

<<高性能计算机上的数值线性代数>>

图书基本信息

书名：<<高性能计算机上的数值线性代数>>

13位ISBN编号：9787302244998

10位ISBN编号：7302244995

出版时间：2011-2

出版时间：清华大学出版社

作者：（美）冬格拉 等著

页数：342

版权说明：本站所提供下载的PDF图书仅提供预览和简介，请支持正版图书。

更多资源请访问：<http://www.tushu007.com>

<<高性能计算机上的数值线性代数>>

内容概要

the purpose of this book is to unify and document in one place many of the techniques and much of the current understanding about solving systems of linear equations on vector and parallel computers. this book is not a textbook, but it is meant to provide a fast entrance to the world of vector and parallel processing for these linear algebra applications. we intend this book to be used by three groups of readers: graduate students, researchers working in computational science, and numerical analysts. as such, we hope this book can serve both as a reference and as a supplement to a teaching text on aspects of scientific computation.

the book is divided into five major parts: (1) introduction to terms and concepts, including an overview of the state of the art for high-performance computers and a discussion of performance evaluation (chapters 1-4); (2) direct solution of dense matrix problems (chapter 5); (3) direct solution of sparse systems of equations (chapter 6); (4) iterative solution of sparse systems of equations (chapters 7-9); and (5) iterative solution of sparse eigenvalue problems (chapters 10-11). any book that attempts to cover these topics must necessarily be somewhat out of date before it appears, because the area is in a state of flux. we have purposely avoided highly detailed descriptions of popular machines and have tried instead to focus on concepts as much as possible; nevertheless, to make the description more concrete, we do point to specific computers.

<<高性能计算机上的数值线性代数>>

作者简介

作者：（美国）冬格拉（Jack j.Dongarra）（美国）Iain S.Duff（美国）Danny C.Sorensen 等
Jack J. Dongarra is a Distinguished Professor of Computer Science at the University of Tennessee and a Distinguished Scientist at Oak Ridge National Laboratory. Iain S. Duff is Group Leader of Numerical Analysis at the CCLRC Rutherford Appleton Laboratory, the Project Leader for the Parallel Algorithms Group at CERFACS in Toulouse, and a Visiting Professor of Mathematics at the University of Strathclyde. Danny C. Sorensen is a Professor of Computational and Applied Mathematics at Rice University. Henk A. van der Vorst is a Professor in Numerical Analysis at Utrecht University in the Netherlands.

<<高性能计算机上的数值线性代数>>

书籍目录

- about the authors
- preface
- introduction
- 1 high-performance computing
 - 1.1 trends in computer design
 - 1.2 traditional computers and their limitations
 - 1.3 parallelism within a single processor
 - 1.3.1 multiple functional units
 - 1.3.2 pipelining
 - 1.3.3 overlapping
 - 1.3.4 risc
 - 1.3.5 vliw
 - 1.3.6 vector instructions
 - 1.3.7 chaining
 - 1.3.8 memory-to-memory and register-to-register
 - organizations
 - 1.3.9 register set
 - 1.3.10 stripmining
 - 1.3.11 reconfigurable vector registers
 - 1.3.12 memory organization
 - 1.4 data organization
 - 1.4.1 main memory
 - 1.4.2 cache
 - 1.4.3 local memory
 - 1.5 memory management
 - 1.6 parallelism through multiple pipes or multiple
- processors
 - 1.7 message passing
 - 1.8 virtual shared memory
 - 1.8.1 routing
 - 1.9 interconnection topology
 - 1.9.1 crossbar switch
 - 1.9.2 timeshared bus
 - 1.9.3 ring connection
 - 1.9.4 mesh connection
 - 1.9.5 hypercube
 - 1.9.6 multi-staged network
 - 1.10 programming techniques
 - 1.11 trends: network-based computing
- 2 overview of current high-performance computers
 - 2.1 supercomputers
 - 2.2 risc-based processors
 - 2.3 parallel processors
 -
- 3 implementation details and overhead

<<高性能计算机上的数值线性代数>>

4 performance: analysis, modeling, and measurements

5 building blocks in linear algebra

6 direct solution of sparse linear systems

7 krylov subspaces: projection

8 iterative methods for linear systems

9 preconditioning and parallel preconditioning

10 linear eigenvalue problems $ax =$

11 the generalized eigenproblem

bibliography

index

章节摘录

版权页：插图：Scalar registers are a form of very-high-speed memory used to hold the most heavily referenced data at any point in the program's execution. These registers can send data to the functional units in one clock cycle; this process is typically an order of magnitude faster than the memory speed. As one might expect, memory with this speed is expensive; and as a result, the amount of high-speed memory supplied as registers is limited. In general, the speed of main memory is insufficient to meet the requirements of the functional units. To ensure that data can be delivered fast enough to the functional units and that the functional units can get rid of their output fast enough, most manufacturers have introduced a special form of very-high-speed memory, the so-called vector register. A vector register is a fixed set of memory locations in a special memory under user control. For example, on Cray computers, a vector register consists of 64 or 128 elements. The rule is for a functional unit to accept complete vector registers as operands. The hardware makes it possible for the elements of the vector register to be fed one by one to the functional unit at a rate of exactly one per clock cycle per register. Also, the register accepts 64 successive output elements of the functional unit at a rate of one element per clock cycle. The locations within a vector register cannot be accessed individually. Loading, storing, and manipulating the contents of a vector register are done under control of special vector instructions. These vector instructions are issued automatically by the compiler, when applicable. Since a functional unit typically involves three vector operands, two for input and one for output, there are more than three vector registers. For example, the Cray processors typically have 8 or 16 vector registers. The seemingly extra registers are used by the system in order to keep intermediate results as operands for further instructions. The programmer can help the compiler exploit the contents of these registers by making it possible to combine suitable statements and expressions.

<<高性能计算机上的数值线性代数>>

媒体关注与评论

"This excellent book successfully achieves the authors' purpose to unify and document in one place many of the techniques and much of the current understanding about solving systems of linear equations [and eigensystems] on vector and parallel computers. We highly recommend it as a very useful reference for both graduate students and practitioners..." — F. A. Smith and R. E. Funderlic, *Statistical software Newsletter* (part of the *Computational Statistics and Data Analysis* journal). "...Numerical Linear Algebra for High-Performance Computers is a major revision to the book entitled *Solving Linear Systems on Vector and Shared Memory Computers*, published by SIAM in 1990. But the current book updates the material focusing on vector and parallel computing for linear algebra and presents new contents on the eigenvalue problem. In short, this reviewer wants to thank the authors for writing such a good book." — Pan Zeng, *Applied Mechanics Reviews*, Vol. 52, No. 7, July 1999. "The present revised and extended volume brings the computing aspects up to date, contains a much more detailed treatment of preconditioning for linear systems, and extends the numerical treatment to the algebraic eigenvalue problem..., many books have been published in numerical linear algebra... However, in my view, the need for the present book has increased rather than diminished..." — Nicholas J. Higham, *University of Manchester, SIAM Review*, Vol. 42, No. 3, September 2000. "This is a volume that every numerical analyst should have, either for updating his lectures or for his own research. I also encourage students, engineers, and practitioners to read it. It is obviously a very valuable addition to the existing literature." — Claude Brezinski, *Numerical Algorithms* 22 (1999).

<<高性能计算机上的数值线性代数>>

编辑推荐

《高性能计算机上的数值线性代数(影印版)》：国际著名数学图书

<<高性能计算机上的数值线性代数>>

版权说明

本站所提供下载的PDF图书仅提供预览和简介，请支持正版图书。

更多资源请访问:<http://www.tushu007.com>