Theory and Numerics of Matrix Eigenvalue Problems

James W. Demmel (University of California, Berkeley)
Nicholas J. Higham (University of Manchester)
Peter Lancaster (University of Calgary)

November 22–27, 2003

Matrix eigenvalue problems arise in many applications in science and engineering, ranging from the dynamical analysis of structural systems such as bridges and buildings to theories of elementary particles in atomic physics. Many current engineering design processes depend on the reliable computation of eigenvalues of matrices or matrix polynomials of possibly huge dimensions and often having particular structures. Underlying these computations are theory and numerical algorithms developed over the last 50 years, since the introduction of digital computers.

The aim of this workshop was to bring together researchers in the theory and numerical solution of eigenvalue problems with a view to surveying the state of the art, promoting collaboration, and making progress on the many challenging problems in this area. Also invited were users of these algorithms in applications. A unique feature of the workshop was that researchers from all parts of the spectrum from core linear algebra to numerical linear algebra to applications were able to interact and work together intensively for the duration of the workshop.

The workshop had 36 attendees. One confirmed participant had to withdraw the week before the meeting due to a family illness. The organizers encouraged everyone to fully participate in the workshop by giving a talk, and indeed every attendee gave a 30-minute presentation.

The attendees spanned the range from a PhD student and a recent PhD of just 2 weeks, to emeritus professors. Four participants were women. About one third of the attendees were from Europe, the rest from North America.

Several themes from the workshop are now discussed. We have not attempted to mention every speaker.

1 Structured Eigenvalue Problems

Eigenvalue problems with structure are increasingly being encountered in applications such as control theory and dynamical systems. Here the matrices may be, for example, (point or block) Toeplitz, Hankel or Hamiltonian, perhaps combined with other properties such as symmetry or skew-symmetry, or may possess group or algebraic structures. The aim is to develop theory and numerical methods that respect the structure and spectral properties of the problem, because this can lead to significantly faster and/or more accurate solutions. For example, an algorithm for solving a real Hamiltonian eigenvalue problem should ideally iterate exclusively with Hamiltonian matrices and should produce computed eigenvalues with symmetry about the real and imaginary axes. On the theoretical side, it is necessary to develop canonical forms, perturbation theory that measures how the eigensystems change when the parameters in the matrix are perturbed, and information about the possible eigensystems.

Bini (Pisa) and Gu (Berkeley) talked about work by their respective groups on efficient implementation of the QR algorithm for companion matrices. A companion matrix is the result of expressing a polynomial root finding problem as a matrix eigenvalue problem (this is how MATLAB finds polynomial roots), and is also of independent interest. The idea is to adapt the QR algorithm to exploit the structure in the Hessenberg iterates that results from the fact that a companion matrix is a low rank perturbation of a matrix with displacement structure. By exploiting the structure it is possible to implement the entire QR algorithm in $O(n^2)$ operations, as opposed to the usual $O(n^3)$ operations for a full matrix. Achieving this order has been a long-standing open problem and this new work may well result in software packages such as MATLAB adopting the new QR implementation.

A different form of structure was studied in the talks by N. Mackey (Kalamazoo), Tisseur (Manchester) and Higham (Manchester), namely, automorphism group and Lie and Jordan algebra structure associated with a scalar product on $\mathbb{R}^n$ or $\mathbb{C}^n$. Mackey described Givens- and Householder-like transformations in an automorphism group for introducing zeros into vectors, which are essential tools in structured matrix computations. She also discussed to what extent when A itself has structure the factors in standard decompositions (polar decomposition, matrix square roots, matrix sign decomposition, singular value decomposition) inherit this structure. Tisseur investigated the structured mapping problem of characterizing all structured matrices A such that $Ax = b$ for given x and b; this problem arises in computing structured backward errors and condition numbers, for example. Higham derived conditions on a matrix function f for it to preserve automorphism group structure. The principal matrix square root was shown to be one such function and he developed methods for computing the square root that exploit the group structure.

2 Polynomial Eigenvalue Problems

Polynomial eigenvalue problems (PEPs) have been attracting growing interest in recent years. The special case of the quadratic eigenvalue problem (QEP) $(\lambda^2 A + \lambda B + C)x = 0$ is fundamental to vibration problems with damping.

Both Bai (U.C. Davis—talk given by Demmel in his absence) and Meerbergen (Free Field Technologies) described Krylov subspace methods for the QEP that exploit the structure of the problem and are more efficient in storage and computation than the application of standard techniques to the generalized eigenvalue problem (GEP) that results when the QEP is linearized. Bai's approach culminates in projecting the QEP to a lower dimensional QEP that retains symmetry or skew-symmetry in the coefficient matrices. These new techniques are expected to have a major impact on the numerical solution of PEPs as well as on model reduction.

Related to the problem of solving the QEP is that of finding a solvent of the matrix polynomial $AX^2 + BX + C = 0$. Guo (Regina) examined a particular structured QEP modelling a gyroscopic system. After a Cayley transform, a solvent can be obtained, by solving a much-studied matrix equation of the form $X + A^T X^{-1} A = Q$. The method of cyclic reduction provides an efficient iterative method for the solution of this equation.

Lancaster (Calgary) discussed an inverse quadratic eigenvalue problem: Given complete data on eigenvalues and eigenvectors for an underdamped system of size n, determine the three real $n \times n$ coefficient matrices. Three approaches were compared using spectral theory, structure preserving similarities, and factorization theory.

S. Mackey (Manchester) discussed "palindromic matrix polynomials", which in the quadratic case have the form $\lambda^2 A + \lambda B + A^T$, where $B = B^T$. These problems arise in some recent applications in SAW filters and in modelling the German railway system. Mackey showed how to construct linearizations that are themselves palindromic and which enable this structure to be exploited numerically.

Mehrmann (Berlin) discussed numerical solution of parametric eigenvalue problems in robust control, emphasizing the importance of H_∞ norm computations and showing how to avoid the use of Riccati equations. Embedding into matrix pencils and using structured eigensolvers were emphasized.

Koev (MIT) described how to accurately compute the eigensystem of a totally nonnegative matrix. Working exclusively with the representation of the matrix into bidiagonal factors, he showed how the matrix can be reduced to (nonsymmetric) tridiagonal form, after which a symmetrization converts to a bidiagonal SVD problem that is solved by the standard QR iteration. This guarantees that all eigenvalues are computed to high relative accuracy, whereas conventional algorithms guarantee no relative accuracy at all.

Watkins (Pullman) presented a unifying viewpoint on GR algorithms for products of matrices. He showed that several methods previously thought of as extensions of the QR algorithm, including an algorithm for computing the periodic Schur form, are in fact special cases of the QR algorithm.

3 Applications

Several speakers described applications of eigenvalue problems in science and engineering.

Liu (SLAC, Stanford) described work in the DOE TOPS SciDAC project on large eigenvalue problems in cavity design for particle accelerators. One important program for solving these problems is Omega3P, a parallel distributed-memory finite-element code for electromagnetic modeling of complex 3D structures. A variety of algorithms for solving the resulting large sparse eigenproblems in parallel were presented and compared.

Frommer (Wuppertal) discussed Krylov subspace methods for computing $\mathrm{sign}(Q)b$, where Q is a very large, sparse, complex matrix arising in lattice chromodynamics computations in particle physics. He compared the use of rational approximations to the sign function developed over the last 15 years by numerical analysts, an old rational approximation of Zolotarev that requires the spectrum to lie in $[-b,-a]\cup[a,b]$ with a and b known, and more recently developed Krylov techniques. He concluded that Lanczos based projection techniques are often close to optimal.

Dhillon (Austin) described inverse eigenvalue problems and matrix nearness problems arising in the timely problem of wireless communications, comparing finite-step methods with alternating projection methods. He introduced the concept of tight frames and explained its relevance.

Carrington (Montreal) discussed large eigenvalue problems in quantum dynamics. He described the chemistry background to the eigenvalue problems he deals with and explained how chemists currently solve these problems.

4 Theory

Most of the talks included theory to a greater or lesser extent. In this section we include talks whose main purpose was to develop new theoretical results or techniques, and which have not already been mentioned in earlier sections.

Byers (U. Kansas) showed that an arithmetic can be defined for matrix pencils $A - \lambda B$, which can lead to more concise and elegant derivations of algorithms.

Zhou (Calgary) considered the eigenvalue problem for non-selfadjoint analytic matrix functions of two complex variables. She showed that, by using Newton diagram methods it is possible to identify derivatives of eigenvalues branching from a multiple semisimple eigenvalue.

Li (William and Mary) described eigenvalue and singular value inequalities and the relation with Littlewood-Richardson coefficients.

Tsatsomeros (U. Washington) gave a theorem on spectrum localization for almost skew-symmetric matrices.

Rodman (William and Mary) described a class of robustness problems in matrix analysis. For the algebraic Riccati equation and the generalized polar decomposition he investigated the stability of the solution/factors under perturbations, according to two different definitions of stability.

5 Sparse Eigenvalue Computations

In many practical applications the matrices of interest are huge, but sparse: relatively few entries are nonzero. Moreover, typically only a few eigenpairs are required, so that much computation time and storage can be saved. The operations that can be performed are limited to matrix-vector products and sparse direct or preconditioned iterative solvers.

Sorensen (Rice) gave a theory of convergence of polynomial restarted Krylov methods for eigenvalue computations, motivated by trying to understand an earlier paper by Beattie, Embree and Rossi. The main questions were to understand how nonnormality and the distribution of the starting vector with respect to the desired subspace affected convergence. By using results from complex approximation theory, asymptotic convergence rates (if not the actual error bounds) as a function of iteration number are accurately predicted.

Lehoucq (Sandia, Albuquerque) gave a numerical comparison of algorithms for computing a large number of eigenvectors of a generalized eigenvalue problem arising from a modal analysis of elastic structures using preconditioned iterative methods. In particular, he examined algebraic-multigrid-based alternatives to shift-invert Lanczos with a direct solver. On a set of difficult numerical examples, Locally Optimal Block Preconditioned Conjugate Gradient method (LOBPCG) emerged as the best alternative, provided a sufficiently large block was used.

Knyazev (Denver, Colorado), asked “Is there life after the Lanczos method?” and concluded positively by describing his LOBPCG method (also discussed by Lehoucq). The basic idea (of the unblocked method) is to minimize the Rayleigh Quotient on the subspace spanned by the current approximation, the current residual, and the previous approximation. Numerical tests indicate that it has the same linear (but not super-linear) convergence speed as Lanczos.

Spence (Bath) discussed the effect of inexact solves on the convergence of inverse iteration, showing that a quadratic rate of convergence can be obtained if residuals of the solves are suitably bounded.

6 Conclusions

The workshop provided an excellent atmosphere and facilities for interaction between attendees and for research collaboration. The ample supply of meeting rooms, computer terminals in every bedroom, ready access to printers, and well-stocked 24-hour lounge and kitchen, made working during the workshop easy and pleasant.

It is a pleasure to thank the staff at BIRS for the excellent facilities and support.

Existing collaborations were continued, and new ones initiated.

One of us (JWD) formulated plans for a new release of the Fortran linear algebra library LAPACK through discussions with participants during the workshop. Another of us (NJH) is involved in two grant proposals joint with other participants begun during the workshop and now submitted. The third of us (PL) initiated plans for collaboration on a book on matrix analysis in indefinite scalar product spaces, and for a BIRS application on a similar topic.

The participants appreciated the chance to interact with people they would not normally meet at other conferences. We feel that our aim of encouraging interaction between researchers across the spectrum from core linear algebra to numerical linear algebra to applications was fully justified and we are confident that in due course, as ideas and collaborations begun at the workshop reach fruition, the workshop will prove to have achieved its aims.