the number of nodes* for many real-world complex networks.

Contributed talks / 13

A message-passing approach to obtain tr(f(A)) with applications to network analysis

Author: Grover Enrique Castro Guzman[1]

[1] *University of Sao Paulo*

Corresponding Author: grover.org.rf@gmail.com

In recent decades, graphs have become widely used for studying technological, biological, and social systems. For this reason, several methods have been proposed to measure the structural and dynamical properties of graphs, aiming to gain insights into the represented system. Among these methods, measures based on matrix functions have been introduced for assessing vertex centrality and communicability. While numerous iterative techniques for computing matrix function traces have been suggested, these methods often necessitate a global perspective or collective coordination, such as a global sum. To address these limitations, we present a message-passing approach designed to operate efficiently within parallel, heterogeneous, distributed, and asynchronous computing systems. Through simulations and analyses on both simulated and real-world networks, we demonstrate the performance of our method and compare it with the current state-of-the-art approach.