GESELLSCHAFT für
ANGEWANDTE MATHEMATIK und MECHANIK e.V.
INTERNATIONAL ASSOCIATION of APPLIED MATHEMATICS and MECHANICS

86th Annual Meeting

of the International Association
of Applied Mathematics and Mechanics

March 23-27, 2015
Lecce, Italy

Book of Abstracts - Extract

2015

UNIVERSITÀ DEL SALENTO

jahrestagung.gamm-ev.de

Scientific Program - Timetable

Sunday 22	Time	Monday 23	Tuesday 24	Wednesday 25	Thursday 26	Friday 27
	9:15		Contributed sessions (15 in parallel)	Plenary Lecture Moritz Diehl	Contributed sessions (15 in parallel)	Contributed sessions (14 in parallel)
	9:30					
	9:45					
	10:15	Registration		von Mises prize lecture		
	10:30					
	10:45					
	11:15		Coffee Break	Coffee Break	Coffee Break	Coffee Break
	11:30					
	11:45					
	12:15		Plenary Lecture Thomas Böhlke	General Assembly	Plenary Lecture Ferdinando Auricchio	Contributed sessions (11 in parallel)
	12:30					
	12:45					
	13:15	Opening	Lunch	Lunch	Lunch	
	13:30					
	13:45	Univ. Chorus Performance				Closing
	14:15	Prandtl Lecture Keith Moffatt	Plenary Lecture Enrique Zuazua	Contributed sessions (15 in parallel)	Plenary Lecture Daniel Kressner	
	14:30					
	14:45					
	15:15	Plenary Lecture Giovanni Galdi	Plenary Lecture Nikolaus Adams		Plenary Lecture Stanislaw Stupkiewicz	
	15:30					
	15:45					
Registration pre-opening	16:15	Coffee Break	Coffee Break Poster session	Coffee Break	Coffee Break Poster session	
	16:30	Minisymposia & Young Reseachers' Minisymposia (10 in parallel)	Contributed sessions (14 in parallel)	Contributed sessions (15 in parallel)	Contributed sessions (15 in parallel)	
	16:45					
	17:15					
	17:30					
	17:45					
	18:15		Public lecture Francesco D'Andria			
	18:30					
	18:45					
	19:15	Opening reception at Castle of Charles V				
	19:30					
	19:45					
	20:15			Conference dinner at Hotel Tiziano		
	20:30					
	20:45					
	21:15					
	21:30					
	21:45					

GAMM 2015

Università del Salento

Table of contents

S17: Applied and numerical linear algebra

The aim of this section is to bring together experts in the field of applied and numerical linear algebra, discussing recent theoretical and algorithmic developments.

Designing rational filter functions for solving eigenvalue problems by contour integration

Marc Van Barel
Department of Computer Science, KU Leuven, Belgium

In this talk, the following *eigenvalue problem* is considered. Given an integer $m \geq 1$, a domain $\Omega \subset \mathbb{C}$ and a matrix-valued function $T : \Omega \to \mathbb{C}^{m \times m}$ analytic in Ω, we want to compute the values $\lambda \in \Omega$ (eigenvalues) and $v \in \mathbb{C}^m$, $v \neq 0$ (eigenvectors) such that

$$T(\lambda)v = 0.$$

Note that this formulation reduces to the linear eigenvalue problem in case $T(z) = A - zB$, and to the polynomial eigenvalue problem when $T(z)$ is a polynomial matrix. If the problem size m is equal to 1, then the problem reduces to that of computing all the zeros λ of the analytic scalar function T inside the domain Ω.

The number of eigenvalues could be large, e.g., when m is large, or in case of a polynomial eigenvalue problem when the degree of the polynomial matrix is large. In several applications, one is not interested in *all* eigenvalues but only in those lying in a certain region(s) of the complex plane. Therefore, we can reduce the original problem of finding *all* eigenvalues into one where we are only interested in those eigenvalues (and corresponding eigenvectors) lying within (or in the neighborhood) of a given closed contour $\Gamma \subset \Omega$.

The approach discussed in this talk is based on (numerical approximations of) contour integrals of the resolvent operator $T(z)^{-1}$ applied to a rectangular matrix $\hat{V}$:

$$\frac{1}{2\pi i} \int_\Gamma f(z) T(z)^{-1} \hat{V} dz \quad \in \mathbb{C}^{m \times q}$$

where $f : \Omega \to \mathbb{C}$ is analytic in Ω and $\hat{V} \in \mathbb{C}^{m \times q}$ is a matrix chosen randomly or in another specified way, with $q \leq m$.

The contour integral is approximated by a quadrature rule with nodes t_j and corresponding weights u_j, i.e.,

$$\int_\Gamma f(z) dz \approx \sum_{j=1}^{N} u_j f(t_j).$$

As was explained in [1], from Keldysh' theorem, we know that the resolvent function $T(z)^{-1}$ can be written (for simple eigenvalues λ) as

$$T(z)^{-1} = \sum_k v_k w_k^H \frac{1}{z - \lambda_k} + R(z)$$

with $R(z)$ an analytic function where $T(z)$ is analytic. Note that if T^{-1} is a matrix-valued strictly proper rational function, the analytic function R is equal to zero. This is the case, for example, if $T(z) = A - zB$ with B nonsingular or if $T(z)$ is a matrix polynomial in z with nonsingular highest degree coefficient.

Hence, applying the quadrature rule on the moments of the resolvent function $z^l T(z)^{-1}$ gives us

$$\begin{aligned} \int_\Gamma z^l T(z)^{-1} dz &\approx \sum_{j=1}^{N} u_j t_j^l T(t_j)^{-1} \\ &= \sum_k v_k w_k^H \sum_{j=1}^{N} \frac{u_j t_j^l}{t_j - \lambda_k} + \sum_{j=1}^{N} u_j t_j^l R(t_j). \end{aligned}$$

The filter functions $b_l(z)$ are defined as the rational functions of degree δ corresponding to the quadrature rule as follows

$$b_l(z) = \sum_{j=1}^{N} \frac{u_j t_j^l}{t_j - z}.$$

In [2], we argued that it is important to have robust and cheap ways to design good filter functions. This design consists in finding the weights u_j and the nodes t_j such that the following conditions are satisfied:

1. $b_l(z) = b_0(z)z^l$;
2. $|b_l(z)|$ is large inside Γ and small outside Γ;
3. $|\sum_{j=1}^{N} u_j t_j^l R(t_j)|$ is small.

We call the function $b_0(z)$ "the filter function" and denote it as $b(z)$. It is easy to see that as long as $l + \nu < N$ with ν the degree of the numerator of the filter function $b(z)$, the first condition is satisfied. To satisfy the second and third condition, we will solve an optimization problem looking for the variables u_j and t_j such that $|b_l(z)|$ is large inside Γ and small outside Γ. The validity of the approach will be illustrated by some numerical experiments.

References

[1] W.-J. Beyn. An integral method for solving nonlinear eigenvalue problems. *Linear Algebra and its Applications*, 436:3839–3863, 2012.

[2] M. Van Barel and P. Kravanja. Nonlinear eigenvalue problems and contour integrals. Technical Report TW656, Department of Computer Science, KU Leuven, October 2014.

Parallel Bidiagonal SVD via the Method of Multiple Relatively Robust Representations

Jan Winkelmann, Paolo Bientinesi
RWTH Aachen

One of the central problems in numerical linear algebra is the computation of the singular value decomposition (SVD) of a real bidiagonal matrix. One way to calculate the bidiagonal SVD is by reducing the problem to a real symmetric tridiagonal eigenproblem. The state of the art eigensolver-based SVD method is the Divide and Conquer SVD algorithm, with a run-time complexity of $\mathcal{O}(n^3)$ in the worst case, and $\mathcal{O}(n^2)$ in the best case. A more modern eigensolver is the method of Multiple Relatively Robust Representations (MR^3), which calculates k eigenpairs with $\mathcal{O}(nk)$ operations in the worst case. A bidiagonal SVD solver based on MR^3 promises to be an improvement over the current state of the art by providing lower asymptotic runtime and the ability to calculate a subset of singular triplets at reduced cost. Currently, no implementation is readily available that uses MR^3 as an SVD solver. We provide such a solver, adapted from of an existing MR^3 implementation.

Using MR^3 as a bidiagonal SVD solver in a numerically stable way is not a trivial task. Willems and Lang [1] have presented a modification of MR^3 that can be used to obtain bidiagonal SVDs in a numerically stable manner. The approach uses MR^3 on Golub-Kahan matrices. Unfortunately, this approach does not work if MR^3 is used in a black-box manner. The reason lies in the way the input matrix is represented internally by the algorithm. MR^3 uses LDL^* factorizations to ensure accuracy of the results, and LDL^* factorizations of definite matrices provably guarantee high accuracy. As high accuracy is always desirable, black-box MR^3 shifts the eigenvalue spectrum of the input matrix to be definite. However, this shift causes problems with the quality of the extracted SVD later on. Willems and Lang solved this problem by avoiding the initial shifting of the spectrum, thus guaranteeing the quality of the obtained SVD.

Indeed, shifting the eigenvalue spectrum of the input matrix is something MR^3 does multiple times. Roughly speaking, MR^3 is similar to bisection and inverse iteration in a manner that guarantees the orthogonality of the eigenvectors. This cannot be accomplished if some of the eigenvalues are clustered, that is, if they have small relative distances. In MR^3, clusters are broken by shifting the spectrum close to the clustered eigenvalues, as this enlarges their relative distances. After the shift, the eigenpairs for the newly "unclustered" eigenvalues can be calculated with high accuracy.

Our contribution consists of the following elements: We extended an existing MR^3 eigensolver implementation (incidentally also by Willems) to be a bidiagonal SVD solver. We equipped our implementation with shared-memory parallelism. Previous work has shown that MR^3 can achieve a high degree of parallelism [2]. We adopted the same task parallel approach. This entails the decomposition into three distinct tasks, (1) shifting the spectrum to break clusters, (2) refinement and classification of the eigenvalues after a shift, and (3) calculating the eigenpair for eigenvalues with large relative distances. In this task we discuss the implementation of our shared-memory parallel SVD solver and present results regarding runtime and the quality of the resulting decomposition.

References

[1] P.R. Willems, B. Lang. The MR^3-GK Algorithm for the Bidiagonal SVD. Electronic Transactions on Numerical Analysis. (2012) Volume 39, 1-21.

[2] M. Petschow, E. Peise, P. Bientinesi. High-Performance Solvers for Dense Hermitian Eigenproblems. SIAM Journal on Scientific Computing. (2013) Volume 35(1), pp. C1-C22.

A Hessenberg reduction algorithm for diagonal plus low rank matrices

Dario A. Bini, Leonardo Robol
University of Pisa
Scuola Normale Superiore, Pisa

We present an algorithm for Hessenberg reduction of matrices of the form $A = D + UV^*$ where D is real diagonal and U, V are in $\mathbb{C}^{n\times k}$. Examples of such problems may be encountered as a first step in the solution of structured eigenvalue problems. Some classes of block companion matrices (see for example [1]) do have this particular structure and thus may benefit from an efficient method for the computation of the Hessenberg form.

The algorithm uses a rank k generalization of the quasiseparable techniques developed in [4], obtaining an asymptotic cost for the reduction of $O(n^2k)$ floating point operations. Other algorithms have been presented in the literature, for example in [3] and [2]. Some of them do have an asymptotic cost in the rank that is cubic, thus making them not effective when the rank k is bigger than $n^{\frac{1}{3}}$. This can be an issue when dealing with linearizations of matrix polynomials that often have a quasiseparable rank not negligible with respect to the size of the problem. For this reason we focused our attention at obtaining a linear complexity in k.

We introduce some theoretical results on rank conservation that allow to sharply bound the quasiseparable ranks of the matrices obtained at each step of the reduction and explicitly represent every matrix without redundancy in the rank. More precisely, we provide an explicit representation of the trailing principal submatrix of $Q_jAQ_j^*$, the matrix obtained at the j-th step of the reduction. This representation can be used to compute the unitary matrix Q_{j+1}. Moreover, we discuss different representations for these intermediate matrices and analyze the numerical stability of the approach. Some strategies are proposed in order to avoid cancellation in the process with the purpose of increase its stability.

References

[1] Bini, Dario A., and Leonardo Robol. On a Class of Matrix Pencils Equivalent to a Given Matrix Polynomial. arXiv preprint arXiv:1406.1025 (2014).

[2] Delvaux, Steven, and Marc Van Barel. A Hessenberg reduction algorithm for rank structured matrices. SIAM Journal on Matrix Analysis and Applications 29.3 (2007): 895-926.

[3] Eidelman, Yuli, Israel Gohberg, and Luca Gemignani. On the fast reduction of a quasiseparable matrix to Hessenberg and tridiagonal forms. Linear algebra and its applications 420.1 (2007): 86-101.

[4] Vandebril, Raf, Marc Van Barel, and Nicola Mastronardi. Matrix computations and semiseparable matrices: linear systems. Vol. 1. JHU Press, 2010.

An extended Hessenberg form for Hamiltonian matrices

Micol Ferranti[1], Bruno Iannazzo[2], Thomas Mach[1], Raf Vandebril[1]
[1]Dept. of Computer Science, KU Leuven, Belgium
Dep[2]t. of Mathematics and Computer Science, Università degli studi di Perugia, Italy

The QR method is one of the most used algorithms to compute eigenvalues of medium sized matrices. It is a two step method: the matrix is first transformed via unitary similarity transformations to a convenient condensed form, whose eigenvalues are then computed via a suitable iterative method.

The typical generic condensed form is the well known Hessenberg form, see, e.g., [4]. However, when the original matrix exhibits particular properties, the preservation of the structure through the whole procedure is advisable. For this reason, some special formulations of the QR method have been developed for specific classes of matrices (i.e., symmetric, unitary, etc.).

On the other hand, it has been shown in [3] that the classic Hessenberg form is not the only possible choice as an intermediate condensed form. A much wider family of matrices can be used within a QR-like algorithm, in order to achieve the same results. This family extends classes such as Hessenberg and Hessenberg-like, and will be named extended Hessenberg form.

A Hamiltonian matrix is represented by

$$H = \begin{bmatrix} A & G \\ F & -A^H \end{bmatrix} \in \mathbb{C}^{2n\times 2n},$$

where $F = F^H$ and $G = G^H$. The eigenvalues of a Hamiltonian matrix are symmetric with respect to the imaginary axis. To preserve this symmetry within a QR algorithm, it would be desirable to work exclusively with Hamiltonian matrices during the iterative process.

The generic Hessenberg form does not retain the Hamiltonian structure. Thus many attempts have been made to design a suitable Hamiltonian condensed form, see, for example, [1, 2]. Until now, a Hamiltonian QR algorithm has been found for matrices H whose bottom-left block F has rank 1 [1]. This algorithm makes use of the so called Hamiltonian Hessenberg form.

In this talk we will show how the arguments presented in [3] can be adapted to the Hamiltonian context, in order to derive a new Hamiltonian condensed form, which extends the classic Hamiltonian Hessenberg form.

References

[1] R. Byers. A Hamiltonian QR-algorithm. *SIAM Journal on Scientific and Statistical Computation*, 7(1):212-229, 1986.

[2] C. C. Paige and C. F. Van Loan. A Schur decomposition for Hamiltonian matrices. *Linear Algebra and its applications*, 41:11-32, 1981.

[3] R. Vandebril. Chasing bulges or rotations? A metamorphosis of the QR-algorithm. *SIAM Journal on Matrix Analysis and Applications*, 32:217-247, 2011.

[4] D. S. Watkins. Francis's algorithm. *American Mathematical Monthly*, 118:387-403,2011.

Rank–revealing decomposition via block anti–triangular factorization

Nicola Mastronardi, Paul Van Dooren
Istituto per le Applicazioni del Calcolo "M. Picone", sede di Bari
Consiglio Nazionale delle Ricerche
Via G. Amendola, 122/D, I-70126 Bari, Italy
Catholic University of Louvain Department of Mathematical Engineering
Avenue Georges Lemaitre 4
B-1348 Louvain-la-Neuve, Belgium

An algorithm for computing the block anti–triangular (BAT) form of a symmetric matrix has been recently introduced [1]. In particular, given a symmetric matrix $A \in \mathbb{R}^{n\times n}$, an orthogonal matrix $Q \in \mathbb{R}^{n\times n}$ is sought such that

$$A = QMQ^T = \begin{bmatrix} \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & Y^T \\ \mathbf{0} & \mathbf{0} & X & Z^T \\ \mathbf{0} & Y & Z & W \end{bmatrix} \begin{matrix} \}n_0 \\ \}n_1 \\ \}n_2 \\ \}n_1 \end{matrix}$$

with n_0, n_1 and n_2 depending on the inertia of A. Moreover, $Q \in \mathbb{R}^{n\times n}$ is an orthogonal matrix , $Z \in \mathbb{R}^{n_1\times n_2}$, $W \in \mathbb{R}^{n_1\times n_1}$ are symmetric, $Y \in \mathbb{R}^{n_1\times n_1}$ is a nonsingular lower anti–triangular matrix and $X \in \mathbb{R}^{n_2\times n_2}$ is symmetric definite.

It has been shown that the latter BAT factorization of a symmetric matrix can be efficiently updated/downdated by a symmetric rank–one modification. Moreover, the updating of a matrix obtained appending to a BAT matrix one more row and, symmetrically, a column, represents the kernel step of the algorithm described in [1] and can be done in $O(n^2)$ floating point operations.

In this talk we will describe an algorithm that factorizes a symmetric matrix as the product QMQ^T, with M in a rank–revealing BAT form, i.e., denoted by $\sigma_i,\ i = 1, \ldots, n$, the singular values of M in a decreasing order, and supposed $\sigma_k \gg \sigma_{k+1}$,

$$M = \begin{bmatrix} M_{11} & M_{12} \\ M_{12}^T & M_{22} \end{bmatrix}, \quad M_{11} \in \mathbb{R}^{(n-k)\times(n-k)},\ M_{11} \in \mathbb{R}^{(k\times k}},\ M_{12} \in \mathbb{R}^{k\times(n-k)},$$

with $0 \le k \le n$, we have

$$\mathrm{cond}(M_{11}) \simeq \sigma_1/\sigma_k \ \text{ and } \ \|M_{11}\|_F^2 + \|M_{12}\|_F^2 \simeq \sigma_{k+1} + \cdots + \sigma_n.$$

Numerical results will be shown together with comparisons with existing algorithms for symmetric rank–revealing factrorizations available in the literature [2].

References

[1] N. Mastronardi, P. Van Dooren. The antitriangular factorization of symmetric matrices, SIAM J. Matrix Anal. Appl., 34 (2013), 173–196.

[2] P.C. Hansen, P.Y.Yalamov, Computing Symmetric Rank-Revealing Decompositions via Triangular Factorization, SIAM Journal on Matrix Analysis and Applications, 23 (2001), 443–458.

On complex J-symmetric eigenproblems

Peter Benner[1], Heike Faßbender[2], Chao Yang[3]
[1]Max Planck Institute for Dynamics of Complex Technical Systems, Magdeburg, Germany
[2]TU Braunschweig, Institut Computational Mathematics,
AG Numerik, Braunschweig, Germany
[3]Lawrence Berkeley National Laboratory, Computational Research Division, Berkely, USA

The eigenproblem $H_C x = \lambda x$ for matrices

$$H_C = \begin{bmatrix} A & C \\ D & -A^T \end{bmatrix} \in \mathbb{C}^{2n\times 2n}, \quad A, C = C^T, D = D^T \in \mathbb{C}^{n\times n}.$$

will be considered. Please note, that here X^T denotes transposition, $Y = X^T, y_{ij} = x_{ji}$, no matter whether X is real or complex, while X^H denotes conjugate transposition, $Y = X^H, y_{ij} = \overline{x_{ji}}$.

For

$$J_n = \begin{bmatrix} 0 & I_n \\ -I_n & 0 \end{bmatrix} \in \mathbb{R}^{2n\times 2n}, \quad I_n \in \mathbb{R}^{n\times n}$$

we have

$$(H_C J)^T = H_C J.$$

Matrices H_C are called complex-J-symmetric. The eigenvalues of H_C display a symmetry: they appear in pairs $(\lambda, -\lambda)$. If x is the right eigenvector corresponding to λ, $H_C x = \lambda x$, than Jx is the left eigenvector corresponding to the eigenvalue $-\lambda$ of H_C, $(Jx)^T H_C = -\lambda (Jx)$.

Any complex J-symmetric matrix X is said to be in structured Schur form if

$$X = \begin{bmatrix} R & B \\ 0 & -R^T \end{bmatrix}, \qquad R, B = B^T \in \mathbb{C}^{n\times n},$$

where the nonzero eigenvalues of R either have positive real part or zero real part and positive imaginary part. We will prove that for any complex J-symmetric matrix H_C there exists a complex symplectic and unitary matrix $W \in \mathbb{C}^{2n\times 2n}$

$$W^T J W = J \qquad W^H W = I,$$

such that $W^H H_C W$ is in structured Schur form.

The most popular way to compute the standard Schur form of a general matrix is the QR algorithm. It is tempting to derive a structured QR algorithm for transforming H_C iteratively into structured Schur form. We will discuss why this is not possible and suggest other methods to compute eigenvalues and eigenvectors of H_C.

Nonlinear eigenvalue problem expressed in Hermite basis

Heike Fassbender, Nikta Shayanfar
AG Numerik, Institut Computational Mathematics, Technische Universität Braunschweig

Polynomial eigenvalue problems have been used extensively in different areas of mechanics, such as machine foundations, electronic model of metal strip, simplified nuclear power plant problem, acoustic wave problem, Whipple bicycle model, etc. A comprehensive collection of the eigenvalue problems from models of real life as well as structured ones is studied in [3]. The standard form of a polynomial eigenvalue problem is to find scalars λ and nonzero vectors x such that they satisfy

$$P(\lambda)x = 0,$$

where $P : \mathbb{C} \to \mathbb{C}^{n\times n}$ is a given matrix-valued polynomial of degree s as

$$P(\lambda) = \sum_{i=0}^{s} P_i \lambda^i, \qquad P_i \in \mathbb{C}^{n\times n}, \quad P_s \neq 0.$$

The variable $\lambda \in \mathbb{C}$ and the nonzero vector $x \in \mathbb{C}^{n\times 1}$ are the sought *eigenvalue* and *eigenvector*, respectively.

The growing attention to the polynomial eigenvalue problems has created the necessity of studying linearization, which is the standard computational approach to solving eigenvalue problems. One of the first papers discussing the main idea is [4]. It is based on finding a linear matrix polynomial $L(\lambda)$ such that the spectral structure of this particular linearization reproduces that of a given matrix polynomial $P(\lambda)$. More precisely, the above equation is converted to $L(\lambda)y = 0$, where $L(\lambda)$ is a larger size matrix polynomial with the same spectral properties, specially same eigenvalues. Then it can be solved by standard techniques developed for linear eigenvalue problems.

On checking the history of linearization, we find that linearization has been done in different bases, see [1] among others. Actually, matrix polynomials expressed in other bases other than monomial, occur in many applications. For instance, Bernstein basis appears in computer-aided geometric design. Legendre basis helps solving the problems in partial differential equations with symmetries in the boundary conditions. Lagrange polynomial interpolation is traditionally viewed as a tool for theoretical analysis, and [2] demonstrates several advantages to computation in the Lagrange basis. Generalization of the Lagrange basis leads to the Hermite interpolation problem. To satisfy functional or aesthetic criteria, designed objects often have to exactly match prescribed data, such as a series of points and derivatives. This idea leads to the presentation of the Hermite basis, for which the linearization has not been comprehensively studied in literature. This contribution presents ideas of solving nonlinear eigenvalue problems in Hermite basis. We introduce a new linearization for an eigenvalue problem in which the matrix polynomial is expressed in Hermite basis, and the property of strong linearization is to be investigated. This linearization shows a reliable behavior and an acceptable accuracy in comparison to the recent results in the literature.

References

[1] A. Amiraslani, R. Corless , P. Lancaster. Linearization of matrix polynomials expressed in polynomial bases. IMA J. Numer. Anal. 29 (2009), 141–157.

[2] J.P. Berrut, L.N. Trefethen, Barycentric Lagrange interpolation. SIAM Rev. 46 (2004), 501–517.

[3] T. Betcke, N.J. Higham, V. Mehrmann, C. Schröder, F. Tisseur. NLEVP: a collection of nonlinear eigenvalue problems. ACM Trans. Math. Software 39 (2013), 1–28.

[4] I. Gohberg, M.A. Kaashoek, P. Lancaster. General theory of regular matrix polynomials and band Toeplitz operators. Integral Equations Operator Theory 11 (1988), 776–882.

Backward error of polynomial eigenvalue problems solved by linearization

Piers W. Lawrence[1], Marc Van Barel[2], Paul Van Dooren[1]
[1]Université Catholique de Louvain, Department of Mathematical Engineering
[2]KU Leuven, Department of Computer Science

It is commonplace in many application domains to utilize polynomial eigenvalue problems to model the behaviour of physical systems. Many techniques exist to compute solutions of these polynomial eigenvalue problems. One of the most frequently used techniques is linearization, in which the polynomial eigenvalue problem is turned into an equivalent linear eigenvalue problem with the same eigenvalues, and with easily recoverable eigenvectors. The eigenvalues and eigenvectors of the linearization are usually computed using a backward stable solver such as the QZ algorithm. Such backward stable algorithms ensure that the computed eigenvalues and eigenvectors of the linearization are exactly those of a nearby linear pencil, where the perturbations are bounded in terms of the machine precision and the norms of the matrices defining the linearization. With respect to the linearization, we may have solved a nearby problem, but we would also like to know if our computed solution is the exact solution of a nearby polynomial eigenvalue problem.

Furthermore, there has recently been a steady increase in the number of distinct linearizations proposed in the literature, depending mainly on the basis in which the polynomial eigenvalue problems are represented. Certainly, the choice of basis can have a dramatic effect on the backward errors, as can the particular choice of linearization. One of the objectives of this work is to develop a framework for analyzing different polynomial bases and linearizations in a uniform way. Thus, we investigate a particular class of linearizations where the polynomial coefficients are separated from the recurrence relations of the polynomial basis employed.

We use one-sided factorization to relate the linearization to the original polynomial in a very particular way. Given a linearization $\mathcal{L}(\lambda)$ of a polynomial matrix $P(\lambda)$, we find a one-sided factorization $\Phi(\lambda)$, such that $\mathcal{L}(\lambda)\Phi(\lambda) = P(\lambda) \otimes e_1$, where e_1 is the first unit vector. Since the QZ algorithm computes the exact solution of a slightly perturbed linearization, we investigate

$$(\mathcal{L}(\lambda) + \Delta\mathcal{L}(\lambda))(\Phi(\lambda) + \Delta\Phi(\lambda)) = \begin{bmatrix} P(\lambda) + \Delta P(\lambda) \\ 0 \end{bmatrix}, \tag{1}$$

to first order. The perturbation $\Delta\Phi(\lambda)$ to the one-sided factorization is chosen in order to maintain the structure in the bottom of (1). For a given specific basis, we utilize the appropriate convolution matrices in order to obtain upper bounds for the norm of the coefficients of the perturbation $\Delta P(\lambda)$.

For some specific polynomial bases and linearizations, we are able to formulate these upper bounds in a simple way. Thus, we obtain the conditions under which the backward error of the solution of the polynomial eigenvalue problems are small.

IgA vs. FEA in the Spectral Approximation: Symbol-Based Analysis

C. Garoni[a,b], T.J.R. Hughes[c], C. Manni[a], A. Reali[d], S. Serra-Capizzano[b], H. Speleers[a]
[a]Department of Mathematics, University of Roma 'Tor Vergata'
[b]Department of Science and High Technology, University of Insubria (Como, Italy)
[c]Institute for Computational Engineering and Sciences, The University of Texas at Austin
[d]Department of Civil Engineering and Architecture, University of Pavia (Italy)

Isogeometric Analysis (IgA) was recently introduced by T.J.R. Hughes and his research team in order to reduce the gap between Computer-Aided Design (CAD) and Finite Element Analysis (FEA). The main idea in IgA is to use the same basis functions provided by CAD systems – usually, B-splines or NURBS – both in the approximation of differential problems and in the description of the geometry of the related physical domains.

After its birth, IgA has gained a lot of attention because of its remarkable spectral approximation properties, which make it superior to FEA. Considering for simplicity the 1D Laplacian eigenvalue problem

$$\begin{cases} -u_\ell'' = \lambda_\ell u_\ell, \quad \text{in } (0,1), \\ u_\ell(0) = u_\ell(1) = 0, \end{cases} \tag{1}$$

whose solutions are the pairs (λ_ℓ, u_ℓ) with $\lambda_\ell = \omega_\ell^2 = (\ell\pi)^2$ and $u_\ell(x) = \sin(\ell\pi x)$, $\ell = 1, 2, \ldots$, what has been observed in the literature [8, 1, 7, 6] is the following.

- The spectrum $\{\omega_{\ell,\mathrm{IgA}}^2,\ \ell = 1, \ldots, N\}$ of p-degree IgA discretization matrices consists of a unique branch, the so-called 'acoustical branch', which provides a good (and convergent) approximation of practically all the exact eigenfrequencies $\{\omega_\ell^2,\ \ell = 1, \ldots, N\}$, except for few outliers.
- The spectrum $\{\omega_{\ell,\mathrm{FEA}}^2,\ \ell = 1, \ldots, N\}$ of p-degree FEA discretization matrices consists of p branches, one 'acoustical' and $p-1$ 'optical'; only the acoustical branch provides a good approximation of the first N/p eigenfrequencies $\{\omega_\ell^2,\ \ell = 1, \ldots, N/p\}$, whereas the other discrete eigenfrequencies $\omega_{\ell,\mathrm{FEA}}^2,\ \ell > N/p$, related to the optical branches, are spurious and diverge with p.

It is then clear that the approximation of the spectrum of the underalying continuous operator – in this case, the (negative) 1D Laplacian – is much better in the IgA case than in the FEA case. Similar observations also holds for differential problems different from (1); see, e.g., [8].

In this contribution, we analyze the above phenomena through the theory of Generalized Locally Toeplitz (GLT) sequences [9, 10] and the related notion of 'spectral symbol' [2] (see also [3, 4, 5]). The identification of the symbol for both IgA and FEA matrices allows us to give a compact description of the asymptotic spectrum of these matrices, and leads to a clean explanation of:

- the presence of a unique spectral branch, when the eigenvalue problem (1) is approximated by IgA;
- the appearance of p spectral branches when p-FEA is used instead of IgA.

More generally, the symbol-based analysis allows one to predict the existence of $p-k$ branches when C^k-continuous basis functions with any intermediate regularity $0 \le k \le p-1$, are used for the approximation of (1). The latter prediction is confirmed by numerical experiments, an can be extended to the d-dimensional setting through tensorization arguments [2, 4, 5].

References

[1] J.A. Cottrell, A. Reali, Y. Bazilevs, T.J.R. Hughes. Isogeometric Analysis of structural vibrations. Comput. Methods Appl. Mech. Engrg. 195 (2006), 5257–5296.

[2] C. Garoni, T.J.R. Hughes, A. Reali, S. Serra-Capizzano, H. Speleers. Smoothness versus polynomial degree: why IgA outperforms FEA in the spectral approximation. In preparation.

[3] C. Garoni, C. Manni, F. Pelosi, S. Serra-Capizzano, H. Speleers. On the spectrum of stiffness matrices arising from Isogeometric Analysis. Numer. Math. 127 (2014), 751–799.

[4] C. Garoni, C. Manni, S. Serra-Capizzano, D. Sesana, H. Speleers. Spectral analysis and spectral symbol of matrices in isogeometric Galerkin methods. In preparation.

[5] C. Garoni, S. Serra-Capizzano, D. Sesana. Spectral analysis and spectral symbol of d-variate $\mathbb{Q}_{\boldsymbol{p}}$ Lagrangian FEM stiffness matrices. Tech. Report 2014-021, Dept. Information Technology, Uppsala University, Sweden.

[6] T.J.R. Hughes, J.A. Evans, A. Reali. Finite Element and NURBS approximations of eigenvalue, boundary-value, and initial-value problems. Comput. Methods Appl. Mech. Engrg. 272 (2014), 290–320.

[7] T.J.R. Hughes, A. Reali, G. Sangalli. Duality and unified analysis of discrete approximations in structural dynamics and wave propagation: comparison of p-method Finite Elements with k-method NURBS. Comput. Methods Appl. Mech. Engrg. 197 (2008), 4104–4124.

[8] A. Reali. An Isogeometric Analysis approach for the study of structural vibrations. Journal of Earthquake Engineering 10 (2006), 1–30.

[9] S. Serra-Capizzano. Generalized locally Toeplitz sequences: spectral analysis and applications to discretized partial differential equations. Linear Algebra Appl. 366 (2003), 371–402.

[10] S. Serra-Capizzano. The GLT class as a generalized Fourier analysis and applications. Linear Algebra Appl. 419 (2006), 180–233.

How to compute efficiently the Markovian Joint Spectral Radius?

A. Cicone, N. Guglielmi, V. Y. Protasov
Università degli Studi dell'Aquila
Moscow State University

Given a finite set of matrices $\mathcal{F} = \{A_i\}_{i=1}^N$, with $A_i \in \mathbb{C}^{d \times d}$, the Joint Spectral Radius (JSR) of $\mathcal{F}$ is given by the generalization of the Gelfand's formula for the spectral radius of a matrix i.e. $\lim_{k\to\infty} \sup_{P\in\mathcal{P}_k(\mathcal{F})} \|P\|^{1/k}$, $k \in \mathbb{N}$, where $\mathcal{P}_k(\mathcal{F})$ is the set of all possible products of length k of matrices in $\mathcal{F}$. In recent works it has been proved that the JSR can be computed exactly, under suitable and general conditions, using polytope norms, see e.g. [1].

In some cases, however, not all the products are allowed, because the matrices in $\mathcal{F}$ are multiplied each other following some Markovian law. Recently Kozyakin [2] showed that it is still possibile to compute Joint Spectral Radius in the Markovian case as the classical JSR of a significantly higher dimensional set of matrices $\widehat{\mathcal{F}} = \left\{\widehat{A}_i\right\}_{i=1}^N$, with $\widehat{A}_i \in \mathbb{C}^{Nd \times Nd}$. This implies that the exact evaluation of the Markovian JSR can be achieved in general using a polytope norm in $\mathbb{C}^{Nd}$, which is a challenge task if N is large.

In this talk we address the question whether it is possible to reduce the computational complexity for the calculation of the Markovian JSR showing that it is possible to transform the problem into the evaluation of N polytope multinorms in $\mathbb{C}^d$.

As an illustrative application we shall consider the zero–stability of variable stepsize 3–step BDF formulas.

2000 MSC: 15A18, 15A60, 65F15, 65F35

References

[1] N. Guglielmi, F. Wirth, and M. Zennaro. Complex polytope extremality results for families of matrices. SIAM journal on matrix analysis and applications. 27.3 (2005), 721–743.

[2] V. Kozyakin. The Berger–Wang formula for the Markovian joint spectral radius. Linear Algebra and its Applications. 448 (2014), 315–328.

On Krylov subspace methods for the time-fractional Schrödinger equation

Roberto Garrappa, Igor Moret, Marina Popolizio
Università degli Studi di Bari
Università degli Studi di Trieste
Università del Salento, Dipartimento di Matematica e Fisica "Ennio De Giorgi", Italy

Among the various applications of the numerical evaluation of matrix functions we recently addressed the numerical solution of the time–fractional Schrödinger equation. This model is a fundamental topic in physics and it differs from the classic Schrödinger equation since the time derivative is replaced by a fractional one, according to the basic definition by Caputo [1]. Its numerical solution is far more difficult than the standard case and it is still an open problem. The approach we discuss grounds on the possibility to express its solution by means of the Mittag–Leffler (ML) function. Once a discretization is used for the spatial derivative, one faces with the problem to compute this ML function with matrix arguments, usually of large dimension. An additional difficulty is represented by the spectrum of the resulting operators since eigenvalues can belong to the imaginary axis. We analyze the *standard* Krylov projection method and the *Shift–and–Invert* approach to approximate this function; their convergence properties are discussed, together with related issues. Numerical tests are presented to confirm the strength of the approach under investigation.

References

[1] K. Diethelm. The analysis of fractional differential equations. Lecture Notes in Mathematics, Springer-Verlag (2010).

[2] R. Garrappa, I. Moret, M. Popolizio. Solving the time-fractional Schrödinger equation by Krylov projection methods. Journal of Computational Physics (2014), in press.

Stability-Preserving Parametric Model Reduction by Matrix Interpolation using Invariance Properties of Krylov Subspaces

Andreas Barthlen, Patrick Lang
Fraunhofer ITWM

With this contribution we present an idea, how one can guarantee preservation of asymptotic stability for parametric model reduction by matrix interpolation for a linear time-invariant system $\Sigma(p) = (I_n, A(p), B(p), C(p)$, based on a procedure proposed in [1]. Parametric model order reduction in general has become a popular research field over the past years, as using regular model reduction for every parameter choice is impractical, at the very least though inefficient. For that purpose methods based on matrix interpolation in particular, introduced e.g. in [2], have proven themselves to be quite useful. The basic idea is to choose a set of sampling points $p^{(i)}$ over the parameter space and to use regular model reduction methods to locally reduce the system for the chosen sampling points. The reduced parametric system can then be obtained by interpolating the matrices of the reduced systems.

Unfortunately, without modification such an interpolation will in general yield meaningless results, let alone an asymptotically stable system. A method to ensure stability was introduced in [1], where one computes matrices P_i such that the original system $\Sigma_i = \Sigma(p^{(i)})$ satisfies

$$PE > 0, \quad PA + (PA)^T < 0$$

which is a sufficient criteria for asymptotic stability. This criteria is preserved under one-sided model reduction and interpolation with nonnegative weights and has been used e.g. in [3], where it is applied to the local reduced systems. We show, that one can use the invariance of the input Krylov subspace under a nonsingular transformation P to directly compute reduced systems $\Sigma_i = (E_{r,i}, A_{r,i}, B_{r,i}, C_{r,i})$ that satisfy the asymptotic stability criteria

$$E_{r,i} > 0, \quad A_{r,i} + A_{r,i}^T < 0$$

without explicitly computing the matrices P_i. This is not only less computationally expensive than explicitly computing the transformation matrices P_i and subsequently transforming the original system, but additionally gives us more freedom regarding the local reduced systems.

References

[1] R. Eid, R. Castañé-Selga, H. Panzer, T. Wolf and B. Lohmann. Stability-preserving parametric model reduction by matrix interpolation. Mathematical and Computer Modelling of Dynamical Systems Vol. 17 No. 4 (2011), 319–335.

[2] H. Panzer, J. Mohring, R. Eid, B. Lohmann. Parametric Model Order Reduction by Matrix Interpolation. Automatisierungstechnik Vol. 58 No. 8 (2010), 475–484.

[3] M. Geuß , H. Panzer, T. Wolf, B. Lohmann. Stability Preservation for Parametric Model Order Reduction by Matrix Interpolation. Proceedings of the European Control Conference, Strasbourg, France (2014), 1098–1103.

Comparison of polynomial Krylov methods with limited memory consumption for approximating Stieltjes matrix functions

Stefan Güttel*, Marcel Schweitzer†
*The University of Manchester
†Bergische Universität Wuppertal

Approximating $f(A)\mathbf{b}$, the action of the matrix function $f(A)$ on a vector $\mathbf{b}$, is an important task in many areas of scientific computing. Especially when A is large and sparse, iterative methods which extract their approximations from a polynomial Krylov subspace $\mathcal{K}_m(A, \mathbf{b})$ are the methods of choice in many applications. One of the drawbacks of these methods is that the amount of reuired memory grows with the number of iterations which are performed, as a full basis of $\mathcal{K}_m(A, \mathbf{b})$ has to be stored. This often limits the feasible number of iterations, in particular in large scale computations. In recent years, different approaches have been employed for circumventing this problem, namely restarting approaches (which allow to only ever store a small, fixed number of vectors at a time) and rational Krylov methods (which typically need far less iterations to find an approximation of a certain accuracy than polynomial methods, but require solving a linear system in each iteration).

In some applications, solving linear systems with A (or shifted versions of it) by direct methods is not feasible, or A is only implicitly available via a routine, which, given a vector $\mathbf{v}$, returns the matrix vector product $A\mathbf{v}$. In these cases, the linear system solves in each iteration of a rational Krylov method can only be performed approximately by another iterative method. In case that A is Hermitian positive definite, a straight-forward choice is the *conjugate gradient method* which only ever requires storing a fixed number of vectors throughout all iterations. When the number of outer (rational) Krylov iterations needed for reaching the desired accuracy is small, such *inner-outer Krylov methods* can also be interpreted as Krylov methods with limited memory consumption, independent of the number of inner iterations.

Given these possible different approaches for approximating $f(A)\mathbf{b}$ in presence of limited available storage, the natural question that arises is the following:

Given a limited amount of memory (storage of at most m vectors of length N), what is the most efficent method to approximate $f(A)\mathbf{b}$?

While "most efficient" can of course have several meanings, we will focus on finding the method which requires the minimal number of matrix-vector products. For the case that f is a Stieltjes function, one can estimate the asymptotic convergence factors of the different methods. These can then be used to give guidelines for judging how well-suited which method is in a given situation. In addition, we compare these guidelines to the performance observed in actual numerical computations for the different methods.

Orthogonal projection vs. Oblique projection in Krylov subspace recycling

Matthias Bolten, Nemanja Božović, Andreas Frommer
Bergische Universität Wuppertal

Many problems in engineering, numerical simulations in physics etc. require the solution of long sequences of slowly changing linear systems. The paper [1] proposes an algorithm that reduces the cost of solving subsequent systems by recycling selected subspaces generated for the previous systems. The subspace which is to be recycled is chosen in a way that the eigenvalues of smallest magnitude are approximately deflated. This can substantially improve the convergence of the method. Their algorithm, called GCRO-DR, uses an orthogonal projection for deflation. What we propose is an algorithm that uses an oblique projection in the spirit of [2] for deflating the eigenvalues of smallest magnitude. This oblique projection requires approximations to right *and left* eigenvectors, and we propose a way of getting the approximations to the left eigenvectors, without having to build a Krylov subspace with respect to A^H, which saves work. We will also show some numerical comparisons.

References

[1] Michael L. Parks, Eric de Sturler, Greg Mackey, Duane D. Johnson and Spandan Maiti, Recycling Krylov subspaces for sequences of linear systems, SIAM Journal on Scientific Computing, 28 (2006), 1651–1674.

[2] Martin H. Gutknecht, Spectral deflation in Krylov solvers: A theory of coordinate space based methods, ETNA Journal, 39 (2012), 156–185.

A new approach for preconditioning discontinuous Galerkin discretizations

Martin J. Gander, Soheil Hajian
University of Geneva

Domain decomposition preconditioners and in particular the additive Schwarz method are favorite preconditioners for classical finite element methods (FEM). There is a huge effort in designing similar preconditioners for discontinuous Galerkin (DG) discretizations. It has been shown that additive Schwarz methods use different mechanism for convergence when applied to a DG discretization compared to the classical FEM. More precisely, additive Schwarz methods, when applied to DG, use a non-overlapping Robin transmission condition for the communication between subdomains. This is exactly the same transmission condition that optimized Schwarz methods (OSM) use to obtain fast convergence. In this talk we present an OSM preconditioner for a particular DG discretization along with theoretical convergence estimates.

References

[1] Martin J. Gander, Soheil Hajian. Analysis of Schwarz Algorithms for a Discontinuous Galerkin Method. `arXiv:1404.4518 [math.NA]`

Rational least squares approximation via RKFIT

Mario Berljafa and Stefan Güttel
The University of Manchester

For given matrices $\{A, F\} \subset \mathbb{C}^{N\times N}$ and a vector $\mathbf{v} \in \mathbb{C}^N$, we consider the problem of finding a rational function R_m^* of type (m, m) such that

$$\|F\mathbf{v} - R_m(A)\mathbf{v}\|_2^2 \text{ is small.}$$

We propose an iterative algorithm called RKFIT for its solution. At each iteration RKFIT constructs a rational Krylov space and manipulates an associated Arnoldi decomposition to find better approximations to the poles of R_m^*. In the special case when A and F are diagonal matrices, we can compare RKFIT to the popular vector fitting algorithm by Gustavsen and Semlyen (1999).

RKFIT is part of a MATLAB Rational Krylov Toolbox available for download from

`http://www.guettel.com/rktoolbox`

A fast nonstationary preconditioning strategy for ill-posed problems, with application to image deblurring

Marco Donatelli
Dipartimento di Scienza e Alta Tecnologia, Università dell'Insubria, 22100 Como, Italy.

We consider the iterative solution of ill-posed equations

$$Tx = y\,, \tag{1}$$

where $T : \mathcal{X} \to \mathcal{Y}$ is a linear operator between two Hilbert spaces $\mathcal{X}$ and $\mathcal{Y}$ with unbounded (Moore-Penrose) generalized inverse. Hence, problem (1) has to be regularized for a numerical solution. As usually often in the applications, we assume that, instead of the exact data $y \in \mathcal{Y}$ of (1), we are only given approximate data $y^\delta \in \mathcal{Y}$ with $\|y^\delta - y\| \le \delta$, where $\delta \ge 0$ is the corresponding noise level.

When linear equations with $T^*T + \alpha I$ or $TT^* + \alpha I$, with $\alpha > 0$, are computationally expensive to solve, taking an appropriate approximation C of T the following algorithm has been proposed in [4]. Starting with an initial guess x_0 we compute, for $n = 0, 1, 2, \ldots$,

$$h_n = C^*(CC^* + \alpha_n I)^{-1} r_n\,, \qquad r_n = y^\delta - Tx_n\,, \tag{2a}$$

and set

$$x_{n+1} = x_n + h_n\,. \tag{2b}$$

Note that the linear equation (2a) is equivalent to minimizing the Tikhonov functional

$$\|Ch_n - r_n\|^2 + \alpha_n \|h_n\|^2 \longrightarrow \min. \tag{3}$$

over $h_n \in \mathcal{X}$, where C is the aforementioned approximation of T, and α_n is the associated regularization parameter.

In the literature on iterative solvers for (usually well-posed) problems of the form (1) the operator

$$P = C^*(CC^* + \alpha_n I)^{-1} \tag{4}$$

in (2a) would be called *preconditioner*. As we will select different regularization parameters α_n in each iteration, it is therefore appropriate to call the new scheme (2) a nonstationary preconditioned iteration. We provide a theoretical analysis of the new scheme, using regularization parameters that are chosen by a certain adaptive strategy. On the other hand, the parameter sequence $(\alpha_n)_n$ could be also simply defined by the geometric sequence

$$\alpha_n = \alpha q^n, \qquad n = 0, 1, 2, \ldots\,, \tag{5}$$

where $\alpha > 0$ and $0 < q \le 1$.

For image deblurring problems the matrix T has a structure depending on the boundary conditions and the approximation C can be a discrete convolution operator that operates entirely in the Fourier domain. Hence the matrix vector product with the preconditioner P in (4) can be computed by two fast Fourier transforms (FFTs). The numerical performance of this method turns out to be superior to state of the art iterative methods for least square problems, including the conjugate gradient iteration for the normal equation, with and without additional preconditioning.

We discuss also possible improvements of the iteration (2). For instance the nonnegative constraint can be easily enforced, preserving the convergence of the method, by simply projecting every iteration into the nonnegative cone. Furthermore, other weighted norms could be used to define h_n in (2a) in order to preserve special features of the computed solution [1].

Regularization preconditioners can be used also inside recent thresholding iterative methods for image deblurring problems. We discuss the synthesis approach, but our proposal can be applied also to the analysis approach. Thanks to the well-known property that every image has a sparse representation in the wavelet domain, some recent deblurring models are based on a regularization term that promotes the sparsity of the wavelets coefficients

$$\hat{x} = Wx,$$

where W defines the wavelet transform. For instance the regularized problem can be formulated in terms of wavelets coefficients as

$$\min_{\hat{x}}\{\mu\|\hat{x}\|_1 + \|\hat{x}\|_2^2 \,:\, TW^T\hat{x} = y^\delta\}, \tag{6}$$

where $\|\cdot\|_p$ denotes the p-norm, $p \geq 1$, and $\mu > 0$ is a regularization parameter to be chosen. The solution of (6) can be computed by the linearized Bregman splitting that converges very slowly. Hence a preconditioning strategy is usually employed, obtaining the modified linearized Bregman iteration based on a regularizing preconditioner diagonalized by FFTs [3]. We show that this approach can provide low quality restored images when appropriate boundary conditions are imposed to reduce possible boundary artifacts. Accordingly we propose a new modified linearized Bregman iteration based on the preconditioner (4) that improves the quality of the restoration and save some computational cost at the same time [2]. Similarly to (2) our simple iteration is

$$\begin{cases} z_{n+1} = z_n + WC^T(CC^T + \alpha_n I)^{-1}(y^\delta - TW^T\hat{x}_n), \\ \hat{x}_{n+1} = S_\mu(z_{n+1}), \end{cases}$$

where S_μ is the soft-thresholding function defined component-wise as

$$S_\mu(\xi) = \operatorname{sgn}(\xi)\,(|\xi| - \mu)_+ \,.$$

Some numerical experiments show that our proposed preconditioners provide accurate and fast restorations when compared with the state of the art methods with sparsity constraints or total variation regularization.

References

[1] A. Buccini, M. Donatelli. Approximated projected generalized iterated Tikhonov. Manuscript 2015.

[2] Y. Cai, M. Donatelli, T. Z. Huang. Regularization preconditioners for frame-based image deblurring without boundary artifacts. Submitted 2014.

[3] J. F. Cai, S. Osher, Z. Shen. Linearized Bregman iterations for frame-based image deblurring. SIAM J. Imaging Sci. 2–1 (2009), 226–252.

[4] M. Donatelli, M. Hanke. Fast nonstationary preconditioned iterative methods for ill-posed problems, with application to image deblurring. Inverse Problems. 29 (2013), 095008.

Iterated fractional Tikhonov regularization

Davide Bianchi, Alessandro Buccini, Marco Donatelli, Stefano Serra-Capizzano
Dipartimento di Scienza e Alta Tecnologia, Università dell'Insubria, 22100 Como, Italy

We consider linear operator equations of the form

$$Kx = y, \tag{1}$$

where $K : \mathcal{X} \to \mathcal{Y}$ is a compact linear operator between Hilbert spaces $\mathcal{X}$ and $\mathcal{Y}$. We assume y to be attainable, i.e., that problem (1) has a solution $x^\dagger = K^\dagger y$ of minimal norm. Here $K^\dagger$ denotes the (Moore-Penrose) generalized inverse operator of K, which is unbounded when K is compact, with infinite dimensional range. Hence problem (1) is ill-posed and has to be regularized in order to compute a numerical solution. We want to approximate the solution $x^\dagger$ of the equation (1), when only an approximation y^δ of y is available with

$$\|y^\delta - y\| \leq \delta,$$

where δ is called the noise level. Since $K^\dagger y^\delta$ is not a good approximation of $x^\dagger$, we approximate $x^\dagger$ with $x_\alpha^\delta := R_\alpha y^\delta$ where $\{R_\alpha\}$ is a family of continuous operators depending on a parameter α. A classical example is the Tikhonov regularization defined by $R_\alpha = (K^*K + \alpha I)^{-1}K^*$, where I denotes the identity and K^* the adjoint of K.

Recently, new Tikhonov based regularization methods have been proposed in [1], [2] and [3], under the name of fractional Tikhonov, to reduce the oversmoothing property of the Tikhonov regularization in standard form, in order to preserve the details of the approximated solution. Their regularization and convergence properties have been previously investigated showing that they are of optimal order.

In this talk, we firstly provide saturation results similar to the well-known saturation result for Tikhonov regularization: let $R(K)$ be the range of K and let Q be the orthogonal projector onto $\overline{R(K)}$, if

$$\sup\left\{\|x_\alpha^\delta - x^\dagger\| : \|Q(y - y^\delta)\| \leq \delta\right\} = o(\delta^{\frac{2}{3}}),$$

then $x^\dagger = 0$, as long as $\overline{R(K)}$ is not closed. Such results motivated us to introduce the iterated versions of fractional Tikhonov methods in the same spirit of the iterated Tikhonov method. We prove that those iterated methods can overcome the afore-mentioned saturation results.

Afterwards, inspired by the works [4, 5] we introduce the nonstationary variants of our iterated methods. Differently from the nonstationary iterated Tikhonov, we have two nonstationary sequences of parameters. In the noise free case, we give sufficient conditions on these sequences to guarantee the convergence providing also the corresponding convergence rates. In the noise case, we show the stability of the proposed iterative schemes proving that they are regularization methods. Finally, few selected examples confirm the previous theoretical analysis, showing that a proper choice of the nonstationary sequences of parameters can provide better restorations compared to the classical iterated Tikhonov.

References

[1] E. Klann, R. Ramlau. Regularization by fractional filter methods and data smoothing. Inverse Problems 24 (2) (2008), 025018.

[2] E. M. Hochstenbach, L. Reichel. Fractional Tikhonov regularization for linear discrete ill-posed problems. BIT, 51 (1) (2011), 197–215.

[3] D. Gert, E. Klann, R. Ramlau, L. Reichel. On fractional Tikhonov regularization. Private notes (2014).

[4] M. Hanke, C. W. Groetsch. Nonstationary iterated Tikhonov regularization. J. Optim. Theory Appl. 98 (1) (1998), 37–53.

[5] M. Brill, E. Schock. Iterative solution of ill-posed problems-a survey. Model optimization in exploration geophysics, ed. A. Vogel, Vieweg, Braunschweig (1987).

Spectral behavior of preconditioned non-Hermitian multilevel block Toeplitz matrices with matrix-valued symbol

M. Donatelli, C. Garoni, M. Mazza, S. Serra-Capizzano, D. Sesana
University of Insubria,
Department of Science and High Technology, Via Valleggio 11, 22100 Como, Italy

This contribution is devoted to preconditioning strategies for non-Hermitian multilevel block Toeplitz linear systems associated with a multivariate Lebesgue integrable matrix-valued symbol. In particular, we consider special preconditioned matrices, where the preconditioner has a band multilevel block Toeplitz structure, and we complement known results on the localization of the spectrum with global distribution results for the eigenvalues of the preconditioned matrices. In this respect, our main result is as follows. Let $I_k := (-\pi, \pi)^k$, let $\mathcal{M}_s$ be the linear space of complex $s \times s$ matrices, and let $f, g : I_k \to \mathcal{M}_s$ be functions whose components f_{ij}, $g_{ij} : I_k \to \mathbb{C}$, $i, j = 1, \ldots, s$, belong to L^∞. Consider the matrices $T_n^{-1}(g)T_n(f)$, where $n := (n_1, \ldots, n_k)$ varies in $\mathbb{N}^k$ and $T_n(f), T_n(g)$ are the multilevel block Toeplitz matrices of size $n_1 \cdots n_k s$ generated by f, g. Then $\{T_n^{-1}(g)T_n(f)\}_{n\in\mathbb{N}^k} \sim_\lambda g^{-1}f$, i.e. the family of matrices $\{T_n^{-1}(g)T_n(f)\}_{n\in\mathbb{N}^k}$ has a global (asymptotic) spectral distribution described by the function $g^{-1}f$, provided g possesses certain properties (which ensure in particular the invertibility of $T_n(g)$ for all n) and the following topological conditions are met: the essential range of $g^{-1}f$, defined as the union of the essential ranges of the eigenvalue functions $\lambda_j(g^{-1}f)$, $j = 1, \ldots, s$, does not disconnect the complex plane and has empty interior. This result generalizes the one obtained in [1], concerning the non-preconditioned case $g = 1$.
Numerical experiments confirm the theoretical analysis and suggest the choice of optimal GMRES preconditioning techniques to be used for the considered linear systems. Moreover, the obtained results can be used for the spectral analysis of the Preconditioned Hermitian/Skew-Hermitian Splitting (PHSS) method applied to multilevel block Toeplitz linear systems with a Hermitian positive definite multilevel block Toeplitz matrix as a preconditioner [3].

References

[1] M. Donatelli, M. Neytcheva, S. Serra-Capizzano. Canonical eigenvalue distribution of multilevel block Toeplitz sequences with non-Hermitian symbols. Oper. Theory Adv. Appl. 221 (2012), 269–291.

[2] M. Donatelli, C. Garoni, M. Mazza, S. Serra-Capizzano, D. Sesana. Spectral behavior of preconditioned non-Hermitian multilevel block Toeplitz matrices with matrix-valued symbol. Appl. Math. Comput. 245 (2014), 158–173.

[3] M. Donatelli, C. Garoni, M. Mazza, S. Serra-Capizzano, D. Sesana. Preconditioned HSS method for large multilevel block Toeplitz linear systems via the notion of matrix-valued symbol. Numer. Linear Algebra Appl., submitted.

Local Fourier Analysis of Pattern Structured Operators

M. Bolten, K. Kahl, H. Rittich
Applied Computer Science Group University of Wuppertal

Multigrid methods [5] are used to compute the solution u of the system of equations

$$Lu = f\,,$$

where L is typically a discretization of a partial different equations (PDE) and f a corresponding, given right hand side. Local Fourier Analysis (LFA) [2, 5, 6] is well known to provide quantitative estimates for the speed of convergence of multigrid methods, by analyzing the involved operators in the frequency domain.

For the initial formulation of LFA [1] it was crucial to assume that all involved operators have constant coefficients. For many PDE operators the coefficients vary continuously in space. Thus if the grid is fine enough the discrete operator L will only vary slightly between neighboring grid points and hence can be well approximated by an operator with *locally* constant coefficients. Thus constant coefficient are often reasonable assumption.

However, when analyzing more complex problems or even the multigrid method as a whole this assumption is too restrictive. Interpolation and restriction operators typically act differently on variables that have a coarse grid representative and those who do not have one. Another example are patter relaxation schemes like the Red-Black Gauß-Seidel method where red points of the grid are treated differently from the black ones.

It is possible to analyze these cases [3, 4] when allowing for interaction of certain frequencies (see also [5, 6]). Even more, it turns out that when we allow for more frequencies to interact we can analyze operators given by increasingly complex patterns. In our talk we will illustrate a general framework for analyzing pattern structured operators, i.e., operators whose action is invariant under certain shifts of the input function. Furthermore, we discuss different applications.

References

[1] A. Brandt. Multi-level adaptive solutions to boundary-value problems. *Math. Comput.*, 31(138):333–390, 1977.

[2] A. Brandt. Rigorous quantitative analysis of multigrid, I: Constant coefficients two-level cycle with L_2-norm. *SIAM J. Numer. Anal.*, 31(6):1695–1730, 1994.

[3] K. Stüben and U. Trottenberg. Multigrid methods: fundamental algorithms, model problem analysis and applications. In *Multigrid methods (Cologne, 1981)*, volume 960 of *Lecture Notes in Math.*, pages 1–176. Springer, Berlin-New York, 1982.

[4] C.-A. Thole and U. Trottenberg. Basic smoothing procedures for the multigrid treatment of elliptic 3D operators. *Appl. Math. Comput.*, 19(1-4):333–345, 1986.

[5] U. Trottenberg, C. W. Oosterlee, and A. Schüller. *Multigrid*. Academic Press, 2001.

[6] R. Wienands and W. Joppich. *Practical Fourier Analysis For Multigrid Methods*. Chapman Hall/CRC Press, 2005.

Fast Recovery and Approximation of Hidden Cauchy Structure

Jörg Liesen, Robert Luce
TU Berlin

A *Cauchy matrix* $C(s,t) \in \mathbb{C}^{m,n}$ is defined by $n+m$ data points $s \in \mathbb{C}^m$, $t \in \mathbb{C}^n$ satisfying $s_i \neq t_j$ for all i,j, the *Cauchy points*, via the relation

$$C(s,t) = [c_{ij}] := \left[\frac{1}{s_i - t_j}\right].$$

One can take advantage of the peculiar structure of Cauchy matrices, for example when computing matrix vector products ("the fast multipole method") or computing an LU factorization, resulting in algorithms that may be faster than general purpose algorithms by an order of magnitude. These algorithms require that the Cauchy points s,t are explicitly given.

In this contribution we first study the problem of *recovering* the Cauchy points s,t, if only the entries of $C(s,t)$ are known. We give a simple $\mathcal{O}(m+n)$ algorithm for that task. Hence it can be checked in $\mathcal{O}(mn)$, whether any given matrix is a Cauchy matrix. The second problem we study in this work is *approximating* a given matrix $A \in \mathbb{C}^{m,n}$ with a Cauchy matrix. For this one would ideally like to solve the nonlinear problem

$$\min_{s\in\mathbb{C}^m, t\in\mathbb{C}^n} \|C(s,t) - A\|_F. \tag{1}$$

We assume $a_{ij} \neq 0$ for all i,j, which is guaranteed when A *is* a Cauchy matrix.

Instead of solving the presumably difficult problem (1), we solve the ordinary linear least squares problem

$$\min_{s\in\mathbb{C}^m, t\in\mathbb{C}^n} \|D(s,t) - A^{[-1]}\|_F, \tag{2}$$

where $D(s,t) := [s_i - t_j]$, and the superscript $[-1]$ denotes the elementwise inverse of a matrix. We derive an $\mathcal{O}(mn)$ algorithm for the optimal solution of (2). Note that the complexity of our algorithm is linear in the input, and is thus of optimal complexity. Further we prove approximation bounds that relate the solutions of (1) and (2). We pay special attention to the problem of approximating Cauchy points s,t from a given *noisy* Cauchy matrix, i.e. where $A = C(s,t) + N$, and the matrix N represents a source of data uncertainty.

Parallel Tensor Sampling

Lars Grasedyck, Christian Löbbert
Institut für Geometrie und Praktische Mathematik, RWTH Aachen

We consider the problem to reconstruct a tensor in the class of hierarchical low rank tensors from a small number of samples. The number of samples is proportional to the number of degrees of freedom in the respective tensor format [1], but there is in general no guarantee that the tensor can be reconstructed unless a certain restricted isometry property is assumed. Practical examples provide evidence that the reconstruction works in general pretty good, e.g. in the case of a parametric elliptic PDE where many parameters may enter into the diffusion coefficient, the right-hand side, the boundary conditions or the domain itself [2]. In such a case the method is competitive with, e.g. adaptive sparse grids or quasi Monte Carlo. Our aim is to parallelize the sampling process. Unlike for Monte Carlo methods, our sampling is sequential in nature and we modify it to allow several levels of parallelism: We can either apply parallelism for a small number of samples (~ 100) without any additional overhead, for a medium number of samples (~ 10000) with a mild overhead, or finally a complete parallelism for all samples ($\gg 10000$) albeit with a considerable overhead [3]. The overhead does not influence the runtime behaviour but rather the total energy consumption. We thus have three gears for three levels of parallelism ranging from energy efficient (first gear) to high speed (third gear).

References

[1] J. Ballani, L. Grasedyck, M. Kluge. Black Box Approximation of Tensors in Hierarchical Tucker Format. Num. Lin. Alg. Appl. 438(2): 639–657 (2013).

[2] J. Ballani, L. Grasedyck: Hierarchical tensor approximation of output quantities of parameter-dependent PDEs. Preprint 385, IGPM, RWTH Aachen, `http://www.igpm.rwth-aachen.de/preprint`.

[3] L. Grasedyck, C. Löbbert: Parallel Tensor Sampling in the Hierarchical Tucker Format. In preparation.

Multigrid methods for tensor structured problems

Matthias Bolten*, Karsten Kahl*, Sonja Sokolovic*
*Bergische Universität Wuppertal

We look at linear systems whose matrix is related to a tensor structure of the kind $A = \sum_i \bigotimes_j E_i^j$. These problems are found, e.g., in stochastic Markov chain models or in the solution of high-dimensional Poisson equations. Due to the tensor structure of the models, the dimension of the operator A grows rapidly for larger models. So the tensor structure has to be exploited for solving these systems efficiently. Multigrid methods have proven to be well-suited for these kind of problems, e.g., the Poisson equation. To be able to also use multigrid methods for high-dimensional models we build a multigrid method which keeps the tensor structure intact to guarantee computational savings on all grids. We investigate the question what kind of smoothing and coarsening have to be chosen to guarantee a nice convergence in this case and present an approach to adapt the algebraic multigrid framework to this tensor setting using tensor truncation techniques.

Hierarchical tensor approximation of parameter-dependent PDEs

Jonas Ballani
EPF Lausanne

Parametric PDEs appear in a large number of applications, as e.g. in uncertainty quantification or optimisation. Typically, the amount of data to approximate and represent the solution scales exponentially in the parameter dimension. Therefore, a crucial task is to develop special numerical techniques that rely on data-sparsity in order to cope even with high parameter dimensions. In this talk, we will discuss low-rank tensor techniques that allow to reduce the complexity to a linear dependence on the parameter dimension. In particular, our aim is to adaptively construct an approximation of the solution in the hierarchical tensor format from a relatively small set of data samples. Once this approximation from an offline computation is available, the evaluation of quantities of interest becomes a cheap online task. Moreover, the explicit tensor representation can be used to compute stochastic properties of the solution in a straightforward way. The potential of this approach is illustrated by numerical examples.

This is joint work with Lars Grasedyck (RWTH Aachen) and Daniel Kressner (EPF Lausanne).

On the block and global methods for linear systems with multiple right hand sides

A. Frommer, S. Rashedi, g. Ebadi
Bergische Universität Wuppertal Germany,
Tabriz University, Iran

We present new block and global methods for solving large nonsymmetric linear systems of equations with multiple right-hand. Dealing with many right hand sides at the same time allows to use multiplications of a sparse matrix times with block-vectors which are particularly efficient on current processor architectures.

The block version of BiCGSTAB introduced in [1] may suffer from numerical instability when the number of right-hand sides is large and we improve numerical stability by using a QR- factorization of the descent and/or the residual matrix. We also compare these new implementations with the global BiCGSTAB method, presented in[1]. Moreover, as observed in [2], for problems arising from partial differential equations and having complex eigenvalues, the block BiCGSTAB and global BiCGSTAB methods have a tendancy to not work well. We therefore introduce a new version of block BiCGSTAB and global BiCGSTAB in the spirit of the BICGSTAB2 method known for single right hand sides.

Several numerical examples demonstrate that the different variants of the global methods can achieve a smoothed residual and may be more competitive than the block solvers.

References

[1] A. El Guennouni, K. Jbilou, H. Sadok, A block version of BiCGSTAB for linear systems with multiple right-hand sides, Electron. Trans. Numer. Anal. 16 (2003), 129–142.

[2] M.H. Gutknecht, Variants of BiCGStab for matrices with complex spectrum, IPS Research Report No.(1991), 91–14.

[1] J. Zhang, H. Dai, Global CGS algorithm for linear systems with multiple right-hand sides, Numer. Math. J. Chin. Univ. 30 (2008) 390–399 (in Chinese).