MATHEMATICS OF COMPUTATION
Volume 68, Number 226, Pages 887–891
S 0025-5718(99)01151-5

REVIEWS AND DESCRIPTIONS OF TABLES AND BOOKS

The numbers in brackets are assigned according to the American Mathematical Society classification scheme. The 1991 Mathematics Subject Classification can be found in the annual subject index of *Mathematical Reviews* starting with the December 1990 issue.

5[65-01]—*Numerical analysis, an introduction*, by Walter Gautschi, Birkhäuser, Boston, 1997, xiii + 506 pp., 24 cm, hardcover, $64.50

This textbook for an introductory course on numerical analysis on the upper undergraduate level is off the beaten track in several respects. First of all, it interprets the word analysis in its title in its literal sense; the material has been restricted to areas form mathematical analysis so that there is no treatment of numerical linear algebra. On the other hand, numerical ordinary differential equations cover nearly half of the volume. Secondly, while a scan of the Table of Contents seems to indicate a simple-minded approach, the actual reading of whatever chapter reveals a host of details in the theorems, remarks, observations, etc., which are of interest even to the expert. This aspect is further expanded by "Notes" following each chapter, which cover some of the history as well as further developments, with handy pointers to references. Also, the text manages to keep the balance between an intuitive, readily understandable style of presentation and careful, precise formulations. Finally, each chapter is followed by "Exercises and Machine Assignments" of an unusual multitude and variety. They will constitute a welcome and valuable source of material also for those who teach from a different text.

Altogether, the volume reflects the insight and the experiences of a lifetime's occupation with numerical computation and with teaching numerical analysis, which is a never-ending challenge. I believe that it will give the students the right attitude toward the science and art of numerical computation.

Hans J. Stetter

6[65-01]—*Afternotes goes to graduate school, lectures in advanced numerical analysis*, by G. W. Stewart, SIAM, Philadelphia, PA, 1998, xii + 245 pp., $25\frac{1}{2}$ cm, softcover, $35.00

This monograph consists of a set of notes on numerical analysis written shortly *after* the author lectured in a graduate course given at the University of Maryland. The notes consist of 26 sections corresponding to the lectures and the topics presented fall into four categories: Approximation (9 lectures, 74 pp.), Splines (2 lectures, 20 pp.), Eigensystems (7 lectures, 59 pp.), and Krylov sequence methods (6 lectures, 44 pp.), with two additional lectures giving some classical results on linear and nonlinear iterative methods. Stewart's presentation is intuitive and rapid but, at the same time, clear with considerable attention paid to motivation of a particular approach or algorithm. He rarely gives a proof in complete detail, but instead enough detail is provided about the essential ideas that a mature reader can fill in between the lines. The careful reader will definitely get his/her hands dirty in the process.

There is a useful index and bibliography to the textbook literature. The book contains no exercises or examples and hence is best suited as a reference or supplement for a graduate level course in numerical linear algebra.

The book begins with a discussion of two simple, concrete examples that lead to L_2 and L_∞ approximation problems. From this Stewart extracts their common features and proceeds to develop the notion of best approximation in a normed vector space. The second of these problems concerns best approximation in $C[0,1]$. In this context Bernstein polynomials are developed and used in a sketch of the proof of the classical Weierstrass approximation theorem. To develop a practical means of solving (or more accurately, nearly solving) the L_∞ best approximation problem, a strategy is developed based on the de la Vallee Poussin theorem and characterization of the best polynomial approximation in terms of sign alternation at points of worst error. The method that emerges consists in using Chebyshev polynomials and economization of power series.

The discrete and continuous least squares approximation problems are treated in several lectures where the author introduces orthogonal polynomials and their three-term recurrence, the Gram-Schmidt process and its more stable modification, the QR factorization of a full rank matrix, Householder matrices and their role in constructing the QR decomposition of a general rectangular matrix. There is an extensive discussion/analysis wherein the solution of the normal equations is compared to the use of the QR decomposition to solve for the best least-squares approximation. This analysis deals with three important practical considerations: efficiency, conditioning, and accuracy. In assessing the condition of the two approaches, the author gives a backward (ala Wilkinson) error bound wherein the perturbation in the computed solution, due to roundoff, is projected back into a perturbation on the data in the original problem. I won't divulge the conclusion—you'll have to read the book to find out which method wins!

The two lectures on linear and cubic splines are fairly standard textbook fare with the exception of the local error estimation for cubic splines.

Professor Stewart is a leading expert in and developer of algorithms for numerical linear algebra. This expertise comes to the fore in the last half of the book, which is devoted to the two main problems of linear algebra.

The eigenvalue problem is introduced via the problem of solving a linear system of differential equations. Following this, the author develops some facts needed in the development of numerical algorithms, such as similarity, the Schur decomposition to triangular form, the real Schur block triangular form, as well as the less practical Jordan canonical form. Considerable attention is paid to the power and inverse power (with shift) methods leading up to the QR algorithm with shift. Since the convergence of the QR algorithm is known only under special circumstances, Stewart gives a local error analysis under certain simplifying assumptions. Gradually, the author leads the reader to a viable method for computing eigenvalues of a general complex matrix A; namely, reduction of A (by Householder transformations) to upper Hessenberg form followed by application of the QR algorithm with Rayleigh quotient shifts. Following this, he deals with the case where A is real and utilizes Arnoldi reduction to establish the unitary similarity of A to an unreduced Hessenberg matrix. The actual computations avoiding complex arithmetic are done using a double shift QR step. The section on eigensystems culminates with the presentation of the singular value decomposition (SVD) and its computation using the QR algorithm. An application of the SVD is the determination of the rank of a matrix (an ill-posed problem). The little known theorem of Schmidt is presented and can be used to lend credence to the estimated rank provided by the computed SVD.

It is Stewart's contention that, in discussing the solution of large sparse linear systems by iterative methods, the emphasis should be placed on Krylov sequence methods such as the method of conjugate gradients and its preconditioned variants. Consequently, the classical Jacobi and Gauss-Seidel iterations are treated briefly in the penultimate lecture of the book.

The basic idea of Krylov sequence methods is to consider approximation by vectors in $S_k = \text{span}\{u, Au, A^2u, \ldots, A^{k-1}u\}$. Clearly this space is related to the power method for A, and the vectors in $\{U, Au, A^2u, \ldots, A^{k-1}u\}$ tend toward the dominant eigenvector of A. Stewart shows how one can orthogonalize the Krylov vectors using Arnoldi reduction resulting in a scheme, Arnoldi's method, for determining an approximate eigenpair for A. He also discusses a variant of this method using implicit restarting that effectively eliminates unwanted eigenvalues. For symmetric matrices the Arnoldi decomposition simplifies (upper Hessenberg becomes tridiagonal) and results in the famous Lanczos algorithm. In Stewart's opinion this method, suitably modified using selective orthogonalization, is the method of choice for large sparse symmetric eigenproblems.

The introduction to Krylov methods for solving symmetric positive definite linear systems, $Ax = b$, is given by the method of steepest descent. This gives a sequence $\{x_k\}$ with $x_{k+1} = x_k + \alpha_k s_k$ where the stepsize α_k and direction s_k are chosen to minimize $\varphi(x) = \frac{1}{2}x^T Ax - x^T b$. Stewart then considers the following related problem: given a set $\{s_1, s_2, \ldots, s_k\}$ of linearly independent directions, set $x_{k+1} = x_1 + \sum_{j=1}^{k} \widehat{\alpha_k} s_k$ and seek the coefficients such that $\varphi(x_1 + \sum_{j=1}^{k} \widehat{\alpha_k} s_k)$ is minimized. The result is that, at the minimum, the residual is orthogonal to $\text{span}\{s_1, s_2, \ldots, s_k\}$. To effectively compute this minimizing vector one uses A-conjugate directions and the result is the method of conjugate gradients. Stewart does a rather complete error analysis of the method of conjugate gradients and establishes a well-known error bound in the A-norm. This error bound reveals that the condition number of A is the key to the rate of convergence of the method. It is also shown that the method is error reducing in the Euclidean norm.

To speed the convergence of the conjugate gradient method, one uses a positive definite preconditioner M and considers the problem $M^{-1}Ax = M^{-1}b$. Clearly the matrix $M^{-1}A$ is no longer symmetric and thus the author considers the equivalent system $M^{-1/2}AM^{-1/2}y = M^{-1/2}b$, where $y = M^{1/2}x$. Since $M^{1/2}$ is generally unavailable, the algorithm for the equivalent system must eventually be recast to see that $M^{1/2}$ is not needed. But there is another approach that eliminates any discussion of $M^{1/2}$. One simply applies the conjugate gradient method directly to $M^{-1}Ax = M^{-1}b$ with respect to a different inner product, namely $[\cdot, \cdot] = (M\cdot, \cdot)$, where $(\cdot, \cdot)$ is the usual Euclidean inner product. In this inner product $M^{-1}A$ is symmetric positive definite and the previously established error analysis applies verbatim. The determination of a suitable preconditioner is not easy and the author considers in detail only incomplete LU factorization. It is demonstrated that incomplete LU factorization is always possible for invertible diagonally dominant matrices, and the algorithm is presented.

For those who teach courses in numerical linear algebra or those who are simply interested in the subject, this is a well-written modern book that deserves a place on the office bookshelf. The book is packed with information and insights from one of the leaders in the field.

J. THOMAS KING
DEPARTMENT OF MATHEMATICAL SCIENCES
UNIVERSITY OF CINCINNATI
CINCINNATI, OH 45221-0025

7[65-02, 65F10]—*Iterative methods for solving linear systems*, by Anne Greenbaum, Frontiers in Applied Mathematics 17, SIAM, Philadelphia, PA, 1997, xiii + 220 pp., $25\frac{1}{2}$ cm, softcover, $41.00

This is a stimulating book. It describes recent developments in the theory of iterative techniques for solving large sparse systems of linear equations. The book is not a complete survey; Axelsson's treatise [1] is more definitive in this respect. Instead, as the author points out in the preface, "With this book, I hope to discuss a few of the most useful algorithms and the mathematical principles behind their derivation and analysis... I have tried to include the most *useful* algorithms... and the most *interesting* analysis from both a practical and a mathematical point of view." This selective treatment of the theory is what sets Greenbaum's book apart.

The body of the book is preceded by an introductory chapter containing a concise overview of the state of the art. The remainder of the book is split into two parts. Chapters 2 through 7 form the first part of the book and describe "basic" Krylov subspace methods. The second part of the book (Chapters 8 through 12) is devoted to preconditioning aspects.

The distinction between Hermitian/symmetric and non-Hermitian/nonsymmetric linear systems is made obvious to the reader in Chapter 1, and is reinforced in Chapters 2 and 3 ("Some iteration methods" and "Error bounds for CG, MINRES and GMRES"). When solving (real-) symmetric systems, Krylov methods like MINRES (or CG in the positive definite case) generate a best approximation from a subspace of increasing dimension with a fixed computational effort at every iteration. Furthermore, the concept of preconditioning has a sound theoretical basis, since the convergence is completely determined by the eigenvalue distribution of the coefficient matrix. The explanation of why the convergence of CG/MINRES is described by the eigenvalue spectrum even in the presence of rounding errors is outlined in Chapter 4, "Effects of finite precision arithmetic". These fundamental results are an outcome of the author's research, and their inclusion in a textbook for the first time is very valuable. The inherent difficulty in the nonsymmetric/non-Hermitian case is unravelled in Chapter 6, which provides the answer to the question, "Is there a short recurrence for a near-optimal approximation?" The upshot is a plethora of "open problems" in the nonsymmetric case; in particular, the optimal choice of Krylov subspace method is problem specific, and preconditioning has a very limited theoretical basis. (A consequence of the latter is that the book has little to offer practitioners interested in solving problems other than the model