The Abstract Book
of
International Workshop on Accurate Solution of Eigenvalue Problems
at
The Pennsylvania State University
University Park
1998

Monday, July 20

Time: 9-10am

Title: Relatively Robust Representations of tridiagonals

Speaker: B.N. Parlett

Abstract:

Let $T-\tau*I = L*O*L^t$ denote the 'Cholesky' factorization of a real symmetric unreduced tridiagonal matrix. In general O is a signed symmetric permutation (SSP) matrix but we assume that it is diag(+/-1). Our question is whether the small eigenvalues are determined to high relative accuracy by bidiagonal L? Let $(s, \lambda - \tau)$ represent a typical eigenpair with $||s|| = 1$. The vector that governs the answer is $f = L^t * s$ and the relative condition number is $||f||^2/|\lambda - \tau|$. We study this question as τ varies. We show that the smallest eigenvalue in magnitude is always determined to high relative accuracy by L. Using this Dhillon and I have shown that, in the presence of roundoff error, our algorithm delivers a computed approximation to s with error angle O(macheps/relgap) whenever the eigenpair is determined to high relative accuracy by L. This result leads us to focus on values of τ close to clusters of close eigenvalues because relgap is a function of τ. I will give detailed numerical examples to illustrate the results.

Time: 10:30-11:30am

Title: QR Factorization with Complete Pivoting and Accurate Computation of the SVD

Speaker: Nicholas J. Higham

Abstract:

A new algorithm of Demmel et al. for computing the singular value decomposition (SVD) to high relative accuracy begins by computing a rank-revealing decomposition by Gaussian elimination with complete pivoting (GECP). We investigate the use of QR factorization with complete pivoting (that is, column pivoting together with row sorting or row pivoting) as an alternative to GECP for graded matrices. We show that a faster SVD algorithm results. We derive a new row and column-wise backward error result for Householder QR factorization and combine it with the theory of Demmel et al. to show when high accuracy in the computed SVD can be expected. Numerical experiments confirm the theoretical predictions.

Time: 11:30-12:30pm

Title: Relative perturbation theory for Hermitian matrices

Speaker: Ivan Slapničar

Abstract:

Relative perturbation theory essentially gives bounds for relative changes in eigenvalues of a given matrix in the case when the matrix or its elements are perturbed in some relative sense. Such perturbations are important since they typically occur during numerical computations, and to ask how many digits can be or are accurately computed is a very natural question. Although derived by applying classical perturbation results, relative perturbation bounds are often much sharper then their classical counterparts in the case of relative perturbations. Numerous results of this kind can be grouped according to the structure of the matrix. So, there exist bounds for positive definite matrices and singular

values, Hermitian matrices, normal matrices, diagonalizable matrices, and certain matrix pairs. Naturally, the stronger the matrix structure, the better and more easily applicable are the bounds which can be obtained.

This talk deals primarily with Hermitian matrices. The types of perturbations $H \to H + \delta H$ that are being considered are:

- elementwise relative perturbation $|\delta H_{ij}| \leq \varepsilon |H_{ij}|$;
- perturbation of a graded matrix $H = D^* A D$ where $\delta H = D^* \delta A D$;
- perturbation of a matrix in factorized form where $H = GAG^*$ is perturbed through $G \to G + \delta G$.

The relative perturbation bounds are derived in two ways: one approach uses the scaled spectral absolute value of H (the positive definite polar factor of H); the other approach is based on Birkhoff's lemma. Several relative perturbation bounds, residual bound, and quadratic residual bound are given for the eigenvalues. The corresponding eigenvector bounds are given in terms of normwise perturbation bounds for the difference of projections onto invariant subspaces. These subspace bounds are essentially relative variants of the $\sin\Theta$ theorems by Davis and Kahan. Most of our bound are readily computable and can be applied at little extra cost to estimate the accuracy of the eigensolution which is computed by certain highly accurate algorithms.

Time: 2-3pm

Title: Rank-Revealing factorizations are High Accuracy factorizations.

Speaker: Ming Gu

Abstract:

In this talk, we discuss recent development in the existence and computation of rank-revealing QR, Cholesky, and LU factorizations. We further show that both absolute error and relative error bounds similar to those for the SVD can be obtained for these rank-revealing factorizations. Such bounds are not available for the conventional QR, Cholesky, and LU factorizations.

Time: 4:30-5pm

Title: Reduction to banded form: Eigenvalues and the Lyapunov equation e= ndtitle

Speaker: Eugene L. Wachspress

Abstract:

Any real nonsymmetric square matrix may be reduced to banded upper Hessen berg form with a succession of permutation and elementary gaussian transformations. The rows and columns are reduced in succession from the head to tail of the matrix. Bandwidth is governed by a prescribed bound on the magnitude of elements added to the unreduced part of the matrix by a proposed gaussian reduction. A matrix of order n is reduced in around $\frac{8}{3}n^3$ flops. Eigenvalues of the reduced matrix are then computed with shifted inverse iteration and deflation which preserves the banded structure. This requires $O(400n^2)$ flops when eigenvalue accuracy of around 10^{-6} is prescribed. When solving the Lyapunov matrix equation, $AX + XA^\top = 3DC$, the transformation matrix G^{-1} and the transformed right-hand side F are stored along with the reduced matrix S, where $S = 3DGAG^{-1}$ and $F = 3DGCG^\top$. This increases the reductio n flops from $\frac{8}{3}n^3$ to $8n^3$. There are matrices for which more than one pass is used

in the reduction with a relaxed gaussian tolerance for each pass. This can on rare occasions increase the reduction flops by a factor of two or three. ADI solution of the reduced equation $SY + YS^{\top} = 3DF$ requires around $500n^2$ flops. Recovery of $X = 3DG^{-1} * Y * (G^{-1})^{\rightarrow} = p$ uses another $2n^3$ flops. The total flops for solution of the Lyapunov matrix equation is thus around $10n^3 + 900n^2$ when only one reduction pass is needed and may grow to around $20n^3 + 1000n^2$ for some problems . This may be compared with the Bartels-Stewart algorithm which requires around $25 - 30n^3$ flops.

Time: 5-5:30pm

Title: Backward error and Condition Number of Quadratic Eigenvalue Problems

Speaker: Françoise Tisseur

Abstract:

We define and evaluate backward errors and condition numbers for the quadratic eigenvalue problem. We consider normwise measures as well as componentwise measures. The common practice when solving the quadratic eigenvalue problem is to reformulate it as a generalized eigenvalue problem of twice the order. They are many ways of performing this transformation. We investigate the sensitivity of a given eigenvalue of the quadratic eigenvalue problem to perturbations in each of the generalized eigenvalue problem formulations. We show that there can be great variation in sensitivity. The analysis is illustrated by numerical experiments.

Time: 5:30-6pm

Title: Perturbation of the Indefinite QR Factorization

Speaker: Sanja Singer, Saša Singer

Abstract:

Let $G \in \mathbf{C}^{m\times n}$, $m \geq n$ and $J \in \mathbf{C}^{m\times m}$, $J = \mathrm{diag}(j_{11}, \ldots, j_{mm})$, $j_{ii} \in \{-1, 1\}$ are given. If matrix $A = G^*JG$ is regular then exists a factorization of G

$$G = P_1 Q \begin{bmatrix} R \\ 0 \end{bmatrix} P_2^* \quad , \quad Q^* P_1^* J P_1 Q = P_1^* J P_1,$$

where P_1 and P_2 are permutation matrices and R is block upper triangular with diagonal blocks of order 1 or 2. Matrix Q is $P_1^*JP_1$–unitary matrix. This factorization is a generalization of well–known QR factorization and it is called the indefinite QR factorization according to the given J.

For some problems matrices G and J are naturally given, i.e. for matrices A of form $A = L^*L - M^*M$.

Accurate one–sided Jacobi algorithm for diaginalization of A works on G^* which usually has more rows then columns. In our case G^* has more colums then rows and we have two possibilities.

1. Multiplication of G^*JG and then symmetric indefinite factorization of G^*JG gives square factor R' and signature matrix J_1' (smaller then J).

2. Indefinite QR factorization of G gives another square factor R and signature matrix J_1.

In some cases indefinite QR factorization of G can be more accurate then multiplication and factorization strategy.

We prove the following theorem:

Theorem.

Let

$$\begin{bmatrix} R \\ 0 \end{bmatrix} = Q^{-1} P_1^* G P_2 \quad , \quad Q^* J_1 Q = J_1 \quad , \quad J_1 = P_1^* J P_1,$$

be exact indefinite factorization of G, with upper triangular R. If matrix $\widetilde{R}$ is upper triangular and computed in floating–point arithmetic, then $\widetilde{R}$ is considered as a matrix given by exact indefinite QR factorization of some perturbed matrix $\widetilde{G} = G + E$,

$$\begin{bmatrix} \widetilde{R} \\ 0 \end{bmatrix} = \widetilde{Q}^{-1} P_1^* (G + E) P_2 \quad , \quad \widetilde{Q}^* J_1 \widetilde{Q} = J_1 \quad , \quad J_1 = P_1^* J P_1$$

where

$$\|P_1^* E P_2 e_k\|_2 \leq \gamma^{k-1} \|P_1^* G P_2 e_k\|_2 \sum_{i=1}^{k} \mathrm{err}(p_i),$$

$$\mathrm{err}(p_i) = e_H (1 + e_G)^{p_i} + \gamma \left[(1 + e_G)^{p_i} - 1 \right].$$

p_i, β and γ are constants and e_G, e_H depend on ε (accuracy of the basic floating–point operations).

This perturbation bound gives

$$|\delta A_{ij}| \leq (2\delta + \delta^2) \omega_i \omega_j \sqrt{|A_{ii}||A_{jj}|}$$

where

$$\delta = \gamma^{n-1} \sum_{i=1}^{n} \mathrm{err}(p_n) \quad , \quad \|g_k\| = \omega_k \sqrt{|e_k^* G^* J G e_k|}$$

which fits into the standard relative perturbation theory for eigenvalues.

Superior accuracy of indefinite QR algorithm (in some cases) will be illustrated by selected numerical examples.

Tuesday, July 21

Time: 9-10am

Title: TBA

Speaker: James Demmel

Abstract:

Time: 10:30-11:30am

Title: Absolute and Relative Perturbation Bounds for Invariant Subspaces of Diagonalisable Matrices

Speaker: Ilse C. F. Ipsen

Abstract:

We derive absolute and relative bounds for angles between invariant subspaces of a diagonalisable matrix A and a perturbed matrix $A + E$, in the two-norm and in the Frobenius norm. Our absolute bounds are extensions of Davis and Kahan's $\sin\theta$ theorem to diagonalisable matrices and invariant subspaces of any dimension. When the perturbed subspace has dimension one, our relative bound is implied by the absolute bound. Thus, relative bounds for invariant subspaces are not necessarily stronger than absolute bounds.

Our relative bounds are more general than existing bounds because they place no restrictions on the perturbed matrix and they hold for the larger class of diagonalisable matrices. Moreover they are simple and easy to interpret.

Time: 11:30-12pm

Title: On constructing matrices with prescribed singular values and eigenvalues

Speaker: Alicja Smoktunowicz

Abstract:

We propose an efficient algorithm for computing a unit lower triangular matrix A with prescribed singular values $\sigma_1 \geq \sigma_2 \geq \ldots \geq \sigma_n$ such that $\sigma_1\sigma_2 \ldots \sigma_n = 1$. This is a solution of the question raised by N.J.Higham.
The singular values of A are the positive square roots of the eigenvalues of A^*A. Due to A.Horn theorem such matrix A exists if and only if the Weyl conditions are satisfied: $\sigma_1\sigma_2 \ldots \sigma_{i-1} \geq 1$ for $i = 2, \ldots, n$ and $\sigma_1\sigma_2 \ldots \sigma_n = 1$.
The main transformation in our algorithm is reduced to find a stable singular value decomposition of a triangular 2×2 matrix. We construct a sequence of unitarily equivalent lower trangular matrices with diagonal matrix $diag(\sigma_1, \sigma_2, \ldots, \sigma_n)$. We consider also a general case for arbitrary triangular matrices. This is joint work with C.-K. Li and R. Mathias.

The second part of our talk jointly with J. Gluchowska – Jastrzebska and A. Grabarski treats on numerical properties of several algorithms for improving the computed eigenvalues or singular values of some matrices. These algorithms are some variants of Newton's method for finding roots of the functions.

Time: 12-12:30pm

Title: A Hybrid Scheme to Improve Jacobi-Davidson Method On Eigenvalue Problems

Speaker: Yong Sun

Abstract:

The Jacobi-Davidson (JD) iterative method has been proposed as an efficient and parallelizable algorithm for solving various types of large eigenproblems because no matrix inversion is required. In our application of the JD method to standard linear eigenproblems, it is recognized that if the starting vector for the JD iteration is not close enough to the actual eigenvector, as is usually the case, the subspace expansion by solving the Jacobi Orthogonal Component Correction equation (JOCC) often leads to unsatisfactory convergence behavior. As an improvement, we suggest that the JD iteration be carried out only after we obtain a fairly good approximation of the desired eigenvector by some less accurate but faster algorithm, which can be some kind of Krylov subspace method. We will show some examples of such a method. The results from our implementations thus far indicate that the hybrid scheme consistently outperforms the original JD method. We expect the scheme to be readily extendable to other types of eigenproblems as well.

Wednesday, July 22

Time: 9-10am

Title: Accuracy for the eigensolution for indefinite Hermitian matrices.

Speaker: Kresimir Veselć

Abstract:

First results on high relative accuracy concern mostly positive definite matrices and their natural pendant - the singular value decomposition. This also covers most applications. Indefinite matrices are another world. Here too the factorized pendant exists, the so called hyperbolic singular value decomposition. This time the relation between the two is more complicated and less straightforward but nontrivial and relevant. The applications include dicretized Sturm-Liouville problems with indefinite potentials as well as quasidefinite matrices.

Time: 10:30-11:30am

Title: New Variational Characterizations for Trapping Eigenvalues

Speaker: Christopher Beattie

Abstract:

Eigenvalue estimates that are optimal in some sense have self-evident appeal and leave estimators with a sense of virtue and economy. Often such estimates may be viewed as natural outcomes of various variational characterizations of the original eigenvalues — the "min-max" principle being the most common example. We'll review a variety of results related to obtaining optimal bounds — some results are well-known; others are less known; and a few are new. Included in the pantheon are Ritz and harmonic Ritz values; right- and left-definite variants of Lehmann (and Kahan) bounds; and a new class of bounds that could be termed "relative Lehmann" bounds which originate from a novel variational characterization of eigenvalues. We'll see how this helps explain why the Hope Diamond is blue and why, sadly, "optimal" does not necessarily mean "good".

Time: 11:30-12:30pm

Title: Accuracy of computed eigenvalues: backward errors and pseudospectra

Speaker: V. Toumazou

Abstract:

We show that backward errors and pseudospectra combined together are useful tools to assess the validity of a computed eigenvalue.

1. Given a set τ of admissible perturbations ΔA on a matrix A and a norm on τ (relative or absolute), the backward error $BE(z)$ for z as a candidate eigenvalue of A, is the smallest size of perturbation ΔA such that z is an exact eigenvalue of $A + \Delta A$.

2. The pseudospectrum associated with a backward error of level ε is $\sigma_\varepsilon(A) = \{z \in C; \quad BE(z) \leq \varepsilon\}$. It contains all the points z which are seen as eigenvalues within an accuracy tolerance of ε.

In this talk we present two practical classes of application:

a) Given A, τ is the set of all ΔA with scaled norm $\alpha\|\Delta A\|$. Therefore $BE(z) = \alpha/\|(A - zI)^{-1}\|$. Often $\alpha = 1/\|A\|$.

b) Given A and E, τ is the set of perturbations $\Delta A = tE$, $t \in C$ with scaled norm $|t|\beta$. Therefore $BE(z) = \beta/\rho(E(A - zI)^{-1})$. Often $\beta = \|E\|/\|A\|$.

Time: 2-3pm

Title: Inexact Symmetric Lanczos Methods for Eigenproblems with Applications

Speaker: H. Zha

Abstract:

At each step of the Lanczos iteration, a matrix-vector multiplication needs to be performed. Here are two possible ways that these multiplications are not computed accurately: i) the operation involves solving a large linear system, and an iterative method is used for its solution; ii) the matrix is deliberately replaced by a simpler matrix to reduce computational complexity. In either case, there is need to understand how Lanczos methods behave when the matrix-vector multiplications are computed inexactly. We develop a method based on some structured perturbation analysis of tridiagonal matrices to explain the behaviors of inexact Lanczos methods, and show how the insights can be used to control the solution accuracy in an inner-outer iteration of inexact Lanczos methods. Numerical examples from several applications will also be discussed.

Time: 4:30-5pm

Title: Accurate computation of singular values of long products of matrices

Speaker: David Stewart

Abstract:

The difficulty is in computing the smaller singular values of the product to high relative accuracy. A number of methods are available to compute the eigenvalues with small backward error for the product; that is, the computed eigenvalues are the exact eigenvalues of the product

$$(A_t + E_t)(A_{t-1} + E_{t-1}) \cdots (A_2 + E_2)(A_1 + E_1).$$

Are the relative perturbations of the singular values $\delta\sigma_i/\sigma_i$ small? The answer is "not always"; a complete answer requires greater understanding of dynamical systems theory. However, it is often possible to obtain highly accurate singular values even when the ratio of the smallest to the largest singular values is, say, 10^{-1000}.

Time: 5-5:30pm

Title: The relative error in Pruess method for Sturm–Liouville problems

Speaker: Przemysław Kosowski

Abstract:

We consider numerical methods for eigenvalue problems of Sturm–Liouville operator. We will look more closely at Pruess method as the one which gives better bounds for the errors of eigenvalues. The talk will focus on the relative spectral continuity of Sturm–Liouville operator and searching suitable approximation.

Numerical examples will be provided which show the behaviour of the relative accuracy of Pruess method.

Time: 5:30-6pm

Title: A Graphical Approach to Design a Parallel Matrix Transformation Algorithm

Speaker: Suely Oliveira

Abstract:

Lately graph theory and other data structures have played a special role in the development of algorithms for various problems in computational science. In this talk I will present a graph theoretical approach for transforming an arrowhead matrix in tridiagonal formula in a more parallelizable way. The graph theoretical approach has been used to show that certain structures maintain the accuracy of eigenvalues. The same approach can be used here.

Based on a algorithm used by Zha, we designed a new chasing algorithm for transforming arrow matrices into tridiagonal form using a graphic-theoretical approach. Although this algorithm has the same sequential computational complexity and backward error properties as the old algorithms, it is better suited for a pipelined approach. I will present the parallel algorithm for this new chasing method, and its performance results on mesh and hypercube architecures. Remarks on the accuracy of the methods will be made.

Thursday, July 23

Time: 9-10am

Title: Singular vectors as solutions of triangular systems of equations

Speaker: Zlatko Drmač

Abstract:

Let A be a triangular matrix that is well–conditioned with respect to inversion and let $A = U\Sigma V^\tau$ be the SVD of A. In the floating–point Jacobi SVD algorithm we compute the decomposition

$$\tilde{U}\tilde{\Sigma} = (A + \delta A)\hat{V}, \tag{1}$$

where δA is the backward error, $\tilde{U}$ approximates the left singular vectors of A, $\tilde{\Sigma}$ approximates the singular values of A and $\hat{V}$ is exact product of certain exact Jacobi rotations. The right singular vectors of A are approximated by the computed product $\overline{V} \approx \hat{V}$ of floating–point Jacobi rotations. The question is whether or not we can restore $\overline{V}$ or $\hat{V}$ a posteriori, using only the matrices A, $\tilde{U}$ and $\tilde{\Sigma}$, and the fact that $V = A^\tau U\Sigma^{-1} = A^{-1}U\Sigma$. It is known that the formula $V = A^\tau U\Sigma^{-1}$ does not give satisfactory results. On the other hand, we show that the formula $V = A^{-1}U\Sigma$ performs surprisingly well. More precisely, if $\tilde{V}$ is the floating–point approximation of $A^{-1}\tilde{U}\tilde{\Sigma}$, then $\tilde{V} = (I + A^{-1}\Delta A)\hat{V}$, where $A^{-1}\Delta A$ is small if $\| \, |A^{-1}| \cdot |A| \, \|_2$ is moderate. (Here the matrix absolute value is defined element–wise and $\|\cdot\|_2$ is the spectral matrix norm.) An important feature of the derived bound is that it is invariant under row scalings of A because of a special structure of the perturbation matrices δA and ΔA. We combine this technique with the Cholesky SVD of Fernando and Parlett to obtain an new implementation of the Jacobi SVD algorithm that reaches the efficiency of the bidiagonalization based QR algorithm. We also discuss some interesting details related to efficient implementation of the Jacobi SVD algorithm.

Time: 10:30-11am

Title: Weyl-type relative perturbation bounds for eigenvalues of hermitian matrices

Speaker: Juan M. Molera

Abstract:

We present a Weyl-type relative bound for eigenvalues of additive hermitian perturbations of (not necessarily definite) hermitian matrices. This bound, which was already known in the definite case, is shown to be also valid in the nondefinite case. The size of the perturbations for which the bound holds is discussed and we show that the bound cannot be valid for perturbations of arbitrary size without additional assumptions. Some connections with previous bounds in the literature are discussed, together with an interpretation of the perturbation bound in terms of matrix differential calculus, as an upper bound of a first derivative of a function of the unperturbed matrix.

Time: 11:11:30am

Title: Accuracy and Orthogonality in the Restarted Lanczos Method

Speaker: J. Wu.

Abstract:

For real symmetric eigenvalue problems, the Lanczos algorithm is an effective method. One way of using it with limited amount of memory is to restart, i.e., use the Lanczos algorithm to build a Krylov subspace basis until the allowed memory is filled then reduce the basis to a smaller size so that the Lanczos iterations may continue again. When the Lanczos algorithm is implemented in exact arithmetic, the basis vectors are orthogonal to each other. However, in finite precision arithmetic, the basis vectors may lose orthogonality. This loss of orthogonality may cause the Lanczos method to generated spurious solutions. To correct this problem, the Gram-Schmidt orthogonalization procedure is invoked in addition to the orthogonalization operations in the Lanczos algorithm. Since the Gram-Schmidt procedure is relatively expensive, we would like to minimize the number of times it is performed. When the Lanczos algorithm does not restart, it can compute accurate eigenvalues by only maintaining semi-orthogonality among the Lanczos vectors. This talk will exam the relation between orthogonality of the basis vectors and the accuracy of the solutions for a particular restarted Lanczos method, the thick-restart Lanczos method. We will show how to maintain semi-orthogonality among the basis vectors and demonstrate that the solutions can be computed accurately when semi-orthogonality is maintained.

Time: 11:30-12pm

Title: A Characterization of the Multiplicity of Real Eigenvalues of General Matrices

Speaker: H.D. Scolnik

Abstract:

We prove that if H is a real $n \times n$ upper Hessenberg matrix without zeroes in the subdiagonal and $B(\lambda) = H - \lambda I$, then the projection of the n–th column of B onto the orthogonal subspace to the one spanned by the remaining columns is zero when λ is any real eigenvalue of H. This result leads to algorithms based on the variable projections method for computing real eigenvalues.

We also present new results for computing multiplicities of real eigenvalues: one is based on new analytical expressions for the high order derivatives of pseudoinverses and projections, and the other on a new theorem concerning a general formula for the eigenvectors of a given eigenvalue. These results may be particularly useful for structural vibration design and control system design problems.

Time:12-12:30

Title: On the solution of generalized quadratic eigenvalue problems

Speaker: Louis Komzsik

Abstract:

Generalized eigenvalue problems appearing in industrial applications usually contain two or three matrices. On the other hand, most eigenvalue algorithms described in the scientific journals only deal with a single matrix case, sometimes with the two matrix pencil and almost never with the 3 matrix case of the quadratic eigenvalue problems.
The topic of this paper is to discuss the specific issues one faces when solving an industrial quadratic eigenvalue problem. These issues include the transformation to the canonical form which will be the basis of the mathematical solution, the recovery and the analysis of the physical solution. The transformation to a canonical form of twice the size of the original problem is never explicitly executed in commercial environments. The paper will describe a very efficient method of implicitly executing the canonical operator update. Some practical results from large industrial quadratic eigenvalue problems solved by MSC/NASTRAN will also be presented.