20(a)[68N15, 68N20, 65Y05].—PHILIP J. HATCHER & MICHAEL J. QUINN *Data-Parallel Programming on MIMD Computers*, The MIT Press, Cambridge, MA, 1991, xiv + 231 pp., $23\frac{1}{2}$ cm. Price $31.50.

20(b)[65–06, 65Y05, 68N99].—PIYUSH MEHROTRA, JOEL SALTZ & ROBERT VOIGT (Editors), *Unstructured Scientific Computation on Scalable Multiprocessors*, Scientific and Engineering Computation Series, The MIT Press, Cambridge, MA, 1992, xviii + 407 pp., $23\frac{1}{2}$ cm. Price $39.95.

20(c)[65–06, 65P05, 76–06].—HORST D. SIMON (Editor), *Parallel Computational Fluid Dynamics: Implementations and Results*, Scientific and Engineering Computation Series, The MIT Press, Cambridge, MA, 1992, xii + 347 pp., $23\frac{1}{2}$ cm. Price $45.00.

These are volumes one, two, and three of the *Scientific and Engineering Computation Series*, edited by Janusz S. Kowalik. Libraries may not have made their acquisition decisions regarding this series, so we will comment on the series objectives and format to some extent. The second two books are the proceedings of conferences held in 1990. The "unstructured" conference was held at Nags Head in October, 1990, and the "CFD" (computational fluid dynamics) conference was held at Indianapolis in May, 1990.

The subject of all three books is parallel computing. Thus one might conclude that the Series will be devoted to parallel computing. However, the Series Foreword by Kowalik indicates broader objectives: "It will include books on theories, methods, and original applications in such areas as parallelism, large-scale simulations, time-critical computing, computer-aided design and engineering, use of computers in manufacturing, visualization of scientific data, and human-machine interface technology." The first three volumes are devoted to the first two topics on this list (parallelism and large-scale simulations). The last five are also topics of much current interest. If done well, it could be a valuable series.

Volume one of the series is devoted to programming languages for parallel computing, and it focuses in particular on *data-parallel* languages (which it defines rigorously). These languages are naturally associated with SIMD parallel computers, so the book emphasizes their effectiveness as well on MIMD Computers. The book is designed for three different uses. First, it is intended as an introduction to high-level parallel computing languages. Secondly, the book is intended to be used by people developing compilers for either distributed- or shared-memory parallel computers. Finally, it is intended to be a reference book for the programming language Dataparallel C.

The book outlines the various choices that are made in designing a parallel programming language, such as whether parallelism should be explicit or implicit, synchronous or asynchronous, etc. There are five different choices to be made, with more than two options in some cases. Dataparallel C represents one particular choice from the more than seventy possible designs. After describing the language in detail, extensive applications are presented, and the efficiency of the compiler is measured. This book sets a high standard for the presentation and evaluation of a parallel programming language. Given the variety of possible parallel languages, such an exposition is essential for the general technical community to be able to form an opinion on the best type of parallel programming model to use for particular applications. At the moment, sequential application codes are written variously in C, Fortran, Ada, Lisp, and

probably many more sequential languages. It is reasonable to expect that there will be equal or greater variety in the ways that parallelism is expressed, and this book makes a clear case for the data-parallel approach.

The "CFD" proceedings consists of 16 papers, averaging 22 pages in length, and the "unstructured" proceedings consists of 20 papers, averaging 20 pages in length. However, the variation in page length is what one would expect in a journal, indicating that no artificial limits were placed on these papers. This is a refreshing alternative to the practice of limiting proceedings papers to a length appropriate only for an extended abstract, or prescribing any arbitrary (target) length.

Although not an important technical issue, the cover design of this series is quite attractive. However, quality control over the typography is varied. The "CFD" proceedings is quite uniform, with each chapter being individually numbered and the section (and subsection) numbers within each article adhering to this numbering. Moreover, the typography appears to be quite uniform as well in the "CFD" proceedings. The typography of the "unstructured" proceedings is less well structured, with some articles using densely spaced, small fonts and others having quite large print. This variation is unnecessary today, as authors can easily conform to typography requirements with minimal effort.

The two proceedings represent two different types of conferences. The "unstructured" conference was an *ad hoc* workshop, whereas the "CFD" conference was the second in a series that held its *n*th meeting in June, 1993, in Paris. Although the intersection of contributors could have been much larger (over one-third of the papers in the "unstructured" proceedings are CFD-related), the intersection consists of only two authors, both associated with RIACS (the Research Institute for Advanced Computer Science) at NASA Ames Research Center, a major center in these research areas. Although both proceedings are quite general and interdisciplinary, the "CFD" proceedings has a flavor more towards aeronautical engineering, whereas the "unstructured" proceedings tilts more toward computer science.

The "unstructured" proceedings, according to its Preface, addresses "unstructured and dynamically varying computations on scalable multiprocessors." Organized into categories, the problems addressed include:

- methods to effectively map fluids and structural mechanics codes that employ unstructured and/or adaptive meshes,
- scalable algorithms for problems in sparse linear algebra,
- scalable tools and compilers designed to handle irregular scientific computations,
- mapping methods for adaptive fast multipole methods, and
- parallelized grid generation and problem partitioning.

The number of papers devoted to sparse linear algebra, adaptive fast multipole methods and parallelized grid generation is quite small, roughly two or less each. Thus the book is less than definitive in these areas; rather it forms a snapshot of some of the research going on at the time.

The "CFD" book is more than a proceedings; according to its Preface "this book has been nurtured by the conference, but includes a number of additional chapters, dedicated to recent research work of relevance to the emerging field of parallel computational fluid dynamics." Most of the major discretization

techniques are addressed (finite-difference, finite-element, spectral and particle methods). It also includes studies of parallel algorithms of relevance beyond fluid dynamics, such as parallel versions of

- iterative methods for solving discretizations of elliptic boundary value problems,
- integer sorting, and
- discrete Fourier transforms.

It also begins with an overview chapter (by the editor and two colleagues), and ends with a chapter (by the editor and a different colleague) discussing hardware and software issues related to projected teraflops computers and their use for parallel CFD.

The audience for these books is large and interdisciplinary. The publication of the proceedings makes them available to a much broader audience than would normally be the case with a proceedings published by an engineering or science disciplinary organization or just by the conference itself. This series is off to a good start, and it is easy to recommend it for library acquisition. The individual books reviewed here will be valuable to researchers specializing in the particular area covered by each book. People interested in programming unstructured fluids codes will want all three.

L. R. S.

21[65-06, 65P05, 35R35].—P. NEITTAANMÄKI (Editor), *Numerical Methods for Free Boundary Problems*, Internat. Ser. Numer. Math., Vol. 99, Birkhäuser, Basel, 1991, xvi + 439 pp., 24 cm. Price $98.00.

This book is the proceedings of a conference held at the University of Jyväskylä, Finland, in July, 1990. The subject is free-boundary problems for partial differential equations. Although this may sound like a unified theme, it is not, owing to the diversity of types of free-boundary problems. Even in the area of fluid dynamics, there are three distinct research areas represented here, having essentially no relationship between them. No synthesis of these different areas has been attempted in this book.

The decision to publish a proceedings such as this can be justified on various grounds. The proceedings [1] intended to stimulate the combination of numerical and analytical techniques for studying singularly perturbed differential equations. One can argue this is a need that is not being met by regular journals. The justification for the proceedings under review seems less strong and rests primarily in the extra visibility provided to the important subject(s) of free-boundary problems for partial differential equations.

The proceedings consists of 39 papers, averaging 11 pages in length. In fact 29 of them are between 8 and 12 pages (with 10 pages being the mode). It would appear that authors were instructed to provide manuscripts ten pages in length, as is typical in many conferences today. This length leads to rather long extended abstracts (or in some cases advertisements) of work to be found elsewhere. Indeed, any restriction is somewhat unnatural; one finds far more variability in length of paper (and longer papers) in a typical journal, such as this one. This criticism is meant for the genre of conference proceedings as a whole, not just this book in particular. In fact, many of the papers are informative at a survey level, even if important details are often missing.