NORTH-HOLLAND MATHEMATICS STUDIES 127

LARGE SCALE EIGENVALUE PROBLEMS

Jane CULLUM

Ralph A. WILLOUGHBY

LARGE SCALE EIGENVALUE PROBLEMS

Proceedings of the IBM Europe Institute Workshop on Large Scale Eigenvalue Problems held in Oberlech, Austria, July 8-12, 1985

Edited by:

Jane CULLUM

and

Ralph A. WILLOUGHBY

Mathematical Sciences Department
IBM Thomas J. Watson Research Center
Yorktown Heights, New York, U.S.A.

NORTH-HOLLAND – AMSTERDAM • NEW YORK • OXFORD • TOKYO

ISBN: 0 444 70074 9

Publishers:
ELSEVIER SCIENCE PUBLISHERS B.V.
P.O. Box 1991
1000 BZ Amsterdam
The Netherlands

Sole distributors for the U.S.A. and Canada:
ELSEVIER SCIENCE PUBLISHING COMPANY, INC.
52 Vanderbilt Avenue
New York, N.Y. 10017
U.S.A.

Library of Congress Cataloging-in-Publication Data

IBM Europe Institute on Large Scale Eigenvalue
Problems (1985 : Oberlech, Austria)
Large scale eigenvalue problems.

(North-Holland mathematics studies ; 127)
Includes bibliographies and index.
1. Eigenvalues--Congresses. 2. Eigenvalues--Data
processing--Congresses. I. Cullum, Jane K., 1938-
II. Willoughby, Ralph A. III. Title. IV. Series.
QA193.I26 1985 512.9'434 86-13544
ISBN 0-444-70074-9

PRINTED IN THE NETHERLANDS

PREFACE

The papers which are contained in this book were presented at the IBM Europe Institute Workshop on Large Scale Eigenvalue Problems which was held at Oberlech, Austria, July 8-12, 1985. This Workshop was one in a series of summer workshops sponsored by the IBM World Trade Corporation for European scientists.

The unifying theme for this Workshop was 'Large Scale Eigenvalue Problems'. The papers contained in this volume are representative of the broad spectrum of current research on such problems. The papers fall in four principal categories:

(1) Novel algorithms for solving large eigenvalue problems
(2) Use of novel computer architectures , vector and parallel
(3) Computationally-relevant theoretical analyses
(4) Science and engineering problems where large scale eigenelement computations have provided new insight.

Most of the papers in this volume are readily accessible to the reader who has some knowledge of mathematics. A few of the papers require more mathematical knowledge. In each case, additional papers on these subjects are available from the authors of these papers. The interested reader can obtain such reprints by writing to the appropriate authors. A complete list of the names and addresses of the authors is included at the end of this book. A corresponding list of the Workshop speakers who were not able to submit papers is also included. Interested readers should consult both lists.

Jane Cullum
Ralph A. Willoughby
Program Organizers
April 1986

TABLE OF CONTENTS

Introduction to the Proceedings. 1
Jane Cullum and Ralph A. Willoughby

High Performance Computers and Algorithms from Linear Algebra. 15
Jack J. Dongarra and Daniel C. Sorensen

The Impact of Parallel Architectures on the Solution of Eigenvalue Problems. 37
Ilse C.F. Ipsen and Youcef Saad

Computing the Singular Value Decomposition on a Ring of Array Processors. 51
Christian Bischof and Charles Van Loan

Quantum Dynamics with the Recursive Residue Generation Method:
Improved Algorithm for Chain Propagators. 67
Robert E. Wyatt and David S. Scott

Eigenvalue Problems and Algorithms in Structural Engineering. 81
Roger G. Grimes, John G. Lewis and Horst D. Simon

A Generalised Eigenvalue Problem and the Lanczos Algorithm. 95
Thomas Ericsson

Numerical Path Following and Eigenvalue Criteria for Branch Switching. 121
Yong Feng Zhou and Axel Ruhe

The Lanczos Algorithm in Molecular Dynamics: Calculation of Spectral Densities. 143
Giorgio Moro and Jack H. Freed

Investigation of Nuclear Dynamics in Molecules by Means of the Lanczos Algorithm. . . . 163
Erwin Haller and Horst Köppel

Examples of Eigenvalue/Vector Use in Electric Power System Problems. 181
James E. Van Ness

A Practical Procedure for Computing Eigenvalues of Large Sparse Nonsymmetric Matrices. 193
Jane Cullum and Ralph A. Willoughby

Computing the Complex Eigenvalue Spectrum for Resistive Magnetohydrodynamics. . . . 241
Wolfgang Kerner

Ill Conditioned Eigenproblems. 267
Francoise Chatelin

Stably Computing the Kronecker Structure and Reducing Subspaces of Singular Pencils A - λ B for Uncertain Data. 283
James Demmel and Bo Kågström

Addresses of Authors and Other Workshop Speakers. 325

Index to Proceedings. 329

Large Scale Eigenvalue Problems
J. Cullum and R.A. Willoughby (Editors)

INTRODUCTION TO PROCEEDINGS

Jane Cullum
Ralph A. Willoughby

IBM T. J. Watson Research Center
Yorktown Heights, New York 10598
USA

We provide a brief summary of each paper contained in this volume, provide some indications of the relationships between these papers, and provide a few additional references for the interested reader.

The papers included in this volume can be classified into the following four categories.

(1) Novel algorithms for solving large eigenvalue problems

See the papers by Zhou and Ruhe, by Kerner, and by Cullum and Willoughby.

(2) Use of novel architectures for solving eigenvalue problems

See the papers by Dongarra and Sorensen, by Ipsen and Saad, and by Bischof and Van Loan. The paper by Dongarra and Sorensen addresses both the question of restructuring the EISPACK library routines [1977] of restructuring the LINPACK library routines [1979] for novel architectures.

(3) Computationally-relevant theoretical analyses

See the papers by Demmel and Kagstrom, by Chatelin, and by Ericsson.

(4) Examples from science and engineering where large scale eigenvalue and eigenvector computations have provided new insight into fundamental properties and characteristics of physical systems, both those existing in nature and those which have been constructed artificially.

See the papers by Grimes, Lewis, and Simon, by Van Ness, by Kerner, by Moro and Freed, and by Haller and Koppel.

Most of the currently active areas of research in modal analysis are represented in this volume. With the exception of the three papers dealing with novel architectures which are presented first, the papers are presented in an ordering which takes us from the 'easiest' problems, the real symmetric eigenvalue problems, to the most difficult ones, the computation of the Kronecker canonical forms of general matrix pencils.

Many engineering and scientific applications yield very large matrices. Historically, the sizes of the matrices which must be used have grown as the computing power has grown. Therefore, there is much interest in understanding how to exploit the new vector and parallel architectures in such computations. The first paper by Dongarra and Sorensen addresses two basic questions dealing with such architectures. First, they look at the types of computers which are currently available and at those which should be available within the next few years. They then examine the question, how do we exploit such architectures for linear algebra computations? They include introductory descriptions of classifications for the various arithmetic engines and storage hierarchies. This discussion is followed by a table of advanced computers, proposed and existing, together with tables of characteristics of these machines.

Dongarra and Sorensen include some discussion of data communication and how that relates to algorithm performance. The cost of algorithm execution can be dominated by the amount of memory traffic rather than by the number of floating point operations involved. A performance classification is given for algorithms on a vector computer; scalar, vector, and super-vector. Data management and synchronization add to the complications in designing algorithms for parallel computers.

The authors also address such issues as program complexity, robustness, ease of use, and portability, each of which plays an important role in the analysis. Certain basic vector and matrix-vector operations are fundamental to many linear algebra algorithms and these are examined very carefully. The authors contend that it is possible to achieve a reasonable fraction of the peak performance on a wide variety of different architectures through the use of program modules that handle certain basic procedures. These matrix-vector modules form an excellent basis for constructing linear algebra programs for vector and parallel processors. With this philosophy, the machine dependent code is isolated to a few modules. The basic routines provided

in the linear equation solving package, LINPACK [1979], and in the EISPACK library [1977] for eigenvalue and eigenvector computations, achieve vector but not super-vector performance.

The paper by Ipsen and Saad presents a brief survey of recent research on multiprocessor (parallel) architectures and on algorithms for numerical linear algebra computations which take advantage of such architectures. The emphasis is on algorithms for solving symmetric eigenvalue problems. Basic terminology is introduced and data communication problems such as start-up times and synchronization are discussed. Three loosely coupled architectures; a processor ring, a two-dimensional processor grid and the hypercube are considered.

The paper by Bischof and Van Loan looks at one of the parallel architectures which is available today, the LCAP configuration designed by Enrico Clementi of IBM and at the problem of implementing a block Jacobi, singular value decomposition algorithm on LCAP. LCAP consists of ten Floating Point Systems FPS-164/Max array processors connected in a ring structure via large bulk memories. The algorithm is a block generalization of a parallel Jacobi scheme which appeared in the paper by Brent, Luk, and Van Loan [1985]. The parallel procedure developed in the Bischof and Van Loan paper could also be applied to the real symmetric eigenvalue problem. The authors however did not achieve the speedups which others had predicted. This paper is a good illustration of the difficulties and considerations encountered in translating an idea for a parallel algorithm into a practical procedure.

Many of the papers in this volume use procedures which rest upon the so-called Lanczos recursion. For more details and background information on this recursion, the reader is referred to Parlett [1980] and Cullum and Willougby [1985]. Both of these books contain bibliographies with references to much of the recent research on Lanczos procedures. A brief survey of recent research in this area is contained in Cullum and Willoughby [1985b].

The paper by Wyatt and Scott provides an example of the use of a real symmetric Lanczos procedure. The objective is to compute time dependent quantal transition probabilities. These transition probabilities are obtainable from differences of survival amplitudes for surviving in a particular state at time t given that we started in that state at time t=0. These survival amplitudes can be computed if all of the eigenvectors of an associated Hamiltonian operator are known. However, an eigenvector decomposition of this operator cannot be obtained easily.

The Lanczos algorithm provides real symmetric tridiagonal representations of this Hamiltonian which reduce (at least theoretically) the survival amplitude computations to computations of the eigenvalues of tridiagonal matrices and computations of the first components of the eigenvectors of these tridiagonal matrices. Both of these computations are reasonable. Specifically, Wyatt and Scott compute

$$\sum_{\alpha=1}^{M} s_{1\alpha}^{2} \exp^{-iE_{\alpha}t} \tag{1.1}$$

where M is the size of the Lanczos matrix being used; $s_{1\alpha}$ denotes the first component of the eigenvector of that tridiagonal matrix corresponding to the eigenvalue E_{α}.

Wyatt and Scott use the Lanczos recursion with no reorthgonalization because they have very large matrices and therefore the amount of computer storage which would be required by the Lanczos methods which require reorthogonalization would be too large. However, if the Lanczos vectors are not reorthogonalized, then extra or 'spurious' eigenvalues can appear among the eigenvalues of the Lanczos matrices. These are not genuine representations of the eigenvalues of the original matrix, and these eigenvalues must be handled appropriately if the results are to have any validity. Such eigenvalues can however be identified very easily, see Cullum and Willoughby [1985], and in earlier papers Wyatt and coauthors were using this identification test. However, the main point of this current paper is that for their particular application it is not necessary to sort these Lanczos eigenvalues. All of the computed quantities can be used in their computations and they still get correct results. They demonstrate this numerically. A partial explanation for this follows directly from the characterization of these spurious eigenvalues given in Cullum and Willoughby [1985]. The spurious eigenvalues are eigenvalues of a particular submatrix of the Lanczos matrix being considered and because of this the first components of their eigenvectors are pathologically small. Therefore, their contributions to the sum in Eqn.(1.1) are pathologically small.

Algorithms exist for computing eigenvalues of large, real symmetric generalized eigenvalue problems $Ax = \lambda Bx$ where A and B are real symmetric and B is positive definite. However, there are many open questions regarding this problem when neither A nor B is not definite. Grimes, Lewis and Simon focus on this problem. They first provide a survey of the types of eigenvalue/eigenvector problems encountered in structural engineering problems. They then

discuss extensions of procedures designed for the standard real symmetric eigenvalue problem to such problems. When B is not definite, these 'symmetric' problems are genuinely nonsymmetric and many numerical difficulties can be experienced in trying to solve them. Unfortunately, in many structural problems, B is only positive semidefinite.

Grimes, Lewis and Simon outline the two most common classes of structural engineering problems, vibration and buckling analyses. In vibration analyses, the higher frequency modes of vibration are not important because it is unlikely that they will be excited. In buckling analyses usually only the smallest positive eigenvalue and corresponding eigenvector are required. Because of the slow convergence of such eigenvalues in the numerical algorithms designed thus far, it is typical to factor one or more of the matrices involved in the eigenvalue computations and to use such factorizations to transform the desired eigenvalues into eigenvalues with dominant magnitudes. When matrices $A - \sigma B$ are factored, the method being used is called a shift and invert method.

Grimes, Lewis, and Simon also list some of the practical considerations which must be faced by any numerical analyst designing algorithms which are to be used within the constraints of structural engineering packages. For example, typically there are restrictions on the way the required data is stored and can be accessed. For this reason block versions of modal algorithms have a number of desirable features for structural engineering calculations. A detailed discussion of a block Lanczos procedure for the problem $Ax = \lambda Bx$, where A and B are real symmetric and B is positive semidefinite, will appear shortly. See reference 7 in the Grimes, Lewis and Simon paper. The authors point out that in some situations it is necessary to use a model which involves large nonsymmetric matrices and/or solve nonlinear eigenvalue problems. However, satisfactory eigenvalue procedures for these nonsymmetric structural problems have not yet been devised. This is an open area for research.

The paper by Ericsson derives some of the computationally-important theoretical properties of generalized eigenvalue problems $Kx = \lambda Mx$, where K and M are real symmetric matrices. In the first part of his paper he develops the analysis which he needs for examining three types of procedures for computing eigenvalues: (1) Inverse Iteration; (2) Power Methods; and (3) Lanczos Methods. In most of his paper he assumes that K is a nonsingular matrix and that the pencil of matrices $(K - \lambda M)$ is nonsingular. Equivalently, this means that K and M do not have

a common null vector. He proves however, that any theorem which is valid under those assumptions must also be valid when K is singular so that there is no loss of generality in his arguments.

Under these conditions these generalized problems can behave very nonsymmetrically and Ericsson illustrates that type of behavior with examples. He then focuses on the problem for M a positive semidefinite matrix. This type of problem is encountered frequently in structural engineering problems. He uses the analysis which he has developed to look at the ability of the three procedures listed above to compute good approximations to the eigenvectors of the given generalized problem. This analysis points out a basic difficulty with both Lanczos methods and with power methods, namely keeping the eigenvector approximations in the proper part of the space. Since M is singular, the generalized eigenvalue problem has infinite eigenvalues. If the starting vector in the Lanczos procedure contains a nonzero projection on the subspace spanned by the eigenvectors corresponding to these infinite eigenvalues, then this projection may grow and produce significant errors in the resulting Ritz vectors computed. This is a serious problem which must be dealt with numerically. Ericsson also shows that this problem can happen when M is nonsingular but very ill-conditioned. He also addresses the question of obtaining error estimates for computed Ritz vectors in the case that M is singular.

As mentioned in the Grimes, Lewis and Simon paper, a more accurate representation of a particular structural problem may require the solution of a nonlinear eigenvalue problem: Find u and λ such that $G(u, \lambda) = 0$ where G is a nonlinear vector function, u is a vector of the same dimension, and λ is a scalar.

Zhou and Ruhe examine such problems, not only in the context of solving nonlinear eigenvalue problems, but as a general path following problem. Typically, the manifold of solutions consists of a curve or path as illustrated in several figures in this paper. Bifurcation points in this curve, places where several paths meet, and turning points, where the curve has the hyperplane $\lambda = c$ (c constant) are of interest. These are points where the Jacobian of G with respect to u is singular. In the linear problem there is a bifurcation point at each eigenvalue. They propose a modification to the Euler-Newton path following algorithm which uses the solution of a linear eigenproblem to give both a prediction of the position of singular points and the direction of bifurcating branches. Several examples illustrating this idea are included.

The next level of difficulty in dealing with algorithms for solving eigenvalue problems is to design procedures which are applicable to complex symmetric problems. For diagonalizable, complex symmetric matrices one can write down a Lanczos recursion which is completely analogous to the real symmetric recursion. The left and the right eigenvectors of a complex symmetric matrix are identical so only one set of Lanczos vectors has to be generated. As is shown in Wilkinson [1965], in the general nonsymmetric case, it is necessary to replace the single Lanczos recursion which is used in the real symmetric case by a set of two such recursions. One of these recursions uses the given matrix and the other recursion uses the transpose of the given matrix. The papers by Haller and Koppel and by Moro and Freed describe applications where eigenvalue and eigenvector computations are used to obtain basic physical properties of molecular systems and the matrices involved are complex symmetric. Haller and Koppel consider both real symmetric and complex symmetric matrices.

Moro and Freed are studying molecular motion. The information obtained from spectroscopic or scattering techniques yields only macroscopic responses to external perturbing influences. The objective in these studies is however, to understand the basic mechanisms controlling the molecular motion. Moro and Freed describe the connections between experimental measurements, spectral density computations and the identification of these basic mechanisms. In practice, different theoretical models for the underlying mechanisms are assumed and then comparisons of the resulting macroscopic quantities are made with experimental measurements. These comparisons require spectral densities.

Under certain assumptions, the computations of the spectral densities can be reduced to the computation of the effect of the resolvent of a certain operator on certain vectors. The form of the operator and of the particular vectors depends upon the system being studied. The authors develop these relationships. They then show how the Lanczos algorithm can be used to obtain a 'tridiagonal' representation of the operator, and how this representation can be used to reduce the required spectral density computations to computations of continued fractions whose coefficients are simply the entries of the tridiagonal matrices generated by the Lanczos procedure.

The main part of the Moro and Freed paper considers problems where the operator can be symmetrized so that it is either a real symmetric or a complex symmetric operator. In both of

these cases the Lanczos recursions reduce to a single recursion and the Lanczos tridiagonal matrices are symmetric, either real symmetric or complex symmetric. In the last section of their paper the authors extend some of their ideas to more general nonsymmetric operators.

Haller and Koppel are also looking at problems in molecular dynamics. They have attacked the very difficult and interesting problem of modeling the vibronic coupling in polyatomic molecules. They have obtained a model which reproduces gross features of complex experimental spectra. From this they can make several inferences. The basic computation which is required is the determination of the spectral distribution. Because of the sizes of the matrices involved, it is not possible to use standard eigenelement algorithms for these computations. The Lanczos algorithm with no reorthgonalization plays a critical role in their computations. In this paper the authors consider matrices up to size 40800. However, they want to consider matrices of size up to 10^6. They use the computed results to support their theoretical models.

Electric power systems problems yield some of the most difficult nonsymmetric eigenvalue/eigenvector problems. Van Ness summarizes and illustrates these types of problems. Modern power systems consist of many generating stations and load centers connected together by an electrical transmission system. Small disturbances in such systems can be studied by using eigenanalysis on linearizations of the system equations around some nominal operating state. The objective of this analysis is to determine whether or not the linearized system has any eigenvalues with positive real parts, and to determine the sensitivities of such eigenvalues and of the eigenvalues with small negative real parts, to perturbations of parameters in the model. Sensitivity analysis requires the computation of eigenvectors. A history of the study of several oscillation problems in power systems is presented.

Nonsymmetric problems are considered in the paper by Cullum and Willoughby. The objective is to devise a Lanczos procedure for computing eigenvalues of large, sparse, diagonalizable, nonsymmetric matrices. The authors propose a two-sided Lanczos procedure with no reorthogonalization which uses both the given matrix A and its transpose. Two sets of Lanczos vectors are generated and the Lanczos matrices are chosen such that they are complex symmetric and tridiagonal matrices. In exact arithmetic the Lanczos vectors generated are biorthogonal. A generalization of the QL algorithm is used to compute the eigenvalues of these matrices. A generalization of the identification test for spurious eigenvalues which was used in

the case of real symmetric problems is found to apply equally well here. Several properties of complex symmetric tridiagonal matrices are derived. Numerical results on matrices of size n = 104 to size n = 2961 demonstrate the strengths of this procedure and the type of convergence which can be expected. For the problems up to size n = 656, the Lanczos results are compared directly with the corresponding results obtained using the relevant EISPACK subroutines. All arithmetic is complex even if the starting matrix is real. Because there is no reorthogonalization, this procedure can be used on very large problems and can be used to compute more than just a few of the extreme eigenvalues of largest magnitude.

Kerner addresses the question of the stability of plasmas which are confined magnetically. Such plasmas play a key role in the research on controlled nuclear fusion. This is an application where large scale eigenvalue and eigenvector computations provide new insight into basic physical behavior. The most dangerous instabilities in a plasma are macroscopic in nature and can be described by the basic resistive magnetohydrodynamic model. A well-chosen discretization of this model transforms this model into generalized eigenvalue problems: $Ax = \lambda Bx$ where A is a general matrix and B is a real symmetric and positive definite matrix. The eigenvalues and eigenvectors of these systems provide knowledge about the behavior of the plasma. Of particular interest are the Alfven modes, and the author studies the effects of the resistivity upon these modes and upon the sound modes. These eigenvalues are found to lie on curves in the plane, and the eigenvalue and eigenvector computations are performed by using a path following technique which uses inverse iteration and a continuation method. Convergence is demonstrated by performing these computations over finer mesh sizes.

Central to the successful computation of eigenelements are both the theoretical stability of the given problem with respect to perturbations in the data and the numerical stability of the algorithm being used to perform the computations. Chatelin addresses both questions. She focuses on defective eigenvalues. Only multiple eigenvalues can be defective. An eigenvalue is defective if its multiplicity as a root of the characteristic polynomial of the given matrix is larger than the dimension of the subspace of eigenvectors associated with that eigenvalue.

Chatelin first looks at the question of condition numbers of eigenvalues and of invariant subspaces. She then shows that the method of simultaneous inverse iteration is not stable if the

subspace being computed corresponds to a defective eigenvalue. She proposes a particular modification of a block Newton method which is stable.

Demmel and Kagstrom present algorithms and error bounds for computing the Kronecker canonical form (KCF) of matrix pencils, $A - \lambda B$, where A and B can be rectangular matrices. For the standard eigenvalue problem, $Ax = \lambda x$, the Jordan canonical form (JCF) provides insight into the behavior of the system under perturbations in the matrix A. The KCF is a generalization of the JCF which can be used to provide similar insight into the generalized eigenvalue problem $Ax = \lambda Bx$ and into general systems involving pencils of matrices $A - \lambda B$ where A and B may be rectangular. The KCF is obtained by applying left and right nonsingular transformations which simultaneously reduce the A and B matrices to block diagonal matrices. Each of these diagonal blocks may have one of three forms. These forms together with the matrix transformations characterize the subspace associated with each block.

Computing features of the KCF can be an ill-posed problem; that is, small changes in the data may result in large changes in the answers. By restricting the class of perturbations allowed, Demmel and Kagstrom look at the question of how much the Kronecker structure can change when perturbations are made in the matrices. These perturbations can occur because of roundoff errors or because of uncertainties in the input data. They then analyze the errors incurred in algorithms for computing Kronecker structures. They are particularly interested in singular pencils, for example when the determinant of $(A - \lambda B) = 0$. The error bounds obtained can be used for determining the accuracy of computed Kronecker features. Their results have applications to control and linear systems and these relationships are discussed. Wilkinson [1979] provides additional introductory comments.

Several other talks were given at the Austrian Workshop which, for a variety of reasons, are not included in this volume. In particular, some of them have already been published elsewhere. G. W. Stewart opened the workshop with a survey of the basic theory and algorithms used in eigenelement analysis and computation. Much of this material can be found, for example, in his book Stewart [1973] and in the book by Golub and Van Loan [1983]. Later in the Workshop program, Stewart presented an algorithm for doing simultaneous iterations on the ring processor, ZMOB, which resides at the University of Maryland.

Kosal Pandey described the Hermitian eigenelement problems which arise in his studies of surface properties of various materials. His results are directly applicable to semiconductor materials. Physical and chemical properties of a surface are determined by its surface structure. Two basic questions addressed are: (1) Determine the atomic structure at the surfaces of such materials; and (2) Determine basic physical characteristics such as how the surface will react with various chemicals. Pandey obtains this type of information by computing eigenfunctions of Schrodinger's equation. Using these computations, he has shown that the accepted buckling reconstruction mechanism for the configuration of atoms at surfaces is valid only for heteropolar surfaces. He has proposed an alternative π – bonding model for homopolar surfaces which fits well with both theoretical arguments and with experimental data. For more details on this work the reader is referred to Pandey [1983, 1983b].

In some applications it is necessary to compute the eigenelements of matrices obtained by simple modifications of a given matrix, for example, a rank one modification. In other applications one or more of the eigenvalues of a system are specified a priori and the user is asked to determine a matrix with those eigenvalues. Gene Golub surveyed some of the work on such problems. Much of his talk is contained in the references Golub [1973] and Boley and Golub [1978].

A key question in any eigenelement computation is how do we know that the answers obtained are meaningful? Beresford Parlett presented the material contained in the paper Kahan, Parlett, and Jiang [1982.] This paper includes error estimates for nonsymmetric problems which are applicable, for example, to Lanczos algorithms for nonsymmetric problems. Parlett gave two talks. The second talk surveyed the recent research on the question of maintaining semi-orthogonality of the Lanczos vectors generated by a Lanczos recursion. The interested reader can find much of this material in the book, Parlett [1980].

A complete list of the authors and coauthors with their full addresses is included at the beginning of this book. A corresponding list of the speakers who do not have papers in this volume is contained at the end of this book. The interested reader can obtain additional references by contacting the authors and speakers directly.