3[76F05, 76F99, 65-02, 65M70, 76-02, 76D05]—*Dynamic Multilevel Methods and the Numerical Simulation of Turbulence*, by Thierry Dubois, Francois Jauberteau, and Roger Temam, Cambridge University Press, New York, New York, 1999, xix+289 pp., $23\frac{1}{2}$ cm, hardcover, \$59.95

This book describes the multilevel methods for time dependent simulations of incompressible turbulence. The authors have performed extensive research in this area and the book is a good source of the state of art in this methodology.

The first three chapters are surveys about the incompressible Navier-Stokes equations, general turbulence theory, and spectral methods. The surveys are brief but references are given for readers desiring more details. Chapter 4 compares DNS (Direct Numerical Simulation) with various turbulence modeling. Long time behavior of the Navier-Stokes equations is discussed in Chapter 5. This is an area in which the third author has done extensive research. This chapter and Chapter 6, which discusses the separation of scales, form the theoretical basis on the applicability of the multilevel methods to incompressible turbulence simulation. In Chapter 7, the basic procedure of multilevel methods is illustrated through a simple system of ODE, which carries many of the essential ingredients of the more complex Navier-Stokes equations but avoids the functional analysis framework. The last three chapters are about the methodology, implementation details, and numerical results of the multilevel methods applied to Navier-Stokes equations in 2D and 3D, both for the periodic cases and for the well-bounded flows. They are based on several recent papers and give the state of the art in the application of this methodology on turbulence simulation.

The authors argue that "new chapters of numerical analysis will have to be written in relation with the multilevel treatment of large evolutionary problems". It is at least safe to say that multilevel or multiresolution methods will play a more important role in large scale scientific computing in the years to come. There is a great challenge to make multilevel methods more efficient. The factor of CPU time saving of the multilevel methods over DNS in this book is about 2 to 2.5. It would certainly appeal to users if this factor can be increased.

CHI-WANG SHU

4[65-01, 65Fxx, 65Y05, 65Y10, 65Y20]—*Numerical linear algebra for high-performance computers*, by Jack J. Dongarra, Iain S. Duff, Danny C. Sorensen, and Henk A. van der Vorst, SIAM, Philadelphia, PA, 1998, xvii+342 pp., $25\frac{1}{2}$ cm, softcover, \$37.00

This book is meant to provide helpful information on state-of-the-art numerical linear algebra techniques to be used in advanced high-performance computation to a rather heterogeneous community, ranging from graduate students to professionals in computational sciences. In my opinion, one is also very likely to find extremely good advice for implementing recent and advanced algorithms on sequential machines. Quoting from the authors' preface, " ... this book is a major revision of a previous edition of the book, entitled *Solving Linear Systems on Vector and Shared Memory Computers*", published in 1991; indeed, it contains a lot of new material that covers the recent advances in the development of parallel architecture and software.

The book is rich with pointers to available software and practical suggestions for its use. The up-to-date bibliography reflects that this is an important and live subject and betrays the fact that the authors are well-known experts in the field who have contributed some of the best-known software packages and algorithms in the area.

Each chapter contains useful introductory sections on the major algorithmic aspects that are followed by more technical and detailed sections. The description of the topics is rather self-contained, therefore the reader is to forced to rely on many other technical books for the algorithmic issues.

Chapter 1 presents the state of the art in parallel architectures, for single and multiple processors. Not surprisingly, much has happened since the previous edition of this book (when vector architectures were still major players). New sections have been devoted to message passing and network-based environments. Chapter 2 has been substantially reduced in scope, limiting the overview of current high-performance computers to MIMD systems with Shared and Distributed Memory.

Chapters 3 and 4 deal with general implementation aspects, such as synchronization, load balancing, indirect addressing and the Message Passing Interface. The section on performance analysis recalls some classical tools while providing helpful suggestions for carrying out practical performance measurements.

The major theme of this book is numerical linear algebra. Its presentation starts in Chapter 5 with dense computations. The major linear algebra kernels (BLAS) are emphasized, which form the basic blocks for LAPACK, the most mature of packages for dense linear algebra. The principal algorithms for solving linear systems and least squares problems are recalled. This chapter also describes ScaLAPACK, whose aim is to provide a linear algebra software standard for distributed memory MIMD computers.

Chapter 6 focuses on direct methods for sparse matrices, for which algorithmic implementations on parallel machines require a mostly ad hoc design. The reader who is not familiar with the area is first introduced to the essential topic of sparse data structures. Then, she/he can follow more closely the technical discussion on sparse orderings, cliques and indirect addressing. Frontal and multifrontal methods make the "sparse" codes amenable to vector and parallel environments: performance results on modern machines give a feeling of the potential of these powerful methods. Useful pointers to public domain and commercial packages are given that will help the user who wants to build applications that employ these packages.

The next three chapters deal with iterative methods of the Krylov subspace variety. This constitutes a major expansion from the previous edition; it is well deserved given that this topic was at the center of one of the most exciting activities in the field of numerical linear algebra and has generated some of the most successful iterative solvers.

The first of these chapters introduces the basic principles and ideas behind Krylov subspaces while Chapter 8 provides a detailed description of several methods that have been proposed in recent years. The basic parallel kernels in iterative methods are mostly BLAS1 operations, together with the step involving the coefficient matrix, usually through a matrix-vector multiplication; most of the analysis of this second phase is postponed to Chapter 10. Chapter 8 also contains useful descriptions of different implementations of the highly successful Conjugate Gradients algorithm. It would have been nice if the authors had also included a discussion on (block) methods for linear systems with multiple right-hand sides, since they are useful in many applications and also make natural use of BLAS2 and BLAS3 computations. The chapter terminates with a discussion on testing iterative methods.

Preconditioning is an important step in iterative approaches; this and parallel implementations are thoroughly presented in Chapter 9. Besides the classical incomplete schemes, a few pages are devoted to the presentation of recent approaches, such as Sparse Approximate Inverse, and Element-by-Element preconditioning, which typically exhibit their best performance in a parallel context. It would have been nice if the authors had opted for a more detailed presentation of domain decomposition methods.

Chapters 10 and 11 describe methods for the standard and generalized eigenvalue problems. After a survey of the most widely used approaches, the very successful package ARPACK is described together with its parallel implementation P_ARPACK (written using MPI). Several important issues are discussed at length, providing the reader with some implementation hints and with a good feeling of the expected performance. Finally, the Appendices gives the necessary information to practically deal with some of the described codes.

In conclusion, in spite of what the authors say in the Preface (" ... Any book that attempts to cover these topics must necessarily be somewhat out of date before it appears"), the book contains a lot of up-to-date material, and I recommend it to computational scientists who deal with linear algebra methods on any of parallel, vector and sequential (!) computer environments. Readers who already own the previous edition will find that this book has been significantly expanded to include recent important advances in numerical linear algebra tools and HPC environments that will make their "HP computational life" much easier.

VALERIA SIMONCINI
INST. DI ANALISI NUM. DEL CNR
CASTEL SAN PIETRO BOLOGNA
ITALY

5[65F15, 65F10]—*ARPACK Users' Guide, Solution of Large-Scale Eigenvalue Problems with Implicitly Restarted Arnoldi Methods*, by R. B. Lehoucq, D. C. Sorensen, and C. Yang, SIAM, Philadelphia, PA, 1998, xv+142 pp., $25\frac{1}{2}$ cm, softcover, $39.00

The chief impediment to solving large eigenvalue problems is lack of sufficient memory—a difficulty that has two aspects. In the first place, if the order of the matrix in question is large, the matrix must be represented in some compact form. This limits what we can do with the matrix to simple operations like multiplying it by a vector or, if we are lucky, factoring it so that we have a representation of its inverse. The second aspect is that we cannot hope to store the entire matrix of eigenvectors and must content ourselves with computing a few selected eigenpairs. We are also limited in the number of extra working vectors that we can use to compute these eigenpairs.

Krylov sequence methods are popular in part because they can be made to respect these limitations. The methods proceed by orthogonalizing a Krylov sequence $u, Au, A^2u, \ldots$. When the resulting vectors are arranged in a matrix