

BOOK OF ABSTRACTS

IWASEP 9

9th International Workshop on
Accurate Solution of Eigenvalue Problems

Napa Valley, June 4 - 7, 2012

Organizers:

Jesse Barlow (Chair), Zlatko Drmac, Volker Mehrmann, Ivan Slapničar,
Kresimir Veselic.

Local Organizing Committee:

Jim Demmel, Esmond G. Ng, Oded Schwartz.

Deflations Preserving Relative Accuracy

William Kahan

Deflation turns a matrix eigenproblem into two of smaller dimensions by annihilating a block of off-diagonal elements. When can this be done without changing the eigenvalues of an Hermitian matrix beyond each one's last significant digit or two no matter how widespread are eigenvalues' magnitudes? We seek practicable answers to this question, particularly for tridiagonals. Answers for bidiagonals' singular values were found by Ren-Cang Li in 1994.

Accurate Computation of Eigenvalues by an Analytical Level Set Method

Pavel Grinfeld

A numerical technique, based on an analytical level set method, is presented. The proposed technique yields eigenvalues for a broad range of smooth shapes with arbitrarily high accuracy. The technique applies to most common linear operators, including the Laplace operator which will be the focus of the presentation. The method is characterized by utmost simplicity, exponential convergence, ease of analysis and the ability to produce an exact solution to an approximate problem.

Orthogonal Functions and Inverse Eigenvalue Problems

Marc Van Barel

Orthogonal polynomials on the real line satisfy a three term recurrence relation. This relation can be written in matrix notation by using a tridiagonal matrix. Similarly, orthogonal polynomials on the unit circle satisfy a Szegő recurrence relation that corresponds to an (almost) unitary Hessenberg matrix. It turns out that orthogonal rational functions with prescribed poles satisfy a recurrence relation that corresponds to diagonal plus semiseparable matrices. This leads to efficient algorithms for computing the recurrence parameters for these orthogonal rational functions by solving corresponding inverse eigenvalue problems. In this talk we will study several of these connections between orthogonal functions and matrix computations and give some numerical examples illustrating the numerical behaviour of these algorithms.

Spectral Dichotomy for Toeplitz Matrices

Alexander Malyshev

The spectral dichotomy algorithms compute invariant (or reducing) subspaces of a matrix (or a pencil) corresponding to the eigenvalues inside a given contour. Quadratically convergent algorithms for the unit circle as a contour are based on the doubling idea: starting from a matrix pencil $\lambda B_0 - A_0$, the pencil $\lambda B_{j+1} - A_{j+1}$ is derived from $\lambda B_j - A_j$ by the formulas $A_{j+1} = X_j A_j$ and $B_{j+1} = Y_j B_j$, where the matrices X_j and Y_j satisfy $X_j B_j + Y_j A_j = 0$ and are such that the matrix $[X_j \ Y_j]$ is nonsingular.

Some applications, e.g., factorization of polynomials, may require A_0 and B_0 with the Toeplitz (or companion) structure. It is attractive to take advantage of the structure and devise faster variants of the doubling algorithms.

A fast implementation will be presented for the iteration, where $X_j = A_j(A_j + B_j)^{-1}$ and $Y_j = B_j(A_j + B_j)^{-1}$, which is an analog of the well-known iteration producing the matrix sign function. Results for other doubling iterations may be announced too.

On the Ritz Values of Normal Matrices

Zvonimir Bujanović

The implicitly restarted Arnoldi method (IRAM) introduced by Sorensen [2] is a well-known algorithm for computing a few eigenpairs of a large, generally non-symmetric sparse matrix. The eigenvalue approximations in this algorithm are usually taken as the Ritz or the harmonic values computed from a Krylov subspace.

The convergence of the IRAM has been a subject of intensive study. While Sorensen proved the convergence when the algorithm is used to compute the extreme eigenvalues of Hermitian matrices, the conditions for the convergence in the general case are still unknown. In particular, Embree [1] constructed a class of matrices for which the algorithm fails to converge, even in the exact arithmetic. A key property that ensures the failure is the non-normality of the example matrices. Furthermore, Tebbens et al. [3] show that an arbitrary convergence history for the Arnoldi method is possible.

In our talk, we discuss the convergence of IRAM for normal matrices. We demonstrate the difficulty in keeping the monotonicity of the Ritz values, which was essential for the convergence in the Hermitian case. A simple condition for a set of complex numbers to appear as Ritz values of a normal matrix is given: it is necessary and sufficient that a certain Cauchy matrix has a positive vector in its kernel. This fact is then used to explore the more complex geometry of Ritz and harmonic Ritz values in the normal case, which in turn has implications on the IRAM's convergence. We also present a variant of the Cauchy interlacing lemma.

References

- [1] Mark Embree: *The Arnoldi Eigenvalue Iteration with Exact Shifts Can Fail*, SIAM J. Matrix Anal. Appl., vol. 31, p. 1–10, 2009.
- [2] Danny C. Sorensen: *Implicit Application of Polynomial Filters in a k -Step Arnoldi Method*, SIAM J. Matrix Anal. Appl., vol. 13, p. 357–385, 1992.
- [3] Jurjen D. Tebbens and Gérard Meurant: *Any Ritz Value Behavior Is Possible for Arnoldi and for GMRES with Any Convergence Curve*, submitted, 2011.

Condition Numbers of Randomly Sampled Orthonormal Matrices

Ilse Ipsen

We consider two strategies for sampling rows from $m \times n$ matrices Q with orthonormal columns. The first strategy samples c rows with replacement, while the second one treats each row as an iid Bernoulli random variable, and samples it with probability $g=c/m$. We present different probabilistic bounds for the condition numbers of the sampled matrices and express them in terms of the coherence of Q . Numerical experiments confirm the accuracy of the bounds, even for small matrix dimensions. We also present algorithms to generate matrices with user-specified coherence, and apply the bounds to the solution of general, full-rank least squares problems with the randomized preconditioner from BLENDENPIK. This is joint work with Thomas Wentworth.

Computing Jordan Blocks of Irreducible Totally Nonnegative Matrices

Plamen Koev

In 2005 Fallat and Gekhtman fully characterized the Jordan Canonical Form of the irreducible totally nonnegative matrices. In particular, all nonzero eigenvalues are simple and the possible Jordan structures of the zero eigenvalues are well understood and described. Starting with the bidiagonal decomposition of these matrices, we present an algorithm for computing all the eigenvalues, including the Jordan blocks, to high relative accuracy in what we believe is the first example of Jordan structure being computed accurately in the presence of roundoff errors.

The Newton Polygon and Structured Eigenvalue Perturbation

Julio Moro

Abstract: The Newton polygon, an elementary geometric construction first devised by Sir Isaac Newton, has been often used in the context of perturbation theory as a tool for deriving explicit first-order eigenvalue perturbation expansions. On one hand, this usually gives useful information on the directions in which perturbed eigenvalues move, something which is crucial in several practical situations when eigenvalues need to be pushed in certain specific directions, or must be moved as fast as possible away from a critical (or dangerous) region by a perturbation, which is typically small. On the other hand, these asymptotic expansions often lead to sharp bounds on the condition number of eigenvalues.

Most of these results, however, are obtained for arbitrary, nonstructured perturbations. If the matrix or operator under study belongs to a specific class of structured operators, it makes sense to consider only perturbations having the same structure, thereby restricting the admissible Newton polygons. So far, it seems that the structures most amenable to such a structured perturbation analysis via the Newton polygon are those defined via indefinite scalar products for which structured canonical forms are available.

In this talk we will both review classic results for unstructured perturbation as well as explore the case of structured perturbations. Taking as a guide a specific example, involving zero eigenvalues of complex skew-symmetric matrices, we will illustrate the interplay between matrix structure and the Newton polygon.

This talk corresponds to joint work with Mara Pelez, Daniel Kressner, Jim Burke, Michael Overton, Rafik Alam, Shreemayee Bora, Michael Karow and Volker Mehrmann.

Accurate Eigenvalue Decomposition of Arrowhead Matrices and Applications

Nevena Jakovčević Stor

Ivan Slapničar

Jesse Barlow

We present a new, improved, algorithm for solving an eigenvalue problem of real symmetric arrowhead matrix. Under certain conditions the algorithm computes all eigenvalues and all components of the corresponding eigenvectors with high relative accuracy in $O(n^2)$ operations. The algorithm is based on shift-and-invert technique and limited use of double precision arithmetic when necessary. Each eigenvalue and the corresponding eigenvector can be computed separately, which makes the algorithm suitable for cases when only part of the spectrum is required and for parallel computing.

We also present perturbation theory, applications to Hermitian arrowhead matrices and symmetric tridiagonal matrices, diagonal-plus-rank-one matrices, and numerical examples.

Smooth Local Bases for Perturbed Invariant Subspaces

G. W. Stewart

Let A be a matrix of order n and let \mathcal{X} be an invariant subspace of A ; i.e., $A\mathcal{X} \subset \mathcal{X}$. Under appropriate conditions it can be shown that if $\tilde{A} = A + E$ is a sufficiently small perturbation of A then \tilde{A} has a unique subspace $\tilde{\mathcal{X}}$ that can be identified with \mathcal{X} in the sense that $\lim_{E \rightarrow 0} \tilde{\mathcal{X}} = \mathcal{X}$.

In applications it is not sufficient simply to know that $\tilde{\mathcal{X}}$ exists. Instead a basis for $\tilde{\mathcal{X}}$ is required, whose construction is not a trivial matter. Kato [2, §II.4], for example, works with a matrix $A(z)$ whose elements are analytic functions of the complex variable z and derives a differential equations whose solution $X(z)$ is a basis for the invariant subspace corresponding to the perturbation $A(z) = A(z_0) + [A(z) - A(z_0)]$. Chatelin [1, §2.9] assumes that $A(z) = A(0) + zH$ and produces bases as power series in z .

In [3, 4] the presenter assumed a general perturbation E and determined bases of the form $X + YP$, where X is an orthonormal basis for \mathcal{X} Y is an orthonormal basis for the orthogonal complement of \mathcal{X} , by solving a nonlinear equation for P . This procedure produces a cluster of bases surrounding X and it is natural to ask about their relation to one another. Are they continuous functions of E ? Differentiable functions? The main result of this work is that if E is small enough then the matrices $P(E)$ are Frechet differentiable with respect to E . An informative bound for $\|P'(E)\|$ is also derived. As an application the result is used to give an elementary proof that if $A(z)$ is analytic then $P[E(z)]$ is also analytic—a result that overlaps with Kato's.

References

- [1] F. Chatelin. *Eigenvalues of Matrices*. Wiley, New York, 1993.
- [2] T. Kato. *Perturbation Theory for Linear Operators*. Springer-Verlag, New York, 1966.
- [3] G. W. Stewart. Error bounds for approximate invariant subspaces of closed linear operators. *SIAM Journal on Numerical Analysis*, 8:796–808, 1971.
- [4] G. W. Stewart. Error and perturbation bounds for subspaces associated with certain eigenvalue problems. *SIAM Review*, 15:727–764, 1973.

Minimization Principle for Linear Response Eigenvalue Problem

Zhaojun Bai and Ren-Cang Li

We present two theoretical results for the linear response eigenvalue problem. The first result is a minimization principle for the sum of the smallest eigenvalues with the positive sign. The second result is Cauchy-like interlacing inequalities. Although the linear response eigenvalue problem is a nonsymmetric eigenvalue problem, these results mirror the well-known trace minimization principle and Cauchy's interlacing inequalities for the symmetric eigenvalue problem.

Efficient Linear Algebra Processing Inside a Database Management System

Carlos Garcia-Alvarado, Carlos Ordonez, Naveen Mohanam
Edgar Martinez-Encarnacion, Predrag T. Tasic

Linear algebra algorithms are of utmost importance in data mining. It has been a common practice to export data sets outside of the database management system (DBMS) for data mining processing. However, this creates a considerable performance bottleneck that is especially restrictive for large data sets. The integration of linear algebra algorithms into a DBMS is a difficult task due to the DBMS's relational model foundation, system architecture, and lack of support for native array and matrix processing, which make the implementation very challenging and hinder efficient performance.

As a consequence, all linear algebra algorithms, starting with the most basic routines such as matrix multiplication, have to be reworked for accuracy and efficient performance inside a DBMS. In our research [1, 2], we present different efficient algorithms that exploit multiple data set layouts, SQL queries, and database extensibility mechanisms (e.g. User-Defined Functions) for solving fast linear regression and singular value decomposition (SVD) inside a DBMS. These algorithms are incorporated into a database system through predefined routines (stored procedures) that can directly analyze relational tables to compute statistical models and store the resulting models as relational tables, as well. The proposed optimizations include table layouts for storing matrices (vertical and horizontal layouts), computation of one-pass sufficient statistics, parallelization of tasks through SQL queries and User-Defined Functions, and efficient memory management. Preliminary work has also been done for integrating well-known external numerical libraries such as LAPACK and Intel's Math Kernel Library. Last but not least, we demonstrate that our algorithms can analyze large and high-dimensional data sets faster than external data mining tools such as R.

References

- [1] C. Ordonez. Statistical model computation with UDFs. In *IEEE Transactions on Knowledge and Data Engineering (TKDE)*, 22(12):1752-1765, 2010.
- [2] C. Ordonez and C. Garcia-Alvarado. A Data Mining System Based on SQL Queries and UDFs for Relational Databases. In *Proc. ACM CIKM*, pages 2521–2524, 2011.

Local Convergence of Inexact Newton-like Methods for Nonlinear Algebraic Eigenvalue Problems

Daniel Szyld and Fei Xue

Nonlinear algebraic eigenvalue problems of the form $T(\lambda)v = 0$ arise naturally in a variety of science and engineering applications. This talk concerns the local convergence rates of several inexact Newton-like algorithms for the solution of a simple eigenpair of the nonlinear eigenvalue problem. We show how the tolerances for the approximate solution of the linear systems arising in these algorithms affect the local convergence rates. In particular, with appropriately chosen tolerances for the inner solves, the inexact algorithms can achieve the same order of convergence as the exact methods. The analysis is illustrated by numerical experiments.

Comparative Survey of Sparse Eigensolvers and Effect of Initialization on Problems
with Low Rank Perturbations of Small Norm

Poster Submission

Tim Mitchell

Michael L. Overton

Josef Sifuentes

In applications and algorithms for stability analysis and control, it is often necessary to solve many related sparse eigenvalue problems in each of which at most only a few eigenvalues are of interest. In particular, the recent algorithms for approximating the pseudospectral abscissa and radius of linear dynamical systems [Guglielmi and Overton] and correspondingly, the H_∞ norm for systems with input and output [Guglielmi, Gürbüzbalaban, and Overton], require only computing the eigenvalue of largest real part or modulus (along with corresponding left and right eigenvectors) for a particular perturbation of the input matrix at every iteration. Though these methods only guarantee a lower bound approximation, they often do compute an exact solution in practice and furthermore are fast when the systems are large and sparse, unlike previous exact methods. By computing a sequence of rank-1 updates to the fixed input matrix, the new methods enable each eigenproblem to be efficiently solved using sparse eigensolver codes, whose main computational cost consists of matrix-vector products. As these perturbation sequences are not only rank-1 but also of low norm and converging, it is natural to ask whether information from a previously solved eigenproblem can effectively accelerate a subsequent one in the iteration. While the implicitly restarted Arnoldi method [Sorensen] is one of the most efficient and robust sparse algorithms available (and specifically, its implementation ARPACK), it may only be initialized with a single starting vector and modifying the algorithm to accept a recycled starting space for general problems appears to be difficult. On the other hand, the Jacobi-Davidson QR/QZ algorithms introduced by Fokkema, Sleijpen and Van der Vorst can both be started with an initial space as well as make use of a preconditioner, though it is not clear how competitive these features are in practice. We present a comparative survey of IRAM and JDQR for solving perturbed problems given information from the unperturbed problem's solution, using a variety of vectors and spaces, as well as preconditioning, as input parameters. Both random matrices and highly non-normal matrices from EigTool are considered, augmented with random rank-2 perturbations and structured rank-1 perturbations of small norm.

Multicore and the Challenges of the Nonsymmetric Eigenvalue Problem

Matthew Nability

The computational challenges of solving the nonsymmetric eigenvalue problem on multicore processors are significant. We survey critical issues of multi-core architectures and programming models. To achieve both high performance and portability across a range of new architectures, several common issues arise. We present recent work on computational kernels integral to several approaches to the nonsymmetric eigenvalue problem.

A modified dqds algorithm

Sheng-Guo Li, Ming Gu, Beresford N. Parlett

We present some new deflation strategies and some modified shift strategies for dqds algorithm. One of the deflation strategies is tailored for some matrices which are currently found difficult for LAPACK code, xLASQ. Through many numerical tests we show each of these techniques helps for most of the matrices. These techniques together make our implementation can be up to 11x faster for the difficult matrices and 20% faster in general.

Automating DEIM for Nonlinear Model Reduction

Danny Sorensen

Recently, we developed a new method for nonlinear model reduction that is based upon proper orthogonal decomposition (POD) combined with discrete empirical interpolation (DEIM). This POD-DEIM approach has provided spectacular dimension and complexity reduction for challenging systems large scale ordinary differential equations (ODEs). Reductions from 15,000 variables to 40 variables in the reduced model with very little loss of accuracy have been achieved. Several examples will be shown.

The DEIM is surprisingly simple and amounts to replacing orthogonal projection with an interpolatory projection of the nonlinear term that only requires the evaluation of a few selected components of the nonlinear term. This is the crux of the improvement over POD and it has potential for being automated as well. The technology used for automatic differentiation can be adapted to enable the automatic generation of a code for the reduced order function and its Jacobian from the code for the original high dimensional problem.

A Residual Replacement Strategy for Improving the Maximum Attainable Accuracy of Communication-Avoiding Krylov Subspace Methods

Erin Carson and James Demmel

Krylov subspace methods (KSMs) are a class of iterative algorithms commonly used to solve linear systems. In each iteration, the recursively computed approximate solution x^k and corresponding residual $r^k = b - Ax^k$ are updated. These updates can be written in the form

$$x^k = x^{k-1} + \alpha_{k-1}p^{k-1}, \quad r^k = r^{k-1} - \alpha_{k-1}Ap^{k-1} \quad (1)$$

or something similar. This encompasses algorithms such as Conjugate Gradient (CG), steepest descent, Biconjugate Gradient (BICG), Conjugate Gradient Squared (CGS), and Stabilized Biconjugate Gradient (BICGSTAB).

It is important to notice that x^k and r^k have different round-off patterns in finite precision arithmetic. Therefore, computational errors made in x^k are not self-correcting. Throughout the iteration, these errors accumulate and cause deviation of the true residual, $b - Ax^k$, and computed residual, r^k . This limits the *maximum attainable accuracy*, which indicates how accurately we can solve the system on a computer with machine precision ϵ . When the algorithm reaches this maximum attainable accuracy, the computed residual will appear to continue decreasing in norm, whereas the norm of the true residual stagnates. This can lead to a very large error in the solution despite the algorithm reporting a very small residual norm.

This motivates the use of strategies such as residual replacement, which limit the deviation of the two residuals by replacing the computed residual by the true residual at select iterations. (see, e.g., [13, 17]). In this way, errors which occur after a replacement do not depend on errors accumulated before the replacement. While performing a residual replacement in each iteration would maintain agreement between the two residuals, it may cause deterioration in the convergence of r^k due to large perturbations to the Lanczos recurrence relation for the residual vectors [13, 17]. Residual replacement strategies must then carefully select iterations where residual replacement takes place to satisfy two objectives: 1) the deviation of the computed and true residual must remain small, and 2) the Lanczos recurrence for the computed residuals is sufficiently maintained. Determining iterations where both requirements are met requires keeping track of an estimate of the accrued rounding error in the algorithm. Van der Vorst and Ye have successfully implemented such a strategy for standard Krylov methods [17]. There are many other analyses of the behavior of various KSMs in finite precision arithmetic (see, e.g., [6, 9, 10, 16, 11]).

The computation that occurs in each iteration in standard Krylov methods, namely, the updates of x^k and r^k , consist of one or more sparse matrix-vector multiply (SpMV) and vector operations. Because there are dependencies between iterations in standard Krylov methods and the main kernels in each iteration have low computation to communication ratios, standard Krylov method implementations are communication-bound on modern computer architectures. This motivated s -step, or, Communication-Avoiding KSMs (CA-KSMs), which are equivalent

to the standard KSM implementations in exact arithmetic. These variants use blocking strategies to perform s computation steps of the algorithm for each communication step, allowing an $O(s)$ reduction in total communication cost (see, e.g., [1, 2, 4, 5, 7, 8, 12, 14, 15]). Despite the potential performance benefits of CA-KSMs, the finite precision error in these variants grows with s , an obstacle to their use in practice. This means the deviation of the true and computed residual observed in standard KSMs is potentially worse for CA-KSMs, further decreasing the accuracy of the method. Although many previous authors have observed this behavior in CA-KSMs (see, e.g., [2, 3, 8, 12, 18]), we are the first, to our knowledge, to provide a quantitative analysis of round-off error in these algorithms which limits the maximum attainable accuracy.

Following the work of Van der Vorst and Ye , we bound the deviation of the true and computed residual in finite precision CA-KSMs, which leads to an implicit residual replacement strategy. We show how to implement our strategy without affecting the asymptotic communication or computation cost of the algorithm. Numerical experiments demonstrate the effectiveness of our residual replacement strategy for both CA-CG and CA-BICG. Specifically, it is shown that accuracy of order $O(\epsilon)\|A\|\|x\|$ can be achieved with a small number of residual replacement steps for an appropriately chosen polynomial basis, which demonstrates the potential for practical use of CA-KSMs.

References

- [1] E. Carson, N. Knight, and J. Demmel, *Avoiding communication in two-sided Krylov subspace methods*, Tech. report, Tech. Report, University of California at Berkeley, Berkeley, CA, USA, 2011.
- [2] A.T. Chronopoulos and C.W. Gear, *On the efficient implementation of preconditioned s -step conjugate gradient methods on multiprocessors with memory hierarchy*, Parallel computing **11** (1989), no. 1, 37–53.
- [3] A.T. Chronopoulos and C.D. Swanson, *Parallel iterative s -step methods for unsymmetric linear systems*, Parallel Computing **22** (1996), no. 5, 623–641.
- [4] J. Demmel, M. Hoemmen, M. Mohiyuddin, and K. Yelick, *Avoiding communication in computing Krylov subspaces*, Tech. report, Technical Report UCB/EECS-2007-123, University of California Berkeley EECS, 2007.
- [5] D.B. Gannon and J. Van Rosendale, *On the impact of communication complexity on the design of parallel numerical algorithms*, Computers, IEEE Transactions on **100** (1984), no. 12, 1180–1194.
- [6] A. Greenbaum, *Estimating the attainable accuracy of recursively computed residual methods*, SIAM Journal on Matrix Analysis and Applications **18** (1997), 535.

- [7] A.C. Hindmarsh and H.F. Walker, *Note on a Householder implementation of the GMRES method*, Tech. report, Lawrence Livermore National Lab., CA (USA); Utah State Univ., Logan (USA). Dept. of Mathematics, 1986.
- [8] M. Hoemmen, *Communication-avoiding Krylov subspace methods*, Thesis. UC Berkeley, Department of Computer Science (2010).
- [9] G. Meurant and Z. Strakos, *The Lanczos and conjugate gradient algorithms in finite precision arithmetic*, Acta Numerica **15** (2006), 471–542.
- [10] G.A. Meurant, *The Lanczos and conjugate gradient algorithms: from theory to finite precision computations*, vol. 19, Society for Industrial Mathematics, 2006.
- [11] C.C. Paige, M. Rozioznik, and Z. Strakos, *Modified Gram-Schmidt (MGS), least squares, and backward stability of MGS-GMRES*, SIAM Journal on Matrix Analysis and Applications **28** (2006), no. 1, 264.
- [12] J.V. Rosendale, *Minimizing inner product data dependencies in conjugate gradient iteration*, (1983).
- [13] G.L.G. Sleijpen and H.A. van der Vorst, *Reliable updated residuals in hybrid Bi-CG methods*, Computing **56** (1996), no. 2, 141–163.
- [14] E. Sturler, *A performance model for Krylov subspace methods on mesh-based parallel computers*, Parallel Computing **22** (1996), no. 1, 57–74.
- [15] S.A. Toledo, *Quantitative performance modeling of scientific computations and creating locality in numerical algorithms*, Ph.D. thesis, Massachusetts Institute of Technology, 1995.
- [16] C.H. Tong and Q. Ye, *Analysis of the finite precision bi-conjugate gradient algorithm for nonsymmetric linear systems*, Math. Comp, Citeseer, 1995.
- [17] H.A. Van der Vorst and Q. Ye, *Residual replacement strategies for Krylov subspace iterative methods for the convergence of true residuals*, SIAM Journal on Scientific Computing **22** (1999), no. 3, 835–852.
- [18] H.F. Walker, *Implementation of the GMRES method using Householder transformations*, SIAM Journal on Scientific and Statistical Computing **9** (1988), 152.

Subspace Iteration with Random Start Matrix

Ming Gu

The power method and subspace iteration method can be used to find a few largest eigenvalues or singular values. It is well-known that their convergence rate critically depends on the separation of the eigenvalues or singular values as well as the start matrix. In this talk, we develop new convergence results for these methods.

On Sensitivity and Perturbation analysis of Eigenproblems

Rafikul Alam

Nonlinear eigenvalue problems such as holomorphic and rational eigenproblems occur in many applications. As such these eigenproblems can be represented in several ways and each such representation plays an important role in influencing the conditioning of the eigenproblem. We discuss conditioning of eigenvalues of holomorphic and rational eigenproblems relative to a given representation. Specially, we discuss conditioning of a rational eigenproblem when represented as a minimal realization of rational matrix function and then discuss the effect of linearization on the conditioning of the eigenvalues. Various notions of separations of matrices and matrix pencils play an important role in perturbation analysis of eigenproblems. We revisit the issues of separations of matrix pencils and discuss solutions of some associated distance problems. We also discuss some local and global perturbation bounds for approximate eigenvalues.

References

- [1] S. S. AHMAD AND R. ALAM, *Sensitivity and perturbation analysis of nonlinear eigenproblems*, preprint
- [2] R. A. SMITH, *The condition numbers of the eigenvalue problem*, Numer. Math., 10(1967), pp.232-240.
- [3] G. W. STEWART AND J. SUN, *Matrix Perturbation Theory*, Academic Press, 1990.
- [4] F. TISSEUR, *Backward error and condition of polynomial eigenvalue problem*, Linear Algebra Appl., 309 (2000), pp. 339–361.
- [5] J. H. WILKINSON, *The Algebraic Eigenvalue Problem*, Oxford University press, 1965.

Accurate SVD of Cauchy–type matrices and applications
Zlatko Drmač

Our current work is focused on the development of efficient and reliable numerical algorithms for model order reduction in the \mathcal{H}_2 norm. As parts of different computational tasks, accurate SVD and eigenvalue decompositions of various matrices are required. Interestingly, in many situations, the key matrices have rich structure. In particular, related to various forms of rational approximations in Hardy spaces, the matrices are Cauchy–like, e.g. the Pick and Löwner matrices, paired Cauchy or Vandermonde matrices etc. For example, least squares approximation by rational functions is based on the matrix LS problem, where the coefficient matrix is $A = (\mathcal{C}_1, \mathcal{C}_2)$ with paired generalized rectangular Cauchy matrices $\mathcal{C}_1, \mathcal{C}_2$. This motivates development of new algorithms for accurate SVD and eigenvalue computations with various generalizations of classical matrices of Cauchy type. Our perturbation theory reveals condition numbers that determine how many digits of accuracy are feasible in floating point computation. We will present some of the new methods capable of achieving the accuracy exactly to the level determined by the perturbation theory.

Optimal Parameters for Damping Optimization in Linear Vibrating Systems

Zoran Tomljanović

A mathematical model of a linear vibrational system will be considered. This model is described with second-order differential equation

$$M\ddot{x} + D\dot{x} + Kx = 0,$$

where the mass matrix M and the stiffness matrix K are positive definite matrices of order n . The damping matrix is $D = C_u + C$, where C_u is positive semidefinite matrix which represents the internal damping and the positive semidefinite matrix C represents the external damping.

The damping parameters of interest are positions of dampers, their viscosities and modal damping parameters.

The problem of determination of optimal damping parameters can be reformulated as the minimization of the trace of the solution of the Lyapunov equation

$$AX + XA^T = -GG^T,$$

where A is the $2n \times 2n$ matrix obtained from linearization. The matrix G depends on the part of the undamped eigenfrequencies to be damped. If $G = I$ then all undamped eigenfrequencies have to be damped, but usually, G has small rank which corresponds with the damping of the part of the undamped spectrum.

Finding the optimal damping positions and viscosities, such that $\text{trace}X(C)$ is minimal, is a very demanding numerical problem, caused by the large number of trace calculations which can be required for larger matrix dimensions.

We propose a different techniques which significantly accelerate the optimization process.

On the other hand the optimal modal damping parameters have been derived as a solutions of certain systems of equations depending on damping models.

These results have been obtained partly in joint work with: Peter Benner (Max Planck Institute for Dynamics of Complex Technical Systems, Magdeburg, Germany), Ivana Kuzmanović, (Department of Mathematics, J.J. Strossmayer University of Osijek, Osijek, Croatia), Ninoslav Truhar (Department of Mathematics, J.J. Strossmayer University of Osijek, Osijek, Croatia)

List of publications:

- P. Benner, Z. Tomljanović, N. Truhar; Optimal Damping of Selected Eigenfrequencies Using Dimension Reduction. Accepted for publication in Numerical linear algebra with applications. (2011) DOI: 10.1002/nla.833
- I. Kuzmanović, Z. Tomljanović, N. Truhar; Optimization of material with modal damping. Applied mathematics and computation. 218 (2012) pp. 7326-7338
- Z. Tomljanović, N. Truhar, K. Veselić; Optimizing a damped system - a case study. International journal of computer mathematics; 11 March 2011 (iFirst) (2011) ; 1-13
- P. Benner, Z. Tomljanović, N. Truhar; Dimension reduction for damping optimization in linear vibrating system; Journal of Applied Mathematics and Mechanics (ZAMM), 91 (2011), 3; 179-191; DOI: 10.1002/zamm.201000077
- N. Truhar, Z. Tomljanović, R.-C. Li; Analysis of the solution of the Sylvester equation using low-rank ADI with exact shifts; Systems and Control Letters, 59 (2010); 248-257. IF=1.797
- N. Truhar , Z. Tomljanović; Estimation of optimal damping for mechanical vibrating systems; International Journal of Applied Mathematics and Mechanics, 5 (2009) , 5; 14-26

Stability Issues of Large-Scale Fast Matrix Multiplication

Grey Ballard
James Demmel
Olga Holtz
Benjamin Lipshitz
Oded Schwartz

Parallel matrix multiplication algorithms based on the classical ($O(n^3)$) algorithm are typically preferred over theoretically faster algorithms for matrix multiplication (such as Strassen's) for two main reasons. The first concern is that the asymptotic advantage is irrelevant for matrix sizes that are used in practice. The second concern is stability. We address these two concerns and suggest that parallel fast matrix multiplication algorithms should be used in practice.

The performance concern is dealt with in our recent paper [1], where we show a communication avoiding parallel fast matrix multiplication that is faster than any classical parallel matrix multiplication algorithms, both in theory (asymptotically) and in practice, for large matrices and for small matrices when many processors are used.

The stability concern is considered by Higham [2] and by Demmel et al. [3]. In this talk we provide experimental tests of the diagonal scaling technique suggested in [3], showing that their stabilizing technique works not just in theory but also in practice.

References

- [1] G. Ballard, J. Demmel, O. Holtz, B. Lipshitz, and O. Schwartz. Communication-optimal parallel algorithm for Strassen's matrix multiplication. Technical Report EECS-2012-32, UC Berkeley, March 2012. To appear in SPAA 2012.
- [2] N. J. Higham. *Accuracy and Stability of Numerical Algorithms*. SIAM, Philadelphia, PA, 2nd edition, 2002.
- [3] J. Demmel, I. Dumitriu, and O. Holtz. Fast linear algebra is stable. *Numerische Mathematik*, 108(1):59–91, 2007.

Communication-Avoiding Non-symmetric Eigensolver using Spectral Divide & Conquer

Grey Ballard
James Demmel
Ioana Dumitriu

Algorithms have two kinds of costs: arithmetic and communication. Communication means moving data, either between levels of a memory hierarchy, or between processors over a network. The cost of communication can greatly exceed the cost of arithmetic, so we seek algorithms that minimize communication. We present parallel and sequential algorithms for the non-symmetric eigenvalue problem that do asymptotically less communication than conventional algorithms, and in fact attain communication lower bounds presented in [1]. We present the algorithms and analyze their convergence properties and communication costs. Our algorithms use randomization in two ways; we know of no numerically stable, communication-optimal non-symmetric eigenvalue algorithm that does not use randomization. The same algorithm extends to the generalized non-symmetric eigenvalue problem, symmetric eigenvalue problem, and singular value decomposition, again doing asymptotically less communication than conventional algorithms, and attaining communication lower bounds. In the talk we will focus in particular on how randomization is used within the algorithms and also on comparisons to standard algorithms. This work is currently available as a technical report [2].

References

- [1] BALLARD, G., DEMMEL, J., HOLTZ, O., AND SCHWARTZ, O. Minimizing communication in numerical linear algebra. *SIAM J. Matrix Analysis Applications* 32, 3 (2011), 866–901.
- [2] BALLARD, G., DEMMEL, J., AND DUMITRIU, I. Communication-optimal parallel and sequential eigenvalue and singular value algorithms. EECS Technical Report EECS-2011-14, UC Berkeley. Feb 2011.

Accurate eigenvalues of totally nonnegative matrices
by means of the discrete hungry Toda equation

Akiko Fukuda
Yusaku Yamamoto
Emiko Ishiwata
Masashi Iwasaki
Yoshimasa Nakamura

The recursion formula of the qd algorithm is known as being just equal to the integrable discrete Toda equation. In this talk, we consider the discrete hungry Toda (dhToda) equation [1],

$$Q_k^{(n+M)} = Q_k^{(n)} + E_k^{(n)} - E_{k-1}^{(n+1)}, \quad k = 1, 2, \dots, m, \quad (2)$$

$$E_k^{(n+1)} = \frac{Q_{k+1}^{(n)}}{Q_k^{(n+M)}} E_k^{(n)}, \quad k = 1, 2, \dots, m-1, \quad E_0^{(n)} := 0, \quad (3)$$

which is a generalization of the discrete Toda equation with a positive integer M . We show in [2] that the dhToda equation is applicable for computing eigenvalues of a class of totally nonnegative (TN) matrices, which are matrices with all minors nonnegative. The algorithm is named the dhToda algorithm, and can be regarded as an extension of the dqd algorithm for computing the eigenvalues of a symmetric tridiagonal matrix. The band width of TN matrix corresponds to the parameter M . The dhToda equation gives a sequence of LR transformations. The differential form is presented for numerical stability, and the shift of origin is introduced for convergence acceleration [3]. Global convergence of the differential shifted dhToda algorithm is also clarified. Error analysis guarantees that the dhToda algorithm computes all the eigenvalues with high relative accuracy. Thus the dhToda algorithm inherits some properties of the dqds algorithm.

References

- [1] T. Tokihiro, A. Nagai and J. Satsuma, Proof of solitonical nature of box and ball systems by means of inverse ultra-discretization, *Inverse Problems*, 15 (1999), 1639–1662.
- [2] A. Fukuda, E. Ishiwata, Y. Yamamoto, M. Iwasaki and Y. Nakamura, Integrable discrete hungry systems and their related matrix eigenvalues, *Annal. Mat. Pura Appl.*, DOI 10.1007/s10231-011-0231-0.
- [3] A. Fukuda, Y. Yamamoto, M. Iwasaki, E. Ishiwata and Y. Nakamura, On a shifted LR transformation derived from the discrete hungry Toda equation, *Monat. Math.*, DOI 10.1007/s00605-012-0404-y.

dqds and Quasiseparable Matrices

Beresford Parlett

We begin by introducing the qds and dqds algorithms, discussing the difference between them and showing how they may be extended (as eigensolvers) beyond the tridiagonal case. As an example we present recent work by P.Zhlobich that extends qds to all quasiseparable (qs) matrices and dqds to those that are also Hessenberg. The latter class includes companion matrices and the former includes CMV and Green's matrices.

Stability theorem of the dqds with aggressive deflation for singular values

Kensuke Aishima, Yuji Nakatsukasa, Ichitaro Yamazaki

The dqds algorithm is the standard method for computing all the singular values of a bidiagonal matrix with high relative accuracy [2]. Its efficient implementation is now available as a LAPACK subroutine DLASQ [4]. Recently, in order to reduce the dqds runtime, we incorporated into dqds a technique called aggressive deflation, which had been applied successfully to the Hessenberg QR algorithm [1]. We proposed an efficient and stable implementation by taking advantage of the bidiagonal structure [3]. Numerical results illustrate that our aggressive deflation strategy often reduces the dqds runtime significantly. In addition, a shift-free version of our algorithm has a potential to be parallelized in a pipelined fashion.

In this talk, we focus on a stability theorem which states that our aggressive deflation strategy preserves high relative accuracy of singular values in floating point arithmetic. The theorem is proved by using a mixed forward-backward error analysis.

References

- [1] K. Braman, R. Byers, and R. Mathias, The multishift QR algorithm. II. Aggressive early deflation, *SIAM J. Matrix Anal. Appl.*, 23 (2002), pp. 948–973.
- [2] K. V. Fernando and B. N. Parlett, Accurate singular values and differential qd algorithms, *Numer. Math.*, 67 (1994), pp. 191–229.
- [3] Y. Nakatsukasa, K. Aishima, I. Yamazaki, dqds with aggressive early deflation, *SIAM J. Matrix Anal. Appl.*, vol. 33 (2012), pp. 22–51.
- [4] B. N. Parlett and O. A. Marques, An implementation of the dqds algorithm (positive case), *Linear Algebra Appl.*, 309 (2000), pp. 217–259.

Biography

September 1983, born in Japan.

March 2006, Bachelor of Engineering from Faculty of Engineering, The University of Tokyo.

March 2008, Master of Information Science and Technology from Graduate School of Information Science and Technology, The University of Tokyo.

March 2011, Doctor of Information Science and Technology from Graduate School of Information Science and Technology, The University of Tokyo.

April 2008 - March 2011, Research Fellow of the Japan Society for the Promotion of Science (DC1).

Since April 2011, Assistant Professor, The University of Tokyo.

Publications

K. Aishima, T. Matsuo, K. Murota, M. Sugihara: On convergence of the dqds algorithm for singular value computation, *SIAM Journal on Matrix Analysis and Applications*, vol. 30 (2008), pp. 522–537.

K. Aishima, T. Matsuo, K. Murota: Rigorous proof of cubic convergence for the dqds algorithm for singular values, *Japan Journal of Industrial and Applied Mathematics*, vol. 25 (2008), pp. 65–81.

K. Aishima, T. Matsuo, K. Murota, M. Sugihara: Superquadratic convergence of DLASQ for computing matrix singular values, *Journal of Computational and Applied Mathematics*, vol. 234 (2010), 1179–1187.

K. Aishima, T. Matsuo, K. Murota, M. Sugihara: A survey on convergence theorems of the dqds algorithm for computing singular values, *Journal of Math-for-Industry*, vol. 2 (2010), pp. 1–11.

K. Aishima, T. Matsuo, K. Murota: A note on the dqds algorithm with Rutishauser’s shift for singular values, *Japan Journal of Industrial and Applied Mathematics*, vol. 28 (2011), pp. 251–262.

K. Aishima, T. Matsuo, K. Murota, M. Sugihara: A Wilkinson-like multishift QR algorithm for symmetric eigenvalue problems and its global convergence, to appear in *Journal of Computational and Applied Mathematics*.

Y. Nakatsukasa, K. Aishima, I. Yamazaki: dqds with aggressive early deflation, *SIAM Journal on Matrix Analysis and Applications*, vol. 33 (2012), pp. 22–51.

Error control for hp -adaptive approximations of semi-definite eigenvalue problems

Luka Grubišić
Stefano Giani
Jeffrey Owall

We present a highly accurate finite element computational method for obtaining eigenvalue/eigenvector approximations of differential operators with known null spaces (e.g. operators having Neumann or periodic boundary conditions). In this context we identify two types of problems. Prototype examples for these problems are the Neumann Laplace operator and the operator describing the cell problem which is obtained by the Floquet analysis of photonic crystals.

In the case of the Laplace operator on a simply connected polygonal domain with the Neumann boundary condition the null vector is contained in the piecewise polynomial finite element space. In the case of the periodic model problem from the modeling of photonic crystals, the zero eigenvector contains exponential function(s) and so it is not contained in a typical piecewise polynomial finite element space.

The problems under study turn out to be rather challenging. We demonstrate this by analyzing the Neumann version of the touching squares example of M. Dauge. This problem can exhibit arbitrarily bad singularities in its eigenfunctions. However, using the penalization approach we are able to present a robust and efficient residual estimators for hp -adaptive approximations for both of these classes of problems.

In particular, we present reliable a-posteriori error estimates for hp -adaptive finite element approximations of semi-definite eigenvalue/eigenvector problems with explicitly known null spaces. We also present detailed numerical experiments confirming our theory and give several benchmark results which could serve the purpose of numerical testing of other adaptive procedures. Our estimators are implemented in the AptoFEM package.

Kogbetliantz-like Method for the Hyperbolic SVD

Sanja Singer
Vedran Novaković

It is well-known that both the ordinary and the hyperbolic Jacobi SVD algorithms are accurate in the relative sense. If a matrix is triangular, the Jacobi SVD algorithm can be replaced by the Kogbetliantz SVD algorithm. For the last one, there is always a choice from which side (left- or right-hand side) the transformations should be applied first, to obtain yet another accurate SVD algorithm.

Our goal was to construct the Kogbetliantz-like hyperbolic SVD algorithm, that should be used in the last phase of the orthogonalization, when the columns of a matrix are almost orthogonal.

To our surprise, there is only one pair of angles (trigonometric from the left, and hyperbolic from the right) that diagonalize a nonsingular 2×2 matrix. The computation of $\tan(2\varphi)$ and $\tanh(2\psi)$ can be made accurate, and the algorithm works satisfactory for small matrices.

Towards Multi-GPU Jacobi (H)SVD

Vedran Novaković

Sanja Singer

We present a design and implementation of the Jacobi-type (hyperbolic) SVD algorithms targeting NVidia GPU clusters. We build on a foundation of the Jacobi-type algorithms for CPU clusters, but these algorithms have to be rethought in the view of GPU computational constraints, such as “cheap” arithmetic, “slow” memory access, and fine granularity of independent tasks. We aim at the large dense matrices—orders of tens of thousands are feasible for the small hardware installations. However, the proposed algorithms scale to bigger problems easily, whenever the available machinery permits.

The GPU-oriented SVD algorithms are a topic of an intensive research, e.g., in such projects as MAGMA and CULA. But, unlike some hybrid CPU–GPU efforts, the inherent parallelism of the Jacobi-type methods shifts the computation almost entirely to the GPUs, while the CPUs have only the scheduling and synchronization roles.

The proposed algorithms are variants of the full-block (hyperbolic) Jacobi SVD, organized in three blocking levels. The blocks of the outermost level are mapped, one-to-one, to the available GPUs. Inside of the each GPU the inner blocking is performed according to the size of the small but fast “shared” memory. In the innermost level a pointwise Jacobi algorithm is performed.

The role of blocking goes beyond the data partition—it ensures full utilization of the fastest memories and communication channels. For data exchange at the outermost level, a class of communication-optimal parallel pivoting strategies are generated. In the inner level, we have to adapt the Cholesky or tall-and-skinny QR factorizations for the shared memory.

In this talk we shall further discuss the details and the choices that have to be made at the each level. Finally, we shall briefly consider future research on Kepler GPUs and larger clusters.

This work is supported in part by NVidia Academic Partnership program.

$UB_{k+1}V$ Block Sparse Householder Decomposition
talk submission

Gary W. Howell

http://www.ncsu.edu/itd/hpc/Documents/Publications/gary_howell/contents.html

This talk describes Householder reduction of a rectangular sparse matrix to small band upper triangular form. Using block Householder transformations gives good orthogonality, is computationally efficient, and has good potential for parallelization. The algorithm is similar to the standard dense Householder reduction used as part of the usual dense SVD computation. For the sparse algorithm, the original sparse matrix is accessed only for sparse matrix dense matrix (SMDM) multiplications. For a triangular bandwidth of $k+1$, the dense matrices are the k rows or columns of a block Householder transformation. Using an initial random block Householder transformation allows reliable computation of a collection of largest singular values.

Recent progress has included

- Improving convergence criteria in terms of requests for a desired number of largest singular values,
- Improving performance of “tall and skinny” and “long and thin” BLAS (the predominant calculation),
- Using the sparse block Householder algorithm to solve sparse least squares problems.

Optimization of Eigenvalues of Hermitian Matrix Functions and Applications

Emre Mengi

The main theme of this talk is a Hermitian matrix function depending on its parameters analytically. We describe a numerical algorithm for the global optimization of a specified eigenvalue of such a Hermitian matrix function over the space of parameters. The algorithm is driven by piece-wise quadratic under-estimators for the eigenvalue function, and locating their global minimizers. In the multi-dimensional case the global minimization of the piece-wise quadratic models can be posed as quadratic programs. The algorithm generates sequences converging to global optimizers at a linear rate in practice. The second part of the talk is devoted to specific distance problems over the space of matrices leading to these Hermitian eigenvalue optimization problems. We discuss problems such as the distance to a nearest defective matrix, and the distances to matrix pencils and matrix polynomials with specified eigenvalues, which are motivated by sensitivity analysis in numerical linear algebra as well as applications in control theory and signal processing.

The talk is based on joint works with Michael Karow, Mustafa Kılıç, Daniel Kressner, Ivica Nakić, Ninoslav Truhar and E. Alper Yıldırım

Eigenvalue Algorithms for Symmetric Hierarchical Matrices

Thomas Mach

We will present three different approaches to compute the eigenvalues of symmetric hierarchical (\mathcal{H} -) matrices. The \mathcal{H} -matrices have been introduced by W. Hackbusch in 1998 [1]. He also introduced a subset of simple structured \mathcal{H} -matrices called the \mathcal{H}_ℓ -matrices.

The main advantage of \mathcal{H} -matrices is their data-sparsity requiring only $\mathcal{O}(n(\log_2 n)^\alpha)$ storage for a matrix $M \in \mathbf{R}^{n \times n}$. Further, many arithmetic operations, like matrix additions, multiplications and factorizations, can be performed with almost linear complexity in the matrix dimension n .

We will use the cheap arithmetic to build cheap eigenvalue solvers that compute either one eigenvalue in $\mathcal{O}(n(\log_2 n)^\alpha)$ or all in $\mathcal{O}(n^2(\log_2 n)^\alpha)$.

First, we try to use the cheap Cholesky factorization for \mathcal{H} -matrices to build an explicit LR Cholesky algorithm, see [2]:

$$\begin{aligned}L_i L_i^T &:= M_i - \mu I \\ M_{i+1} &:= L_i^T L_i + \mu I.\end{aligned}$$

We observe that this algorithm does *not* lead to a structure preserving eigenvalue algorithm, since the local block-wise ranks grow too much.

We will show that the structure of \mathcal{H}_ℓ -matrices is almost preserved. The LR Cholesky algorithm for \mathcal{H}_ℓ -matrices is of almost quadratic complexity.

Second, we will use a bisectioning method to compute the eigenvalues of \mathcal{H}_ℓ -matrices in a more efficient way. The third approaches compute the smallest or some of the eigenvalues by preconditioned inverse iteration.

This is joint work with Peter Benner (MPI Magdeburg).

References

- [1] W. HACKBUSCH, *A sparse matrix arithmetic based on \mathcal{H} -matrices. Part I: Introduction to \mathcal{H} -matrices*, Computing, 62 (1999), pp. 89–108.
- [2] H. RUTISHAUSER, *Solution of eigenvalue problems with the LR-transformation*, Nat. Bur. Standards Appl. Math. Ser., 49 (1958), pp. 47–81.

Fast and accurate con-eigenvalue algorithm for optimal rational approximations

Terry S. Haut

The need to compute small con-eigenvalues, and the associated con-eigenvectors, of positive-definite Cauchy matrices naturally arises when constructing rational approximations with a near optimally small L^∞ error. Specifically, given a rational function with n poles in the unit disk, a rational approximation with $m \ll n$ poles in the unit disk can be obtained from the m th con-eigenvector of an $n \times n$ Cauchy matrix, where the associated con-eigenvalue $\lambda_m > 0$ gives the approximation error in the L^∞ norm. Unfortunately, standard algorithms do not accurately compute small con-eigenvalues (and the associated con-eigenvectors) and, in particular, yield few or no correct digits for con-eigenvalues smaller than the machine roundoff.

In this talk, I will first describe the theory behind constructing near optimal rational approximations, and connect it with the problem of computing small con-eigenvalues of Cauchy matrices. I will then present a new fast and accurate algorithm for computing con-eigenvalues and con-eigenvectors of positive-definite Cauchy matrices, yielding even the tiniest con-eigenvalues with high relative accuracy. The algorithm computes the m th con-eigenvalue in $\mathcal{O}(m^2n)$ operations and, since the con-eigenvalues of positive-definite Cauchy matrices decay exponentially fast, yields near optimal rational approximations in $\mathcal{O}(n(\log \delta^{-1})^2)$ operations, where δ is the approximation error in the L^∞ norm. I will also present error bounds demonstrating the high relative accuracy of the computed con-eigenvalues and the high accuracy of the unit con-eigenvectors. Finally, I will briefly describe recent applications of this algorithm for computing near optimal rational approximations of functions with singularities and sharp transitions, where approximation errors close to machine precision are reliably obtained. In one application I will present, a numerical calculus based on this algorithm is used to compute highly accurate solutions of viscous Burgers' equation over long time intervals, with viscosity as small as 10^{-5} .

Reproducible Parallel Floating-Point Computations

Hong Diep Nguyen

James Demmel

Because of rounding errors, floating-point operations such as addition and multiplication are not associative, so computed results depend also on the order of computation. Therefore we cannot get the same answer from run-to-run even on the same machine with varying numbers of available processors. That makes understanding the reliability of output harder, especially with the increasing level of parallelism. In the specific field of testing and debugging for example, each debug run should produce the same trace as the previous one. On the other hand, the running time should not be too slow compared to a non-debug run. Our goals are to design reproducible algorithms which are as efficient as the non-reproducible counter-parts.

Currently, our works focus on designing repeatable algorithms on shared-memory multicore architecture, which is a subset of our overall goal, yet it is still difficult to achieve the same performance as non-repeatable counter-parts. Our first approaches include using a reproducible reduction tree, and using higher precision.

A set of parallel algorithms use a reduction tree to produce the final result. This reduction tree is subject to change with varying numbers of computing threads, which is the source of non-reproducibility. Our first approach is based on using a reproducible reduction tree, which is fixed ahead of computing time and does not depend on available resources at run time. Input data are split into fixed-size chunks, and the reduction tree is imposed over these chunks instead of threads. Experimental results showed that the overhead of computing time introduced by the reproducible reduction tree is negligible in comparison with the conventional reduction tree.

Non-reproducibility of floating-point computation comes from rounding-errors, so another solution might be to eliminate these rounding errors. Using a higher precision reduces the effect of rounding errors and might produce reproducible results given that the problem is not too ill-conditioned. Nevertheless, using a higher precision (for ex. double-double precision) can drastically increase the computing time. Our objective is to obtain reproducible results, not accurate results, so we do not need to use higher precision everywhere, but only at some selective points of computation. This technique has been embraced in testing and debugging. Experimental results showed that using higher precision helps to provide reproducible results for some set of not too ill-conditioned problems with a minor overhead of computing time.

Numerical Solution of Two-Parameter Differential Eigenvalue Problems

B. Chanane (Joint work with A. Boucherif)

In an interesting paper published in 1963, F. M. Arscott [1] showed that the method of separation of variables used in solving boundary value problems for Laplace's equation leads to a two-parameter eigenvalue problem for ordinary differential equations with the auxiliary requirement that the solutions satisfy boundary conditions at several points. This has led to an extensive development of multiparameter spectral theory for linear operators. In the paper [2] the authors give an overview results on two-parameter eigenvalue problems for second order linear differential equations. See [3] for latest account on the subject. Consider the two-parameter Sturm-Liouville problem

$$\begin{cases} -y'' + qy = (\mu_1^2 w_1 + \mu_2^2 w_2)y, & 0 < x < 1 \\ y(0) = 0, & y(c) = 0, & y(1) = 0 \end{cases}$$

where w_1, w_2 are positive and in $C^2[0, 1]$ and $q \in L[0, 1]$ and $c \in (0, 1)$ some given constant.

By an eigenvalue of the above problem, we mean a value of the couple (μ_1, μ_2) for which problem () has a non trivial solution. Conditions that insure the existence of eigenvalue can be found in [3]. Our objective is to effectively compute the eigenpairs (μ_1^2, μ_2^2) using the *Regularized Sampling Method* (RSM) introduced by the author in [4] and tested on broad classes of Sturm-Liouville problems (Singular, Non-Self-Adjoint, Non-Local, Impulsive,...). We should point out that we have restricted our attention to Dirichlet boundary conditions in order to eliminate technical details that might obscure the ideas.

We shall consider the associated initial value problem,

$$\begin{cases} y'' + (\mu_1^2 w_1 + \mu_2^2 w_2)y = qy, & 0 < x < 1 \\ y(0) = 0, & y'(0) = 1 \end{cases}$$

and deal first with the unperturbed case ($q = 0$) then with the perturbed case ($q \neq 0$). Let PW_{β_1, β_2} denote the Paley-Wiener space,

$$PW_{\beta_1, \beta_2} = \left\{ \begin{array}{l} h(z_1, z_2) \text{ entire} / |h(z_1, z_2)| \leq C \exp \{ \beta_1 |z_1| + \beta_2 |z_2| \}, \\ \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} |h(z_1, z_2)|^2 dz_1 dz_2 < \infty \end{array} \right\}$$

• In the unperturbed case ($q = 0$),

Theorem The solution φ_1 is an entire function of $(\mu_1, \mu_2) \in C^2$ for each fixed $x \in (0, 1]$ of order $(1, 1)$ and type $(\sigma_1(x), \sigma_2(x))$ and satisfies the estimate,

$$|\varphi_1(x)| \leq K_1 \exp [\sigma_1(x) |\mu_1| + \sigma_2(x) |\mu_2|] , \quad (\mu_1, \mu_2) \in C^2$$

for each fixed $x \in (0, 1]$ where, $\sigma_i(x) = 2 \{ x \int_0^x w_i(\xi) d\xi \}^{\frac{1}{2}}$ for $i = 1, 2$. Moreover, the function φ_1 satisfies the estimate

$$|\varphi_1(x)| \sim \frac{K_4(x)}{|\mu_1| + |\mu_2|}, \text{ as } |\mu_1| + |\mu_2| \rightarrow \infty, (\mu_1, \mu_2) \in R^2$$

where K_4 is bounded and is independent of (μ_1, μ_2) .

Thus, φ_1 as a function of (μ_1, μ_2) is not in a Paley-Wiener space.

Lemma The function α defined by

$$\alpha(x) = (\text{sinc} \{ \sigma_1(x)\mu_1 + \sigma_2(x)\mu_2 \})^m$$

where $\text{sinc}(z) = z^{-1} \sin z$ and m is a positive integer, is an entire function of $(\mu_1, \mu_2) \in C^2$ for each fixed $x \in (0, 1]$ of order $(1, 1)$ and type $(\sigma_1(x), \sigma_2(x))$. Furthermore, α satisfies the estimate

$$|\alpha(x)| \sim \frac{K_3(x)}{|\mu_1| + |\mu_2|}, \text{ as } |\mu_1| + |\mu_2| \rightarrow \infty, (\mu_1, \mu_2) \in R^2.$$

where K_3 is bounded and is independent of (μ_1, μ_2) .

Theorem $\alpha(x)\varphi_1(x)$, as a function of (μ_1, μ_2) belongs to the Paley-Wiener space PW_{β_1, β_2} where $(\beta_1, \beta_2) = ((m+1)\sigma_1(x), (m+1)\sigma_2(x))$ for each fixed $x \in (0, 1]$.

• In the perturbed case ($q \neq 0$), we show that the solution $y(x)$ is an entire function of (μ_1, μ_2) for each $x \in (0, 1]$, of order $(1, 1)$ and type $(\sigma_1(x), \sigma_2(x))$ but not in a Paley-Wiener space. However multiplication by α gives a function $\alpha(x)y(x)$ of (μ_1, μ_2) in a Paley-Wiener space PW_{β_1, β_2} for each $x \in (0, 1]$. More specifically, we have the following,

Theorem The function \tilde{y} defined by $\tilde{y}(x) = \alpha(x)y(x)$, belongs to PW_{β_1, β_2} where $(\beta_1, \beta_2) = ((m+1)\sigma_1(x), (m+1)\sigma_2(x))$ as a function of $(\mu_1, \mu_2) \in C$ for each $x \in (0, 1]$, and satisfies the estimate,

$$|\alpha(x)y(x)| \sim \frac{K(x)}{(|\mu_1| + |\mu_2|)^{m+1}}, \text{ as } |\mu_1| + |\mu_2| \rightarrow \infty, (\mu_1, \mu_2) \in R^2.$$

where K is bounded and is independent of (μ_1, μ_2) .

Hence, $\tilde{y}(x; \mu_1, \mu_2)$ can be recovered at each $x \in (0, 1]$ from its samples at the lattice points $(\mu_{1j}, \mu_{2k}) = (j \frac{\pi}{(m+1)\sigma_1(x)}, k \frac{\pi}{(m+1)\sigma_2(x)})$, $(j, k) \in Z^2$ using the rectangular cardinal series,

$$f(\mu_1, \mu_2) = \sum_{j=-\infty}^{\infty} \sum_{k=-\infty}^{\infty} f(\mu_{1j}, \mu_{2k}) \frac{\sin \beta_1(\mu_1 - \mu_{1j})}{\beta_1(\mu_1 - \mu_{1j})} \frac{\sin \beta_2(\mu_2 - \mu_{2k})}{\beta_2(\mu_2 - \mu_{2k})}$$

for $f \in PW_{\beta_1, \beta_2}$. The convergence of the series being uniform and in $L^2_{d\mu_1 d\mu_2}(R^2)$, and $\mu_{mn} = n\pi/\beta_m$, $m = 1, 2$, $n \in Z$. The eigenpairs are therefore (μ_1^2, μ_2^2) where (μ_1, μ_2) solve the nonlinear

system $B_1(\mu_1, \mu_2) := \tilde{y}(1; \mu_1, \mu_2) = 0$, $B_2(\mu_1, \mu_2) := \tilde{y}(c; \mu_1, \mu_2) = 0$.

To illustrate the efficiency of the method, we shall find some of the eigenpairs of the two-parameter Sturm-Liouville problem with three-point boundary conditions given by

$$\begin{cases} -y'' = (\mu_1^2 + \mu_2^2 x)y, & 0 < x < 1 \\ y(0) = y(0.7) = y(1) \end{cases}$$

The general solution y of the differential equation can be expressed in terms of the Airy functions of the first and second kind Ai and Bi and their first derivatives as

$$y(x; \mu_1, \mu_2) = \frac{(-1)^{2/3} \left(\text{Ai} \left(\frac{\sqrt[3]{-1}\mu_1^2}{\mu_2^{4/3}} \right) \text{Bi} \left(\frac{\sqrt[3]{-1}(\mu_1^2 + \mu_2^2 x)}{\mu_2^{4/3}} \right) - \text{Ai} \left(\frac{\sqrt[3]{-1}(\mu_1^2 + \mu_2^2 x)}{\mu_2^{4/3}} \right) \text{Bi} \left(\frac{\sqrt[3]{-1}\mu_1^2}{\mu_2^{4/3}} \right) \right)}{\mu_2^{2/3} \left(\text{Ai}' \left(\frac{\sqrt[3]{-1}\mu_1^2}{\mu_2^{4/3}} \right) \text{Bi} \left(\frac{\sqrt[3]{-1}\mu_1^2}{\mu_2^{4/3}} \right) - \text{Ai} \left(\frac{\sqrt[3]{-1}\mu_1^2}{\mu_2^{4/3}} \right) \text{Bi}' \left(\frac{\sqrt[3]{-1}\mu_1^2}{\mu_2^{4/3}} \right) \right)}$$

The eigenpairs (μ_1^2, μ_2^2) can be obtained from the solutions (μ_1, μ_2) of the system $\tilde{y}(1; \mu_1, \mu_2) = 0$, $\tilde{y}(c; \mu_1, \mu_2) = 0$. For numerical purposes we have truncated the associated series to $|j|, |k| \leq N = 50$ and took $m = 5$ in the function α . The next table shows the exact eigenpairs together with their approximations using the *Regularized Sampling Method* (RSM).

μ_1 (exact)	μ_2 (exact)	μ_1 (RSM)	μ_2 (RSM)
7.788149097670813	7.5908224365786845	7.78814883048352	7.590823605399021
4.194056047341936	22.273796542861913	4.194057357011064	22.273795699572364
15.597607295800684	15.165889997583099	15.597605904429917	15.1658929384375
13.87615754699624	30.597837626955204	13.876157718920586	30.59783713797321

No effort has been exerted to get higher precisions, thus all of the errors are between 10^{-6} and 10^{-7} . A simple shift in (μ_1, μ_2) could have allowed us having much higher precision.

Acknowledgment The support of KFUPM is greatly appreciated.

References

- [1] F. M. Arscott, 1963, Paraboloid coordinates and Laplace's equation, Proc. Royal Soc. Edinburgh 66, pp. 129-139.
- [2] P. Binding, H. Volkmer, 1996, Eigencurves for two-parameter Sturm- Liouville equations, SIAM Review 38, pp. 27-48.
- [3] F. V. Atkinson, Angelo B. Mingarelli, Multiparameter problems Sturm-Liouville Theory, CRC Press, 2011.
- [4] B. Chanane, Computation of the eigenvalues of Sturm-Liouville Problems with parameter dependent boundary conditions using the regularized sampling method, Math. Comput. 74 (252) (2005) 1793-1801, published electronically S 0025-5718(05)01717-5.