Books

Highly Parallel Computing, G. S. Almasi and A. Gottlieb, Benjamin-Cummings Publishers, Redwood City, CA, 1989. 600 pp. (ISBN 0-8053-0177-1).

This is a nicely constructed book that covers a broad range of essential topics about parallel computing. There are many particularly attractive features that should make it the preferred book to use as an introduction to the area. It is particularly appropriate for professionals and graduate students who are already familiar with aspects of computing. The authors discuss critical topics in machine organization, parallel application concepts, language constructs, and appropriate systems environments for various topologies of application and machine structure.

This reviewer has been concerned with these topics for some time, and so it is particularly satisfying that the authors (while describing material that is absolutely current) have also made the effort to preserve pioneering efforts that are now quite old, such as ILIAC IV and Cmmm (Carnegie Mellon Multiprocessor), and to recapture the insights and discoveries, frustrations and concerns that have been associated with parallel computing from the beginning. So much recent work in parallel computing presents the issues of the area (tractable and intractable) as though they have been recently discovered. The historical dimension of the book adds to its richness and completeness, and suggests some of the difficulties that have been encountered in achieving high degrees of parallelism. Whereas many of these difficulties are being relaxed because of technology, many remain areas of argument and constraint. And although it is surely possible to construct machines that consist of large populations of microprocessors, it is still not clear what their range of general applicability might be. Nor is it clear what range of applications can be conceptually parallelized and how that parallelism should be discovered and expressed.

The book starts with an excellent overview of the area, introducing fundamental concepts of synchronization and parallel models of computing. It then presents a set of sample applications drawn from scientific applications, engineering applications, and systems applications. It is particularly informative in its discussion of the parallel model in database management, an area of critical interest because it seems to be the "special" application of most potential immediate impact on systems structure. Perhaps a mild criticism of the approach to applications might be made here, because no consideration is given to the possibility that there may be high degrees of parallel potential on some standard commercial applications and that an investigation of these applications might reveal a vastly enlarged opportunity for parallelism.

The book presents a nice discussion of the technology basis for distributed processing. The reader comes away with a good understanding of the logic, memory, and I/O technologies and the role these play in constraining or enabling parallel processing. This discussion is followed by a consideration of the classic computational models of various forms of parallelism.

There is well considered and complete material on language approaches and language constructs with elegant examples of the problems addressed by the language features. There is an equally excellent review of the methods used by language translators to discover parallelism implicit in a program. One of the argued issues of the future of parallel computing has (always) been the degree to which programmers must specify parallel computation as opposed to the degree by which compiling instruments can decompose programs to discover parallel opportunity without explicit programmer statements. Examples, al-

though largely based on vectorization, provide an insight into what is involved in discovering parallel opportunity. The book also discusses some of the operating systems issues and presents some classical models for operating systems appropriate for parallel environments.

Finally, parallel architectures are discussed with a nicely selected set of example machines. The achievement of the authors here is that the systems they discuss cover a very complete range of conceptual models of parallel computation. The section on parallel architectures, therefore, serves not only to familiarize the reader with systems that have been prototyped or marketed but to reinforce the conceptual principles underlying the work as well.

We have been waiting for highly parallel computing for some time now, arguing about its potential efficiency and potential costs, about its proper conceptual models, and about its source of applications. It now seems that the technology makes this area seem immediate and critical enough that more aggressive efforts will be made to bring families of highly parallel systems to commercial reality. This book is a timely and considered statement of the state of the art, its history and its possible future. For those of us who have encountered aspects of the topic over time, it will enrich our understanding of "how it turned out" and "what they still don't know." For those of us who are becoming familiar with the area, it is a splendid synthesis by authors clearly at ease in both hardware and software issues.

Harold R. Lorin
IBM Thornwood Education Center
Thornwood
New York