Householder Symposium XVII
Book of Abstracts



Zeuthen, Germany
June 1-6, 2008

Householder Symposium XVII
June 1-6, 2008, Zeuthen, Germany

The Householder Symposium on Numerical Linear Algebra in Zeuthen, Germany, is the seventeenth in a series that was initiated by Alston S. Householder (1904-1993), Director of the Mathematics Division of Oak Ridge National Laboratory and Ford Professor at the University of Tennessee. These international meetings were devoted to matrix computations and linear algebra and were held in Gatlinburg, Tennessee. The last "Gatlinburg" conference was held in 1969 on the occasion of Householder's retirement. At the time, it was decided to continue the meetings but vary the place. Since then meetings have been held at three year intervals in a variety of venues.

The conferences are run in tandem by the permanent Householder Committee and a Local Organizing Committee. This year the chair of the Householder Committee is Angelika Bunse-Gerstner (University of Bremen), and the Local Organizing Committee consists of Jörg Liesen, Volker Mehrmann and Reinhard Nabben (all TU Berlin). The complete list of the Householder Symposia is as follows:

	Year	Place	Local organizers
I	1961	Gatlinburg, USA	A. S. Householder
II	1963	Gatlinburg, USA	A. S. Householder, F. W. J. Olver
III	1964	Gatlinburg, USA	A. S. Householder
IV	1969	Gatlinburg, USA	A. S. Householder
V	1972	Los Alamos, USA	R. S. Varga
VI	1974	Hopfen am See, BRD	F. L. Bauer
VII	1977	Asilomar, USA	G. H. Golub
VIII	1981	Oxford, England	L. Fox, J. H. Wilkinson
IX	1984	Waterloo, Canada	J. A. George
X	1987	Fairfield Glade, USA	G. W. Stewart, R. C. Ward
XI	1990	Tylosand, Sweden	Å. Björck
XII	1993	Lake Arrowhead, USA	T. F. Chan, G. H. Golub
XIII	1996	Pontresina, Switzerland	W. Gander, M. H. Gutknecht
XIV	1999	Whistler, B.C., Canada	J. M. Varah
XV	2002	Peebles, Scotland	P. Knight, A. Ramage, A. Wathen
XVI	2005	Seven Springs, USA	J. Barlow, D. Szyld, H. Zha
XVII	2008	Zeuthen, Germany	J. Liesen, V. Mehrmann, R. Nabben

The meeting is also the occasion for the award of the Householder Prize for the best thesis in Numerical Linear Algebra. This prize was established at the Gatlinburg Symposium in 1969 to recognize the outstanding contributions of Alston S. Householder. This year the chair of the Householder Prize Committee is Charles Van Loan (Cornell University).

Local Organizing Committee

Jörg Liesen (TU Berlin)
Volker Mehrmann (TU Berlin)
Reinhard Nabben (TU Berlin)

Program Committee

Angelika Bunse-Gerstner (Chair, University of Bremen)
Tony Chan (UCLA)
Alan Edelman (MIT)
Nick Higham (University of Manchester)
Ilse Ipsen (North Carolina State University)
Roy Mathias (University of Birmingham)
Michael Overton (Courant Institute, New York University)
Zdeněk Strakoš (Czech Academy of Sciences)
Paul Van Dooren (Catholic University of Louvain)
Charles Van Loan (Cornell University)

Householder Prize Committee

James Demmel (University of California, Berkeley)
Howard Elman (University of Maryland)
Sabine Van Huffel (K.U. Leuven)
Volker Mehrmann (TU Berlin)
Charles Van Loan (Chair, Cornell University)

Sponsorship

We are grateful to the following organizations for their sponsorship of the meeting.

Walter de Gruyter GmbH & Co. KG
Deutsche Forschungsgemeinschaft
Gesellschaft für Angewandte Mathematik und Mechanik e.V.
International Linear Algebra Society
The MathWorks, Inc.
The Numerical Algorithms Group Ltd (NAG)
National Science Foundation (NSF)
Technische Universität Berlin
Wiley InterScience

List of Abstracts

Pierre-Antoine Absil
Differential-Geometric Foundations of Jacobi-Davidson Methods 12
Athanasios C. Antoulas
Model reduction by means of fast frequency sweep 14
Mario Arioli
Matrix square root and interpolation spaces 17
Owe Axelsson
An additive subspace correction method combined with an ADI-type splitting 18
Zhaojun Bai
Robust and Efficient Methods for Multi-Length Scale Matrix Computations and Applications in Quantum Mechanical Simulations 18
Ulrike Baur
Control-oriented Model Reduction for Parabolic Systems 19
Peter Benner
Solving Algebraic Riccati Equations for Stabilization of Incompressible Flows 21
Michele Benzi
Localization Phenomena in Matrix Functions: Theory, Algorithms, Applications 23
Timo Betcke
Improving computed eigenpairs of polynomial eigenvalue problems by scaling and iterative refinement 26
David Bindel
Error Bounds and Error Estimates for Nonlinear Eigenvalue Problems 28
Åke Björck
Band Householder/Lanczos Methods for Least Squares with Multiple Right Hand Sides 28
Matthias Bollhöfer
Algebraic multigrid for Helmholtz equations 29
Shreemayee Bora
Structured eigenvalue condition number and backward error of a class of polynomial eigenvalue problems 30
Karen Braman
Middle Deflations in the QR Algorithm 33
Angelika Bunse-Gerstner
Model reduction for large scale MIMO-systems via rational tangential Hermite interpolation 36
Xiao-Wen Chang
Solving Integer Linear Least Squares Problems 36
John M. Conroy
The Linear Algebra of a Multi-document Summarization System 38
Achiya Dax
From Eigenvalues to Singular Values: The Rectangular Quotients Equality and Related Issues 40

Eric de Sturler
Matrix Analysis Problems arising in Fast Implementations of the Quantum Monte Carlo Method 41
James Demmel
Communication avoiding algorithms for linear algebra 44
Inderjit S. Dhillon
The Log-Determinant Divergence and Its Applications 44
Sue Dollar
Sparse direct methods need not suffer under dense rows 45
Marco Donatelli
The Antireflective Algebra with Applications to Image Deblurring 45
Froilán M. Dopico
An Orthogonal and Symmetric High Relative Accuracy Algorithm for the Symmetric Eigenproblem 48
Petros Drineas
The Column Subset Selection Problem 50
Iain Duff
Practical Rank Determination for Square and Rectangular Sparse Matrices . 52
Ioana Dumitriu
Fast Linear Algebra is Stable 54
Alan Edelmann
Mathematical Software Going Forward 55
Lars Eldén
Solving Ill-Posed Cauchy Problems using Rational Krylov Methods 55
Howard C. Elman
Preconditioners for Eigenvalue Problems Arising in Linear Stability Analysis 57
Mark Embree
Damped Mechanical Systems: Spectra, Pseudospectra, Structured Perturbations 59
Yogi A. Erlangga
Multilevel Krylov-Multigrid Method for the 2D Indefinite Helmholtz Equation 60
Oliver G. Ernst
A Posteriori Error Estimators for Krylov Subspace Approximations of Matrix Functions 63
Heike Faßbender
A structured polynomial eigenproblem arising in the analysis of time delay systems and related polynomial eigenproblems 64
Melina A. Freitag
Inexact Preconditioned Shift-and-Invert Arnoldi's Method and Implicit Restarts for Eigencomputations 66
Andreas Frommer
Algebraic Multigrid Methods for Laplacians of Graphs 67
Efstratios Gallopoulos
Linear Algebra Tools for Link Analysis and Dimensionality Reduction 69
Martin J. Gander
An Optimal Complexity Algorithm for Non-Matching Grid Projections 70

Luc Giraud
Convergence in backward error of relaxed GMRES 72
Robert Granat
Parallel Eigenvalue Reordering in Real Schur Forms 73
Serge Gratton
Quasi-Newton formula, matrix nearness problems and preconditioning 76
Anne Greenbaum
Upper and Lower Bounds on Norms of Functions of Matrices 78
Chen Greif
A Conjugate Gradient Method for Skew-Symmetric Matrices 79
Luka Grubišić
Numerical linear algebra for spectral theory of block operator matrices 79
Serkan Gugercin
Structured perturbation theory for inexact Krylov projection methods in model reduction 80
Chun-Hua Guo
Convergence Rate of the Cyclic Reduction Algorithm for a Nonlinear Matrix Equation 81
Nicholas J. Higham
Computing the Fréchet Derivative of the Matrix Exponential, with an application to Condition Number Estimation 83
Michiel Hochstenbach
On the computation of real and purely imaginary eigenvalues of large sparse matrices 84
Thomas Huckle
Frobenius Norm Minimization and Probing for Preconditioning 85
Bruno Iannazzo
Rational Matrix Iterations in Numerical Linear Algebra 86
Ilse Ipsen
Perturbation Bounds for Determinants and Characteristic Polynomials 87
Charles Johnson
A question about Eigenvalues of Tridiagonal Matrices and Motivation for it . 88
Bo Kågström
Product Eigenvalue Problems in Applications: Computing Periodic Deflating Subspaces Associated with a Specified Set of Eigenvalues 88
Misha E. Kilmer
A Multilevel, Modified Total Least Norm Approach to Signal and Image Deblurring 90
Andrew V. Knyazev
Rayleigh-Ritz majorization error bounds with applications to FEM and subspace iterations 91
Plamen Koev
Computing Eigenvalues of Random Matrices 92
Christos Kravvaritis
On the Growth Factor for Hadamard Matrices 92
Daniel Kressner
Computation of Structured Pseudospectra 95

James V. Lambers
Robust Computation of Off-Diagonal Elements of Functions of Matrices . . . 96
Julien Langou
Reduce factorizations . . . 99
Jörg Liesen
On nonsymmetric saddle point matrices that allow conjugate gradient iterations . . . 100
Niloufer Mackey
Theory and Algorithms for Structured Matrix Problems . . . 101
D. Steven Mackey
Linearizations of Singular Matrix Polynomials and the Recovery of Minimal Indices . . . 104
Michael W. Mahoney
Statistical Leverage and Improved Matrix Algorithms . . . 106
Carla D. Martin
A Tensor SVD . . . 107
Nicola Mastronardi
On Things of Little Interest: How to Compute Low Paying Annuities . . . 108
Roy Mathias
Generalising Results about the Hermitian Eigenvalue Problems to the Quadratic Eigenvalue Problem . . . 109
Karl Meerbergen
Recycling Ritz vector for the solution of parameterized linear systems . . . 110
Christian Mehl
Hamiltonian Jacobi Methods: Sweeping Away Convergence Problems . . . 111
Volker Mehrmann
Numerical computation of Sacker-Sell exponents for differential-algebraic equations . . . 112
Hermann Mena
Numerical Solution of Large-Scale Differential Riccati Equations . . . 113
Emre Mengi
¿From Wilkinson distance to nearest non-minimal dynamical systems : Lipschitz-based optimization of singular values . . . 114
Gérard Meurant
New results on the Lanczos and Conjugate Gradient algorithms . . . 117
Ronald B. Morgan
Solution of Linear Equations for Quantum Chromodynamics . . . 118
Julio Moro
Directional perturbation in structured eigenproblems . . . 120
Reinhard Nabben
Multilevel Krylov Methods based on Projections derived from Deflation, Domain Decomposition, and Multigrid Iterations . . . 121
James G. Nagy
Separable Nonlinear Least Squares Problems in Image Processing . . . 122
Maya Neytcheva
Exploiting the Finite Element framework to construct two-by-two block preconditioners . . . 124

Esmond G. Ng
Solving Sparse Least-Squares Problems Using Perturbed QR Factorizations . 127
Nancy K. Nichols
Model Reduction in Variational Data Assimilation 128
Dianne P. O'Leary
Regularization by Residual Periodograms 129
Michael L. Overton
Preconditioning for Semidefinite Programming 130
Archara Pacheenburawana
A One-sided Jacobi Method for Computing the Symplectic SVD 131
Chris Paige
Orthonormal Completion of an array of Unit Length Vectors 131
Bob Plemmons
Nonnegative Least Squares and Extensions 132
Bor Plestenjak
Jacobi–Davidson Method for Two-Parameter Eigenvalue Problems 133
Marielba Rojas
LSTRS: MATLAB Software for Large-Scale Trust-Region Subproblems and Regularization 135
Miro Rozložník
Numerical behavior of saddle–point solvers 136
Axel Ruhe
Getting orthogonal eigenvectors with no reorthogonalization 139
Yousef Saad
Efficient linear algebra methods in data mining 140
Berkant Savas
Grassmann-Quasi-Newton Algorithms for Best Multilinear Rank Approximation of Tensors 140
Christian Schröder
Numerical Solution of Palindromic Eigenvalue Problems 142
Jennifer A. Scott
The World of (HSL) Sparse Direct Solvers 143
Stefano Serra-Capizzano
Google PageRanking problem: the model and the analysis 144
Valeria Simoncini
Approximating functions of large matrices: computational aspects and applications 145
Alicja Smoktunowicz
Block Classical Gram–Schmidt Algorithm with Reorthogonalization 145
Danny C. Sorensen
Computation and Application of Balanced Model Order Reduction 146
Alastair Spence
Preconditioning the shift-invert transform in large sparse eigenvalue computations 147
Andreas Stathopoulos
An unrestarted, one-pass Lanczos with on-the-fly updating of eigenvectors . 148

G. W. Stewart
Analysis of the Residual Arnoldi Method 150
Zdeněk Strakoš
Sensitivity of Gauss-Christoffel quadrature, and sensitivity of the Jacobi matrix to small perturbations of the spectral data 151
Gilbert Strang
Compressed Sensing and Random Matrices 154
Brian D. Sutton
Computing the Complete CS Decomposition 154
Daniel B. Szyld
Application of Inexact Krylov Subspace Methods to Control Problems Governed by PDEs 156
Jurjen Duintjer Tebbens
An Updated Preconditioner for Sequences of General Nonsymmetric Linear Systems 156
Petr Tichý
On a New Proof of the Faber-Manteuffel Theorem 158
Françoise Tisseur
How to Detect and Solve Hyperbolic Quadratic Eigenvalue Problems 160
David Titley-Peloquin
Stopping Criteria for LSQR 162
Lloyd N. Trefethen
Continuous analogues of QR, LU, SVD and Schur 165
Joel A. Tropp
Algorithms for Matrix Column Selection 165
Miroslav Tůma
Balancing Incomplete Factorizations for Preconditioning 166
Eugene Tyrtyshnikov
Tensor rank estimates and tensor-based algorithms 168
Marc Van Barel
A Generalization of the $QR-$Algorithm for Semiseparable Plus Diagonal Matrices 169
Paul Van Dooren
A manifold approach for comparing two matrices 170
Martin van Gijzen
IDR(s): a family of simple and fast algorithms for solving large nonsymmetric systems of linear equations 172
Sabine Van Huffel
Tensor-Based Biosignal Processing 174
Charles Van Loan
Multilinear Algebra Computations in Quantum Chemistry 175
Raf Vandebril
A Rational Generalization of the QR-algorithm 175
Andras Varga
Numerical Methods for Continuous-time Periodic Systems 175
Stephen A. Vavasis
Algorithms and Complexity for Nonnegative Matrix Factorization 178

Elena Virnik
Positive Descriptor Systems 179
Kees Vuik
On complex shifted Laplace preconditioners for the vector Helmholtz equation 180
Justin W.L. Wan
Iterative Methods for Solving Partial Integro-Differential Equations Arising from Option Pricing 182
Wolfgang Wülling
On Clustered Ritz Values in the Lanczos Method 183
Hongguo Xu
Perturbation of Purely Imaginary Eigenvalues of Hamiltonian Matrices under Structured Perturbations 184
Krystyna Ziętak
On Some Known and Open Matrix Nearness Problems 184
List of participants with addresses 187

Differential-Geometric Foundations of Jacobi-Davidson Methods

P.-A. Absil, Michiel E. Hochstenbach

The most popular field of application of differential geometry is without doubt relativity theory: spacetime is a manifold equipped with a pseudo-Riemannian metric related to the stress-energy tensor through the Einstein field equations, and inertial motion occurs along geodesics induced by the pseudo-Riemannian metric.

More surprisingly perhaps, techniques of differential and Riemannian geometry have found applications in several areas of science and engineering, such as crystallography, thermodynamics and information theory.

Differential-geometric techniques have also been utilized to design and analyze eigenvalue algorithms since the work of Shub in 1986, if not earlier. That the Jacobi-Davidson approach, as well as a block version thereof, is closely related to Newton's method on Grassmann manifolds, was pointed out in [3].

The area of numerical computations on manifolds has made important progress since [3]. Several publications have parted with the "Riemannian dogma" of systematically making use of Riemannian geodesics. *Retractions* and *locally smooth families of parameterizations* provide frameworks to relax the Riemannian exponential into a broader class of mappings that offer opportunities to reduce the numerical burden while preserving convergence properties. Likewise, the Newton method on manifolds stated and analyzed in [2] allows for using *any* affine connection, instead of only the Riemannian one. Moreover, concerning the Grassmann manifold $\mathrm{Grass}(p, n)$ of p-planes in $\mathbf{C}^n$, results in [1] make it possible to use arbitrary bases to represent p-dimensional subspaces, whereas previous work required orthonormal bases or projectors.

As a consequence, *the* Newton method on Riemannian manifolds, as stated in [4], has now evolved into a large class of geometric Newton methods, which vary according to the choice of an affine connection and a retraction while retaining the local quadratic convergence property.

After an introduction to the modern formulation of Newton methods on manifolds, this talk shows how, when applied to a well chosen vector field, these Newton methods produce well-known forms of the Jacobi equation for the generalized eigenvalue problem. An advantage of this unification under the framework of manifold-based Newton methods is that the superlinear convergence of these Jacobi-Davidson methods follows directly from the convergence theory developed at the level of abstract manifolds.

More specifically, consider the quotient space

$$\mathbf{C}_*^{n\times p}/\mathrm{GL}_p := \{[Y] : Y \in \mathbf{C}_*^{n\times p}\},$$

where $\mathbf{C}_*^{n\times p}$ denotes the set of all full-column-rank $n \times p$ matrices, GL_p denotes the set of all invertible $p \times p$ matrices and

$$[Y] := Y\,\mathrm{GL}_p := \{YM : M \in \mathrm{GL}_p\}$$

is the equivalence class of all $n \times p$ matrices that have the same column space as Y. The set $\mathbf{C}_*^{n\times p}/\mathrm{GL}_p$ is in one to one correspondence with the set $\mathrm{Grass}(p, n)$ of all p-dimensional subspaces in $\mathbf{C}^n$; the correspondence is given by $Y \mapsto \mathsf{col}(Y)$, where $\mathsf{col}(Y)$ denotes the *column space* of Y, i.e., the subspace of $\mathbf{C}^n$ spanned by the columns of Y. Endowed with its unique "natural" manifold structure, $\mathrm{Grass}(p, n)$ is called a *Grassmann manifold*.

To define a horizontal distribution $\mathcal{H}$, pick a function $U \in \mathbf{C}_*^{n\times p} \mapsto \tilde{U} \in \mathbf{C}_*^{n\times p}$ such that $U^H\tilde{U}$ is invertible and $\mathsf{col}(\tilde{U})$ is constant on each equivalence class, and let

$$\mathcal{H}_U := \{Z \in \mathbf{C}^{n\times p} : \tilde{U}^H Z = 0\}. \tag{1}$$

The purpose of the horizontal distribution is to provide uniquely defined representations of tangent vectors to $\mathrm{Grass}(p,n)$. Indeed, given $\mathcal{U} = \mathsf{col}(U) \in \mathrm{Grass}(p,n)$ and a tangent vector $\xi_{\mathcal{U}} \in T_{\mathcal{U}}\mathrm{Grass}(p,n)$, there exists a unique $\overline{\xi}_U \in \mathcal{H}_U$, called the *horizontal lift* of $\xi_{\mathcal{U}}$ at U, such that $\mathrm{D}\,\mathsf{col}(U)[\overline{\xi}_U] = \xi_U$. Note that a horizontal vector field $\overline{\xi}$ defines a horizontal lift if and only if it satisfies the compatibility relation

$$\overline{\xi}_{UM} = \overline{\xi}_U M \tag{2}$$

for all $M \in \mathrm{GL}_p$.

Let A and B be arbitrary $n \times n$ matrices with B invertible. Consider the vector field ξ on $\mathrm{Grass}(p,n)$ defined by the horizontal lift

$$\overline{\xi}_U = P_{BU,\tilde{U}} AU, \tag{3}$$

where

$$P_{U,V} := I - U(V^HU)^{-1}V^H$$

denotes the projector along the columns space of U into the orthogonal complement of the column space of V. Note that $\tilde{U}^HBU$ is required to be invertible for the projector to be well defined.

A key observation is that $\overline{\xi}_U = 0$ if and only if $\mathsf{col}(U)$ is an invariant subspace of $B^{-1}A$. Hence, searching for a p-dimensional invariant subspace of $B^{-1}A$ amounts to searching for a zero of the vector field ξ on $\mathrm{Grass}(p,n)$ defined by the horizontal lift (3).

The Newton equation for computing a zero of a vector field ξ on an abstract manifold $\mathcal{M}$ equipped with an affine connection ∇ is

$$\nabla_{\eta_x}\xi = -\xi_x, \tag{4}$$

$x \in \mathcal{M}$, where ∇ is an affine connection on $\mathcal{M}$. We select an affine connection on $\mathrm{Grass}(p,n)$ as follows. Consider a function $U \in \mathbf{C}_*^{n\times p} \mapsto \tilde{W} \in \mathbf{C}_*^{n\times p}$ such that $\mathsf{col}(\tilde{W})$ is constant on each equivalence class $[U]$ and such that $\tilde{U}^H\tilde{W}$ is invertible for all $U \in \mathbf{C}_*^{n\times p}$. We choose the affine connection on $\mathrm{Grass}(p,n)$ defined by

$$\overline{(\nabla_{\eta_{\mathsf{col}(U)}}\xi)}_U = P_{\tilde{W},\tilde{U}}\mathrm{D}\overline{\xi}(U)[\overline{\eta}_U], \tag{5}$$

where ξ is a vector field on $\mathrm{Grass}(p,n)$ and $\eta_{\mathsf{col}(U)}$ is a tangent vector to $\mathrm{Grass}(p,n)$ at $\mathsf{col}(U)$. When $\mathcal{M}$ is the Grassmann manifold $\mathrm{Grass}(p,n) = \mathbf{C}_*^{n\times p}/\mathrm{GL}_p$, with the choice (1) for the horizontal space, the choice (5) for the affine connection, and the choice (3) for the vector field ξ, then the left-hand side of the Newton equation (4) becomes

$$\begin{aligned} & P_{\tilde{W},\tilde{U}}\mathrm{D}\overline{\xi}(U)[Z] \\ = \;& P_{\tilde{W},\tilde{U}}\left(P_{BU,\tilde{U}}AZ - BZ(\tilde{U}^HBU)^{-1}\tilde{U}^HAU - BU(...)\right), \end{aligned}$$

where Z stands for $\overline{\eta}_U$ to simplify the notation. Finally, we choose

$$\tilde{W} := BU \tag{6}$$

to eliminate the $BU(...)$ term, and the Newton equation (4) becomes

$$P_{BU,\tilde{U}}(AZ - BZ(\tilde{U}^H BU)^{-1}\tilde{U}^H AU) = P_{BU,\tilde{U}}AU, \quad \tilde{U}^H Z = 0. \tag{7}$$

It follows from the convergence theory of the geometric Newton method that the iteration on $\mathrm{Grass}(p,n)$ defined by

$$\mathsf{col}(U) \mapsto \mathsf{col}(U+Z), \tag{8}$$

where Z solves (7), converges locally, at least quadratically, to the spectral invariant spaces of $B^{-1}A$.

References

[1] P.-A. Absil, R. Mahony, and R. Sepulchre. Riemannian geometry of Grassmann manifolds with a view on algorithmic computation. *Acta Appl. Math.*, 80(2):199–220, January 2004.

[2] P.-A. Absil, R. Mahony, and R. Sepulchre. *Optimization Algorithms on Matrix Manifolds.* Princeton University Press, Princeton, NJ, January 2008.

[3] Eva Lundström and Lars Eldén. Adaptive eigenvalue computations using Newton's method on the Grassmann manifold. *SIAM J. Matrix Anal. Appl.*, 23(3):819–839, 2001/02.

[4] Steven T. Smith. Optimization techniques on Riemannian manifolds. In *Hamiltonian and gradient flows, algorithms and control*, volume 3 of *Fields Inst. Commun.*, pages 113–136. Amer. Math. Soc., Providence, RI, 1994.

Model reduction by means of fast frequency sweep

A.C. Antoulas

In this note we propose the use of *dominant system poles* as the basis for a new model reduction algorithm which makes use of a fast frequency sweep (FFS) method. The goal of this procedure is to compute a reduced order model by evaluating the transfer function of the original high-order system, in as few frequency points as possible.

1 Problem formulation

Starting point is a linear dynamical system described by a set of high order differential equations. It is convenient here to describe the system in the frequency domain as follows:

$$K(s)X(s) = BU(s),\ Y(s) = CX(s), \quad K(s) := \left[K_p s^p + K_{p-1}s^{p-1} + \cdots + K_1 s + K_0\right],$$

where s is the complex frequency, $U(s)$, $X(s)$, $Y(s)$ – the Laplace transforms of $u(t)$, $x(t)$, $y(t)$ – are the input, an internal variable, and the output, of the system. The internal variable is of dimension n, while the input and the output are assumed here to be scalar; K_i, B, C are constant matrices of appropriate size. The transfer function of such a system is $H(s) = CK^{-1}(s)B$. An important special case is that of first order systems where $K(s) = Es - A$,

with E, A, constant $n \times n$ matrices. For the applications of interest the dimension n is high (in the thousands and above). Such cases arise typically when systems described by partial differential equations are spatially discretized, for instance in electromagnetic and circuit simulations, as well as in many other areas [1].

The model reduction problem consists in computing approximants of the form

$$\hat{K}(s)\hat{X}(s) = \hat{B}U(s), \quad \hat{Y}(s) = \hat{C}\hat{X}(s), \quad \hat{K}(s) := \left[\hat{K}_p s^p + \hat{K}_{p-1}s^{p-1} + \cdots + \hat{K}_1 s + \hat{K}_0\right].$$

The input $U(s)$ and the output $\hat{Y}(s)$ are still scalars while the new internal variable $\hat{X}(s)$ has reduced dimension k, where typically $k \ll n$. The main requirement is that the behavior of the reduced system be close to that of the original in some appropriate sense. Quite often the error consists of some norm of the difference $H - \hat{H}$ between the two transfer functions (e.g. the 2-norm on the imaginary axis). For an overview of model reduction methods and error norms, see [1].

Two of the usual families of model reduction methods are the SVD based methods (e.g. balanced truncation) and the Krylov based methods (e.g. rational Lanczos).

In this note we will be concerned with a different approach to model reduction sometimes referred to (mostly in the electromagnetics simulation area) as fast frequency sweep (FFS). It consists in using measurements of the transfer function $H(i\omega_j)$ at certain frequencies ω_j on the imaginary axis, in order to construct approximants. In this method, since the evaluations of the transfer function are costly (they involve the inversion of $K(i\omega_j)$) we are looking for a judicious choice of a small number of frequencies ω_j, which will yield a good quality approximant.

2 The FFS model reduction method

• One ingredient of the new approach is the concept of **dominant system poles**. Given a rational function $H(s)$ (assumed scalar), let $s_0 \in C$ be a (simple) pole, i.e. the expression $(s - s_0)H(s)$ has no singularity at $s = s_0$. The *residue* of this pole is defined as $\rho_0 := (s - s_0)H(s)|_{s=s_0}$. There are various ways of defining *dominance* of system poles. One is to use as measure the proximity to the imaginary axis, that is $\frac{1}{|Re(s_0)|}$. An often better measure is the following. Given a pole s_i and the corresponding residue ρ_i, we define

$$\mu_i := \frac{|\rho_i|}{|Re(s_i)|}.$$

The pole s_i is called *dominant* with respect to the pole s_j, if $\mu_i > \mu_j$. Notice that for poles s_i, s_j which are not too close to each other, μ_i, μ_j, give the (local) maxima of the amplitude of the frequency response, which are attained at the frequencies $\omega_i = Im(s_i)$, $\omega_j = Im\,(s_j)$, respectively.

• The second main ingredient is the **data-based model reduction method** which was developed in [2]. Therein, using the Loewner and the shifted Loewner matrices as main tools reduced order models in terms of (E, A, B, C) quadruples are readily determined from the data. Constraints like stability and passivity of the approximant can be imposed in a natural way.

• **The proposed FFS model reduction method** consists of the following steps:

1. Compute some of the most dominant system poles $s_j = \sigma_j + i\omega_j$.

2. Evaluate the transfer function of the system at the frequencies ω_j, that is at the imaginary parts of the computed dominant poles.

3. Finally, compute a low order model of the underlying system using the Loewner matrix based, data-to-reduced-model method developed in [2],

• **Computational issues**. For small first-order systems described by means of a quadruple (E, A, B, C) and the transfer function $H(s) = C(sE - A)^{-1}B$, it is feasible to compute all poles s_j of a given system as eigenvalues of the pencil (A, E), as well as the corresponding dominance indices μ_i. Then after ordering, some dominant poles s_j are chosen and the value of the transfer function $H(i\omega_j)$, is determined.

The major issue for the applicability of the proposed method to large-scale systems is the determination of a few dominant poles of (A, E), without the need to compute all the poles first. This can be achieved using the method presented in [3]. This is an iterative algorithm called SADPA (subspace accelerated dominant pole algorithm).

• **Higher-order systems and nonlinear eigenvalue problems**. If the system under consideration is described by means of higher order equations of the form $K(s)X(s) = BU(s)$, $Y(s) = CX(s)$, where $K(s)$ is an $n \times n$ polynomial matrix (with n large) the method of dominant poles leads to *nonlinear eigenvalue problems*. These have been studied in detail (see for instance [4], [5], [6]). They are often based on some form of linearization, involving the determination of a pencil (A, E) whose eigenvalues are the same as those of $K(s)$.

• **Problems and goals**. **(a)** Our first goal is a complete analysis of this new measument-based (FFS) model reduction method together with comparisons with methods like balanced truncation, rational Krylov, etc. **(b)** We are currently investigating the extension of the *subspace accelerated dominant pole algorithm* (SADPA) of [3] from first order $K(s) = sE - A$, to the case of higher order polynomial matrices $K(s)$, thereby taking into account any structure (symmetries) which may be present. Our second goal is to present efficient computational algorithms for the determination of dominant poles of polynomial matrices, together with applications to large scale electromagnetic problems.

References

[1] A.C. Antoulas, *Approximation of large-scale dynamical systems*, Advances in Design and Control, **DC-06**, SIAM, Philadephia (2005).

[2] A.J Mayo and A.C. Antoulas, *A framework for the solution of the generalized realization problem*, Linear Algebra and Its Applications, **425**: 634-662 (2007).

[3] J. Rommes and N. Martins, *Efficient computation of transfer function dominant poles using subspace acceleration*, IEEE Trans. on Power Systems, **21**: 1218-1226 (2006).

[4] I.C.F. Ipsen, *Accurate Eigenvalues for Fast Trains*, SIAM News, Volume 37, Number 9, November 2004.

[5] D.S. Mackey, N. Mackey, C. Mehl, and V. Mehrmann, *Vector spaces of linearizations for matrix polynomials*, SIAM J. Matrix Anal. Appl., 28: 971-1004, 2006.

[6] D.S. Mackey, N. Mackey, C. Mehl, and V. Mehrmann, *Structured polynomial eigenvalue problems: good vibrations from good linearizations*, SIAM J. Matrix Anal. Appl., 28: 1029-1051, 2006.

Matrix square root and interpolation spaces

Mario Arioli, Daniel Loghin

Fractional powers of a matrix

$$A^\theta, \quad \theta \in (-1, 1)$$

are related to the approximation of the scalar product of the order θ Hilbert interpolation space

$$[\mathcal{H}, \mathcal{M}]_\theta$$

between the two Hilbert spaces $\mathcal{H}$ and $\mathcal{M}$ with $\mathcal{H} \subset \mathcal{M}$ and $\mathcal{H}$ dense in $\mathcal{M}$(see [2] for the theory of the interpolation spaces of Sobolev spaces). The interpolation spaces play a key role in several computational fields ranging from Domain Decomposition (DD) [1, 4, 5], solution of biharmonic equations [3], and image denoising.

Let $\mathcal{H}_N \subset \mathcal{H}$ and $\mathcal{M}_N \subset \mathcal{M}$ be two two finite N-dimensional Hilbert spaces such that the projection operators π_1 and π_2, and the immersion operators i_1 and i_2

$$\begin{cases} \pi_1 : \mathcal{H} \to \mathcal{H}_N \\ i_1 : \mathcal{H}_N \to \mathcal{H} \\ \pi_2 : \mathcal{M} \to \mathcal{M}_N \\ i_2 : \mathcal{M}_N \to \mathcal{M} \end{cases}$$

are linear and continuous, i.e. $\pi_1 \in \mathcal{L}(\mathcal{H}, \mathcal{H}_N)$, $i_1 \in \mathcal{L}(\mathcal{H}, \mathcal{L}_N)$, $\pi_2 \in \mathcal{L}(\mathcal{M}, \mathcal{M}_N)$, and $i_2 \in \mathcal{L}(\mathcal{M}, \mathcal{M}_N)$. Conforming finite-element approximation of classical Sobolev spaces are included in this framework.

$\mathcal{H}_N$ and $\mathcal{M}_N$ are isometric to $\mathbb{R}^N$ using the scalar products defined respectively by the symmetric and positive matrices H and M, i.e.

$$\begin{array}{llll} \forall u, v \in \mathcal{H}_N & \exists \mathbf{u}, \mathbf{v} \in \mathbb{R}^N & such \;\; that & \langle u, v \rangle_{\mathcal{H}_N} = \mathbf{u}^T H \mathbf{v} \\ \forall w, z \in \mathcal{M}_N & \exists \mathbf{w}, \mathbf{z} \in \mathbb{R}^N & such \;\; that & \langle w, z \rangle_{\mathcal{M}_N} = \mathbf{w}^T M \mathbf{z} \end{array}$$

Following the definition of interpolation space given in [2], we introduce the space

$$[\mathcal{H}_N, \mathcal{M}_N]_\theta$$

which is also isometric to $\mathbb{R}^N$ with the scalar product defined by the matrix

$$H_\theta = M + M(M^{-1}H)^\theta.$$

Despite its unsymmetric look, we prove that H_θ is symmetric and positive definite by the properties of M, H, and of the generalized eigenvalue problem

$$H\mathbf{u} = \lambda M \mathbf{u}.$$

The computation of H_θ by a vector or of H_θ^{-1} by a vector are the elementary operations needed in order to precondition several kinds of operators (the Steklov-Poincaré operator in DD is a classical example [4, 5]). $H_\theta \mathbf{y}$ and $H_\theta^{-1}\mathbf{z}$ can be computed by the generalized Lanczos algorithm followed by the computation of the θ-power of the tridiagonal matrix. Moreover, the products $H_\theta \mathbf{y}$ and $H_\theta^{-1}\mathbf{z}$ need only a good approximation of the θ-power function in an interval reasonably distant from zero.

We will illustrate the usefullness of this approach using examples coming from the solution of elliptic variational problems by DD methods and from the solution of biharmonic problems.

References

[1] P. E. Bjorstad and Olof B. Widlund. Iterative methods for the solution of elliptic problems on regions partitioned into substructures. *SIAM J. Numer. Anal.*, 23:1097–1120, 1986.

[2] J. L. Lions and E. Magenes. Non-homogeneous boundary values problems and applications. I Springer, Berlin-Heidelberg-New York, 1972.

[3] P. Peisker Numerical solution of the first biharmonic equations. *Mathematical Modelling and Numerical Analysis*, 22:655–676, 1988.

[4] A. Quarteroni and A. Valli Domain Decomposition Methods for Partial Differential Equations. Oxford University Press, Oxford, UK 1999.

[5] A. Toselli and O. Widlund Domain Decomposition Methods. Springer Verlag, Berlin-Heidelberg, 2005.

An additive subspace correction method combined with an ADI-type splitting

Owe Axelsson, Maya Neytcheva

The additive subspace method in a multilevel setting requires an efficient solution method for the block diagonal matrices which arises in a two-and multilevel context and corresponds to the fine-mesh nodes. A new such method is described here, which is based on an ADI-type splitting of the matrix.

Robust and Efficient Methods for Multi-Length Scale Matrix Computations and Applications in Quantum Mechanical Simulations

Zhaojun Bai, Wenbin Chen, Roger Lee, Richard Scalettar and Ichitaro Yamazaki

Forefront problems of scientific and engineering computing often require solutions involving matrix computations. Great progress has been made both in general procedures and also in more focused situations which exploit specific matrix structure and sparsity patterns. A generation of excellent algorithms has been developed and presented in textbooks and research monographs. Vast amounts of public and proprietary software libraries, packages and templates have been made widely accessible, such as LAPACK and ScaLAPACK among many others. The impact of progress in large scale matrix computations to the advancement of computational science and engineering cannot be overstated.

However, challenges of the advancement of multi-physics and multi-scale mathematical modeling and simulations to numerical linear algebra theory and algorithms have not been widely addressed by the research community. Most solvers are not designed in ways that are

robust, efficient and self-adapting for underlying multi-scale simulations. This is an emerging research area. A broad range of scientific and engineering modeling and simulation problems involve multiple scales for which traditional monoscale approaches have proven to be inadequate, even with the largest supercomputers.

In this talk, we describe a range of our synergistic activities on development of robust and efficient linear algebra solvers which are specially designed for multi-length scale matrix computation problems arising from quantum mechanical simulations of materials. This includes hybrid design and use of existing and new algorithms and development of high-performance software. We report our progress on (a) the study of the dynamics of conditioning, eigenvalue distributions and other critical properties of multi-length scale matrices, (b) the design and analysis of robust and efficient self-adapting linear algebra solvers, and (c) the development of high performance software solving real multi-scale physical phenomena. We will demonstrate the feasibility, through the use of newly developed linear algebra solvers, of 1000-electron quantum monte carlo simulations on modern desktop machines and extreme large-scale simulations on supercomputers. Such simulations would allow physicists and chemists to address important questions concerning the magnetic and transport properties of materials with strong electron-electron interactions.

The results of such robust and efficient linear algebra solvers have more general impacts beyond the application discussed above. A host of other applications awaits the development of multiscale linear algebra solvers. As one example, the methodology described has close connections to problems in lattice gauge theory, which shares many of the matrix structures and algorithms with condensed matter QMC simulations. Dynamical mean field theory and and Hatano-Nelson localization are other examples.

The first phase of our research efforts started more than three years ago, through an NSF ITR grant. The current second phase is supported by DOE SciDAC (Scientific Discovery through Advanced Computing) program for the project "Modeling Materials at the Petascale".

Related publication: Z. Bai, W. Chen, R. Scalettar and I. Yamazaki. *Lecture Notes on Advances of Numerical Methods for Quantum Monte Carlo Simulations of the Hubbard Model.* 100+ pages, submitted for publication, 2007

Control-oriented Model Reduction for Parabolic Systems

Ulrike Baur

We consider linear time-invariant (LTI) stable systems of the following form

$$\Sigma : \left\{ \begin{array}{rcll} \delta x(t) & = & Ax(t) + Bu(t), & t > 0, \qquad x(0) = x^0, \\ y(t) & = & Cx(t) + Du(t), & t \geq 0, \end{array} \right.$$

with $A \in \mathbf{R}^{n \times n}$, $B \in \mathbf{R}^{n \times m}$, and $C \in \mathbf{R}^{p \times n}$ arising, e.g., from the discretization and linearization of parabolic PDEs. If $t \in \mathbf{R}$ and $\delta x(t) = \dot{x}(t)$, the system is a continuous-time system; it describes a discrete-time system if $t \in \mathbf{N}_0$ and $\delta x(t) = x(t+1)$ is the forward shift operator. We will assume that the system Σ is large-scale, having order $n = \mathcal{O}(10^5)$ with $n \gg m$, p. We allow the system matrix A to be dense, provided that a *data-sparse* representation [1] exists.

In practical applications, e.g. in VLSI chip design and in control of PDEs, the associated large-scale LTI systems must be solved very often. A reduction of the state-space dimension

is absolutely necessary to attack these problems in reasonable time. On the other hand, in modern (LQG-, $\mathcal{H}_2$-, $\mathcal{H}_\infty$-) feedback control the controller has at least the dimension of the underlying parabolic control system. Therefore, real-time control is only possible for systems with reduced number of states resulting in controllers of lower complexity.

To preserve important system properties as stability and to obtain a global error bound, we focus on model order reduction methods based on balanced truncation (BT), the most commonly used technique in control theory. The major computational task in BT and related methods is the solution of two large-scale matrix equations, the so-called Lyapunov equations,

$$A\mathcal{P} + \mathcal{P}A^T + BB^T = 0, \quad A^T\mathcal{Q} + \mathcal{Q}A + C^TC = 0,$$

for continuous-time systems and of two Stein equations

$$A\mathcal{P}A^T - \mathcal{P} + BB^T = 0, \quad A^T\mathcal{Q}A - \mathcal{Q} + C^TC = 0$$

for systems in discrete time. A variant of the classical BT method, the cross-Gramian approach, is based on the solution of one Sylvester equation. During the last 17 years a lot of methods were derived for the solution of large-scale matrix equations. But we want to point out that all methods are of cubic complexity when applied to dense problems.

We have developed a collection of solvers for different types of matrix equations which break this curse of dimensionality for the practically relevant class of data-sparse systems [2, 3, 4, 5]. These iterative solvers are based on the *sign function method* for continuous-time systems and on the *squared Smith method* in discrete time. By using data-sparse matrix approximations, hierarchical ($\mathcal{H}$) matrix formats, and the corresponding formatted arithmetic we obtain efficient implementations of BT and related methods having all linear-polylogarithmic complexity.

We will demonstrate the accuracy and applicability of the derived methods on large-scale systems coming from 2D and 3D FEM and BEM discretizations. For example, a comparison of the accuracy and efficiency of the new sign function iteration with an implementation in standard (dense) arithmetic is shown in Table 1. The solvers are applied to a dense system arising from a boundary element discretization of the Laplace equation in $\Omega \subset \mathbf{R}^3$. The approximate BT method applied to the largest system (with $n = 32,768$) computes a very small reduced-order model ($r = 10$) for an absolute error tolerance of 10^{-4}. This becomes only possible through the combination of hierarchical-matrix approximation and model reduction.

n	# it.	rank($\mathcal{P}$)		storage in MB		time [sec]		rel. residual		rel. error
		$\mathcal{H}$	full	$\mathcal{H}$	full	$\mathcal{H}$	full	$\mathcal{H}$	full	
2048	11	10	10	34	64	228	282	9.4e-07	9.4e-07	2.1e-06
8192	11	11	11	220	1025	1769	17,071	4.6e-07	4.6e-07	2.6e-06
32,768	12	12	-	1371	* 16,404	15,306	-	7.7e-07	-	-

Table 1: Accuracy, complexity and rank($\mathcal{P}$) of the solution factors computed with sign function method in formatted ($\mathcal{H}$) and in dense arithmetic (full) for different problem sizes n (* is extrapolated).

References

[1] W. Hackbusch, *A sparse matrix arithmetic based on $\mathcal{H}$-matrices. I. Introduction to $\mathcal{H}$-matrices*, Computing, 62 (1999), pp. 89–108.

[2] U. Baur, *Low Rank Solution of Data-Sparse Sylvester Equations*, Preprint #266, MATHEON, DFG Research Center "Mathematics for Key Technologies", Berlin, FRG, http://www.math.tu-berlin.de/DFG-Forschungszentrum, Oct. 2005. To appear in Numer. Lin. Alg. Appl.

[3] U. Baur and P. Benner, *Factorized solution of Lyapunov equations based on hierarchical matrix arithmetic*, Computing, 78 (2006), pp. 211–234.

[4] U. Baur and P. Benner, *Gramian-Based Model Reduction for Data-Sparse Systems*, Preprint CSC 07-01, Fakultät für Mathematik, TU Chemnitz, Chemnitz, FRG, http://www-user.tu-chemnitz.de/∼ubaur/pub/Hmodelreduction.pdf, Jan. 2007.

[5] U. Baur, *Control-oriented model reduction for parabolic systems*, PhD Thesis, in preparation, 2007.

Solving Algebraic Riccati Equations for Stabilization of Incompressible Flows

Peter Benner

We will discuss the numerical solution of algebraic Riccati equations (AREs) arising in the exponential stabilization of flow problems. This stabilization is to be achieved by action of boundary control using feedback mechanisms. In very recent work by Raymond [11] and earlier attempts by Barbu, Lasiecka, and Triggiani [2, 3], it is shown analytically that it is possible to construct a linear-quadratic optimal control problem associated to the Navier-Stokes equation so that the resulting feedback law, applied to the instationary Navier-Stokes equation, exponentially stabilizes unstable solution trajectories under reasonable assumptions.

The novelty of the approach by Raymond [11] is that for the first time a linear feedback solution for this problem is provided that allows localized Dirichlet boundary control in normal directions which is of major importance in practical applications. (We will demonstrate this using the well-known benchmark problem of the backward facing step.) Moreover, Raymond's proof of existence of an exponentially stabilizing linear feedback law is constructive: it provides a method for the computation of Dirichlet boundary controls which stabilize an incompressible flow described by the Navier-Stokes equations on a connected and bounded domain Ω in the sense that its solution is controlled to a desired steady-state solution. In our talk, we will briefly outline this construction for the 2D case—the 3D case is treated similarly, but the derivation is quite more involved. The computation of the exponentially stabilizing feedback law is based on a modified Oseen approximation to the difference between the actual and desired flows. The feedback law then takes the form $u(t) = -R_A^{-1} b B^* \Pi P z(t)$, where $z(t)$ is the current difference between the actual and the desired flow, A is the Oseen operator, $P : L_2(\Omega) \mapsto V_n^0(\Omega)$ is the Helmholtz projection[1], $B := (\lambda_0 I - A) D_A$ is the control operator with D_A denoting the Dirichlet operator associated with $\lambda_0 I - A$ ($\lambda_0 > 0$), and $R_A := b D_A^* (I - P) D_A b + I$. The Riccati operator $\Pi = \Pi^* \in \mathcal{L}(V_n^0(\Omega))$ is the unique nonnegative semidefinite weak solution of the operator Riccati equation

$$0 = I + (A + \omega I)^* \Pi + \Pi (A + \omega I) - \Pi (B_\tau B_\tau^* + B_n R_A^{-1} B_n^*) \Pi, \tag{1}$$

[1] $V_n^0(\Omega)$ is the set of divergence-free $L_2(\Omega)$ functions with zero normal component.

where B_τ and B_n correspond to the projection of the control action in the tangential and normal directions derived from the control operator B.

Thus, the numerical solution of the stabilization problem for the Navier-Stokes equation requires the numerical solution of the operator equation (1). Our solution strategy will be based on a finite-element discretization of the Oseen equation using Taylor-Hood elements. This results in a large-scale ARE which is used to compute an approximating solution to (1). For sufficient accuracy, the dimension of the FEM ansatz space will in general be too large to solve the ARE with classical methods based on associated eigenvalue problems.

For linear-quadratic optimal control problems related to linear parabolic and hyperbolic partial differential equations, the corresponding large-scale AREs are usually solved with some variant of the Newton-Kleinman iteration (see, e.g., [1, 10, 4]) or multigrid methods [12, 7]. The latter variants so far have shown their efficiency only for the heat equation on simple domains, therefore we concentrate here on the Newton-Kleinman method. Unfortunately, this solution strategy will in general fail here due to the following reasons:

1. In each step of the Newton-Kleinman iteration, a Lyapunov equation of the form

$$A_k^T X_k + X_k A_k + C_k C_k^T = 0$$

has to be solved for the unknown X_k. This can be done very efficiently if C_k has few columns, using, e.g., the low-rank Cholesky factor ADI iteration [9, 10]. In our situation, after spatial discretization of the modified Oseen equation, the constant term I in (1) will become the mass matrix M_h of the finite element basis so that $C_k C_k^T = M_h + G_k G_k^T$ for some matrix G_k depending on the current Newton iterate and the control operator. Thus, C_k will have rank n and the cost for the Newton iteration will be at least $\mathcal{O}(n^2)$ per step which is way too high for practical applications.

2. For linear parabolic PDEs, the Newton-Kleinman iteration can be started from $X_0 = 0$ as the discretization of the elliptic operator of spatial derivatives will lead to a stable stiffness matrix. In this situation, it is known that the Newton-Kleinman iteration converges quadratically to the desired solution. Here, $A + \omega I$ will in general lead to an unstable stiffness matrix. Thus, to start the iteration, a stabilizing initial guess needs to be computed first. So far, no efficient method to achieve this for large-scale systems is known.

3. In order to discretize the Oseen operator in such a way that a standard ARE results, it would be necessary to use divergence-free finite elements. As there are no feasible ansatz spaces with this property, we would have to project the FEM basis functions explicitly onto $V_n^0(\Omega)$, requiring each time the application of the Helmholtz projector. This will be too expensive in general.

In our talk we will discuss strategies to overcome the above difficulties. We will present a multilevel idea that may be used to solve both 1. and 2. when combined with a variant of Newton-Kleinman also employed in [1, 5]. Another approach to 2. is to use the solution of a homogeneous ARE, also called algebraic Bernoulli equation, as initial guess X_0. We will discuss possibilities to obtain this solution efficiently for large-scale systems, some of them related to [6]. In order to avoid the explicit computation of the Helmholtz projection, we will employ a strategy to solve Lyapunov equations associated with the Oseen equation proposed in [8]. That is, rather than discretizing $(A + \omega I)P$, we will use standard finite-elements to

discretize $A + \omega I$ and then obtain the Helmholtz projected states Px as part of the solutions of saddle point problems. Hence, in an ADI iteration for solving the Lyapunov equation in each Newton-Kleinman step, a saddle point problem corresponding to $A + \omega_k I$ will be solved from which we can recover the update of the current Newton iterate.

References

[1] H.T. Banks and K. Ito. A numerical algorithm for optimal feedback gains in high dimensional linear quadratic regulator problems. *SIAM J. Cont. Optim.*, 29(3):499–515, 1991.

[2] V. Barbu. Feedback stabilization of the Navier-Stokes equations. *ESAIM: Control Optim. Calc. Var.*, 9:197–206, 2003.

[3] V. Barbu, I. Lasiecka, and R. Triggiani. Tangential boundary stabilization of Navier-Stokes equations. *Memoirs of the AMS* 181, 2006.

[4] P. Benner. Efficient algorithms for large-scale quadratic matrix equations. *Proc. Appl. Math. Mech.*, 1(1):492–495, 2002.

[5] P. Benner, V. Hernández, and A. Pastor. On the Kleinman iteration for nonstabilizable systems. *Math. Control, Signals, Sys.*, 16:76–93, 2003.

[6] K. Gallivan, X. Rao, and P. Van Dooren. Singular Riccati equations stabilizing large-scale systems. *Linear Algebra Appl.*, 415(2–3): 359–372, 2006.

[7] L. Grasedyck. Nonlinear multigrid for the solution of large scale Riccati equations in low-rank and $\mathcal{H}$-matrix format. *Numer. Lin. Alg. Appl.*, to appear.

[8] M. Heinkenschloss, D.C. Sorensen, and K. Sun. Balanced Truncation Model Reduction for a Class of Descriptor Systems with Application to the Oseen Equations. Technical Report CAAM TR07-02, Rice University, `http://www.caam.rice.edu/caam/trs/2007.html`, January 2007.

[9] J.-R. Li and J. White. Low rank solution of Lyapunov equations. *SIAM J. Matrix Anal. Appl.*, 24(1):260–280, 2002.

[10] T. Penzl. *Numerische Lösung großer Lyapunov-Gleichungen.* Logos–Verlag, Berlin, Germany, 1998. (Dissertation, Fakultät für Mathematik, TU Chemnitz, 1998.)

[11] J.-P. Raymond. Local boundary feedback stabilization of the Navier-Stokes equations. In *Control Systems: Theory, Numerics and Applications, Rome, 30 March – 1 April 2005*, Proceedings of Science. SISSA, `http://pos.sissa.it`, 2005.

[12] I.G. Rosen and C. Wang. A multi–level technique for the approximate solution of operator Lyapunov and algebraic Riccati equations. *SIAM J. Numer. Anal.*, 32(2):514–541, 1995.

Localization Phenomena in Matrix Functions: Theory, Algorithms, Applications

Michele Benzi, Nader Razouk

Many physical phenomena are characterized by strong *localization*, that is, rapid decay outside of a small spatial or temporal region. Frequently, such localization can be expressed as decay in the entries of a function $f(A)$ of an appropriate sparse or banded matrix A that encodes a description of the system under study. Important examples include the decay in the entries of the density matrix for non-metallic systems in quantum chemistry (a function of the

Hamiltonian), the localization of the eigenvectors in the Anderson model, and the decay behavior of the inverse of a banded symmetric positive definite matrix; see, e.g., [1,7,9,12,17,21]. Localization also plays an important role in quantum information processing [6,22] and in signal processing [13].

Localization phenomenona are of fundamental importance both at the level of the physical theory and at the computational level, because they open up the possibility of approximating relevant matrix functions in $O(n)$ time, where n is a measure of the size of the system. With few exceptions (e.g., [2,3,18,20]), numerical analysts have not systematically exploited this property in the design of fast approximation algorithms.

We present theoretical results on decay rates for the entries of rather general functions of banded or sparse matrices, and we use this theory to guide the development of fast algorithms for computing approximations to matrix functions. We are motivated by the ongoing need for fast algorithms for approximating matrix functions that occur in important areas of computational science and engineering, like electronic structure calculations in quantum chemistry and computational material science. There is also a need for numerical analysts to provide a sound basis for existing $O(n)$ algorithms developed by physicists and to develop new, more efficient ones.

We make use of results from classical approximation theory in the complex plane (e.g., [10]) to obtain decay bounds for the entries of analytic functions of sparse diagonalizable matrices. Our results [5] generalize and unify several previous bounds known in the literature for special classes of functions and matrices, such as the inverse, the resolvent, and the exponential; see, e.g., [4,8,9,15,16,19,24]. The theory can be used to determine the bandwidth or sparsity pattern outside of which the entries of $f(A)$ are guaranteed to be so small that they can be neglected without exceeding a prescribed error tolerance. We discuss sufficient conditions for the possibility of $O(n)$ approximations to within a prescribed error in terms of the spectral properties of the matrices involved and of the possible singularities of f. We describe methods based on classical approximation theory, including Chebyshev [11], Faber [10], and Newton polynomials [23], and related approaches. Together with truncation strategies for neglecting small elements, these techniques form the basis for the approximation algorithms we use. These algorithms can also be used to estimate the trace and the determinant of certain functions of A (see, e.g., [2,3,18,20]), and to construct explicit preconditioners for $Ax = b$. Numerical experiments illustrating the promising behavior of the proposed approximation schemes will be presented.

References

[1] R. Baer and M. Head-Gordon, *Sparsity of the density matrix in Kohn–Sham density functional theory and an assessment of linear system-size scaling methods*, Phys. Rev. Lett., 79 (1997), pp. 3962–3965.

[2] Z. Bai and G. H. Golub, *Bounds for the trace of the inverse and the determinant of symmetric positive definite matrices*, Ann. Numer. Math., 4 (1997), pp. 29–38.

[3] C. Bekas, E. Kokiopoulou, and Y. Saad, *An estimator for the diagonal of a matrix*, Appl. Numer. Math., 57 (2007), pp. 1214–1229.

[4] M. Benzi and G. H. Golub, *Bounds for the entries of matrix functions with applications to preconditioning*, BIT, 39 (1999), pp. 417–438.

[5] M. Benzi and N. Razouk, *Decay bounds and $O(n)$ algorithms for approximating functions of sparse matrices*, Electronic Transactions on Numerical Analysis, 28 (2007), pp. 16–39.

[6] M. Cramer, J. Eisert, M. B. Plenio and J. Dreissig, *Entanglement-area law for general bosonic harmonic lattice systems*, Phys. Rev. A, 73 (2006), Art. No. 012309.

[7] D. R. Bowler and M. J. Gillan, *Density matrices in $O(N)$ electronic structure calculations: theory and applications*, Comp. Phys. Comm., 120 (1999), pp. 95–108.

[8] N. Del Buono, L. Lopez and R. Peluso, *Computation of the exponential of large sparse skew-symmetric matrices*, SIAM J. Sci. Comput., 27 (2005), pp. 278–293.

[9] S. Demko, W. F. Moss and P. W. Smith, *Decay rates for inverses of band matrices*, Math. Comp., 43 (1984), pp. 491–499.

[10] S. W. Ellacott, *Computation of Faber series with application to numerical polynomial approximation in the complex plane*, Math. Comp., 40 (1983), pp. 575–587.

[11] L. Fox and I. B. Parker, *Chebyshev Polynomials in Numerical Analysis*, Oxford University Press, Oxford, UK, 1968.

[12] S. Goedecker, *Linear scaling electronic structure methods*, Rev. Mod. Phys., 71 (1999), pp. 1085–1123.

[13] J. Gutiérrez-Gutiérrez, P. M. Crespo, and A. Böttcher, *Functions of banded Hermitian block Toeplitz matrices in signal processing*, Linear Algebra Appl., 422 (2007), pp. 788–807.

[14] N. J. Higham, *Functions of Matrices: Theory and Computation*, SIAM, Philadelphia, forthcoming.

[15] A. Iserles, *How large is the exponential of a banded matrix?*, New Zealand J. Math., 29 (2000), pp. 177–192.

[16] S. Jaffard, *Propriétés des matrices "bien localisées" près de leur diagonale et quelques applications*, Ann. Inst. H. Poincaré Anal. Non Linéaire, 7 (1990), pp. 461–476.

[17] W. Kohn, *Density functional and density matrix method scaling linearly with the number of atoms*, Phys. Rev. Lett., 76 (1996), pp. 3168–3171.

[18] D. J. Lee and I. C. F. Ipsen, *Zone determinant expansions for nuclear lattice simulations*, Phys. Rev. C, 68 (2003), pp. 064003-1/8.

[19] R. Nabben, *Decay rates of the inverse of nonsymmetric tridiagonal and band matrices*, SIAM J. Matrix Anal. Appl., 20 (1999), pp. 820–837.

[20] A. Reusken, *Approximation of the determinant of large sparse symmetric positive definite matrices*, SIAM J. Matrix Anal. Appl., 23 (2002), pp. 799–818.

[21] V. E. Sacksteder, *$O(N)$ algorithms for disordered systems*, Numer. Linear Algebra Appl., 12 (2005), pp. 827–838.

[22] N. Schuch, J. I. Cirac and M. M. Wolf, *Quantum states on harmonic lattices*, Comm. Math. Phys., 267 (2006), pp. 65–92.

[23] H. Tal-Ezer, *Polynomial approximation of functions of matrices and applications*, J. Scientific Comput., 4 (1989), pp. 25–60.

[24] L. N. Trefethen, M. Contedini and M. Embree, *Spectra, pseudospectra, and localization for random bidiagonal matrices*, Comm. Pure Appl. Math., 54 (2001), pp. 595–623.

Improving computed eigenpairs of polynomial eigenvalue problems by scaling and iterative refinement

Timo Betcke

We consider the polynomial eigenvalue problem (PEP) of the form

$$P(\lambda)x := (\lambda^\ell A_\ell + \lambda^{\ell-1}A_{\ell-1} + \cdots + A_0)x = 0, \quad A_k \in \mathbb{C}^{n\times n},\ A_\ell \neq 0,\ \ell \geq 1. \tag{1}$$

Every $\lambda \in \mathbb{C}$ for which there exists a solution $x \in \mathbb{C}^n\backslash\{0\}$ of $P(\lambda)x = 0$ is called an eigenvalue of P with associated right eigenvector x. Correspondingly, a left eigenvector $y \in \mathbb{C}^n\backslash\{0\}$ is defined by $y^*P(\lambda)x = 0$.

In recent years there has been a growing interest in the solution of PEPs, with applications in diverse areas such as structural mechanics, molecular dynamics or optical waveguide design. A major problem for the solution of PEPs is the lack of standard software for the reliable solution of general PEPs.

In this talk we will discuss two building blocks that are important steps towards the goal of a numerically reliable general purpose PEP solver, namely scaling to improve condition numbers of eigenvalues and iterative refinement to improve the backward and forward error of computed eigenpairs.

Scaling of PEPs

Rescaling a matrix A by a nonsingular diagonal matrix D to $D^{-1}AD$ is a standard technique to improve the sensitivity of the eigenvalues of A to perturbations in the data. In the first part of the talk we extend this strategy to the PEP (1).

Let $P(\lambda)$ be a given matrix polynomial. We define a scaling of $P(\lambda)$ to be the polynomial $D_1P(\lambda)D_2$, where

$$D_1, D_2 \in \mathcal{D}_n := \{D \in \mathbb{C}^{n\times n} :\ D \text{ is diagonal and } \det(D) \neq 0\}.$$

Let $\kappa_P(\lambda)$ be the normwise condition number of an eigenvalue λ of $P(\lambda)$. We show that the minimal achievable condition number $\kappa_{D_1PD_2}(\lambda)$ under scaling with $D_1, D_2 \in \mathcal{D}_n$ can be bounded below and above by the componentwise condition number of the eigenvalue. Furthermore, if the left and right eigenvectors y and x associated with λ have nonzero elements an almost optimal scaling is given by $D_1 = \operatorname{diag}(|y|)$ and $D_2 = \operatorname{diag}(|x|)$.

This does not only give a theoretical bound on the condition number of eigenvalues after scaling but also leads to a cheaply computable test for whether the accuracy of a computed eigenvalue might have been influenced by bad scaling of $P(\lambda)$.

Furthermore, once we have a computed an approximate eigentriplet $(\tilde{y}, \tilde{x}, \tilde{\lambda})$ we can use this information to rescale $P(\lambda)$ with $D_1 = \operatorname{diag}(|\tilde{y}|)$ and $D_2 = \operatorname{diag}(|x|)$ before the start of an iterative refinement procedure.

In order to find scaling matrices D_1 and D_2 before an eigenvalue computation we propose a heuristical algorithm that is a generalization of a scaling method by Lemonnier and Van Dooren for generalized eigenvalue problems [3]. Our algorithm scales $P(\lambda)$ according to the condition

$$\sum_{k=0}^{\ell} \omega^{2k}\|D_1A_kD_2e_i\|_2^2 = 1, \quad \sum_{k=0}^{\ell} \omega^{2k}\|e_j^*D_1A_kD_2\|_2^2 = 1, \quad i,j = 1,\ldots,n, \tag{2}$$

where $\omega > 0$ is a weighting parameter that is ideally chosen close to the magnitude of the wanted eigenvalues. Numerical examples demonstrate that for extreme eigenvalues this can crucially improve the condition numbers under scaling. Further details about these results are contained in [1].

Iterative Refinement

The standard way to solve the PEP (1) for $\ell \geq 2$ is to convert $P(\lambda)$ into a linear pencil $L(\lambda) = \lambda X + Y$ having the same spectrum as $P(\lambda)$ and then solve the eigenproblem for $L(\lambda)$. Using a backward stable solver for the linearization we can expect to compute an eigenpair with a backward error in the order of machine precision for the linearization $L(\lambda)$. However, in [2] it is shown that the backward error of the corresponding approximate eigenpair $(\tilde{x}, \tilde{\lambda})$ for $P(\lambda)$ might be much larger.

Therefore, it is often useful to improve the eigenpair $(\tilde{x}, \tilde{\lambda})$ by iterative refinement, which for the PEP (1) takes the form of applying some steps of Newton iteration to [4]

$$F(x, \lambda) := \begin{bmatrix} P(\lambda)x \\ e_k^T x - 1 \end{bmatrix} = 0.$$

This can be combined with rescaling $P(\lambda)$ using the approximate left and right eigenvectors $\tilde{y}$ and $\tilde{x}$.

We discuss the forward and backward error of an eigenvalue and eigenpair after iterative refinement of a PEP and give results on the computational complexity of the iterative refinement procedure.

This work is part of an ongoing project about "Numerical Algorithms for Polynomial Eigenvalue Problems" which is supported by Engineering and Physical Sciences Research Council grant EP/D079403/1.

References

[1] T. Betcke, *Optimal scaling of generalized and polynomial eigenvalue problems*, MIMS EPrint 2007.120, Manchester Institute for Mathematical Sciences, University of Manchester, Manchester, UK, October 2007. Submitted to SIAM J. Matrix Anal. Appl.

[2] N. J. Higham, R.-C. Li, and F. Tisseur, *Backward error of polynomial eigenproblems solved by linearization*, MIMS EPrint 2006.137, Manchester Institute for Mathematical Sciences, University of Manchester, Manchester, UK, June 2007. To appear in SIAM J. Matrix Anal. Appl.

[3] D. Lemonnier and P. Van Dooren, *Balancing regular matrix pencils*, SIAM J. Matrix Anal. Appl., 28 (2006), pp. 253–263.

[4] F. Tisseur, *Newton's method in floating point arithmetic and iterative refinement of generalized eigenvalue problems*, SIAM Journal on Matrix Analysis and Applications, 22 (2001), pp. 1038–1057.

Error Bounds and Error Estimates for Nonlinear Eigenvalue Problems

David Bindel

Nonlinear eigenvalue problems occur naturally when looking at differential equation models that include damping, radiation, or delay effects. Such problems also arise when we reduce a linear eigenvalue problem, as occurs when we rewrite a PDE eigenvalue problem in terms of a boundary integral equation. In each of these cases, we are interested in studying the zeros of a meromorphic matrix-valued function $A : \mathbf{C} \to \mathbf{C}^{n\times n}$.

When A is a linear function of z, there are standard results that describe regions containing $\Lambda(A)$. Many of these results can be extended to the case of nonlinear eigenvalue problems. For example, let us define the nonlinear spectrum $\Lambda(A) := \{z \in \mathbf{C} : A(z) \text{ singular}\}$ and the nonlinear pseudospectrum $\Lambda_\epsilon(A) := \{z \in \mathbf{C} : \|A(z)^{-1}\| \geq \epsilon^{-1}\}$. Then the following facts hold:

- If A is analytic on a simply connected domain Ω and $D \subset \Omega$ is a connected component of $\Lambda_\epsilon(A)$, then $D \cap \Lambda(A) \neq \emptyset$.
- *Generalized Rouché theorem*: If A and E are analytic on a bounded simply connected domain Ω, A is invertible on a Jordan curve $\Gamma \subset \Omega$, and $\|A(z)^{-1}E(z)\| < 1$ for all $z \in \Gamma$, then A and $A + E$ have the same number of nonlinear eigenvalues (counting multiplicity) inside Γ.
- *Generalized pseudospectral containment*: If A and E are analytic on a simply connected domain Ω and $\|E\| \leq \gamma$ on Ω, then $\Lambda(A + E) \cap D \subset \Lambda_\gamma(A)$.
- *Generalized Gershgorin theorem*: $\Lambda(A) \subset \cup_j\{z : |a_{jj}(z)| \leq r_j(z)\}$, where $r_j(z) = \sum_{i\neq j} |a_{ij}(z)|$.

These simple facts can be used as the basis of eigenvalue counts and rigorous error bounds. If λ is an isolated eigenvalue of A and w^* and v are left and right null vectors for $A(\lambda)$, we can also estimate the change $\delta\lambda$ due to an analytic perturbation $\delta A(z)$ using an analogue of the Rayleigh quotient formula:

$$\delta\lambda = -\frac{w^*(\delta A(\lambda))v}{w^*A'(\lambda)v}.$$

In this talk, we briefly discuss these facts, then show how we use them to obtain error estimates and rigorous error bounds in computations of damping in models of electromechanical resonators and in computations of resonance poles for 1D Schrödinger operators.

Band Householder/Lanczos Methods for Least Squares with Multiple Right Hand Sides

Åke Björck

Golub and Kahan showed that any rectangular matrix A can be decomposed as $A = UBV^T$, where B is a lower (upper) bidiagonal matrix and U and V are orthogonal matrices. The reduction to upper bidiagonal form

$$U^T(b\ AV) = \begin{pmatrix} \beta_1 e_1 & B \\ 0 & 0 \end{pmatrix},$$

where e_1 is the first unit vector, is the core decomposition in the LSQR algorithm of Paige and Saunders. It provides projections of the solution onto a nested sequence of Krylov subspaces. This decomposition provides an elegant way to extract a core problem for both linear least squares and total least squares problems.

In this talk generalizations of the bidiagonal reduction are given for solving problems with $p > 1$ right-hand sides $B = (b_1, \ldots, b_n)$. The matrix (BA) is now reduced to upper band form with $p + 1$ bands. A stable direct method that uses Householder transformations can be devised. Deflation and premature termination of the process is discussed. Finally a Lanczos algorithm is given that generalizes LSQR.

Algebraic multigrid for Helmholtz equations

Matthias Bollhöfer, Marcus Grote, Olaf Schenk

We discuss preconditioning methods for Helmholtz equations in two and three spatial dimensions with high wave number in inhomogeneous media. In particular we present an algebraic multigrid algorithm that is based on three major algorithmic components: Symmetric maximum weight matchings [2] to increase the block diagonal dominance of the system, algebraic inverse–based pivoting [1] to drive the algebraic coarsening process and, filtering based on selected sample frequencies derived from the wave number of the underlying Helmholtz equation. These components can be seen as an efficient black-box approach for the ill-conditioned symmetric indefinite Helmholtz matrices. We present theoretical results and demonstrate the algorithmic effectiveness through inhomogeneous model problems and two- and three dimensional seismic subsurface geology simulations.

The Helmholtz equation is given by a partial differential equation

$$-\Delta u - k^2 u = f$$

in two or three spatial dimensions. Here we are interested in the case, when the wave number k is considerably large such that we are faced with a highly indefinite selfadjoint problem. Typically, boundary conditions of Sommerfeld type lead to complex–valued equations and thus the underlying linear system is complex–symmetric. For the numerical solution we propose a multilevel ILU approach, where on each level the system matrix A is factored as

$$A = \begin{pmatrix} L_F & 0 \\ L_C & I \end{pmatrix} \begin{pmatrix} D_F & 0 \\ 0 & S_C \end{pmatrix} \begin{pmatrix} L_F^T & L_C^T \\ 0 & I \end{pmatrix} + E.$$

Here E refers to the error that is introduced by dropping entries. The main property of this approach consists of controlling the decompositon by a pivoting strategy which keeps $\| \begin{pmatrix} L_F & 0 \\ L_C & I \end{pmatrix}^{-1} \| \leq M$, where M is chosen in advance. We will present theoretical results that are induced by this approach. In particular, the inverse triangular factor will be approximately sparse and furthermore it can be shown that the eigenvalues of the coarse grid system S_C reveal the small eigenvalues in modulus from the original system. In the real symmetric case strict bounds for the relations between the eigenvalues λ_j of the original system and μ_j of the coarse grid system will be derived. We will see that for positive eigenvalues we have

$$\frac{1}{M^2}\frac{1}{\lambda_j} - \|D_F^{-1}\| \leq \frac{1}{\mu_j} \leq \frac{1}{\lambda_j} + M^2\|D_F^{-1}\|.$$

For negative eigenvalues an analogous inequality applies. This justifies driving our coarsening process with a small M such as $M = 3$ in order to ensure that the coarse grid system will pick up the eigenvalues with small modulus from the original system. Beside keeping the inverse triangular factor bounded, the approximate multilevel decomposition is stabilized using sample frequencies depending on the wave number k. If we discretize the domain using finite differences, then we must ensure that sufficiently many grid points are used to resolve the relevant frequencies properly. This induces a sample frequency with alternating sign pattern and we keep the approximate decomposition exact along this sample frequency. We demonstrate the algorithmic effectiveness through inhomogeneous model problems and two- and three dimensional seismic subsurface geology simulations. It is shown that the Krylov-subspace iteration using our algebraic multilevel preconditioner depends relatively mildly on the wave number and that the algorithmic components have a highly beneficial impact on efficiency and convergence.

2D exterior scattering wave problem on a disc

	grid	ILUT+BICGSTAB			AMG+sQMR		
$freq.$	points	$\frac{nnz(LU)}{nnz(A)}$	its.	sec.	$\frac{nnz(LDL^T)}{nnz(A)}$	its.	sec.
π	60×480	2.7	93	14	3.1	29	2
$3 \times \pi$	180×1440	2.6	491	642	2.7	145	47
$6 \times \pi$	360×2880	2.7	1027	5066	2.8	377	481

References

[1] M. Bollhöfer and Y. Saad. Multilevel preconditioners constructed from inverse–based ILUs. *SIAM J. Sci. Comp.*, 27(5):1627–1650, 2006.

[2] I.S. Duff and S. Pralet. Strategies for Scaling and Pivoting for Sparse Symmetric Indefinite Problems. *SIAM J. on Matr. Anal. Appl.*, 27(2):314–340, 2005.

Structured eigenvalue condition number and backward error of a class of polynomial eigenvalue problems

Shreemayee Bora

Given a complex matrix polynomial $P(\lambda) := \Sigma_{k=0}^{m} \lambda^k A_k$, we consider the polynomial eigenvalue problem of finding a vector x and a scalar λ such that $P(\lambda)x = 0$ under the additional assumption $P(\lambda)$ has $\star$-palindromic, $\star$-antipalindromic, $\star$-even or $\star$-odd structure. Here $\star$ denotes the transpose T or the conjugate transpose $*$. The name 'palindromic' refers to a word or a phrase which remains unchanged upon writing the letters in the reverse order. For a polynomial $P(\lambda) = \Sigma_{k=0}^{m} \lambda^k A_k$, its reversal denoted by rev$P(\lambda)$ is the polynomial rev$P(\lambda) = \lambda^m P(1/\lambda) = \Sigma_{k=0}^{m} \lambda^{m-k} A_k$. Also the adjoint $P^\star(\lambda)$ is given by $P^\star(\lambda) = \Sigma_{k=0}^{m} \lambda^k A_k^\star$. In the context of polynomials, $\star$-palindromic structure implies that we get back the original polynomial upon taking its reversal and then its adjoint. On the other hand $\star$-even polynomials are such that the original polynomial may be obtained by taking its adjoint and replacing λ by $-\lambda$. Table 2 gives the precise definitions of all these polynomials as

Table 2: Some matrix polynomial structures and corresponding spectral symmetries

Basic structure	Property	Spectral symmetry
T-palindromic	$\mathrm{rev}P^T(\lambda) = P(\lambda)$	$(\lambda, 1/\lambda)$
T-antipalindromic	$\mathrm{rev}P^T(\lambda) = -P(\lambda)$	$(\lambda, 1/\lambda)$
$*$-palindromic	$\mathrm{rev}P^*(\lambda) = P(\lambda)$	$(\lambda, 1/\overline{\lambda})$
$*$-antipalindromic	$\mathrm{rev}P^*(\lambda) = -P(\lambda)$	$(\lambda, 1/\overline{\lambda})$
T-even	$P^T(-\lambda) = P(\lambda)$	$(\lambda, -\lambda)$
T-odd	$P^T(-\lambda) = -P(\lambda)$	$(\lambda, -\lambda)$
$*$-even	$P^*(-\lambda) = P(\lambda)$	$(\lambda, -\overline{\lambda})$
$*$-odd	$P^*(-\lambda) = -P(\lambda)$	$(\lambda, -\overline{\lambda})$

The algebraic, geometric and partial multiplicities of the eigenvalues in each pair listed in the table are equal (for a proof see [4]).

well as the associated spectral symmetries. We observe that the spectral symmetry of $*$-palindromic/antipalindromic polynomials is with respect to the unit circle, while that of the $*$-even/odd is with respect to the imaginary axis. These polynomials came into focus in the work of Mackey, Mackey, Mehl and Mehrmann in [4] where they develop a systematic approach for constructing structure preserving linearizations of these polynomials and list some important applications.

We consider the condition number of simple eigenvalues as well as the backward error of approximate eigenvalues of these polynomials with respect to structure preserving perturbations, the setting of norms on the coefficient matrices and the polynomial being the same as the one considered in [9]. The main contribution of this work is that it obtains characterisations of simple eigenvalues that have equal structured and unstructured condition numbers. In the process we derive formulae for the structured condition numbers for $*$-palindromic/antipalindromic and T-even/odd polynomials. For the other structures we obtain lower and upper bounds on the structured condition number. Numerical experiments suggest that these bounds are fairly tight and that the structured condition number is almost equal to the unstructured one. We also investigate the structured and unstructured backward errors of approximate eigenvalues of these polynomials with the aim of obtaining conditions under which these quantities are equal. In particular we derive sets of conditions that are necessary for the two backward errors to be equal and then show that some of these conditions are also sufficient for such an equality to occur. In the process we obtain complete characterisations of complex numbers that have the same structured and unstructured backward error for the $*$-even/odd, T-even/odd polynomials as well as $*$-palindromic/antipolindromic polynomials of even degree.

The results indicate that simple eigenvalues of complex $*$-palindromic/antipalindromic polynomials have the same structured and unstructured normwise condition numbers while the same holds for simple eigenvalues of complex $*$-even/odd polynomials only if they are purely imaginary or have orthogonal left and right eigenvectors. We also observe that the sets of approximate eigenvalues of complex $*$-palindromic/antipalindromic polynomials that have the same structured and unstructured backward errors include the unit circle while those for the complex $*$-even/odd polynomials include the imaginary axis.

Structure preserving perturbations of smallest size that cause $*$-palindromic and $*$-even polynomials to have eigenvalues on the imaginary axis and the unit circle respectively are very important for control theory applications as these may result in loss of spectral symmetry and uniqueness in the choice of the deflating subspace associated with the eigenvalues in the left half plane (respectively the open unit disk) [2, 5, 6, 7, 8]. It is well known that such information may be read off from the structured ϵ-pseudospectra of polynomials [10]. The results for the backward error show that the same information may be obtained from the corresponding unstructured ϵ-pseudospectra. A similar conclusion holds for the $*$-even polynomials vis a vis the imaginary axis.

The presence of the numbers 1 and -1 in the spectra of T-palindromic and T-antipalindromic polynomials may lead to difficulties in finding structure preserving linearizations of these polynomials (see Section 6, [4]). The formulae for the structured condition number of simple eigenvalues of these polynomials indicate that the eigenvalues that display equal sensitivity to structure preserving and arbitrary perturbations are essentially real numbers. Therefore 1 and -1 have the same structured and unstructured condition numbers when they are simple eigenvalues of such polynomials except under certain conditions when these numbers always occur in the spectra of these polynomials (in which case their structured condition numbers are 0). We further show that the sets of approximate eigenvalues of T-palindromic/antipalindromic polynomials having equal structured and unstructured backward errors include the entire real axis in some cases and either the positive or negative part of it in others, depending upon the structure and the degree of the polynomial. However, in all cases, these sets always include the numbers 1 and -1.

References

[1] S. Bora. Structured eigenvalue conditioning and backward error of a class of poynomial eigenvalue problems. Preprint 417, DFG Centre MATHEON (Mathematics for Key Technologies), TU Berlin, Germany.

[2] G. Freiling, V. Mehrmann and H. Xu. Existence, uniqueness and parametrization of Lagrangian invariant subspaces. *SIAM J. Matrix Anal. Appl.,* Vol 23, No. 4, pp. 1045-1069.

[3] N. J. Higham, D. S. Mackey and F. Tisseur. The conditioning of linearizations of matrix polynomials, MIMS eprint 2005.23, School of Mathematics, University of Manchester. To appear in *SIAM J. Matrix Anal. Appl.*

[4] D. S. Mackey, N. Mackey, C. Mehl and V. Mehrmann. Palindromic polynomial eigenvalue problems,: Good vibrations from good linearizations, *SIAM J. Matrix Anal. Appl.*, to appear.

[5] V. Mehrmann, *The Autonomous Linear Quadratic Control Problem, Theory and Numerical Solution,* Number 163 in Lecture Notes in Control and Information Sciences, Springer-Verlag, Heidelberg, July 1991.

[6] A.C.M. Ran and L. Rodman. Stability of invariant maximal semidefinite subspaces. *Linear Algebra Appl.,* 62: 51-86, 1984.

[7] A.C.M. Ran and L. Rodman. Stability of invariant Lagrangian subspaces i. *Operator Theory: Advances and Applications (I. Gohberg ed.),* 32: 181-218, 1988.

[8] A.C.M. Ran and L. Rodman. Stability of invariant Lagrangian subspaces ii. *Operator Theory: Advances and Applications (H.Dym, S. Goldberg, M.A. Kaashoek and P. Lancaster eds.)*, 40: 391-425, 1989.

[9] F. Tisseur. Backward error and condition of polynomial eigenvalue problems, *Linear Algebra Appl.,* 309(2000), pp. 309-339.

[10] F. Tisseur and N. J. Higham. Structured pseudospectra for polynomial eigenvalue problems, with applications. *SIAM J. Matrix Anal. Appl.,* 23(10): 187-208, 2001.

Middle Deflations in the QR Algorithm

Karen Braman

Since its introduction in 1961, the QR algorithm [4, 7] has remained the primary workhorse for the solution of non-symmetric eigenvalue problems. Recent work has improved the performance of the QR algorithm by applying an aggressive early deflation strategy [3]. Like conventional small-subdiagonal deflation, aggressive early deflation creates zero subdiagonal entries near the lower-right-hand corner of a Hessenberg matrix. This causes the eigenvalue problem to split into two asymmetric parts: a small Hessenberg matrix in the lower-right-hand corner and a much larger Hessenberg matrix in the upper-left-hand corner. Ideally, a deflation procedure would produce a zero subdiagonal near the middle of the Hessenberg matrix, thus splitting the eigenvalue problem more symmetrically into two subproblems of nearly equal size. If successful, such a middle deflation would cut the work remaining from $O(n^3)$ [5] on the entire matrix to $O\left(\frac{n^3}{8}\right)$ on each of the two remaining sub-matrices. If it was to cost $O(n^2)$ flops to obtain a middle deflation and such deflations could be found consistently, then the entire cost for an n-by-n eigenvalue problem drops to $O(n^2 \log n)$. Although an asymptotic flop count under $O(n^3)$ may not be achievable, there remains potential for a substantial reduction in execution time. Thus, middle deflations are valuable enough that one can spend a fair amount of work (say $O(k^3)$ for $k \ll n$) and still see an overall increase in performance.

Implementation: The procedure for searching for middle deflations can be described as follows: Beginning with H, a QR iterate in Hessenberg form, consider H_{22}, a k-by-k block situated near the middle of H. Compute a Schur factorization $H_{22} = UTU^T$. If one was to perform a similarity transformation on H using $\hat{U} = I_{m_1} \oplus U \oplus I_{m_2}$ with I_{m_1} and I_{m_2} identity matrices of appropriate sizes, then the resulting matrix $\hat{H} = \hat{U}^T H \hat{U}$ would have a *Hessenberg-plus-spikes form*. That is, it would be in Hessenberg form except for the $(2,1)$ block which would have non-zero elements in its last column (the *vertical spike*) and the $(3,2)$ block which would have non-zero elements in its first row (the *horizontal spike*). Also, the middle k-by-k block would be quasi-triangular, with the eigenvalues of H_{22} in 1-by-1 and 2-by-2 blocks along its diagonal. Note that each spike is simply a scaling of one of the rows of U. The vertical spike is $h_{i,i-1}(u_{1,:})^T$ and the horizontal spike is $h_{j+1,j}(u_{k,:})$, where $(u_{p,:})$ denotes the pth row of U, and H_{22} lies in rows and columns i through $j = i + k - 1$. The presence of middle deflations is determined by the number of negligible spike elements, working from the

bottom of the vertical spike and from the left of the horizontal spike. If the total number of negligible spike tip elements is greater than k, a middle deflation reducing perturbation can be formulated. Computationally, the similarity transformation described above is only necessary if it will produce deflations, and so the procedure initially only computes the spikes (by scaling the first and last rows of U appropriately) and then checks for small tip elements using the window-norm-stable criterion as described in [3]. Let r and s denote the number of deflatable tip elements in the horizontal and vertical spikes, respectively. If $r + s > k$, then setting the negligible tip elements to zero creates, in theory, a reducing perturbation P (see [3]). Completing the middle deflation process requires returning $\hat{H} + P$ to Hessenberg form via orthogonal similarity transformations and doing so in such a manner that only affects the first $k - r$ rows and columns and the last $k - s$ rows and columns of the midddle deflation window. This can be done effectively with Householder reflectors, working column-wise on the upper portion of the window and row-wise on the lower portion. Accumulating these reflectors along with U into W gives $\tilde{H} \equiv W^T(H+P)W$ where $\tilde{H}$ now has $r+s-k$ deflations near its middle row.

Numerical Experiments: The process above has been implemented in FORTRAN in such a way as to provide a careful comparison with LAPACK's [1] most recent version of `DHSEQR` (LAPACK 3.1.1) which includes aggressive early deflation. The new code, `DHSMQR`, does not use the aggressive early deflation strategy at all. Rather it applies the procedure above to check for middle deflations after each QR sweep.

The `DHSEQR` and `DHSMQR` codes were run on various matrices from the collection of non-Hermitian eigenvalue problems [2]. For some of these matrices the middle deflation strategy provides a significant improvement over aggressive early deflation and for others it does not. For instance, `DHSMQR` is 48% faster than `DHSEQR` on **tols2000** (2.621 minutes vs. 5.084 minutes), about the same on **olm2000** (1.681 minutes vs. 1.689 minutes), and 57% slower on **bwm2000** (8.792 minutes vs. 5.592 minutes). These differences can be partially attributed to the number of QR sweeps applied to the entire matrix before middle deflations were found. For example, in the case of **bwm2000**, 9 sweeps were required before middle deflations were found, which means the code spent time looking for middle deflations that weren't there 8 times. However, the expense of failed searches does not tell the whole story either, as `DHSMQR` won big on **tols2000** after applying 4 QR sweeps before finding deflations, but essentially tied with `DHSEQR` on **olm2000** even though only one sweep was necessary. Certainly understanding the factors that influence execution time is a complex area which warrants further investigation.

One way to minimize the expense of failed middle deflation searches is to quit looking as soon as failure is obvious. That approach is implemented in the results above by using a variation of `DHSEQR` to compute the Schur factorization of the middle deflation window. `DHSPQR` 'peeks' through the window by monitoring the computation of the orthogonal Schur factor of the middle deflation window and abandoning that computation as soon as that factor will no longer have the necessary number of negligible elements in the first and last rows.

Some Areas of Continuing Investigation: As currently implemented, `DHSMQR` computes one Schur factorization of the middle window $H_{22} = UTU^T$, checks the first and last rows of U for small-norm tip elements and then gives up if an insufficient number have been found. That is by no means an exhaustive search. It is possible that a different Schur factorization would have an orthogonal factor with the necessary structure. One way to check other factorizations is to reorder the eigenvalues on the diagonal of T. While it has not yet been included in the FORTRAN code, preliminary investigations using Matlab show that

'swapping' can coax a middle deflation out of a window that did not show any at first glance.

The major cost of searching for middle deflations comes from computing the Schur factorization of the deflation window. Even with the strategy of 'quit looking when failure is obvious' (`DHSPQR`), it can happen that nearly the entire factorization process is completed before the search is abandoned. It would be beneficial, then, to have a means of predicting, with a relatively small amount of work, the success or failure of the deflation process before the computation of the Schur factor. Recalling the results on **tols2000**, **olm2000**, and **bwm2000**, it would be useful to have a heuristic to predict *a priori* which of these should be approached with middle deflations and which should be left to aggressive early deflation.

In addition to an effective heuristic, a theoretical explanation for the middle deflation process is desirable. Kressner [6] has recently shown how the aggressive early deflation strategy is equivalent to extracting converged Ritz pairs from a certain Krylov subspace. Since the middle deflation strategy is an extension of the aggressive deflation approach, it should be possible to develop a similar convergence theory, possibly involving certain left and right Krylov subspaces.

Conclusions: Middle deflation shows much promise for improving the performance of the QR algorithm. Indeed, it has been seen to perform better even than aggressive early deflation on some of the NEP matrices. A deeper understanding of the convergence theory, a careful implementation of 'swapping' and discovery of an effective heuristic for predicting middle deflations would all be steps toward the elusive $O(n^2 \log n)$ flop count.

References:

[1] E. Anderson, Z. Bai, C. Biscof, S. Blackford, J. Demmel, J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, A. McKenney, and D. Sorensen. *LAPACK User's Guide.* Society for Industrial and Applied Mathematics, Philadelphia, PA, Third Edition, 1999. FORTRAN source available for version 3.1.1 online from http://www.netlib.org/lapack/index.html.

[2] A. Bai, D. Day, J. Demmel and J. Dongarra. *A Test Matrix Collection for Non-Hermitian Eigenvalue Problems.*, 1996. Available online from http://www.cs.ucdavis.edu/~bai/NEP/.

[3] K. Braman, R. Byers, and R. Mathias. The Multishift QR Algorithm. II. Aggressive Early Deflation. *SIAM J. Matrix Anal. Appl.*, 23(4):948-973, 2002.

[4] J.G.F. Francis. The QR Transformation. I and II. *Comput. J.*, 4:265-271 and 332-345, 1961/1962.

[5] G.H. Golub and C.F. Van Loan. *Matrix Computations.* The Johns Hopkins Press, Baltimore, Maryland, Third edition, 1996.

[6] D. Kressner. The Effect of Aggressive Early Deflation on the Convergence of the QR Algorithm. Technical Report, Umea University, Department of Computing Science, 2006.

[7] V.N. Kublanovskaya. On Some Algorithms for the Solution of the Complete Eigenvalue Problem. *USSR Journal of Computational Mathematics and Physics.*, 3:637-657, 1961.

Model reduction for large scale MIMO-systems via rational tangential Hermite interpolation

Angelika Bunse-Gerstner, Dorota Kubalinska, Daniel Wilczek, Georg Vossen

Modelling strategies in various fields of application often result in dynamical systems of very high dimension where the high complexity makes it impossible to work with these systems in further computational processes. Model reduction methods can then be helpful as they provide automatic processes which find systems of the same form but of lower complexity, whose input-output-behaviour approximates the behaviour of the original system. VLSI-design, chemical reactors, micro-electronic-mechanical systems or climate and ocean modelling are some of the areas of application.

For linear time invariant systems with multiple input and multiple output (MIMO systems) the input-output-behaviour can be characterized by the transfer function $H(z)$ of the system, which is a $p \times m$ dimensional rational matrix function, where m is the number of inputs and p the number of outputs of the system.

This talk considers rational tangential Hermite interpolation to compute projections for large scale MIMO systems. They reduce the system's dimension such that the $p \times m$ transfer function of the reduced system is close to the transfer function of the large system. First order necessary conditions are presented for the interpolation points and directions to give optimal approximations in the H_2 or h_2-norm. Analogous results are shown for more general norms. These new results for general norms allow to handle also unstable systems, which for instance typically occur in weather modelling, in a similar manner as the stable ones on the basis of the usual H_2 or h_2-norm norms.

Computational methods are shown to compute the projections for interpolation data satisfying these first order conditions. Numerical experiments demonstrate the approximation quality of the method.

Solving Integer Linear Least Squares Problems

Xiao-Wen Chang

Integer least squares (ILS) problems may arise from many applications such as communications, cryptoanalysis, lattice design, Monte Carlo second-moment estimation, radar imaging, global positioning systems etc, see, e.g., [1], [8], [9], and [10]. Unlike real LS problems, ILS problems are NP-hard. A typical algorithm to solve an ILS problem consists of two phases: reduction and search. The reduction process transforms the original problem to a new problem to make the search process easier and more efficient and we can cast its crucial part as a matrix factorization. The search process is to generate and examine the integer points over a specific region and find the optimal solution.

In this talk, we will present some of our recent work on various ILS problems, see [2]-[7]. The key in our approaches is to incorporate column pivoting strategies in the QR factorization in the reduction process, which can make the search process much more efficient.

In the following, we use our work on box-constrained underdetermined ILS problems to illustrate the main ideas involved in solving ILS problems. Given a matrix $\mathbf{A} \in \mathcal{R}^{m \times n}$ with

full row rank, $\mathbf{b} \in \mathcal{R}^m$, $\mathbf{l} \in \mathcal{Z}^n$ and $\mathbf{u} \in \mathcal{Z}^n$, a box-constrained underdetermined ILS problem has the form:

$$\min_{\mathbf{x}\in\mathcal{Z}^n} \|\mathbf{b} - \mathbf{A}\mathbf{x}\|_2, \quad \text{subject to} \quad \mathbf{l} \le \mathbf{x} \le \mathbf{u}. \tag{1}$$

This underdetermined problem is more complicated than an overdetermined problem. Suppose we compute the QR factorization of $\mathbf{A}$ with some column pivoting strategy, giving $\mathbf{AP} = \mathbf{QR}$, where $\mathbf{R}$ is trapezoidal. Then with $\bar{\mathbf{b}} \equiv \mathbf{Q}^T\mathbf{b}$, $\mathbf{z} \equiv \mathbf{P}^T\mathbf{x}$, $\bar{\mathbf{l}} \equiv \mathbf{P}^T\mathbf{l}$ and $\bar{\mathbf{u}} \equiv \mathbf{P}^T\mathbf{u}$, the problem (1) is transformed to

$$\min_{\mathbf{z}\in\mathcal{Z}^n} \|\bar{\mathbf{b}} - \mathbf{R}\mathbf{z}\|_2, \quad \text{subject to} \quad \bar{\mathbf{l}} \le \mathbf{z} \le \bar{\mathbf{u}}. \tag{2}$$

Suppose that the optimal solution $\mathbf{z} \in \mathcal{Z}^n$ satisfies $\|\bar{\mathbf{b}} - \mathbf{R}\mathbf{z}\|_2 < \beta$, where β can be chosen by taking the nearest integer solution to the corresponding box constrained *real* LS problem, see [3]. The search process will find the optimal solution within this hyper-ellipsoid. Then we have

$$(\bar{b}_m - \sum_{j=m}^{n} r_{mj}z_j)^2 < \beta^2, \tag{3}$$

$$(\bar{b}_k - \sum_{j=k}^{n} r_{kj}z_j)^2 < \beta^2 - \sum_{i=k+1}^{m} (\bar{b}_i - \sum_{i=j}^{n} r_{ij}z_j)^2, \quad k = m-1:-1:1. \tag{4}$$

From (3) and (4) as well as the box constraint in (2) we can establish a sequence of inequalities

$$\lambda_n \le z_n \le \mu_n, \quad \lambda_k(z_{k+1},\ldots,z_n) \le z_k \le \mu_k(z_{k+1},\ldots,z_n), \quad k = n-1:-1:1,$$

where the lower and upper bounds on z_k depend on $z_{k+1},\ldots,z_n$, then based on these inequalities, a first-depth tree search process can be developed to find the optimal solution, see [5].

The efficiency of the search process depends on $\mathbf{R}$ in (2). In order to make it fast, we propose some column pivoting strategies in the QR factorization of $\mathbf{A}$, see [2] and [5].

Numerical results will be given to show our algorithm can be much faster than the existing algorithms in the literature.

References

[1] E. Agrell, T. Eriksson, A. Vardy, and K. Zeger, *Closest point search in lattices*, IEEE Trans. Inform. Theory, 48 (2002), pp. 2201–2214.

[2] X.-W. Chang and Q. Han, *Solving box-constrained integer least squares problems*, to appear in IEEE Transactions on Wireless Communications.

[3] X.-W. Chang and X. Yang, *A new fast generalized sphere decoding algorithm for underdetermined MIMO systems*, in Proceedings of 23rd Biennial Symposium on Communications, Kingston, Ontario, Canada, May 30 - June 1, 2006, pp. 18-21.

[4] X.-W. Chang and X. Yang, *An efficient regularization spproach for underdetermined MIMO system decoding*, in Proceedings of International Wireless Communications and Mobile Computing Conference 2007, Honolulu, Hawaii, August 12-16, 2007, pp. 349-353.

[5] X.-W. Chang and X. Yang, *An efficient tree dearch decoder for underdetermined MIMO systems*, to appears in Proceedings of IEEE GLOBECOM 2007, November 26-30, 2007, Washington, D.C.

[6] X.-W. Chang, X. Yang, and T. Zhou, *MLAMBDA: A modified LAMBDA method for integer least-squares estimation*, J. of Geodesy, 79 (2005), pp. 525–565.

[7] X.-W. Chang and T. Zhou, *MILES: MATLAB package for solving Mixed Integer LEast Squares problems*, to appear in GPS Solutions.

[8] M. O. Damen, H. El Gamal, and G. Caire, *On maximum-likelihood detection and the search for the closest lattice point*, IEEE Trans. Inform. Theory, 49 (2003), pp. 2389–2402.

[9] C. P. Schnorr and M. Euchner, *Lattice basis reduction: improved practical algorithms and solving subset sum problems*, Mathematical Programming, 66 (1994), pp. 181–199.

[10] P. Teunissen, *The least-squares ambiguity decorrelation adjustment: a method for fast GPS ambiguity estitmation*, J. of Geodesy, 70 (1995), pp. 65–82.

The Linear Algebra of a Multi-document Summarization System

John M. Conroy, Dianne P. O'Leary, Judith D. Schlesinger

Multi-Document summarization is a top-down technique to cope with information overload. Instead of simply listing a set of documents retrieved by a query, a short summary for one or more clusters is given to either answer the user's request or help the user determine if the documents are relevant to the request. Much attention has been given to this problem in recent years, and performance of systems, as measured by automatic techniques as well as human evaluation, has greatly increased. Two model problems have been studied and evaluated in a series of competitions:

1. Given a cluster of 5 to 10 documents on the same topic in multiple languages, generate a 100 word paragraph giving the the essential information in the documents.

2. Given 25 documents in English and a topic description giving the interest of the user, generate a 250 word topic-focused summary.

Our system, CLASSY (Clustering, Linguistics [Linear], And [Algebra] Statistics for Summarization Yield), performs extremely well and far exceeds its peers in the first task. In fact, it achieved human-level performance in the 2006 Multi-lingual Summarization Evaluation (http://research.microsoft.com/ lucyv/MSE2006.htm).
No system has achieved human performance at the second task; however, CLASSY is in the top performing group of systems on this task, as measured in the most recent NIST evaluation (http://duc.nist.gov/).

In this talk we focus on the linear algebra employed by CLASSY, which includes the truncated SVD, pivoted QR, and non-negative matrix factorization.

We now give a brief overview of CLASSY to define the underlying linear algebraic problems. CLASSY is divided into four major parts:

1. Linguistic preprocessing defines sentence boundaries, tags sentences by their type (e.g., headline, dateline, body), and then reduces their length using linguistic heuristics to shorten sentences when appropriate.

2. Sentences are given a score that approximates the fraction of terms it shares with an unknown "model" human summary.

3. A subset selection problem is solved to find a "non-redundant" collection of high scoring sentences.

4. The final summary is generated by arranging selected sentences to improve "flow." This is achieved by approximately solving a Traveling Salesman Problem.

We now elaborate on parts 2 and 3, which are computational linear algebra problems key to recent performance gains.

For a given cluster of documents, a term-sentence matrix is constructed. The columns of the matrix represent the sentences of the document and the rows represent the terms (loosely the words) that occur in the cluster. There are two distinguished subsets of terms, "query-terms," which are extracted from the topic description, and "signature-terms." Signature terms are terms that occur more frequently in the cluster than expected, as computed by a mutual information statistic.

Sentences are scored for inclusion in a summary using an "approximate oracle score" that estimates the fraction of terms a sentence has in common with unknown human model summaries. The score is built upon an estimate of the probability that a term t will be included in a human summary given a topic, τ. This probability is denoted $P(t|\tau)$. It is approximated using the query terms, signature terms, and the distribution of the terms in the relevant document cluster. Specifically, we estimate our target probability by a weighted average of three distributions: the impulse distributions for the query and signature terms and the probability that a term occurs in the sentences to be considered for extraction:

$$P_{qs\rho}(t|\tau) = \frac{1}{4}q_t(\tau) + \frac{1}{4}s_t(\tau) + \frac{1}{2}\rho_t(\tau)$$

where $s_t(\tau)$=1 if t is a signature term for topic τ and 0 otherwise, $q_t(\tau) = 1$ if t is a query term for topic τ and 0 otherwise, and $\rho_t(\tau)$ is the maximum likelihood estimate of the probability that term t occurs in a sentence in the topic τ.

Based on the above estimate, the approximate oracle score for sentence x is computed as

$$\omega_{qs\rho}(x) = \frac{1}{|x|}\sum_{t\in T} x(t)P_{qs\rho}(t|\tau),$$

where $x(t) = 1$ if the sentence x contains the term t and 0 otherwise. This estimate is quite good and typically correlates with an "oracle score" based on human generated summaries with Pearson correlation coefficients in the range of 0.62 to 0.98 depending on the cluster of documents.

Computing the score $\omega_{qs\rho}(x)$ can be viewed as a matrix-vector multiply if we define $p_i = P_{qs\rho}(t = i|\tau)$ and $\hat{A}$ the column stochastic term-sentence matrix with

$$\hat{A}_{ij} = \frac{1}{|x|}x(t = i).$$

Thus the vector of scores for each sentence is given by

$$\omega_{qs\rho} = p' \hat{A}.$$

Our score identifies the sentences that are likely to contain terms that would occur in a human summary. The goal is to select a subset of the sentences that have the proper coverage (distribution) of terms for the summary. We can model this task as a classic subset selection problem. We initiate this task by first constructing a term-sentence matrix A containing the top scoring sentences from the cluster, with the columns weighted proportionally to their approximate oracle score. The goal is then to select a non-redundant subset of the sentences. We have explored two approaches to this problem:

1. the pivoted QR factorization, and
2. a truncated singular value decomposition followed by a direct method for approximate non-negative matrix factorization.

The second approach was demonstrated to produce summaries that more closely match the distribution of terms included in human generated summaries. The essential part of the algorithm, the part that gives the biggest gain, is respecting the non-negativity. It was this approach that produced summaries in the multi-lingual summarization task that were judged to have content comparable to the human generated summaries. The talk will contrast our direct method for approximate non-negative factorization with the iterative methods such as those proposed by Thomas Hoffman, Michael Berry, Inderjit Dhillon, and Haesun Park.

From Eigenvalues to Singular Values: The Rectangular Quotients Equality and Related Issues

Achiya Dax

The question that initiates our talk is how to extend the Rayleigh quotient in order to estimate a singular value of a general rectangular matrix. Let $\mathbf{A}$ be a real $m \times n$ matrix and let $\mathbf{u}$ and $\mathbf{v}$ be a pair of nonzero vectors. It is shown that the Rectangular quotient,

$$\sigma = \sigma(\mathbf{u}, \mathbf{A}, \mathbf{v}) = \mathbf{u}^T \mathbf{A} \mathbf{v} / (\|\mathbf{u}\|_2 \, \|\mathbf{v}\|_2),$$

solves a number of least norm problems that characterize the "corresponding" singular value. Also, as with the Rayleigh quotient, there exist bounds on the distance between σ and the closest singular value.

A second aim of this talk is to introduce a new matrix equality, the Rectangular Quotients Equality (RQE), and to discuss its consequences. Recall that the celebrated Eckhart-Young theorem considers a minimum norm problem, in which a given matrix is approximated by another matrix of a lower rank. The RQE illuminates Eckhart-Young theorem in a new light, turning it from a minimum norm problem into an equivalent maximum problem. The new objective function to be maximized is composed from a sum of squares of rectangular quotients.

The symmetric version of the RQE considers traces of symmetric matrices, and adds similar insight into Ky Fan's extremum principles. In particular, we show that Ky Fan's

maximum problem is equivalent to a problem of minimizing the trace of a difference between two symmetric matrices, where one matrix has a restricted rank. A comparison of the two cases reveals a remarkable similarity between Eckhart-Young theorem and Ky Fan's maximum principle.

The extremum principles of Ky Fan consider sums and products of eigenvalues of Hermitian matrices. The third part of the talk derives extended rules for rectangular matrices. In this case singular values take the role of eigenvalues. Yet, as we shall see, there are interesting differences between the two cases. The proofs are based on 'rectangular' versions of Cauchy interlace theorem and Poincaré separation theorem.

The minimum norm approach, which motivates the Rectangular quotient, enables us to retrieve a left singular vector from a right one, and vice versa. The ultimate retrieval rule is both a minimum norm solution and a result of an orthogonalization process: Given an estimate, $\mathbf{u}$, of a left singular vector, then the corresponding right singular vector is obtained by orthogonalizing the columns of $\mathbf{A}$ against $\mathbf{u}$. Similarly, given an estimate, $\mathbf{v}$, of a right singular vector, then the corresponding left singular vector is obtained by orthogonalizing the rows of $\mathbf{A}$ against $\mathbf{v}$. These observations lead to a new orthogonalization method, one that is called "Orthogonalization via Deflation". The new method constructs an SVD-type decomposition, which is useful for calculating low-rank approximations of large sparse matrices. Another favourable situation arises in problems with missing data. If time allows I will outline the details of the new scheme.

Matrix Analysis Problems arising in Fast Implementations of the Quantum Monte Carlo Method[1]

Eric de Sturler[2], David M. Ceperley[3], Kapil Ahuja[2], Bryan Clark[3]

The most accurate methods for computing material properties, requiring no prior assumptions, are the Quantum Monte Carlo Methods (QMC), such as Variational Monte Carlo. Unfortunately, so far, these methods are prohibitively expensive for large numbers of particles, as the evaluation of probabilities in each step of a Monte Carlo (MC) integration requires the computation of a determinant (see below), or more precisely the ratio of two determinants. As a result, the computational complexity of the algorithm is $O(N^3)$, where N is the number of particles and the complexity is measured as the (average) cost of N single particle moves. Nevertheless, these methods are important because alternative methods (such as density functional theory) do not provide sufficient accuracy for many problems.

The problems in deriving fast algorithms ($O(N)$ or close) for QMC are, essentially, matrix analysis questions. There are also intriguing links with higher-dimensional interpolation. I will provide a basic introduction to QMC methods, describe the most important problems, and pose these as matrix analysis or numerical linear algebra problems. For several of these problems I will provide solutions, but for others I would like to probe the collective knowledge and intelligence of the numerical linear algebra community. The solution to these problems may have a huge impact on computational physics, chemistry, and materials science. I will also

[1]This research was supported, in part, by NSF grants EAR-0530643 and DMR 0325939 through the Materials Computation Center at UIUC.

[2]Dept. of Mathematics, Virginia Tech

[3]Dept. of Physics, UIUC

demonstrate numerically that for several classes of simple but physically meaningful problems our methods achieve linear scaling. This is important, because there are some differences between the theoretical formulation of the problem and its practical implementation.

In the remainder of this abstract, I provide some details about problems in fast QMC algorithms. The computational physics aspects are taken from [5] and [3]. We are looking for a wave function that describes the probability distribution of electron positions and spin given the positions of the nuclei (which are treated separately). This wave function is often chosen to be the product of relatively simple single particle wave functions $\Psi(\mathbf{x_1}, \ldots, \mathbf{x_N}) = \mathbf{\Pi_{i=1}^N} \psi_\mathbf{i}(\mathbf{x_i})$, where the particle coordinates $\mathbf{x_i}$ represent 3D position and spin. However, QMC can treat arbitrary wave functions. For electrons we need the wave function to be antisymmetric. An antisymmetric wave function that is a linear combination of (the minimum number of) terms of the form above can be represented compactly as a Slater determinant

$$\Psi_{\mathrm{AS}}(\mathbf{x_1}, \ldots, \mathbf{x_N}) = \frac{\mathbf{1}}{\sqrt{\mathbf{N!}}} \begin{vmatrix} \psi_1(\mathbf{x_1}) & \psi_2(\mathbf{x_1}) & \cdots & \psi_N(\mathbf{x_1}) \\ \psi_1(\mathbf{x_2}) & \psi_2(\mathbf{x_2}) & \cdots & \psi_N(\mathbf{x_2}) \\ \vdots & \vdots & \ddots & \vdots \\ \psi_1(\mathbf{x_N}) & \psi_2(\mathbf{x_N}) & \cdots & \psi_N(\mathbf{x_N}) \end{vmatrix}. \tag{1}$$

The Slater determinants form a basis in the space of antisymmetric wave functions. So, we consider a linear combination of many particle wave functions, compute the expectation value of the energy, and vary the coefficients following some minimization algorithm to find the ground state. The computation of the expected energy forms the computational bottleneck, integrating for all particles over a 3D region. This step is typically carried out using MC integration. Since the wave function, Ψ_{AS}, takes very small values over most of configuration space, we use an algorithm that preferentially samples areas with larger values of the wave function (Metropolis). The basic step in such a MC integration moves a single particle, evaluates the ratio $p = (\Psi(R')/\Psi(R))^2$, where R' is the modified configuration and R is the old configuration, and accepts the move with probability $\min(1, p)$. If the move is accepted, we must also update the Slater matrix A (that defines the determinant). In addition, at selected steps we need the inverse of the Slater matrix or its action on several vectors for the computation of certain observables. Finally, the single particle wave functions are often optimized to be (nearly) zero for far away particles, which makes the Slater matrix sparse. However, this is not possible for all materials.

The most important problem is the rapid computation of the ratio of the two Slater determinants. If we move particle k, the matrix changes as $A' = A + e_k u^T$ and this ratio is given by $(1 + u^T A^{-1} e_k)$, where $u = (\psi_1(x'_k) \ldots \psi_N(x'_k))^T - (\psi_1(x_k) \ldots \psi_N(x_k))^T$. We do not need to compute this ratio exactly, but an estimate needs to be unbiased and sufficiently accurate that the convergence of MC does not slow down too much.

Clearly, for sparse matrices, we can get estimates for $A^{-1}e_k$ or $u^T A^{-1} e_k$ in a number of ways, including (sparse) approximations to the inverse, (preconditioned) iterative solvers, and methods to estimate bilinear forms [1, 2]. As we need to do this repeatedly for many moves, efficient algorithms to update various approximations play an important role. In addition, we can exploit the fact that the Slater matrices with large determinant are quite special. If the matrix is not sparse, we must consider alternative methods. For example, we can try to approximate the ratio by computing the determinant ratio for small, selected, principal submatrices that depend on the particle move,

$$|\hat{A}'|/|\hat{A}|. \tag{2}$$

We have observed that this generally leads to sufficiently accurate approximations except for a few steps, which need to be recognized.

In [1] Golub and Bai have proposed several methods for estimating determinants. Unfortunately, in several tests these methods turned out to be too expensive for use in QMC. They did provide favorable scaling for model problems, which might prove useful once very large problems become feasible.

Several important questions in deriving efficient algorithms and proving their effectiveness relate to the Slater matrix itself. (1) Under which conditions can its inverse be approximated sufficiently accurately by a sparse matrix, and how do we update such an approximation efficiently? (2) How sensitive is the (approximate) inverse, or what is the condition number of the Slater matrix? (3) Does the structure of the approximations play a role, and (4) how accurate do we need to approximate the inverse (or solve a linear system)? (5) Since we preferentially sample configurations with large determinants, it is important to characterize Slater matrices for which the determinant is maximum or nearly so (and for wave functions near the ground state). (6) What are the locations of and bounds for the eigenvalues in this case? It turns out that the eigenvalues are more or less uniformly distributed in an annulus centered around the origin. This, obviously, has important implications for iterative solvers. (7) What are good preconditioners in this case? (8) What are the locations of and bounds for the singular values? (9) Can we characterize how eigenvalues and singular values change for single particle updates? Such changes may determine whether we can successfully use Krylov subspace recycling [4, 6] to speed up convergence of iterative methods.

Although there is no formal link with interpolation, the Slater matrix looks like a matrix arising from interpolation, leading to several connections that can be exploited. The accuracy of (2) can be interpreted in terms of the sensitivity of a certain interpolation w.r.t the abscissas, and accurate interpolation in higher dimensions has a certain inherent ill-conditioning. Finally, an important computer graphics problem is to interpolate a point cloud by a surface represented at each point by a few localized basis functions. Hence, the inverse of the interpolation matrix must be accurately approximated by a sparse matrix. The underlying theory is potentially important for the choice of functions to represent the spin orbitals, $\psi_k(.)$.

References

[1] Z. Bai, M. Fahey, and G. H. Golub, *Some large-scale matrix computation problems*, Journal of Computational and Applied Mathematics, 74 (1996), pp. 71–89.

[2] G. Golub and Z. Strakoš, *Estimates in quadratic formulas*, Numerical Algorithms, 8 (1994), pp. 241–268.

[3] B. Hammond, W. L. Jr, and P. Reynolds, *Monte Carlo Methods in Ab Initio Quantum Chemistry*, World Scientific, Singapore, New Jersey, London, Hong Kong, 1994.

[4] M. Parks, E. de Sturler, G. Mackey, D. Johnson, and S. Maiti, *Recycling Krylov subspaces for sequences of linear systems*, SIAM Journal on Scientific Computing, 28 (2006), pp. 1651–1674.

[5] J. M. Thijssen, *Computational Physics*, Cambridge University Press, 2001.

[6] S. Wang, E. de Sturler, and G. Paulino, *Large-scale topology optimization using preconditioned Krylov subspace methods with recycling*, International Journal for Numerical Methods in Engineering, 69 (2007), pp. 2441–2468.

Communication avoiding algorithms for linear algebra

James Demmel, Mark Hoemmen

We survey results on designing numerical algorithms to minimize the largest cost component: communication. This could be bandwidth and latency costs between processors over a network, or between levels of a memory hierarchy; both costs are increasing exponentially compared to floating point. We describe novel algorithms in sparse and dense linear algebra, for both direct methods (like QR and LU) and iterative methods that can minimize communication. A number of these algorithms are optimal in the sense of minimizing communication. This represents joint work with Laura Grigori, Julien Langou, Marghoob Mohiyuddin, and Kathy Yelick.

The Log-Determinant Divergence and Its Applications

Inderjit S. Dhillon

The Log-Determinant (LogDet) Divergence between two $N \times N$ positive definite matrices X and Y is given by:

$$\begin{aligned} D_{\ell d}(X,Y) &= \text{trace}XY^{-1} - \log\det(XY^{-1}) - N, \\ &= \sum_{i=1}^{N}\sum_{j=1}^{N}(\mathbf{v}_i^T\mathbf{u}_j)^2\left(\frac{\lambda_i}{\theta_j} - \log\frac{\lambda_i}{\theta_j} - 1\right), \end{aligned} \tag{1}$$

where $X = V\Lambda V^T$ and $Y = U\Theta U^T$ are their eigendecompositions. The LogDet divergence can be used to measure the distance between two positive definite matrices. Although this measure is not symmetric, it has several intriguing properties, for example, scale invariance:

$$D_{\ell d}(X,Y) = D_{\ell d}(\alpha X, \alpha Y), \quad \alpha \geq 0.$$

In fact, for any invertible M

$$D_{\ell d}(X,Y) = D_{\ell d}(M^T X M, M^T Y M).$$

Equation (1) indicates that the LogDet divergence is determined by the relative magnitude of eigenvalues, thus having a propensity to preserve both large and small eigenvalues when used as a distance measure for matrix nearness problems.

The LogDet divergence has been previously used, mainly as a tool for analysis, in multivariate statistics and numerical optimization. It is called Stein's loss [3] in statistics and has been used to show superlinear convergence of the well known BFGS quasi-Newton method [1, 2]. It also has an intimate connection with the multivariate normal distribution, and distance between covariance matrices.

In this talk, I will discuss recent uses of the LogDet divergence as a computational tool in various applications in data analysis. As formula (1) indicates, the LogDet divergence preserves low-rank structures and conditioning in matrix nearness problems. We have exploited this property to develop fast and effective algorithms for row-lank kernel and metric learning, which are two key problems in semi-supervised data analysis.

References

[1] Richard H. Byrd and Jorge Nocedal. A tool for the analysis of quasi-Newton methods with application to unconstrained minimization. *SIAM J. Numer. Anal.*, 26:727–739, 1989.

[2] R. Fletcher. A new variational result for quasi-Newton formulae. *SIAM Journal on Optimization*, 1(1):18–21, 1991.

[3] W. James and C. Stein. Estimation with quadratic loss. In *Proceedings of the Fourth Berkeley Symposium on Math., Stat. and Prob.*, pages 361–379, 1961.

Sparse direct methods need not suffer under dense rows

Patrick Amestoy, Sue Dollar, John Reid, Jennifer Scott

The efficiency of the analyse phase of sparse direct methods may be dramatically hampered if the sparse matrix being considered contains some rows with a significantly larger number of nonzero entries than that of the majority of the rows. We will discuss how the presence of these (quasi-)dense rows may be overcome and, hence, enable the efficient formation of high quality elimination orderings.

Numerical examples from a variety of applications will reveal that the CPU time needed to form the required ordering can drop by several orders of magnitude but the number of nonzeros in the factors will remain largely unaffected.

The Antireflective Algebra with Applications to Image Deblurring

Antonio Aricò, Marco Donatelli, James G. Nagy, Stefano Serra-Capizzano

We study the properties of the Antireflective algebra introduced in [2]. In particular, using the spectral decomposition of the matrices belonging to such algebra, we will show that it is very useful for implementing filtering methods in image restoration problems [3]. The necessity to define the Antireflective algebra arises from the imposition of antireflective boundary conditions (BCs) to noisy and blurred images introduced by Serra-Capizzano in [1]. Indeed, when the point spread function (PSF) is symmetric with respect to each direction independently (in such case we say that the PSF is strongly symmetric), antireflective BCs lead to matrices belonging to the Antireflective algebra.

We count zero Dirichlet, periodic, reflective (also called Neumann), and anti-reflective BCs (see [5]). In the noise free case, it has been shown that the antireflective BCs produce better restoration, i.e., the ringing effects are negligible with respect to the other BCs [1]. Moreover, from a computational point of view, the antireflective BCs show essentially the same favorable behavior as the reflective BCs. Indeed, despite the apparent complicate expression of the resulting matrix, the product is again possible by two-level fast Fourier transforms, while the solution of a linear system can be obtained by fast sine transforms more lower order computations for strongly symmetric PSFs. In the case of noisy data, imposing antireflective

BCs and applying standard regularizing procedures, the negative surprise was a partial loss of the favorable computational features and, in some cases, a deterioration of the reconstruction quality also with a moderate level of noise [6]. In order to improve the quality of the reconstruction and to recover the efficiency of the related numerical procedures for antireflective BCs, we introduced the re-blurring strategy. The re-blurring essentially replaces the transposition operator with the correlation operator, i.e., the discrete re-blurring operator is obtained imposing the BCs to the PSF flipped along each direction [7]. It is proved that the re-blurring approach is crucial to preserve the algebra structure for least square regularizing methods applied to deconvolution problems with antireflective BCs and strongly symmetric PSFs.

For notational simplicity we describe the Antireflective algebra only in the 1D case. However the multidimensional case can be simply treated by tensor products. Let Q be the discrete sine transform of type I (DST-I) matrix of order n with entries $[Q]_{i,j} = \sqrt{\frac{2}{n+1}} \sin\left(\frac{ij\pi}{n+1}\right)$, $i,j = 1,\dots,n$. It is known that Q is orthogonal and symmetric. For any $\mathbf{v} \in \mathbb{R}^n$, the matrix-vector product $Q\mathbf{v}$ can be computed in $O(n\log(n))$ operations by using the fast sine transform. We define $\tau_n = \{QDQ : D = \mathrm{diag}(\lambda_1,\dots,\lambda_n),\ \lambda_i \in \mathbb{R}\}$ and $\mathcal{C}_l$ the set of real-valued cosine polynomials of degree at most l. The eigenvalues of a τ matrix are given by a $v \in \mathcal{C}_{n-1}$ sampled at $y \in G_n \equiv \left\{\frac{k\pi}{n+1},\ k = 1,\dots,n\right\}$ (i.e., $\lambda_k = v(\frac{k\pi}{n+1})$, $k = 1,\dots,n$). In such case the τ matrix of order n is denoted by $\tau_n(v)$ and it is called the τ matrix generated by the symbol v.

Let $h \in \mathcal{C}_l$, we define the $AR_n(\cdot)$ operator

$$AR_n(h) = \begin{pmatrix} h(0) & & \\ \mathbf{v}_{n-2}(h) & \tau_{n-2}(h) & J\mathbf{v}_{n-2}(h) \\ & & h(0) \end{pmatrix}, \tag{1}$$

where J is the flip matrix, $\mathbf{v}_{n-2}(h) = \tau_{n-2}(\phi(h))\mathbf{e}_1$, with $\mathbf{e}_1$ being the first vector of the canonical basis, and $(\phi(h))(y) = \frac{h(y)-h(0)}{2(\cos(y)-1)}$. The Antireflective algebra of order n is defined as $\mathcal{AR}_n = \{M \in \mathbb{R}^{n\times n} \,|\, M = AR_n(h),\ h \in \mathcal{C}_{n-2}\}$. We observe that a matrix in $\mathcal{AR}_n$ and its spectrum are univocally determined by the associated symbol h like in the circulant algebra case. Moreover the set $\mathcal{AR}_n$ is an algebra since it holds

1. $\alpha AR_n(h_1) + \beta AR_n(h_2) = AR_n(\alpha h_1 + \beta h_2)$,
2. $AR_n(h_1)AR_n(h_2) = AR_n(h_1 h_2)$,

for $\alpha,\beta \in \mathbb{R}$ and for $h_1, h_2 \in \mathcal{C}_{n-2}$. The product $h = h_1 \cdot h_2$ is the unique polynomial that satisfies the interpolation conditions $h(y) = h_1(y)h_2(y)$, for $y \in G_{n-2}$. Unfortunately the Antireflective algebra is not closed for transposition and this gives problems for least square solvers that resort to normal equations. Moreover the matrices in the Antireflective algebra are not normal and hence they can not be diagonalized by unitary transforms. However, it holds

$$AR_n(h) = T_n \mathrm{diag}_{y\in\{0\}\cup G_{n-2}\cup\{0\}}\ (h(y))\ T_n^{-1}, \tag{2}$$

where the eigenvector matrix, called Antireflective transform, is "near" to be unitary:

$$T_n = \left[\, 1 - \frac{y}{\pi}\,,\ \sin(y)\,,\ \cdots,\ \sin((n-2)y)\,,\ \frac{y}{\pi}\,\right]\Big|_{G_{n-2}\cup\{0,\pi\}}. \tag{3}$$

The inverse Antireflective transform T_n^{-1} has a structure analog to T_n and, like the computation of the eigenvalues, it is mainly obtained by DST-I. We note that from (2) and (3) the spectral decomposition of antireflective matrices preserves the classical frequency decomposition property, typical of image restoration problems: large eigenvalues are associated to low frequencies, while small eigenvalues are associated to high frequencies. Indeed, $h(0) = 1$, that thanks to the normalization condition of the PSF is the largest eigenvalue, is associated to a linear function.

Combining the re-blurring approach with the spectral decomposition (2) we can obtain filtering methods analogous to those associated with Tikhonov regularization and truncated singular value decomposition. Denoting by

$$A = AR_n(h) = T_n D_n T_n^{-1}\,, \qquad T_n = \left[\begin{array}{cccc} \mathbf{t}_1 & \mathbf{t}_2 & \cdots & \mathbf{t}_n \end{array}\right], \qquad T_n^{-1} = \left[\begin{array}{c} \tilde{\mathbf{t}}_1^T \\ \vdots \\ \tilde{\mathbf{t}}_n^T \end{array}\right]$$

the spectral decomposition of A, a spectral filter solution is given by

$$\mathbf{f}_{\rm reg} = \sum_{i=1}^{n} \phi_i \frac{\tilde{\mathbf{t}}_i^T \mathbf{g}}{d_i} \mathbf{t}_i\,, \tag{4}$$

where $\mathbf{g}$ is the observed image, $D_n = \mathrm{diag}(d_1, d_2, \ldots, d_n)$ with $d_i = h\left(\frac{(i-1)\pi}{n-1}\right)$, $i = 2, \ldots, n-2$, $d_1 = d_n = h(0)$, and ϕ_i are the filter factors. We recall that for two well known filtering methods, truncated spectral value decomposition (TSVD) and Tikhonov regularization, the filter factors are

$$\phi_i^{\rm tsvd} = \left\{\begin{array}{ll} 1 & \text{if } d_i \geq \delta \\ 0 & \text{if } d_i < \delta \end{array}\right. \qquad \text{and} \qquad \phi_i^{\rm tik} = \frac{d_i^2}{d_i^2 + \lambda}\,,\ \lambda > 0, \tag{5}$$

where the problem dependent regularization parameters δ and λ must be chosen. We note that in (5) we used a slightly abuse of notation since the filtering regularization strategies in (4) implement exactly the re-blurring approach with the Jordan canonical form (2), instead of the classic singular values decomposition. Imposing $\phi_1 = \phi_n = 1$, the solution $\mathbf{f}_{\rm reg}$ in (4) is exactly that obtained by the homogeneous antireflective BCs proposed in [4]. Concluding, by using the Antireflective algebra, in the case of strongly symmetric PSFs it is possible to implement classic filtering methods for the antireflective BCs with the re-blurring approach in a fast and stable way. This is similar to the reflective BCs case, but the obtained restorations are better in terms of the quality of the restored images, as confirmed by a wide numerical experimentation.

References

[1] S. Serra-Capizzano, *A note on anti-reflective boundary conditions and fast deblurring models*, SIAM J. Sci. Comput. 25-3 (2003), pp. 1307–1325.

[2] A. Aricò, M. Donatelli and S. Serra-Capizzano, *Spectral analysis of the anti-reflective algebra*, Linear Algebra Appl., in press.

[3] A. Aricò, M. Donatelli, J. Nagy and S. Serra-Capizzano, *The Anti-Reflective Transform and Regularization by Filtering*, submitted.

[4] M. Christiansen and M. Hanke, *Deblurring methods using antireflective boundary conditions*, SIAM J. Sci. Comput., in press.

[5] P. C. Hansen and J. G. Nagy and D. P. O'Leary *Deblurring Images: Matrices, Spectra, and Filtering*, SIAM, Philadelphia, PA, 2006.

[6] M. Donatelli and S. Serra Capizzano. *Anti-reflective boundary conditions and re-blurring*, Inverse Problems, 21 (2005) pp. 169–182.

[7] M. Donatelli, C. Estatico, A. Martinelli, and S. Serra-Capizzano, *Improved image deblurring with anti-reflective boundary conditions and re-blurring*, Inverse Problems, 22 (2006) pp. 2035–2053.

An Orthogonal and Symmetric High Relative Accuracy Algorithm for the Symmetric Eigenproblem

Froilán M. Dopico, Plamen Koev, Juan M. Molera

The factorization $A = XDX^T$ of a symmetric matrix A is said to be a *rank revealing decomposition* (RRD) of A if the matrix X is well conditioned and D is diagonal and nonsingular. It is well known that a RRD of A determines accurately its eigenvalues and eigenvectors, i.e., small componentwise perturbations of D and small normwise perturbations of X produce small relative variations in the eigenvalues of A, and small variations in the eigenvectors with respect the eigenvalue relative gap [2, 5].

Based on this property, several algorithms have been proposed in the last years to compute *accurate* eigenvalues and eigenvectors of symmetric RRDs. These algorithms allow us to compute accurate eigenvalues and eigenvectors for many classes of symmetric matrices, since RRDs may be accurately computed for many types of matrices [1, 2, 5, 8].

The existing algorithms that compute accurate eigenvalues and eigenvectors of symmetric RRDs are very satisfactory in the positive definite case. In this case, one simply computes accurately the SVD of $X\sqrt{D}$ by using one-sided Jacobi [3, 6, 7]. However, for indefinite symmetric RRDs the available algorithms are not fully satisfactory, although they work well in practice.

The first algorithm for symmetric indefinite RRDs was proposed in [11], and has been carefully analyzed and developed in [9, 10]. The main drawback of this algorithm is that it uses hyperbolic transformations, while it is well known that symmetric matrices can be diagonalized by using only orthogonal transformations. This fact implies that the error bounds for eigenvalues and eigenvectors include some terms that are not guaranteed to be small. The second algorithm was introduced in [4]. The main drawback of this algorithm is that the symmetry of the matrix is not respected in the process.

The purpose of this work is to present the first algorithm that computes accurate eigenvalues and eigenvectors of symmetric indefinite RRDs *using only orthogonal transformations and respecting the symmetry of the problem.* This algorithm is very simple: in exact arithmetic, it is essentially the classical Jacobi algorithm with the stopping criterion introduced in [3]. In finite precision, the key point is to apply the Jacobi procedure implicitly to the factors of the RRD. The backward and forward error analyses of the algorithm are straightforward consequences of new theoretical properties of quasi-diagonal symmetric RRDs, and the relative

error bounds for the eigenvalues and eigenvectors are the best ones to be expected from the sensitivity of the problem.

References

[1] J. Demmel, *Accurate singular value decompositions of structured matrices*, SIAM J. Matrix Anal. Appl., 21 (1999), pp. 562-580.

[2] J. Demmel, M. Gu, S. Eisenstat, I. Slapničar, K. Veselić and Z. Drmač, *Computing the singular value decomposition with high relative accuracy*, Linear Algebra Appl. 299 (1999), pp. 21-80.

[3] J. Demmel and K. Veselić, *Jacobi's method is more accurate than QR*, SIAM J. Matrix Anal. Appl., 13 (1992), pp. 1204-1245.

[4] F. M. Dopico, J. M. Molera and J. Moro, *An orthogonal high relative accuracy algorithm for the symmetric eigenproblem*, SIAM J. Matrix Anal. Appl., 25 (2003), pp. 301-351.

[5] F. M. Dopico and P. Koev, *Accurate Symmetric Rank Revealing and Eigendecompositions of Symmetric Structured Matrices*, SIAM J. Matrix Anal. Appl., 28 (2006), pp. 1126-1156.

[6] Z. Drmač, *Accurate computation of the product-induced singular value decomposition with applications*, SIAM J. Numer. Anal., 35 (1998), pp. 1969-1994.

[7] R. Mathias, *Accurate eigensystem computations by Jacobi methods*, SIAM J. Matrix Anal. Appl., 16 (1995), pp. 977-1003.

[8] M. J. Peláez and J. Moro, *Accurate factorization and eigenvalue algorithms for symmetric DSTU and TSC matrices*, SIAM J. Matrix Anal. Appl., 28 (2006), pp. 1173-1198.

[9] I. Slapničar, *Accurate Symmetric Eigenreduction by a Jacobi Method*, PhD Thesis, Fachbereich Mathematik Fernuniversität, Gesamthochschule Hagen, Germany, 1992.

[10] I. Slapničar, *Highly accurate symmetric eigenvalue decomposition and hyperbolic SVD*, Linear Algebra Appl., 358 (2003), pp. 387-424.

[11] K. Veselić, *A Jacobi eigenreduction algorithm for definite matrix pairs*, Numer. Math., 64 (1993), pp. 241-269.

The Column Subset Selection Problem

Christos Boutsidis, Michael W. Mahoney, Petros Drineas

The topic of selecting a small number of "representative" columns (or rows) of an $m \times n$ matrix A has attracted significant attention in the Numerical Linear Algebra and, much more recently, in the Theoretical Computer Science community. In particular, it is well-known [7] that the best rank-k approximation of A may be expressed as

$$A_k = U_k U_k^T A,$$

where U_k is an $m \times k$ matrix containing the top k left singular vectors of A. Notice that A_k is an $m \times n$ matrix that minimizes $\|A - D\|_{2,F}$ over all matrices D of rank at most k. (Here $\|\cdot\|_2$ and $\|\cdot\|_F$ denote the spectral and Frobenius norms of a matrix respectively.) In words, the "best" rank-k approximation is obtained by projecting a matrix on the subspace spanned by its top k left singular vectors.

A natural question is whether the top k left singular vectors could be replaced by k carefully chosen columns of the matrix A without too much loss in accuracy. In particular, researchers have considered the following Column Subset Selection Problem (CSSP):

Definition 1. *Given an $m \times n$ matrix A and a positive integer k, pick k columns of A forming an $m \times k$ matrix C such that the residual*

$$\|A - P_C A\|_\xi ,$$

is minimized over all possible $\binom{n}{k}$ choices for the matrix C. Here P_C denotes the projection onto the k-dimensional space spanned by the columns of C and $\xi = 2$ or F denotes the spectral norm or Frobenius norm.

Thus, the CSSP is to find a subset of exactly k columns of A that "captures" as much of A as possible, with respect to the spectral and/or Frobenius norm, in a projection sense. This problem has been studied extensively in numerical linear algebra and recently it has been studied in theoretical computer science. Algorithms for this problem have found applications in, e.g., scientific computing [2], as well as data mining and information retrieval; see [11, 9, 4] and references therein.

In our work (see [1] for a much expanded version of this abstract) we present a novel two-stage algorithm for this problem. In the first stage of our algorithm, we randomly select $O(k \log k)$ columns of V_k^T, i.e., of the transpose of the $n \times k$ matrix consisting of the top k right singular vectors of A, according to a judiciously-chosen probability distribution that depends on information in the top-k right singular subspace of A. Then, in the second stage, we apply a deterministic column-selection procedure to select exactly k columns from the set of columns of V_k^T selected by the first phase. The algorithm then returns the corresponding k columns of A. For our main algorithm and its analysis, we use an algorithm described by Pan [10] in the second stage. We prove the following theorem.

Theorem 1. *There exists an algorithm that approximates the solution to the CSSP and takes as input an $m \times n$ matrix A of rank ρ and a positive integer k; runs in $O(\min\{mn^2, m^2 n\})$*

time; and returns as output an $m \times k$ matrix C consisting of exactly k columns of A such that with probability at least $1 - 10^{-20}$:

$$\begin{aligned} \|A - P_C A\|_2 &\leq O\left(k^{\frac{3}{4}} \log^{\frac{1}{2}}(k)\,(\rho - k)^{\frac{1}{4}}\right) \|A - A_k\|_2 \\ \|A - P_C A\|_F &\leq O\left(k\sqrt{\log k}\right) \|A - A_k\|_F\,. \end{aligned}$$

P_C denotes a projection onto the column span of the matrix C, and A_k denotes the best rank-k approximation to the matrix A.

Solving the CSSP exactly is a hard combinatorial problem, and thus research has historically focused on computing approximate solutions to the CSSP. Since $\|A - A_k\|_\xi$ provides an immediate lower bound for $\|A - P_C A\|_\xi$, for $\xi = 2, F$ and for any choice of C, a large number of approximation algorithms have been proposed to select a subset of k columns of A such that the resulting matrix C satisfies

$$\|A - A_k\|_\xi \leq \|A - P_C A\|_\xi \leq p(k,n)\,\|A - A_k\|_\xi$$

for some function $p(k,n)$. Within the numerical linear algebra community, most of the work on the CSSP has focused on spectral norm bounds and is related to the so-called Rank Revealing QR (RRQR) factorization. In particular, the work of Golub on pivoted QR factorizations [6] was followed by much research addressing the problem of constructing an efficient RRQR factorization. It is rather straightforward to observe that any algorithm that constructs an RRQR factorization of the matrix A with provable guarantees also provides provable guarantees for the CSSP; see Foster and Liu [5], who provide a survey and an empirical evaluation of some of these algorithms.

For $\xi = 2$, the best previously-known bound for the CSSP is $p(k,n) = \sqrt{1 + k(n-k)}$ [8], while for $\xi = F$, the best bound is $p(k,n) = \sqrt{k+1}$ [3]. Note that the former result is algorithmic, in that there exists an efficient (not exponential in k) algorithm to compute C, whereas the latter result is purely existential. Thus, our approach asymptotically improves the best previously-known result for the spectral norm version of the CSSP by a factor of $n^{1/4}$. (Here we assume that k is a small constant, as is typically the case in data applications of these techniques.) Our approach also provides an algorithmic bound for the Frobenius norm version of the CSSP that is only $O(\sqrt{k \log k})$ worse than that best previously-known existential result. A novel feature of the algorithm that we present in this paper is that it combines in a nontrivial manner recent algorithmic developments in the theoretical computer science community with more traditional techniques from the numerical linear algebra community in order to obtain improved bounds for the CSSP.

References

[1] http://www.cs.rpi.edu/~boutsc/CSSProblem/

[2] T.F. Chan and P.C. Hansen. Some applications of the rank revealing QR factorization. *SIAM J Sci and Stat Comp*, 13:727–741, 1992.

[3] A. Deshpande, L. Rademacher, S. Vempala, and G. Wang. Matrix approximation and projective clustering via volume sampling. *SODA*, 2006

[4] P. Drineas, I. Kerenidis, and P. Raghavan. Competitive recommendation systems, *Symposium on Theory of Computing*, 2002.

[5] L.V. Foster and Xinrong Liu. Comparison of rank revealing algorithms applied to matrices with well defined numerical ranks. *manuscript*, 2006.

[6] G. H. Golub. Numerical methods for solving linear least squares problems. *Numer Math*, 7:206–216, 1965.

[7] G.H. Golub and C.F. Van Loan. *Matrix Computations*. Johns Hopkins University Press, Baltimore, 1989.

[8] M. Gu and S.C. Eisenstat. Efficient algorithms for computing a strong rank-revealing QR factorization. *SIAM J Sci Comp*, 17:848–869, 1996.

[9] M.W. Mahoney, M. Maggioni, and P. Drineas. Tensor-CUR decompositions and data applications, *SIAM J. Matrix Anal Appl*, in press.

[10] C. T. Pan. On the existence and computation of rank-revealing LU factorizations. *Linear Algebra Appl*, 316:199–222, 2000.

[11] J. Sun, Y. Xie, H. Zhang, and C. Faloutsos. Less is more: Compact matrix decomposition for large sparse graphs, *SIAM Conference on Data Mining*, 2007.

Practical Rank Determination for Square and Rectangular Sparse Matrices

Iain Duff

We study the use of techniques for the practical computation of the rank of a sparse matrix and the selection of a corresponding set of linearly independent rows and columns. We will consider the computation of a basis for the null space of a sparse matrix.

The application area that prompted our recent work was in limit analysis calculations from geotechnical engineering, stimulated by a visit to Lyamin and Sloan at the University of Newcastle in New South Wales. Limit analysis is a very powerful tool for estimating the maximum load capacity of geostructures such as tunnels, retaining walls, dams, and building foundations.

Here the need is to solve the problem

$$\begin{pmatrix} \mathbf{H} & \mathbf{A^T} & \mathbf{G^T} \\ \mathbf{A} & \mathbf{0} & \mathbf{0} \\ \mathbf{G} & \mathbf{0} & \mathbf{F} \end{pmatrix} \begin{pmatrix} \mathbf{d} \\ \mu \\ \lambda \end{pmatrix} = \begin{pmatrix} \mathbf{c} \\ \mathbf{0} \\ \mathbf{0} \end{pmatrix}$$

where the rectangular matrix $\mathbf{A}$ for the equality constraints is rank deficient, often significantly so.

In two dimensions, the rank deficiencies can be removed before assembly but this is difficult and potentially very costly for three dimensions. The engineers would therefore like to reduce the matrix $\mathbf{A}$ to a smaller (square) matrix of full rank.

Lyamin and Sloan were using the HSL multifrontal code MA57 to solve the augmented system and their familiarity with this encouraged them to determine the rank and the independent rows and columns by using MA57 on $\mathbf{AA^T}$.

However, the effect of cancellation and a huge number of very small entries caused problems for the multifrontal solver particularly because their need to identify linearly independent rows and columns led to the factorization being conducted on the larger of the two matrices $\mathbf{AA^T}$ and $\mathbf{A^TA}$.

We show, in Table 3, results from a run of the original MA57 code on the matrix AA^T from one of their test matrices where the matrix A is 17424 × 9410 with 30143 entries. Note that, as we increase the value which we require for a legitimate pivot, the factorization has increasing difficulty in choosing pivots resulting in more work and storage until it fails for lack of the latter. Note also that as the largest the rank could be is 9410, we are far from computing it correctly.

Code	Tol	Time (secs)	Entries in factors	# delayed	max frt	rank
MA57	10^{-20}	5.7	1530839	12009	1450	16315
	10^{-16}	55.2	2541906	28118	3409	14315
	10^{-15}		Insufficient storage			

Table 3: Runs of MA57 on DEC 650 on Lyamin and Sloan test problem.

Our first solution was to modify MA57 to flush very small entries from the frontal matrices by removing immediately any unpivoted fully summed blocks all of whose entries are less than a tolerance (Tol). These entries are ordered immediately to the end, are removed from the front, the factorization is continued. In this case, we see that, as the dropping parameter is increased, more entries are dropped resulting in a much more efficient factorization. The rank of 9409 obtained with a drop tolerance of 10^{-12} has been verified independently by Lyamin and Sloan as being "correct" for this test problem.

We have incorporated code using this dropping strategy in the new release of MA57 that has been included in HSL 2007. We plan to examine the use of this code on large scale three dimensional problems from Lyamin and Sloan and also to examine the possible use of this code as a preconditioner.

Code	Tol	Time (secs)	Entries in factors	# delayed	max frt	rank
Analyse		0.0	787347	-	319	-
MA57m	10^{-20}	2.1	836583	2868	350	16296
	10^{-12}	0.3	381568	8015	319	9409

Table 4: Run of modified version of MA57 on test problem.

We compared this with a direct approach of using a rook-pivoting version (RK48) of the (non-multifrontal) HSL unsymmetric code MA48 directly on the rectangular system.

This required substantial modification MA48 so that whenever a row is encountered we immediately discard any row in the reduced matrix with all entries less than the tolerance. These rows, called "zero rows" in Table 5 are immediately ordered to the end and we continue with the factorization. As in the case of the modified version of MA57 a tolerance of 10^{-12} is used to determine the "zero rows". We see from the results in Table 5 that the modified

# rows		2295	6143	6911	17424	26064	41472	89999
# columns		1990	1090	2402	9410	23232	38015	14643
# entries		11100	9214	10366	30143	129184	141118	119998
Time for analysis								
	MA48	0.05	2.57	3.39	180.5	7.7	67.1	163.1
	RK48	0.03	0.02	0.06	0.7	2.4	0.7	3.5
"Zero" rows								
	RK48	301	5020	4469	8007	2672	5087	75350
Computed rank								
	MA48	1990	1089	2401	9409	23228	35455	14642
	RK48	1990	1089	2401	9409	23228	35455	14642

Table 5: Modified version of MA48 used to determine rank and independent rows and columns. Runs are on a Dell 650

code substantially outperforms the original MA48 and in all cases, as independently verified by Lyamin and Sloan, obtains the required rank.

We have redesigned the modified code as a new routine MC58 that will be included in HSL 2007 and we plan to experiment with this on larger problems from Lyamin and Sloan and to use it in identifying a nonsingular block of the rectangular constraint matrix in a saddle-point matrix formulation for constrained optimization problems.

We will also study the effects of three different kinds of scaling techniques in this context and will compare these approaches with work with my French colleagues in Toulouse on determining a null-space basis in a multifrontal approach by identifying dependent rows and columns and ordering them to the root node where robust dense algorithms can be used for rank and null space basis determination.

Fast Linear Algebra is Stable

James Demmel, Ioana Dumitriu, Olga Holtz

In our paper "Fast Multiplication is Stable" (Jim Demmel, Ioana Dumitriu, Olga Holtz, Robert Kleinberg, *Numerische Mathematik*, **106**(2)(2007), 199-224) we showed that a large class of fast recursive matrix multiplication algorithms is stable in a normwise sense, and that in fact if multiplication of n-by-n matrices can be done by any algorithm in $O(n^{\omega+\eta})$ operations for any $\eta > 0$, then it can be done stably in $O(n^{\omega+\eta})$ operations for any $\eta > 0$.

This talk will describe recent results extending the above to show that essentially all standard linear algebra operations, including LU decomposition, QR decomposition, linear equation solving, matrix inversion, solving least squares problems, (generalized) eigenvalue problems and the singular value decomposition can also be done stably (in a normwise sense and with high probability) in $O(n^{\omega+\eta})$ operations.

Mathematical Software Going Forward

Alan Edelmann

When the Householder meetings and precursors began, it seems unlikely that anyone could have imagined the extensive use of mathematical software as we know it today. In addition to the research from our community, the incredible speeds of modern computers at commodity prices is the packaging of mathematical software. Our community's research has been showcased often first in MATLAB(R), and this has been a win-win for our community and the MathWorks. In addition to MATLAB, same generation proprietary software such as Mathematica and Maple, new generation proprietary software such as Labview(R) and PVWave,(R) and new generation open source software such as Octave, Scilab, Python and R and their packages have made the landscape increasingly interesting. In this talk I will show some of the beginnings of a comparison of software languages and the beginnings of some thinking as to what we as a community might ask for going forward. This work is in collaboration with the Numerical Mathematics Consortium. Some questions are how accurate should a computation be? How careful should we be in disclosing such accuracy? Can we put together a "Consumer Reports" style review of numerical algorithms? What about a "Rosetta Stone" to move between packages and to understand the "Faux Amis" or false cognates between languages. I like to think I'm inspired by much of Velvel Kahan's life work, but I wouldn't want to promise he'd agree with everything I might say. Years ago I jokingly mentioned that in software manuals I wanted on page 2 the command to quit which was usually very hard to find in those days. Today I want to know how to enter the matrix [3 1;4 2] and invert it without reading tons of pages of manuals. I'd like to know if it's okay that beta(1,x) in MATLAB yields 1 for very large x even if the correct answer is exactly 1/x. Is it okay that octave gives $log2(2^n) - n$ as nonzero? That Cholesky is sometimes upper and sometimes lower triangular? And nearly everyone is sloppy with empty arrays. Also while our community has for a long time been thinking about linear algebra accuracy for parallelism, what about everything else?

Solving Ill-Posed Cauchy Problems using Rational Krylov Methods

Lars Eldén, Valeria Simoncini

Let Ω be a connected domain in R^2 with smooth boundary $\partial\Omega$, and assume that L is a linear, self-adjoint, and negative definite elliptic operator in defined in Ω. We will consider the ill-posed Cauchy problem,

$$\begin{aligned} &u_{zz} + Lu = 0, \qquad (x,y,z) \in \Omega \times [0, z_1], \\ &u(x,y,z) = 0, \qquad (x,y,z) \in \Omega \times [0, z_1], \\ &u(x,y,0) = g(x,y), \qquad (x,y) \in \Omega, \\ &u_z(x,y,0) = h(x,y), \qquad (x,y) \in \Omega. \end{aligned}$$

The problem is to determine the values of u on the upper boundary,
$f(x,y) = u(x,y,z_1),\ (x,y) \in \Omega.$

This is an ill-posed problem in the sense that the solution (if it exists), does not depend continuously on the data. It is a variant of a classical problem considered originally by Hadamard.

Since the domain is cylindrical with respect to z, we can use a separation of variables approach, and write the solution formally as

$$u(x,y,z) = \exp(Az) \begin{pmatrix} g \\ h \end{pmatrix} \qquad A = \begin{pmatrix} 0 & I \\ -L & 0 \end{pmatrix}$$

where I is the identity operator. The exponential function can be expressed in terms of the eigenvalue expansion of L. Due to the fact that L is unbounded, the operator A is unbounded, and, in the computation of $\exp(Az)$, any data errors or rounding errors would be blown up, leading to a meaningless approximation of the solution.

The problem can be stabilized (regularized) if the operator L is replaced by a bounded approximation. In previous papers [3, 4] this approach has been used for another ill-posed Cauchy problem, the sideways heat equation, where wavelet and spectral methods were used for approximating the time derivative. A similar procedure for a Cauchy problem for the Laplace equation was studied in [2]. It is not clear to us how a spectral or wavelet approximation of derivatives can be used in cases when the domain is complicated so that, e.g., a finite element procedure is used for the numerical approximation of the 2-D operator L.

One common approach for solving this type of Cauchy problems is to define an energy functional, and minimize that using an iterative procedure, where in each iteration a 3D elliptic (well-posed) problem is solved [1]. In this talk we will describe an iterative procedure with fast convergence, where in each iteration we only need to solve a 2D elliptic (well-posed) problem.

Naturally, since it is the large eigenvalues of L (those that tend to infinity) that are associated with the ill-posedness, a regularization method can be devised that is based on computing the smallest eigenvalues of L, and discarding the components of the solution corresponding to large eigenvalues. Such a method will be compared to the one that is the main contribution of this paper: We will derive a method for approximating

$$u(x,y,z) = \exp(Az) \begin{pmatrix} g \\ h \end{pmatrix},$$

based on projection onto a subspace generated by a rational Krylov sequence involving the operator A^{-1}, perhaps shifted. In each step of the Krylov method we need to solve a standard, well-posed two-dimensional elliptic problem involving L. This means that a standard (black box) 2D elliptic solver can be used. The exponential function of the restriction of the operator A onto a low-dimensional subspace will computed.

The convergence of the approximation of the exponential function will be discussed, and a criterion for terminating the Krylov sequence will be given. Numerical experiments show that the Krylov space method requires far less work (2D elliptic solves) than the method based on computing the eigenvalue expansion of L.

We also demonstrate that the method can be applied to inverse problems for parabolic equations and for the Helmholtz equation.

References

[1] S. Andrieux, T. N. Baranger, and A. Ben Abda. Solving Cauchy problems by minimizing an energy-like functional. *Inverse Problems*, 22:115–133, 2006.

[2] F. Berntsson and Lars Eldén. Numerical solution of a Cauchy problem for the Laplace equation. *Inverse Problems*, 17:839–854, 2001.

[3] T. Regińska and L. Eldén. Solving the sideways heat equation by a wavelet-Galerkin method. *Inverse Problems*, 13:1093–1106, 1997.

[4] T. Regińska and L. Eldén. Stability and convergence of a Wavelet-Galerkin method for the sideways heat equation. *J. Inverse Ill-Posed Problems*, 8:31–49, 2000.

Preconditioners for Eigenvalue Problems Arising in Linear Stability Analysis

Howard C. Elman

This talk concerns computational methods for solving generalized eigenvalue problems of the type that arise in stability analysis of models of incompressible flow. A main source of these problems is the incompressible Navier-Stokes equations

$$\begin{aligned} &\alpha \mathbf{u}_t - \nu \nabla^2 \mathbf{u} + (\mathbf{u} \cdot \text{grad})\mathbf{u} + \text{grad}\, p = \mathbf{f} \\ &-\text{div}\, \mathbf{u} = 0 \end{aligned} \qquad \text{on } \mathcal{D} \subset \mathbf{R}^d,\ d = 2 \text{ or } 3.$$

Given a solution $\begin{pmatrix} \mathbf{u} \\ p \end{pmatrix}$ to the steady problem ($\alpha = 0$ above), the issue is the stability of this solution with respect to perturbations. *Linear stability analysis* treats perturbations of the form $e^{\lambda t}\begin{pmatrix} \mathbf{v} \\ q \end{pmatrix}$, which either grow or decay exponentially in time. This leads to an eigenvalue problem

$$\begin{pmatrix} \mathcal{L}\mathbf{v} + \text{grad}\ p \\ -\text{div}\ \mathbf{v} \end{pmatrix} = \lambda \begin{pmatrix} \mathcal{M}\mathbf{v} \\ 0 \end{pmatrix}, \tag{1}$$

where the operators $\mathcal{L}$ and $\mathcal{M}$ are

$$\mathcal{L}\mathbf{v} = -\nu \nabla^2 \mathbf{v} + (\mathbf{u} \cdot \text{grad})\mathbf{v} + (\mathbf{v} \cdot \text{grad})\mathbf{u}, \quad \mathcal{M}\mathbf{v} = -\mathbf{v}.$$

The problems are parameterized by the *Reynolds number* R, which is inversely proportional to ν, and the aim is to find the largest Reynolds number R_L for which the rightmost eigenvalue of (1) has negative real part. For $R > R_L$, the steady solution is unstable with respect to at least one infinitesmal disturbance.

Discretization of (1) yields a real algebraic eigenvalue problem of the form

$$\begin{pmatrix} A & B^T \\ B & 0 \end{pmatrix} \begin{pmatrix} v \\ q \end{pmatrix} = \lambda \begin{pmatrix} M & 0 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} v \\ q \end{pmatrix}, \quad \text{or} \quad \mathcal{A}x = \lambda \mathcal{B}x. \tag{2}$$

Eigenvalues of small real part are sought, and efficient eigenvalue algorithms will require the solution of linear systems of equations in which the coefficient matrix is (a possibly shifted version of) the matrix on the left of (2),

$$\mathcal{A}_\theta = \begin{pmatrix} A - \theta M & B^T \\ B & 0 \end{pmatrix} = \mathcal{A} - \theta \mathcal{B}. \tag{3}$$

The key to development of effective solution algorithms is to combine good iterative eigenvalue solvers with effective preconditioners for (3).

We will describe three approaches to these questions.

- A technique called *tuning*, developed by Freitag and Spence [3] can be used to enhance the effect of preconditioners designed for $\mathcal{A}_0 = \mathcal{A}$ when they are applied in the setting of eigenvalues. Suppose inverse iteration is used for (2), leading to a sequence of systems of the form

  $$(\mathcal{A} - \theta_k \mathcal{B}) y^{(k)} = \mathcal{B} x^{(k)}. \tag{4}$$

  Given a preconditioner $\mathcal{Q}_\mathcal{A}$ for $\mathcal{A}$, let

  $$\hat{\mathcal{Q}}_\mathcal{A} = \mathcal{Q}_\mathcal{A} + (w^{(k)} - \mathcal{Q}_\mathcal{A} x^{(k)}) x^{(k)T},$$

  where $w^{(k)}$ is a "tuning" vector to be specified. If $w^{(k)} = \mathcal{A} x^{(k)}$, then $\mathcal{B} x^{(k)}$ tends to an approximate eigenvector of the preconditioned matrix $(\mathcal{A} - \theta \mathcal{B}) \hat{\mathcal{Q}}_\mathcal{A}^{-1}$. That is, as the inverse iterate $x^{(k)}$ tends to an eigenvector of (2), the preconditioner used within this computation is adapted so that the right-hand side in the system *simultaneously* converges to an eigenvector of the preconditioned matrix. This enhances the convergence of the linear system solver for (4).

- As an eigenvalue iteration such as inverse iteration proceeds, estimates for the eigenvalues of the preconditioned matrix $(\mathcal{A} - \theta_{k-1} \mathcal{B}) \hat{\mathcal{Q}}_\mathcal{A}^{-1}$ can be computed and used to find polynomials that implicitly filter out unwanted components of the right-hand side of (4) in later steps. This will further speed convergence of Krylov subspace methods and offset a decline in performance often seen for inner iterations during the latter stages of eigenvalue iterations [5].

- The discrete divergence free condition $Bv = 0$ implies that the augmented problem

  $$\begin{pmatrix} A + B^T W^{-1} B & B^T \\ B & 0 \end{pmatrix} \begin{pmatrix} v \\ q \end{pmatrix} = \lambda \begin{pmatrix} M & 0 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} v \\ q \end{pmatrix}$$

  has the same solution as (2). Olshanskii and Benzi [1, 4] have shown that addition of the *augmented Lagrangian* $B^T W^{-1} B$ leads to improved performance of preconditioned Navier-Stokes solvers; in the context of eigenvalue iterations, the shift tends to make the system (4) harder to solve, but the augmenting term appears to counteract this difficulty. We have seen that the choice of the scaling matrix W has some impact on this question.

We show that these ideas can be combined with efficient preconditioning methods developed for the linearized Navier-Stokes equations [2], and we will describe their impact on the performance of solution algorithms for stability analysis of incompressible flows.

References

[1] M. Benzi and M. A. Olshanskii. An augmented Lagrangian-based approach to the Oseen problem. *SIAM J. Sci. Comput.*, 28:2095–2113, 2006.

[2] H. C. Elman, D. J. Silvester, and A. J. Wathen. *Finite Elements and Fast Iterative Solvers: with Applications in Incompressible Fluid Dynamics.* Oxford University Press, Oxford, UK, 2005.

[3] M. A. Freitag and A. Spence, Convergence rates for inexact inverse iteration with application to preconditioned iterative solves *BIT*, 47:27–44, 2007.

[4] M. A. Olshanskii and M. Benzi. An augmented Lagrangian Approach to Linearized Problems in Hydrodynamic Stability. Technical Report, Emory University Department of Mathematics and Computer Science, 2007.

[5] V. Simoncini and L. Eldén. Inexact Rayleigh quotient-type methods for eigenvalue computations. *BIT*, 42:159–182, 2002.

Damped Mechanical Systems: Spectra, Pseudospectra, Structured Perturbations

Steven J. Cox, Mark Embree, Jeffrey M. Hokanson

A variety of vibrating mechanical systems are modeled by the second order differential equation

$$\mathbf{M}\ddot{\mathbf{x}} + \mathbf{G}\dot{\mathbf{x}} + \mathbf{K}\mathbf{x} = \mathbf{0},$$

with positive definite mass and stiffness matrices $\mathbf{M}$ and $\mathbf{K}$. Such systems arise from discrete networks or the discretization of damped wave equations. The evolution of $\mathbf{x}(t)$ is analyzed via the quadratic eigenvalue problem

$$\lambda^2\mathbf{M}\mathbf{x} + \lambda\mathbf{G}\mathbf{x} + \mathbf{K}\mathbf{x} = \mathbf{0},$$

which has been the object of considerable study over the past fifty years; see, e.g., [2, 4]. In the absence of damping ($\mathbf{G} = \mathbf{0}$), such systems are skew-adjoint in the physically-relevant inner product, and energy is conserved. One would like to design a damping matrix $\mathbf{G}$ that is optimal within some prescribed class, in the sense that it minimizes the spectral abscissa or greatest total energy.

This talk shall first address a particular instance of this problem, arising from a model of a wire vibrating in a magnetic field. This system is governed by the integro-differential equation

$$u_{tt}(x,t) = u_{xx}(x,t) - a(x)\int_0^\pi a(\xi)u_t(\xi,t)\,d\xi$$

for $x \in [0,\pi]$ with homogeneous Dirichlet boundary conditions. Though this damping, whose discretization gives a rank-1 matrix $\mathbf{G}$, is considered rather weak (in that the continuous problem has eigenvalues arbitrarily close to the imaginary axis for typical choices of a), we shall illustrate how particular singular a damp the spectrum quite effectively [1].

Does moving the spectrum far into the left half-plane correspond to rapid energy decay for the dynamical system? Pseudospectra provide a natural tool for investigating this question. We shall explore several alternatives, including a direct analysis of the quadratic matrix polynomial [5] and various approaches to the linearized problem

$$\begin{bmatrix} \mathbf{0} & \mathbf{I} \\ -\mathbf{K} & -\mathbf{G} \end{bmatrix} \begin{bmatrix} \mathbf{x} \\ \mathbf{y} \end{bmatrix} = \lambda \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{M} \end{bmatrix} \begin{bmatrix} \mathbf{x} \\ \mathbf{y} \end{bmatrix}.$$

This provides a natural setting in which to examine the suitability of structured versus unstructured perturbation analysis for describing the behavior of dynamical systems.

We shall close by presenting high-precision physical experiments conducted in our vibrating string laboratory. The spectrum can be estimated from displacement measurements, which we collect from photodetectors. The real parts of the eigenvalues prove particularly difficult to estimate. We use our data to test state-of-the-art methods for this problem from the signal processing and nuclear resonance magnetic imaging communities (see, e.g., [3, 6]). Such methods, including the Hankel Total Least Squares (HTLS) algorithm, prove effective for a limited number of modes. As expected, performance degrades for data in which more eigenvalues are evident (e.g., we measure more than 100 from a vibrating continuous string).

[1] Steven J. Cox, Mark Embree, Jeffrey M. Hokanson, and Walter A. Kelm. "Magnetic damping of an elastic conductor," in preparation.

[2] F. R. Gantmacher and M. G. Krein, *Oscillation Matrices and Kernels and Small Vibrations of Mechanical Systems*, revised ed., AMS Chelsea, Providence, 2002.

[3] Yingbo Hua and Tapan K. Sarkar. "Matrix pencil method for estimating parameters of exponentially damped/undamped sinusoids in noise," *IEEE Trans. Acous., Speech, Sig. Proc.* 38 (1990) 814–824.

[4] Peter Lancaster, *Lambda-Matrices and Vibrating Systems*, Pergamon, Oxford, 1966.

[5] Françoise Tisseur and Nicholas J. Higham, "Structured pseudospectra for polynomial eigenvalue problems, with applications," *SIAM J. Matrix Anal. Appl.* 23 (2001) 187–208.

[6] L. Vanhamme, Tomas Sundin, Paul Van Hecke, and Sabine Van Huffel. "MR spectroscopy quantization: a review of time-domain methods," *NMR Biomed.* 14 (2001) 233–246.

Multilevel Krylov-Multigrid Method for the 2D Indefinite Helmholtz Equation

Yogi A. Erlangga, Reinhard Nabben

In Erlangga and Nabben [2], a class of multilevel method which uses nested Krylov iterations is introduced to solve $Au = b$. The construction of the method, called "Multilevel Krylov Iteration" (MKI), is based on a projection process, which shifts some small eigenvalues of the matrix A to the largest one. This projection results in a projected system, which is better

conditioned than the original system. From an abstract point of view, this projection is done by the action of a preconditioning matrix associated with the generalized Wielandt deflation process in eigenvalue computations. The information of the shifted, small eigenvalues is contained in the Galerkin matrix, associated with the matrix A, which has to be inverted. From the talk of Reinhard Nabben, we have suggested that the convergence of a Krylov method is insensitive to the inexact solve of the Galerkin system. In this respect, the Galerkin system can then be solved only approximately by using (inner) Krylov iterations. To obtain a fast convergence of the inner iteration, the same projection preconditioner with proper dimension is applied to the Galerkin system. By repeating the same procedure in the Galerkin solves, a multilevel, nested iteration results. For 2D Poisson and convection-diffusion problems, the convergence of the method can be made independent of mesh size and physical parameters involved.

In this talk, we extend this method to solve the preconditioned system

$$AM^{-1}\hat{u} = b, \qquad \hat{u} = Mu, \qquad A, M \in \mathbb{C}^{n\times n}, \tag{1}$$

where A is the matrix associated with the discretization of the highly indefinite Helmholtz equation

$$\frac{\partial^2\phi}{\partial x^2} + \frac{\partial^2\phi}{\partial y^2} + k^2\phi = f, \qquad \text{with some boundary conditions,} \tag{2}$$

and M is the preconditioner. The indefinite Helmholtz equation is an example where many (if not all) efficient iterative methods for standard problems become useless. Within the context of Krylov methods, one effective preconditioner for (2) is the matrix associated with the discretization of the operator [4], called "shifted Laplacian":

$$\frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} + (\alpha + \hat{\imath}\beta)k^2, \qquad \alpha, \beta \in \mathbb{R}, \qquad \hat{\imath} = \sqrt{-1}. \tag{3}$$

The resultant preconditioning matrix M is then inverted approximately by using one multigrid iteration. The convergence of a Krylov method is then independent of the gridsize, h, and linearly dependent of wavenumber, k, with a small proportionality constant. For this particular problem and preconditioning, reference [3] provides some spectral bounds. While the maximum eigenvalue is nicely bounded and this bound is independent of the wavenumber k, the small eigenvalues are of order $\mathcal{O}(1/k)$. This small eigenvalues will be handled by projection.

For (1) the projection is defined as

$$P = I - Z\widehat{E}^{-1}Y^TAM^{-1} + \lambda_{\max}Z\widehat{E}^{-1}Y^TAM^{-1}, \tag{4}$$

where $\widehat{E} = Y^TAM^{-1}Z$ is called the Galerkin matrix, $\lambda_{\max}$ the largest eigenvalue (in magnitude) of AM^{-1}, and $Z, Y \in \mathbb{C}^{n\times r}$, $r \ll n$, the proper deflation subspaces. For the Helmholtz equation preconditioned by the shifted Laplacian, we have that $\lambda_{\max}(AM^{-1}) \to 1$ [5, 3].

The immediate complication is that, following [4], the matrix M is approximately inverted by one multigrid iteration. Since in this case the inverse of M is not explicitly known, the Galerkin matrix $\widehat{E} = Y^TAM^{-1}Z$ can not be explicitly determined, and therefore the correct shift of small eigenvalues of AM^{-1} can not be done. Furthermore, even if one chooses to explicitly invert M, such a computation is too expensive, and amounts to the same cost as solving the system $Au = b$ by a direct solver.

In order to construct an MKI algorithm for (1), we propose to approximate the Galerkin matrix $\widehat{E}$ as follows[1]:

$$\widehat{E} \approx (Y^T AZ)(Y^T MZ)^{-1} Y^T Z. \tag{5}$$

In this form, $\widehat{E}$ is approximated by a product of Galerkin matrices associated with A, M, and the identity I. Since $r \ll n$, the product $Y^T MZ =: \hat{M}$ is much easier to invert, and $\widehat{E}$ can therefore be determined. Once the Galerkin matrix $\widehat{E}$ is known, the associated Galerkin system is then solved by a Krylov method. In order to obtain high convergence rate, projection similar to (4) is applied to the Galerkin system, namely

$$\widehat{P} = I - \widehat{Z}\widehat{\widehat{E}}^{-1}\widehat{Y}^T\widehat{A}\widehat{M}^{-1}\widehat{B} + \widehat{\lambda}_{\max}\widehat{Z}\widehat{\widehat{E}}^{-1}\widehat{Y}^T\widehat{A}\widehat{M}^{-1}\widehat{B}, \qquad \widehat{\widehat{E}} = \widehat{\widehat{Y}}^T\widehat{E}\widehat{\widehat{Z}}, \tag{6}$$

with $\widehat{A} = Y^T AZ$ and $\widehat{B} = Y^T Z$.

The advantage of the approximation (5) now becomes apparent.

Since we solve the Galerkin system associated with $\widehat{E}$ with a Krylov method, the inverse of $\widehat{E}$ needs not be computed explicitly. Instead, $\widehat{M}$ and $\widehat{B}$ can be considered as another form of preconditioning to $\widehat{A}$. The multiplication with $\hat{B}$ can be carried out easily, and the inversion of $\widehat{M}$ can be established via a multigrid iteration. Since $\widehat{M}$ is the coarse-grid approximation to M, the same multigrid components used for M are also used for $\widehat{M}$, but for $\widehat{M}$ with one multigrid level reduced. This implies that the deflation matrices Z and Y resemble the interpolation matrix in multigrid. Hence, because we use multigrid to invert M, no extra information is needed for Z and Y. With this setting of Z and Y, $\widehat{A}$ is a coarse-grid approximation to A.

By repeating the same procedure to the Galerkin system associated with $\widehat{\widehat{E}}$ in (6), we obtain an MKI algorithm. On every level, a multigrid iteration of reduced level is run to approximately invert the associated shifted Laplacian matrix.

We will show numerical results from solving the 2D high wavenumber Helmholtz equation with this method, and compare them with results based on some ILU preconditioners. For the MKI method, the convergence can be made practically independent of the wavenumber and mesh size. With an increase in wavenumber, gains in CPU time become significant.

While in this talk we are merely focusing on the Helmholtz equation, this approach can be regarded as the framework for general cases with preconditioner.

This work is financially supported by the *Deutsche Forschungsgemeinschaft* (DFG), Project Number NA248/2-2.

References

[1] Y.A. Erlangga, R. Nabben, *On a multilevel Krylov method for the Helmholtz equation preconditioned by shifted Laplacian*, 2007, submitted.

[2] Y.A. Erlangga, R. Nabben, *Multilevel projection-based nested Krylov iteration for boundary value problems*, 2007, to appear in SIAM J. Sci. Comput.

[3] M.B. van Gijzen, Y.A. Erlangga, C. Vuik, *Spectral analysis of the discrete Helmholtz operator preconditioned with a shifted Laplacian*, SIAM J. Sci. Comput. 29 (5), 2007, pp. 1942-1958.

[4] Y.A. Erlangga, C.W. Oosterlee, C. Vuik, *A novel multigrid based preconditioner for heterogeneous Helmholtz problems*, SIAM J. Sci. Comput. 27, 2006, pp. 1471-1492.

[5] Y.A. Erlangga, C. Vuik, C.W. Oosterlee, *On a class of preconditioners for solving the Helmholtz equation*, Appl. Num. Math., 50, 2004, pp. 409-425.

A Posteriori Error Estimators for Krylov Subspace Approximations of Matrix Functions

Oliver G. Ernst, Michael Eiermann, Stefan Güttel

The evaluation of a matrix function times a vector $f(A)\mathbf{b}$ is a computational task that occurs frequently in applications, the most well-known instance being the use of the matrix exponential in exponential integrators for the numerical solution of initial value problems. Building on our recent work [3, 1, 2] on restarted Krylov subspace approximations of $f(A)\mathbf{b}$, we give an overview of how a posteriori error estimators may be constructed for such approximations. Such error estimators are crucial for devising reliable stopping criteria (see e.g. [4, 5]).

The techniques we consider are based on contour integration and Gaussian quadrature and, alternatively, on interpolation series. The latter are generalizations of Saad's *corrected approximation* [7] and the *k-extended schemes* of Philippe and Sidje [6]. We also present improved estimates which can be derived for the class of Stieltjes functions. Finally, we explore the relationship of our estimators with those of Strakos and Tichy [8] for solving linear systems of equations.

References

[1] M. Afanasjew, M. Eiermann, O. G. Ernst, and S. Güttel. Implementation of a restarted Krylov subspace method for the evaluation of matrix functions. *(submitted)*, 2007.

[2] M. Afanasjew, M. Eiermann, O. G. Ernst, and S. Güttel. The method of steepest descent for matrix functions. *(submitted)*, 2007.

[3] M. Eiermann and O. G. Ernst. A restarted Krylov subspace method for the evaluation of matrix functions. *SIAM Journal on Numerical Analysis*, 44(6):2481–2504, 2006.

[4] A. Frommer and V. Simoncini. Stopping criteria for rational matrix functions of Hermitian and symmetric matrices. Preprint BUW-SC 2007/1, Bergische Universität Wuppertal, 2007.

[5] M. Ilić, I. W. Turner, and D. P. Simpson. A restarted Lanczos approximation to functions of a symmetric matrix. *(submitted to IMA Journal of Numerical Analysis)*, 2007.

[6] B. Philippe and R. B. Sidje. Transient solutions of Markov processes by Krylov subspaces. Research Report, RR-1989, INRIA, 1995.

[7] Y. Saad. Analysis of some Krylov subspace approximations to the exponential operator. *SIAM Journal on Numerical Analysis*, 29(1):209–228, February 1992.

[8] Z. Strakoš and P. Tichý. On error estimation in the conjugate gradient method and why it works in finite precision computations. *ETNA*, 13:56–80, 2002.

A structured polynomial eigenproblem arising in the analysis of time delay systems and related polynomial eigenproblems

Heike Faßbender, Steve Mackey, Nil Mackey, Christian Schröder

We will discuss matrix polynomial eigenproblems of the form

$$Q(\lambda)v = 0$$

where

$$Q(\lambda) = \sum_{i=0}^{k} \lambda^i B_i, \qquad B_k \neq 0 \tag{1}$$

with $B_i \in \mathbb{C}^{n\times n}, B_i = P\overline{B}_{k-i}P, i = 0, \ldots, k$, for an involutory matrix P, ($P^{-1} = P$). Here $\overline{B}$ denotes the componentwise conjugation of the matrix entries. With

$$\overline{Q}(\lambda) = \sum_{i=0}^{k} \lambda^i \overline{B}_i,$$

and

$$\mathrm{rev}(Q(\lambda)) := \lambda^k Q(\frac{1}{\lambda}) = \sum_{i=0}^{k} \lambda^{k-i} B_i, \tag{2}$$

(1) satisfies

$$Q(\lambda) = P \cdot \mathrm{rev}(\overline{Q}(\lambda)) \cdot P. \tag{3}$$

Eigenproblems as in (1) arise for example in the analysis of time-delay systems. A neutral linear time-delay system (TDS) with m constant delays $h_0 = 0 < h_1 < \ldots < h_m$ is described by

$$\begin{array}{lll} \sum_{k=0}^{m} D_k \dot{x}(t-h_k) & = \sum_{k=0}^{m} A_k x(t-h_k), & t > 0, \\ x(t) & = \varphi(t), & t \in [-h_m, 0], \end{array} \tag{4}$$

with $x : [-h_m, \infty) \to \mathbb{R}^n$ and $A_k, D_k \in \mathbb{R}^{n\times n}$ for $k = 1, \ldots, m$. The stability of a TDS can be determined by the solution of an eigenproblem as in (3), see [1].

Matrix polynomials which satisfy (3) are reminiscent of the different palindromic polynomials defined in [2]:

- palindromic: $\mathrm{rev}(Q(\lambda)) = Q(\lambda)$
- anti-palindromic: $\mathrm{rev}(Q(\lambda)) = -Q(\lambda)$
- $\star$-palindromic: $\mathrm{rev}(Q^\star(\lambda)) = Q(\lambda)$
- $\star$-anti-palindromic: $\mathrm{rev}(Q^\star(\lambda)) = -Q(\lambda)$
- even, odd: $Q(-\lambda) = \pm Q(\lambda)$

- $\star$-even, $\star$-odd: $Q^\star(-\lambda) = \pm Q(\lambda)$

where $\star$ is used as an abbreviation for transpose T in the real case and either T or conjugate transpose $*$ in the complex case. The different palindromic matrix pencils have the property that reversing the order of the coefficient matrices, followed perhaps by taking their transpose or conjugate transpose, leads back to the original matrix polynomial. Palindromic matrix polynomials can be thought of as generalizations of symplectic matrices, because they are strongly related (via the Cayley transformation) to even/odd matrix polynomials which represent generalizations of Hamiltonian matrices. Due to their intimate relationship odd/even and palindromic matrix polynomials can be studied in a similar way.

As will be discussed here, strongly related to the matrix polynomial (3) are matrix polynomials $Q(\lambda)$ with

$$Q(\lambda) = -P \cdot \mathrm{rev}(\overline{Q}(\lambda)) \cdot P, \tag{5}$$

$$Q(\lambda) = P \cdot \overline{Q}(-\lambda) \cdot P, \tag{6}$$

and

$$Q(\lambda) = -P \cdot \overline{Q}(-\lambda) \cdot P. \tag{7}$$

The matrix polynomials considered here are not in the classes of the different palindromic matrices considered in [2], but they have very similar properties.

The classical approach to investigate or numerically solve polynomial eigenvalue problems is linearization. The companion forms provide the standard examples of linearization for a matrix polynomial $Q(\lambda)$. Let $X_1 = X_2 = diag(B_k, I_n, \ldots, I_n)$,

$$Y_1 = \begin{bmatrix} B_{k-1} & B_{k-2} & \cdots & B_0 \\ -I_n & 0 & \cdots & 0 \\ & \ddots & \ddots & \vdots \\ 0 & & -I_n & 0 \end{bmatrix}, \quad \text{and} \quad Y_2 = \begin{bmatrix} B_{k-1} & -I_n & \cdots & 0 \\ B_{k-2} & 0 & \cdots & \\ \vdots & \vdots & \ddots & -I_n \\ B_0 & 0 & \cdots & 0 \end{bmatrix}.$$

Then $C_1(\lambda) = \lambda X_1 + Y_1$ and $C_2(\lambda) = \lambda X_2 + Y_2$ are the first and second companion forms for $Q(\lambda)$. These linearizations do not reflect the structure present in the matrix polynomial Q. Therefore only standard numerical methods can be applied to solve the eigenproblem. In a finite precision environment this may produce physically meaningless results [3], e.g., loss of symmetries in the spectrum. Hence, it is important to construct linearizations that reflect the structure of the given polynomial, and then to develop numerical methods for the resulting linear eigenvalue problem that properly address these structures.

Here we will analyze the spectral symmetry of the structured polynomials introduced, the relationship between them and structured linearizations. In particular, we will discuss a Schur-type decomposition of a generalized eigenvalue problem with PCP structure. Numerical examples will be presented comparing the eigenvalue conditioning of the structured linearizations with that of the companion form ones and a structure preserving eigenvalue method with the standard one.

References

[1] E. Jarlebring, *The spectrum of delay-differential equations: Computational analysis*, PhD thesis, TU Braunschweig, Institute Computational Mathematics, 2008.

[2] D.S. Mackey, N. Mackey, C. Mehl, and V. Mehrmann, *Structured polynomial eigenvalue problems: Good vibrations from good linearizations*, SIAM J. Matrix Anal. Appl., 28 (2006), pp. 1029–1051.

[3] F. Tisseur and K. Merbergen, *A survey of the quadratic eigenvalue problem*, SIAM Review, 43 (2001), pp. 234–286.

Inexact Preconditioned Shift-and-Invert Arnoldi's Method and Implicit Restarts for Eigencomputations

Melina A. Freitag

Consider the computation of a few eigenvectors and corresponding eigenvalues of a large sparse nonsymmetric matrix using the Arnoldi algorithm. In order to compute eigenvalues in an isolated cluster around a given value, shift-and-invert Arnoldi's method with and without implicit restarts may be applied. The shifted linear system arising at each step of Arnoldi's method is solved iteratively using GMRES. The costs of the inexact solves are measured by the number of inner iterations needed by the iterative solver at each outer step of the algorithm.

In this talk it is firstly shown how the relaxation strategy developed by Simoncini [3] can be extended to implicitly restarted Arnoldi's method which yields an improvement in the overall costs of the method.

Secondly, a new preconditioning strategy which was successfully applied in the context of inexact inverse iteration [1,2] is proposed for the inner solver. For inexact shift-and-invert Arnoldi's method a system of the form

$$(A - \sigma I)y = q_k \tag{1}$$

needs to be solved at each step, where q_k denotes the previously calculated basis vector of the Krylov subspace. It was shown in [1,2] that for inexact inverse iteration a preconditioner $\mathbf{P}$ which satisfies $\mathbf{P}x = Ax$ and is just a rank-one change of a standard preconditioner P reduces the total number of inner iterations significantly. This strategy can be applied to shift-invert-Arnoldi's method with and without restarts, so that instead of (1) a system of the form

$$(A - \sigma I)\mathbf{P}_k^{-1}\tilde{y} = q_k, \quad \mathbf{P}_k^{-1}\tilde{y},$$

has to be solved where $\mathbf{P}_k^{-1}$ is a rank-k change of the standard preconditioner and satisfies $\mathbf{P}_k Q_k = AQ_k$, where Q_k contains the basis vectors of the Krylov subspace. It is shown that this small rank change of the preconditioner, called the tuned preconditioner, can produce considerable savings in the total number of iterations. Numerical experiments indicate savings of the total number of iterations between 30 and 50 per cent.

Finally, a comparison of Arnoldi's method with this new tuned preconditioner to the preconditioned Jacobi-Davidson method is presented. Numerical experiments illustrate the theory.

References

[1] M.A. Freitag and A. Spence. Convergence rates for inexact inverse iteration with application to preconditioned iterative solves, BIT, 47 (2007), pp. 27-44.

[2] M.A. Freitag and A. Spence. A tuned preconditioner for inexact inverse iteration applied to Hermitian eigenvalue problems, to appear in IMA J. Numer. Anal.

[3] V. Simoncini. Variable Accuracy of Matrix-Vector Products in Projection Methods for Eigencomputation, SIAM J. Numer. Anal. v. 43, n.3 (2005), pp. 1155-1174.

Algebraic Multigrid Methods for Laplacians of Graphs

Andreas Frommer

The graph partioning problem is at the heart of the parallelization of many computations, the graph G representing a (finite element, e.g.) discretization or, in a more abstract manner, the interdependence of different computational tasks. There are two goals: The work load must be distributed evenly between processors, and the amount of communication induced by the partitioning should me small.

Very recently, as an attempt to improve over standard algorithms as implemented in packages like JOSTLE or METIS, *shape optimizing* graph partitioning heuristics have been proposed, see [1, 3, 4]. The idea is to try to keep the number of boundary vertices small among the partitionings rather than the number of cut edges. This corresponds to the observation that the amount of communication is more likely to be determined by the number of boundary vertices and not by the number of cut edges.

The FOS/C heuristic of [3, 4] is a shape optimizing heuristics where a *diffusive process* – load is distributed iteratively between neighboring vertices – is applied to the graph G. The final distribution of the load then determines the partitioning. Mathematically, the result of this process is the solution of a linear system of the form

$$L_G x = b \tag{1}$$

where L_G is the Laplacian of the graph, a singular diagonally dominant M-matrix with zero row sums, and b is a certain vector orthogonal to $e = (1, \ldots, 1)$.

For the efficiency of the whole shape optimizing approach, it is crucial to be able to solve the system (1) efficiently. This is why, going beyond the multilevel approach in [4] which uses a grid hierarchy for the purposes of 'direct' grid partitioning, we study an algebraic multigrid method for (1). It is based on classical ('Ruge-Stüben') algebraic multigrid ideas with the following ingredients:

- smoothing is via Gauss-Seidel or relaxed Jacobi
- coarse variables are selected as a maximal independent set of vertices in a reduced graph (weakly coupling edges are removed from G)
- interpolation and restriction operators are constructed from L_G
- the coarse variable system is obtained by a Galerkin projection

In this talk, we prove results on the two-level and multilevel convergence of this approach, also addressing the question of bounds on the convergence rate which are independent of the size of the graph G. Our results follow from transporting the theory presented in [5] for non-singular M-matrices to the singular case. An important tool in the analysis is the following result.

Proposition: Let A be symmetric and positive semidefinite and denote $\langle x, y\rangle_A = \langle x, Ay\rangle$. Let $A = M - N$ be a splitting and put $H = I - M^{-1}A$. Assume that

$$e \notin \text{null}(A) \Rightarrow \langle He, He\rangle_A < \langle e, e\rangle_A.$$

Then H is semi-convergent, i.e. $\lim_{k\to\infty} H^k$ exists.

Note that this result has some interest on its own. For example, we get an elegant proof of an old result by Keller [2]: Assume that A is symmetric and positive semidefinite and $M + M^T - A$ is positive definite. Then for all $e \notin \text{null}(A)$

$$\langle He, He\rangle_A = \langle e, e\rangle_A - \underbrace{\langle M^{-1}Ae, (M + M^H - A)\cdot M^{-1}Ae\rangle}_{>0 \text{ since } M^{-1}Ae\neq 0}.$$

So, by the proposition, H is semiconvergent. The proposition also helps in the analysis of algebraic Schwarz methods.

We will present several numerical results for a collection of test graphs with up to half a million nodes and several million vertices arising from applications. These results show that the algebraic multigrid method for (1) is uniformly efficient, particularly when used as a preconditioner in the CG iteration. We then uniformly achieve linear convergence factors of about approximately 0.3.

Optionally, i.e. if the time allowed for this talk permits, we will compare the approach outlined so far with an algebraic multigrid method where the coarse variables are obtained as aggregates of the fine variables induced by a matching of G.

References

[1] Diekmann, R, Preis, R., Schlimbach, F. and Walshaw, C., Shape-optimized mesh partitioning and load balancing for parallel adaptive FEM, J. Parallel Computing 26, 1555-1581 (2000)

[2] Keller, H.B., On the solution of singular and semidefinite linear systems by iteration, J. Soc. Ind. Appl. Math., Ser. B, Numer. Anal. 2, 281-290 (1965)

[3] Meyerhenke, H., Monien, B. and Schamberger, S.: Accelerating shape optimizing load balancing for parallel FEM simulations by algebraic multigrid, 20th International Parallel and Distributed Processing Symposium (IPDPS 2006), IEEE Conference Proceeding (2006)

[4] Meyerhenke, H., Schamberger, S.: A Parallel shape optimizing load balancer, in Euro-Par 2006 (Nagel, W., Walter, W., and Lehner, W. eds, Lecture Notes in Computer Science 4128, Springer (2006), 232-242

[5] Stüben, K.: Algebraic multigrid (AMG): an introduction with applications, in Trottenberg, U., Oosterlee, C., and Schüller, A.: Multigrid, Academic Press (2000)

Linear Algebra Tools for Link Analysis and Dimensionality Reduction

Efstratios Gallopoulos

Internet Algorithmics and Information Retrieval (IR) have become grand challenge application areas for Numerical Linear Algebra. Two important problems are *i*) ranking webpages based on link analysis of the Web graph, and *ii*) dimensionality reduction of very large datasets. Specifically, the PageRank algorithm [8] introduces content-neutral ranking over Web pages. This ranking is applied on matrices derived from the link adjacencies resulting from a crawl, or to the set of pages returned by the Google search engine in response to posting a search query. Two viewpoints are to order the pages based on the nonnegative eigenvector or the nonnegative solution of a linear system. A stochastic approach is to order based on the steady state distribution of a Markov chain. We recently established a new interpretation of the series approach [5]. Specifically, the finite versions of the power series interpretation of PageRank and other related ranking functions ([2]) can be represented by means of products of Google-type matrices, each built from the same link matrix but, typically, different damping factors. We describe this "multidamping" approach and discuss its properties and implications from its use. Matrices in Web and Information Retieval applications are enormous, so dimensionality reduction is an essential step for their effective manipulation.

As with the Google matrix, image and text collections are encoded by matrices whose elements are frequently nonnegative. This motivates the construction of low rank approximations whose constituent factors are also nonnegative matrices. This leads to better interpretability ([7]) but also causes a jump in the mathematical and computational complexity of the problem relative to the SVD, e.g. it was recently shown that the exact nonnegative matrix factorization (NMF) is NP-complete [9]. We have been considering means to compute fast nonnegative approximate factorizations directly from the unit rank terms in the SVD expansion of the matrix. Both the partial and full SVD of the positive section of each such term is easily computable and can be used to generate columns and rows for the left and right factors of the NMF. Continuing recent work in [3], we present some important features of this approach that can be readily applied to initialize NMF algorithms. If time permits, we will also describe powerful new tools for researchers from numerical linear algebra working desiring rapid prototyping of algorithms and dataset development in IR and Web-IR.

The work reported in this presentation is conducted in collaboration with Christos Boutsidis and Giorgios Kollias.

References

[1] R. Baeza-Yates, P. Boldi and C. Castillo: "Generic Damping Functions for Propagating Importance in Link-Based Ranking". Internet Mathematics (to appear).

[2] P. Boldi, R. Posenato, M. Santini and S. Vigna. Traps and pitfalls of topic-biased PageRank". In WAW 2006. Fourth Workshop on Algorithms and Models for the Web-Graph, Lecture Notes in Computer Science. SpringerVerlag, 2007.

[3] C. Boutsidis and E. Gallopoulos, SVD based initialization: A head start for nonnegative matrix factorization", Pattern Rcognition, (to appear).

[4] M. Brinkmeier, Pagerank revisited", ACM Trans. Inter. Tech. 6 (2006), no. 3, 282-301.

[5] D. Gleich, P. Glynn, G. Golub and C. Greif, Three results on the PageRank vector: eigenstructure, sensitivity, and the derivative", Web Information Retrieval and Linear Algebra Algorithms, Dagstuhl Seminar Proceedings 07071: Web Information Retrieval and Linear Algebra Algorithms, A. Frommer, M. W. Mahoney, D. B. Szyld (Eds.), 2007.

[6] G. Kollias and E. Gallopoulos, Multidamping simulation framework for link-based ranking", Dagstuhl Seminar Proceedings 07071: Web Information Retrieval and Linear Algebra Algorithms, A. Frommer, M. W. Mahoney, D. B. Szyld (Eds.), 2007.

[7] D. D. Lee, H. S. Seung, Learning the Parts of Objects by Non-Negative Matrix Factorization", Nature 401 (1999) 788-791.

[8] L. Page, S. Brin, R. Motwani, and T. Winograd, The PageRank Citation Ranking: Bringing Order to the Web, Tech. report, Stanford Univ., 1998.

[9] S.A. Vavasis, On the complexity of nonnegative matrix factorization, In http://arxiv.org/abs/0708.4149, 2007.

An Optimal Complexity Algorithm for Non-Matching Grid Projections

Martin J. Gander

Non-matching grids are becoming more and more common in scientific computing. Two particular examples are the mortar methods in domain decomposition by Bernardi, Maday and Patera [1], and the patch method for local refinement by Glowinski, He, Lozinski, Rappaz and Wagner [2], and Picasso, Rappaz and Rezzonico [3], which is also known under the name 'numerical zoom', see Pironneau [4]. The former method permits an entirely parallel grid generation and adaptation, where naturally the processor that will be computing approximations to the solution on a subdomain is also generating and adapting the grid for that subdomain. This can be performed without communication, if non-matching grids are permitted, which is possible with the mortar method without loss of accuracy. In the patch method, one has a large scale solver for a particular partial differential equation, and wants to add more precision in certain areas, without having to change the large scale code. One thus introduces refined patches in these regions, and uses a residual correction iteration between solutions on the patches and solutions on the entire domain, to converge to a more refined solution in the patch regions. The patch method also permits the use of non-matching grids which is very convenient for the adaptive local refinement.

In both examples, an essential step in the algorithm is the projection of a finite element function from one grid onto the finite element space of the other, non-matching grid: in the mortar method and generalizations thereof, this projection is needed in every iteration of the domain decomposition method, in order to exchange interface conditions, and in the patch method, the residual is computed on one grid, and then has to be used as a right hand side for the other grid, and vice versa, at each iteration. One therefore needs to compute these grid projections. For a problem in two dimensions, the mortar method has one dimensional

interfaces, and a fast algorithm based on the ideas of merge sort computes the projection in $O(n)$ steps, where n is the number of elements touching the interface, see [5]. For the patch method however in two dimensions, a two dimensional projection is already needed. The same is the case for the 3d mortar method, where the interfaces are now two dimensional. The situation is even harder for the patch method in 3d, since three dimensional residuals need to be projected.

There are two problems that need to be addressed in order to obtain an efficient projection algorithm, a combinatorial one and a numerical one: the combinatorial one stems from the fact that in principle, every element of one grid could be intersecting with every element of the other grid, and hence the naive approach would immediately lead to an $O(n^2)$ algorithm, where n is the number of elements. This is well known in the community of specialists; for example in Flemisch, Kaltenbacher and Wohlmuth [6], the authors state: "One possible realization is given in Algorithm 1; we remark that the naive implementation of this algorithm for 2d-problems is of order $O(n^2)$". The numerical difficulty is related to the calculation of the intersection of two finite elements: it is numerically difficult to calculate this intersection, because one needs to take numerical decisions whether two lines intersect or not, and whether one point is in an element or not, see Shewchuk [7], where a special, adaptive precision floating point arithmetic is developed for general decision problems of this type. For the patch method, the authors in [3] state: "Some difficulties remain though since we must compute integrals involving shape functions that are defined on non-compatible meshes", and they use a midpoint rule and compute in which elements the barycenters of the elements of the other grid lie to obtain an approximation. In [6], the authors state: "the intersection of two arbitrary triangles is quite more complex than the situation in 1d", and a long, sophisticated program with many special cases is used to compute the intersection [8]. The influence of the error of a quadrature rule for this problem is studied in [9]. For the patch method in 3d, such an approximation even makes the numerical evaluation of the method difficult [10].

We will describe an optimal complexity algorithm which overcomes both the combinatorial and numerical problems. It uses an advancing front technique and neighboring element information, which is often available in a finite element grid, in order to reduce the combinatorial complexity to $O(n)$ where n is the number of elements. This result is independent of the dimension of the problem, and the combinatorial part of the algorithm is surprisingly short, it is a little more than half a page in Matlab. The numerical problems for calculating the intersections are overcome by an inclusive strategy in the computation of intersections of lines, and points within elements, which turns out to be numerically robust. For the 2d case, this procedure gives the intersecting polygon of two elements, and for the 3d case the intersecting polyhedron, face by face with counter-clock wise ordered corners. We will show numerical experiments in two and three dimensions to illustrate the algorithm, and also show scaling experiments which confirm the optimal complexity of the new algorithm.

References

[1] C. Bernardi, Y. Maday and A. T. Patera, Domain decomposition by the mortar element method, in Asymptotic and Numerical Methods for Partial Differential Equations with Critical Parameters, editors: H.G. Kaper and M. Garbey, N.A.T.O. ASI, Kluwer Academic Publishers, pp. 269-286, 1993.

[2] R. Glowinski, J. He, A. Lozinski, J. Rappaz and J. Wagner, Finite element approximation of multi-scale elliptic problems using patches of elements. Numerische Mathematik 101(4), pp. 663-687, 2005.

[3] M. Picasso, J. Rappaz, V. Rezzonico, Multiscale Algorithm with Patches of Finite Elements, Communications in Numerical Methods in Engineering, 2007.

[4] O. Pironneau, Numerical Zoom for Multiscale Problems with an Application to Flows Through Porous Media, plenary presentation given at the conference in honor of J. Rappaz 60th birthday, Lausanne, September 2007.

[5] M.J. Gander, C. Japhet, Y. Maday and F. Nataf, A New Cement to Glue Nonconforming Grids with Robin Interface Conditions : The Finite Element Case, in Domain Decomposition Methods in Science and Engineering, Lecture Notes in Computational Science and Engineering 40, Springer Verlag, pp. 259–266, 2004.

[6] B. Flemisch, M. Kaltenbacher and B. I. Wohlmuth, Elasto-Acoustic and Acoustic-Acoustic Coupling on Nonmatching Grids, Int. J. Num. Meth. Eng., 2002.

[7] J. R. Shewchuk, Adaptive Precision Floating-Point Arithmetic and Fast Robust Geometric Predicates, technical report CMU-CS-96-140, Carnegie Mellon University, 1996.

[8] C. Hager, private communication.

[9] Y. Maday, F. Rapetti and B. I. Wohlmuth, The Influence of Quadrature Formulas in 2d and 3d Mortar Element Methods, preprint, 2007.

[10] V. Rezzonico, Multiscale Algorithm with Patches: Applications and Extensions, plenary presentation given on the Swiss Numerical Analysis Day, April 4th, 2007.

Convergence in backward error of relaxed GMRES

L. Giraud, S. Gratton, J. Langou

This work is the follow-up of the experimental study presented in [1]. It is based on and extends some theoretical results in [3, 4]. In a backward error framework we study the convergence of GMRES when the matrix-vector products are performed inaccurately. This inaccuracy is modeled by a perturbation of the original matrix. We prove the convergence of GMRES when the perturbation size is proportional to the inverse of the computed residual norm; this implies that the accuracy can be relaxed as the method proceeds which gives rise to the terminology relaxed GMRES. As for the exact GMRES we show under proper assumptions that only happy breakdowns can occur. Furthermore the convergence can be detected using a by-product of the algorithm. We explore the links between relaxed right-preconditioned GMRES and flexible GMRES. In particular this enables us to derive a proof of convergence of FGMRES.

Many situations in scientific computing may benefit from such a scheme. For instance a natural application of this idea occurs in computational electromagnetic, where the fast multipole method provides approximations of the matrix-vector product within a user-defined

accuracy; the less accurate the matrix-vector, the faster the computation. The key point is then to design a criterion to control the accuracy of the matrix-vector product so that the iterates achieve a satisfactory convergence level. Another example arises in non-overlapping domain decomposition where the matrix-vector involving the Schur complement can be approximated. In classical implementation the local subproblems are solved using a sparse direct solver, that can be replaced by an iterative scheme which stopping criterion is selected to comply with the relaxation strategy for GMRES applied to the interface problem.

Results on numerical experiments to illustrate the behaviour of the relaxed GMRES monitored by the proposed relaxation strategies will be reported.

More details on this work can be found in [2].

References

[1] A. Bouras and V. Frayssé. Inexact matrix-vector products in Krylov methods for solving linear systems: a relaxation strategy. *SIAM Journal on Matrix Analysis and Applications*, 26(23):660–678, 2005.

[2] L. Giraud, S. Gratton, and J. Langou. Convergence in backward error of relaxed GMRES. *SIAM J. Scientific Computing*, 29(2):710–728, 2007.

[3] V. Simoncini and D. B. Szyld. Theory of inexact Krylov subspace methods and applications to scientific computing. *SIAM J. Scientific Computing*, 25(2):454–477, 2003.

[4] J. van den Eshof and G. L. G. Sleijpen. Inexact Krylov subspace methods for linear systems. *SIAM Journal on Matrix Analysis and Applications*, 26(1):125–153, 2004.

Parallel Eigenvalue Reordering in Real Schur Forms

Robert Granat, Bo Kågström, Daniel Kressner

Given a general square matrix $A \in R^{n\times n}$, computing the *Schur decomposition* (see, e.g., [3])

$$Q^T A Q = T \tag{1}$$

is the standard approach to solving *non-symmetric eigenvalue problems* (NEVPs), that is, computing eigenvalues and invariant subspaces (or eigenvectors) of a general dense matrix. In Equation (1), $T \in R^{n\times n}$ is quasi-triangular with diagonal blocks of size 1×1 and 2×2 corresponding to real and complex conjugate pairs of eigenvalues, respectively, and $Q \in R^{n\times n}$ is orthogonal. The matrix T is called the *real Schur form* of A and its diagonal blocks, corresponding to the eigenvalues of A, can occur in any order along the diagonal.

For any decomposition of the form

$$Q^T A Q = T \equiv \begin{bmatrix} T_{11} & T_{12} \\ 0 & T_{22} \end{bmatrix}, \tag{2}$$

with $T_{11} \in R^{p\times p}$ for some integer p and provided that $T(p+1,p)$ is zero, the first p columns of the matrix Q span an *invariant subspace* of A corresponding to the p eigenvalues of T_{11}. Invariant subspaces are important since they often eliminate the need of explicit eigenvector calculations in applications. However, notice that if $T(2,1) = 0$, $v_1 = Q(1:n,1)$ is an

(orthonormal) eigenvector of A corresponding to the eigenvalue $\lambda_1 = T(1,1)$. If $T(2,1) \neq 0$, the leading 2×2 block corresponds to a complex conjugate pair of eigenvalues and in order to retain to real arithmetic we have to compute their eigenvectors simultaneously.

Computing the *ordered* real Schur form (ORSF) of a general $n \times n$ matrix is vital in many applications, for example, in *stable-unstable separation* for solving Riccati matrix equations [11]. It can also be used for computing explicit eigenvectors of A by reordering each eigenvalue (block) of interest to the leading position in T and reading off the corresponding first (two) column(s) in Q, as illustrated above. This approach is not utilized in practice for the standard (or generalized) Schur form but is useful in computing eigenvectors to certain matrix products arising in periodic systems analysis and design (see, e.g., [6]).

In [1], a direct algorithm for reordering adjacent eigenvalues in the real Schur form (2) is proposed. For the special case of a tiny matrix T where $p, n-p \in \{1,2\}$, the method is as follows:

- Solve the continuous-time Sylvester (SYCT) equation

$$T_{11}X - XT_{22} = \gamma T_{12}, \tag{3}$$

  where γ is a scaling factor to avoid overflow in the right hand side.

- Compute the QR factorization

$$\begin{bmatrix} -X \\ \gamma I \end{bmatrix} = QR \tag{4}$$

  using Householder transformations (elementary reflectors).

- Apply Q in the similarity transformation of T:

$$\tilde{T} = Q^T T Q \tag{5}$$

- Standardize 2×2 block(s) if any exists.

In the method above, each swap is performed tentatively to guarantee backward stability by rejecting each swap that appears unstable with respect to a stability criterion. By applying a bubble-sort procedure based on the adjacent swap method, where all selected eigenvalues (usually pointed to by a *select* vector) are moved step by step towards the top-left corner of T, ordered Schur forms can be computed. We remark that for $n = 2$ the adjacent swaps are performed using Givens rotations.

Following the ideas from [10], we may *delay* and *accumulate* all updates from inside a *computational window* until the reordering is completed locally. Then the *update regions* of T (and the corresponding regions of Q) are updated, preferably in level 3 operations. The window is then moved towards the top-left part of the matrix, such that the cluster now ends up in the bottom-right part of the current window, and the reordering process is repeated.

We present parallel algorithms [7] for blocked variants [10] of the method in [1]. Furthermore, we apply our techniques to parallel eigenvalue reordering in the generalized real Schur form [8, 9]. The novel parallel algorithms adopt computational windows and delay multiple outside-window updates until each window has been completely reordered locally. By using multiple concurrent windows, a high level of concurrency is obtained if computations and communications are organized such that synchronization overhead is minimized.

Experimental results for ScaLAPACK-style Fortran 77 implementations on a Linux cluster confirm the efficiency and scalability of our algorithms in terms of more than 16 times of parallel speedup using 64 processor for large scale problems. Even on a single processor our implementations delivers another 14 times of speedup if compared to the state-of-the-art serial implementation for the standard Schur form from LAPACK.

It will be demonstrated how the parallel algorithms in combination with the existing functionality of ScaLAPACK and the recently developed software library SCASY [4, 5] can be used for parallel computation of stable invariant subspaces and associated reciprocal condition number estimates [2] of Hamiltonian matrices on distributed memory platforms. Moreover, we will also outline how the parallel reordering method can be generalized to large-scale periodic (or product) eigenvalue problems.

References

[1] Z. Bai and J. W. Demmel. On swapping diagonal blocks in real Schur form. *Linear Algebra Appl.*, 186:73–95, 1993.

[2] Z. Bai, J. W. Demmel, and A. McKenney. On computing condition numbers for the nonsymmetric eigenproblem. *ACM Trans. Math. Software*, 19(2):202–223, 1993.

[3] G. H. Golub and C. F. Van Loan. *Matrix Computations*. Johns Hopkins University Press, Baltimore, MD, third edition, 1996.

[4] R. Granat and B. Kågström. Parallel Solvers for Sylvester-type Matrix Equations with Applications in Condition Estimation, Part I: Theory and Algorithms. UMINF-07.15, Dept. Computing Science, Umeå University. Submitted to *ACM Trans. Math. Software*, 2007.

[5] R. Granat and B. Kågström. Parallel Solvers for Sylvester-type Matrix Equations with Applications in Condition Estimation, Part II: the SCASY Software. UMINF-07.16, Dept. Computing Science, Umeå University. Submitted to *ACM Trans. Math. Software*, 2007.

[6] R. Granat, B. Kågström and D. Kressner Computing Periodic Deflating Subspaces Associated with a Specified Set of Eigenvalues. *BIT Numerical Mathematics* (accepted), 2007.

[7] R. Granat, B. Kågström and D. Kressner Parallel Eigenvalue Reordering in Real Schur Forms. UMINF-07.20, Dept. Computing Science, Umeå University. Submitted to *Concurrency and Computation: Practice and Experience*, September 2007. Also as LAPACK Working Note #192.

[8] B. Kågström. A direct method for reordering eigenvalues in the generalized real Schur form of a regular matrix pair (A, B). In *Linear algebra for large scale and real-time applications (Leuven, 1992)*, volume 232 of *NATO Adv. Sci. Inst. Ser. E Appl. Sci.*, pages 195–218. Kluwer Acad. Publ., Dordrecht, 1993.

[9] B. Kågström and P. Poromaa. Computing eigenspaces with specified eigenvalues of a regular matrix pair (A, B) and condition estimation: theory, algorithms and software. *Numer. Algorithms*, 12(3-4):369–407, 1996.

[10] Daniel Kressner. Block algorithms for reordering standard and generalized Schur forms. *ACM Trans. Math. Software*, 32(4):521–532, December 2006.

[11] P. Van Dooren. A generalized eigenvalue approach for solving Riccati equations. *SIAM J. Sci. Statist. Comput.*, 2(2):121–135, June, 1981.

Quasi-Newton formula, matrix nearness problems and preconditioning

S. Gratton, A. Sartenaer, Ph.L. Toint, J. Tshimanga

The history of quasi-Newton methods for optimization is rich and long. Starting with the Davidon-Fletcher-Powell (DFP) method, exemplified by the famous Broyden-Fletcher-Goldfarb-Shanno (BFGS) update and excellently explained in the classical book by Dennis and Schnabel [3], they have played an important role in the solution of practical optimization problems. When these techniques are applied to quadratic problems with coefficient matrix A, they provide estimations of A and of A^{-1} based on low rank modifications of the identity matrix. They are therefore closely related to two topics that play an important role in numerical linear algebra : matrix nearness problems and preconditioning by low rank update of the identity matrix.

In this work we study the properties of approximations of the inverse of a matrix derived by generalizing the BFGS formula. We suppose that A is a symmetric and positive definite matrix of order n, and denote by S a $n \times k$ matrix of full rank k. Letting $Y = AS$ and $B = Y^T S$, for any nonzero real parameter σ, we define the bloc generalization of the BFGS formula by

$$H = \left[I_n - SB^{-1}Y^T\right]\left[I_n - YB^{-1}S^T\right] + \sigma SB^{-1}S^T.$$

We are interested in the properties of H considered as an approximation of A^{-1}, when the set S is varied. Even when A is sparse, H is usually a dense matrix ; however matrix-vector products involving H can be performed without forming H, just using matrix vector products by A and vector operations involving S. By analogy with the celebrated limited memory BFGS update, we call H a *limited memory preconditioner* (LMP) of A.

We first show basic properties on H. Under the full rank assumption on S, we show that the matrix H is positive definite, and we derive expressions for the matrix vector product by H^{-1}. We also show that H admits a factorization $H = GG^T$, where G is not the Cholesky factor of H; however the action of G on a vector is again easily obtained with limited $O(kn)$ storage and operations.

We then show that H has interesting properties that makes it reasonable to consider it as a possible preconditioner of A. Readily, $HAS = \sigma S$, which indicates that H/σ solves exactly any linear system $Ax = b$ with right-hand side in the range of AS. We also show that, *for any full rank* S the condition number of AH is smaller than that of A, and that the eigenvalues of the preconditioned matrix HA exhibit larger clusters than that of A. We also show that for particular choices of S, the limited memory preconditioner is mathematically equivalent to the deflation technique described in [5, 11]. Note that, in order to be efficient and keep the dimension of the low dimensional space reasonably small, limited memory preconditioners are generally used in combination with a first level preconditioner that does a good job of

clustering most eigenvalues near to σ with relatively few outliers near the origin [1, 2, 11, 13, 15].

As an application of the LMP theory, we consider the following preconditioning technique for the solution of symmetric positive definite linear systems with a constant symmetric and positive definite matrix A and multiple right-hand sides b_k, $k = 1, 2, \ldots$ available in sequence. We assume that the systems are sequentially solved with the conjugate gradient algorithm (CG). For the k linear system, we decide to introduce in S vectors that are by-products obtained from the application of CG on the $k-1$ first linear systems. More precisely, we consider three examples of LMP: the spectral-LMP, the Quasi-Newton-LMP and the Ritz-LMP, where the columns of S are respectively eigenvectors of A, descent directions of CG and Ritz vectors; the spectral-LMP [1, 2, 8] and Quasi-Newton-LMP [12], are well known preconditioners, and the Ritz-LMP is new. The common particularity shared by these three LMPs is that the column of S are conjugate with respect to the matrix A. We next consider the practical issue of using Ritz-information to approximate eigen-information [4] and show that, using this Ritz-information within the spectral- LMP yields a perturbed form of the Ritz-LMP. We address the question of selecting the column of S to have the most efficient Ritz-LMP and show that our preconditioner is competitive with the one proposed in [12] . A generalization of these results is proposed for a sequence of linear systems with slowly varying matrices.

To illustrate all the above theoretical results, numerical experiments are shown both on academic problems and on a real life data assimilation system [4]. In the latter case, non-linear least-squares problems with more than 10^7 unknowns is solved using a Gauss-Newton approach [9]. The sequence of linear least-squares problems involved in the nonlinear solution scheme is solved by CG. The correspondence between CG and Lanczos [6] is exploited on each linear problem to extract approximate spectral information. This spectral information is used to design a spectral preconditioner for the subsequent linear least-squares problem. The numerical experiments demonstrate the superiority of the Ritz-LMP preconditioner we designed, over the other existing variants. Finally we define a computationally efficient way to generate a useful matrix S for discretized problems by using smoothing properties of stationary iterative methods and show preliminary results on academic examples from the calculus of variations [10].

The results of this work are presented in a recent technical report [14].

References

[1] B. Carpentieri, I. S. Duff, and L. Giraud. A class of spectral two-level preconditioners. *SIAM Journal on Scientific Computing*, 25:749–765, 2003.

[2] B. Carpentieri, L. Giraud, and S. Gratton. Additive and multiplicative spectral two-level preconditioners for general linear systems. *SIAM J. Scientific Computing*, 29(4):1593–1612, 2007.

[3] J. Dennis, R. Schnabel. Numerical Methods for Uncostrained Optimization and Nonlinear Equations. Classics in Applied Mathematics. SIAM. 2000.

[4] M. Fisher. Minimization algorithms for variational data assimilation. In *Proc. ECMWF seminar "Recent developments in numerical methods for atmospheric modelling*, pages 364–385, September 1998.

[5] J. Frank, C. Vuik. On the construction of deflation-based preconditioners. *SIAM Journal on Scientific Computing*, 23:442–462, 2001.

[6] G. H. Golub and C. F. Van Loan. *Matrix Computations*. Johns Hopkins University Press, Baltimore, MD, USA, third edition, 1996.

[7] L. Giraud, S. Gratton. On the sensitivity of some spectral preconditioners. *SIAM J. Matrix Analysis and Applications*, 27(04),1089–1105, 2006.

[8] L. Giraud, S. Gratton, E. Martin. Incremental spectral preconditioners for sequences of linear systems. *Applied Numerical Mathematics*, (57)11-12:1164–1180, 2007.

[9] S. Gratton, A. Lawless, N. Nichols. Approximate Gauss-Newton methods for nonlinear least squares problems. *SIAM J. on Optimization*, 18:106–132, 2007.

[10] S. Gratton, A. Sartenaer, Ph.L. Toint. Recursive Trust-Region Methods for Multiscale Nonlinear Optimization. *Accepted for publication in SIAM J. on Optimization.*

[11] R. Nabben, C. Vuik. A comparison of deflation and coarse grid correction applied to porous media flow. Report 03-10, Delft University of Technology, Department of Applied Mathematical Analysis, Delft, 2003.

[12] J. Nocedal, J.L. Morales. Automatic Preconditioning by Limited Memory Quasi-Newton Updating. *SIAM J. Optimization*, 10(4):1079–1096,2000.

[13] A. Padiy, O. Axelsson, and B. Polman. Generalized augmented matrix preconditioning approach and its application to iterative solution of ill-conditioned algebraic systems. *SIAM Journal on Matrix Analysis and Applications*, 22(3):793–818, 2000.

[14] J. Tshimanga, S. Gratton, A. Weaver, A. Sartenaer. Limited-memory preconditioners for incremental four-dimensional variational data assimilation. Technical Report TR/PA/07/95, CERFACS, Toulouse, France, 2007. Submitted to the Quarterly Journal of the Royal Meteorological Society.

[15] H. Waisman, J. Fish, R. S. Tuminaro, and J. Shadid. The generalized global basis method. *Int. Journal of Numerical Methods in Engineering*, 61(8):1243–1269, 2004.

Upper and Lower Bounds on Norms of Functions of Matrices

Anne Greenbaum

Given an n by n matrix A, we look for a set S in the complex plane and positive scalars m and M such that for all polynomials (or analytic functions) p, the inequalities

$$m \cdot \inf\{\|f\|_{\mathcal{L}^\infty(S)} : \ f(A) = p(A)\} \leq \|p(A)\| \leq M \cdot \inf\{\|f\|_{\mathcal{L}^\infty(S)} : \ f(A) = p(A)\}$$

hold, where $\|\cdot\|$ denotes the operator norm corresponding to the 2-norm for vectors. We try to find a set S for which M/m is of moderate size; e.g., much smaller than the condition number of an eigenvector matrix of A, which is a ratio one finds when A is diagonalizable and S is taken to be the spectrum of A. We show that for 2 by 2 matrices, if S is the field of values, then one can take $m = 1$ and $M = 2$. We show that for a perturbed Jordan block – a matrix A that is an n by n Jordan block with eigenvalue 0 except that its $(n, 1)$-entry is $\nu \in (0, 1)$ – if S is the unit disk, then $m = M = 1$. More generally, we show that if A is a companion matrix whose eigenvalues lie in the open unit disk $\mathcal{D}$, then $m = 1$ if $S = \mathcal{D}$. We argue, however, that because of the possibly wild behavior of even minimal-norm interpolating functions when the number of interpolation points is large and the interpolation values are chosen perversely, for many matrices it may be difficult/impossible to find such a set S.

A Conjugate Gradient Method for Skew-Symmetric Matrices

Chen Greif, James Varah

A conjugate gradient method for a skew-symmetric matrix A can be derived from a Lanczos-type iteration by using a stable block decomposition of the skew tridiagonal matrix generated by the algorithm. For even k, the Lanczos & Galerkin conditions can be broken down into a series of recurrences much like those of the symmetric positive definite case. Indeed, the resulting algorithm is mathematically equivalent to CG applied to the positive definite matrix $(-A^2)$, but does not involve squaring the matrix. Various preconditioners can be applied to the basic iteration. We describe some of these, including a two-sided preconditioner based on the Bunch decomposition of a skew-symmetric matrix, which appears to work well in practice.

Numerical linear algebra for spectral theory of block operator matrices

Luka Grubišić

We present a perturbation theory for sign-indefinite quadratic forms in a Hilbert space. Under an additional qualitative assumption on the structure of the form, which is in analogy to the structure of the so called quasi-definite matrices from Linear Algebra, we prove an operator representation theorem. Special emphasis has been set on analysis of quadratic forms which are unbounded at both ends and which can be tackled by our algebraic theory in a natural way.

Furthermore, with the help of weakly formulated Riccati equations we obtain subspace perturbation theorems for these "quasi-definite operators" and present accompanying estimates on the perturbation of the spectra. Our estimates for the rotation of invariant subspaces have a form of a "relative" $\tan 2\Theta$ theorem and are a new result even in the standard matrix case. The example of the Stokes block matrix operator—which is associated to the Cosserat eigenvalue problem—is used to illustrate the theory and to show that our estimates can be attained. As another application we present estimates for the strong solution of the Stokes system in a Lipschitz domain and favorably compare these results with considerations from [1, Section 6 and 7]. The title of the presentation has been motivated by the paper [5], which uses linear algebra in the spectral theory in another context. The present work continues our previous studies of weakly formulated operator equations and relative perturbation theory from [2, 3, 4]

This is a joint work with Vadim Kostrykin, Konstantin Makarov and Krešimir Veselić

References

[1] M. Dindoš and M. Mitrea, *The stationary Navier-Stokes system in nonsmooth manifolds: The Poisson problem in Lipschitz and C^1 domains*, Arch. Ration. Mech. Anal. **174** (2004), 1 – 47.

[2] L. Grubišić. On eigenvalue estimates for nonnegative operators. *SIAM J. Matrix Anal. Appl.*, 28(4):1097–1125, 2006.

[3] L. Grubišić and K. Veselić. On Ritz approximations for positive definite operators I (theory). *Linear Algebra and its Applications*, 417(2-3):397–422, 2006.

[4] L. Grubišić and K. Veselić. On weakly formulated Sylvester equation and applications. *Integral Equations and Operator Theory*, 58(2):175–204, 2007.

[5] J. Sjoestrand and M. Zworski. Elementary linear algebra for advanced spectral problems. 2003. http://www.citebase.org/abstract?id=oai:arXiv.org:math/0312166.

Structured perturbation theory for inexact Krylov projection methods in model reduction

Chris Beattie, Serkan Gugercin

Dynamical systems are the basic framework for modeling and control of an enormous variety of complex systems. Direct numerical simulation of the associated models has been one of the few available means when goals include accurate prediction or control of complex physical phenomena. However, the ever increasing need for improved accuracy requires the inclusion of ever more detail in the modeling stage, leading inevitably to ever larger-scale, ever more complex dynamical systems.

Simulations in such large-scale settings can be overwhelming and make unmanageably large demands on computational resources; this is the main motivation for model reduction, which has as its goal production of a much lower dimensional system having the same input/output characteristics as the original system.

Rational Krylov subspaces are often capable of providing nearly optimal approximating subspaces for model reduction, although the main cost to generate the required bases arises from solving linear systems of the form $(\sigma \mathbf{I}_n - \mathbf{A})\mathbf{v} = \mathbf{b}$. Since the need for more detail and accuracy in the modeling stage causes the system dimension, n, to reach levels on the order of millions, direct solvers for this linear system are no longer feasible. In practice, iterative solvers become necessary and choice of termination criteria and effect of preconditioning influences the quality of final reduced order models.

In this talk, we introduce the use of inexact solves in a Krylov-based model reduction setting and present the resulting structured perturbation effects on the underlying model reduction problem. For a selection of interpolation points that satisfy first-order necessary $\mathcal{H}_2$-optimality conditions, a primitive basis remains remarkably well-conditioned and errors due to inexact solves does not tend to degrade the reduced order models. Conversely, for poorly selected interpolation points, errors can be greatly magnified through the model reduction process.

When inexact solves are performed within a Petrov-Galerkin framework, the resulting reduced order models are backward stable with respect to the approximating transfer function. As a consequence, Krylov-based model reduction with well chosen interpolation points is robust with respect to the structured perturbations due to inexact solves. General bounds on the $\mathcal{H}_2$ and $\mathcal{H}_\infty$ system error associated with an inexact reduced order model are introduced that provide a new tool to understand the structured backward error and stopping criteria.

Convergence Rate of the Cyclic Reduction Algorithm for a Nonlinear Matrix Equation

Chun-Hua Guo, Wen-Wei Lin

A discrete-time quasi-birth-death (QBD) process is a Markov chain with state space $\{(i,j)\,|\,i\geq 0, 1\leq j\leq m\}$, and with a transition probability matrix of the form

$$P=\begin{bmatrix} B_0 & B_1 & 0 & 0 & \cdots \\ A_0 & A_1 & A_2 & 0 & \cdots \\ 0 & A_0 & A_1 & A_2 & \cdots \\ 0 & 0 & A_0 & A_1 & \cdots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix},$$

where B_0, B_1, A_0, A_1, and A_2 are $m\times m$ nonnegative matrices such that P is stochastic. In particular, $(A_0+A_1+A_2)e=e$, where $e=(1,1,\ldots,1)^T$.

We assume that the matrix P and the matrix $A=A_0+A_1+A_2$ are both irreducible. Thus, $A_0\neq 0$ and $A_2\neq 0$. Moreover, there exists a unique vector $\alpha>0$ with $\alpha^T e=1$ and $\alpha^T A=\alpha^T$.

The minimal nonnegative solution G of the matrix equation

$$G=A_0+A_1G+A_2G^2 \tag{1}$$

plays an important role in the study of the QBD process. We also need the dual equation

$$F=A_2+A_1F+A_0F^2,$$

and let F be its minimal nonnegative solution. It is known that the minimal solutions G and F are stochastic or substochastic.

The cyclic reduction (CR) algorithm for (1), or for $-A_0+(I-A_1)G-A_2G^2=0$, is the following:

Set $T_0=A_0, \quad U_0=I-A_1, \quad V_0=A_2, \quad W_0=I-A_1$.
For $k=0,1,\ldots$, compute

$$\begin{aligned} T_{k+1} &= T_kU_k^{-1}T_k, \\ U_{k+1} &= U_k-T_kU_k^{-1}V_k-V_kU_k^{-1}T_k, \\ V_{k+1} &= V_kU_k^{-1}V_k, \\ W_{k+1} &= W_k-V_kU_k^{-1}T_k. \end{aligned}$$

It is shown in [1] that the algorithm is well defined. When $\alpha^TA_0e\neq\alpha^TA_2e$, the sequence $\{W_k\}$ converges quadratically to a nonsingular matrix W_* and $W_*^{-1}A_0=G$. When $\alpha^TA_0e=\alpha^TA_2e$, in which case the QBD is said to be null recurrent and the solutions G and F are both stochastic, determining the convergence rate of $\{W_k\}$ has been a difficult problem.

In [3] the convergence of a closely related algorithm is shown to be linear with rate 1/2, after a long analysis and under two additional assumptions.

In this talk, we show that the convergence of the CR algorithm is at least linear with rate 1/2 in the null recurrent case, without those two additional assumptions.

The much simpler analysis is based on a connection between the CR algorithm and the doubling algorithm. Such a connection is observed earlier in [5] for a different nonlinear matrix equation.

It is easily verified that

$$L_0 \begin{bmatrix} I \\ A_2G \end{bmatrix} = M_0 \begin{bmatrix} I \\ A_2G \end{bmatrix} G,$$

where

$$L_0 = \begin{bmatrix} 0 & I \\ A_0 & 0 \end{bmatrix}, \quad M_0 = \begin{bmatrix} A_2 & 0 \\ I - A_1 & -I \end{bmatrix},$$

and that $L_0 - \lambda M_0$ is a linearization of $-A_0 + \lambda(I - A_1) - \lambda^2 A_2$.

Using the doubling technique, we can define

$$L_k = \begin{bmatrix} -Q_k & I \\ T_k & 0 \end{bmatrix}, \quad M_k = \begin{bmatrix} V_k & 0 \\ P_k & -I \end{bmatrix}$$

for all $k \geq 0$ by the following doubling algorithm:

Set $T_0 = A_0, \quad P_0 = I - A_1, \quad Q_0 = 0, \quad V_0 = A_2.$

For $k = 0, 1, \ldots$, compute

$$\begin{aligned} T_{k+1} &= T_k(P_k - Q_k)^{-1}T_k, \\ P_{k+1} &= P_k - T_k(P_k - Q_k)^{-1}V_k, \\ Q_{k+1} &= Q_k + V_k(P_k - Q_k)^{-1}T_k, \\ V_{k+1} &= V_k(P_k - Q_k)^{-1}V_k. \end{aligned}$$

It is readily seen that the CR algorithm is recovered from the doubling algorithm by letting $U_k = P_k - Q_k$ and $W_k = W_0 - Q_k$. So the doubling algorithm is also well defined.

We also have

$$\widehat{L}_0 \begin{bmatrix} I \\ A_0F \end{bmatrix} = \widehat{M}_0 \begin{bmatrix} I \\ A_0F \end{bmatrix} F,$$

where

$$\widehat{L}_0 = \begin{bmatrix} 0 & I \\ A_2 & 0 \end{bmatrix}, \quad \widehat{M}_0 = \begin{bmatrix} A_0 & 0 \\ I - A_1 & -I \end{bmatrix}.$$

It is easily seen that $\widehat{M}_0 - \lambda\widehat{L}_0$ is also a linearization of $-A_0 + \lambda(I - A_1) - \lambda^2 A_2$.

The convergence rate of the doubling algorithm in the null recurrent case is determined by using

1. an important result [2] about the spectral properties of the quadratic pencil $-A_0 + \lambda(I - A_1) - \lambda^2 A_2$ and of the matrices G and F,
2. the Kronecker form for matrix pairs (M_0, L_0) and $(\widehat{L}_0, \widehat{M}_0)$,
3. the doubling technique applied to these two matrix pairs.

We have the following result:

$$P_k - (I - A_1 - A_0F) = O(2^{-k}), \quad Q_k - A_2G = O(2^{-k}).$$

So $\lim P_k = P_* = I - A_1 - A_0F$, which is nonsingular [4]. We have $P_*^{-1}A_2 = F$, since $A_2 = (I - A_1 - A_0F)F$. Similarly, $\lim Q_k = Q_* = A_2G$, with $I - A_1 - A_2G$ nonsingular [4]. We have $(I - A_1 - Q_*)^{-1}A_0 = G$, since $A_0 = (I - A_1 - A_2G)G$.

For the CR algorithm, we have $W_k = W_0 - Q_k = I - A_1 - Q_k$. So we now know that, in the null recurrent case, $\{W_k\}$ converges to $W_* = I - A_1 - A_2G$ at least linearly with rate 1/2 and that $G = W_*^{-1}A_0$.

With our convergence rate analysis for the null recurrent case, we have gained more insights for the convergence behaviour for the CR algorithm for nearly null recurrent cases, regardless of how many eigenvalues of G or F are on or near the unit circle.

References

[1] D. Bini and B. Meini, *On the solution of a nonlinear matrix equation arising in queueing problems*, SIAM J. Matrix Anal. Appl., 17 (1996), pp. 906–926.

[2] H. R. Gail, S. L. Hantler, and B. A. Taylor, *Spectral analysis of $M/G/1$ and $G/M/1$ type Markov chains*, Adv. Appl. Probab., 28 (1996), pp. 114–165.

[3] C.-H. Guo, *Convergence analysis of the Latouche–Ramaswami algorithm for null recurrent quasi-birth-death processes*, SIAM J. Matrix Anal. Appl., 23 (2002), pp. 744–760.

[4] G. Latouche and V. Ramaswami, *Introduction to Matrix Analytic Methods in Stochastic Modeling*, SIAM, Philadelphia, PA, 1999.

[5] W.-W. Lin and S.-F. Xu, *Convergence analysis of structure-preserving doubling algorithms for Riccati-type matrix equations*, SIAM J. Matrix Anal. Appl., 28 (2006), pp. 26–39.

Computing the Fréchet Derivative of the Matrix Exponential, with an application to Condition Number Estimation

Nicholas J. Higham, Awad H. Al-Mohy

The matrix exponential is a much-studied matrix function having many applications. The Fréchet derivative of the matrix exponential describes the first order sensitivity of e^A to perturbations in A and its norm determines a condition number for e^A. Among the numerous methods for computing e^A the scaling and squaring method is probably the most widely used, with the state of the art implementation being that in [2]. We show how Higham's implementation can be extended to compute both e^A and the Fréchet derivative at A in the direction E, denoted by $L(A, E)$, at a cost about three times that for computing e^A alone. The algorithm is derived from the scaling and squaring method by differentiating the Padé approximants and the squaring recurrence, re-using quantities computed during the evaluation of the Padé approximant, and intertwining the recurrences in the squaring phase. To guide the choice of algorithmic parameters an extension of the existing backward error analysis for the scaling and squaring method is developed which shows that, modulo rounding errors, the approximations obtained are $e^{A+\Delta A}$ and $L(A + \Delta A, E + \Delta E)$, with the *same* ΔA

in both cases, and with computable bounds on $\|\Delta A\|$ and $\|\Delta E\|$. The algorithm for $L(A, E)$ is used to develop an algorithm for simultaneously computing both e^A and an estimate of its condition number.

In addition to results specific to the exponential, we develop some results and techniques for arbitrary functions. We show how a matrix iteration for $f(A)$ yields an iteration for the Fréchet derivative and show how to efficiently compute the Fréchet derivative of a power series. We also show that a matrix polynomial and its Fréchet derivative can be evaluated at a cost at most three times that of computing the polynomial itself and give a general framework for evaluating a matrix function and its Fréchet derivative via Padé approximation.

References

[1] A. H. Al-Mohy and N. J. Higham. Computing the Fréchet derivative of the matrix exponential, with an application to condition number estimation. MIMS EPrint, Manchester Institute for Mathematical Sciences, The University of Manchester, UK, 2008.

[2] Nicholas J. Higham. The scaling and squaring method for the matrix exponential revisited. *SIAM J. Matrix Anal. Appl.*, 26(4):1179–1193, 2005.

On the computation of real and purely imaginary eigenvalues of large sparse matrices

Michiel Hochstenbach

The question whether a large sparse matrix has *purely imaginary* eigenvalues is an important one in many applications, e.g., stability considerations and bifurcation analysis; it may also be useful in the computation of pseudospectra (see for instance [2]).

The harmonic and refined Rayleigh–Ritz extraction methods are valuable alternatives to standard Rayleigh–Ritz if interior eigenvalues are wanted and a target value is given for the eigenvalues of interest [5, 6, 1, 7]. However, for the task of computing purely imaginary eigenvalues, a sensible target is often not available.

In [3], extraction methods were proposed that are suitable to find rightmost eigenvalues. These techniques are, as a result, promising for purely imaginary eigenvalues, as long as they are the rightmost ones.

In this talk we will present extraction methods for the challenging open problem of how to efficiently find *purely imaginary eigenvalues that are in the interior* of the spectrum of a large sparse matrix.

We will also consider the "sister problem" of how to effectively find *real* eigenvalues. In several applications, real eigenvalues are the physically relevant ones. We will show that similar methods can be exploited as for the purely imaginary eigenvalues, but also that for real matrices, this problem can be handled with cheaper approaches.

If time allows, we will also give a wider overview of recent developments in harmonic Rayleigh–Ritz techniques and their applications, together with some other new extraction processes (e.g., [4] and work in progress).

Part of this talk is joint with Pierre-Antoine Absil (Louvaine-la-Neuve) and Elias Jarlebring (TU Braunschweig).

References

[1] C. Beattie, *Harmonic Ritz and Lehmann bounds*, Electron. Trans. Numer. Anal., 7 (1998), pp. 18–39.

[2] J. V. Burke, A. S. Lewis, and M. L. Overton, *Robust stability and a criss-cross algorithm for pseudospectra*, IMA J. Numer. Anal., 23 (2003), pp. 359–375.

[3] M. E. Hochstenbach, *Generalizations of harmonic and refined Rayleigh–Ritz*, Electron. Trans. Numer. Anal., 20 (2005), pp. 235–252.

[4] ——, *Variations on harmonic Rayleigh–Ritz for standard and generalized eigenproblems*, Preprint.

[5] R. B. Morgan, *Computing interior eigenvalues of large matrices*, Linear Algebra Appl., 154/156 (1991), pp. 289–309.

[6] C. C. Paige, B. N. Parlett, and H. A. van der Vorst, *Approximate solutions and eigenvalue bounds from Krylov subspaces*, Num. Lin. Alg. Appl., 2 (1995), pp. 115–133.

[7] G. W. Stewart, *Matrix algorithms. Vol. II*, Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 2001.

Frobenius Norm Minimization and Probing for Preconditioning

Thomas Huckle, Alexander Kallischko

In this paper we introduce a new method for defining preconditioners for the iterative solution of a system of linear equations $Ax = b$. The new method can be seen as a generalization of the class of preconditioners derived by Frobenius norm minimization by minimizing the more general problem

$$\min \| \begin{pmatrix} C \\ f^T \end{pmatrix} M - \begin{pmatrix} B \\ e^T \end{pmatrix} \|_F$$

for givern sparsity pattern in M, e.g. with $C = A$, $B = I$, and $f^T = e^T A$. To this aim we augment the Frobenius minimization by an additional norm minimization improving the preconditioner relative to a chosen subspace. This new method is closely related to preconditioners related to weighted Frobenius norm minimization (e.g. FSAI, SPAI, SPAI with target matrix), and also to modified preconditioners like MILU or MICC (where the chosen subspace is related to the vector $(1, 1, ..., 1)$). Furthermore, this augmented Frobenius norm minimization can be seen as a regularization of the interface probing for Schur complement systems allowing any set of probing vectors.

We develop a toolbox for computing preconditioners that are improved relative to a given small probing subspace. Furthermore, by this MSPAI (modified SPAI) probing approach we can improve any given preconditioner with respect to this probing subspace. All the computations are embarrassingly parallel. Additionally, for symmetric linear system we introduce new techniques for symmetrizing preconditioners.

The new approach can be used to derive improved preconditioners e.g. for Schur complement problems, but also for sparsification of dense or thick matrices where we are looking for a sparse approximation of the original problem.

Many numerical examples, e.g. from PDE applications such as domain decomposition and Stokes problem, show that these new preconditioners often lead to faster convergence and smaller condition numbers.

Rational Matrix Iterations in Numerical Linear Algebra

Bruno Iannazzo

Iterations of the kind $x_{k+1} = \varphi(x_k),\ x_0 \in \mathbf{C}$, where φ is a rational function, are called rational iterations and are commonly used in the numerical computation of the roots of polynomials. In fact, most of the root-finding algorithms for polynomials are based on rational iterations, among them there is the noticeable Newton method.

It is not surprising that matrix iterations, formally similar to rational iterations, arise in many problems of numerical linear algebra with important applications like solving matrix equations and computing matrix functions.

There are many ways to generalize a rational iteration to the matrix case. The most straightforward is to consider the iteration $X_{k+1} = \varphi(X_k), X_0 = A \in \mathbf{C}^{n\times n}$, where the function $\varphi(X_k)$ is applied to a square matrix as a matrix function. We call an iteration of this kind, *pure rational matrix iteration*.

Pure rational matrix iterations are used to define and to approximate the matrix sign function and for computing the matrix sector function. However, we give some restrictions on the use of that class: it is proved that the limit of a pure rational matrix iteration converging at least linearly is a diagonalizable matrix and can have only a limited number of eigenvalues independent of A.

In view of these restrictions, we consider a larger class of matrix iterations defined as $X_{k+1} = \psi(X_k, A)$, where ψ is a rational function of both its arguments, and X_0 is a function of A. This class reduces to the class of rational iterations in the scalar case and each iterate is a rational function of A. Iterations of this kind have been largely used to approximate matrix roots.

These iterations are attractive for their simplicity, but hide many difficulties; some of them are inherited from the scalar case; in fact, studying the local and global convergence of a rational iteration even in the scalar case is very difficult. But the matrix case needs a deeper analysis and raises new numerical problems due to the noncommutativity of matrix multiplication and to the possibility of having an infinite number of fixed points.

A thorough and general analysis of convergence and on the numerical behavior of these iterations has been started very recently. Two main problems are encountered.

1. Can matrix convergence be reduced to scalar convergence on eigenvalues?

2. Are matrix iterations numerically stable?

The answer to the first question is "yes" if the convergence is at least linear, this result is proved in [3] and we present another approach which gives more insights into the problem and can be generalized to holomorphic matrix iterations [5].

The answer to the second question is subtle because some iterations which are proved to converge quadratically, show instability when the matrix sequence approaches to the limit and this implies numerical divergence.

This phenomenon leads to the definition of *iteration stable in a neighborhood of a fixed point*, given in [1]. An iteration is said stable in a neighborhood of a fixed point S if its differential at S has bounded powers.

Using this definition, we prove that any pure rational matrix iteration is stable in a neighborhood of a fixed point (to which convergence is at least linear). This theorem provides a general strategy to make a rational iteration stable, an example is given in [4].

As particular topic, we present a class of rational iterations for computing matrix roots, based on the König family, a family of root-finding iterations due to Schröder [8]. We investigate convergence, stability and we give a theorem which singles out those iterations of the König family which preserve the structure of group of automorphisms associated with a scalar product.

We show also an unexpected application of this theorem to the theory of generally convergent root-finding algorithms. In fact, using the results of [2] based on the work of McMullen [6], we prove that the König family provides generally convergent root-finding algorithms of arbitrary high order of convergence for cubic polynomials.

[1] Sheung Hun Cheng, Nicholas J. Higham, Charles S. Kenney, and Alan J. Laub. *Approximating the logarithm of a matrix to specified accuracy.* SIAM J. Matrix Anal. Appl., 22(4):1112–1125, 2001.

[2] Jane M. Hawkins. *McMullen's root-finding algorithm for cubic polynomials.* Proceedings of the AMS, 130(9):2583–2592, 2002.

[3] N. Higham. *Functions of Matrices: Theory and Computation.* Book in preparation.

[4] B. Iannazzo. *On the Newton method for the matrix pth root.* SIAM J. Matrix Anal. Appl., 28(2):503–523, 2006.

[5] B. Iannazzo. *Notes on rational matrix iterations.* Preprint, Dipartimento di Matematica, Università di Pisa, 2007.

[6] B. Iannazzo. *Numerical solution of certain nonlinear matrix equations.* PhD Thesis, 2007.

[7] Curt McMullen. *Braiding of the attractor and the failure of iterative algorithms.* Invent. math. 91, 259–272, 1988.

[8] E. Schröder. *Ueber unendlich viele Algorithmen zur Auflösung der Gleichungen.* Mathematische Annalen, 1870.

Perturbation Bounds for Determinants and Characteristic Polynomials

Ilse Ipsen, Rizwana Rehmann

We derive absolute perturbation bounds for the coefficients of the characteristic polynomial of a complex matrix. The bounds consist of elementary symmetric functions of singular values, and suggest that coefficients of normal matrices are better conditioned with regard to absolute perturbations than those of general matrices. When the matrix is Hermitian positive-definite, the bounds can be expressed in terms of the coefficients themselves. We also improve absolute and relative perturbation bounds for determinants. The motivation for

the bounds is an application in quantum physics, where the coefficients of the characteristic polynomial supply information about thermodynamic properties of fermionic systems.

This is joint work with Rizwana Rehmann.

A question about Eigenvalues of Tridiagonal Matrices and Motivation for it

Charles Johnson

Given real numbers $t(1) < t(2) < ..... < t(2k-1)$, does there exist a k-by-k tridiagonal matrix with k of these numbers as eigenvalues and the remaining $k-1$ as eigenvalues, one each, of each of the $k-1$ proper leading principal submatrices.

Product Eigenvalue Problems in Applications: Computing Periodic Deflating Subspaces Associated with a Specified Set of Eigenvalues

Robert Granat, Bo Kågström, Daniel Kressner

Let us consider a *linear discrete-time descriptor system* of the form

$$\begin{aligned} E_k x_{k+1} &= A_k x_k + B_k u_k \\ y_k &= C_k x_k + D_k u_k \end{aligned}$$

with state, input and output vectors x_k, u_k, and y_k, respectively. The coefficient matrices A_k, B_k, C_k, D_k, E_k are supposed to be of matching dimensions. In addition, we assume that the system is *periodic* for some period $p \geq 1$, i.e., $A_{k+p} = A_k, B_{k+p} = B_k, C_{k+p} = C_k, D_{k+p} = D_k, E_{k+p} = E_k$ for all integers k. Computational tasks for such systems, which arise naturally from processes that exhibit seasonal or periodic behavior, can often be addressed by solving *product (or periodic) eigenvalue problems* [9]. For example, if all E_k are square and invertible then the system is *asymptotically stable* if and only if all eigenvalues of the monodromy matrix $\Pi = E_p^{-1} A_p E_{p-1}^{-1} A_{p-1} \cdots E_1^{-1} A_1$ lie strictly inside the unit disk.

Forming the product Π explicitly and applying a standard eigensolver may lead to disastrous numerical results. The most viable way to solve a general product eigenvalue problem is to compute a *generalized periodic real Schur form* (GPRSF), using the *periodic QZ algorithm* [2, 7]. The algorithm orthogonally transforms the matrix sequences E_k and A_k into upper (quasi-)triangular form to a cost only increasing linearly with p. In many applications, it is necessary to have the eigenvalues along the diagonal of the GPRSF in a certain order.

We present a direct method for reordering eigenvalues in the GPRSF of a regular K-cylic matrix pair sequence (A_k, E_k) [3, 4, 5]. Following and generalizing existing approaches, reordering consists of consecutively computing the solution to an associated periodic Sylvester-like matrix equation and constructing K pairs of orthogonal matrices. These pairs define an orthogonal K-cyclic equivalence transformation that swaps adjacent diagonal blocks in the GPRSF. An error analysis of this swapping procedure is presented, which extends existing results for reordering eigenvalues in the generalized real Schur form of a regular pair (A, E). Our direct reordering method is used to compute periodic deflating subspace pairs

corresponding to a specified set of eigenvalues. Associated perturbation theory is presented in [8, 1]. This computational task arises in various applications, e.g., solving discrete-time periodic Riccati equations [7]. We present computational experiments that confirm the stability and reliability of the eigenvalue reordering method, including periodic eigenvalue problems from real applications (see also the talk by Stefan Johansson).

Finally, we show how the methods discussed so far can be extended to general matrix products of the form

$$A_p^{s_p} A_{p-1}^{s_{p-1}} \cdots A_1^{s_1},$$

with the indices $s_1, \ldots, s_p \in \{1, -1\}$ and matching dimensions of $A_1, \ldots, A_p$. Here, we admit matrices A_k corresponding to an index $s_k = -1$ to be singular or even rectangular, in which case the matrix product should only be understood in a formal sense. Matlab tools for solving product eigenvalue problems of this most general form will also be reviewed [6].

References

[1] P. Benner, V. Mehrmann, and H. Xu. Perturbation analysis for the eigenvalue problem of a formal product of matrices. *BIT*, 42(1):1–43, 2002.

[2] A. Bojanczyk, G. H. Golub and P. Van Dooren (1992). The periodic Schur decomposition; algorithm and applications. In: *Proc. SPIE Conference*. Vol. 1770. pp. 31–42, 1992.

[3] R. Granat and B. Kågström. Direct Eigenvalue Reordering in a Product of Matrices in Periodic Schur Form. *SIAM J. Matrix Anal. Appl.* **28**(1), 285–300, 2006.

[4] R. Granat, B. Kågström and D. Kressner. Reordering the Eigenvalues of a Periodic Matrix Pair with Applications in Control. In: *Proc. of 2006 IEEE Conference on Computer Aided Control Systems Design (CACSD)*. pp. 25–30, 2006. ISBN:0-7803-9797-5.

[5] R. Granat, B. Kågström and D. Kressner. Computing Periodic Deflating Subspaces Associated with a Specified Set of Eigenvalues. *BIT Numerical Mathematics* (accepted), 2007.

[6] R. Granat, B. Kågström and D. Kressner. MATLAB Tools for Solving Periodic Eigenvalue Problems. In *Proc. of the 3rd IFAC Workshop on Periodic Control Systems*, St Petersburg, Russia, 2007.

[7] J. J. Hench and A. J. Laub. Numerical solution of the discrete-time periodic Riccati equation. *IEEE Trans. Automat. Control* **39**(6), 1197–1210, 1994.

[8] W.-W. Lin and J.-G. Sun. Perturbation analysis for the eigenproblem of periodic matrix pairs. *Linear Algebra Appl.* **337**, 157–187, 2001.

[9] A. Varga and P. Van Dooren (2001). Computational methods for periodic systems - an overview. In: *Proc. of IFAC Workshop on Periodic Control Systems, Como, Italy*. pp. 171–176, 2001.

A Multilevel, Modified Total Least Norm Approach to Signal and Image Deblurring

Misha E. Kilmer, Malena I. Español, Dianne P. O'Leary

In some signal and image deblurring applications, certain information about the blurring process is unavailable, so it is impossible to precisely model the blurring operator. Thus, the inverse problem of recovering the deblurred signal or image in the presence of noise is complicated by the need to model the uncertainty in the blurring operator, adding to the computational complexity of the inversion. In the case of an invertible blurring operator, the forward problem modeling the blurring of the true signal x is

$$(A + E)x \approx b_t + \eta = b, \qquad b_t = (A + E)x \tag{1}$$

where η denotes (unknown) noise corrupting the perfect data b_t, E is an unknown perturbation of the assumed (known) blurring operator A, and b denotes the blurred, noisy signal.

The goal is to recover an approximation to x given the perturbations to both A and b_t. One difficulty is that the matrix A is ill-conditioned with no gap in the singular value spectrum. Regularized total least squares (R-TLS) and total least norm (R-TLN) approaches have been presented in the literature which require the minimization of a functional with respect to the unknown perturbation in the blurring operator, the residual, and the desired image. If Tikhonov regularization is employed, one wishes to solve [1, 2]

$$\min_{x,E} \left\| \begin{array}{c} r \\ \mathrm{vec}(E) \end{array} \right\|_p^p + \lambda^q \|Lx\|_q^q, \tag{2}$$

where $r = b - (A + E)x$, L is a regularization operator, such as a discrete first derivative operator, and $\lambda > 0$ is the regularization parameter. This equation describes the R-TLS problem when $p = q = 2$ and the R-TLN problem otherwise. The regularization term has a good deal of influence over the quality of the estimates of x and E.

Much of the work to date in solving (2) in imaging applications has required the perturbation operator to have special structure (e.g. sparsity structure or Toeplitz type structure) in order to make the minimization problem more computationally tractable (see, for example, [3, 2]) since E can then be written in terms of a small number of parameters [4]. The long-term goal of our research is to lift this operator restriction while preserving computational feasibility and improving the quality of the solutions to both E and x.

As a first step towards attaining this goal, we present a multilevel approach to finding regularized solutions to (1). We rely on a Haar wavelet decomposition to move between grids. By transforming (1) into the wavelet domain and exploiting properties of the initial estimate, A, of the blurring operator, we derive equations governing the coarse grid solve and the residual correction. We show that solving the coarse grid and residual correction problems each involve solving a problem of the form (2), except that r is replaced by a slightly perturbed definition.

We then show that one of the benefits of using Haar wavelets to move between grids is that when one does wish to enforce a Toeplitz-type structure on E, this structure is preserved on all levels. Exploiting the assumption that E has Toeplitz structure, we give a computationally efficient algorithm for the residual-correction modified-TLS problem that must be solved at each level. Finally, we present results that indicate the promise of this approach on deblurring signals with edges.

References

[1] Gene Golub, Per Christian Hansen and Dianne P. O'Leary, "Tikhonov Regularization and Total Least Squares," SIAM J. on Matrix Analysis and Applications, Vol. 21, No. 1, (1999), Pages 185-194.

[2] Armin Pruessner and Dianne P. O'Leary, "Blind Deconvolution Using a Regularized Structured Total Least Norm Approach," SIAM J. on Matrix Analysis and Applications, Vol. 24, (2003), Pages 1018-1037.

[3] Haoying Fu and Jesse Barlow, "A Regularized Structured Total Least Squares Algorithm for High-Resolution Image Reconstruction," Linear Algebra and its Applications, Vol. 391 (2003) Pages 75-98. Special Issue on Linear Algebra in Signal and Image Processing.

[4] J. B. Rosen, H. Park, J. Glick, "Total Least Norm Formulation and Solution for Structured Problems," SIAM J. on Matrix Analysis and Applications, Vol. 17, (1996), Page 110-126.

Rayleigh-Ritz majorization error bounds with applications to FEM and subspace iterations

Andrew V. Knyazev, Merico E. Argentati

The Rayleigh-Ritz method finds the stationary values, called Ritz values, of the Rayleigh quotient on a given trial subspace as approximations to eigenvalues of a Hermitian operator A. If the trial subspace is A-invariant, the Ritz values are some of the eigenvalues of A. Given two subspaces X and Y of the same finite dimension, such that X is A-invariant, the absolute changes in the Ritz values of A with respect to X compared to the Ritz values with respect to Y represent the absolute eigenvalue approximation error. A recent paper [1] bounds the error in terms of the principal angles between X and Y using weak majorization. Here, we improve the constant in this bound and extend it to the practically important case where $dimX < dimY$. We present our Rayleigh-Ritz majorization error bound in the context of the finite element method (FEM), and show how it can improve known FEM eigenvalue error bounds. We derive a new majorization-type convergence rate bound of subspace iterations and combine it with the previous result to obtain a similar bound for the block Lanczos method. This presentation is based mostly on [2]. A corresponsing result where neither X nor Y is A-invariant can be found in [3]. The case of infinite dimensional subspaces is considered in [4].

References.

[1] M. E. Argentati, A. V. Knyazev, C. C. Paige, and I. Panayotov, Bounds on changes in Ritz values for a perturbed invariant subspace of a Hermitian matrix, http://arxiv.org/abs/math/0610498.

[2] A. V. Knyazev and M. E. Argentati, Rayleigh-Ritz majorization error bounds with applications to FEM and subspace iterations, http://arxiv.org/abs/math/0701784.

[3] A. V. Knyazev and M. E. Argentati, Majorization for Changes in Angles Between Subspaces, Ritz values, and graph Laplacian spectra, SIAM Journal on Matrix Analysis and Applications, 29 (2006), no. 1, pp. 15-32.

[4] A. V. Knyazev, A. Jujunashvili, and M. E. Argentati, Angles Between Infinite Dimensional Subspaces with Applications to the Rayleigh-Ritz and Alternating Projectors Methods, http://arxiv.org/abs/0705.1023.

Computing Eigenvalues of Random Matrices

Plamen Koev

While the subject of numerically computing the eigenvalues of matrices is quite mature, the problem of computing eigenvalue *distributions* of *random* matrices is still in its infancy. The problem is to compute the distributions of the eigenvalues given the distributions of the matrix entries for certain random matrix classes, e.g., Wishart, Jacobi, etc.

Somewhat surprisingly, explicit formulas for these eigenvalue distributions have been available since the 1960s, but only in terms of the *hypergeometric function of a matrix argument* – a notoriously slowly converging series of multivariate orthogonal polynomials called Jack functions. The efficient evaluation of such distributions had eluded the researchers for over 40 years even though the matrices involved are sometimes as small as 3×3.

The proposed talk will be accessible to anyone with numerical linear algebra knowledge as I will present the (fairly straightforward) derivation of the explicit formulas for the eigenvalue distributions of random matrices. In the main part of the talk, I will present the connections to combinatorics, representation theory and fast Fourier transforms, which we used to recently develop the first efficient algorithms for these distributions.

Our new algorithms had immediate effect to applications in multivariate statistics as techniques that had long been developed in theory could finally be used in practice: Multivariate statistics is concerned with the extraction of meaningful information from enormous datasets, most often organized as a matrix X of n observations on m objects. Various questions of correlation, interdependency, and classification of the data are critical in applications ranging from genomics to biostatistics and wireless communications to image and target recognition.

As a mathematical model for the above setup one considers a matrix of normal entries with arbirtary covariances $A \sim N_m(0, \Sigma)$. The matrix $A^T A$ is called Wishart.

The statistics used to test various hypothesis on the original datasets are usually the extreme eigenvalues (or functions thereof) of the matrix $X^T X$. They are tested against the theoretical predictions for the Wishart matrix $A^T A$.

Thus the importance of being able to compute the theoretically predicted eigenvalue distributions.

The applications that have been enabled by our new algorithms range from genomics to target classification which I will also briefly describe.

I will also outline a number of prominent computational problems that remain open.

On the Growth Factor for Hadamard Matrices

Christos Kravvaritis, Marilena Mitrouli

Gaussian Elimination (GE) with some pivoting technique is the most famous method for solving linear systems. However, its stability is governed by the growth factor, whose values are not specified generally [11]. For a *completely pivoted* matrix A (CP, i.e. no row and column

exchanges are needed during GE with complete pivoting), the backward error analysis for GE is expressed in terms of the *growth factor*

$$g(n, A) = \frac{\max\{p_1, p_2, \ldots, p_n\}}{|a_{11}^{(1)}|},$$

where $p_1, p_2, \ldots, p_n$ are the pivots of A.

It is interesting to study the growth factor of Hadamard matrices, as they are the only known matrices that attain growth factor equal to their order. Besides, they are an attractive tool for several applications [6, 12]. A *Hadamard matrix* H of order n (denoted by H_n) has entries ± 1 and satisfies $HH^T = H^T H = nI_n$. They are defined only for $n = 1, 2$ and $n \equiv 0 \bmod 4$, however the existence of relatively small Hadamard matrices ($n = 668$) is still not known [6].

In 1964 Tornheim [10] proved that $g(n, H) \geq n$ for a CP $n \times n$ Hadamard matrix H. In 1968 Cryer [1] conjectured that

"$g(n, A) \leq n$, with equality iff A is a Hadamard matrix".

20 years later, in 1988 Day and Peterson [2] proved the equality only for the Hadamard-Sylvester class [6]. They provided some experimental results for pivot patterns of 16×16 Hadamard matrices and conjectured that the fourth pivot from the end is $n/4$, which was shown to be false in 1998 by Edelman and Friedman [4], who found the first H_{16} with fourth pivot from the end $n/2$. In 1991 Gould discovered a 13×13 matrix with growth factor 13.0205 [5]. Thus the first part of Cryer's conjecture was shown to be false. The second part of the conjecture concerning the growth factor of Hadamard matrices still remains an open problem.

Over all these years small progress has been made towards the Hadamard part of Cryer's conjecture. In 1995 Edelman and Mascarenhas [3] proved that the growth factor of H_{12} is 12 by demonstrating its unique pivot pattern. As it is already stated in [3],

"it appears very difficult to prove that $g(16, H_{16}) = 16$".

A great difficulty arises at the study of this problem because H-equivalence operations (row and column interchanges and multiplications by -1) do not preserve pivots, i.e. H-equivalent matrices do not have necessarily the same pivot pattern. So, for the case of proving the pivot structures of H_{16}, a naive computer exhaustive search performing all possible H-equivalence operations would require $(16!)^2(2^{16})^2 \approx 10^{36}$ trials.

The following lemma constitutes a powerful tool for this research, since it offers a possibility for calculating pivots in connection with minors.

Lemma 1. [1],[2] *Let A be a CP matrix. The magnitude of the pivots appearing after application of GE operations on A is given by*

$$p_j = \frac{A(j)}{A(j-1)}, \quad j = 1, 2, \ldots, n, \quad A(0) = 1,$$

where $A(j)$ denotes the absolute value of the $j \times j$ principal minor of A.

So, it is obvious that the calculation of minors of Hadamard matrices is important in order to study pivot structures, and moreover the growth factor for CP Hadamard matrices.

The first known effort for calculating minors of Hadamard matrices was accomplished in 1907 by Sharpe [9] for $n-1$, $n-2$ and $n-3$ minors. In 2001 Koukouvinos et al. [7] developed

theoretically a method for evaluating all possible $(n-j)\times(n-j)$, $j \geq 1$, minors of Hadamard matrices, which could be generalized as an algorithm. The existing results combined with new, sophisticated techniques of computational linear algebra yielded minors of H_{16}. So, in 2007 [8] all 34 possible pivot patterns of H_{16} were demonstrated and the complete pivoting conjecture for H_{16} was proved.

In this talk we will present briefly the ideas developed in [8] and explain how it was possible to dodge the above mentioned forbidding complexity. We will conclude by discussing the following open problems.

Relative open problems.

- It is known that H_{16} matrices can be classified with respect to the H-equivalence in 5 classes I,...,V, see [12]. An open problem is to classify theoretically the appearing pivot patterns in the appropriate equivalence classes. A possible approach is to deal with this problem in connection with the existence of D-optimal designs (± 1 matrices with maximum determinant) embedded in H_{16} matrices. For example, in order to prove the non occurrence of $p_6 = 10/3$ in I-Class (Sylvester class), it is equivalent to show that the D-optimal design of order 6 cannot exist inside a matrix of this class. Or, whenever $p_5 = 3$, equivalently this means that the D-optimal design of order 5 exists embedded in these H_{16}.

- The methods used so far can be the basis for calculating the pivot pattern of Hadamard matrices of higher orders, such as H_{20}, H_{24} etc. The complexity of such problems points out the need for developing algorithms that can implement very effectively the proposed ideas, or other, more elaborate ideas. For example, a parallel implementation, which would limit significantly the computational time needed by the algorithms, is under investigation.

- Finally, it is a mystery why the value 8 as fourth pivot from the end appears only for matrices from the Hadamard-Sylvester class, and specifically only in one pivot pattern. It is interesting to study this issue and to find out whether the construction properties of this class can explain this exceptional case.

References

[1] Cryer, C.W.: Pivot size in Gaussian elimination. Numer. Math. **12**, 335–345 (1968)

[2] Day, J., Peterson, B.: Growth in Gaussian Elimination. Amer. Math. Monthly **95**, 489–513 (1988)

[3] Edelman, A., Mascarenhas, W.: On the complete pivoting conjecture for a Hadamard matrix of order 12. Linear Multilinear Algebra **38**, 181–187 (1995)

[4] Edelman, A., Friedman, D.: A counterexample to a Hadamard matrix pivot conjecture. Linear Multilinear Algebra **44**, 53–56 (1998)

[5] Gould, N.: On growth in Gaussian elimination with pivoting. SIAM J. Matrix Anal. Appl. **12**, 354–361 (1991)

[6] Horadam, K.J.: Hadamard matrices and their appplications, Princeton University Press, Princeton (2007)

[7] Koukouvinos, C., Mitrouli, M., Seberry, J.: An algorithm to find formulae and values of minors of Hadamard matrices. Linear Algebra Appl. **330**, 129–147 (2001)

[8] Kravvaritis, C., Mitrouli, M.: On the growth problem for a Hadamard matrix of order 16. (2007), submitted

[9] Sharpe, F.R.: The maximum value of a determinant. Bull. Amer. Math. Soc. **14**, 121–123 (1907)

[10] Tornheim, L.: Pivot size in Gauss reduction. Tech. Report, Calif. Res. Corp., Richmond, Calif., February 1964

[11] Trefethen, L.N.: Three mysteries of Gaussian elimination. ACM SIGNUM Newsletter **20**, 2–5 (1985)

[12] Wallis, W.D., Street, A.P., Wallis, J.S.: Combinatorics: Room Squares, Sum-Free Sets, Hadamard Matrices (Lectures notes in mathematics 292). Springer-Verlag, New York (1972)

Computation of Structured Pseudospectra

Michael Karow, Daniel Kressner

The standard definition of the ϵ-pseudospectrum,

$$\Lambda_\epsilon(A) = \cup\{\Lambda(A+E) \,|\, E \in \mathbb{C}^{n\times n}, \|E\| \le \epsilon\},$$

imposes no structure on the perturbations of A. We list some applications below, where it is more reasonable to restrict $A+E$ to a linear subspace or submanifold $\mathcal{M}$ of $\mathbb{C}^{n\times n}$:

$$\Lambda_\epsilon^{\mathcal{M}}(A) = \cup\{\Lambda(A+E) \,|\, A+E \in \mathcal{M}, \|E\| \le \epsilon\}.$$

Eigenvalue Computation A backward stable eigenvalue solver computes the exact eigenvalue of a slightly perturbed matrix $A+E$. In many instances, there is more information besides $\|E\| \le \epsilon$ available. For example, if A is real, most solvers (with the exception of standard Jacobi-Davidson and other inexact inverse iteration techniques working with a single shift) satisfy a real backward error. Using $\Lambda_\epsilon^{\mathcal{M}}(A)$ with an appropriate choice of $\mathcal{M}$ instead of $\Lambda_\epsilon(A)$ then admits tighter estimates of the expected forward error.

Interconnected systems Perturbing only the subsystems of an interconnected linear time-invariant control system corresponds to a block diagonal perturbation of the system matrix. Investigating for which value of ϵ the set $\Lambda_\epsilon^{\text{block diagonal}}(A)$ crosses the imaginary axis is a classical question in robust control theory.

Pseudospectra of nonstandard eigenvalue problems Polynomial and product eigenvalue problems can be embedded into standard eigenvalue problems by means of linearization and lifting, respectively. This process induces structure in A, which should be taken into account by the pseudospectrum in order to admit conclusions on the eigenvalue sensitivities of the original problem. Also, comparing the sizes of the corresponding $\Lambda_\epsilon^{\mathcal{M}}(A)$ and $\Lambda_\epsilon(A)$ reveals whether this process can be expected to lead to numerical difficulties.

If time permits, we will also discuss the use of structured pseudospectra for Hamiltonian systems whose flow is approximated on a symplectic submanifold.

After computing the Schur decomposition of A, it requires $O(n^2)$ operations to check whether a point in the complex plane belongs to $\Lambda_\epsilon(A)$. This technique is successfully employed in the software package Eigtool. The computational cost for checking whether a point belongs to $\Lambda_\epsilon^{\mathcal{M}}(A)$ heavily depends on the structure. It can range from $O(n^2)$ (linearization of matrix polynomials) to NP-hard (block diagonal with more than two blocks). In many instances, a computational procedure is not even known.

In this talk, we present $O(n^2)$ algorithms for real, symmetric, skew-symmetric, and Hermitian structures. It is planned to include these algorithms in a structured variant of Eigtool. To this end, some additional issues need to be addressed. In particular, the lack of Lipschitz continuity forces the use of adaptive refinement techniques to avoid artificial spikes when plotting structured pseudospectra.

Robust Computation of Off-Diagonal Elements of Functions of Matrices

James V. Lambers

1 Introduction

In [2] Golub and Meurant describe a method for computing quantities of the form

$$\mathbf{u}^T f(A)\mathbf{v}, \tag{1}$$

where $\mathbf{u}$ and $\mathbf{v}$ are N-vectors, A is an $N \times N$ symmetric positive definite matrix, and f is a smooth function. Let $a = \lambda_N$ be the smallest eigenvalue, and $b = \lambda_1$ be the largest eigenvalue. Then, the quantity (1) can be viewed as a Riemann-Stieltjes integral

$$\mathbf{u}^T f(A)\mathbf{v} = I[f] = \int_a^b f(\lambda)\, d\alpha(\lambda). \tag{2}$$

provided that the measure $\alpha(\lambda)$, which is defined in terms of the spectral decomposition of A, is positive and increasing.

As discussed in [2], the integral $I[f]$ can be bounded using either Gauss, Gauss-Radau, or Gauss-Lobatto quadrature rules, whose nodes, as well as the weights, can be obtained using the symmetric Lanczos algorithm if $\mathbf{u} = \mathbf{v}$, and the unsymmetric Lanczos algorithm if $\mathbf{u} \neq \mathbf{v}$ (see [4]). In the case $\mathbf{u} \neq \mathbf{v}$, there is the possibility that the weights may not be positive, which destabilizes the quadrature rule (see [1] for details). One way to handle this case is to rewrite (1) using a decomposition such as

$$\mathbf{u}^T f(A)\mathbf{v} = \frac{1}{\delta}[\mathbf{u}^T f(A)(\mathbf{u} + \delta\mathbf{v}) - \mathbf{u}^T f(A)\mathbf{u}], \tag{3}$$

where δ is a small constant.

While a sufficiently small δ will ensure positive weights, a choice that is too small can lead to catastrophic cancellation, as in the approximation of derivatives by finite differences. In fact, this decomposition leads to an approximation of the derivative

$$\frac{d}{d\delta}\left[\mathbf{u}^T(\mathbf{u}+\delta\mathbf{v})\mathbf{e}_1^T f[T(\delta)]\mathbf{e}_1\right]\Big|_{\delta=0}, \tag{4}$$

where $T(\delta)$ is the tridiagonal matrix output by the unsymmetric Lanczos algorithm applied to the matrix A with starting vectors $\mathbf{u}$ and $(\mathbf{u}+\delta\mathbf{v})$ (which reduces to the symmetric Lanczos algorithm for $\delta = 0$). Our goal is to develop an efficient algorithm for computing this derivative analytically.

2 Derivatives of the Nodes and Weights

For a given δ, let λ_j, $j = 1, \ldots, K$, be the nodes of the K-point Gaussian rule obtained by applying the unsymmetric Lanczos algorithm to A with starting vectors $\mathbf{u}$ and $(\mathbf{u}+\delta\mathbf{v})$. Let w_j, $j = 1, \ldots, K$, be the corresponding weights. Then, letting $\delta \to 0$, we obtain the following, assuming all required derivatives exist:

$$\mathbf{u}^T f(A)\mathbf{v} \approx \mathbf{u}^T\mathbf{v}\sum_{j=1}^{K} w_j f(\lambda_j) + \mathbf{u}^T\mathbf{u}\sum_{j=1}^{K}[w_j' f(\lambda_j) + w_j\lambda_j' f'(\lambda_j)] \tag{5}$$

where the $'$ denotes differentiation with respect to δ, and evaluation of the derivative at $\delta = 0$. Equivalently, these derivatives are equal to the length of $\mathbf{v}$ times the directional derivatives of the nodes and weights, in the direction of $\mathbf{v}$, and evaluated at the origin.

¿From the Lanczos algorithm, $T(\delta)$ has entries that are functions of δ. Because the nodes and weights are obtained from the eigenvalues and eigenvectors of this matrix, it is desirable to use these relationships to develop efficient algorithms for computing the derivatives of the nodes and weights in terms of those of the recursion coefficients. Such algorithms are described in [6].

3 Derivatives of the Recursion Coefficients

Let A be a symmetric positive definite $n \times n$ matrix and let $\mathbf{r}_0$ be an n-vector. Suppose that we have already carried out the symmetric Lanczos iteration to obtain orthogonal vectors $\mathbf{r}_0, \ldots, \mathbf{r}_K$ and the Jacobi matrix T_K. Now, suppose that we wish to compute the entries of the modified matrix $\hat{T}_K$ that results from applying the unsymmetric Lanczos iteration with the same matrix A and the initial vectors $\mathbf{r}_0$ and $\mathbf{r}_0 + \mathbf{f}$, where $\mathbf{f}$ is a given perturbation. A Lanczos-like iteration, introduced in [8] and based on ideas from [3], produces these values directly from a Krylov subspace whose initial vector is $\mathbf{f}$.

In [8], this iteration was used to efficiently obtain the recursion coefficients needed to approximate $\mathbf{u}^T f(A)(\mathbf{u}+\delta\mathbf{v})$ from those used to approximate $\mathbf{u}^T f(A)\mathbf{u}$, where $\mathbf{u}$ depended on a parameter. Here, we use this iteration for a different purpose. From the expressions for the entries of $\hat{T}_K$, the derivatives of the recursion coefficients α_j, $j = 1, \ldots, K$, and β_j, $j = 1, \ldots, K-1$, can be obtained by setting $\mathbf{r}_0 = \mathbf{u}$ and $\mathbf{f} = \delta\mathbf{v}$. By differentiating the recurrence relations in the abovementioned iteration with respect to δ and evaluating at $\delta = 0$, we obtain a new set of recurrence relations that produce the derivatives (see [6] for details).

4 Applications

In [6], this approach to computing (1) was used to reformulate Krylov subspace spectral (KSS) methods. These methods, introduced in [7], [5] are designed to solve time-dependent variable-coefficient PDE (both parabolic and hyperbolic). The reformulation improves the numerical stability of these methods, and also reveals that they are actually high-order operator splittings. This perspective facilitates stability analysis, as demonstrated in [6], where it is shown that under reasonable assumptions on the coefficients of the spatial operator, 1-node KSS methods are unconditionally stable, and the 2-node method for parabolic problems has an approximate solution operator that is bounded independently of the time step and grid size.

Current work is focusing on the efficient construction of approximations to products of functions of A and $\mathbf{v}$, such as $A^{-1}\mathbf{v}$ or $e^{-At}\mathbf{v}$. If we consider the approximation of (1), where $\mathbf{u}$ is set to each member of a basis, then such an approximation can be computed from a Krylov subspace

$$\mathcal{K}(A, \mathbf{v}, K) = \{\mathbf{v}, A\mathbf{v}, A^2\mathbf{v}, \ldots, A^{K-1}\mathbf{v}\}$$

and diagonal matrices obtained from the nodes and weights, as well as their derivatives, for each choice of $\mathbf{u}$, as opposed to the approach of computing such approximations from a linear combination of vectors in a basis for $\mathcal{K}(A, \mathbf{v}, K)$. As demonstrated in [6] and [8], for the case where A is a spectral discretization of a differential operator, the greater flexibility in our approach has resulted in an explicit time-stepping scheme for PDE that exhibits the stability of implicit methods. Therefore, it makes sense to consider whether this approach can be applied to more general matrices, especially those that are ill-conditioned, to develop algorithms for computing $f(A)\mathbf{v}$ that are more accurate and robust than a straightforward use of $\mathcal{K}(A, \mathbf{v}, K)$.

References

[1] K. Atkinson, *An Introduction to Numerical Analysis, 2nd Ed.* (1989) Wiley.

[2] G. H. Golub and G. Meurant, "Matrices, Moments and Quadrature", in *Proceedings of the 15th Dundee Conference*, June-July 1993, D. F. Griffiths and G. A. Watson (eds.), Longman Scientific & Technical, 1994.

[3] G. H. Golub and M. H. Gutknecht, "Modified Moments for Indefinite Weight Functions", *Numer. Math.* **57** (1989), p. 607-624.

[4] G. H. Golub and J. Welsch, "Calculation of Gauss Quadrature Rules", *Math. Comp.* **23** (1969), 221-230.

[5] P. Guidotti, J. V. Lambers and K. Sølna, "Analysis of 1-D Wave Propagation in Inhomogeneous Media", *Num. Funct. Anal. Opt.* **27** (2006), p. 25-55.

[6] J. V. Lambers, "Derivation of High-Order Spectral Methods for Variable-Coefficient PDE Using Modified Moments", *Electr. Trans. Num. Anal.* **28** (2007), to appear.

[7] J. V. Lambers, "Krylov Subspace Spectral Methods for Variable-Coefficient Initial-Boundary Value Problems", *Electr. Trans. Num. Anal.* **20** (2005), p. 212-234.

[8] J. V. Lambers, "Practical Implementation of Krylov Subspace Spectral Methods", *J. Sci. Comput.* **32** (2007), p. 449-476.

Reduce factorizations

Jim Demmel, Laura Grigori, Hoemmen, Julien Langou

QR factorizations of tall and skinny matrices with their data partitioned vertically across several processors arise in a wide range of applications. Various methods exist to perform the QR factorization of such matrices: Gram-Schmidt, Householder, or CholeskyQR. In this talk, we present the Allreduce Householder QR factorization. This method is stable and performs, in our experiments, from four to eight times faster than ScaLAPACK routines on tall and skinny matrices.

The Cholesky algorithm first forms the normal equations ($C = A^T A$), then computes the Cholesky factor of the normal equations ($R = \texttt{chol}(C)$), finally, if the Q-factor is desired, computes Q with $Q = A/R(= AR^{-1})$. This algorithm is briefly explained in Gander [2], then in more details in Björck [1]. Stathopoulos and Wu [5] presents an extensive bibliography, some applications and recent computer experiments to motivate the use of the CholeskyQR algorithm. The CholeskyQR algorithm has been widely popularized by recent block eigenvalue software (e.g. ANASAZI, BLOPEX and PRIMME).

The CholeskyQR algorithm has three main advantages:

1. it is extremely fast and scalable. There are two reasons for this: first, it only uses fast kernel operations (`SYRK`, `CHOL`, and `GEMM`). And, second, it requires only one synchronization point, which makes it extremely attractive on latency limited machines,

2. the algorithm relies on efficient and widespread kernels (MPI, BLAS, and LAPACK), so that its performance is portable,

3. the algorithm is extremely simple and can be encoded in no more than four lines.

Unfortunately, the CholeskyQR algorithm is unstable. In large eigenvalue computation, it has been observed [3, 4] that eigensolvers can fail to converge or have a considerably slower convergence when the orthogonality of the Ritz vectors is poor.

In this paper, we generalize the *one-synchronization-point* property of CholeskyQR algorithm to the Householder QR factorization. Since our algorithm will be based on Householder transformations, it will be stable. The algorithm is based on reduce elimination tree as Cholesky QR, so the name Reduce Householder factorization.

References

[1] Å. Björck. *Numerical Methods for Least Squares Problems.* Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 1996.

[2] W. Gander. Algorithms for the QR decomposition. Tech. Report 80-02, ETH, Zürich, Switzerland, 1980.

[3] U. Hetmaniuk and R. Lehoucq. Basis selection in LOBPCG. *Journal of Computational Physics*, 218:324–332, 2006.

[4] A. V. Knyazev, M. Argentati, I. Lashuk, and E. E. Ovtchinnikov. Block locally optimal preconditioned eigenvalue xolvers (BLOPEX) in HYPRE and PETSc). Technical report, University of Colorado at Denver and Health Sciences Center, 2006. Submitted to SISC,

[5] A. Stathopoulos and K. Wu. A block orthogonalization procedure with constant synchronization requirements. *SIAM Journal on Scientific Computing*, 23(6):2165–2182, 2002.

[6] L. N. Trefethen and D. Bau III. *Numerical Linear Algebra*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 1997.

On nonsymmetric saddle point matrices that allow conjugate gradient iterations

Jörg Liesen, Beresford N. Parlett

Many applications in science and engineering require solving large linear algebraic systems in saddle point form; see [1] for an extensive survey. Often the system matrix is of the form

$$\begin{bmatrix} A & B^T \\ B & -C \end{bmatrix}, \tag{1}$$

where $A = A^T \in \mathbf{R}^{n\times n}$ is positive definite ($A > 0$), $B \in \mathbf{R}^{m\times n}$ has rank $r \leq m \leq n$, and $C = C^T \in \mathbf{R}^{m\times m}$ is positive semidefinite ($C \geq 0$). The matrix in (1) is congruent to the block diagonal matrix $[A\ 0;\ 0\ S]$, where $S = -(C + BA^{-1}B^T)$, so that $-S = -S^T \geq 0$. Hence this matrix is symmetric *indefinite*, with n positive and $\mathrm{rank}(S)$ negative eigenvalues. Typically, $\mathrm{rank}(S)$ is close to m and therefore, unless m is very small, the matrix in (1) is highly indefinite. This feature is a major challenge for iterative solvers such as Krylov subspace methods.

It has been noted by several authors (see [1, p. 23] for references) that, for solving linear algebraic systems, there is no harm in negating the second block row in (1), and using as system matrix

$$\mathcal{A} \equiv \begin{bmatrix} A & B^T \\ -B & C \end{bmatrix}. \tag{2}$$

Symmetry has been abandoned but what, if anything, has been gained?

For a moment, suppose that $B = 0$. Then, due to the assumptions on A and C, the matrix $\mathcal{A}$ in (2) is symmetric *positive semidefinite*, while the saddle point matrix in (1) still is symmetric indefinite. Hence, for $B = 0$, the eigenvalues of $\mathcal{A}$ are real and nonnegative, which from the point of view of an iterative solver is certainly more desirable than having eigenvalues on both sides of the origin. From this point of view, negating the second block row appears to be the canonical approach[2]. Of course, as $\|B\|$ rises from zero, the matrix

[2]Prof. Karl Peter Grotemeyer (08.09.1927 – 30.07.2007), a German mathematician who served as president of my alma mater, the University of Bielefeld, for 22 years, once defined "canonical approach" as follows: "Everyone who is sane would do that same thing."

$\mathcal{A}$ becomes nonsymmetric, and naturally at some point the eigenvalues are no longer real (and/or $\mathcal{A}$ is no longer diagonalizable). But, until then, $\mathcal{A}$ is a nice matrix. In this talk we will discuss how to exploit these observations.

First, we will show that there exists a "magic" hidden bilinear form with respect to which $\mathcal{A}$ is symmetric. We will discuss necessary and sufficient conditions so that this bilinear form is positive definite and thus defines a proper inner product. These conditions turn out to be relations between $\|B\|$ and the extreme eigenvalues of A and C, which should be expected from the above motivation. For example, it is sufficient that

$$2\|B\| \ < \ \lambda_{\min}(A) - \lambda_{\max}(C)\,,$$

where $\lambda_{\min}$ and $\lambda_{\max}$ denote largest and smallest eigenvalue, respectively. When the conditions are satisfied, a well defined conjugate gradient (CG) method for solving linear systems with $\mathcal{A}$ exists. We will show an efficient implementation of such method, and discuss practical issues such as error bounds, preconditioning and numerical examples. Many of these results can be found in [2].

Next, we will present results of our current investigations of the cases when the conditions are no longer satisfied, i.e. when $\|B\|$ is too large to make $\mathcal{A}$ a nice matrix. Here our bilinear form is not an inner product, and our CG method may break down. However, even if the CG method breaks down at some step k, the CG iterate at step $k+1$ may be well defined. Thus, from a successful step $k-1$ it is possible to compute step $k+1$ and bypass the ill-defined step k. We call this approach "skip ahead CG". We will describe the derivation of the corresponding algorithm, and study it numerically.

References

[1] M. Benzi, G. H. Golub, and J. Liesen, *Numerical solution of saddle point systems*, Acta Numer., 14 (2005), pp. 1–137.

[2] J. Liesen and B. N. Parlett, *On nonsymmetric saddle point matrices that allow conjugate gradient iterations*, submitted, 2007 (available online at `www.math.tu-berlin.de/∼liesen`).

Theory and Algorithms for Structured Matrix Problems

Niloufer Mackey

Polynomial eigenproblems (PEPs) are playing an increasingly important role in the engineering and applied sciences, especially in extreme designs. These designs often lead to eigenproblems with poor conditioning, while the physics of the system leads to algebraic structure that numerical methods should preserve if they are to provide physically meaningful results. One of the aims of my research is to increase our theoretical understanding of these structures to aid in the development of effective algorithms for structured PEPs that arise in applications.

The standard way to solve PEPs is by linearization, that is, by finding a linear polynomial with the same spectrum, with the most frequently used linearizations being the first and second companion forms. Together with D. S Mackey, C. Mehl and V. Mehrmann, a new theory

of linearizations has been developed showing how to systematically construct a continuum of linearizations by generalizing the two companion forms to form two subspaces, denoted by $\mathbb{L}_1(P)$ and $\mathbb{L}_2(P)$, of the space of $kn \times kn$ pencils [8]. Almost all pencils in these roughly $\frac{1}{2}k^2n^2$-dimensional subspaces are linearizations of the original matrix polynomial $P(\lambda)$, and their k-dimensional intersection, $\mathbb{DL}(P) = \mathbb{L}_1(P) \cap \mathbb{L}_2(P)$ has proved to be a particularly fertile hunting ground for finding structure-preserving linearizations for palindromic, alternating, and symmetric polynomials [3], [9]. Linearizations in these subspaces that have nearly optimal conditioning have been identified for certain cases [4].

We are currently investigating new sources of linearizations in the following contexts:

- The analysis in [9] shows that the presence of certain exceptional eigenvalues may prevent some structured polynomials from having any structured linearization in $\mathbb{L}_1(P)$ and $\mathbb{L}_2(P)$. Furthermore, the presence of an eigenvalue λ close to an exceptional eigenvalue λ_0 will result in λ being poorly conditioned in the structured linearization.

- Recently discovered linearizations in [1] also derive from the companion forms, but in a different way. Finding a common generalization that encompasses both approaches could lead to new structured linearizations that avoid the difficulties associated with exceptional eigenvalues.

- Recent work of Jarlebring [5], [6] on linear (retarded and neutral) differential-delay equations has brought to light a new type of structured QEP with properties strongly reminiscent of palindromic structure. These polynomials satisfy $Q \operatorname{rev} \overline{P}(\lambda) Q = P(\lambda)$, where Q is any involution, and rev reverses the coefficients of the polynomial. Together with D. S. Mackey and H. Faßbender, we are working on classifying structure-preserving linearizations for this new class of polynomials, and identifying those that are well-conditioned.

Another essential step is to find useful condensed forms for these structured linearizations. Recent investigations [10] for palindromic pencils have led to an anti-triangular Schur form that seems well adapted to numerical computations.

Structured Mapping Problems: Let $\mathbb{S} \subset \mathbb{K}^{n\times n}$ be a class of structured matrices and consider the following structured mapping problems:

Existence: For which vectors x, b does there exist some $A \in \mathbb{S}$ such that $Ax = b$?

Characterization: Determine the set $\mathcal{S} = \{ A \in \mathbb{S} : \ Ax = b \}$ of all structured mappings taking x to b.

One of our motivations for studying these problems stems from the analysis of structured backward errors in the solutions to structured linear systems and structured eigenproblems. Together with D. S. Mackey and F. Tisseur, a complete and unified solution for these two problems was presented in [11] when $\mathbb{S}$ is the subspace of self-adjoint or skew-adjoint matrices associated with an orthosymmetric scalar product. These $\mathbb{S}$ include, for example, symmetric and skew-symmetric, Hermitian, pseudo-Hermitian and skew-Hermitian, Hamiltonian, persymmetric and perskew-symmetric matrices. Our approach gives easy-to-check conditions for the existence problem and an explicit solution for the characterization problem. We also identify structured mappings with various minimality properties. These results have already found application in the analysis of structured condition numbers and backward errors [7], [14], and constitute the first step towards characterizing the set $\{ A \in \mathbb{S} : \ AX = B \}$, where X and B are given matrices.

We are currently working on extending our results to the case when $\mathbb{S}$ is the automorphism (or isometry) group associated with the scalar product. Examples of such $\mathbb{S}$ are the unitary groups and the symplectic groups. Since we are now dealing with a nonlinear structure in $\mathbb{K}^{n\times n}$, different techniques are needed to solve the corresponding structured mapping problems.

Computing the Symplectic SVD: Since symplectic matrices have an SVD in which all three factors are also symplectic, it is natural to look for a structure-preserving algorithm to compute such an SVD. Using symplectic orthogonal transformations first described in [12], which allow us to target every off-diagonal entry, a one-sided Jacobi algorithm was developed together with A. Pacheenburawana [2], [13]. On well-conditioned matrices, the algorithm has good convergence behavior, even though the off-diagonal norm does not decrease monotonically. We are currently investigating the convergence properties of this algorithm more closely.

References

[1] E. N. Antoniou and S. Vologiannidis. A new family of companion forms of polynomial matrices. *Electr. J. Lin. Alg.*, 11:78–87, 2004.

[2] Archara Chaiyakarn. Structure Preserving Algorithms for Computing the Symplectic Singular Value Decomposition. PhD Thesis, Western Michigan University, 2005.

[3] Nicholas J. Higham, D. Steven Mackey, Niloufer Mackey, and Françoise Tisseur. Symmetric linearizations for matrix polynomials. *SIAM J. Matrix Anal. Appl.*, 29:143–159, 2006.

[4] Nicholas J. Higham, D. Steven Mackey, and Françoise Tisseur. The conditioning of linearizations of matrix polynomials. *SIAM J. Matrix Anal. Appl.*, 28:1005–1028, 2006.

[5] Elias Jarlebring. A quadratic eigenproblem in the analysis of a time delay system. In *Proc. Applied Math. and Mechanics*, 2006. Special Issue: GAMM Annual Meeting 2006 – Berlin. Preprint available from "http://www.tu-bs.de/ jarlebri".

[6] Elias Jarlebring. On critical delays for linear neutral delay systems. Preprint for European Control Conference 2007, available from "http://www.tu-bs.de/ jarlebri", 2006.

[7] Michael Karow and Daniel Kressner and Françoise Tisseur. Structured eigenvalue condition numbers. *SIAM J. Matrix Anal. Appl.*, 28:1052–1068, 2006.

[8] D. Steven Mackey, Niloufer Mackey, Christian Mehl, and Volker Mehrmann. Vector spaces of linearizations for matrix polynomials. *SIAM J. Matrix Anal. Appl.*, 28:971–1004, 2006.

[9] D. Steven Mackey, Niloufer Mackey, Christian Mehl, and Volker Mehrmann. Structured polynomial eigenvalue problems: Good vibrations from good linearizations. *SIAM J. Matrix Anal. Appl.*, 28:1029–1051, 2006.

[10] D. Steven Mackey, Niloufer Mackey, Christian Mehl, and Volker Mehrmann. Numerical methods for palindromic eigenvalue problems: Computing the anti-triangular Schur form. Submitted.

[11] D. Steven Mackey, Niloufer Mackey, and Françoise Tisseur. Structured mapping problems for matrices associated with scalar products, Part I: Lie and Jordan algebras. To appear in *SIAM J. Matrix Anal. Appl.*

[12] D. Steven Mackey, Niloufer Mackey and Françoise Tisseur. Structured tools for structured matrices. *Electr. J. Lin. Alg.*, 10:106–145, 2003.

[13] Niloufer Mackey and Archara Pacheenburawana. A one-sided Jacobi method for computing the symplectic SVD. In preparation.

[14] Françoise Tisseur and Stef Graillat. Structured condition numbers and backward errors in scalar product spaces. *Electr. J. Lin. Alg.*, 15:159–177, 2006.

Linearizations of Singular Matrix Polynomials and the Recovery of Minimal Indices

D. Steven Mackey
Joint work with F. De Terán and F. Dopico

A linearization of an $n \times n$ matrix polynomial $P(\lambda) = \sum_{j=0}^{k} \lambda^j A_j$ of degree k is a $kn \times kn$ pencil $L(\lambda) = \lambda X + Y$ such that $E(\lambda)L(\lambda)F(\lambda) = \operatorname{diag}\big(P(\lambda), I_{k(n-1)}\big)$ for some $kn \times kn$ matrix polynomials $E(\lambda)$ and $F(\lambda)$ that are unimodular, i.e. $\det E(\lambda)$ and $\det F(\lambda)$ are nonzero constants independent of λ. The use of such pencils $L(\lambda)$ is now well-established as a theoretical and computational tool for studying the properties of matrix polynomials [5], [10], [11].

For any polynomial $P(\lambda)$ there are infinitely many linearizations, but in practice the two most commonly used linearizations are the first and second (Frobenius) companion pencils $\mathcal{C}_1(\lambda) = \lambda X + Y_1$ and $\mathcal{C}_2(\lambda) = \lambda X + Y_2$, where $X = \operatorname{diag}(A_k, I_{k(n-1)})$ and

$$Y_1 = \begin{bmatrix} A_{k-1} & A_{k-2} & \cdots & A_0 \\ -I_n & 0 & \cdots & 0 \\ \vdots & \ddots & \ddots & \vdots \\ 0 & \cdots & -I_n & 0 \end{bmatrix}, \quad Y_2 = \begin{bmatrix} A_{k-1} & -I_n & \cdots & 0 \\ A_{k-2} & 0 & \ddots & \vdots \\ \vdots & \vdots & \ddots & -I_n \\ A_0 & 0 & \cdots & 0 \end{bmatrix}.$$

Recent papers [1],[11] have generalized these companion pencils in several different directions, providing large sets of easily available alternatives to the Frobenius companion pencils; having such alternatives allows the possibility of structure-preserving linearizations [1], [7], [12] as well as improved conditioning and backward error properties [6], [8], [9].

Almost all of the analyses in this recent work have been restricted to the case where P is regular, i.e. when $\det P(\lambda) \neq 0$. By contrast this talk will focus on $n \times n$ *singular* polynomials $P(\lambda)$, with $\det P(\lambda) \equiv 0$. The first issue to be discussed is the existence of large classes of readily available linearizations for singular P that are easy to construct from the data in P. It is well known that the Frobenius companion pencils $\mathcal{C}_1(\lambda)$ and $\mathcal{C}_2(\lambda)$ are linearizations for P even when P is singular. But which, if any, of the generalized companion pencils introduced in [1] and [11] are also linearizations when P is singular? In the first part of the talk we will identify many that are still linearizations and many that are not.

One of the key structural features of a singular pencil are its *minimal indices*, which encode information about the "singular part" of the Kronecker canonical form (KCF) of a

pencil. In [4], Gantmacher defines the right (or column) minimal indices of an $n \times n$ singular pencil $L(\lambda) = \lambda X + Y$ by viewing L as a linear operator on $\mathbb{F}^n$, where $\mathbb{F} = \mathbb{C}(\lambda)$ is the field of rational functions in the variable λ, and considering the (necessarily nontrivial) right nullspace $N \subseteq \mathbb{F}^n$ of $L(\lambda)$. Among all polynomial bases for N are certain distinguished "G-minimal" bases constructed in the following "greedy" manner. First choose any nonzero vector polynomial $\mathbf{v}_1(\lambda) \in N$ with minimal degree in N. Then among all vector polynomials in N that are linearly independent of $\mathbf{v}_1(\lambda)$, pick any $\mathbf{v}_2(\lambda)$ of minimal degree. Next choose any $\mathbf{v}_3(\lambda)$ of minimal degree among all vector polynomials in N that are linearly independent of $\{\mathbf{v}_1(\lambda), \mathbf{v}_2(\lambda)\}$. Continue extending in a similar manner to a "G-minimal" polynomial basis $\{\mathbf{v}_1(\lambda), \mathbf{v}_2(\lambda), \ldots, \mathbf{v}_p(\lambda)\}$ for N. Letting $\varepsilon_i = \deg \mathbf{v}_i(\lambda)$, Gantmacher shows that the degree sequence $\varepsilon_1 \leq \varepsilon_2 \leq \cdots \leq \varepsilon_p$ of such a G-minimal basis is independent of all the choices made in this construction, and so constitutes an intrinsic property of the pencil $L(\lambda)$, its right minimal indices. A similar construction of a G-minimal polynomial basis for the left nullspace in $\mathbb{F}^n$ of the operator $L(\lambda)$ generates another intrinsic degree sequence $\eta_1 \leq \eta_2 \leq \cdots \leq \eta_p$, the left (or row) minimal indices of $L(\lambda)$. Gantmacher then shows that these left and right minimal indices are the same as the sizes of the "singular blocks" in the KCF of $L(\lambda)$, thus proving that these sizes are uniquely determined. For the purposes of this talk, the key point about this construction is that it is not really about pencils or the KCF at all, it's about polynomial bases of *any* subspace of $\mathbb{F}^n$. Thus we can define left and right minimal indices of *any* singular polynomial $P(\lambda)$ as the G-minimal indices of the left and right nullspaces of $P(\lambda)$, respectively.

An alternative approach to defining minimal indices can be found in [3], using a different minimality property to identify distinguished polynomial bases in subspaces of $\mathbb{F}^n$. Given a subspace $W \subseteq \mathbb{F}^n$, a polynomial basis $\{\mathbf{w}_1(\lambda), \mathbf{w}_2(\lambda), \ldots, \mathbf{w}_\ell(\lambda)\}$ for W will be called "F-minimal" if the total degree $\sum_{j=1}^{\ell} \deg \mathbf{w}_j(\lambda)$ of the basis is minimal among all polynomial bases for W. Forney shows in [3] that every F-minimal basis for W has the same degree sequence $\alpha_1 \leq \alpha_2 \leq \cdots \leq \alpha_\ell$, thus uniquely determining a set of F-minimal indices $\{\alpha_i\}$ for W. (Note that in [3] these α_i are termed the "invariant dynamical indices" of W.) Are the F-minimal indices and the G-minimal indices of W related? In [2, Lemma 2.4] it is shown that these two apparently different sets of indices are always identical. Thus both approaches lead to the same values for the left and right minimal indices of a singular polynomial, and agree with the classical notion for singular pencils derived from the Kronecker canonical form.

The second issue to be addressed in this talk is the relationship between the minimal indices of a singular polynomial P and the minimal indices of any of its linearizations L. Can the minimal indices of P be recovered from the minimal indices of L in a systematic and uniform way? We consider this question for all the pencils identified as linearizations in the first half of the talk, and show how the answer depends on the particular linearization chosen.

References

[1] E. N. Antoniou and S. Vologiannidis. A new family of companion forms of polynomial matrices. *Electr. Journ. Lin. Alg.*, 11:78–87, 2004.

[2] F. De Terán and F.M. Dopico. Low rank perturbation of Kronecker structures without full rank. *SIAM Matrix Anal. Appl.*, 29(2):496–529, 2007.

[3] G.D. Forney, Jr. Minimal bases of rational vector spaces, with applications to multivariable linear systems. *SIAM J. Control*, 13(3):493–520, 1975.

[4] F. R. Gantmacher. *Theory of Matrices, Vol. II*, Chelsea, New York, 1959.

[5] I. Gohberg, P. Lancaster, and L. Rodman. *Matrix Polynomials*. Academic Press, 1982.

[6] N. J. Higham, R.-C. Li, and F. Tisseur. Backward error of polynomial eigenproblems solved by linearization. MIMS EPrint 2006.137, Manchester Institute for Mathematical Sciences, The University of Manchester, UK, June 2006. To appear in *SIAM J. Matrix Anal. Appl.*

[7] N. J. Higham, D. S. Mackey, N. Mackey, and F. Tisseur. Symmetric linearizations for matrix polynomials. *SIAM Matrix Anal. Appl.*, 29(1):143–159, 2006.

[8] N. J. Higham, D. S. Mackey, and F. Tisseur. The conditioning of linearizations of matrix polynomials. *SIAM Matrix Anal. Appl.*, 28(4):1005–1028, 2006.

[9] N.J. Higham, D. S. Mackey, F. Tisseur and S.D. Garvey. Scaling, sensitivity and stability in the numerical solution of quadratic eigenvalue problems. MIMS EPrint 2006.406, Manchester Institute for Mathematical Sciences, The University of Manchester, UK, November 2006. To appear in *Int. J. of Numerical Methods in Engineering*.

[10] P. Lancaster. *Lambda-matrices and Vibrating Systems*. Pergamon Press, 1966.

[11] D. S. Mackey, N. Mackey, C. Mehl, and V. Mehrmann. Vector spaces of linearizations for matrix polynomials. *SIAM Matrix Anal. Appl.*, 28(4):971–1004, 2006.

[12] D.S. Mackey, N. Mackey, C. Mehl, and V. Mehrmann. Structured polynomial eigenvalue problems: Good vibrations from good linearizations. *SIAM Matrix Anal. Appl.*, 28(4):1029–1051, 2006.

Statistical Leverage and Improved Matrix Algorithms

Michael W. Mahoney

Given an $m \times n$ matrix A and a rank parameter k, define the leverage of the i-th row of A to be the i-th diagonal element of the projection matrix onto the span of the top k left singular vectors of A. Historically, this statistical concept (and generalizations of it) has found extensive applications in, e.g, diagnostic regression analysis. Very recently, this concept has been central in the development of improved algorithms for several fundamental matrix problems. Two examples of this will be described. The first problem is the least squares approximation problem, in which there are n constraints and d variables. Classical algorithms, dating back to Gauss and Legendre, use $O(nd^2)$ time. We describe a randomized algorithm that uses only $O(nd \log d)$ time to compute a relative-error, i.e., $1 \pm \epsilon$, approximation. After an initial pre-processing with a randomized Hadamard transform, the algorithm uniformly randomly samples $r = O(d \log d \log n/\epsilon)$ constraints, and the least squares approximation problem on r constraints and d variables is solved exactly using a classical algorithm. An

important part of the analysis is to show that the pre-processing by a randomized Hadamard transform approximately uniformizes the leverage of each constraint, in which case uniform sampling is nearly optimal. The second problem is the problem of selecting a "good" set of exactly k columns from an $m \times n$ matrix, and the algorithm of Gu and Eisenstat provides the best previously existing result. We describe a two-stage algorithm that improves on their result. In particular, it provides the strongest known accuracy bounds (assuming that k is small) for the spectral norm of the error $A - QQ^T A$, where Q is an orthogonal basis spanning the k chosen columns. The algorithm is a two-phase hybrid algorithm. In the first phase, $O(k \log k)$ columns of the matrix A are randomly selected using a non-uniform probability distribution that depends on leverage information. Then, in the second phase, exactly k of these $O(k \log k)$ columns are deterministically selected using the Gu-Eisenstat procedure. Recent application of these ideas in the context of modern statistical data analysis will be briefly described.

A Tensor SVD

Carla D. Martin, Misha E. Kilmer, Lisa Perrone

The Singular Value Decomposition (SVD) of a matrix gives us important information about a matrix such as its rank, an orthonormal basis for the column or row space, and reduction to diagonal form. In applications, especially those involving multiway data analysis, information about the rank and reduction of tensors to have fewer nonzero entries are useful concepts to try to extend to higher dimensions. However, many of the powerful tools of linear algebra such as the SVD do not, unfortunately, extend in a straight-forward way to tensors of order three or higher. As such, there are multiple ways to extend the matrix SVD to higher-dimensions. The two most widely used representations are the CANDECOMP-PARAFAC model [3], and the TUCKER3 model [5] also referred to as the *Higher-Order Singular Value Decomposition* [2].

Traditionally, extending the SVD to higher-order tensors involves a representation using the outer product of vectors. These outer products for tensors can be written in terms of the n-mode product [4], which can also be used to describe multiplication of two tensors. In paricular, the product of an $\ell \times n_2 \times n_3$ tensor and an $\ell \times m_2 \times m_3$ tensor in the first mode is an $n_2 \times n_3 \times m_2 \times m_3$ tensor. Unfortunately, since one of the dimensions collapses under this definition of tensor-tensor product, it implies that the set of third-order tensors is not closed under this type of operation. Furthermore, this type of operation does not allow us to specify a notion of inverse, since the collapsing operation is not invertible.

In this talk, we present another generalization of the SVD where the terms in the decomposition are all tensors. In order to define our *tensor SVD*, we define a notion of tensor-tensor multiplication that is closed under multiplication, and address the issue of invertibility. A major motivation for considering this new type of tensor multiplication was to attempt to devise rank-revealing types of factorizations for tensors, reminiscent of matrix factorizations, which we could then use for applications such as tensor compression.

Similar to matrices, we formulate a definition such that the product of two $n \times n \times n$ tensors yields an $n \times n \times n$ tensor. We can show the operation is associative, identify the notion of an identity operator and give properties for an inverse. It follows that the set of all such invertible $n \times n \times n$ tensors under this operation forms a group. We expand our

list of definitions to include the concepts of transpose and orthogonality. We show that it is possible to use all these definitions to derive tensor decompositions that have rank-revealing characteristics in the sense that they reveal when it is possible to approximate an $n \times n \times n$ tensor by another tensor of rank $\leq n$.

References

[1] J.D. Carroll and J. Chang, Analysis of individual differences in multidimensional scaling via an N-way generalization of "Eckart-Young" decomposition. *Psychometrika*, 35 (1970) 283-319.

[2] L. De Lathauwer, B. De Moor, and J. Vandewalle, A multilinear singular value decomposition. *SIAM Journal of Matrix Analysis and Applications*, 21 (2000) 1253-1278.

[3] R.A. Harshman, Foundations of the PARAFAC procedure: Model and conditions for an 'explanatory' multi-mode factor analysis. *UCLA Working Papers in phonetics*, 16 (1970) 1-84.

[4] T.G. Kolda and B. Bader, Tensor Decompositions and Applications, Submitted August 2007.

[5] L.R. Tucker, Some mathematical notes on three-mode factor analysis. *Psychometrika*, 31 (1966) 279-311.

On Things of Little Interest: How to Compute Low Paying Annuities

Nicola Mastronardi, Marc Van Barel, Raf Vandebril

The problem of computing the nullspace of (block–)Toeplitz/Hankel matrices arises in a variety of applications, such as system theory, linear prediction and parameter estimation.

In this talk a numerical method, based on the generalized Schur algorithm, to compute the nullspace of the aforementioned matrices will be presented, and its stability properties will be discussed.

We will show that the proposed method can be used to solve the well known *"realization problem"* in system identification. More precisely, given a sequence of $p \times m$ matrices $h_k,\ k > 0$, the *realization problem* consists of finding a positive number n and constant matrices (A, B, C) such that

$$h_k = CA^{k-1}B, \quad C \in R^{p\times n}, A \in R^{n\times n}, B \in R^{n\times m},\ k = 1, 2, \ldots$$

The triple (A, B, C) is called a *realization* of the sequence $\{h_k\}$ and the latter is called a realization sequence. (A, B, C) is a *minimal realization* if, among all realizations of the sequence, its dimension is the smallest possible.

The proposed algorithm computes the triple (A, B, C) from the nullspace of the block–Hankel matrix

$$H_{j,k} = \begin{bmatrix} h_1 & h_2 & h_3 & \cdots & h_k \\ h_2 & h_3 & \cdots & \cdots & h_{k+1} \\ b_3 & \cdots & \cdots & \cdots & \vdots \\ \vdots & \cdots & \cdots & \cdots & h_{j+k-2} \\ h_j & h_{j+1} & \cdots & h_{j+k-2} & h_{j+k-1} \end{bmatrix},$$

with $k, j \gg n$.

The rank of the matrix $H_{j,k}$ provides the smallest possible dimension of the realization and is equal to the number of non–zero entries of the first computed vector of the nullspace of $H_{j,k}$. Moreover, the matrix A is computed as a companion matrix whose entries are retrieved from the latter vector.

Generalising Results about the Hermitian Eigenvalue Problems to the Quadratic Eigenvalue Problem

Roy Mathias

Let H_n denote the space of $n \times n$ Hermitian matrices. For $X \in H_n$ and $i = 1, \ldots, n$, let $\lambda_i(X)$ denote the ith largest eigenvalue of X. The eigenvalues of Hermitian matrices enjoy 3 remarkable properties

1. A min-max represeantion.
2. Monotonicity. That is,
$$A \geq B \ \Rightarrow \lambda_i(A) \geq \lambda_i(B).$$
3. Interlacing (under rank one perturbations and extraction of submatrices).

These properties allow one to very easily prove perturbation results, and to develop algorithms. Unitary, entry-wise nonnegative, and totally positive matrices have some, though not all, of these properties.

In this talk I will consider generalisations of these properties to the quadratic eigenvalue problem

$$\det(t^2 A + tB + C) = 0$$

where A, B, C are Hermitian.

Note that these 3 properties are closely related. It is natural to ask whether one of the 3 properties should be viewed as the fundamental one, and if so, which one.

Recycling Ritz vector for the solution of parameterized linear systems

Karl Meerbergen, Zhaojun Bai

The solution of the parameterized linear system

$$(K - \omega^2 M)x = f$$

where K and M are large sparse real symmetric matrices, and M is positive definite, arises in many engineering applications. The parameter ω is the frequency or the wave number. The solution is wanted for a large number of ω's selected from the frequency range $[\omega_{\min}, \omega_{\max}]$.

This problem is studied for several years in the context of model reduction and shifted linear system solvers, see e.g. [1] [2] [8] [9] [4] [5]. For this application, the Lanczos method preconditioned by $(K - \sigma M)^{-1}$ is very suitable. The method is intimately connected to the spectral transformation Lanczos method.

Recently, Ko and Bai [3] discussed the improvements of the AMLS frequency sweeping method by recycling eigenvectors.

The improvements of iterative methods by recycling Ritz vectors is understood [7] and is particularly useful for solving linear systems with multiple right-hand sides. We apply similar ideas to parameterized problems, i.e. to recycle Ritz vector chosen so that we can compute x for a large number of ω's in significantly smaller computation time than the parameterized Lanczos method. This is different from [6] since we use the fact that the matrix has the special form $K - \omega^2 M$. We show convergence theory.

We illustrate the method and the theory by a numerical example with multiple right-hand sides, arising from the 3D discretization of a layered windscreen, where the Ritz vectors obtained for the solution of one right-hand side are used for the solution of the remaining right-hand sides. For this example, recycling these Ritz pairs significantly improves the convergence of the parameterized Lanczos method for the remaining right-hand sides. The reason is that all eigenvalues in $[\omega_{\min}, \omega_{\max}]$ are computed with the first right-hand side and particularly these eigenvalues delay convergence.

References

[1] Z. Bai and R. Freund. A partial Padé-via-Lanczos method for reduced-order modeling. *Linear Alg. Appl.*, 332–334:141–166, 2001.

[2] A. Frommer and U. Glässner. Restarted GMRES for shifted linear systems. *SIAM J. Sci. Comput.*, 19:15–26, 1998.

[3] Jin-Hwan Ko and Zhaojun Bai. High-frequency response analysis via algebraic substructuring. Technical Report CSE-2007-18, UC Davis, CA. USA, 2007.

[4] K. Meerbergen. The solution of parameterized symmetric linear systems. *SIAM J. Matrix Anal. Appl.*, 24(4):1038–1059, 2003.

[5] K. Meerbergen. Fast frequency response computation for Rayleigh damping. Technical Report TW480, K.U.Leuven, Department of Computer Science, 2006. Accepted for publication in Intern. J. Numerical Methods in Engineering.

[6] M.L. Parks, E. De Sturler, G. Mackey, D. Johnson, and S. Maiti. Recycling Krylov subspaces for sequences of linear systems. Technical Report UIUCDS-R-2004-2421 (CS), UILU-ENG-2004-1722 (ENGR), University of Illinois at Urbana-Champaign, 2004.

[7] Y. Saad, M. Yeung, J. Erhel, and F. Guyomarc'h. A deflated version of the conjugate gradient algorithm. 21(5):1909–1926, 2000.

[8] V. Simoncini and E. Gallopoulos. Transfer functions and resolvent norm approximation of large matrices. *ETNA*, 7:190–201, 1998.

[9] V. Simoncini and F. Perotti. On the numerical solution of $(\lambda^2 A + \lambda B + C)x = b$ and application to structural dynamics. *SIAM J. Sci. Comput.*, 23(6):1876–1898, 2002.

Hamiltonian Jacobi Methods: Sweeping Away Convergence Problems

Christian Mehl

The Hamiltonian eigenvalue problem has been intensively studied in the last 25 years due to its numerous applications like, e.g., in linear-quadratic optimal control and H_∞-control, and due to its relation to generalized eigenvalue problems with Hermitian, symmetric, or skew-symmetric coefficient matrices. Many algorithms have been proposed for the solution of the Hamiltonian eigenvalue problem, but the generalization of one of the most successful algorithms in computational mathematics (the QR algorithm) encounters severe efficiency problems due to the lack of a "Hamiltonian version" of the preliminary reduction to Hessenberg form - a phenomenon that has usually been referred to as "Van Loan's curse" in the Numerical Linear Algebra community.

On the other hand, Jacobi's method for the diagonalization of a symmetric matrix is a famous, successful, and easy to implement algorithm for the computation of eigenvalues of symmetric matrices. No wonder that this algorithm has been generalized or adapted to many other classes of eigenvalue problems. In particular, Hamiltonian Jacobi methods for the Hamiltonian eigenvalue problem have been investigated by Byers (generalizing a nonsymmetric Jacobi method proposed by Stewart) and Bunse-Gerstner/Faßbender (generalizing a nonsymmetric Jacobi method proposed by Greenstadt and Eberlein).

Surprisingly, the Hamiltonian Jacobi method proposed by Bunse-Gerstner/Faßbender only shows linear asymptotic convergence despite the fact that asymptotic quadratic convergence can be observed by the Greenstadt/Eberlein Jacobi method. In this talk, we give an explanation for this unexpected behavior through a detailed investigation of sweep patterns. We show that for nonsymmetric Jacobi algorithms, sweep patterns that lead to convergence of the algorithm need not automatically imply quadratic asymptotic convergence - a phenomenon emphasizing a fundamental difference in the behavior of symmetric and nonsymmetric Jacobi methods. Then we identify sweep patterns that guarantee asymptotic quadratic convergence for nonsymmetric Jacobi methods and show how the sweep pattern of the algorithm by Bunse-Gerstner/Faßbender has to be modified in order to ensure the desired convergence behavior. Finally, we propose a slightly different Hamiltonian Jacobi method, investigate its convergence behavior, and discuss it advantages and disadvantages over the method proposed by Bunse-Gerstner/Faßbender.

Numerical computation of Sacker-Sell exponents for differential-algebraic equations

Vu Hoang Linh, Volker Mehrmann

Lyapunov and exponential dichotomy spectral theory is a very important topic in the stability analysis of dynamical systems. This theory extends the classical concepts of eigenvalues to linear systems with variable coefficients and nonlinear systems. In this talk the classical theory is generalized from ordinary differential equations (ODEs) to nonautonomous differential-algebraic equations (DAEs). By using orthogonal changes of variables, the original DAE system is transformed into an appropriate condensed form, for which concepts such as Lyapunov exponents, Bohl exponents, exponential dichotomy and spectral intervals of various kinds can be analyzed via the resulting underlying ODE. Some essential differences between the spectral theory for ODEs and that for DAEs are pointed out.

After presenting the theory, numerical methods for computing the spectral intervals associated with Lyapunov, Bohl and Sacker-Sell (exponential dichotomy) spectra are derived by modifying and extending those methods proposed for ODEs by Dieci and van Vleck. The key issue in the computation of these spectra is the computation of generalized Schur forms for pairs of matrix valued functions of one variable. We present numerical methods to compute these generalized Schur forms and also discuss perturbation theory and error analysis.

Finally, some numerical examples are presented to illustrate the theoretical results and the properties of the numerical methods.

Keywords: differential-algebraic equations, strangeness index, Lyapunov exponent, Bohl exponent, Sacker-Sell spectrum, exponential dichotomy, spectral interval, smooth QR factorization, continuous QR algorithm, discrete QR algorithm, kinematic equivalence, Steklov function

AMS(MOS) subject classification: 65L07, 65L80, 34D08, 34D09

References

[1] P. Bohl. Über Differentialungleichungen. *J. f. d. Reine und Angew. Math.*, 144:284–313, 1913.

[2] W. A. Coppel. *Dichotomies in Stability Theory.* Springer-Verlag, New York, NY, 1978.

[3] J.L. Daleckii and M.G. Krein. *Stability of solutions of differential equations in Banach spaces.* American Mathematical Society, Providence, RI, 1974.

[4] L. Dieci and T. Eirola. On smooth decompositions of matrices. *SIAM J. Matr. Anal. Appl.*, 20:800–819, 1999.

[5] L. Dieci and E. S. Van Vleck. Lyapunov and other spectra: a survey. In *Collected lectures on the preservation of stability under discretization (Fort Collins, CO, 2001)*, pages 197–218. SIAM, Philadelphia, PA, 2002.

[6] L. Dieci and E. S. Van Vleck. Lyapunov spectral intervals: theory and computation. *SIAM J. Numer. Anal.*, 40:516–542, 2002.

[7] L. Dieci and E. S. Van Vleck. Orthonormal integrators based on Householder and Givens transformations. *Future Generation Computer Systems*, 19:363–373, 2003.

[8] L. Dieci and E. S. Van Vleck. Perturbation theory for approximation of Lyapunov exponents by QR methods. *J. Dyn. Diff.Eq.*, 18:815–842, 2006.

[9] P. Kunkel and V. Mehrmann. *Differential-Algebraic Equations. Analysis and Numerical Solution*. EMS Publishing House, Zürich, Switzerland, 2006.

[10] P. Kunkel and V. Mehrmann. Stability properties of differential-algebraic equations and spin-stabilized discretization. *Electr. Trans. Num. Anal.*, (to appear), 2007.

[11] A. M. Lyapunov. The general problem of the stability of motion. Translated by A. T. Fuller from Edouard Davaux's French translation (1907) of the 1892 Russian original. *Internat. J. Control*, pages 521–790, 1992.

[12] Vu Hang Linh and V. Mehrmann. Spectral Intervals for Differential Algebraic Equations and their Numerical Approximations, Technical report 402, DFG Research Center MATHEON, *Mathematics for key technologies* in Berlin, TU Berlin. Str. des 17. Juni 136, D-10623 Berlin, Germany

Numerical Solution of Large-Scale Differential Riccati Equations

Hermann Mena, Peter Benner

The differential Riccati equations (DREs) arises in several applications, especially in control theory. Particularly, parabolic partial differential equations (PDEs) constrained optimization problems can be formulated as abstract Cauchy problems. Imposing a quadratic cost functional we obtain a linear quadratic regulator (LQR) problem for an infinite-dimensional system. The optimal control is then given in feedback form in terms of the solution of the operator differential Riccati equation on finite time horizons. We derive an approximation result enabling us to solve the problem numerically. Then, large-scale DREs resulting from the semi discretization in space of the corresponding PDEs have to be solved. Nonlinear control problems in model predictive control context require the solution of DREs with time-varying coefficients. Hence, we consider time-varying symmetric DREs of the form

$$\begin{aligned} \dot{X}(t) &= Q(t) + X(t)A(t) + A^T(t)X(t) - X(t)S(t)X(t), \\ X(t_0) &= X_0, \end{aligned}$$

where $t \in [t_0, t_f]$ and $Q(t)$, $A(t)$, $S(t)$ are piecewise continuous locally bounded $n \times n$ matrix-valued functions. Moreover, in most control problems, fast and slow modes are present. This implies that the associated DRE will be fairly stiff which in turn demands for implicit methods to solve such DREs numerically. Therefore, we will focus here on the stiff case. Typically the coefficient matrices of the resulting DRE there have a given structure (e.g. sparse, symmetric, low rank ...). We develop numerical methods capable of exploiting this structure.

Among the linear multistep methods and one step methods for solving ordinary differential equations (ODEs) that can deal with stiffness efficiently, the backward differentiation formulae

(BDF) and Rosenbrock type methods, respectively, are commonly used. In this research we develop efficient matrix valued versions of these ODE methods suitable for large-scale DREs. Low rank approximations and perturbations are the key ingredients to establish this task. The crucial question of suitable step size and order selection strategies is also addressed.

Solving the DRE using BDF methods requires the solution of an algebraic Riccati equation (ARE) in every step. The ARE is a nonlinear matrix equation so it is natural to apply the Newton's method. The application of the Newton's method yields a Lyapunov equation

$$F^T X + XF = -WW^T$$

to be solved in every step, where F is stable (i.e., its eigenvalues lie in the left half complex plane). Analogously, the application of an s-stage Rosenbrock method to the DRE requires the solution of one Lyapunov equation in each stage, as for the BDF methods, we solve the Lyapunov equation by the low rank version of the alternating direction implicit (ADI) algorithm. In both cases F can be represented as the sum of a sparse matrix and a low rank perturbation. This implies the usage of the Sherman-Morrison-Woodbury formula in the solver, to increase the efficiency.

The convergence of the ADI algorithm strongly depends on the set of shift parameters. The selection of the shift parameters leads to a rational min-max-problem. Since the knowledge about the shape of the complex spectrum is crucial for computing the optimal solution, this is infeasible for the large-scale problems. Therefore, we use an new procedure which has a good performance in practice. The computational cost depends on the computation of an Arnoldi process for the matrix involved and on the computation of elliptic integrals.

We discuss the numerical solution of optimal control problems for instationary heat, convection-diffusion and diffusion-reaction equations formulating them as abstract LQR problems. We also study optimal control problems with nonlinear PDEs. Tracking and stabilization type control problems are treated with model predictive control techniques where we need to solve linearized problems on small time frames.

¿From Wilkinson distance to nearest non-minimal dynamical systems : Lipschitz-based optimization of singular values

Emre Mengi

The distance of a matrix $A \in \mathbf{C}^{n\times n}$ to the closest one with a multiple eigenvalue has been introduced by Wilkinson as a measure of sensitivity of the eigenvalues to perturbations. This distance, commonly referred as the Wilkinson distance, has recently been reduced to the singular value optimization problem

$$\inf_{\lambda\in\mathbf{C}} \sup_{\gamma\in\mathbf{R}_+} \sigma_{2n-1}\left(\begin{bmatrix} A-\lambda I & \gamma I \\ 0 & A-\lambda I \end{bmatrix}\right). \tag{1}$$

by Malyshev. Above and in what follows σ_j denotes the jth largest singular value. Malyshev's derivation is based on the observation that A has λ as a multiple eigenvalue if and only if the matrix $(A-\lambda I)^2$ is of rank $n-2$ or smaller. More generally one can show that if $\lambda_1, \lambda_2, \ldots, \lambda_r$ are eigenvalues of A with the dimension of the associated eigenspace r or greater (that is if the sum of the geometric multiplicites of $\lambda_1, \lambda_2, \ldots, \lambda_r$ is at least r), then

$$\operatorname{rank}[(A-\lambda_1 I)(A-\lambda_2 I)\ldots(A-\lambda_r I)] \leq n-r. \tag{2}$$

Conversely, if the rank condition above holds, a subset of $\lambda_1, \lambda_2, \ldots, \lambda_r$ are eigenvalues of A with the sum of the geometric multiplicities r or greater.

A seemingly unrelated problem concerns the dynamical system

$$\dot{x}(t) = Ax(t) + Bu(t), \quad y(t) = Cx(t) + Du(t), \tag{3}$$

where $A \in \mathbf{C}^{n\times n}, B \in \mathbf{C}^{n\times k}, C \in \mathbf{C}^{p\times n}, D \in \mathbf{C}^{p\times k}$, $n \geq k$ and $n \geq p$. We consider the distance to the closest system with the dimension of space of controllable states equal to r or smaller,

$$\tau_r(A, B) = \inf\{\|[\Delta A \ \Delta B]\| : \operatorname{rank}(\mathcal{C}(A + \Delta A, B + \Delta B)) \leq r\}$$

with

$$\mathcal{C}(A, B) = [B \ \ AB \ \ A^2B \ \ \ldots \ \ A^{n-1}B].$$

Above and throughout this work $\|\cdot\|$ denotes the spectral norm. Analogously the quantity $\tau_r(A^*, C^*)$ can be defined as the distance to the closest system with the dimension of space of observable states equal to r or smaller. These quantities combined measure how reducible the system is to order r. We call the dynamical system (3) reducible to order $r < n$ if there exists a dynamical system of order r with the coefficient matrices $A_r \in \mathbf{C}^{r\times r}, B_r \in \mathbf{C}^{r\times m}, C_r \in \mathbf{C}^{p\times r}, D$ that maps the input space u to the output space y in the same manner as (3). (Formally the operators associated with these systems from the input space to the output space are equal.) The controllability and observability canonical forms reveal that when $\operatorname{rank}(\mathcal{C}(A, B)) \leq r$ or $\operatorname{rank}(\mathcal{C}(A^*, C^*)) \leq r$, the system is reducible to order r. Therefore if either of the distances $\tau_r(A, B), \tau_r(A^*, C^*)$ is small, there is a nearby system reducible to order r with the associated operator behaves like the operator of the system (3), and preserving the relation between the input and output space well. First we establish the equivalence of $\operatorname{rank}(\mathcal{C}(A, B)) \leq r$ and

$$\exists \lambda_1, \lambda_2, \ldots, \lambda_{n-r} \in \mathbf{C} \ \ \operatorname{rank}([(A - \lambda_1 I)(A - \lambda_2 I)\ldots(A - \lambda_{n-r} I) \ \ B]) \leq r. \tag{4}$$

Notice that when $r = n - 1$, the condition above reduces to the Popov-Belevitch-Hautus test $\exists \lambda \in \mathbf{C} \ \ \operatorname{rank}([A - \lambda I \ \ B]) \leq n - 1$ for uncontrollability. Then we deduce a lower bound on $\tau_r(A, B)$

$$\inf_{\lambda \in \mathbf{C}^{n-r}} \sup_{\gamma \in \mathbf{R}_+^{n-r+1}} \sigma_{(n-r)(n-1)+1} \left(\begin{bmatrix} A - \lambda_1 I & B & & & & \\ \gamma_1 I & 0 & A - \lambda_2 I & B & & \\ 0 & 0 & \gamma_2 I & 0 & & \\ & & & & \ddots & \\ & & & & & A - \lambda_{n-r} I \quad B \end{bmatrix} \right)$$

$$\leq \tau_r(A, B), \tag{5}$$

where $\lambda = [\lambda_1 \ \ldots \ \lambda_{n-r}]$ and $\gamma = [\gamma_1 \ \ldots \ \gamma_{n-r+1}]$. We conjecture that $\tau_r(A, B)$ is equal or within a constant factor of the singular value optimization problem above on the left. Specifically when $r = n - 1$, the quantity $\tau_r(A, B)$ becomes the distance to uncontrollability and the optimization problem reduces to well-known singular value characterization for the distance to uncontrollability $\inf_{\lambda \in \mathbf{C}} \sigma_n([A - \lambda I \ \ B])$ due to Eising.

Both of the minimization problems in (1) and (5) can be rearranged so that they are in the form

$$\inf_{\lambda \in \mathbf{R}^s} \sup_{\gamma \in \mathbf{R}_+^r} \sigma_j(P \otimes I_{l,l} + I_{n,m} \otimes Q(\lambda, \gamma))) \tag{6}$$

where $P \in \mathbf{C}^{n\times m}$ with $m \geq n$, $Q(\lambda,\gamma) \in \mathbf{C}^{l\times l}$, $\otimes$ denotes the Kronecker product and $I_{n,m}$ is the matrix obtained by concatenating $m-n$ zero columns to the identity matrix of size n. The entries of the matrix $Q(\lambda,\gamma) \in \mathbf{C}^{l\times l}$ satisfy the following two special properties;

1. Each entry is a sum $c_1 e^{i\theta} + c_2 e^{i(\theta-\pi/2)}$ for some θ where each of c_1 and c_2 is zero or otherwise equal to one of the entries of λ or γ. Note that $|c_1 e^{i\theta} + c_2 e^{i(\theta-\pi/2)}| = \|[c_1\ c_2]\|$.

2. Each entry of λ or γ appears at most once as c_1 or c_2 in the sum in 1. in a column and in a row of $Q(\lambda,\gamma)$.

Furthermore, we assume that the inner maximization problem is unimodal, that is when λ is fixed, σ_j as a function of γ has only one local extremum which happens to be the global maximum. In particular the inner maximization problem in (1) is unimodal.

For the solution of the problems of kind (6) we exploit the fact that singular values are insensitive to perturbations. In particular we show that the functions

$$f(\lambda,\gamma) = \sigma_j(P \otimes I_{l,l} + I_{n,m} \otimes Q(\lambda,\gamma))) \text{ and } g(\lambda) = \sup_{\gamma\in\mathbf{R}^r_+} f(\lambda,\gamma).$$

are Lipschitz with Lipschitz constant one, that is

$$|f(\lambda+\delta\lambda,\gamma+\delta\gamma) - f(\lambda,\gamma)| \leq \|[\delta\lambda\ \ \delta\gamma]\| \text{ and } |g(\lambda) - g(\lambda+\delta\lambda)| \leq \|\delta\lambda\|.$$

The Lipschitzness properties above are simple consequences of the Weyl's theorem (which states that the eigenvalues of a symmetric matrix are Lipschitz with respect to perturbations to the matrix) and the specific structure of $Q(\lambda,\gamma)$.

Numerically to solve the non-convex Lipschitz optimization problem in (6) we use the algorithm DIRECT introduced by Jones, Perttunen and Stuckmann and its freely available implementation by Gablonsky. DIRECT is the multi-variable generalization of an algorithm by Schubert and Piyavskii for Lipschitz functions of single variable. Schubert-Piyavskii algorithm constructs a piecewise linear approximation called a (lower) envelope for a Lipschitz function h with Lipschitz constant c (c is one for the singular value optimization problem (6)). The envelope bounds h from below. At each iteration the function is evaluated at the minimum of the envelope, say at $\tilde{x}$. Now the piecewise linear function $d(x) = h(\tilde{x}) - c|x-\tilde{x}|$ lies below h and the envelope for h is updated accordingly. The process of updating the envelope continues until the minimum of the envelope and the minimum of the evaluated function values differ by less than a tolerance.

For our specific problem the function we minimize by DIRECT is $g(\lambda)$ over s real variables. The inner maximization problem over γ, which yields us $g(\lambda)$ for a given λ, can be solved effectively by means of a technique such as BFGS or golden section search. Indeed we employ BFGS in practice because it converges faster and once the function value $f(\lambda,\gamma)$ is evaluated, which requires the calculation of a singular value, we can use the analytic formula for the derivatives in terms of the corresponding singular vectors. The costly operation involved in the solution of (6) in this manner is the function evaluations of $g(\lambda)$ and therefore $f(\lambda,\gamma)$ at various λ and γ. We describe how the singular value σ_j in (6) can be typically computed at each λ and γ in $O(ln^2)$ time (recall that n, l are the number of rows of the matrices $P, Q(\lambda,\gamma)$, respectively) once a Schur factorization of P is performed initially by using an iterative eigenvalue solver, for instance the shift-and-invert restarted Arnoldi implementation in ARPACK.

We have implemented the approach summarized above for the Wilkinson distance. In our limited experiments we have obtained the exact distance within a few digit precision efficiently. The sample running times in seconds on random matrices of various sizes with entries selected from the normal distribution with zero mean and variance equal to one are listed in the table below.

10	20	40	80	100	200	400	800
3	3	3	9	16	47	124	454

We should note that the running times are greatly affected by how rapidly the singular value function changes. In general for matrices with flat singular values the computations may take considerably longer compared to sharper functions with slopes close to one.

New results on the Lanczos and Conjugate Gradient algorithms

Gérard Meurant, Hassane Sadok

When using the Lanczos algorithm on a symmetric matrix A starting from a vector v^1, one builds a matrix V_k, $k = 1, \ldots$ whose columns are the Lanczos vectors v^k by a three-term recurrence. The Lanczos tridiagonal matrix T_k is $T_k = V_k^T A V_k$. Its eigenvalues are the Ritz values which are approximations of the eigenvalues of A. In this paper we study the matrix $\hat{T}_k = (V_k^T A^{-1} V_k)^{-1}$ and we prove that

$$\hat{T}_k = T_k + \tau_k e^k (e^k)^T,$$

where τ_k is a real number and e^k is the last column of the identity matrix of order k. Therefore, $\hat{T}_k$ is tridiagonal and equal to T_k except for the bottom right element. The eigenvalues of $\hat{T}_k$ that we call Petrov-Galerkin values are also approximations of the eigenvalues of A. They satisfy interlacing relations with the Ritz values. Moreover, they are equal to the harmonic Ritz values if the starting vector is chosen as $v^1 = Av$ for a given vector v. We will show some examples and comparisons of the Ritz values, the harmonic Ritz values and the Petrov-Galerkin values and will discuss how the latter ones can be computed. We will also give some characterization of τ_k.

It is well known that the Lanczos algorithm is tightly coupled with the Conjugate Gradient algorithm. If the CG error at iteration k is denoted by ε_k, we will prove that

$$\frac{\|\varepsilon_k\|_A^2}{\|\varepsilon_{k-1}\|_A^2} = \tau_k (e_k)^T T_k^{-1} e^k.$$

The values $(e_k)^T T_k^{-1} e^k$ are easily seen to be equal to $1/\delta_k$ where δ_k is the last diagonal value of the LDL^T factorization of T_k.

Unfortunately, the parameter τ_k is not directly computable. We will show how it can be approximated by introducing a delay as it was done to approximate the A-norm of the error, see [G. Meurant, *The Lanczos and Conjugate Gradient algorithms, from theory to finite precision computations*, SIAM, 2006].

Solution of Linear Equations for Quantum Chromodynamics

Ronald B. Morgan

1 Introduction

Solving linear equations is an essential part of most problems in quantum chromodynamics (QCD). Usually there are multiple right-hand sides for each matrix and sometimes there are also multiple shifts of the matrix. We will discuss several techniques for some of the various linear equations problems that arise in QCD and give new convergence theory. These techniques use approximate eigenvectors to deflate eigenvalues and improve convergence. As mentioned in the next section, these eigenvectors can be computed while simultaneously solving linear equations.

2 GMRES with Deflated Restarting (GMRES-DR)

GMRES-DR [1] is a deflated version of restarted GMRES that simultaneously solves a system of linear equations and computes eigenvalues. It uses the subspace

$$Span\{\tilde{y}_1, \tilde{y}_2, \ldots \tilde{y}_k, r_0, Ar_0, A^2r_0, A^3r_0, \ldots, A^{m-k-1}r_0\}, \tag{1}$$

where r_0 is the residual vector at the start of the new cycle and $\tilde{y}_1, \ldots \tilde{y}_k$ are harmonic Ritz vectors computed at the end of the previous cycle. This subspace has approximate eigenvectors added to a Krylov porion.

3 Multiple Right-hand Sides

For systems of equations with multiple right-hand sides, we use the approximate eigenvectors generated during solution of the first right-hand side with GMRES-DR to help solve the other right-hand sides. We give an approach called GMRES-Proj [2] that has cycles of GMRES alternated with minimum residual projections over the approximate eigenvectors. GMRES-Proj is efficient because the approximate eigenvectors come from GMRES-DR in the Arnoldi-like recurrence

$$AV_k = V_{k+1}\bar{H}_k, \tag{2}$$

where V_k is a n by k matrix whose columns span the subspace of approximate eigenvectors, V_{k+1} is the same except for an extra column and $\bar{H}_k$ is a full $k+1$ by k matrix. In experiments with large complex QCD matrices, the convergence for the second and subsequent right-hand sides is usually an order of magnitude or more faster for GMRES(20)-Proj(30) than for regular GMRES(20).

4 Multiple shifts

For shifted systems $(A - \sigma_i)x^i = b$, GMRES-DR can solve all shifted systems simultaneously. For multiple right-hand sides, all shifted systems need to be restarted with parallel residual vectors. However, the deviation from having parallel residuals can be pushed in the direction

of one vector, specifically the last column of V_{k+1} from (2). And the error from this deviation can later be corrected. The solution of one additional system with v_{k+1} as right-hand side is required. We now leave the topic of multiple shifts and return to just multiple right-hand sides.

5 Deflated BiCGStab for subsequent right-hand sides

For the second and subsequent right-hand sides, a deflated BiCGStab can be used instead of GMRES-Proj. However, the projection over approximate eigenvectors can only be done once, before the BiCGStab iteration begins. This is not generally effective, because the minimum residual projection does not reduce enough the components of the residual vectors in the directions of eigenvectors corresponding to small eigenvalues. One way to fix this is to restart BiCGStab and do more projections. Another approach is to apply a better projection using both right and *left* eigenvectors. For certain QCD problems, this is particularly applicable, because the left eigenvectors can be easily determined from the right eigenvectors (using the QCD $\gamma 5$ matrix). Tests with large QCD matrices show that this deflated BiCGStab is a significant improvement over regular BiCGStab and is sometimes better than GMRES-Proj.

6 Deflated Lanczos for Hermitian Problem

By using the $\gamma 5$ matrix or the normal equations, QCD systems can be turned into Hermitian problems (with very different spectra). This allows the conjugate gradient method to be used, but deflation is still essential. We are developing a Lanczos-DR method that simultaneously solves linear equations and computes eigenvalues and eigenvectors. The eigenvalue portion of this algorithm can be mathematically equivalent to implicitly restarted Arnoldi and to thick restarted Lanczos or it can use a minres/harmonic approach. To deal with roundoff error, we investigate a selective orthogonalization approach that reorthogonalizes against the approximate eigenvectors that are saved at the restart. This is more natural than the original Parlett-Scott approach, because the eigenvectors are already being computed. This reorthogonalization is a compromise between full reorthogonalization and Simon's partial reorthogonalization, giving good accuracy at a medium price. For multiple right-hand sides, the eigenvectors are used to deflate the small eigencomponents from the conjugate gradient method.

7 Parlett's trick

Parlett [3] has an efficient way of combining Krylov subspaces from two different right-hand sides. We use this with the first subspace being the Krylov subspace of approximate eigenvectors generated by Lanczos-DR. This is more practical than the original approach, because it does not require a large basis to be stored. The implementation must be carefully done because of possible roundoff error.

8 Deflated Restarted BiCG

Next, we consider QCD systems for which the left eigenvectors are not automatically available. For example, this is the case in twisted mass calculations. We are developing a deflated restarted Lanczos/BiCG algorithm that computes both right and left eigenvectors while it solves a system of linear equations. As mentioned earlier, both right and left eigenvectors are needed for an efficient BiCGStab method for multiple right-hand sides. This restarting may also be useful to control some of the instability of nonsymmetric Lanczos. We also plan to look at deflated restarted QMR.

9 Convergence results

Convergence results will be given for the methods mentioned above and for implicitly restarted Arnoldi and thick restarted Lanczos. The results use the fact that while the subspace combines an approximate eigenvector portion with a Krylov portion, the whole thing is a Krylov subspace (with a different starting vector). Also used is the fact that this subspace contains Krylov subspaces with each Ritz vector as starting vector.

References

[1] R. B. Morgan. GMRES with Deflated Restarting. *SIAM J. Sci. Comput.*, 24:20–37, 2002.

[2] R. B. Morgan and W. Wilcox. Deflated iterative methods for linear equations with multiple right-hand sides. 2004.

[3] B. N. Parlett. A new look at the Lanczos algorithm for solving symmetric systems of linear equations. *Linear Algebra Appl.*, 29:323–346, 1980.

Directional perturbation in structured eigenproblems

Julio Moro

In certain situations coming up in applications, the eigenvalues of a structured matrix or operator A must be perturbed within the structure in a rather specific way, choosing either the direction in which (possibly some of) the perturbed eigenvalues should move, or the optimal directions E in matrix space so that the eigenvalues of $A + E$ behave in some desired way. In several cases, the goal of the perturbation is to move the eigenvalues away from some dangerous region (a resonance band, for instance, or a boundary of stability). This talk will be focused on such kind of problems.

In these situations, standard eigenvalue perturbation theory is usually not enough, since the structure of A must be somehow incorporated into the perturbation analysis. In this talk we review some of the fundamental tools at our disposal in order to adapt the perturbation theory to some particular structures, illustrating the basic features of this perturbation analysis through two specific kinds of problems: frequency isolation problems in structural mechanics, and the passivation problem in control theory.

The content of this talk corresponds to joint work with Shreemajee Bora, Froilán Dopico, Juan Egaña, Michael Karow, Daniel Kressner, Volker Mehrmann, Maria Peláez and Fernando de Terán.

Multilevel Krylov Methods based on Projections derived from Deflation, Domain Decomposition, and Multigrid Iterations

Yogi A. Erlangga, Reinhard Nabben

For various applications, it is well-known that a two-level-preconditioned Krylov method is an efficient method for solving large and sparse linear systems. Beside a traditional preconditioner like incomplete Cholesky decomposition, a projector has been included as preconditioner to get rid of a number of small and large eigenvalues of the matrix. In literature, various projection methods are known coming from the fields of deflation, domain decomposition (BPS and BNN) and multigrid.

Here we give an algebraic formulation of projection methods. This formulation allows a comparison between these technics. By giving a detailed introduction into the deflation technic we establish analogies between the abstract balancing, the additive coarse grid correction method, used in domain decomposition methods, and the deflation technic. We establish spectral properties of these methods as well as comparisons of the related condition numbers. Moreover we prove that deflation technic generates iteration vectors whose A-norms are less than the A-norms of the iteration vectors generated by the abstract balancing preconditioner.

Surprisingly, we also obtain similar theoretical results for these methods applied to nonsymmetric systems.

These detailed analysis end up with the development of a new multilevel deflation method or a multilevel, projection based, Krylov method for arbitrary systems.

The projection of this method is constructed in a similar way as in deflation, but shifts small eigenvalues to the largest one instead of to zero. In contrast with deflation, however, the convergence rate of a Krylov method combined with this new projection method is insensitive to the inaccurate solve of the Galerkin matrix, which with some particular choice of deflation subspaces is closely related to the coarse-grid solve in multigrid or domain decomposition methods. Such an insensitivity allows the use of inner iterations to solve the Galerkin problem. An application of a Krylov subspace method to the associated Galerkin system with the projection preconditioner leads to a multilevel, nested Krylov iteration. In this multilevel projection Krylov subspace method, information of small eigenvalues to be projected is contained implicitly in the Galerkin system associated with the matrix of the linear system to be solved. These small eigenvalues, from a Krylov method point of view, are responsible for slow convergence. In terms of projection methods, this is conceptually similar to multigrid, but different in the sense that in multigrid the projection is done by the smoother. Furthermore, with the only condition that the deflation matrices are full rank, we have in principle more freedom in choosing the deflation subspace. Intergrid transfer operators used in multigrid are some of the possible candidates.

We present numerical results from solving the Poisson equation and the convection-diffusion equation, both in 2D. The latter represents the case where the related matrix of coefficients is nonsymmetric. By using a simple piece-wise constant interpolation as the basis for constructing the deflation subspace, we obtain the following results: (i) h-independent

convergence for the Poisson equation, and (ii) almost independent of h and the Péclet number for the convection-diffusion equation.

This new multilevel Krylov method works also very well for the indefinite linear systems arising from a discretization of the Helmholtz equation. Numerical results for high wavenumbers show the effectiveness of the method for solving Helmholtz problems. Not only can the convergence be made almost independent of grid size h, but also only mildly dependent of the wavenumber k. These results will be presented in detail in the talk by Yogi Erlangga.

This work is supported by the *Deutsche Forschungsgemeinschaft* (DFG), Project Number NA248/2-2. The talk is based on the following references:

References

[1] Y.A. Erlangga and R. Nabben, Deflation and balancing preconditioners for Krylov subspace methods applied to nonsymmetric matrices, *submitted* 2007,

[2] Y.A. Erlangga and R. Nabben, Multilevel projection based Krylov methods for boundary value problems, to appear in SIAM J. Sci. Comput.,

[3] Y.A. Erlangga and R. Nabben, On a multilevel Krylov method for the Helmholtz equation preconditioned by shifted Laplacian, *submitted* 2007

[4] R. Nabben and C. Vuik, A comparison of Deflation and Coarse Grid Correction applied to porous media flow, *SIAM J. Numer. Anal.*, 42:1631–1647, 2004.

[5] R. Nabben and C. Vuik, A Comparison of Deflation and the Balancing Preconditioner, *SIAM J. Sci. Comput.*, 27, 1742–1759, 2006.

[6] R. Nabben and C. Vuik, A comparison of abstract versions of deflation, balancing and additive coarse grid correction preconditioners , *submitted* 2006.

[7] J.M. Tang and R. Nabben and K. Vuik and Y.A. Erlangga Theoretical and Numerical Comparison of Projection Methods derived from Deflation, Domain Decomposition and Multigrid Methods, *submitted* 2007.

Separable Nonlinear Least Squares Problems in Image Processing

James G. Nagy

We consider the following general problem that arises in image processing applications:

$$\mathbf{b} = \mathbf{A}(\mathbf{y})\,\mathbf{x} + \mathbf{e} \tag{1}$$

where

- $\mathbf{b}$ is a vector representing known image data,

- $\mathbf{x}$ is a vector representing an unknown image we wish to reconstruct,
- $\mathbf{A}$ is an ill-conditioned matrix that depends on (often unknown) parameters, and these parameters are defined by a vector $\mathbf{y}$,
- $\mathbf{e}$ is a vector that represents unknown additive noise in the measured data.

We consider applications where equation (1) is obtained from discretization of an ill-posed inverse problem, and thus the matrix $\mathbf{A}(\mathbf{y})$ is severely ill-conditioned. This means that the noise term, $\mathbf{e}$, cannot be ignored, and that without some form of *regularization*, computed solutions will be dominated by high frequency noise. This type of problem arises in many imaging problems, including image deblurring, image resolution enhancement, and data fusion.

The general problem given in equation (1) requires computing approximations of $\mathbf{x}$ and $\mathbf{y}$. One approach to compute these approximations is to formulate a nonlinear least squares problem,

$$\min_{\mathbf{x},\mathbf{y}} \|\mathbf{A}(\mathbf{y})\,\mathbf{x} - \mathbf{b}\|_2^2 \ . \tag{2}$$

We assume that $\mathbf{x} \in \mathcal{R}^n$ and $\mathbf{y} \in \mathcal{R}^p$, and in our applications $p \ll n$. Equation (2) is called a *separable* nonlinear least squares problem; notice the separate dependence on linear parameters $\mathbf{x}$ and nonlinear parameters $\mathbf{y}$.

As previously mentioned, regularization must be incorporated into the problem to avoid magnifying noise in the computed reconstructions. One approach is to incorporate regularization into the optimization problem,

$$\min_{\mathbf{x},\mathbf{y}} \left\{ \|\mathbf{A}(\mathbf{y})\,\mathbf{x} - \mathbf{b}\|_2^2 + \lambda_x R(\mathbf{x}) + \lambda_y S(\mathbf{y}) \right\}$$

where R and S are regularization functionals and λ_x and λ_y are regularization parameters. Alternatively we can use a Levenberg-Marquardt type approach, where we develop an optimization scheme to solve the basic nonlinear least squares problem given by equation (2), and then regularize the linear system solves that arise during the iteration process. We have had recent success using a weighted generalized cross validation scheme within a Lanczos hybrid bidiagonalization-regularization method [2] for this purpose. (Obtaining estimates of good regularization parameters is one of the most difficult aspects of solving ill-posed inverse problems.)

Although we discuss regularization issues, the main focus of this presentation is on efficient computational approaches to compute estimates of $\mathbf{x}$ and $\mathbf{y}$. For example, it is possible to exploit structure of the problem; note that equation (2) is linear in $\mathbf{x}$, and (in the applications we consider) there are relatively few parameters in $\mathbf{y}$. This kind or structure can be exploited, for example, by using a variable projection scheme; see, e.g., [3, 4, 5, 6]. For large scale problems in image processing, it is also important to exploit structure of the mathematical model, and in particular, in the matrix $\mathbf{A}(\mathbf{y})$. This is necessary not only when solving (for a fixed $\mathbf{y}$) the linear least squares problem (2), but also when solving linear systems involving the Jacobian matrix. Examples of efficient algorithms such as this have been developed for certain image processing applications; see, for example [1, 7].

In this presentation we describe how the model problem (1) arises in image processing applications, and discuss the challenges to efficiently computing accurate approximations of $\mathbf{x}$ and $\mathbf{y}$. In addition, we use specific examples in image resolution enhancement and image deblurring to show how structure of the application can be exploited for efficient computations.

References

[1] J. M. Chung, E. Haber, and J. G. Nagy. Numerical methods for coupled super-resolution. *Inverse Problems*, 22:1261–1272, 2006.

[2] J. M. Chung, J. G. Nagy, and D. P. O'Leary. A weighted GCV method for Lanczos hybrid regularization. Technical Report TR-2007-004-A, Mathematics and Computer Science Department, Emory University, 2007. To appear in ETNA.

[3] G. Golub and V. Pereyra. The differentiation of pseudo-inverses and nonlinear least squares whose variables separate. *SIAM J. Numer. Anal.*, 10:413–432, 1973.

[4] G. Golub and V. Pereyra. Separable nonlinear least squares: the variable projection method and its applications. *Inverse Problems*, 19:R1–R26, 2003.

[5] L. Kaufman. A variable projection method for solving separable nonlinear least squares problems. *BIT*, 15:49–57, 1975.

[6] A. Ruhe and P. Wedin. Algorithms for separable nonlinear least squares problems. *SIAM Review*, 22:318–337, 1980.

[7] C. R. Vogel, T. Chan, and R. J. Plemmons. Fast algorithms for phase diversity-based blind deconvolution. In *Adaptive Optical System Technologies*, volume 3353. SPIE, 1998.

Exploiting the Finite Element framework to construct two-by-two block preconditioners

Maya Neytcheva, Owe Axelsson

In this talk we discuss possibilities to construct efficient two-by-two block factorized preconditioners for general matrices, arising from finite element discretizations of scalar or vector partial differential equations.

The work to be presented is based on two milestones. The first one is the elegancy of the finite element method (FEM) which is not only a powerful discretization tool but also offers the possibility to perform much analysis on a local, element level. There, since the corresponding matrices and vectors have small dimensions, even exact and symbolic computations are feasible to perform. Furthermore, adaptive finite element methods (AFEM) have undergone a substantial development in theory and implementation, offering the advantage to reduce the required degrees of freedom by 'zooming' only in those parts of the problem domain where it turns out to be necessary.

The second milestone is the ever increasing scale of linear systems to be solved within various application settings, entailing the use of iterative solution methods and, thus, the need for efficient preconditioners. Constructing good preconditioners has been a topic of intensive research and implementation over the last thirty years and has still not lost its importance. Despite the significant success in development of general algebraic preconditioning techniques, it still remains true that most efficient preconditioners are those, where in the construction we incorporate more information than what is contained in the system matrix only, i.e. if the preconditioners are more 'problem-dependent'. As an example we mention the so-called

AMGe methods where a powerful algebraic multigrid method is build having some knowledge about the element matrices or mesh topology behind the system matrix to be solved, see, for instance [5].

To fix some notations, we assume that the target problem to be solved is a linear system of algebraic equations $A\mathbf{x} = \mathbf{b}$ and A admits the form

$$A = \begin{bmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{bmatrix}.$$

We consider here a class of potentially very powerful algebraically constructed preconditioners, based on the above presumed block two-by-two structure of the system matrix and the related to that exact block-factorization of A. It is well known that the two crucial issues for two-by-two block-factorized preconditioners to be efficient are to construct good approximations of the main pivot block A_{11} and of the Schur complement matrix $S_A = A_{22} - A_{21}A_{11}^{-1}A_{12}$, or of their inverses. In general, these are not easy tasks because,for instance, both A_{11}^{-1} and S_A could be dense while we would like to have sparse approximations in order to keep the computational complexity of the method as low as possible.

There exists already a broad expertize in constructing such preconditioners and the question "Is there anything left to contribute?" is relevant. Our answer here is positive and we support it by bringing in focus the advantages of fully utilising the finite element discretization framework. There we are able to obtain approximations of the matrix blocks by assembly of matrices of small dimension, obtained by manipulation of local finite element matrices. In particular, within the construction phase we compute simultaneously sparse approximations of a matrix inverse, a Schur complement matrix and a matrix product, which latter is a novel idea and can significantly reduce the complexity of the action of the so-constructed preconditioner. Furthermore, we show that the proposed preconditioning technique is applicable within the adaptive finite element framework. Namely, we discuss the following preconditioner M to A,

$$M = \begin{bmatrix} B_{11} & 0 \\ A_{21} & S \end{bmatrix} \begin{bmatrix} I_1 & Z_{12} \\ 0 & I_2 \end{bmatrix}, \tag{1}$$

with S being an approximation of S_A, B_{11} (B_{11}^{-1}) being an approximation (approximate inverse) of A_{11} and Z_{12} - an approximation of the block matrix product $A_{11}^{-1}A_{12}$.

Here we draw the attention to the following three issues.

- As (1) indicates, we will assume that Z_{12} is a (sparse) approximation of the whole block $A_{11}^{-1}A_{12}$, which will reduce the computational complexity of applying the block-factorized preconditioner, avoiding one solution with B_{11}.
- As is clearly seen, all three matrix blocks B_{11} (B_{11}^{-1}), S and Z_{12}, depend directly or indirectly on A_{11}^{-1}. Our intention is now to decouple the approximations of these blocks by permitting different approximations of A_{11}^{-1} for each case.
- Since we will consider matrices which arise from finite element discretizations of (systems of) partial differential equations, to construct the above block matrix approximations we will use the fact that the element matrices will admit also a block two-by-two structure as the global matrix itself, namely, that

$$A = \sum_{k=1}^{M} R_k^T A_k R_k, \quad \text{where} \quad A_k = \begin{bmatrix} A_{k,11} & A_{k,12} \\ A_{k,21} & A_{k,22} \end{bmatrix}.$$

Here M is the number of (macro-)elements and R_k is a Boolean matrix which maps local to global numbering of the degrees of freedom.

Then we will advocate the following way to construct the block matrix approximations:

$$S = \sum_{k=1}^{M} R_k^T S_k R_k = \sum_{k=1}^{M} R_k^T (A_{k,22} - A_{k,21} A_{k,11}^{-1} A_{k,12}) R_k \tag{2}$$

$$B_{11}^{-1} = \sum_{k=1}^{M} R_k^T \widehat{A}_{k,11}^{-1} R_k \tag{3}$$

$$Z_{12} = \sum_{k=1}^{M} R_k^T \widehat{A}_{k,11}^{-1} A_{k,12} R_k \tag{4}$$

The element contributions $\widehat{A}_{k,11}^{-1}$ are inverted and properly scaled restrictions of the assembled global A_{11} block onto the kth (macro-)element.

The spectral equivalence (with a rather small constant) of S_A and S as in (2) has been shown for symmetric positive definite matrices in [6] and [1]. Clustering of the eigenvalues of $S^{-1}S_A$ around the unit number has been experimentally observed for non-symmetric saddle point problem arising from viscoelasticity in [3].

The quality of the approximations (3) and (4) for problems with discontinuous coefficients and anisotropies, discretized on arbitrary triangular meshes using linear finite elements, have been analysed in [4]. It is found that an approximation of A_{11}, assembled from restriction-based and scaled element-wise inverses does not depend on jumps in the problem coefficients and thus will never deteriorate as the classical element-by-element techniques do (observed and analysed in [7], for example).

How do we obtain a proper two-by-two block splitting of the system matrix? Such a structure is already present in a natural way for certain problems, such as for Navier-Stokes or Navier equations of elasticity formulated in coupled displacements and hydrostatic pressure variables. If we have only a scalar equation, we can permit the use of two consecutive nested mesh refinements and associate the matrix splitting with partitioning the nodes into coarse (belonging to the coarse mesh) and fine (or newly appeared) points, which belong to the fine mesh only. The latter, referred to as the two-level framework, is well understood for regular refinements. Not much is analysed in the AFEM setting, however.

In this talk we will describe the construction of the approximations in some detail. We will present theoretical estimates for the quality of these approximations for more general matrices. We will also discuss the applicability of the two-level preconditioner in an AFEM setting.

We point out that (i) by construction we guarantee sparse approximations and a set-up phase with high degree of parallelism, (ii) the approach is applicable in two or three dimensions as well as for scalar equations or systems of equations, (iii) can be applied to selfadjoint or non-selfadjoint problems, and (iv) the technique is naturally extendable to a multilevel setting.

The efficiency of the preconditioner (1) in standard FEM and in AFEM setting will be demonstrated on matrices arising from various applications, among which are problems with discontinuous coefficients, parabolic problems and problems formulates in terms of the so-called phase-field model, such as capillary driven flows or binary alloy solidification.

References

[1] O. Axelsson, R. Blaheta, M. Neytcheva. Preconditioning of boundary value problems using elementwise Schur complements. TR 2006-048, Department of Information Technology, Uppsala University, Sweden. Submitted.

[2] E. Bängtsson, M. Neytcheva. Preconditioning of nonsymmetric saddle point systems as arising in modelling of visco-elastic problems. 2006. Submitted.

[3] E. Bängtsson, M. Neytcheva. An agglomerate multilevel preconditioner for linear isostacy saddle point problems. In S. Margenov, I. Lirkov, J. Wasnievski (editors) *Proceedings of the Fifth International Conference on Large Scale Scientific Computations (LSSC'05)*, June 2005, Sozopol, Bulgaria, Springer *Lecture notes in Computer Science*, Vol. 3743, 2006, 113–120.

[4] E. Bängtsson, M. Neytcheva. Finite element block-factorized preconditioners. TR 2007-008, Department of Information Technology, Uppsala University, Sweden.

[5] J. Jones and P. Vassilevski, AMGe Based on Element Agglomeration, *SIAM Journal on Scientific Computing*, 23 (2001), pp. 109-133.

[6] J. Kraus. Algebraic multilevel preconditioning of finite element matrices using local Schur complements. *Numerical Linear Algebra with Applications*, 13 (2006), 49-70.

[7] H-C Lee, A.J. Wathen. On Element-by-Element preconditioning for general elliptic problems, *Computer Methods in Applied Mechanics and Engineering*, 92 (1991), 215-229.

Solving Sparse Least-Squares Problems Using Perturbed QR Factorizations

Esmond G. Ng

Let A be a large sparse $m \times n$ matrix and b be an m-vector, with $m \geq n$. We consider the solution of the linear least-squares problem $\min_x ||Ax - b||_2$ using an orthogonal factorization of A. Denote the orthogonal factorization by $A = Q \begin{bmatrix} R \\ O \end{bmatrix}$, where Q is an $m \times m$ orthogonal matrix and R is an $n \times n$ upper triangular matrix.

Assume for the moment that A has full column rank. Then R is the Cholesky factor of the symmetric and positive definite matrix $A^T A$, at least mathematically. Any rows in A that are relatively dense correspond to dense submatrices in $A^T A$, and hence will potentially make R dense. Thus, for sparsity reasons, it may not be desirable to compute Q and R in this case. A possibility is to ignore the dense rows in A when we compute Q and R, resulting in an orthogonal factorization of a smaller matrix.

On the other hand, if A is rank-deficient or ill-conditioned, then column pivoting is needed to ensure numerical stability in computing Q and R. However, when A is sparse, column pivoting may not be desirable, since the factorization algorithms and data structures will have to be flexible enough to accommodate the dynamic changes in sparsity structure

during numerical factorization. A possible solution in this case is to ignore column pivoting altogether, and to perturb the diagonal entry of R when a small element (in magnitude) is encountered. The perturbations can be accomplished by adding rows to the matrix A, resulting in an orthogonal factorization of a larger matrix.

Thus, it is sometimes desirable to consider the orthogonal factorization of a perturbed matrix $\bar{A}$, where $\bar{A}$ is obtained by adding rows to, or removing rows from, the initial matrix A. Let $\bar{A} = \bar{Q} \begin{bmatrix} \bar{R} \\ O \end{bmatrix}$ be the orthogonal factorization of the perturbed matrix $\bar{A}$. In previous work [1, 2, 3], $\bar{Q}$ and $\bar{R}$ are used to solve a correspondingly modified least-squares problem. Then the solution is modified via an updating procedure to provide a solution to the original least-squares problem. It is considered a direct approach.

Alternatively, we can consider an iterative approach. For example, we can solve the original least-squares problem using LSQR [4], with $\bar{R}$ as a preconditioner. We have investigated the quality of $\bar{R}$ as a preconditioner by analyzing the spectral properties of the preconditioned matrix. The results have showed that $\bar{R}$ can be an effective preconditioner for LSQR. In this talk, we summarize our findings.

This is joint work with Haim Avron and Sivan Toledo at Tel-Aviv University.

References

1. Bjorck, Ake, "A General Updating Algorithm for Constrained Linear Least Squares Problems", SIAM J. Sci. Stat. Comput., Vol. 5, No. 2 (1984), pp. 394-402.

2. Heath, Michael T., "Some Extensions of an Algorithm for Sparse Linear Least Squares Problems", SIAM J. Sci. Stat. Comput., Vol. 3, No. 2 (1982), pp. 223-237.

3. Ng, Esmond, "A Scheme for Handling Rank-Deficiency in the Solution of Sparse Linear Least Squares Problems", SIAM J. Sci. Stat. Comput., Vol. 12, No. 5 (1991), pp. 1173-1183.

4. Paige, Christopher C. and Michael A. Saunders, "LSQR: An Algorithm for Sparse Linear Equations and Sparse Least Squares", ACM Trans. Math. Softw., Vol. 8, No. 1 (1982), pp. 43-71.

Model Reduction in Variational Data Assimilation

N.K. Nichols, A. Bunse-Gerstner, C. Boess, A. Kubalinska

For the very large systems that arise in the environmental sciences, the available data are not sucient to initiate a complex computational forecasting model. The technique of data assimilation enables measured observations (over time) to be combined with dynamical model predictions to generate accurate estimates of the system states - both current and future. The problem of data assimilation can be treated either by sequential techniques based on optimal ltering methods, or by variational techniques that aim to solve an optimization problem subject to the dynamical system equations. For the very large inverse problems arising in the prediction of atmosphere and ocean circulations and other environmental systems,

approximate methods are needed for computational eciency in order to produce forecasts in real-time. Model reduction techniques from control theory provide an attractive method for improving the computational eciency of data assimilation schemes. Recently we have investigated the use of reduced order models together with Gauss-Newton iteration for solving four-dimensional variational data assimilation problems. Results obtained using a balanced-truncation model reduction method in a shallow water test case are presented. The results illustrate the benets of this approach and show that the same accuracy in the solution can be obtained more eciently by model reduction than by the standard procedures. A new model reduction technique using rational interpolation has also been tested for use in this context and is shown to have similar performance.

Regularization by Residual Periodograms

Dianne P. O'Leary, Bert W. Rust

Consider an ill-posed problem transformed if necessary so that the errors in the data are independent identically normally distributed with mean zero and variance one. We write the problem as

$$b = Ax^* + \eta\ , \quad \eta \sim N(\ 0\ ,\ I_m\ )\ .$$

where A is an $m \times n$ matrix, b is an $m \times 1$ vector of observations, and η is an unknown vector of errors.

We survey regularization and parameter selection from a linear algebra and statistics viewpoint and compare the statistical distributions of regularized estimates of the solution and the residual. By comparing the formulas for the truncated SVD regularization and the Tikhonov regularization, for example, we can see how each method increases the bias to reduce the variance in the estimators.

We discuss methods for choosing a regularization parameter in order to assure that the residual for the model is statistically plausible. Ideally, as proposed by Rust (1998, 2000), the results of candidate parameter choices should be evaluated by plotting the resulting residual $r = b - Ax$ along with its periodogram and its cumulative periodogram, and using the following three diagnostic criteria:

Diagnostic 1. The residual norm-squared should be within two standard deviations of the expected value of $\|\eta\|^2$, which is m. This is a way to quantify the Morozov discrepancy principle.

Diagnostic 2. The elements of η are drawn from a $N(0, 1)$ distribution, and a graph of the elements of the residual should look like samples from this distribution. (In fact, a histogram of the entries of the residual should look like a bell curve.)

Diagnostic 3. We consider the elements of both η and r as time series, with the index i ($i = 1, \ldots, m$) taken to be the time variable. Since $\eta \sim N(0, I_m)$, the η_i form a white noise series. Therefore the residuals r for an acceptable estimate should constitute a realization of such a series.

Sometimes an automated choice is needed, and we evaluate a method for choosing the regularization parameter that makes the residuals as close as possible to white noise, using the third diagnostic test, based on the periodogram of the residual. We compare this method with standard techniques such as the discrepancy principle, the L-curve, and generalized cross

validation, showing that it performs better on two new test problems as well as a variety of standard problems.

We show how our parameter choice method can be used for large-scale problems and for nonlinear models.

References

[1] M. Bertero and P. Boccacci, *Introduction to Inverse Problems in Imaging*, IOP, Bristol, 1998.

[2] W. Burrus and V. Verbinski, *Fast neutron spectroscopy with thick organic scintillators*, Nuclear Instruments and Methods, 67 (1969), pp. 181–196.

[3] H. Engl, M. Hanke, and A. Neubauer, *Regularization of Inverse Problems*, Kluwer Academic Publishers, 1996.

[4] W. A. Fuller, *Introduction to Statistical Time Series*, John Wiley & Sons, New York, 1996.

[5] P. C. Hansen, *Regularization tools: a Matlab package for analysis and solution of discrete ill-posed problems*, Numer. Algorithms, 6 (1994), pp. 1–35.

[6] ——, *Rank-deficient and discrete ill-posed problems*, SIAM Monographs on Mathematical Modeling and Computation, Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1998. Numerical aspects of linear inversion.

[7] P. C. Hansen, M. E. Kilmer, and R. H. Kjeldsen, *Exploiting residual information in the parameter choice for discrete ill-posed problems*, BIT, 46 (2006), pp. 41–59.

[8] M. E. Kilmer and D. P. O'Leary, *Choosing regularization parameters in iterative methods for ill-posed problems*, SIAM J. on Matrix Analysis and Applications, 22 (2001), pp. 1204–1221.

[9] B. W. Rust, *Parameter selection for constrained solutions to ill-posed problems*, Computing Science and Statistics, 32 (2000), pp. 333–347.

[10] ——, *Truncating the singular value decomposition for ill-posed problems*, Tech. Report NISTIR 6131, National Institute of Standards and Technology, U.S. Dept. of Commerce, Gaithersburg, MD, July 1998.

Preconditioning for Semidefinite Programming

Chen Greif, Ming Gu, Michael L. Overton

Semidefinite programming has become a valuable paradigm whose practical impact is mainly limited by the ill-conditioned large dense systems of linear equations that arise when implementing primal-dual interior-point methods. We propose preconditioners for solving these systems by iterative methods.

A One-sided Jacobi Method for Computing the Symplectic SVD

Archara Pacheenburawana, Niloufer Mackey

The SVD is one of the most useful decompositions in the theory of matrices. Recent work [1], [2] has shown that every symplectic matrix has a symplectic SVD, that is, $A = U\Sigma V$ where all three factors are symplectic, in addition to U and V being unitary, and Σ being a positive diagonal matrix of singular values. It is natural, therefore, to look for a structure preserving algorithm for computing a symplectic SVD of a symplectic matrix.

In this talk we present a structure preserving one-sided Jacobi algorithm to compute the symplectic SVD of real and complex symplectic matrices, i.e., matrices $A \in \mathbb{C}^{n\times n}$ such that $AJA^T = J$, where $J = \begin{bmatrix} 0 & I_n \\ -I_n & 0 \end{bmatrix}$.

Our algorithm uses symplectic orthogonal transformations that were first described in [3]. These transformations allow us to access off-diagonal entries that cannot be reached by 2×2 based symplectic plane rotations, or by double symplectic rotations. Sorting versions can be efficiently computed at each iteration so that the singular values, which are reciprocally paired, appear in "symplectic order" on the diagonal at convergence. We show how to design effective sweep patterns for our algorithms, and present numerical experiments showing that our algorithms provide well structured factors, and have good convergence behavior [4], [5], even though the off-diagonal norm is not guaranteed to decrease monotonically.

References

[1] H.Fassbender and K.Ikramov. Several observations on symplectic, Hamiltonian, and skew-Hamiltonian matrices. *Linear Algebra Appl.*, 400:15–29, 2005.

[2] Hongguo Xu. An SVD-like matrix decomposition and its applications. *Linear Algebra Appl.*, 368:1–24, 2003.

[3] D. Steven Mackey, Niloufer Mackey and Françoise Tisseur. Structured Tools for Structured Matrices. *The Electronic Journal of Linear Algebra*, 10:106–145, 2003.

[4] Archara Chaiyakarn. Structure Preserving Algorithms for Computing the Symplectic Singular Value Decomposition. PhD Thesis, Western Michigan University, 2005.

[5] Archara Pacheenburawana and Niloufer Mackey. A one-sided Jacobi method for computing the symplectic SVD. In preparation.

Orthonormal Completion of an array of Unit Length Vectors

Chris Paige

Charles Sheffield [6] pointed out that the modified Gram–Schmidt (MGS) orthogonalization algorithm for the QR factorization of $B \in \mathbf{R}^{m\times n}$ can be interpreted as a QR factorization applied to the matrix B augmented with a square matrix of zero elements on top. This is

true in theory for *any* method of QR factorization, but for Householder's method *it is true in the presence of rounding errors as well*. Björck and Paige [2] proved this rigorously, and used it to further analyze the behavior of the MGS algorithm and some of its applications, and in that paper and [3] described several practical uses of this fascinating knowledge. Paige, Rozložník, and Strakoš [4] used the same idea to prove the numerical backward stability of the MGS-GMRES method of Saad and Schultz [5], a result that had been sought for decades.

Sheffield's result has thus been of great value in our understanding of widely used numerical algorithms. It turns out that his result can be reformulated and generalized—to become the "Orthonormal Completion of *any* sequence of Unit Length Vectors" (not just MGS vectors). This new idea appears to be important theoretically, and to facilitate the analysis and understanding of other practical algorithms. We will present and illustrate this idea, show its relationship to Householder transformations, indicate some theoretical results that follow, and show how this can make some previously difficult analyses short and understandable, offering hope for the successful analysis of other important algorithms. For example the paper by Barlow, Bosner and Drmač [1] contains a thorough but difficult 24 page rounding error analysis. The new idea here allows us to analyze their bidiagonalization algorithm quickly and clearly.

References

[1] J.L. Barlow, N. Bosner and Z. Drmač, *A new stable bidiagonal reduction algorithm*, Lin. Alg. Appl. 397:35–84, 2005.

[2] Å. Björck and C.C. Paige, *Loss and recapture of orthogonality in the modified Gram-Schmidt algorithm,* SIAM J. Matrix Anal. Appl. 13:176–190, 1992.

[3] Å. Björck and C.C. Paige, *Solution of Augmented Linear Systems using Orthogonal Factorizations*, BIT 34:1–24, 1994.

[4] C.C. Paige, M. Rozložník, and Z. Strakoš, *Modified Gram-Schmidt (MGS), Least Squares, and backward stability of MGS-GMRES*, SIAM J. Matrix Anal. Appl. 28:264–284, 2006.

[5] Y. Saad and M. H. Schultz, *GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems*, SIAM J. Sci. Statist. Comput. 7:856–869, 1986.

[6] Charles Sheffield (June 25, 1935 – November 2, 2002), *Comment to Gene Golub.*

Nonnegative Least Squares and Extensions

Bob Plemmons

A brief overview of the development of direct and iterative algorithms for nonnegative least squares (NNLS) computations is given, from Lawson and Hansen (1974) to present. Beginning in the 1990s, NNLS computations have been generalized to approximate nonnegative matrix and tensor factorizations (NMF and NTF), in order to obtain low-dimensional representations of arrays of nonnegative data. A suitable representation for nonnegative data is essential to applications in fields such as statistics, signal and image processing, machine learning, and data mining. Low rank constraints on high dimensional massive data sets are prevalent in dimensionality reduction and data analysis across numerous scientific disciplines. This talk is concerned recent work on algorithms for NMF and their extensions to NTF. Applications to space object identification using hyperspectral data, with numerical results, are provided.

Jacobi–Davidson Method for Two-Parameter Eigenvalue Problems

Michiel E. Hochstenbach, Bor Plestenjak

Overview. We show how to modify the Jacobi–Davidson method so that it can be applied to two-parameter eigenvalue problems. This approach, in particular when it is combined with the recently developed harmonic and refined extraction methods [7], is very efficient in computing eigenvalues close to a given target.

Introduction. We consider the *two-parameter eigenvalue problem*

$$\begin{aligned} A_1x_1 &= \lambda B_1x_1 + \mu C_1x_1 \\ A_2x_2 &= \lambda B_2x_2 + \mu C_2x_2, \end{aligned} \tag{1}$$

where A_i, B_i, and C_i are given $n_i \times n_i$ complex matrices, $\lambda, \mu \in \mathbf{C}$, and nonzero $x_i \in \mathbf{C}^{n_i}$ for $i = 1, 2$. A pair (λ, μ) is an *eigenvalue* and $x_1 \otimes x_2$ is the corresponding *eigenvector*. Problems of this kind arise in a variety of applications [1], particularly when the method of separation of variables is used to solve boundary value problems [11]; see [6] for several other applications.

In the tensor product space $S := \mathbf{C}^{n_1} \otimes \mathbf{C}^{n_2}$ we can define *operator determinants*

$$\begin{aligned} \Delta_0 &= B_1 \otimes C_2 - C_1 \otimes B_2 \\ \Delta_1 &= A_1 \otimes C_2 - C_1 \otimes A_2 \\ \Delta_2 &= B_1 \otimes A_2 - A_1 \otimes B_2 \end{aligned}$$

and show that (1) is equivalent to the coupled generalized eigenvalue problem

$$\begin{aligned} \Delta_1 z &= \lambda \Delta_0 z \\ \Delta_2 z &= \mu \Delta_0 z \end{aligned} \tag{2}$$

for the decomposable tensor $z = x_1 \otimes x_2$ [2]. We say that (1) is *nonsingular* if Δ_0 is nonsingular and in this case $\Delta_0^{-1}\Delta_1$ and $\Delta_0^{-1}\Delta_2$ commute.

Instead of (1) one can solve the associated coupled generalized eigenvalue problem (2). In the right definite case, where matrices A_i, B_i, C_i are Hermitian and Δ_0 is positive definite, this can be achieved by numerical methods for simultaneous diagonalization of commutative symmetric matrices [3, 10], while an algorithm for the general nonsingular case using QZ algorithm is presented in [6]. However, this approach is only feasible for problems of low dimension.

For larger problems, if we want to compute all the eigenvalues, the continuation method can be applied [8, 9], while if we are interested in a part of the eigenvalues close to a given target, the Jacobi–Davidson method [5, 6, 7] gives very good results.

Harmonic extraction. Let the columns of the matrix $U_{ik} \in \mathbf{C}^{n_i \times k}$ form an orthonormal basis for the search space $\mathcal{U}_{ik}$ for the ith equation (1) for $i = 1, 2$. The Ritz conditions

$$\begin{aligned} (A_1 - \theta B_1 - \eta C_1)\, u_1 &\perp \mathcal{U}_{1k} \\ (A_2 - \theta B_2 - \eta C_2)\, u_2 &\perp \mathcal{U}_{2k} \end{aligned} \tag{3}$$

lead to the smaller projected two-parameter eigenvalue problem

$$\begin{aligned} U_{1k}^T A_1 U_{1k} c_1 &= \theta\, U_{1k}^T B_1 U_{1k} c_1 + \eta\, U_{1k}^T C_1 U_{1k} c_1 \\ U_{2k}^T A_2 U_{2k} c_2 &= \theta\, U_{2k}^T B_2 U_{2k} c_2 + \eta\, U_{1k}^T C_2 U_{2k} c_2, \end{aligned}$$

where $u_i = U_{ik}c_i$, $c_i \in \mathbf{C}^k$ for $i = 1, 2$, and $\theta, \eta \in \mathbf{C}$. Ritz values (θ, η) and Ritz vectors $u_1 \otimes u_2$ are approximations to the eigenpairs of (1). This extraction, combined with an appropriate correction equation in a Jacobi–Davidson type subspace method, works well for exterior eigenvalues of right definite problems [5]. For non right definite problems we get better results if we apply the two-sided Jacobi–Davidson method [6].

For interior eigenvalues close to a target (σ, τ) one should use harmonic extraction. We were able to generalize harmonic Ritz vectors by imposing the Galerkin conditions

$$\begin{array}{lcl} (A_1 - \theta B_1 - \eta C_1)\, u_1 & \perp & (A_1 - \sigma B_1 - \tau C_1)\, \mathcal{U}_{1k} \\ (A_2 - \theta B_2 - \eta C_2)\, u_2 & \perp & (A_2 - \sigma B_2 - \tau C_2)\, \mathcal{U}_{2k}, \end{array} \tag{4}$$

which again lead to a smaller projected two-parameter eigenvalue problem. This approach works well for interior eigenvalues and we can also generalize Saad type theorems [7].

Algorithm. We give a sketch of the Jacobi–Davidson type method for the nonsingular two-parameter eigenvalue problem. Details can be found in [5, 6, 7].

1. *Start.* Choose initial vectors u_1, u_2, set $U_{11} = [u_1]$, $U_{12} = [u_2]$, and $k = 1$.
2. *Iterate.* Until convergence or $k > k_{\max}$ do:
   a) Compute Ritz or harmonic Ritz pairs of the small two-parameter eigenvalue problem obtained from (3) or (4) using the associated coupled generalized eigenvalue problem.
   b) Select an appropriate eigenpair approximation $((\theta_k, \eta_k),\ u_{1k} \otimes u_{2k})$.
   c) Stop if $\|(A_i - \theta_k B_i - \eta_k C_i)u_{ik}\| \le \varepsilon$ for $i = 1, 2$.
   d) Solve approximately correction equations and obtain new directions s_1 and s_2.
   e) Expand. Set $U_{i,k+1} = \mathrm{MGS}\,(U_{ik}, s_i)$ for $i = 1, 2$.
   f) Set $k = k + 1$.
   g) Restart. If the number of columns of U_{ik} exceeds $l_{\max}$, then replace U_{ik} with new orthonormal basis of dimension $l_{\min}$ for $i = 1, 2$.

Future work. In several practical applications (see, e.g., [4]) two-parameter eigenvalue problems are singular. Some theoretical results and numerical methods for singular two-parameter eigenvalue problems will be presented.

References

[1] F. V. Atkinson, *Multiparameter spectral theory*, Bull. Amer. Math. Soc., 74 (1968), pp. 1–27.

[2] F. V. Atkinson, *Multiparameter eigenvalue problems*, Academic Press, New York, 1972.

[3] A. Bunse-Gerstner, R. Byers, V. Mehrmann, *Numerical methods for simultaneous diagonalization*, SIAM J. Matrix Anal. Appl., 14 (1993), pp. 927–949.

[4] N. Cottin, *Dynamic model updating — a multiparameter eigenvalue problem*, Mech. Syst. Signal Pr., 15 (2001), pp. 649–665.

[5] M. E. Hochstenbach and B. Plestenjak, *A Jacobi–Davidson type method for a right definite two-parameter eigenvalue problem*, SIAM J. Matrix Anal. Appl., 24 (2002), pp. 392–410.

[6] M. E. Hochstenbach, T. Košir, and B. Plestenjak, *A Jacobi–Davidson type method for the nonsingular two-parameter eigenvalue problem*, SIAM J. Matrix Anal. Appl., 26 (2005), pp. 477–497.

[7] M. E. Hochstenbach and B. Plestenjak, *Harmonic Rayleigh–Ritz extraction for the multiparameter eigenvalue problem*, to appear in ETNA.

[8] B. Plestenjak, *A continuation method for a right definite two-parameter eigenvalue problem*, SIAM J. Matrix Anal. Appl., 21 (2000), pp. 1163–1184.

[9] B. Plestenjak, *A continuation method for a weakly elliptic two-parameter eigenvalue problem*, IMA J. Numer. Anal., 21 (2001), pp. 199–216.

[10] T. Slivnik, G. Tomšič, *A numerical method for the solution of two-parameter eigenvalue problems*, J. Comput. Appl. Math., 15 (1986), pp. 109–115.

[11] H. Volkmer, *Multiparameter Eigenvalue Problems and Expansion Theorems*, Springer-Verlag, Berlin, 1988.

LSTRS: MATLAB Software for Large-Scale Trust-Region Subproblems and Regularization

Marielba Rojas, A. Santos, Danny C. Sorensen

We describe a MATLAB implementation [6] of the LSTRS method [5] for the large-scale trust-region subproblem:

$$\min \quad \frac{1}{2}x^T Hx + g^T x \quad \text{subject to (s.t.)} \quad \|x\|_2 \leq \Delta, \tag{1}$$

where H is an $n \times n$, real, large, symmetric matrix, g is an n-dimensional real vector, and Δ is a positive scalar. Problem (1) arises in connection with the trust-region globalization strategy in optimization. A special case of problem (1), namely, a least squares problem with a norm constraint, is equivalent to Tikhonov regularization [7] for discrete forms of ill-posed problems.

LSTRS is based on a reformulation of the trust-region subproblem as a parameterized eigenvalue problem, and consists of an iterative procedure that finds the optimal value for a scalar parameter. The adjustment of the parameter requires the solution of a large-scale eigenvalue problem at each step. The method relies on matrix-vector products only and has low and fixed storage requirements, features that make it suitable for large-scale computations. In the MATLAB implementation, the Hessian matrix of the quadratic objective function can be specified either explicitly, or in the form of a matrix-vector multiplication routine. Therefore, the implementation preserves the matrix-free nature of the method. The MATLAB implementation offers several choices for the eigenvalue calculation and it also allows the users to specify their own eigensolver routine.

We present a brief description of the LSTRS method from [5] and describe the main components and features of the MATLAB software, with special emphasis on the large eigenvalue computations. We include comparisons with the following state-of-the-art techniques for solving large-scale instances of problem (1): the Semidefinite Programming approach of Fortin and Wolkowicz [1], the Sequential Subspace Method of Hager [3], and the Generalized Lanczos Trust Region method of Gould et al. [2] as implemented in the HSL library [4]. We present examples of use of the LSTRS software as well as results for large-scale regularization problems.

References

[1] C. Fortin and H. Wolkowicz. The trust-region subproblem and semidefinite programming, 2003. *Optimization Methods and Software*, to appear in the special issue dedicated to Jochem Zowe's 60th birthday.

[2] N. Gould, S. Lucidi, M. Roma and Ph. L. Toint. Solving the trust–region subproblem using the Lanczos method, *SIAM J. Optim.*, 9(2):504–525, 1999.

[3] W. W. Hager. Minimizing a quadratic over a sphere, *SIAM J. Optim.*, 12:188-208, 2001.

[4] HSL 2004: A collection of Fortran codes for large scale scientific computation. Available from `www.cse.clrc.ac.uk/nag/hsl`.

[5] M. Rojas, S.A. Santos and D.C. Sorensen. A new matrix-free algorithm for the large-scale trust-region subproblem, *SIAM J. Optim.*, 11(3):611-646, 2000.

[6] M. Rojas, S.A. Santos and D.C. Sorensen. LSTRS: MATLAB Software for Large-Scale Trust-Region Subproblems and Regularization, *ACM Trans. Math. Software*, to appear.

[7] A.N. Tikhonov. Regularization of incorrectly posed problems, *Soviet Math.*, 4:1624-1627, 1963.

Numerical behavior of saddle–point solvers

Miro Rozložník, Pavel Jiránek

Saddle point problems have recently attracted a lot of attention and appear to be a time-critical component in the solution of large-scale problems in many applications of computational science and engineering. A large amount of work has been devoted to the construction of various iterative solvers and preconditioning techniques for their solution. For a recent survey on applications, methods, and results on saddle point problems, we refer to [1] and to numerous references therein. Significantly less attention has been paid to numerical stability issues related to saddle point solvers. In this contribution we study numerical behavior of several Krylov subspace solvers applied to

$$\begin{pmatrix} A & B \\ B^T & 0 \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} f \\ 0 \end{pmatrix}, \tag{1}$$

where A is a square nonsingular matrix and B has full column rank. Two main approaches are analyzed: the Schur complement reduction method, based on the elimination of primary variables, and the null-space projection method which relies on a basis for the null-space for the constraints. In practical situations, there is a need for cheap approximate solutions to the inner systems with the matrix A or to the associated least squares problem with the block B. Numerous inexact schemes have been used and analyzed, including Uzawa algorithms, null-space methods, multilevel or multigrid methods, domain decomposition methods or inner-outer iterations [3, 4, 5, 8, 2]. These works have focused mainly on the analysis of the convergence delay caused by the inaccurate solution of inner problems. In this talk we concentrate on the complementary question. We analyze the best possible accuracy which we can get from inexact schemes solving either the Schur complement or null-space projected system.

If we consider the Schur complement reduction method and discuss algorithms which compute simultaneously the approximate solutions x_{k+1} and y_{k+1} to the unknowns x and y, satisfying ideally $Ax_{k+1} + By_{k+1} = f$ with y_{k+1} obtained from solving the Schur complement system $B^T A^{-1} By = B^T A^{-1} f$. Without specifying any particular method in the outer iteration process, we assume that y_{k+1} is computed via $y_{k+1} = y_k + \alpha_k p_k^{(y)}$. We distinguish between three mathematically equivalent back-substitution formulas

$$x_{k+1} = A^{-1}(f - By_{k+1}), \tag{2}$$

$$x_{k+1} = x_k + \alpha_k(-A^{-1}Bp_k^{(y)}), \tag{3}$$

$$x_{k+1} = x_k + A^{-1}(f - Ax_k - By_{k+1}). \tag{4}$$

These schemes have been used and studied in the context of the classical Uzawa algorithm, two-level pressure correction approach, or inner-outer iteration method; see, e.g., the schemes using (2) in [4], (3) in [8] or (4) in [3], respectively. Because the solves with the matrix A in formulas (2)-(4) are expensive, these systems are in practice solved only approximately. Our analysis [6, 7] is based on the assumption that every solution of a system with A is replaced by an approximate (inexact) solution produced by some (possibly iterative) method. The resulting vector is then interpreted as an exact solution of the system with the same right-hand side vector but with a perturbed matrix $A + \Delta A$. We always consider that the relative norm of the perturbation matrix is bounded by $\|\Delta A\| \leq \tau \|A\|$, where τ represents the normwise relative backward error measuring the level of inexactness in the computation. Note that if $\tau = O(u)$ then we have a backward stable method for solving the systems with A. When considering the outer iteration process for solving the Schur complement system it appears that the maximum attainable accuracy of the computed approximate solutions $\bar{y}_{k+1}$ in terms of the residual $B^T A^{-1} f - B^T A^{-1} B\bar{y}_{k+1}$ is mainly given by the inexactness in solving the inner problems (measured by τ) and it is not further magnified by rounding errors. The situation is different when looking at the behavior of residuals $f - A\bar{x}_{k+1} - B\bar{y}_{k+1}$ and $-B^T\bar{x}_{k+1}$ for the original saddle point system, which describe how accurately the two block equations of (1) are satisfied. Indeed, the maximum attainable accuracy then depends significantly on the back-substitution formula used for computing the vectors $\bar{x}_{k+1}$. Our results confirm that independent of the fact that the inner systems are solved inexactly (τ is usually much larger than the machine precision u), back-substitution formulas (3) and (4) lead ultimately to the residuals on the $O(u)$ level. Depending on whether we use (3) or (4), the computed iterates may satisfy either the first or the second block equation of (1) to the working accuracy, while the most straightforward scheme (2) leads to both residuals

$f - A\bar{x}_{k+1} - B\bar{y}_{k+1}$ and $-B^T\bar{x}_{k+1}$ on the $O(\tau)$ level.

Analogous results have been obtained also for the null-space projection method. We assume that in the outer iteration the projected system $(I-BB^\dagger)A(I-BB^\dagger)x = (I-BB^\dagger)f$ is solved using the recursion $x_{k+1} = x_k + \alpha_k p_k^{(x)}$ so that the approximate solutions x_{k+1} satisfy $B^T x_{k+1} = 0$. Again, we distinguish between three back-substitution formulas

$$y_{k+1} = B^\dagger(f - Ax_{k+1}), \tag{5}$$

$$y_{k+1} = y_k + B^\dagger(r_k^{(x)} - \alpha_k A p_k^{(x)}),\ r_{k+1}^{(x)} = (I - BB^\dagger)(r_k^{(x)} - \alpha_k A p_k^{(x)}), \tag{6}$$

$$y_{k+1} = y_k + B^\dagger(f - Ax_{k+1} - By_k). \tag{7}$$

Similar inexact schemes have been considered in the null-space methods for solving quadratic programming problems, in multigrid methods, or as constraint preconditioners [2, 5, 9]. The pseudoinverses $B^\dagger$ in (5)-(7) are not computed explicitly and the operation is performed by solving the associated least squares problems with the block B. These problems are in practice solved inexactly and we interpret the computed results in terms of the backward error τ. It appears that the residuals $f - A\bar{x}_{k+1} - B\bar{y}_{k+1}$ for the schemes (6) and (7) will finally reach $O(u)$, while the scheme (5) leads to the accuracy that is only $O(\tau)$. The null-space approach has been studied also from another perspective in the framework of the indefinite constraint preconditioner applied to the saddle point system (1). The numerical behavior of the preconditioned conjugate gradient (PCG) algorithm has been studied in [9] and it was shown that the residual update strategy (6) (see also [5]) leads to highly accurate residuals $f - A\bar{x}_{k+1} - B\bar{y}_{k+1}$ when using PCG for the case A symmetric and positive definite. Theory for a more complicated situation, when nonsymmetric iterative methods are applied, was developed in [7].

Our results are an example of rather general fact that it is advantageous to use the update formulas in the form "new_value = old_value + small_correction". Indeed, the generic and actually the cheapest implementations (3) or (6) deliver the approximate solutions which satisfy either the second or the first block equation to the working accuracy. In addition, the schemes with a corrected direct substitution (4) and (7) are also numerically very attractive. This behavior was probably observed or is already tacitly known. However, we haven't found any explicit reference to these issues and our results thus provide a necessary theoretical explanation. The implementations that we pointed out as optimal are actually those which are widely used and suggested in applications. We believe that ideas on inexact solution of inner problems and similar backward analysis can be successfully used also in a different context, namely for inexact preconditioning of saddle point or general linear systems.

References

[1] M. Benzi, G.H. Golub, and J. Liesen, *Numerical solution of saddle point problems*, Acta Numerica 14 (2005), pp. 1–137.

[2] D. Braess, P. Deuflhard, K. Lipnikov, *A subspace cascadic multigrid method for mortar elements*, Computing 69 (2002), pp. 205-225.

[3] J.H. Bramble, J.E. Pasciak, and A.T. Vassilev, *Inexact Uzawa algorithms for non-symmetric saddle point problems*, Math. Comp. 69 (2000), pp. 667–689.

[4] H.C. Elman, G.H. Golub, *Inexact and preconditioned Uzawa algorithms for saddle point problems*, SIAM J. Numer. Anal. 31 (1994), pp. 1645–1661.

[5] N.I.M. Gould, M.E. Hribar, J. Nocedal, *On solution of equality constrained quadratic programming problems arising in optimization*, SIAM J.Comput. 23(2001), 1375-1394.

[6] P. Jiránek, M. Rozložník, *Maximum attainable accuracy of inexact saddle point solvers.* SIAM J. Matrix Anal., to appear.

[7] P. Jiránek, M. Rozložník, *Limiting accuracy of segregated solution methods for non-symmetric saddle point problems.* Journal od Comput. and Appl. Math., to appear.

[8] A. Ramage, A.J. Wathen, *Iterative solution techniques for the Stokes and Navier-Stokes equations*, Internat. J. Numer. Methods Fluids 19 (1994), pp. 67–83.

[9] M. Rozložník and V. Simoncini, *Krylov subspace methods for saddle point problems with indefinite preconditioning*, SIAM J. Matrix Anal. Appl., 24 (2002), pp. 368–391.

Getting orthogonal eigenvectors with no reorthogonalization

Axel Ruhe

We have looked for a method to compute accurate eigenvalues and reasonably orthogonal eigenvectors to a large symmetric matrix, without doing any time consuming reorthogonalizations of vectors.

We have used the Lanczos algorithm [4] without reorthogonalization, followed by the Multiple Relative Robust Representation MRRR algorithm [2], on the resulting tridiagonal matrix T_m. We have found that a variant of the Cullum device [1] can be used to choose those eigenvalues θ_i of the tridiagonal matrix T_m that give good Ritz approximations to the eigenvalues λ_i of the original matrix A. The behaviour of the Lanczos algorithm without reorthogonalization has been studied by Greenbaum [3], who proved that all converged Ritz values θ_i are in small neighbourhoods of eigenvalues λ_i of A, even in cases of lost orthogonality in the basis V, but left open the question whether *all* eigenvalues λ_i would eventually get any close converged Ritz value θ_i.

Our variant of the Cullum device [1], is to choose those Ritz values θ_i, whose components $s_{1,i}$ in the starting vector v_1 are nonzero. We found that the residuals of these Ritz vectors were at the rounding error level, while their orthogonality to the other converged Ritz vectors was larger by a factor. Not surprisingly, we found that the dominating computing time was spent in the matrix product $X = VS$ getting eigenvectors X for the original matrix A from those S of the tridiagonal matrix T.

The reader is referred to the report [5] for further details.

References

[1] J. Cullum and R. A. Willoughby, *Computing eigenvalues of very large symmetric matrices — an implementation of a Lanczos algorithm with no reorthogonalization*, Journal of Computational Physics, 44 (1984), pp. 329–358.

[2] I. S. Dhillon, B. N. Parlett, and C. Vömel, *The design and implementation of the MRRR algorithm*, ACM Trans. Math. Softw., 32 (2006), pp. 533–560.

[3] A. Greenbaum, *Behavior of slightly perturbed Lanczos and conjugate-gradient recurrences*, Linear Algebra Appl., 113 (1989), pp. 7–63.

[4] C. Lanczos, *An iteration method for the solution of the eigenvalue problem of linear differential and integral operators*, Journal of Research of the National Bureau of Standards, 45 (1950), pp. 255–282.

[5] J. Rydh, *Very large symmetric eigenvalue problems solved by Lanczos and multiple relatively robust representations for tridiagonals*, master's thesis, Computer Science and Communications, Royal Institute of Technology, Stockholm, 2007.

Efficient linear algebra methods in data mining

Yousef Saad

Linear algebra techniques constitute the tool of choice in many disciplines loosely grouped under the area of "Data Mining". This area includes information retrieval, pattern recognition, classification, learning, image analysis, recommender systems, etc.. It is the source of many new, interesting, and sometimes challenging linear algebra problems. In fact, one can argue that while the current era of numerical linear algebra has been shaped largely by problems arising from PDEs and CFD, these new disciplines are now starting to shape the "next chapter" in numerical linear algebra. Two common themes in all of these disciplines are those of (1) clustering (trying to group similar objects together) and (2) dimension reduction (finding some low-dimensional representation of some data). We will begin by illustrating some of these these problems and then show a few numerical algorithms for solving them. We will begin by showing how approximation theory can be used to solve the query problem in information retrieval. Then we will discuss graph-based methods and show how the SVD and the Lanczos algorithm come to the rescue in these methods.

Grassmann-Quasi-Newton Algorithms for Best Multilinear Rank Approximation of Tensors

Berkant Savas, Lek-Heng Lim

We consider the problem of approximating a given tensor $\mathcal{A} \in \mathbf{R}^{J\times K\times L}$ by another tensor $\mathcal{B}$ of equal dimensions but of lower multilinear rank [7],

$$\min_{\mathcal{B}} \|\mathcal{A} - \mathcal{B}\|_F \,. \tag{1}$$

In the matrix case, the solution to the corresponding problem is given by the truncated SVD [1]. But truncating the Higher Order SVD will in general not give the best low rank approximation [5, 6]. Assuming $\operatorname{rank}(\mathcal{B}) = (r_1, r_2, r_3)$ we can express $\mathcal{B}$ as product of a core tensor $\mathcal{C} \in \mathbf{R}^{r_1\times r_2\times r_3}$ and orthonormal matrices $X \in \mathbf{R}^{I\times r_1}$, $Y \in \mathbf{R}^{J\times r_2}$ and $Z \in \mathbf{R}^{K\times r_3}$,

i.e. $[\mathcal{B} = (X, Y, Z) \cdot \mathcal{C} \quad \text{where} \quad \mathcal{B}(i,j,k) = b_{ijk} = \sum_{\lambda\mu\nu} x_{i\lambda} y_{j\mu} z_{k\nu} c_{\lambda\mu\nu}.]$ Using this factorization one can show that the minimization problem (1) is equivalent to $[\max_{X,Y,Z} \|\mathcal{A} \cdot (X, Y, Z)\|_F \quad \text{subject to} \quad X\mathsf{TX} = \mathsf{I}, \mathsf{YTY} = \mathsf{I}, \mathsf{ZTZ} = \mathsf{I},]$ where $\mathcal{A}\cdot(X, Y, Z) \equiv (X\mathsf{T}, \mathsf{YT}, \mathsf{ZT})\cdot\mathcal{A}$ and the core tensor $\mathcal{C}$ is no longer present. We define the objective function to be maximized, $[\Phi(X, Y, Z) = \|\mathcal{A} \cdot (X, Y, Z)\|_F^2 = \sum_{ijk} A_{ijk}^2, \quad A_{ijk} = \sum_{\lambda\mu\nu} a_{\lambda\mu\nu} x_{\lambda i} y_{\mu j} z_{\nu k}.]$ Due to the invariance of the norm under orthogonal transformations we have $\Phi(X, Y, Z) = \Phi(XU, YV, ZW)$ for orthogonal matrices U, V and W. The maximization problem is over-parameterized, i.e. we are interested in the subspaces spanned by X, Y, Z and not specific basis for the different subspaces. This means that the objective function is maximized over a product Grassmann manifolds [2, 3].

In this talk we will present three Quasi-Newton algorithms for solving the best multilinear rank approximation problem. All three algorithms exploit the fact that the problem may be viewed as an optimization problem over a product of Grassmann manifolds. The first algorithm uses BFGS updates in global coordinates, i.e. the computations are made in the Euclidean spaces in which the Grassmann manifolds are embedded. The second algorithm uses BFGS updates also but in local coordinates. Here the computations are made on the tangent spaces to various points of the three manifolds. And the last algorithm is based on Limited memory BFGS updates in global coordinates suited for large scale problems.

We will also present experimental results and discuss the performance of the different algorithms in terms of computational efficiency and memory usage. Tensor approximation problems occur in various applications involving signal processing, pattern classification, biomedicine and multidimensional data, e.g. [4, 8].

References

[1] C. Eckart and G. Young. The approximation of one matrix by another of lower rank. *Psychometrika*, 1:211–218, 1936.

[2] Alan Edelman, Tomás A. Arias, and Steven T. Smith. The geometry of algorithms with orthogonality constraints. *SIAM J. Matrix Anal. Appl.*, 20(2):303–353, 1998.

[3] Lars Eldn and Berkant Savas. A newton-grassmann method for computing the best multilinear rank-(r_1, r_2, r_3) approximation of a tensor. Technical Report LITH-MAT-R-2007-6-SE, Department of Mathematics, Linköpings Universitet, 2007.

[4] Lieven De Lathauwer, L. Hoegaerts, and J. Vandewalle. A grassmann-rayleigh quotient iteration for dimensionality reduction in ica. In *Proc. 5th Int. Workshop on Independent Component Analysis and Blind Signal Separation (ICA 2004)*, pages 335–342, Granada, Spain, Sept. 2004.

[5] Lieven De Lathauwer, Bart De Moor, and Joos Vandewalle. A multilinear singular value decomposition. *SIAM Journal on Matrix Analysis and Applications*, 21(4):1253–1278, 2000.

[6] Lieven De Lathauwer, Bart De Moor, and Joos Vandewalle. On the best rank-1 and rank-$(r_1, r_2, ..., r_n)$ approximation of higher-order tensors. *SIAM Journal on Matrix Analysis and Applications*, 21(4):1324–1342, 2000.

[7] Vin De Silva and Lek-Heng Lim. Tensor rank and the ill-posedness of the best low-rank approximation problem. *to appear in SIAM Journal on Matrix Analysis and Applications*, 2006. (accepted October 2006).

[8] M. Alex O. Vasilescu and Demetri Terzopoulos. Multilinear subspace analysis of image ensembles. In *Proc. Computer Vision and Pattern Recognition Conf. (CVPR'03)*, volume 2, pages 93–99, Madison WI, 2003.

Numerical Solution of Palindromic Eigenvalue Problems

Christian Schröder

The polynomial palindromic eigenvalue problem [4] is defined as

$$P(\lambda)x = (\sum_{i=0}^{k} \lambda^i A_i)x = 0 \quad \text{with} \quad A_i = A_{k-i}^T, \text{ for } i = 0, \ldots, n, \tag{1}$$

where A_i, $i = 0, \ldots, n$ are given complex n-by-n matrices, whereas the nonzero n-vector x and the scalar λ are wanted. It arises in a variety of applications, for example, in the optimal control problem or in the noise modeling of high speed trains [4]. The name 'palindromic' stems from the observation that the polynomial $P(\lambda)$ is invariant under reversing the order of the coefficients and transposing, and palindromes are words or sentences that can be read in either direction, like 'rotor' or 'A man, a plan, a canal, Panama.'

The symmetry in the coefficients of $P(\lambda)$ causes the eigenvalues to appear in reciprocal pairs $(\lambda, \frac{1}{\lambda})$. This fact motivates the research for numerical algorithms that compute correctly paired eigenvalues in finite precision arithmetic.

The polynomial problem (1) can be linearized in a structure preserving way [4] yielding a linear palindromic eigenvalue problem

$$Ax = \lambda A^T x. \tag{2}$$

Known algorithms for this class of problems include variants of the QR-, and Jacobi algorithms [3, 2]. These methods aim at transforming A to palindromic Schur form [2], i.e., at computing a unitary matrix Q such that

$$T := Q^T A Q$$

is anti triangular, that is, $t_{ij} = 0$ whenever $i + j \leq n$. The eigenvalues are then given by $\lambda_i = t_{i,n-i+1}/t_{n-i+1,i}$, which are paired as $\lambda_i = 1/\lambda_{n-i+1}$.

Another method, the anti-triangular URV-algorithm [4] produces two unitary matrices U, V such that

$$U^T A V, \quad U^T(A - A^T)U, \quad \text{and} \quad V^T(A - A^T)V$$

are anti triangular. These products can be interpreted as periodic Schur form of a product eigenvalue problem related to (2) and thus also reveal the desired eigenvalues.

The talk will be on recent developments in the area of palindromic eigenvalue computations, like an implicit version of the palindromic QR algorithm using bulge chasing, or a method to transform the anti triangular URV decomposition into a palindromic Schur form.

References

[1] D. Steven Mackey, Niloufer Mackey, Christian Mehl, and Volker Mehrmann. Structured polynomial eigenvalue problems: Good vibrations from good linearizations. SIAM Journal on Matrix Analysis and Applications, 28(4):1029–1051, 2006.

[2] D. Steven Mackey, Niloufer Mackey, Christian Mehl, and Volker Mehrmann. Numerical methods for palindromic eigenvalue problems: Computing the anti-triangular Schur form. Technical Report 409, DFG Research Center Matheon, Mathematics for key technologies in Berlin, TU Berlin, Str. des 17. Juni 136, D-10623 Berlin, Germany, 2007. url: http://www.matheon.de.

[3] Christian Schröder. A QR-like algorithm for the palindromic eigenvalue problem. Preprint 388, TU Berlin, Matheon, Germany, 2007.

[4] Christian Schröder. URV decomposition based structured methods for palindromic and even eigenvalue problems. Preprint 375, TU Berlin, Matheon, Germany, March 2007.

The World of (HSL) Sparse Direct Solvers

Jennifer A. Scott

The robustness and generality of sparse direct methods makes them popular for solving the linear systems that arise in a wide variety of applications in computational science and engineering. Software that implements such methods has been available since the 1970s and has been continually evolving ever since. In recent years, problem sizes have increased rapidly while new computer architectures have presented direct solvers with new challenges; new algorithms have been proposed and new solvers have been designed and developed. The aim of this talk is to present an overview of the current world of direct solvers [1], looking in particular at recent developments within the mathematical software library HSL [3]. These include efficient out-of-core solvers [4, 5] and, most recently, a mixed precision sparse symmetric solver [2]. Motivated by the advantages of using single precision on modern architectures, this new solver computes a matrix factorization using the single precision version of a direct solver and then seeks to regain higher precision using a simple iterative method.

References

[1] N.I.M. Gould, Y. Hu and J.A. Scott. A numerical evaluation of sparse direct solvers for the solution of large, sparse, symmetric linear systems of equations *ACM Transactions on Mathematical Software*, 33, Article 10, 32 pages.

[2] J.D. Hogg and J.A. Scott. On the use of mixed precision for the fast and robust solution of sparse symmetric linear systems on a multicore machine. Rutherford Appleton Laboratory Technical Report, to appear.

[3] HSL. A collection of Fortran codes for large-scale scientific computation. See http://www.cse.scitech.ac.uk/nag/hsl/

[4] J.K. Reid and J.A. Scott. An out-of-core sparse Cholesky solver. Rutherford Appleton Laboratory Technical Report RAL-TR-2006-013. Submitted to *ACM Transactions on Mathematical Software*.

[5] J.K. Reid and J.A. Scott. An efficient out-of-core multifrontal solver for large-scale unsymmetric element problems. Rutherford Appleton Laboratory Technical Report RAL-TR-2007-014. To appear in *Inter. Journal on Numerical Methods in Engineering*.

Google PageRanking problem: the model and the analysis

Stefano Serra-Capizzano

Let A be a given n-by-n complex matrix with eigenvalues $\lambda, \lambda_2, \ldots, \lambda_n$. Suppose there are nonzero vectors $x, y \in \mathbb{C}^n$ such that $Ax = \lambda x$, $y^*A = \lambda y^*$, and $y^*x = 1$. Let $v \in \mathbb{C}^n$ be such that $v^*x = 1$, let $c \in \mathbb{C}$, and assume that $\lambda \neq c\lambda_j$ for each $j = 2, \ldots, n$. Define $A(c) := cA + (1-c)\lambda x v^*$. The eigenvalues of $A(c)$ are $\lambda, c\lambda_2, \ldots, c\lambda_n$. Every left eigenvector of $A(c)$ corresponding to λ is a scalar multiple of $y - z(c)$, in which the vector $z(c)$ is an explicit rational function of c. If a standard form such as the Jordan canonical form or the Schur triangular form is known for A, we show how to obtain the corresponding standard form of $A(c)$.

The web hyper-link matrix $G(c)$ used by Google for computing the PageRank is a special case in which A is real, nonnegative, and row stochastic (taking into consideration the dangling nodes), $c \in (0,1)$, x is the vector of all ones, and v is a positive probability vector. The PageRank vector (the normalized dominant left eigenvector of $G(c)$) is therefore an explicit rational function of c. Extrapolation procedures on the complex field may give a practical and efficient way to compute the PageRank vector when c is close to 1.

A discussion on the model, on its adherence to reality, and on possible variations is also considered.

Key words. Google matrix, rank-one perturbation, Jordan canonical form, extrapolation formulae.

AMS subject classifications. 65F10, 65F15, 65Y20, 15A18, 15A21, 15A51.

References

[1] C. Brezinski, M. Redivo Zaglia, *Extrapolation Methods. Theory and Practice*, North-Holland, Amsterdam, 1991.

[2] R. Horn, C. Johnson, *Matrix Analysis*, Second Edition, Cambridge University Press, New York, to appear 2006.

[3] R. Horn, S. Serra-Capizzano, *Standard anonical forms for certain rank one perturbations and an application to the (complex) Google PageRanking problem*, Internet Math., to appear.

[4] S. Serra-Capizzano, *Jordan canonical form of the Google matrix: a potential contribution to the PageRank computation*, SIAM J. Matrix Anal. Appl., 27 (2005) pp. 305–312.

Approximating functions of large matrices: computational aspects and applications

Valeria Simoncini

Given a square matrix A, the problem of numerically approximating the action of the matrix function $f(A)$ on a vector v is of great importance in a wide range of applications. Over the years, several methods have been proposed for approximating the action of matrix functions such as the exponential, rational, sign and trigonometric functions. More recently, great attention has been paid to the computational aspects of this problem when large dimensions occur, due to the possible advantages of using matrix function evaluations over standard discretization techniques. In fact, the use of matrix functions is the core of many exponential integrators for solving systems of ordinary differential equations or time-dependent partial differential equations.

When A has large dimension, the computation of $f(A)v$ may be effectively approximated by projecting the problem onto a subspace of possibly much smaller dimension. The Krylov subspace

$$K_k(A, v) = \text{span}\{v, Av, \ldots, A^{k-1}v\}$$

has been extensively used to this purpose, due to its favorable computational and approximation properties.

In this talk, we discuss several aspects associated with the efficient and accurate computation of $f(A)v$ by projection methods, such as stopping criteria and acceleration procedures, when A is large and symmetric positive definite. If time allows, we also report on new results in application problems where the use of matrix function evaluations is particularly effective.

Block Classical Gram–Schmidt Algorithm with Reorthogonalization

Jesse L. Barlow and Alicja Smoktunowicz

An error analysis result is given for a block classical Gram–Schmidt factorization of a full rank matrix A into $A = QR$ where Q is left orthogonal (has orthonormal columns) and R is upper triangular.

Classical Gram-Schmidt (CGS) with reorthogonalization has the potential to become a good algorithm for computing a QR factorization in modern computing environments, even though, without reorthogonalization, it has very poor error analysis properties (cf. [1], [3]).

Since CGS is based upon BLAS-2 operations and block CGS is based upon BLAS-3 operations, it leads to better performance in memory hierarchy based architectures and parallel architectures than does Householder or modified Gram-Schmidt (MGS) algorithms. The numerical properties of the block classical Gram-Schmidt algorithm are not precisely known. A goal of this talk is to establish the error analysis properties of the block Gram–Schmidt algorithm.

We prove that the proposed Reorthogonalized Block Classical Gram-Schmidt is numerically stable under some reasonable assumptions. In floating point arithmetic with machine unit ε_M, we show that if R satisfies certain conditions, then the computed Q and R satisfy

$$\|A - QR\|_F \leq \varepsilon_M f_1(m,n)\|A\|_F + O(\varepsilon_M^2), \quad (1)$$

$$\|I_n - Q^T Q\|_F \leq \varepsilon_M f_2(m,n) + O(\varepsilon_M^2) \quad (2)$$

for modestly growing functions $f_1(\cdot)$ and $f_2(\cdot)$.

At the core of our block CGS method is a routine **local_qr**, where for a matrix B where $p \leq n \leq m$ produces

$$[Q_B, R_B] = \textbf{local_qr}(B)$$

where

$$\begin{aligned} \|B - Q_B R_B\|_F &\leq \varepsilon_M f_3(m,p)\,\|B\|_F + O(\varepsilon_M^2), & (3)\\ \|I_p - Q_B^T Q_B\|_F &\leq \varepsilon_M f_4(m,p) + O(\varepsilon_M^2). & (4) \end{aligned}$$

The routine **local_qr** may be produced using Householder or Givens Q–R factorization.

We propose modification of our algorithm that allows for the used of modified Gram–Schmidt (MGS) instead of **local_qr**. We perform numerical tests on our algorithm to confirm our numerical analysis.

The Reorthogonalized Block Classical Gram-Schmidt algorithm that we develop is a generalization of the CGS2 algorithm of Giraud et al. [2]. A similar algorithm is presented by Stewart [4], but no rank or condition assumptions are made about A or R.

References

[1] J. L. Barlow, A. Smoktunowicz, and H. Erbay, Improved Gram-Schmidt downdating methods, BIT 45 (2005) 259– 285.

[2] L. Giraud, Langou, M. Rozložnik, and J. Van Den Eshof, Rounding error analysis of the classical Gram-Schmidt orthogonalization process, Numer.Math. 101 (2005) 87–100.

[3] A. Smoktunowicz, J. L. Barlow, and J. Langou, A note on the error analysis of Classical Gram-Schmidt, Numer.Math. 105 (2) (2006) 299–313.

[4] G. W. Stewart, Block Gram-Schmidt orthogonalization, Technical Report TR-4823, Department of Computer Science, The University of Maryland, College Park, MD, August 2006.

Computation and Application of Balanced Model Order Reduction

Danny C. Sorensen

Model reduction seeks to replace a large-scale system of differential or difference equations by a system of substantially lower dimension that has nearly the same response characteristics. Balanced truncation is a dimension reduction technique which provides a means to approximate a linear time invariant dynamical system with a reduced order system that has excellent approximation properties. These include a global a priori error bound and preservation of asymptotic stability. To construct the transformations required to obtain balanced reduction, large scale Lyapunov equations must be approximately solved in low rank factored form. This intensive computational task has recently been made tractable through algorithmic advances.

We have developed new methods for solving Lyapunov equations that are capable of solving problems on the order of 1 million state variables on a high end workstation.

These techniques have been extended to a class of descriptor systems which includes the semidiscrete Oseen equations with time independent advection. This is accomplished by eliminating the algebraic equation using a projection. However, we have been able to develop a solution scheme that does not require explicit construction and application of this projector.

We have also applied balanced truncation to a compartmental neuron model to construct a low dimensional, morphologically accurate, quasi-active integrate and fire model. The complex model of dimension 6000 is faithfully approximated with a low order model with 10 variables. More precisely, 10 variables suffice to reliably capture the soma voltage for a number of true morphologies over a broad (in space and time) class of synaptic input patterns. This savings will permit, for the first time, one to simulate large networks of biophysically accurate cells over realistic time spans.

Preconditioning the shift-invert transform in large sparse eigenvalue computations

Alastair Spence

This talk is concerned with the efficient preconditioned iterative solution of shifted linear systems arising from inexact inverse iteration applied to large sparse eigenvalue problems. Extensions are made to inverse subspace iteration and the simplified Jacobi-Davidson method.

Inexact inverse iteration is an "inner-outer" iterative method, with the inner iteration being the (preconditioned) iterative solve of a shifted linear system and the outer iteration being the usual inverse iteration step (see, for example, Golub & Ye, BIT, (2000), **40**). First, we prove that a sequence of rank-one changes to a standard preconditioner, which we call "tuning" the preconditioner, can result in significant savings in inner iteration costs. In particular, for the generalised eigenvalue problem with a fixed shift and decreasing tolerance for the inner solves, using the tuned preconditioner produces no increase in the number of inner iterations as the outer iteration proceeds. This is in contrast to using a standard preconditioner when the number of inner iterations increases as the outer iteration proceeds. Numerical results are presented for the computation of a complex eigenvalue arising from a mixed finite element discretisation of the linearised Navier-Stokes equations with a 2-level domain decomposition preconditioner. Second, for the symmetric eigenvalue problem with a Rayleigh quotient shift, we compare the use of a tuned preconditioner with the approach of Simoncini & Eldén (BIT, (2002), **42**) where the right-hand side is altered to improve the efficiency.

We extend the tuning results to the case of inexact inverse subspace iteration. If the subspace is of dimension p, then tuning the preconditioner is achieved by a sequence of rank p changes to the preconditioner. Again tuning the preconditioner is shown, both in theory and practice, to provide significant computational gains over the un-tuned preconditioner.

Finally, we provide an equivalence result for inexact inverse iteration and simplified Jacobi-Davidson when preconditioned iterative solves are applied to a Galerkin-Krylov inner solver. This extends a result of Simoncini & Eldén.

The work on inexact inverse iteration and the equivalence with Jacobi-Davidson is joint work with Melina Freitag (Bath) and is to appear in IMAJNA and LAA. The work on inexact

subspace iteration is joint with Mickaël Robbé and Miloud Sadkane (Brest) and is submitted to SIMAX.

An unrestarted, one-pass Lanczos with on-the-fly updating of eigenvectors

Andreas Stathopoulos

In recent years, large strides have been made toward the solution of the large, sparse, Hermitian eigenvalue problem. In particular, finding a small number (less than 5) of eigenvalues is well understood theoretically and there exist subspace iterative methods and software that compute these eigenvalues and their eigenvectors with nearly optimal convergence. By optimal convergence we mean the convergence of a corresponding method that uses unlimited memory and does not restart. Yet, some problems in materials science now require the computation of more than 20,000 eigenpairs of matrices of size several million. In lattice QCD, a few hundred eigenvectors are needed for even larger matrices. These applications are pushing the computational frontiers to where no methods are known with both optimal convergence and limited memory.

In this abstract, we discuss a modification of the unrestarted Lanczos method that updates a window of the required Ritz vectors on the fly, thus having the potential to alleviate some of the problems encountered in cases of large number of eigenpairs or interior eigenpairs.

In the case of one, extreme eigenpair, near optimal methods can be developed by approaching the problem as a constrained optimization one. The use of inexact Newton methods leads to a class of inner-outer methods such as JDCG or JDQMR, for which the inner iteration can be stopped at optimal points, leading to convergence within a small factor (usually less than 2) from optimal. Alternatively, quasi-Newton methods accelerate the locally optimal Conjugate Gradient method (a variant of the non-linear Conjugate Gradient), leading to the subspace accelerated GD+k which has demonstrated convergence almost indistinguishable from optimal.

Single-vector Newton methods, however, target a large number of eigenpairs individually, which is clearly suboptimal to the global convergence of the unrestarted Lanczos. Block Newton methods can be used instead (block GD+k, LOBPCG, trace-minimization, etc.) but it is unclear why k polynomials of degree m should be superior to a single polynomial of degree km. Practitioners continue to observe that when a sufficient large number of eigenvalues is needed, various forms of unrestarted Lanczos are the best choice.

The optimality of unrestarted Lanczos comes with several disadvantages. First, the size of the projected tridiagonal matrix may become too large. Second, without orthogonalization (Cullum and Willoughby version) ghost eigenvalues must be managed and, most importantly, the iteration number and thus the size of the tridiagonal matrix may grow inordinately large. Third, selective or partial orthogonalization could significantly increase the cost. Finally, Lanczos vectors must be stored or recomputed during a second pass.

The idea in this research is to use the optimal restarting strategies from GD+k to update a window of Ritz vectors but without restarting the Lanczos process. Consider the Lanczos method on matrix A, and a set of $m > nev$ vectors, V, that keep track of the required nev Ritz vectors and the most recent Lanczos vectors. Initially, V is composed of the first m Lanczos vectors with the usual tridiagonal projection matrix $T_m = V^H AV$. At this point, we

apply the Rayleigh-Ritz on T_m and restart V with the current nev Ritz vectors $u_i^{(m)}$. The critical feature of our algorithm is that we also include in V the nev Ritz vectors from step $m-1$. Thus, the restarted V is computed as an orthonormal set of vectors spanning:

$$[u_1^{(m)}, u_2^{(m)}, \ldots, u_{nev}^{(m)}, u_1^{(m-1)}, \ldots, u_k^{(m-1)}]. \quad (1)$$

From this point forth, the algorithm deviates from thick restarted Lanczos or GD+k because it does not use the residual of $u_1^{(m)}$ to build a new Lanczos iteration. It keeps using the next $(m-2nev)$ Lanczos vectors, v_i, from the original, unrestarted process to append to V:

$$[u_1^{(m)}, u_2^{(m)}, \ldots, u_{nev}^{(m)}, u_1^{(m-1)}, \ldots, u_k^{(m-1)}, v_{m+1}, \ldots, v_{m+m-2nev}].$$

Note that orthogonality need not be maintained as the Lanczos vectors v_i are orthogonal (in exact arithmetic) to all previous Lanczos vectors and therefore to any Ritz vectors in V. The projection matrix for the resulting V is $H = V^H A V = \begin{bmatrix} \Theta & C^H \\ C & T \end{bmatrix}$, where Θ is a diagonal matrix of size $2nev$ containing the nev Ritz values, T is the section from $m+1$ to $2m-2nev$ of the tridiagonal matrix of the unrestarted Lanczos, and only the first row of C is non-zero. Thus, H can be computed from the Lanczos coefficients with negligible additional cost. At this point, the Rayleigh-Ritz is applied again and V is restarted with the nev Ritz vectors from H and the previous nev Ritz vectors from $H(1:m-1, 1:m-1)$. Then the update algorithm repeats, while the original Lanczos continues.

In theory, the use of Rayleigh Ritz and thick restart guarantee monotonic convergence of the nev Ritz values in V, but not necessarily to any eigenvalues. To our surprise, in all of our experiments with $nev > 4$, the nev updated vectors converge to the required eigenpairs at a rate identical to that of unrestarted Lanczos! Note that any information that is discarded from V when we restart cannot be reintroduced in V because future Lanczos vectors v_i are orthogonal to it. By restarting through eq. (1) almost all Lanczos information regarding the nev eigenpairs is kept in condensed form in V. In fact, with thick restarting alone (i.e., without the $u_i^{(m-1)}$ directions) the algorithm stagnates.

In our experiments, nev values as small as 4 or 5 are adequate for convergence of the updated Ritz pairs to almost machine precision and at the unrestarted Lanczos rate. The value of m seems to be much less important for convergence, but larger values of m reduce computational costs. In particular, for $m \gg nev$, the computational cost of the algorithm is similar to that of unrestarted Lanczos of the same number of steps, but without the solution of a large tridiagonal system and without the need to store the Lanczos vectors v_i or to recompute them in a second pass. Instead, we have to solve a series of $m \times m$ matrices H and store m vectors.

A basic version of this algorithm was first used on top of the CG algorithm, not Lanczos, to approximate eigenvectors while solving a linear system of equations. These eigenvectors were later used to deflate subsequent right hand sides, with dramatic improvement of the condition number. In this talk, we explore algorithmic issues and applicability when the computation of the eigenvalues is the eventual goal.

Loss of orthogonality in v_i will eventually corrupt the orthogonality of V, preventing further convergence of the Ritz vectors. This occurs when at least one eigenvector has started to converge, and therefore places a limit on how large nev can be used in a meaningful way. However, the algorithm lends itself naturally to selective orthogonalization of v_i as all the

required Ritz vectors are available at no additional cost. Ghost eigenpairs that are not part of V do not seem to affect convergence. Besides the choice of m and nev, another important issue is deciding when to stop the algorithm. The usual Lanczos estimates of the residual norms are freely available from our algorithm, but loss of orthogonality must also be estimated. Finally, even when the ultimate achievable convergence is limited (by orthogonality or by poor choice of m, nev) our method can be wrapped inside an outer method like Davidson.

Analysis of the Residual Arnoldi Method

Che-Rung Lee, G. W. Stewart

Let A be a matrix of order n. The classical Arnoldi method produces a nested sequence of orthonormal bases

$$U_k = (u_1 \ u_2 \ \cdots \ u_k)$$

by orthonormalizing Au_k against the columns of U_k. The resulting bases satisfy the **Arnoldi relation**

$$AU_k = U_k\hat{H}_k + \rho_k u_{k+1}\mathbf{e}_k^*,$$

where $\hat{H}_k$ is upper Hessenberg.

The spaces spanned by the U_k frequently contain increasingly good approximations to eigenpairs corresponding to eigenvalues on the periphery of the spectrum of A. These approximations can usually be recovered by the following Rayleigh–Ritz procedure, which computes an approximation (μ, z) to a **target** eigenpair (λ, x) of A.

1. Compute the **Rayleigh quotient** $H_k = U_k^* A U_k$.
2. Of the eigenpairs of (μ_i, w_i) of H_k choose a **candidate** pair, say (μ, w), such that μ approximates the target eigenvalue λ.
3. Let $(\mu, z) = (\mu, U_k w)$.

In many cases the candidates converge to their respective targets (although the convergence theory is weak and often difficult to apply). However, if we introduce relative errors of size ϵ into the u_k the convergence stagnates at an accuracy proportional to ϵ. (Even a single error in u_2 is sufficient to cause this behavior). The **residual Arnoldi method** chooses a target eigenpair and expands U_k by orthonormalizing the candidate residual $r_k = Az_k - \mu z_k$ against the columns of U_k.

This procedure is equivalent to the usual Arnoldi procedure when there is no error. When error is present and ϵ is sufficiently small, the candidates for the specified target continue to converge at essentially the same rate as in the error free method. However, the candidates for the other targets stagnate.

The analysis of the residual Arnoldi method is long and detailed. But two aspects are worth pointing out.

First, we are able to show that the U_k computed from the process with error satisfy a perturbed Arnoldi relation of the form

$$(A + E_k)U_k = U_k H_k + \rho_k u_{k+1}\mathbf{e}_k^*.$$

Unfortunately, the perturbations E_k do not get smaller with increasing k. But assuming the norms of the E_k are bounded, we are able to show that for the target eigenvector x the products $E_k x$ do get smaller. This is a key lemma in our analysis.

Second, we finesse the problem of proving the proving convergence of the Ritz pairs by assuming that we are dealing with an error free Arnoldi process that already converges. Specifically, we suppose that without error the candidate Ritz vectors $z_k^{(0)}$ satisfy

$$\|z_k^{(0)} - x\| \leq \kappa_k,$$

where the κ_k decrease. We are then able to show that the Ritz vectors $z_k^{(\epsilon)}$ generated with error satisfy

$$\|z_k^{(\epsilon)} - x\| \leq \tilde{\kappa}_k,$$

where the relative error in $\tilde{\kappa}_k$ as an approximation to κ_k is $O(\epsilon)$. This says that we will see essentially the same convergence curves when ϵ is small.

Finally, we note that there is a residual shift/invert Arnoldi algorithm. The payoff with this algorithm can be quite large, since one can solve the linear systems in the invert phase to low relative error. This shift/invert residual Arnoldi method has an analogous convergence result.

Sensitivity of Gauss-Christoffel quadrature, and sensitivity of the Jacobi matrix to small perturbations of the spectral data

Zdeněk Strakoš, Dianne P. O'Leary, Petr Tichý

In numerical computations the question how much does a function change under perturbations of its arguments is of central importance.

Consider a problem of sensitivity and numerically stable computation of the entries of the Jacobi matrix from its eigenvalues and the first elements of its normalized eigenvectors [3, 14, 2]. This represents the classical inverse problem in the discrete Gauss-Christoffel quadrature calculations [11, 8, 5, 6, 7, 1, 9, 10], and its perturbation analysis is equivalent to the analysis of sensitivity of the Lanczos recurrences to the spectrum and components of the starting vector in the individual invariant subspaces [21, 20, 18, 4].

With connection to orthogonal polynomials and the Gauss-Christoffel quadrature, the problem has a long and very interesting history; for a brief review of the literature see [19]. It has been investigated in many independent lines of development, often without considering the relationship to the results developed in a different field or using a different language. This can be illustrated by the fact, that the fundamental contribution by Dirk Laurie [15, 16], motivated by the earlier work by Gragg and Harrod [12], has practically been unnoticed in the numerical linear algebra literature.

This contribution briefly reviews different lines of development and puts the problem into a broader context of sensitivity of the Gauss-Christoffel quadrature with respect to small perturbations of the distribution function [19], and of numerical stability of the conjugate gradient method [13, 17].

In numerical quadrature, a definite integral is approximated by a finite sum of functional values evaluated at given quadrature nodes and multiplied by given weights. Consider a sufficiently smooth integrated function uncorrelated with the perturbation of the distribution function. Then it seems natural *that given the same number of function evaluations,* the difference between the quadrature approximations is of the same order as the difference between the (original and perturbed) approximated integrals. That is perhaps one of the reasons why, to our knowledge, the sensitivity question has not been formulated and addressed in the literature, though, as mentioned above, other sensitivity problems, motivated, in particular, by computation of the quadrature nodes and weights from moments, have been thoroughly studied by many authors. We show that *even a small perturbation of a distribution function can cause large differences in Gauss-Christoffel quadrature estimates.*

This sensitivity in Gauss-Christoffel quadrature can be observed for discontinuous, continuous, and even analytic distribution functions, and for analytic integrands uncorrelated with changes in the distribution functions and with no singularity close to the interval of integration.

The sensitivity of the Gauss-Christoffel quadrature illustrated in the contribution is related to the difference in the *size* of the support of the original and of the perturbed distribution functions. For a discrete distribution function, the size is the number of points of increase, and for a continuous distribution function it is the length (measure) of the union of intervals containing points at which the distribution function increases. In general, different supports of the *same size* do not exhibit sensitivity in quadrature results.

The sensitivity of Gauss-Christoffel quadrature cannot be explained using existing analysis based on modified moments. In our examples, if the support of the original distribution function differs in size from the support of the auxiliary (perturbed) distribution function, then the matrices of both modified and mixed moments become highly ill-conditioned. The same is true if the supports are different but of the same size. But only in the case of different size of the supports are the recurrence coefficients (i.e., the entries of the Jacobi matrix) and the Gauss-Christoffel quadrature estimates highly sensitive to the perturbation.

References

[1] B. Beckermann and E. Bourreau, *How to choose modified moments?*, J. Comput. Appl. Math., 98 (1998), pp. 81–98.

[2] D. Boley and G. H. Golub, *A survey of matrix inverse eigenvalue problems*, Inverse Problems, 3 (1987), pp. 595–622.

[3] C. de Boor and G. H. Golub, *The numerically stable reconstruction of a Jacobi matrix from spectral data*, Linear Algebra Appl., 21 (1978), pp. 245–260.

[4] V. Druskin, L. Borcea, and L. Knizhnerman, *On the sensitivity of Lanczos recursions to the spectrum*, Linear Algebra Appl., 396 (2005), pp. 103–125.

[5] S. Elhay, G. H. Golub, and J. Kautsky, *Updating and downdating of orthogonal polynomials with data fitting applications*, SIAM J. Matrix Anal. Appl., 12 (1991), pp. 327–353.

[6] ——, *Jacobi matrices for sums of weight functions*, BIT, 32 (1992), pp. 143–166.

[7] H.-J. Fischer, *On generating orthogonal polynomials for discrete measures*, Z. Anal. Anwendungen, 17 (1998), pp. 183–205.

[8] W. Gautschi, *A survey of Gauss-Christoffel quadrature formulae*, in E. B. Christoffel: The Influence of His Work on Mathematics and Physics, P. L. Butzer and F. Fehér, eds., Birkhäuser, Basel, 1981, pp. 72–147.

[9] ——, *Orthogonal polynomials: computation and approximation*, Numerical Mathematics and Scientific Computation, Oxford University Press, New York, 2004.

[10] ——, *Orthogonal polynomials (in Matlab)*, J. Comput. Appl. Math., 178 (2005), pp. 215–234.

[11] G. H. Golub and J. H. Welsch, *Calculation of Gauss quadrature rules*, Math. Comp. 23 (1969), 221-230; addendum, ibid., 23 (1969), pp. A1–A10.

[12] W. B. Gragg and W. J. Harrod, *The numerically stable reconstruction of Jacobi matrices from spectral data*, Numer. Math., 44 (1984), pp. 317–335.

[13] W. Greenbaum, *Behavior of slightly perturbed Lanczos and conjugate gradient recurrences* , Linear Algebra Appl., 113 (1989), pp. 7–63.

[14] J. Kautský and G. H. Golub, *On the calculation of Jacobi matrices*, Linear Algebra Appl., 52/53 (1983), pp. 439–455.

[15] D. P. Laurie, *Accurate recovery of recursion coefficients from Gaussian quadrature formulas*, J. Comput. Appl. Math., 112 (1999), pp. 165–180.

[16] ——, *Computation of Gauss-type quadrature formulas*, J. Comput. Appl. Math., 127 (2001), pp. 201–217.

[17] G. Meurant and Z. Strakoš, *The Lanczos and conjugate gradient algorithms in finite precision arithmetic*, Acta Numerica, 15 (2006), pp. 471–542.

[18] F. Natterer, *Discrete gel'fand-levitan theory, www page of the author*, 1999.

[19] D. P. O'Leary, Z. Strakoš, and P. Tichý, *On sensitivity of Gauss-Christoffel quadrature*, Numerische Mathematik, 107 (2007), pp. 147–174.

[20] C. C. Paige and P. Van Dooren, *Sensitivity analysis of the Lanczos reduction*, Numer. Linear Algebra Appl., 6 (1999), pp. 29–50.

[21] S. F. Xu, *A stability analysis of the Jacobi matrix inverse eigenvalue problem*, BIT, 33 (1993), pp. 695–702.

Compressed Sensing and Random Matrices

Gilbert Strang

Compressed sensing acquires only data to be used, rather than depending on high compression at a later step. It is assumed that the signal is sparse in the right basis. The sensor measures inner products with k "random" vectors and looks for a sparse signal with those inner products. True sparsity would be an optimization in the L^0 norm (NP-hard) but the key to the new theory is that L^1 optimization can give a perfect answer.

The linear algebra problem is to estimate the smallest singular value for k by k submatrices S of the n by k matrix A of random sampling vectors. A could be a random orthogonal matrix from `qr(randn)` or a random $+1/-1$ matrix or random columns of the Fourier (DFT) matrix.

We report on a 1997 conjecture (Goreinov-Tyrtyshnikov in LAA) that every A with k orthonormal columns has a k by k submatrix S with $\|S^{-1}\| \leq \sqrt{n}$. Numerical experiments by Cy Chan show that this best possible constant is seldom approached. But the distribution of $\|S^{-1}\|$ for all submatrices is very different—many S are nearly singular. We show histograms of the distribution of $\sigma_{\min}$ and discuss also the problem of constructing A to maximize $\sigma_{\min}$ for the *worst* submatrix S—as compressed sensing needs.

Computing the Complete CS Decomposition

Brian D. Sutton

We present an algorithm for computing the complete CS decomposition (CSD) of a partitioned unitary matrix. Although the existence of the CSD has been recognized since 1977, existing algorithms compute only a reduced version (the 2-by-1 CSD), which is equivalent to two simultaneous singular value decompositions. Our algorithm computes the complete 2-by-2 CSD, which is equivalent to four simultaneous—and highly interrelated—SVD's. The algorithm appears to be the only fully specified algorithm available. The computation occurs in two phases. In the first phase, the unitary matrix is reduced to bidiagonal block form, as described by Sutton and Edelman. In the second phase, the blocks are simultaneously diagonalized using techniques from bidiagonal SVD algorithms of Golub, Kahan, and Demmel.

References

[1] E. Anderson, Z. Bai, C. Bischof, L. S. Blackford, J. Demmel, Jack J. Dongarra, J. Du Croz, S. Hammarling, A. Greenbaum, A. McKenney, and D. Sorensen. *LAPACK Users' guide (third ed.)*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 1999.

[2] Z. Bai. The CSD, GSVD, their applications and computations. Preprint Series 958, Institute for Mathematics and its Applications, University of Minnesota, April 1992.

[3] Z. Bai and J. Demmel. Computing the generalized singular value decomposition. *SIAM J. Sci. Comp.*, 14(6):1464–1486, 1993.

[4] Chandler Davis and W. M. Kahan. Some new bounds on perturbation of subspaces. *Bull. Amer. Math. Soc.*, 75:863–868, 1969.

[5] Chandler Davis and W. M. Kahan. The rotation of eigenvectors by a perturbation. III. *SIAM J. Numer. Anal.*, 7:1–46, 1970.

[6] James Demmel and W. Kahan. Accurate singular values of bidiagonal matrices. *SIAM J. Sci. Statist. Comput.*, 11(5):873–912, 1990.

[7] Alan Edelman and Brian D. Sutton. The beta-Jacobi matrix model, the CS decomposition, and generalized singular value problems. *Found. Comput. Math.* DOI 10.1007/s10208-006-0215-9. http://www.springerlink.com/content/r46712153613x777.

[8] G. Golub and W. Kahan. Calculating the singular values and pseudo-inverse of a matrix. *J. Soc. Indust. Appl. Math. Ser. B Numer. Anal.*, 2:205–224, 1965.

[9] Gene H. Golub and Charles F. Van Loan. *Matrix computations*. Johns Hopkins Studies in the Mathematical Sciences. Johns Hopkins University Press, Baltimore, MD, third edition, 1996.

[10] V. Hari. Accelerating the SVD block-Jacobi method. *Computing*, 75(1):27–53, 2005.

[11] Camille Jordan. Essai sur la géométrie à n dimensions. *Bull. Soc. Math. France*, 3:103–174, 1875.

[12] C. C. Paige and M. A. Saunders. Towards a generalized singular value decomposition. *SIAM J. Numer. Anal.*, 18(3):398–405, 1981.

[13] C. C. Paige and M. Wei. History and generality of the CS decomposition. *Linear Algebra Appl.*, 208/209:303–326, 1994.

[14] G. W. Stewart. On the perturbation of pseudo-inverses, projections and linear least squares problems. *SIAM Rev.*, 19(4):634–662, 1977.

[15] G. W. Stewart. Computing the CS decomposition of a partitioned orthonormal matrix. *Numer. Math.*, 40(3):297–306, 1982.

[16] Brian D. Sutton. Matlab implementation of an algorithm for computing the complete CSD. `http://faculty.rmc.edu/bsutton/`.

[17] Brian D. Sutton. *The stochastic operator approach to random matrix theory*. PhD thesis, Massachusetts Institute of Technology, Cambridge, MA 02139, June 2005.

[18] Charles Van Loan. Computing the CS and the generalized singular value decompositions. *Numer. Math.*, 46(4):479–491, 1985.

[19] David S. Watkins. Some perspectives on the eigenvalue problem. *SIAM Rev.*, 35(3):430–471, 1993.

[20] J. H. Wilkinson. Global convergence of tridiagonal QR algorithm with origin shifts. *Linear Algebra and Appl.*, 1:409–420, 1968.

Application of Inexact Krylov Subspace Methods to Control Problems Governed by PDEs

Daniel B. Szyld

In many circumstances, a known good preconditioner may not be easily computable. Instead, an approximation to it may be available. This is the case, for example, when the preconditioner has an inverse associated with it, such as in Schur complements (e.g., in saddle point problems), or in the reduced Hessian in some control problems. For parabolic control problems, the reduced Hessian may be of the form $H = N^T E^{-T} Q^T Q E^{-1} N$, and thus it has two inverses as factors.

A natural question is: how accurately does one need to solve the intermediate (or inner) linear systems (with E and E^T) when solving the linear system with H using a Krylov subspace method, such as CG, FOM, or GMRES.

In the last few years, analysis of Inexact Krylov methods have led to computable (inner) stopping criteria that guarantee the convergence of the overall method. It turns out that as the outer iterations progress, the inner linear systems can be solved less and less accurately. By solving these linear systems with a more relaxed tolerance, computational savings can be obtained. In this talk, we first review these results, and them apply them to parabolic control problems, where two (possibly different) inner stopping criteria are needed.

Numerical results using inexact CG and inexact truncated FOM are presented illustrating the applicability of these results to various control problems, and the accompanying reduction in computational effort with little or no degradation of the approximations to the solution.

An Updated Preconditioner for Sequences of General Nonsymmetric Linear Systems

Jurjen Duintjer Tebbens

Efficient solution of *sequences* of linear systems is a task arising in numerous applications in engineering and scientific computing. Some examples are simulation of processes in fluid dynamics where every time step requires the solution of a system of nonlinear equations, operation research problems where the linear programs are solved with the simplex method or solution of Helmholtz equations for the propagation of time-harmonic waves. Depending on the linear solver and the properties of the system matrices, several techniques to share part of the computational effort throughout the sequence may be used. Apart from the very simple idea of using hot starts with the solution of the previous system, in some cases exact updating of the system matrix is feasible even for large problems. Rank-one updates of LU factorizations have been used since decades in the simplex method where the change of one system matrix to another is restricted to one column [17]. General rank-one updates of an LU decomposition are discussed in [16]. If the linear solver is a Krylov subspace method, strategies to recycle information gained from previously generated subspaces have shown to be beneficial in many applications [10, 13]. In the special case where shifted linear systems with identical right hand sides have to be solved, Krylov subspace methods allow advantageous implementations based on the fact that all systems generate the same subspace [8]. In nonlinear systems solved with a Newton-type method a well-known strategy is to skip evaluations of the (approximate) Jacobian during some iterations, leading to Shamanskii's

combination of the chord and Newton method [5], [14]. Several subsequent systems of the sequence then differ only by their right hand side and linear solving techniques with multiple right hand sides can be exploited [15]. Another way to save costs is to allow changing the system matrices but freezing the preconditioner [9]. To enhance the power of a frozen preconditioners one may also use *approximate* updates; in [11] we find approximate updates of incomplete Cholesky factorizations and in [1, 3] banded updates were proposed for both symmetric positive definite approximate inverse and incomplete Cholesky preconditioners. In Quasi-Newton methods the difference between system matrices is of small rank and preconditioners may be efficiently adapted with approximate small-rank updates; this has been done in the symmetric positive definite case, see [2, 12].

Our contribution considers a new black-box approximate update scheme for factorized preconditioners introduced in [6] and extended in [4, 7]. It is designed for general nonsymmetric linear systems solved by arbitrary iterative solvers and hence it can be combined with some of the techniques for more specific systems and solvers mentioned before. The basic idea is to combine an incomplete reference factorization with a Gauss-Seidel type of approximation of the difference between the current and the reference matrix. The technique tends to correct deteriorating numbers of iterations needed to solve with a frozen preconditioner by reducing them to an acceptable level. Moreover, the updated factorizations can be more powerful than preconditioners computed from scratch; this may happen, for instance, when the updates are related to more stable reference preconditioners generated earlier in the sequence. Since the updated factorizations are both cheap to compute and cheaply applied as a preconditioner, their use is of considerable interest for practice. This is especially true when preconditioner computations are expensive, like in matrix-free environment where the matrix has to be estimated first. In this talk we give a brief description of the basis update technique and then we present new theoretical results on the quality of the updates. Apart from providing new insight they may serve as guidelines in deciding whether to switch-on the usage of updates in black-box solvers for sequences of linear systems. We also present numerical experiments with a new implementation strategy for matrix-free environment which is based on mixed matrix-free/explicit triangular solves.

This is joint work with Miroslav Tůma (Institute of Computer Science, Academy of Sciences of the Czech Republic). The work is supported by the National Program of Research "Information Society" under project 1ET400300415.

References

[1] M. Benzi and D. Bertaccini, *Approximate Inverse Preconditioning for Shifted Linear Systems*, BIT, 43 (2003), pp. 231–244.

[2] L. Bergamaschi, R. Bru, A. Martínez and M. Putti, *Quasi-Newton Preconditioners for the Inexact Newton Method*, ETNA, 23 (2006), pp. 76–87.

[3] D. Bertaccini, *Efficient Preconditioning for Sequences of Parametric Complex Symmetric Linear Systems*, ETNA, 18 (2004), pp. 49–64.

[4] Ph. Birken, J. Duintjer Tebbens, A. Meister and M. Tůma, *Preconditioner Updates Applied to CFD Model Problems*, to appear in Applied Numerical Mathematics.

[5] R. P. Brent, *Some Efficient Algorithms for Solving Systems of Nonlinear Equations*, SIMUM, 10 (1973), pp. 327–344.

[6] J. Duintjer Tebbens and M. Tůma, *Efficient Preconditioning of Sequences of Nonsymmetric Linear Systems*, SIAM J. Sci. Comput., 29 (2007), pp. 1918–1941.

[7] J. Duintjer Tebbens and M. Tůma, *Improving Triangular Preconditioner Updates for Nonsymmetric Linear Systems*, to appear in LNCS 4818 (proceedings of the 6th International Conference on Large-Scale Scientific Computations), Springer.

[8] A. Frommer and U. Glässner, *Restarted GMRES for Shifted Linear Systems*, SIAM J. Sci. Comput., 19 (1998), pp. 15–26.

[9] D.A. Knoll and D.E. Keyes, *Jacobian-free Newton-Krylov Methods: A Survey of Approaches and Applications*, J. Comput. Physics, 193 (2004), pp. 357–397.

[10] D. Loghin, D. Ruiz and A. Touhami, *Adaptive Preconditioners for Nonlinear Systems of Equations*, J. Comput. Appl. Math., 189 (2006), pp. 326–374.

[11] G. Meurant, *On the Incomplete Cholesky Decomposition of a Class of Perturbed Matrices*, SIAM J. Sci. Comput., 23 (2001), pp. 419–429.

[12] J. Morales and J. Nocedal, *Automatic Preconditioning by Limited-Memory Quasi-Newton Updates*, SIAM J. Opt., 10 (2000), pp. 1079–1096.

[13] M. L. Parks, E. de Sturler, G. Mackey, D. D. Johnson and S. Maiti, *Recycling Krylov Subspaces for Sequences of Linear Systems*, SIAM J. Sci. Comput., 28 (2006), pp. 1651–1674.

[14] V. E. Shamanskii, *A Modification of Newton's Method*, Ukrain. Mat. Z., 19 (1967), pp. 1333–1338.

[15] V. Simoncini and E. Gallopoulos, *An Iterative Method for Nonsymmetric Systems with Multiple Right Hand Sides*, SIAM J. Sci. Comput., 16 (1995), pp. 917–933.

[16] P. Stange, A. Griewank and M. Bollhöfer, *On the Efficient Update of Rectangular LU Factorizations Subject to Low Rank Modifications*, Preprint 2005/27, TU Berlin.

[17] L. M. Suhl and U. H. Suhl, *A Fast LU Update for Linear Programming*, Ann. Oper. Res., 43 (1993), pp. 33–47.

On a New Proof of the Faber-Manteuffel Theorem

Vance Faber, Jörg Liesen, Petr Tichý

A fundamental theorem in the area of iterative methods is the Faber-Manteuffel Theorem [2]. It shows that a short recurrence for orthogonalizing Krylov subspace bases for a matrix A exists if and only if the adjoint of A is a low degree polynomial in A. This result is important, since it characterizes all matrices, for which an optimal Krylov subspace method with short recurrences can be constructed. Here optimal means that the error is minimized in the norm induced by the given inner product. Of course, such methods are highly desirable, due to convenient work and storage requirements for generating the orthogonal basis vectors. Examples are the CG method [3] for solving systems of linear algebraic equations with a symmetric positive definite matrix A, or the MINRES method [6] for solving symmetric but indefinite systems.

Now we briefly describe the result of Faber and Manteuffel. Let A be a nonsingular matrix and v be a vector of grade d (d is the degree of the uniquely determined monic polynomial of smallest degree that annihilates v). For theoretical as well as practical purposes it is often convenient to orthogonalize the basis $v, \ldots, A^{d-1}v$ of the cyclic subspace $\mathcal{K}_d(A, v)$. The classical approach to orthogonalization is to use the Arnoldi method, that produces mutually orthogonal vectors $v_1, \ldots, v_d$ satisfying $\mathrm{span}\{v_1, \ldots, v_n\} = \mathrm{span}\{v, \ldots, A^{n-1}v\}$, $n = 1, \ldots, d$. The algorithm can be written in a matrix form

$$v_1 = v\,, \tag{1}$$

$$A \underbrace{[v_1, \ldots, v_{d-1}]}_{\equiv V_{d-1}} = \underbrace{[v_1, \ldots, v_d]}_{\equiv V_d} \underbrace{\begin{bmatrix} h_{1,1} & \cdots & h_{1,d-1} \\ 1 & \ddots & \vdots \\ & \ddots & h_{d-1,d-1} \\ & & 1 \end{bmatrix}}_{\equiv H_{d,d-1}}, \tag{2}$$

$$(v_i, v_j) = 0 \quad \text{for } i \neq j\,, \quad i, j = 1, \ldots, d\,. \tag{3}$$

As described above, for efficiency reasons, it is desirable to generate such an orthogonal basis with a short recurrence, meaning that in each iteration step only a few of the latest basis vectors are required to generate the new basis vector. This corresponds to the situation when the matrix $H_{d,d-1}$ in (2) is, for each starting vector v_1, low-band Hessenberg matrix. Note that an unreduced upper Hessenberg matrix is called $(s+2)$-band Hessenberg, when its s-th superdiagonal contains at least one nonzero entry, and all its entries above its s-th superdiagonal are zero. We say that *A admits an optimal $(s+2)$-term recurrence* if $H_{d,d-1}$ is for each starting vector at most $(s+2)$-band Hessenberg and, moreover, there exists an initial vector such that $H_{d,d-1}$ is exactly $(s+2)$-band Hessenberg (the s-th superdiagonal contains at least one nonzero entry). The fundamental question is, what properties are necessary and sufficient for A to admit an optimal $(s+2)$-term recurrence. This question was answered by Faber and Manteuffel [2].

THEOREM (Faber-Manteuffel) *Let A be a nonsingular matrix with minimal polynomial degree $d_{\min}(A)$. Let s be a nonnegative integer, $s + 2 < d_{\min}(A)$. Then A admits an optimal $(s+2)$-term recurrence if and only if $A^* = p(A)$, where p is a polynomial of smallest degree s having this property (i.e. A is normal(s)).*

While the sufficiency of the normal(s) condition is rather easy to prove, the proof of necessity given by Faber and Manteuffel is based on a clever, highly nontrivial construction by using results from mathematical analysis ("continuous function"), topology ("closed set of smaller dimension") or multilinear algebra ("wedge product").

In [5], Liesen and Strakoš discuss and clarify the existing important results in the context of the Faber-Manteuffel Theorem. They suggest that, in light of the fundamental nature of the result, it is desirable to find an alternative, and possibly simpler proof of the necessity part.

In our recent paper [1] we address this issue and give two new proofs of the necessity part, which use more elementary tools. In particular, we avoid topological and analytical arguments, and use linear algebra tools instead. In this talk we will present the proof which

employs elementary Givens rotations to prove the necessity part by contradiction. In particular, the proof is based on orthogonal transformations ("rotations") of upper Hessenberg matrices. We will also discuss why the proof of the necessity part is much more difficult than the proof of sufficiency.

References

[1] V. Faber, J. Liesen and P. Tichý, *The Faber-Manteuffel Theorem for linear operators*, SIAM J. Numer. Anal., accepted (2007).

[2] V. Faber and T. Manteuffel, *Necessary and sufficient conditions for the existence of a conjugate gradient method*, SIAM J. Numer. Anal., 21 (1984), pp. 352–362.

[3] M. R. Hestenes and E. Stiefel, *Methods of conjugate gradients for solving linear systems*, J. Research Nat. Bur. Standards, 49 (1952), pp. 409–436.

[4] J. Liesen and P. E. Saylor, *Orthogonal Hessenberg reduction and orthogonal Krylov subspace bases*, SIAM J. Numer. Anal., 42 (2005), pp. 2148–2158.

[5] J. Liesen and Z. Strakoš, *On optimal short recurrences for generating orthogonal Krylov subspace bases*, SIAM Rev., (to appear).

[6] C. C. Paige and M. A. Saunders, *Solution of sparse indefinite systems of linear equations*, SIAM J. Numer. Anal., 12 (1975), pp. 617–629.

How to Detect and Solve Hyperbolic Quadratic Eigenvalue Problems

Chun-Hua Guo, Nicholas J. Higham, Françoise Tisseur

Two recent research themes in numerical linear algebra are

- the use of cyclic reduction for solving matrix Riccati equations and related nonlinear equations,
- the identification and exploitation of structure in solving polynomial eigenvalue problems.

This work combines both themes: it shows how cyclic reduction can be used to test whether a given quadratic matrix polynomial is hyperbolic (or overdamped) and then, in the case of a positive test, to initiate a solution procedure that fully exploits the hyperbolicity. In particular, we give the first practical procedure for linearizing a hyperbolic quadratic into a definite pencil $\lambda X + Y$ for which X or Y is a definite matrix.

A quadratic matrix polynomial

$$Q(\lambda) = \lambda^2 A + \lambda B + C, \quad A, B, C \in \mathbb{C}^{n\times n} \tag{1}$$

is hyperbolic if A, B, and C are Hermitian, $A > 0$, and

$$(x^*Bx)^2 > 4(x^*Ax)(x^*Cx) \quad \text{for all nonzero } x \in \mathbb{C}^n. \tag{2}$$

It is overdamped if it is hyperbolic with $B > 0$ and $C \geq 0$. The eigenvalues of a hyperbolic Q are real and those of an overdamped Q are real and nonpositive.

Overdamped quadratics arise in overdamped systems in structural mechanics [5, Sec. 7.6] and also are closely related to gyroscopic systems. A hyperbolic Q can be shifted to make it overdamped, which enables testing for hyperbolicity to be reduced to testing for overdamping of a suitably shifted Q.

The basis of our test for overdamping is the matrix iteration, based on cyclic reduction,

$$\begin{aligned} &S_0 = B, \quad A_0 = A, \quad B_0 = B, \quad C_0 = C, \\ &S_{k+1} = S_k - A_k B_k^{-1} C_k, \\ &A_{k+1} = A_k B_k^{-1} A_k, \\ &B_{k+1} = B_k - A_k B_k^{-1} C_k - C_k B_k^{-1} A_k, \\ &C_{k+1} = C_k B_k^{-1} C_k. \end{aligned} \tag{3}$$

We show that for a version of this quadratically convergent iteration augmented with a certain scaling parameter α_k, Q is overdamped if and only if for some $m \geq 0$, $B_k > 0$ for $k = 1\colon m-1$ and $Q(-\mu_m) < 0$, where μ_m depends on the α_k [1].

This test is attractive for two main reasons.

1. The definiteness test, which is carried out in practice by attempting a Cholesky factorization, is on the original Q, which means that a positive test can be interpreted irrespective of rounding errors in the iteration.

2. A μ_m resulting from a positive test enables us to construct a $2n \times 2n$ linearization $\lambda X + Y$ with a positive definite X or Y; this generalized eigenproblem can then be solved by using a Cholesky factorization of the definite matrix to reduce the problem to a standard Hermitian eigenproblem, thereby guaranteeing real computed eigenvalues.

How does our approach compare with earlier work? Guo and Lancaster [2] were the first to propose the use of cyclic reduction for testing overdamping. Their approach requires iteration of cyclic reduction to converge, and the computation of two (nonsymmetric) solvents of the associated quadratic matrix equation $Q(X) = 0$. Our test, based on a more complete understanding of the iteration, does not necessarily require the iteration to be run to convergence or the computation of solvents, even for a positive test result, and so is much more efficient. Moreover, we have analyzed the convergence for weakly overdamped quadratics, for which the strict inequality in (2) is replaced by a weak inequality ($\geq$). We find that for weakly overdamped Q with equality in (2) for some nonzero x, the iteration is linearly convergent with constant at worst $1/2$ in the generic case. A reasonable speed of convergence can therefore be expected in almost all practically important cases. Veselić [6] and Higham, Tisseur, and Van Dooren [4] have previously shown that every hyperbolic quadratic can be reformulated as a definite pencil $L(\lambda) = \lambda X + Y$ of twice the dimension, and this connection is explored in detail and in more generality by Higham, Mackey and Tisseur [3], but these papers did not give a practical procedure for ensuring that X or Y is definite.

References

[1] C.-H. Guo, N. J. Higham, and F. Tisseur. Detecting and solving hyperbolic quadratic eigenvalue problems. MIMS EPrint 2007.117, Manchester Institute for Mathematical Sciences, The University of Manchester, UK, Sept. 2007.

[2] C.-H. Guo and P. Lancaster. Algorithms for hyperbolic quadratic eigenvalue problems. *Math. Comp.*, 74(252):1777–1791, 2005.

[3] N. J. Higham, D. S. Mackey, and F. Tisseur. Notes on hyperbolic matrix polynomials and definite linearizations. MIMS EPrint 2007.97, Manchester Institute for Mathematical Sciences, The University of Manchester, UK, July 2007.

[4] N. J. Higham, F. Tisseur, and P. M. Van Dooren. Detecting a definite Hermitian pair and a hyperbolic or elliptic quadratic eigenvalue problem, and associated nearness problems. *Linear Algebra Appl.*, 351–352:455–474, 2002.

[5] P. Lancaster. *Lambda-Matrices and Vibrating Systems.* Pergamon Press, Oxford, 1966. Reprinted by Dover, New York, 2002.

[6] K. Veselić. A Jacobi eigenreduction algorithm for definite matrix pairs. *Numer. Math.*, 64:241–269, 1993.

Stopping Criteria for LSQR

Xiao-Wen Chang, Chris Paige, David Titley-Peloquin

Given $A \in \mathcal{R}^{m\times n}$ with full column rank and $b \in \mathcal{R}^m$, the linear least squares (LS) problem is:

$$\min_x \|b - Ax\|_2. \tag{1}$$

Most practical LS problems contain uncertainties in the data, and instead of solving (1) with ideal data A and b we can at best solve

$$\min_x \|(b+f) - (A+E)x\|_2, \quad \text{where} \quad \|E\|_F \le \alpha\|A\|_F \quad \text{and} \quad \|f\|_2 \le \beta\|b\|_2 \tag{2}$$

for some hopefully approximately known $\alpha \ge 0$ and $\beta \ge 0$. Even if we have perfect data A and b we cannot expect to solve (1) exactly in floating point arithmetic; the best we can generally hope to do is solve a problem of the form (2) with $\alpha = \mathcal{O}(u)$ and $\beta = \mathcal{O}(u)$, u being the machine's unit roundoff.

We say that a given vector $y \in \mathcal{R}^n$ is a *numerically acceptable LS solution* if it is the exact LS solution of a problem within the accepted range of relative errors in the data. In other words, y is a numerically acceptable LS solution if and only if there exist perturbations E and f such that y satisfies the normal equations:

$$(A+E)^T[(b+f) - (A+E)y] = 0, \quad \text{with} \quad \|E\|_F \le \alpha\|A\|_F, \quad \|f\|_2 \le \beta\|b\|_2. \tag{3}$$

If here $\alpha = O(u)$ and $\beta = O(u)$, then y is a *backward stable LS solution.*

In many applications we are given an approximate LS solution y and we would like to verify if it is indeed a numerically acceptable (or a backward stable) LS solution. For example y might be the output of a fast algorithm for structured LS problems whose stability is unknown, or perhaps an iterate from an iterative method for large sparse LS problems. In many practical applications (but not necessarily for ill-posed problems) if y is a numerically acceptable LS solution as defined in (3) then we can be satisfied with it.

Obviously (3) holds if and only if $\hat{\xi} \leq 1$, where

$$\hat{\xi} \equiv \min_{E,f,\xi} \left\{ \xi \ : \ (A+E)^T[(b+f)-(A+E)y] = 0, \ \|E\|_F \leq \xi\alpha\|A\|_F, \ \|f\|_2 \leq \xi\beta\|b\|_2 \right\}. \quad (4)$$

Unfortunately the analytical solution of (4) remains an open question. Some easily computable upper bounds for $\hat{\xi}$ are known, and these can be used to give sufficient conditions for $\hat{\xi} \leq 1$. Such conditions are commonly used as stopping criteria for the iterative solution of large sparse LS problems. For example with $r \equiv b - Ay$ the stopping criteria

$$\|r\|_2 \leq \alpha\|A\|_F\|y\|_2 + \beta\|b\|_2 \quad \text{and} \quad \|A^T r\|_2/\|r\|_2 \leq \alpha\|A\|_F, \quad (5)$$

used in algorithms LSQR in [5] and MINRES and MINRES-QLP in [2], are based on sufficient conditions for $\hat{\xi} \leq 1$. See also the comments in [1, p.309].

Our first contribution is the following: we analyze the behaviour of the quantities $\|r\|_2$ and $\|A^T r\|_2/\|r\|_2$ in minimum residual iterative methods. We explain the previously puzzling observation that $\|A^T r\|_2/\|r\|_2$ initially remains roughly constant or oscillates for ill-conditioned problems, and does not start to decrease until $\|r\|_2$ has reached its minimum. See the website *www.cs.mcgill.ca/˜dtitle/householder* for some sample figures. Our analysis and numerical examples demonstrate that the stopping criteria in (5) can be much too pessimistic, leading to many unnecessary iterates, and sometimes even failing to detect a numerically acceptable (or a backward stable) LS solution when it exists.

Consider now the following two minimization problems, the first for some $\theta > 0$:

$$\begin{aligned} \mu_\theta(y) &\equiv \min_{\Delta A, \Delta b} \left\{ \|[\Delta A, \theta\Delta b]\|_F \ : \ (A+\Delta A)^T[(b+\Delta b)-(A+\Delta A)y] = 0 \right\}, \\ \mu_\infty(y) &\equiv \lim_{\theta\to\infty} \mu_\theta(y) \ = \ \min_{\Delta A} \left\{ \|\Delta A\|_F \ : \ (A+\Delta A)^T[b-(A+\Delta A)y] = 0 \right\}. \end{aligned}$$

These problems were solved in [6] by Waldén, Karlson and Sun.

The question is how to use $\mu_\theta(y)$ or $\mu_\infty(y)$ to determine if y is a numerically acceptable LS solution as defined in (3) and (4). Gu [4] showed that when $\alpha = \mathcal{O}(\tau)$ and $\beta = \mathcal{O}(\tau)$ for any scalar τ, the given y is a numerically acceptable LS solution if and only if

$$\mu_\infty(y) \leq \mathcal{O}(\tau)\|A\|_F.$$

So in this case $\mu_\infty(y)$ can be used to determine if y is a numerically acceptable LS solution, in particular a backward stable LS solution.

Gu's results can easily be extended to the case $\alpha \gg \beta$. Unfortunately when $\alpha \ll \beta$ the condition $\mu_\infty(y) \leq \mathcal{O}(\alpha)\|A\|_F$ is only sufficient for y to be an acceptable solution, whereas $\mu_\infty(y) \leq \mathcal{O}(\beta)\|A\|_F$ is necessary but not sufficient. Note that the case $\alpha \ll \beta$ often occurs in practical LS applications, since b is often a measurement vector and usually subject to a much larger measurement error than the matrix A. Our second contribution is the following: we show how to use $\mu_\theta(y)$ with an appropriate θ to determine if y is a numerically acceptable

LS solution for *any* choice of α and β in (4). As a consequence of this, we now have a computable necessary and sufficient condition to determine if a given vector y is in fact a numerically acceptable LS solution. To our knowledge, this is the only known such criterion.

Computing $\mu_\theta(y)$ explicitly costs $\mathcal{O}(m^3)$ flops, making the result of Waldén, Karlson and Sun too expensive to be useful directly in practical applications. Our final contribution is to give a new estimate of $\mu_\theta(y)$,

$$\tilde{\mu}_\theta(y) \equiv \frac{\eta}{\sqrt{\eta^2 + \|A\|_2^2}} \frac{\|A^T r\|_2}{\|r\|_2},$$

where $r \equiv b - Ay$ and $\eta \equiv \theta\|r\|_2/\sqrt{1+\theta^2\|y\|_2^2}$. Note that $\tilde{\mu}_\theta(y)$ can be estimated cheaply provided an estimate of $\|A\|_2$ is available or can be computed cheaply, as is usually the case. We show that asymptotically, in the limit as y approaches the true LS solution $\hat{x}$,

$$1 \leq \frac{\mu_\theta(y)}{\tilde{\mu}_\theta(y)} \leq \sqrt{\frac{\eta^2 + \sigma_{\max}^2(A)}{\eta^2 + \sigma_{\min}^2(A)}}, \tag{6}$$

where $\sigma_{\max}(A)$ and $\sigma_{\min}(A)$ represent the largest and smallest singular values of A, respectively. The above demonstrates that $\tilde{\mu}_\theta(A)$ is asymptotically tight, except for ill-conditioned LS problems in which $\eta \ll \|A\|_2$. In practice we find that for well-conditioned problems, $\tilde{\mu}_\theta(y)$ is usually of the same order of magnitude as $\mu_\theta(y)$, even when $\|y - \hat{x}\|_2$ is large.

Recent literature contains many bounds or estimates for $\mu_\infty(y)$ (see for example [3] and all the references therein), some of which can be generalized to give bounds or estimates for $\mu_\theta(y)$. Our estimate $\tilde{\mu}_\theta(y)$ is to our knowledge the only one which *both* has a theoretical bound of the form (6) *and* can be computed in $\mathcal{O}(mn)$ flops. Our estimate $\tilde{\mu}_\theta(y)$ can therefore be used to design an ideal stopping criterion for LSQR or any other iterative method for large sparse LS problems.

References

[1] A. Björck, *Numerical Methods for Least Squares Problems*, SIAM, Philadelphia PA, 1996.

[2] S.-C. Choi, *Iterative Methods for Singular Linear Equations and Least-Squares Problems*, Ph.D. Thesis, Scientific Computing & Computational Mathematics, Stanford University, 2006.

[3] J. F. Grcar, *Optimal sensitivity analysis of linear least squares*, Technical Report LBNL-52434. Lawrence Berkeley National Laboratory, 2003.

[4] M. Gu, *Backward perturbation bounds for linear least squares problems*, SIAM J. Matrix Anal. Appl., 20 (1998), pp. 363–372.

[5] C.C. Paige and M.A. Saunders, *LSQR: An Algorithm for Sparse Linear Equations and Sparse Least Squares*, ACM Transactions on Mathematical Software, 8:43–71 (1982).

[6] B. Waldén, R. Karlson and J.-G. Sun, *Optimal backward perturbation bounds for the linear least squares problem*, Numerical Linear Algebra with Applications, 2:271–286 (1995).

Continuous analogues of QR, LU, SVD and Schur

Lloyd N. Trefethen

Numerical linear algebra is built on algorithms to compute matrix factorizations such as the QR, LU, SVD, Schur, Cholesky, and eigenvalue decompositions. We apply these algorithms to matrices which are often obtained by discretizing continuous functions. This raises the question, what would the analogous factorizations and algorithms be if one dealt with the continuous objects directly? For example, suppose A is a "quasimatrix" (Stewart's term) with n columns, each of which is a continuous function defined on an interval $[a, b]$. What is the appropriate definition of an SVD of A? How can one compute it?

Undoubtedly many people have asked such questions over the years. Beginning about five years ago with a collaboration between myself and my former student Zachary Battles, they took on a concrete form as we developed the chebfun system in Matlab for computing with continuous functions instead of discrete vectors. We were able to develop Gram-Schmidt algorithms for QR and SVD factorizations, which can be used for example to compute condition numbers and pseudoinverses of quasimatrices. Other processes encountered obstacles, however, including Householder algorithms, LU factorizations, eigenvalue decompositions, and Krylov subspace iterations.

More recently the chebfun project has moved into a new phase of rapid development with the addition to the project of Ricardo Pachón, Rodrigo Platte, and Toby Driscoll. This has led to new ideas and implementations of some of these factorizations and algorithms—with some surprising twists along the way. In this talk we report on these developments and consider what significance they may have for scientific computing.

Reference: Z. Battles and L. N. Trefethen, An extension of MATLAB to continuous functions and operators, *SIAM J. Sci. Comput.* 25 (2004), 1743–1770.

Algorithms for Matrix Column Selection

Joel A. Tropp

Suppose that A is an $m \times n$ matrix whose columns have unit ℓ_2 norm. A deep result of Bourgain and Tzafriri states that this matrix contains a collection τ of columns for which

$$|\tau| \geq \frac{cn}{\|A\|^2}$$

and

$$0.5\|x\| \leq \|Ax\| \leq 2\|x\|, \qquad \operatorname{supp}(x) \subset \tau.$$

The number c is a universal constant. In words, every matrix with normalized columns contains a large column submatrix that is exceptionally well conditioned.

This talk describes a randomized, polynomial-time algorithm for producing a subset τ that satisfies the foregoing conditions.

Balancing Incomplete Factorizations for Preconditioning

Miroslav Tůma, Rafael Bru, José Marín, José Mas

We consider the construction of robust preconditioners for iterative methods. In particular, we are interested in algebraic procedures for incomplete factorizations. The preconditioners are used for the solution of the system

$$Ax = b$$

for right-hand side vector $b \in I\!R^n$, vector $x \in I\!R^n$ of unknowns, and large sparse system matrix $A \in I\!R^{n \times n}$. We would like to get robust preconditioners rather by *new formulation* of the algorithmic scheme of the factorization than by the means of matrix preprocessings or ad hoc modifications of matrix factors.

Incomplete factorizations represent a class of algebraic preconditioners which is important from both theoretical and practical points of view. Early incomplete factorizations were tightly connected to particular discretizations of partial differential equations. They were often considered as specialized interpolation techniques. During their further development, the factorizations achieved a considerable degree of efficiency, and their algorithms were accompanied by numerous theoretical results for model problems. Nevertheless, new large problems, arising, e.g., from discretizations of complicated operators on unstructured grids still require more robust preconditioners. High-quality incomplete factorizations are also needed inside some multilevel schemes and in non-PDE applications.

Two important general classes of techniques for increasing robustness of preconditioners are *matrix preprocessings* and *modifications of factors.* The former class includes, for example, apriori diagonal modifications of A, matrix prefiltrations, symmetric and nonsymmetric reorderings with possible scalings applied to A. The latter class covers various systematic or ad hoc modifications of incomplete factors by dropping or changing their entries. The modifications may be performed locally, e.g., directed by decisions on magnitudes of computed entries, but they can be also controlled globally by prescribing or extending sparsity structure of the factors or by pre-specifying safe dropping rules. Several sophisticated algorithms leading to high quality incomplete factors were developed over time. The new improvements are often motivated by new theoretical results. In addition, efficient preprocessings and modification strategies are often combined.

One recent group of incomplete factorizations can be linked to factorized approximate inverses. For example, information from the inverse factors forms a crucial part of condition estimators and such procedures may be successfully used for dropping control in the incomplete factorization [1], [2]. Another step in this direction is represented by the computation of breakdown-free incomplete factorization RIF [3] via an intermediate factorized approximate inverse.

In this talk we will describe the new algorithm which computes both direct and inverse incomplete factors of A at the same time. The order of computation of intermediate quantities in this algorithm may enable to monitor conditioning of these factors, and, in addition, the algorithm can be combined with various dropping throughout the incomplete factorization. To get a flavour of the result, consider here, the simplest case of this factorization of the symmetric and positive definite A. Let $A = LDL^T$ be the factorization, where L is unit lower triangular and D is diagonal. The algorithm then computes the factors L, D, as well as the inverse of L. The basic overview of the result can be found in [4]. This result was originally

derived using the theory for Sherman-Morrison-based approximate inverses [5] but it can be motivated in a different way as well.

Consider the left-looking algorithm which computes the columns v_k of V for $k = 1, \ldots, n$ as

$$v_k = a_k - e_k - \sum_{i=1}^{k-1} \frac{(a_k - e_k)^T u_i}{v_i^T e_i} v_i, \tag{1}$$

where a_k denotes k-th column of A and u_i denotes the i-th column of $triu(L - VD^{-1})$. It can be proved that $D = diag(V + I)$ and $L = tril(V + I)D^{-1}$. Clearly, the columns of U needed in the computation can be readily obtained from the previously computed part of V. In addition, it can be shown that $U = L^{-T}D^{-1}$, i.e., U is the scaled inverse factor of A.

During the course of the algorithm, the partially computed direct factor L is used in computation of L^{-1}, and partially computed L^{-1} is used to compute L. Moreover, both L and L^{-1} are involved in decisions related to the incompleteness of the evaluation. In the following we will show one possible set of dropping rules. The rules can be cheaply implemented due to the favorable order of evaluation of the factors.

Let the incomplete algorithm provide approximate quantities as $\hat{L}$, $\hat{D}$, $\widehat{L^{-1}}$, possibly in a scaled form. Denote $\hat{l}_{jk} = (\hat{L})_{jk}$, $\hat{\ell}_{jk} = (\widehat{L^{-1}})_{jk}$ for appropriate indices j, k. Motivated by [1], we drop the entries $\hat{l}_{jk}$ of the column k of $\hat{L}$ if they satisfy

$$|\hat{l}_{jk}| \parallel e_k^T \widehat{L^{-1}} \parallel \leq \tau_1 \tag{2}$$

for a given drop tolerance τ_1. Similarly, we drop entries $\hat{\ell}_{jk}$ of the column of k of $\widehat{L^{-1}}$ if they satisfy

$$|\hat{\ell}_{jk}| \parallel e_k^T \hat{L} \parallel \leq \tau_2 \tag{3}$$

The new algorithm allows straightforward embedding of both these rules. Let us remind that the left-looking algorithm computes in its major step one column of V. In particular, one diagonal entry, the scaled $k-$th row of $\widehat{L^{-1}}$ and the scaled $k-$th column of $\hat{L}$ are computed in the $k-$th major step. All these items are stored in V. Then the conditions (2) and (3) imply that dropping can be fully described via the entries of V which have been computed so far. Moreover, the norms used in the dropping rules can be cheaply computed on-the-fly.

We believe that the availability of a computed approximate inverse in the new algorithm contributes to its high practical robustness, which we experimentally confirmed. In addition, this "two-way" dropping possibility, strongly different from the "one-way" dropping in RIF, may be also partially responsible for fast computation of the factorization in the fully sparse case. We are able to compute the new factorization much faster than RIF (for problems with hundreds of thousands of variables) [5] and we still believe in further improvements.

We call the basic algorithm in the special case of SPD A the *balanced* incomplete factorization since the algorithm may monitor conditioning of both factors (approximations to L and L^{-1}) and balance their sizes (not necessarily uniformly) based on this information. Nevertheless, by the title of the talk we would like to express more a *computational approach* than to discuss one specific algorithm. The algorithm reminds us how much effort may be still needed for better understanding of basic matrix decompositions. We are strongly persuaded that in order to move on with incomplete factorizations towards more efficient preconditioned iterative methods for solving difficult large and sparse problems, it is necessary to develop and analyze more robust preconditioners. In some cases, only algebraic methods are available. Combining direct and inverse factors still seem to have enough potential in this effort.

In our talk we would like not only to describe our results for SPD matrices, but to discuss additional insight into the decomposition and an extension to the nonsymmetric case. Last but not least, we would like to mention that fully sparse implementations of the algorithms like the one discussed here are straightforward, but not simple. In addition, they typically need further modifications. Finally, we would like to discuss some details which are often hidden, but which may contribute to usefulness of the algorithms of the considered class.

[1] M. Bollhöfer and Y. Saad. On the relations between ILUs and factored approximate inverses. *SIAM J. Matrix Anal. Appl.*, 24(1):219–237, 2002.

[2] M. Bollhöfer and Y. Saad. Multilevel preconditioners constructed from inverse-based ILUs. *SIAM J. Sci. Comput.*, 27(5):1627–1650, 2006.

[3] M. Benzi and M. Tůma. A robust incomplete factorization preconditioner for positive definite matrices. *Numer. Linear Algebra Appl.*, 10(5-6):385–400, 2003.

[4] R. Bru, J. Cerdán, J. Marín, and J. Mas. Preconditioning sparse nonsymmetric linear systems with the Sherman-Morrison formula. *SIAM J. Sci. Comput.*, 25(2):701–715, 2003.

[5] R. Bru, J. Marín, J. Mas, and M. Tůma. Balanced incomplete factorization. *to appear in SIAM J. Sci. Comput.*, see the revised version at http://www.cs.cas.cz/˜tuma/ps/bif.pdf, 2008.

Tensor rank estimates and tensor-based algorithms

Ivan Oseledets, Eugene Tyrtyshnikov, Nickolai Zamarashkin

Tensor structures are becoming popular and powerful tools in construction of low-complexity algorithms. For many practical cases we observe that the results of matrix operations can be well-approximated by matrices of low tensor rank (compared to the size). Exact tensor representations and bounds, all the same, are usually not known or trivial, e.g.

$$(A \otimes B)^{-1} = A^{-1} \otimes B^{-1}.$$

Thus, if a matrix is invertible and of tensor rank 1, then the inverse is also of tensor rank 1. We are interested to suggest some less trivial results of the same type.

Let A be a matrix of order $N = n^2$ and the tensor rank be defined as the minimal possible number of Kronecker (tensor) products of $n \times n$ matrices whose sum is A. In general, the tensor rank of A^{-1} may depend not only on the tensor rank of A but also on its size N. A trivial upper bound is, of course, N. However, in some cases it is never attained.

Theorem 1. *If an invertible matrix A of order $N = n^2$ is of tensor rank 2, then the tensor rank of A^{-1} does not exceed $n = \sqrt{N}$. This bound is sharp.*

Now, denote by $M(n, k)$ the class of matrices of the form

$$A = \sum_{i=1}^{k} B_i \otimes C_i,$$

where B_i and C_i are matrices of order n and, in addition,

$$\min(\mathrm{rank} R_i, \mathrm{rank} S_i) \leq 1, \quad 2 \leq i \leq k.$$

We can suggest some estimates where tensor ranks do not depend upon size.

Theorem 2. *If a matrix $A \in M(n,3)$ is invertible then $A^{-1} \in M(n,5)$.*

Theorem 3. *If a matrix $A \in M(n,k)$ is invertible then $A^{-1} \in M(n, O(k^2))$.*

These results provide some ground for approximate tensor-rank computations. In particular, they allow us to describe a class of invertible two-level Toeplitz matrices of low tensor rank, for which the inverses are well-approximated by matrices of low tensor rank. Thus we obtain a rigorous base for fast approximate computation of the inverses to two-level Toeplitz matrices in the low-tensor-rank format. Special attention is still needed to capture and maintain a viable structure in the tensor factors. We present some algorithms of sublinear complexity (with respect to the size) and some results of using a low-displacement-rank structure and a circulant-plus-low-rank structure.

We also discuss some extensions of the results to the case of three and more tensor factors.

A Generalization of the QR-algorithm for Semiseparable plus Diagonal Matrices

Marc Van Barel, Raf Vandebril, Nicola Mastronardi

In this talk we propose a new iterative method based on a structured rank factorization for computing eigenvalues of semiseparable and semiseparable plus diagonal matrices. Also the case of higher order semiseparability ranks is included. More precisely, instead of the traditional QR-iteration, a QH-iteration will be used with Q an orthogonal matrix and H a Hessenberg-like matrix (often called lower semiseparable matrix) having the lower triangular part of semiseparable form, i.e., all submatrices that one can take out of the lower triangular part of the matrix have maximum rank p.

It will be shown that this iteration is very effective for computing the eigenvalues of structured rank matrices. Whereas the traditional QR-method applied onto semiseparable (plus diagonal) and Hessenberg-like matrices uses similarity transformations involving $2p(n-1)$ Givens transformations (when p denotes the semiseparability rank), this iteration only needs $p(n-1)$ iterations, which is comparable to the tridiagonal and Hessenberg situation when p equals one. It will also be shown that this method can, in some sense, be interpreted as an extension of the traditional QR-method for Hessenberg matrices, i.e., the traditional case also fits into this framework.

Based on convergence results for a generalized QR-type iteration, it follows that this iteration inherits also an extra type of convergence behavior w.r.t. the traditional QR-method.

We indicate how the algorithm can be implemented in an implicit way, based on the Givens weight representation of the involved structured rank matrices. Numerical experiments will show the viability of the approach. It will be shown that the new approach needs less iteration steps and also gives rise to more accurate results. Moreover, the traditional QR-method for semiseparable (plus diagonal) matrices is 50% more expensive than this new iteration.

A manifold approach for comparing two matrices

P. Van Dooren, P.A. Absil, T. Cason, C. Fraikin, Y. Nesterov

Extended abstract

We go over a number of optimization problems defined on a manifold in order to compare two matrices possibly of different order. When comparing two matrices A and B it is often natural to allow for a class of transformations acting on these matrices. For instance, when comparing adjacency matrices A and B of two graphs with an equal number of nodes, one can allow symmetric permutations $P^T AP$ on one matrix in order to compare it to B, since this is merely a relabeling of the nodes of A. The so-called comparison then consists in finding the best match between A and B under this class of transformations.

A more general class of transformations would be that of unitary similarity transformations Q^*AQ, where Q is a unitary matrix. This leaves the eigenvalues of A unchanged but rotates its eigenvectors, which will of course play a role in the comparison between A and B. If A and B are of different order, say m and n, one may want to consider their restriction on a lower dimensional subspace:

$$U^*AU \quad \text{and} \quad V^*BV, \tag{1}$$

yielding two square matrices of equal dimension $k \leq \min(m,n)$, which can again be compared. Here U and V belong to $St(k,m)$ and $St(k,n)$ respectively, and $St(k,m) = \{U \in \mathbb{C}^{m\times k} : U^*U = I_k\}$ denotes the *compact Stiefel manifold*.

Several measures of comparison between these restrictions have been proposed in the literature. Fraikin et *al.* [1] propose to maximize the inner product between the *isometric projections*, U^*AU and V^*BV, namely:

$$\arg \max_{\substack{U^*U=I_k \\ V^*V=I_k}} \langle U^*AU, V^*BV\rangle := \Re\,\mathrm{tr}((U^*AU)^*\,(V^*BV)). \tag{2}$$

They show this is also equivalent to

$$\arg \max_{\substack{X=VU^* \\ U^*U=I_k \\ V^*V=I_k}} \langle XA, BX\rangle = \Re\,\mathrm{tr}(A^*X^*BX), \tag{3}$$

and eventually show how this problem is linked to the notion of graph similarity introduced by Blondel et *al.* in [3]. The graph similarity matrix S introduced in that paper also proposes a measure of comparing two matrices A and B via the fixed point of a particular iteration. But it is shown in [2] that this is equivalent to the optimization problem

$$\arg \max_{\|S\|_F=1} \langle SA, BS\rangle = \Re\,\mathrm{tr}\big(A^T S^T BS\big) \tag{4}$$

or also

$$\arg \max_{\|S\|_F=1} \langle S, BSA^*\rangle = \Re\,\mathrm{tr}\big(A^T S^T BS\big). \tag{5}$$

We can also use a distance measure rather than an inner product to compare two matrices :

$$\mathrm{dist}\,(M,N) = \|M-N\|_F^2 = \mathrm{tr}((M-N)^*(M-N)).$$

We analyze distance minimization problems that are essentially the counterparts of the similarity measures defined above. These are

$$\arg \min_{\substack{U^*U=I_k \\ V^*V=I_k}} \operatorname{dist}\left(U^*AU, V^*BV\right), \tag{6}$$

$$\arg \min_{\substack{X=VU^* \\ U^*U=I_k \\ V^*V=I_k}} \operatorname{dist}\left(XA, BX\right), \tag{7}$$

and

$$\arg \min_{\substack{X=VU^* \\ U^*U=I_k \\ V^*V=I_k}} \operatorname{dist}\left(X, BXA^*\right). \tag{8}$$

for the problems involving two isometries U and V. Notice that these three distance problems are not equivalent although the corresponding inner product problems are equivalent. Similarly, we analyze the two problems

$$\arg \min_{\|S\|_F=1} \operatorname{dist}(SA, BS) = \operatorname{tr}\left((SA-BS)^*(SA-BS)\right) \tag{9}$$

and

$$\arg \min_{\|S\|_F=1} \operatorname{dist}(SA, BS) = \operatorname{tr}\left((SA-BS)^*(SA-BS)\right), \tag{10}$$

for the problems involving a single matrix S. Again, these are not equivalent in their distance formulation. We will develop optimality conditions for those different problems, indicate their relations with existing problems from the literature and give an analytic solution for particular matrices A and B.

References

[1] C. Fraikin, Y. Nesterov, and P. Van Dooren. *Optimizing the Coupling Between Two Isometric Projections of Matrices*, SIAM J. Matrix Anal. Appl., 2007.

[2] C. Fraikin and P. Van Dooren. *Graph matching with type constraints on nodes and edges*, Dagstuhl Seminar N 07071 on Web Information Retrieval and Linear Algebra Algorithms, Feb. 2007.

[3] V. D. Blondel, A. Gajardo, M. Heymans, P. Senellart, and P. Van Dooren, *A measure of similarity between graph vertices: applications to synonym extraction and Web searching*, SIAM Rev., 46 (2004), pp. 647–666.

[4] P.-A. Absil, R. Mahony, and R. Sepulchre. *Optimization Algorithms on Matrix Manifolds*, Princeton University Press, 2008

IDR(s): a family of simple and fast algorithms for solving large nonsymmetric systems of linear equations

Peter Sonneveld, Martin van Gijzen

1 Introduction

We consider the linear system

$$\boldsymbol{A}\boldsymbol{x} = \boldsymbol{b}$$

with $\boldsymbol{A} \in \mathcal{C}^{N\times N}$ a large, sparse and non-Hermitian matrix. In 1980, Sonneveld [6] proposed the iterative method IDR for solving such systems. IDR has several favorable features: it is simple, uses short recurrences, and computes the exact solution in at most $2N$ steps (matrix-vector multiplications) in exact arithmetic. Analysis of IDR revealed a close relation with Bi-CG [1]. It was shown in [6] that the iteration polynomial constructed by IDR is the product of the Bi-CG polynomial with another, locally minimizing polynomial. Sonneveld's observation that the Bi-CG polynomial could be combined with another polynomial without matrix-vector multiplications with $\boldsymbol{A}^H$ led to the development first of CGS [3] and later of Bi-CGSTAB [5].

Over the years, CGS and Bi-CGSTAB have completely overshadowed IDR. This is unfortunate since, although there is a clear relation between Bi-CG-type methods and the original IDR method, the underlying ideas are completely different. This suggests that by exploiting the differences new methods may be developed.

Bi-CG, CGS, Bi-CGSTAB and their descendants are essentially based on the computation of two mutually bi-orthogonal bases for the Krylov subspaces $\mathcal{K}^n(\boldsymbol{A}, r_0)$ and $\mathcal{K}^n(\boldsymbol{A}^H, \tilde{r}_0)$. The IDR family, on the other hand, generates residuals that are forced to be in nested subspaces of decreasing dimension. These subspaces are related by $\mathcal{G}_j = (\boldsymbol{I} - \omega_j \boldsymbol{A})(\mathcal{S} \cap \mathcal{G}_{j-1})$, where $\mathcal{S}$ is a fixed proper subspace of $\mathcal{C}^N$, and the ω_j's are nonzero scalars. It can be shown [4] that in general the reduction in dimension of $\mathcal{G}_j$ is the co-dimension s of $\mathcal{S}$. Hence, after at most N/s dimension reductions $\mathcal{G}_m = \{\boldsymbol{0}\}$, for some $m \leq \lceil N/s \rceil$.

In the next section, we describe IDR(s), a generalisation of the original IDR algorithm.

2 IDR(s)

The residual $\boldsymbol{r}_{n+1} = \boldsymbol{b} - \boldsymbol{A}\boldsymbol{x}_{n+1}$ corresponding to the iterate $\boldsymbol{x}_{n+1}$ after $n+1$ iterations is in $\mathcal{G}_{j+1}$ if

$$\boldsymbol{r}_{n+1} = (\boldsymbol{I} - \omega_{j+1}\boldsymbol{A})\boldsymbol{v} \quad \text{with } \boldsymbol{v} \in \mathcal{G}_j \cap \mathcal{S}\ .$$

We may assume, without loss of generality, that the space $\mathcal{S}$ is the left nullspace of some $N \times s$ matrix $\boldsymbol{P}$, i.e.:

$$\boldsymbol{P} = (\boldsymbol{p}_1\ \boldsymbol{p}_2\ \ldots\ \boldsymbol{p}_s), \quad \mathcal{S} = \mathcal{N}(\boldsymbol{P}^H)\,.$$

The vector $\boldsymbol{v}$ is a linear combination of the residuals in $\mathcal{G}_j$, and can be written as

$$\boldsymbol{v} = \boldsymbol{r}_n - \sum_{j=1}^{s} \gamma_j \Delta \boldsymbol{r}_{n-j}, \tag{1}$$

with $\Delta \boldsymbol{r}_l = \boldsymbol{r}_{l+1} - \boldsymbol{r}_l$. Here, we introduced the residual difference vectors $\Delta \boldsymbol{r}_l$ to be able to update the solution $\boldsymbol{x}_{n+1}$ together with the residual $\boldsymbol{r}_{n+1}$.

Since $\boldsymbol{v}$ is in $\mathcal{S} = \mathcal{N}(\boldsymbol{P}^H)$, we also have

$$\boldsymbol{P}^H \boldsymbol{v} = \boldsymbol{0}\,. \tag{2}$$

Combining (1) and (2) yields an $s \times s$ linear system which can be solved for the γ_j's, so that we can determine $\boldsymbol{v}$ and $\boldsymbol{r}_{n+1}$.

Since $\mathcal{G}_{j+1} \subset \mathcal{G}_j$, repeating these calculations will produce new residuals $\boldsymbol{r}_{n+2}, \boldsymbol{r}_{n+3}, \cdots$, in $\mathcal{G}_{j+1}$. Once $s+1$ residuals in $\mathcal{G}_{j+1}$ have been computed, we can expect the next residual to be in $\mathcal{G}_{j+2}$. This process is repeated until $\boldsymbol{r}_n \in \mathcal{G}_m = \{0\}$ for some n and m. This implies that the exact solution can be computed in exact arithmetic using at most $\lceil N/s \rceil\,(s+1)$ matrix-vector products.

Although we may choose ω_{j+1} freely in the calculation of the first residual in $\mathcal{G}_{j+1}$, the same value must be used in the calculation of the subsequent residuals in $\mathcal{G}_{j+1}$. A suitable choice for ω_{j+1} is the value that minimizes the norm of $\boldsymbol{r}_{n+1}$. These ideas lead to the IDR(s) algorithm. For a detailed description of the algorithm we refer to [4].

It can be shown that the most basic variant of our algorithm, IDR(1), is mathematically equivalent to Bi-CGSTAB. Higher-order IDR(s) methods are related to the Bi-CGSTAB generalisation ML(k)BiCGSTAB of Yeung and Chan [7].

3 Practical issues and performance

The IDR(s) algorithm is simple and easy to implement. The memory requirements are modest and the overhead for vector operations is limited. IDR(1) is about as expensive as Bi-CGSTAB in terms of computations and memory requirements, and in our experience is just as stable. Increasing s makes the algorithm slightly more expensive per iteration, but in all our experiments, increasing s also yields a significant decrease in the number of iterations. In practice, we found that for most problems $s = 4$ is a good choice. We have also observed in our tests that choosing the columns of $\boldsymbol{P}$ randomly improved the stability of the method.

We have performed extensive experiments with test problems from many applications. In all our examples, IDR(s) performed as well as or better than Bi-CGSTAB. In particular, for (indefinite) Helmholtz problems, we found that IDR(s) often performs much better than Bi-CGSTAB. For example, in [4] we describe a 3D Helmholtz problem for which IDR(6) outperforms Bi-CGSTAB by a factor of six in terms of CPU time, and Bi-CGstab(8) [2] by a factor of more than two.

4 Concluding remarks

Based on our experience, IDR(s) seems quite promising and it also seems to outperform the state-of-the-art Bi-CG-type methods for important classes of problems. A very important fact in our opinion is that the algorithm that we have implemented is only one member of a family that can be derived using the IDR idea. Therefore, the IDR mechanism offers a new approach for the development of iterative methods for linear systems, different from the classical Bi-CG- or GMRES-based approaches.

References

[1] R. Fletcher. Conjugate gradient methods for indefinite systems. *Lecture notes in Mathematics* 506, Springer-Verlag, Berlin, Heidelberg, New York, pp. 73–89, 1976.

[2] G.L.G. Sleijpen and D.R. Fokkema. BiCGstab(ℓ) for linear equations involving matrices with complex spectrum. *ETNA*, 1:11–32, 1994.

[3] P. Sonneveld. CGS: a fast Lanczos-type solver for nonsymmetric linear systems. *SIAM J. Sci. and Statist. Comput.*, 10:36–52, 1989.

[4] P. Sonneveld and M.B. van Gijzen. IDR(s): a family of simple and fast algorithms for solving large nonsymmetric linear systems. Technical Report 07-07, Department of Applied Mathematical Analysis, Delft University of Technology, Delft, The Netherlands, 2007.

[5] H.A. van der Vorst. Bi-CGSTAB: A fast and smoothly converging variant of Bi-CG for the solution of nonsymmetric linear systems. *SIAM J. Sci. Comp.*, 13:631–644, 1992.

[6] P. Wesseling and P. Sonneveld. Numerical Experiments with a Multiple Grid- and a Preconditioned Lanczos Type Method. *Lecture Notes in Mathematics* 771, Springer-Verlag, Berlin, Heidelberg, New York, pp. 543–562, 1980.

[7] M-C. Yeung and T.F. Chan. ML(K)BiCGSTAB: A BiCGSTAB variant based on multiple Lanczos starting vectors. *SIAM J. Sci. Comput.*, 21(4):1263–1290, 1999.

Tensor-Based Biosignal Processing

Sabine Van Huffel

This contribution deals with tensor decompositions and their benefits in biomed- ical signal processing. After a brief introduction of tensors and their differences with matrices, different decompositions are described and their computational algorithms are introduced as generalizations of matrix-based counterparts. In particular, we will focus on 'Parallel Factor Analysis' (Parafac, also known as Candecomp or the CP model), the most popular tensor decomposition, and overview its mathematical properties. The CP model decomposes in a unique way a higher-order tensor in a minimal sum of rank-1 'atoms'. Furthermore, we will give an overview of biomedical applications, in which tensors are useful. In particular, we will focus on the presurgical evaluation of refractory partial epilepsy for the delineation of the irritative and ictal onset zones using long-term electroencephalographic (EEG) recordings . We will extract the potential distribution of the ictal activity using the higher-order CP model and show that only one atom is related to the seizure activity.

References

[1] De Lathauwer L, De Moor B, and Vandewalle J., Computation of the Canonical Decomposition by means of a simultaneous generalized Schur de- composition, SIAM J. Matrix Anal. Appl. 26:295-327, 2004.

[2] M. De Vos, A. Vergult, L. De Lathauwer, W. De Clercq, S. Van Huffel, P. Dupont, A.

Palmini, W. Van Paesschen, Canonical decomposition of ictal scalp EEG reliably detects the seizure onset zone. NeuroImage, to appear.

Multilinear Algebra Computations in Quantum Chemistry

C. Van Loan

The simulation of complex quantum chemical systems of interacting electrons leads to a number of high dimensional tensor computations. In this talk we will highlight the advantages of low-rank approximation and block tensor data structures.

A Rational Generalization of the QR-algorithm

Raf Vandebril, Marc Van Barel, Nicola Mastronardi

In this talk a generalization of the QR-algorithm will be presented. Each iteration step of the new iterative method consists of two substeps: first an RQ-factorization of the initial matrix $A - \sigma I = RQ$ will be computed, followed by a QR-factorization of the matrix $(A - \kappa I)Q^H$. Applying a step of the QR-method with the unitary similarity transformation defined by the Q-factor of the QR-factorization of the transformed matrix $(A - \kappa I)Q^H$, will yield interesting convergence properties. It will be shown that the convergence behavior is related to a subspace iteration based on a rational function in A namely $(A-\sigma I)(A-\kappa I)^{-1}$. Convergence properties of this new iteration will be investigated and examples will be presented, illustrating the effectiveness of this approach.

Numerical Methods for Continuous-time Periodic Systems

Andras Varga

In the recent years, there is an increased interest to address periodic control problems for continuous-time periodic systems of the form

$$\begin{array}{rcl} \dot{x}(t) & = & A(t)x(t) + B(t)u(t) \\ y(t) & = & C(t)x(t) + D(t)u(t) \end{array} \tag{1}$$

where $A(t) \in R^{n\times n}$, $B(t) \in R^{n\times m}$, $C(t) \in R^{p\times n}$, and $D(t) \in R^{p\times m}$ are real periodic matrices of period T. For a recent account of main theoretical developments see Colaneri's survey paper: "Periodic control systems: theoretical aspects" (in Proc. of IFAC Workshop on Periodic Systems, Yokohama, Japan, 2005) and references cited therein.

Among the first methods for continuous-time periodic systems which can be labelled "computational", we mention just two classes of methods proposed to evaluate system norms. In the first class of methods proposed by Zhou and Hajiwara (2002), a so-called frequency-lifting is employed to build equivalent infinite-dimensional system representations. Approximations obtained using various finite truncations of infinite-dimensional Toeplitz matrices are then employed to evaluate the $\mathcal{H}_2$- and $\mathcal{H}_\infty$-norms using complex-valued computation based algorithms (e.g., to solve Lyapunov and Riccati equations). Besides the need of representing

the system data as Fourier series, these methods provide basically only approximate results. High accuracy can be expected only if sufficiently large order approximations are used, which certainly leads to numerical difficulties.

The second class of methods proposed by Lampe and Rosenwasser (2004) relies on closed-form analytic formulas to compute the $\mathcal{H}_2$-norm. Performing the actual computations involves the evaluation of several integrals and the integration of several matrix differential equations. In the view of recent developments to be presented in this talk, this method can not be considered as real alternative to the new generation of computational methods.

The development of satisfactory computational algorithms for continuous-time periodic systems strongly relies on the progress of computational techniques for discrete-time periodic systems in the last 15 years. We will illustrate one of the first developments by showing how to perform Floquet-analysis, a well-known technique of stability analysis, in a numerically satisfactory way.

Let $\Phi_A(t,\tau)$ denote the transition matrix corresponding to a T-periodic $A(t)$ satisfying

$$\frac{\partial \Phi_A(t,\tau)}{\partial t} = A(t)\Phi_A(t,\tau), \quad \Phi_A(\tau,\tau) = I \tag{2}$$

The eigenvalues of the *monodromy* matrix $\Phi_A(T,0)$ are called characteristic multipliers and their magnitudes characterize the asymptotic stability of the periodic system (1). A Floquet exponent (sometimes called a characteristic exponent), is a complex value μ such that $e^{\mu T}$ is a characteristic multiplier of the system. The real parts of the Floquet exponents are called Lyapunov exponents. The system (1) is asymptotically stable if all Lyapunov exponents are negative, Lyapunov stable if the Lyapunov exponents are nonpositive and unstable otherwise.

To understand the numerical difficulties associated with the computation of characteristic exponents consider the following simple example with

$$A(t) = \begin{bmatrix} 0 & 1 \\ -2\dot{\alpha}(t) & 6 - 2\alpha(t) \end{bmatrix},$$

where $\alpha(t) = 15 + 5\sin t$ is periodic of period $T = 2\pi$. The analytic solution can be easily computed via Floquet theory. With the transformation matrix

$$P(t) = \begin{bmatrix} -3 + \alpha(t) & 1/2 \\ 1 & 0 \end{bmatrix}$$

we obtain the Lyapunov similar matrix

$$\widetilde{A}(t) := P(t)A(t)P^{-1}(t) + \dot{P}(t)P^{-1}(t) = \begin{bmatrix} 0 & 0 \\ 2 & 6 - 2\alpha(t) \end{bmatrix}$$

whose characteristic exponents are

$$\mu_1 = 0, \quad \mu_2 = 6 - 2\alpha_0 = -24,$$

where α_0 is the mean value of α over one period. To determine $\Phi(T,0)$, we used the standard Runge-Kutta integration method on the ordinary differential equation (2) satisfied by $\Phi(t,0)$ with a relative tolerance of 10^{-10}. We obtained

$$\overline{\mu}_1 = -3.9 \cdot 10^{-7}, \quad \overline{\mu}_2 = -2.05.....$$

and, as can be seen, there is no even a single accurate digit in the computed value of μ_2!

The basis of several new techniques for solving continuous-time computational problems is the product form representation of the *monodromy* matrix

$$\Phi_A(T,0) = \Phi_A(T,T-\Delta)\cdots\Phi_A(2\Delta,\Delta)\Phi_A(\Delta,0)$$

where $\Delta = T/N$ for a suitably chosen integer period N. The matrix $F_k := \Phi_A(k\Delta,(k-1)\Delta)$ for $k = 1,2,\ldots,$ is obviously N-periodic in a discrete-time setting. Then, the eigenvalues of $\Phi_A(T,0)$ can be alternatively computed using the periodic Schur form based algorithm applied to the N-periodic matrix F_k. Note that in this algorithm, the factors of the matrix product $F_N\cdots F_2F_1$ are individually manipulated, and the product itself is never explicitly computed. The main advantage of using product form representations of monodromy matrices is that potentially unstable integration of differential equations over the whole period can be completely avoided.

By employing this approach, we get the following values of the characteristic exponents for increasing values of N:

N	$\lvert\overline{\mu}_1\rvert$	$\overline{\mu}_2$
1	$3.9\cdot10^{-7}$	-2.05...
2	$2.6\cdot10^{-7}$	complex !!
5	$4.5\cdot10^{-7}$	-11.2653
10	$2.1\cdot10^{-9}$	-19.7407
25	$7.7\cdot10^{-12}$	-23.9921
50	$2.4\cdot10^{-15}$	-23.9982
100	$2.7\cdot10^{-15}$	-23.99995
200	$3.2\cdot10^{-15}$	-23.9999993
500	$1.9\cdot10^{-14}$	-23.999999998

As can be seen, 10-digits accuracy is possible by working directly on the product form representation (instead of forming this product explicitly).

Using similar approaches, recently, a new algorithmic paradigm called *multi-shot method* evolved for solving continuous-time problems (e.g., various periodic matrix differential equations). Using appropriate *exact* discretizations, the continuous-time problems are reduced to equivalent single- or multi-point discrete-time periodic problems for which reliable computational algorithms are available. By solving the discrete-time problems, so-called *periodic generators* are computed (e.g., initial or multi-point conditions), which serve to determine the continuous periodic solutions (usually by integrating the underlying ordinary matrix differential equations with known initial or multi-point conditions).

In the proposed talk we will discuss the solution of periodic Lyapunov, Sylvester or Riccati differential equations by employing product form representations of appropriate Hamiltonian matrices. Multi-shot methods will be also presented to evaluate Hankel, $\mathcal{H}_2$-, or $\mathcal{H}_\infty$-norms (Varga,2005) or to efficiently compute the expression of function and gradient when solving continuous-time periodic output feedback linear-quadratic problems. The experimental evaluation of the multi-shot method to solve periodic Riccati equations in conjunction with special structure preserving (symplectic) numerical integrators has shown the high effectiveness of these techniques.

Many new algorithmic developments are expected in the coming years for continuous-time periodic problems, as for example, efficient computation of frequency-responses, solution of

stabilization and pole assignment problems, order reduction, solution of various controller synthesis problems. Among new tools to be investigated in the future is the applicability of continuous matrix decompositions to solve various computational problems for time-varying periodic systems.

Algorithms and Complexity for Nonnegative Matrix Factorization

Stephen A. Vavasis

Nonnegative matrix factorization (NMF) has emerged as a powerful tool for information retrieval. Since it was popularized by Lee and Seung [6] in 1999, hundreds of papers have been published applying to text retrieval [4], image retrieval [6], analysis of biochemical experiments [2], and even music.

NMF refers to the following computation. Given an $m \times n$ matrix A all of whose entries are nonnegative, and given an integer parameter k such that $k \leq \min(m, n)$, find two matrices $W \in \mathbb{R}^{m \times k}$ and $H \in \mathbb{R}^{k \times n}$ such that $A \approx WH$, and such that W and H both have nonnegative entries. In the absence of this latter constraint, this problem would be solved by the singular value decomposition.

As an example, A might be a term-document matrix for a text corpus. This means that each column of A is in correspondence with a document from a certain collection (corpus) of documents, and each row is in correspondence with a term (meaning a word, except that common words such as articles and conjunctions are omitted). The (i, j) entry of A is determined by the frequency of term i in document j; if the term does not appear, then $A(i, j) = 0$. Typical term-document matrices are extremely sparse. When an NMF algorithm is applied to such an A, the typical outcome is as follows. The method discovers k topics in the corpus; the matrix W encodes the probability that a particular topic will use a particular term, and the matrix H indicates the prevalence of each topic in each document. The discovery of these topics is automatic; the factorization algorithm does not have any prior knowledge about the contents of the corpus.

The algorithm proposed by Lee and Seung and most others in the literature consists of making an initial guess for W and H either randomly or according a heuristic, and then using a local-search improvement to obtain better choices for W and H. These algorithms often work well in practice, but may have difficulties with convergence, and their output is dependent on the initial guess. Several authors use sophisticated optimization tools [5] to seek an optimum fit of WH to A.

The first part of my talk, which is joint with Ali Ghodsi and Michael Biggs, both of the Department of Statistics of University of Waterloo, will discuss another class of algorithms based on greedy rank-one downdating of A. Algorithms of this type have been discovered independently by several groups (including ours) over the past few years [1, 2]. An algorithm of this type is generally much faster than local search, does not have convergence issues (because it is not iterative), and does not depend on an initial guess. We will report on computational experiments with our algorithm, and we will also present theoretical results proving that for simple models of corpora, the algorithm is able to correctly classify documents.

No algorithm for NMF proposed in the literature (including ours) comes with a guarantee of optimality, i.e., a guarantee that the computed W and H globally minimize $\|A - WH\|$

in some norm. In contrast, the SVD has such a global optimality property for the 2- and Frobenius-norm. The second part of my talk presents a complexity result showing that solving NMF to optimality in general is a hard problem. A problem called exact NMF is defined as follows. Given an $m \times n$ nonnegative matrix A whose rank is exactly k, does there exist $W \in \mathbb{R}^{m \times k}$ and $H \in \mathbb{R}^{k \times n}$, both nonnegative, such that $A = WH$? This problem is a special case of optimal NMF in any norm (since an optimal algorithm, when applied to a matrix of rank exactly k, would have to determine whether there is an exact factorization with zero residual). It is also a special case of nonnegative rank determination, a problem considered by Cohen and Rothblum [3]. Our first result is that exact NMF is equivalent to a problem in polyhedral combinatorics that we call Intermediate Simplex. Our second result is that Intermediate Simplex is NP-hard, and hence it is unlikely that an efficient optimal NMF algorithm will be found.

References

[1] Nasimeh Asgarian and Russell Greiner. Using rank-1 biclusters to classify microarray data. Department of Computing Science, and the Alberta Ingenuity Center for Machine Learning,University of Alberta, Edmonton, AB, Canada, T6G2E8, 2006.

[2] Sven Bergmann, Jan Ihmels, and Naama Barkai. Iterative signature algorithm for the analysis of large-scale gene expression data. *Physical Review E*, 67:031902, 2003.

[3] J. Cohen and U. Rothblum. Nonnegative ranks, decompositions and factorizations of nonnegative matrices. *Linear Algebra and its Applications*, 190:149–168, 1993.

[4] T. Hofmann. Probabilistic latent semantic analysis. In Kathryn B. Laskey and Henri Prade, editors, *UAI '99: Proceedings of the Fifteenth Conference on Uncertainty in Artificial Intelligence, Stockholm, Sweden, July 30-August 1, 1999*, pages 289–296. Morgan Kaufmann, 1999.

[5] H. Kim and H. Park. Sparse non-negative matrix factgorizations via alternating non-negativity-constrained least squares for microarray data analysis. Bioinformatics (to appear), 2007.

[6] D. Lee and H. Seung. Learning the parts of objects by non-negative matrix factorization. *Nature*, 401:788–791, 1999.

Positive Descriptor Systems

Elena Virnik

We consider linear time-invariant control systems of the form

$$\begin{aligned} E\dot{x} &= Ax + Bu,\ x(t_0) = x_0 \\ y &= Cx, \end{aligned}$$

where A, B, C, D are real constant coefficient matrices of appropriate size. The state x, input u and output y are real-valued vector functions. Positive systems arise, when economical, biological or chemical systems are modelled by descriptor systems, in which the state x describes concentrations, populations of species, or numbers of cells. System positivity, then, refers to the restriction that the solution and the output have to be nonnegative vector functions for appropriately chosen initial state and input. It is well known that stability properties of standard systems, where $E = I$, are closely related to the spectral properties of the system matrix A. If the dynamics of the system, however, is described by an implicit differential equation, then stability properties are determined by the eigenvalues and eigenvectors associated with the matrix pencil $\lambda E - A$, or just the matrix pair (E, A). In the case of standard positive systems, stability properties are mainly determined by the well-known Perron-Frobenius theory for matrices and take a simple form compared to the unconstrained case. We present a new extension of the Perron-Frobenius theory to regular matrix pairs (E, A) and show that in the descriptor case it also determines the stability properties of positive systems.

On complex shifted Laplace preconditioners for the vector Helmholtz equation

C. Vuik

1 Introduction

In this paper, the time-harmonic wave equation in heterogeneous media is solved numerically. The underlying equation governs wave propagations and scattering phenomena arising in many area's, e.g. in aeronautics, geophysics, and optical problems. Recently, it appears that the newly developed preconditioner can also be applied to the vector wave equation.

Our applications are motivated by seismic research for the oil industry and a description of laser light in order to read and write optical storage disks. Another application is the determination of the the radar cross section (RCS) of a modern fighter aircraft. For the contribution of the jet engine air intake this is done by determination of the total electric field on the aperture. In order to do this the vector wave equation has to be solved, which is derived from the Maxwell equations.

We restrict ourselves first to the Helmholtz equation discretized by finite difference discretizations. Since the number of grid points per wavelength should be sufficiently large for accurate solutions, for very high wavenumbers the discrete problem becomes extremely large, thus prohibiting the use of direct solvers. However, since the coefficient matrix is sparse, iterative solvers are an interesting alternative.

In many geophysical applications that are of our interest, an unbounded domain is used. In our model we approximate such a domain by a bounded domain, where appropriate boundary conditions are used to prevent spurious reflections. As boundary conditions we compare the following possibilities: Dirichlet, Neumann, Sommerfeld, Absorbing Layer and Perfect Matched Layer. Due to the boundary conditions and damping in the heterogeneous medium, the coefficient matrix is complex-valued.

2 The shifted Laplace preconditioner

It appears that standard iterative solvers (ILU preconditioned Krylov solver, Multigrid, etc.) fail for the Helmholtz equation, if the wavenumber becomes sufficiently high. In this paper we present and analyze Krylov subspace solution methods combined with a novel preconditioner for high wavenumbers. The preconditioner is based on the inverse of an Helmholtz operator, where an artificial damping term is added to the operator. This preconditioner can be approximated by multigrid. This is somewhat surprising as multigrid, without enhancements, has convergence troubles for the original Helmholtz operator at high wavenumbers.

Since our publication [3], which describes this new multigrid based preconditioner for the Helmholtz equation, a number of other publications has appeared, which show that this preconditioner can be used for many applications, as there are geophysics [2], sonar [5], acoustic scattering [1], etc.

3 Spectral analysis

Currently, we are investigating the best choice of the damping term [4]. If the damping term is small Bi-CGSTAB converges fast, but it is difficult to use multigrid for the preconditioner. On the other hand, if the damping term is large the multigrid approximation is very good, but the convergence of Bi-CGSTAB is slow. So a compromise is required to obtain an efficient solver. To find a good value of the damping term we study the spectral properties of the preconditioned matrix. It appears that an eigenvalue analysis of this matrix can be used to predict the convergence of GMRES. In practice it appears that these insights can also be used for the convergence of Bi-CGSTAB. We conclude that Bi-CGSTAB combined with the novel preconditioner converges satisfactorily for all choices of the boundary conditions. In Table 1 we give some results comparing the iterations with the optimal shift with the minimum number of iterations.

Damping	Optimal shift	Iterations "optimal" shift	Min. number of iterations
$\nu = 0$	1	56	54
$\nu = 0.1$	1.005	42	41
$\nu = 0.5$	1.118	20	20
$\nu = 1$	1.4142	13	13

Table 6: Number of iterations for the "optimal" shift and minimum number of iterations

4 Vector Helmholtz problems

Recently, we also use a block version of our preconditioner for vector Helmholtz problems. It has been demonstrated that the shifted Laplace preconditioner can be used effectively to speed up the convergence rate of the GCR method for solution of the linear system resulting from higher order finite element discretization of the vector wave equation. However, the efficiency of the preconditioned Krylov method strongly depends on the amount of work involved in solution of the preconditioner system. Currently, a basic approach has been implemented that uses GCR to solve the preconditioner system, as opposed to the multigrid method advocated

in the original publication [3]. The current approach already gives satisfactory results for intermediate system sizes up, however to speed up the solver further, a combination with an algebraic multigrid method should be made.

References

[1] T. Airaksinen and E. Heikkola and A. Pennanen and J. Toivanen, An algebraic multigrid based shifted-Laplacian preconditioner for the Helmholtz equation, Journal of Computational Physics, 226, pp. 1196-1210, 2007.

[2] I. Duff and S. Gratton and X. Pinel and X. Vasseur, Multigrid based preconditioners for the numerical solution of two-dimensional heterogeneous problems in geophysics, International Journal of Computer Mathematics, 84, pp. 1167-1181, 2007.

[3] Y.A. Erlangga and C.W. Oosterlee and C. Vuik, A Novel Multigrid Based Preconditioner For Heterogeneous Helmholtz Problems, SIAM Journal on Scientific Computing, 27, pp. 1471-1492, 2006.

[4] M. B. van Gijzen, Y. A. Erlangga, and C. Vuik, Spectral Analysis of the Discrete Helmholtz Operator Preconditioned with a Shifted Laplacian, SIAM Journal on Scientific Computing, 29, pp. 1942-1958, 2007.

[5] E. Turkel and Y.A. Erlangga, Preconditining a finite element solver of the exterior Helmholtz equation, European Conference on Computational Fluid Dynamics, 2006.

Iterative Methods for Solving Partial Integro-Differential Equations Arising from Option Pricing

Iris R. Wang, Justin W.L. Wan, Peter Forsyth

Partial integro-differential equations (PIDEs) arise from a wide variety of applications which range from image deblurring to option pricing in finance. There has been active research in developing fast solvers for PIDEs. The challenge is how to deal with the fact that the discretization matrix is a sum of two matrices which correspond to the discretization of the differential term and the integral term. There are two main sources of difficulties. On the matrix side, the discrete differential operator matrix typically has a sparse structure whereas the discrete integral matrix generally is dense. On the other hand, the integral matrix may have a Toeplitz structure that the differential matrix does not have. On the operator side, the discrete differential and integral operators exhibit different characteristics. For instance, they have very different eigenvalues and eigenvectors compositions. Thus, it would be tricky to apply methods such as multigrid in which spectral properties are closely related to the effectiveness of the smoothers on the fine grid. Separately, fast solvers exist for PDE (e.g. multigrid, domain decomposition) and for integral equations (e.g. fast multipole, wavelets). But the fast PDE solvers are often not applicable or not optimal for integral equations, and vice versa. One top of these, since the matrix is dense, one would not construct it explicitly and hence standard preconditioning techniques such as ILU would not be applied.

The motivated application in this talk is the PIDE arising from pricing European options in computational finance. The unknown is the option value $V(S,\tau)$ which depends on the underlying asset price S and time τ. Based on the classic Black-Scholes model, the option value satisfies the following equation:

$$\begin{aligned}\frac{\partial V}{\partial \tau} &= \frac{\sigma^2}{2}S^2\frac{\partial^2 V}{\partial S^2} + rS\frac{\partial V}{\partial S} - rV \\ &\quad + \int_{-\infty}^{\infty} \nu(y)[V(Se^y,\tau) - V(S,\tau) - S(e^y-1)\frac{\partial V}{\partial S}]\, dy\end{aligned}$$

where σ is a constant and r is the interest rate. The integral term comes from the modelling of the underlying asset price by a Lévy process. It essentially describes how often a jump occurs in the asset price. The case where the jump process has infinite activity and infinite variation is of particular interest since the measure function $\nu(y)$ will be singular at the origin and hence standard discretization would not lead to an accurate solution. Note also that this integral is rather complicated and it has no particular structures such as circulant or Toeplitz.

In this talk, we discuss the issues of designing iterative methods for solving the linear systems arising from the above PIDE. For instance, how one may modify the equation to obtain a dense matrix which would possess a nicer matrix structure. We also present convergence analysis of some iterative methods, taking advantage of the M-matrix property. Numerical results are given to demonstrate the effectiveness of the methods. Finally, if time allows, we extend the framework to pricing American options which leads to a linear complementarity problem.

On Clustered Ritz Values in the Lanczos Method

Wolfgang Wülling

In [1], Zdeněk Strakoš and Anne Greenbaum posed two conjectures on the clustering of Ritz values occuring in finite precision computations. In this presentation I will briefly show that the conjectures are true and how to prove them.

Consider the Lanczos method for the real symmetric eigenvalue problem $Au = \lambda u$ in finite precision arithmetic. Due to rounding errors the Lanczos vectors may lose their orthogonality even for a small number of iterations. Consequently, it can be numerically observed that clusters of Ritz values are generated: multiple copies of Ritz values all approximating a single eigenvalue of A.

The first conjecture is as follows: Has any Ritz value in a tight, well separated cluster stabilized to within a small quantity δ which is proportional to the square root of the length of the cluster interval and the gap in the spectrum?

The second conjecture is concerned with the concept of *stabilization of weights*. Any nonzero vector together with the eigenvectors and eigenvalues of A defines a weight function and hence a Riemann-Stieltjes integral. The Lanczos method implicitly computes a sequence of Gauss quadrature approximations. This situation leads to the following question: If a cluster of Ritz values closely approximates an eigenvalue λ of A, does the sum of weights corresponding to the clustered Ritz values closely approximate the original weight belonging to the eigenvalue λ?

[1] Z. Strakoš and A. Greenbaum. *Open questions in the convergence analysis of the Lanczos process for the real symmetric eigenvalue problem.* IMA Research Report 934, IMA Workshop on Iterative Methods, University of Minnesota, Minneapolis, MN, 1992.

Perturbation of Purely Imaginary Eigenvalues of Hamiltonian Matrices under Structured Perturbations

Volker Mehrmann, Hongguo Xu

We discuss the perturbation theory for purely imaginary eigenvalues of Hamiltonian matrices under Hamiltonian and non-Hamiltonian perturbations. We show that there is a substantial difference in the behavior under these perturbations. The spectrum of a Hamiltonian matrix is symmetric about the imaginary axis in the complex plane. We show that with a certain sign characteristic condition, a semi-simple purely imaginary eigenvalue of a Hamiltonian matrix may be perturbed by a small Hamiltonian perturbation, but will stay on the imaginary axis. On the other hand, a small non-Hamiltonian perturbation may pull such an imaginary eigenvalue off the imaginary axis. We also discuss the perturbation of real eigenvalues of real skew-Hamiltonian matrices under structured perturbations. We use these results to analyze the properties of the symplectic URV method of computing the eigenvalues of Hamiltonian matrices.

On Some Known and Open Matrix Nearness Problems

Krystyna Ziętak

A survey of matrix nearness problems is given in [8]. Applications in areas including control theory, numerical analysis and statistics are also outlined in [8]. Difficulties of these problems first of all depend on selected norms which measure a distance of a given matrix from some given class of matrices. We use the unitarily invariant norms, in particular the c_p-Schatten norms and the spectral norm.

In this talk we develop the following matrix nearness problems:

- approximation of $A \in \mathcal{C}^{m\times n}$ by subunitary matrices with respect to any arbitrary unitarily invariant norm (see [3]),
- a minimal rank approximation of $A \in \mathcal{C}^{m\times n}$ with respect to the spectral norm (see [3]),
- approximation of $A \in \mathcal{C}^{n\times n}$ by matrices the spectrum of which is in a strip, with respect to the spectral norm (see [12]),
- strict spectral approximation of a matrix and some related problems (see [15], [16]).

We present properties of the nearest matrices and algorithms for computing them. We also discuss two open problems. The first problem is raised in [12] and the second one follows from [15], [16]. Now we shortly describe the above four problems.

A matrix $Q \in \mathcal{C}^{m\times n}, m \geq n$, is subunitary if and only if $QQ^HQ = Q$. We determine by means of the generalized polar decomposition of $A \in \mathcal{C}^{m\times n}$ matrices E solving for arbitrary unitarily invariant norm the following nearest problem

$$||A - E|| \leq ||A - Q|| \qquad \text{for all } Q \text{ such that} \tag{1}$$

(i) Q has orthonormal columns,

(ii) Q is subunitary of a given rank $r \leq n$,

(iii) Q is subunitary.

The presented in [3] characterizations of matrices E solving the three cases of the problem (1) generalize the results of Ky Fan and Hoffman [2], and Maher [13]. We propose algorithms for computing the subunitary approximant E, in which we use Gander's method [11] and Higham's method [6] for computing polar decomposition of a matrix (for algorithms for computing polar decomposition see also [9]). In the algorithm applying Gander's method we compute the unitary polar factor of A. However, in the algorithm based on Higham's method one computes the polar decomposition of A and the unitary polar factor of the shifted Hermitian polar factor of A.

A minimal rank approximation of $A \in \mathcal{C}^{m\times n}$ is a matrix B for which the minimum

$$\min_{B \text{ minimal rank}} ||A - B|| < \eta \tag{2}$$

is reached, where η is given. This problem was developed by Golub in [4] for the Frobenius norm. In (2) we choose the spectral norm. We show how to compute the minimal rank approximant B without computing the `SVD` of A. In the proposed algorithm we compute the unitary polar factor of the shifted Hermitian polar factor of A.

An approximation of a matrix by matrices with spectra in a given closed convex subspace of the complex plane was initiated by the Halmos paper [5] on an approximation of an operator by positive operators. In [12] we correct some theorem of Khalil and Maher [10] on approximation by matrices with spectra in a strip. In the corrected theorem we characterize solution of the following problem for the spectral norm:

$$\min_{X\in\mathcal{S}_a} ||A - X||, \tag{3}$$

where $\mathcal{S}_a$ is the set of all matrices of the form $X = X_1 + iX_2$, where X_1 is a Hermitian positive definite matrix and X_2 is a Hermitian matrix with the spectrum in $[0, a]$. For $a = 0$ the problem (3) is equivalent to the Halmos problem [5]. In the characterization of the solution of the problem (3) some matrix appers which is not determined explicitly. We disscus some conjecture which determines this matrix completely. Our conjecture is motived by numerical experiments in which we use Higham's algorithm [7] for computing a nearest real symmetric positive semidefinite matrix (Halmos approximant), which can be easily extended to the complex case.

In the last part of the talk we deal with our old conjecture. Let $\mathcal{M}$ be a linear subspace of $\mathcal{C}^{m\times n}$ and let $A \in \mathcal{C}^{m\times n}$. We consider the problem of finding a matrix B from $\mathcal{M}$ such that for the spectral norm we have

$$||A - B|| = \min_{X\in\mathcal{M}} ||A - X||. \tag{4}$$

In general case the matrix B in (4) is not unique. Therefore in [15] we have introduced a strict spectral approximation of A which is the best from the best and it is unique (see also [16]). The strict spectral approximation of a matrix is a generalization of the strict (Chebyshev) approximation of a vector, introduced by Rice in 1962. It is proven by Descloux [1] that the strict approximation of a real vector is the limit of the l_p-approximation when $p \to \infty$.

Since many years we have conjectured that the strict spectral approximation has an analogous property. A similar conjecture was formulated by Rogers and Ward [14] for approximation by positive approximants. In the talk we give some motivations for our conjecture.

The talk is based on [15], [16] and the joint papers [3], [12] with Beata Laszkiewicz.

References

[1] J. Descloux, Approximation in L^p and Chebyshev approximation *J. Soc. Ind. Appl. Math.*, 11:1017–1026 (1963).

[2] K. Fan, A. Hoffman, Some metric inequalities in the space of matrices, *Proc. Amer. Math. Soc.*, 6: 111–116 (1955).

[3] W. Gander, Algorithms for the polar decomposition, *SIAM Sci. Stat. Comput.*, 11: 1102–1115 (1990).

[4] G.H. Golub, Least squares, singular values and matrix approximation, *Applikace Matematiky*, 13 : 44–51 (1968).

[5] P.R. Halmos, Positive approximants, *Indiana Univ. Math. J.*, 21: 951–960 (1972).

[6] N.J. Higham, Computing the polar decomposition - with applications, *SIAM J. Sci. Stat. Comput.*, 7: 1160–1174 (1986).

[7] N.J. Higham, Computing a nearest symmetric positive semidefinite matrix, *Linear Algebra Appl.*, 103: 101–118 (1988).

[8] N.J. Higham, Matrix nearness problems and applications, in M.J.C. Gover, S. Barnett (eds), *Applications of Matrix Theory*, Oxford Univ. Press, 1989, 1–27.

[9] N.J. Higham, *Functions of Matrices: Theory and Computations*, Society for Industrial and Applied Mathematics, book to appear, Philadelphia 2008.

[10] R. Khalil, P. Maher, Spectral approximation in $L(H)$, *Numer. Funct. Anal. Optimiz.*, 21: 693–713 (2000).

[11] B. Laszkiewicz, K. Ziętak, Approximation of matrices and family of Gander methods for polar decomposition, *BIT Numer. Math.*, 46:345–366 (2006).

[12] B. Laszkiewicz, K. Ziętak, Approximation by matrices with restricted spectra, *Linear Algebra Appl.*, in press.

[13] P.J. Maher, Partially isometric approximation of positive operators, *Illinois J. Math.*, 33: 227–243 (1989).

[14] D.D. Rogers, J.D. Ward, C_p-minimal positive approximants, *Acta Sci. Math. (Szeged)*, 43: 109–115 (1981).

[15] K. Ziętak, Strict approximation of matrices, *SIAM J. Matrix Anal. Appl.*, 16:232–234 (1995).

[16] K. Ziętak, Strict spectral approximation of a matrix and some related problems, *Applicationes Mathematicae*, 24: 267–280 (1997).

Pierre-Antoine Absil	Département d'ingénierie mathématique Université catholique de Louvain Bâtiment Euler 1348 Louvain-la-Neuve, Belgium PA.Absil@uclouvain.be http://www.inma.ucl.ac.be/∼absil/	12
Athanasios C. Antoulas	Dept. of Electrical and Computer Eng. Rice University Houston, TX 77251, USA aca@rice.edu www.ece.rice.edu/~aca	14
Mario Arioli	Rutherford Appleton Laboratory Chilton, Didcot, Oxfordshire, OX11 0QX, UK m.arioli@rl.ac.uk	17
Owe Axelsson	Uppsala University Department of Information Technology Box 337 SE-75105 Uppsala, Sweden Owe.Axelsson@it.uu.se http://user.it.uu.se/∼owea	18
Zhaojun Bai	Department of Computer Science and Department of Mathematics University of California Davis, CA 95616, USA bai@cs.ucdavis.edu http://www.cs.ucdavis.edu/∼bai	18
Ulrike Baur	Fakultät für Mathematik Technische Universität Chemnitz Reichenhainer Straße 41 09107 Chemnitz, Germany baur@mathematik.tu-chemnitz.de http://www-user.tu-chemnitz.de/∼ubaur/	19
Peter Benner	Fakultät für Mathematik Technische Universität Chemnitz D-09107 Chemnitz, Germany benner@mathematik.tu-chemnitz.de http://www.tu-chemnitz.de/∼benner	21

Michele Benzi	Department of Mathematics and Computer Science Emory University 400 Dowman Dr., Atlanta, GA 30322, USA benzi@mathcs.emory.edu http://www.mathcs.emory.edu/∼benzi	23
Timo Betcke	The University of Manchester School of Mathematics Oxford Road, M13 9PL, Manchester, UK timo.betcke@maths.man.ac.uk http://www.maths.man.ac.uk/~tbetcke	26
Paolo Bientinesi	Department of Computer Science Duke University P.O.Box 90129 Durham, NC 27708-0129, USA pauldj@cs.utexas.edu http://www.cs.utexas.edu/users/pauldj	
David Bindel	Department of Mathematics, Courant Institute of Mathematical Sciences New York University, NY 10012, USA dbindel@cims.nyu.edu http://www.cims.nyu.edu/~dbindel	28
Åke Björck	Department of Mathematics Linköping University SE-581 83 Linköping, Sweden akbjo@math.liu.se	28
Matthias Bollhöfer	Inst. of Computational Mathematics TU Braunschweig, Germany m.bollhoefer@tu-bs.de http://www-public.tu-bs.de/∼bolle/	29
Shreemayee Bora	Department of Mathematics Indian Institute of Technology Guwahati 781039, Assam, INDIA shbora@iitg.ernet.in	30
Karen Braman	501 E. St Joseph St. Dept of Mathematics South Dakota School of Mines Rapid City, SD 57701, USA Karen.Braman@sdsmt.edu	33

Angelika Bunse-Gerstner	Zentrum für Technomathematik Fachbereich 3, Mathematik und Informatik Universität Bremen Postfach 330 440, D-28334 Bremen, FRG Bunse-Gerstner@math.uni-bremen.de http://www.math.uni-bremen.de/~angelika/	35
Xiao-Wen Chang	School of Computer Science McGill University 3480 University Street Montreal, Quebec, H3A 2A7, Canada chang@cs.mcgill.ca www.cs.mcgill.ca/∼chang	36
John M. Conroy	IDA Center for Computing Science 17100 Science Drive Bowie, MD 20708, USA conroy@super.org	38
Achiya Dax	Hydrological Service P.O.B. 36118 Jerusalem 91360, ISRAEL dax20@water.gov.il	40
Eric de Sturler	Department of Mathematics, 544 McBryde Hall, Virginia Tech Blacksburg, VA 24061, USA sturler@vt.edu http://www.math.vt.edu/people/sturler/	41
Steven Delvaux	Department of Computer Science Katholieke Universiteit Leuven Celestijnenlaan 200A B-3001 Leuven (Heverlee), Belgium steven.delvaux@cs.kuleuven.be	
James Demmel	737 SODA HALL Computer Science Devision University of California at Berkeley Berkeley, CA 94720-1776, USA demmel@cs.berkeley.edu http://www.cs.berkeley.edu/~demmel/	44

Inderjit S. Dhillon	The University of Texas at Austin Department of Computer Sciences 1 University Station C0500 Austin, TX 78712-0500, USA inderjit@cs.utexas.edu http://www.cs.utexas.edu/users/inderjit	44
Sue Dollar	Computational Science & Engineering Dept. Rutherford Appleton Laboratory Chilton, Oxfordshire, OX11 0QX, UK s.dollar@rl.ac.uk www.numerical.rl.ac.uk/people/hsd/	45
Marco Donatelli	Università dell'Insubria Dipartimento di Fisica e Matematica Via Valleggio 11, 22100 Como, Italy marco.donatelli@uninsubria.it http://scienze-como.uninsubria.it/mdonatelli/	45
Froilán M. Dopico	Department of Mathematics Universidad Carlos III de Madrid 28911 Leganés, Madrid, Spain dopico@math.uc3m.es	48
Petros Drineas	Computer Science Department Rensselaer Polytechnic Institute 110 8th Street, Troy, NY 12180, USA drinep@cs.rpi.edu http://www.cs.rpi.edu/~drinep/	50
Iain Duff	STFC Rutherford Appleton Laboratory Harwell Science and Innovation Campus Didcot OX11 0QX, UK i.s.duff@rl.ac.uk www.numerical.rl.ac.uk/people/isd/isd.html	52
Ioana Dumitriu	University of Washington Department of Mathematics BOX 354350, Seatle, WA 98195, USA dumitriu@math.washington.edu http://www.math.washington.edu/~dumitriu	54
Alan Edelmann	Massachusetts Institute of Technology 77 Massachusetts Avenue Cambridge, MA 02139, USA Email: edelman AT mit.edu	55

Lars Eldén	Department of Mathematics Linköping University SE-58183 Linkping, Sweden laeld@math.liu.se http://www.mai.liu.se/~laeld/	55
Howard C. Elman	Department of Computer Science University of Maryland College Park, MD 20742, USA elman@cs.umd.edu http://www.cs.umd.edu/∼elman/	57
Ludwig Elsner	Fakultät für Mathematik Postfach 100131, 33501 Bielefeld, Germany elsner@Mathematik.Uni-Bielefeld.DE	
Mark Embree	Computational and Applied Mathematics Rice University 6100 Main Street – MS 134 Houston, TX 77005–1892, USA embree@rice.edu http://www.caam.rice.edu/~embree	59
Yogi Ahmad Erlangga	Technische Universität Berlin Institut für Mathematik, MA 3-3 Straße des 17. Juni 136, 10623 Berlin, Germany erlangga@math.tu-berlin.de http://www.math.tu-berlin.de/~erlangga	60
Oliver Ernst	Institut für Numerische Mathematik und Optimierung TU Bergakademie Freiberg, 09596 Freiberg, Germany ernst@math.tu-freiberg.de http://www.math.tu-freiberg.de/~ernst/	63
Heike Faßbender	AG Numerik Institut Computational Mathematics, TU Braunschweig Pockelsstr. 14, 38106 Braunschweig, Germany h.fassbender@tu-bs.de http://www.tu-bs.de/∼hfassben	64
Melina A. Freitag	Department of Mathematical Sciences University of Bath Claverton Down, Bath BA2 7AY, United Kingdom m.freitag@maths.bath.ac.uk http://www.bath.ac.uk/~mamamf/	66

Andreas Frommer	Fachbereich Mathematik und Naturwissenschaften Universität Wuppertal 42097 Wuppertal, Germany frommer@math.uni-wuppertal.de www.math.uni-wuppertal.de/~frommer	67
Efstratios Gallopoulos	Computer Engineering and Informatics Department University of Patras, Greece stratis@ceid.upatras.gr	69
Martin J. Gander	Department of Mathematics, University of Geneva Rue du Lievre 2-4, CP 64 CH-1211 Geneva 4, Switzerland martin.gander@math.unige.ch http://www.unige.ch/~gander	70
Luc Giraud	ENSEEIHT-IRIT 2 rue Camichel, 31 000 Toulouse, France giraud@enseeiht.fr www.enseeiht.fr/~giraud	72
Robert Granat	Dept Computing Science and HPC2N Umeå University SE-90187 UMEÅ, Sweden granat@cs.umu.se http://www.cs.umu.se/∼granat/	73
Serge Gratton	CNES 18 Av. E. Belin, 31 401 Toulouse, France serge.gratton@cnes.fr www.cerfacs.fr/~gratton	76
Anne Greenbaum	Math Department, Box 354350 University of Washington Seattle, WA 98195, USA greenbau@math.washington.edu www.math.washington.edu/~greenbau	78
Chen Greif	Department of Computer Science The University of British Columbia 2366 Main Mall, Vancouver, BC V6T 1Z4, Canada greif@cs.ubc.ca http://www.cs.ubc.ca/~greif/	78

Luka Grubišić	Department of Mathematics University of Zagreb Bijenička 30, 10020 Zagreb, Croatia luka.grubisic@math.hr http://web.math.hr/~luka/	79
Serkan Gugercin	Dept. of Mathematics Virginia Tech. 460 McBryde Hall, Blacksburg, VA 24061-0123, USA gugercin@math.vt.edu http://www.math.vt.edu/people/gugercin/	80
Chun-Hua Guo	Department of Mathematics and Statistics University of Regina Regina, SK S4S 0A2, Canada chguo@math.uregina.ca http://www.math.uregina.ca/~chguo/	80
Martin H. Gutknecht	Seminar for Applied Mathematics ETH Zürich ETH-Zentrum, HG G68.2 CH-8092 Zürich, Switzerland martin.gutknecht@sam.math.ethz.ch http://www.sam.math.ethz.ch/∼mhg	
Nicholas J. Higham	School of Mathematics The University of Manchester Manchester M13 9PL, UK higham@ma.man.ac.uk http://www.ma.man.ac.uk/~higham/	83
Michiel Hochstenbach	Dept. of Mathematics and Computer Science TU Eindhoven PO Box 513, 5600 MB Eindhoven, The Netherlands m.e.hochstenbach tue.nl http://www.win.tue.nl/∼hochsten/	84
Thomas K. Huckle	Fakultät für Informatik, TU München Boltzmannstr. 3, 85748 Garching, Germany huckle@in.tum.de http://www5.in.tum.de/persons/huckle.html	85
Bruno Iannazzo	Dipartimento di Fisica e Matematica Via Valleggio 11, I-22100 Como, Italy bruno.iannazzo@uninsubria.it http://www.dm.unipi.it/~iannazzo/	86

Ilse Ipsen	Department of Mathematics North Carolina State University Raleigh, NC 27695-8205, USA ipsen@ncsu.edu	87
Elias Jarlebring	AG Numerik Institut Computational Mathematics TU Braunschweig Pockelsstr. 14, 38106 Braunschweig, Germany e.jarlebring@tu-bs.de http://www.icm.tu-bs.de/∼jarlebri/	
Charles Johnson	Mathematics Department College of William & Mary P.O. Box 8795, Williamsburg, VA 23187-8795 crjohnso@math.wm.edu http://as.wm.edu/Faculty/Johnson.html	88
Bo Kågström	Dept Computing Science and HPC2N Umeå University SE-90187 UMEÅ, Sweden bokg@cs.umu.se http://www.cs.umu.se/∼bokg/	88
Misha E. Kilmer	Department of Mathematics Tufts University 503 Boston Ave., Medford, MA 02155, USA misha.kilmer@tufts.edu http://www.tufts.edu/~mkilme01	90
Andrew Knyazev	Department of Mathematical Sciences University of Colorado at Denver and Health Sciences Center Downtown Denver Campus P.O. Box 173364, Campus Box 170, Denver, CO 80217-3364 andrew.knyazev@cudenver.edu http://math.cudenver.edu/~aknyazev/	91
Plamen Koev	NCSU 232B Harrelson, Raleigh NC 27695, USA pskoev@ncsu.edu www4.ncsu.edu/~pskoev	92

Christos Kravvaritis	Department of Mathematics University of Athens Panepistemiopolis 15784, Athens, Greece ckrav@math.uoa.gr http://ckravv.googlepages.com	92
Daniel Kressner	Seminar für angewandte Mathematik ETH Zürich, Switzerland kressner@math.ethz.ch http://www.math.ethz.ch/~kressner/	95
James V. Lambers	Department of Energy Resources Engineering Stanford University, 367 Panama St Rm 065 Stanford, CA 94305-2220, USA lambers@stanford.edu http://www.stanford.edu/~lambers/	96
Julien Langou	The University of Colorado at Denver and Health Sciences Center Denver, Colorado julien.langou@cudenver.edu http://www-math.cudenver.edu/~langou/	99
Lek-Heng Lim	Department of Mathematics University of California, Berkeley 873 Evans Hall Berkeley, CA 94720-3840, USA lekheng@math.berkeley.edu http://math.berkeley.edu/∼lekheng/	
Jörg Liesen	Institute of Mathematics, MA 4-5, TU Berlin Straße des 17. Juni 136, 10623 Berlin, Germany liesen@math.tu-berlin.de www.math.tu-berlin.de/∼liesen	100
Niloufer Mackey	Department of Mathematics Western Michigan University Kalamazoo, MI 49008, USA. nil.mackey@wmich.edu	101
D. Steven Mackey	Department of Mathematics Western Michigan University Kalamazoo, MI 49008, USA. steve.mackey@wmich.edu	104
Michael W. Mahoney	Yahoo! Research, 701 First Ave., Sunnyvale, CA 94089, USA mahoney@yahoo-inc.com http://www.cs.yale.edu/homes/mmahoney/	106

Ivan Markovsky	ISIS Research Group School of Electronics and Computer Science University of Southampton, Highfield Southampton, SO17 1BJ, UK im@ecs.soton.ac.uk http://users.ecs.soton.ac.uk/im/homepage.html	
Carla D. Martin	James Madison University Department of Mathematics and Statistics MSC 1911, Harrisonburg, VA 22807, USA carlam@math.jmu.edu http://www.math.jmu.edu/∼carlam	107
Nicola Mastronardi	Istituto per le Applicazioni delCalcolo "M.Picone", CNR via G. Amendola 122/D, I-70126 Bari, Italy n.mastronardi@ba.iac.cnr.it http://www.ba.cnr.it/~irmanm21/Welcome.html	108
Roy Mathias	School of Mathematics, University of Birmingham, Birmingham, B15 2TT, England mathias-AT-maths.bham.ac.uk	109
Karl Meerbergen	KU Leuven Department of Computer Science Celestijnenlaan 200A, 3001 Heverlee (Leuven), Belgium Karl.Meerbergen@cs.kuleuven.be http://www.cs.kuleuven.be/~karlm/	110
Christian Mehl	School of Mathematics University of Birmingham Edgbaston, Birmingham B15 2TT, UK mehl@maths.bham.ac.uk http://web.mat.bham.ac.uk/~mehlc/	111
Volker Mehrmann	TU Berlin, Institut für Mathematik, MA 4-5 Strasse des 17. Juni 136, 10623 Berlin, Germany mehrmann@math.tu-berlin.de http://www.math.tu-berlin.de/~mehrmann	112
Hermann Mena	Technische Universität Chemnitz Fakultät für Mathematik Chemnitz, Germany hermann.mena@mathematik.tu-chemnitz.de http://www-user.tu-chemnitz.de/~hmena/	113

Emre Mengi	University of California, San Diego 9500 Gilman Drive, La Jolla, CA 92093-0112 USA emengi@math.ucsd.edu http://www.math.ucsd.edu/∼emengi/	114
Gérard Meurant	CEA BP12, 91680 Bruyeres le Chatel, France gerard.meurant@gmail.com	117
Cleve Moler	The MathWorks, Inc. 3 Apple Hill Drive, Natick, MA 01760-2098 USA moler@mathworks.com	
Ronald B. Morgan	Baylor University, Department of Mathematics Waco, TX 76712-7328, USA Ronald_Morgan@baylor.edu www.baylor.edu/~Ronald_Morgan/	118
Julio Moro	Departamento de Matemáticas Universidad Carlos III de Madrid Avenida Universidad 30, 28911 Leganés, Spain jmoro@math.uc3m.es	120
Reinhard Nabben	TU Berlin, Institut für Mathematik, MA 3-3 Strasse des 17. Juni 136, 10623 Berlin, Germany nabben@math.tu-berlin.de http://www.math.tu-berlin.de/~nabben/	121
Raj Nadakuditi	Massachusetts Institute of Technology 77 Massachusetts Avenue, Cambridge, MA 02139, USA raj@math.mit.edu	
James G. Nagy	Math/CS Department, Emory University Atlanta, GA 30322, USA nagy@mathcs.emory.edu www.mathcs.emory.edu/∼nagy	122
Maya Neytcheva	Uppsala University Department of Information Technology Box 337, SE-75105 Uppsala, Sweden Maya.Neytcheva@it.uu.se http://user.it.uu.se/∼maya	124
Esmond G. Ng	Lawrence Berkeley National Laboratory One Cyclotron Road Mail Stop 50F1650, Berkeley, CA 94720-8139, USA EGNg@lbl.gov	127

Nancy K. Nichols	Department of Mathematics The University of Reading Box 220, Whiteknights, Reading Berks RG6 6AX, UK n.k.nichols@rdg.ac.uk	128
Dianne P. O'Leary	Computer Science Department University of Maryland College Park, MD 20742, USA oleary@cs.umd.edu http://www.cs.umd.edu/users/oleary	129
Michael L. Overton	Courant Institute of Mathematical Sciences 251 Mercer St., New York, NY 10012, USA overton@cs.nyu.edu http://www.cims.nyu.edu/	130
Archara Pacheenburawana	Department of Mathematics and Statistics Faculty of Science and Technology Thammasat University Prathum Thani, 12121 Thailand archara@mathstat.sci.tu.ac.th http://mathstat.sci.tu.ac.th/~archara/	130
Chris Paige	McGill University School of Computer Science 3480 University Street, Montreal H3A 2A7, Canada paige@cs.mcgill.ca http://www.cs.mcgill.ca/~chris/	131
Bob Plemmons	Wake Forest University Dept. Computer Science, Box 7311 Winston Salem, NC 27109, USA plemmons@wfu.edu http://www.wfu.edu/~plemmons/	132
Bor Plestenjak	Department of Mathematics University of Ljubljana Jadranska 19, SI-1000 Ljubljana, Slovenia bor.plestenjak@fmf.uni-lj.si http://www-lp.fmf.uni-lj.si/plestenjak/bor.htm	133
Marielba Rojas	Informatics and Mathematical Modelling Technical University of Denmark 2800 Lyngby, Denmark mr@imm.dtu.dk http://www.imm.dtu.dk/∼mr	135

Miro Rozložník	Institute of Computer Science, Academy of Sciences of the Czech Republic, Pod vodárenskou věží 2, CZ-182 07 Prague, Czech Republic miro@cs.cas.cz www.cs.cas.cz/~miro	136
Axel Ruhe	School of Computer Science and Communication (CSC) Royal Institute of Technology (KTH) SE-100 44 STOCKHOLM, Sweden ruhe@kth.se http://www.csc.kth.se/~ruhe/	139
Yousef Saad	University of Minnesota saad@cs.umn.edu http://www.cs.umn.edu/~saad/	140
Michael Saunders	Dept. of Management Science and Engineering Stanford University 380 Panama Way, Stanford, CA 94305-4026, USA saunders@stanford.edu http://www.stanford.edu/∼saunders	
Berkant Savas	Department of Mathematics Linköping University SE-58183 Linköping, Sweden besav@math.liu.se http://www.mai.liu.se/~besav/	140
Christian Schröder	TU Berlin, Institut für Mathematik, MA 4-5 Straße des 17. Juni 136, 10623 Berlin, Germany schroed@math.tu-berlin.de http://www.math.tu-berlin.de/∼schroed	142
Jennifer A. Scott	Computational Science and Engineering Department Rutherford Appleton Laboratory Oxfordshire OX11 0QX, England j.a.scott@rl.ac.uk http://www.cse.scitech.ac.uk/nag/members.shtml	143
Stefano Serra-Capizzano	Dipartimento di Fisica e Matematica Università dell'Insubria - Sede di Como, Via Valleggio 11, 22100 Como, Italy stefano.serrac@uninsubria.it	144

Valeria Simoncini	Dipartimento di Matematica Università di Bologna, Italy valeria@dm.unibo.it www.dm.unibo.it://~simoncin	145
Alicja Smoktunowicz	Faculty of Mathematics and Information Science University of Technology Plac Politechniki 1, 00-661 Warsaw, Poland smok@mini.pw.edu.pl http://www.mini.pw.edu.pl/~smok/	145
Danny C. Sorensen	Rice University Dept. of Computational & Applied Mathematics 6100 Main St., Houston, TX 77005-1892, USA sorensen@rice.edu caam.rice.edu/~sorensen	146
Alastair Spence	Department of Mathematical Sciences University of Bath Bath, BA2 7AY, United Kingdom a.spence@bath.ac.uk http://www.maths.bath.ac.uk/~as/	147
Andreas Stathopoulos	College of William and Mary Department of Computer Science P.O. Box 8795, Williamsburg, VA 23187-8795, USA andreas@cs.wm.edu html://www.cs.wm.edu/ andreas	148
G. W. (Pete) Stewart	Department of Computer Science University of Maryland College Park MD 20742, USA stewart@cs.umd.edu http://www.cs.umd.edu/∼stewart	150
Zdeněk Strakoš	Institute of Computer Science, Academy of Sciences of the Czech Republic, Pod vodárenskou věží 2, CZ-182 07 Prague, Czech Republic strakos@cs.cas.cz www.cs.cas.cz/~strakos	151
Gilbert Strang	Department of Mathematics, MIT Cambridge MA 02139, USA gs@math.mit.edu math.mit.edu/∼gs	154

Brian D. Sutton	Department of Mathematics Randolph-Macon College PO Box 5005, Ashland, VA 23005, USA briandavidsutton@gmail.com http://faculty.rmc.edu/bsutton/	154
Daniel B. Szyld	Department of Mathematics (038-16) Temple University, 1805 N Broad Street Philadelphia, PA 19122-6094, USA szyld@math.temple.edu http://www.math.temple.edu/szyld	155
Jurjen Duintjer Tebbens	Institute of Computer Science Academy of Sciences of the Czech Republic Pod Vodárenskou věží 2 182 07 Libeň - Praha 8 Czech Republic tebbens@cs.cas.cz www.cs.cas.cz/duintjertebbens	156
Petr Tichý	Institute of Computer Science Czech Academy of Sciences Pod vodárenskou věží 2 18207 Prague, Czech Republic tichy@cs.cas.cz http://www.cs.cas.cz/~tichy/	158
Françoise Tisseur	School of Mathematics The University of Manchester Manchester, M13 9PL, UK ftisseur@ma.man.ac.uk http://www.ma.man.ac.uk/~ftisseur/	160
David Titley-Peloquin	School of Computer Science McGill University, 3480 University Street Montreal, Quebec, H3A 2A7, Canada dtitle@cs.mcgill.ca www.cs.mcgill.ca/~dtitle	162
Lloyd N. Trefethen	Oxford University Computing Laboratory Wolfson Bldg., Parks Rd., Oxford OX1 3QD, UK LNT@comlab.ox.ac.uk http://www.comlab.ox.ac.uk/nick.trefethen/	165
Joel A. Tropp	Caltech/ACM 1200 E. California Blvd., Pasadena, CA 91125-5000, USA jtropp@acm.caltech.edu http://www.acm.caltech.edu/ jtropp	165

Miroslav Tůma	Institute of Computer Science Academy of Sciences of the Czech Republic Pod Vodárenskou věží 2, 182 07 Praha 8 - Libeň, Czech Republic tuma@cs.cas.cz http://www.cs.cas.cz/~tuma/	166
Eugene Tyrtyshnikov	Institute of Numerical Mathematics Russian Academy of Sciences Gibkin Street, 8, Moscow 119333, Russia tee@inm.ras.ru http://bach.inm.ras.ru/chair.htm	168
Marc Van Barel	Department of Computer Science, Katholieke Universiteit Leuven Belgium marc.vanbarel@cs.kuleuven.be http://www.cs.kuleuven.be/∼marc/	169
Paul Van Dooren	Université catholique de Louvain B-1348 Louvain-la-Neuve, Belgium	170
Martin van Gijzen	Faculty EWI, Delft University of Technology Mekelweg 4, 2628 CD Delft, The Netherlands M.B.vanGijzen@ewi.tudelft.nl http://ta.twi.tudelft.nl/nw/users/gijzen/	172
Sabine van Huffel	Department of Electrical Engineering (ESAT-SCD), Katholieke Universiteit Leuven, Kasteelpark Arenberg 10, B-3001 Leuven, Belgium sabine.vanhuffel@esat.kuleuven.be	174
Charles F. Van Loan	Department of Computer Science, 5153 Upson Hall, Cornell University Ithaca, New York 14853, USA cv@cs.cornell.edu	175
Raf Vandebril	Department of Computer Science Katholieke Universiteit Leuven Belgium raf.vandebril@cs.kuleuven.be http://www.cs.kuleuven.be/∼raf/	175
Andras Varga	DLR Oberpfaffenhofen Institute of Robotics and Mechatronics 82234 Wessling, Germany Andras.Varga@dlr.de http://www.robotic.dlr.de/Andras.Varga/	175

James Varah	Department of Computer Science The University of British Columbia 2366 Main Mall, Vancouver, BC V6T 1Z4, Canada jim@varah.com http://www.cs.ubc.ca/∼varah/	
Stephen A. Vavasis	Department of Combinatorics & Optimization University of Waterloo 200 University Ave. W., Waterloo, Ontario Canada N2L 3G1 vavasis@math.uwaterloo.ca http://www.math.uwaterloo.ca/∼vavasis	178
Elena Virnik	TU Berlin, Institut für Mathematik, MA 4-5 Straße des 17. Juni 136, 10623 Berlin, Germany virnik@math.tu-berlin.de www.math.tu-berlin.de/~virnik	179
Kees Vuik	Delft University of Technology Delft Institute of Applied Mathematics P.O. Box 5031, 2600 GA Delft, The Netherlands c.vuik@tudelft.nl http://ta.twi.tudelft.nl/nw/users/vuik/	180
Justin W.L. Wan	University of Waterloo Cheriton School of Computer Science 200 University Avenue West Waterloo, Ontario N2L 3G1, Canada jwlwan@uwaterloo.ca http://www.cs.uwaterloo.ca/~jwlwan/	182
Wolfgang Wülling	Basler Versicherungen Aeschengraben 21, 4002 Basel, Switzerland wolfgang.wuelling@baloise.ch	183
Hongguo Xu	Department of Mathematics, University of Kansas Lawrence, KS 66045, USA xu@math.ku.edu http://www.math.ku.edu/~xu/	184
Krystyna Ziętak	Wrocław University of Technology Inst. Math. Comput. Science Wrocław, Poland krystyna.zietak@pwr.wroc.pl http://www.im.pwr.wroc.pl/~zietak/	184