

An Introduction to Parallel and Vector Scientific Computing

In this text, students of applied mathematics, science, and engineering are introduced to fundamental ways of thinking about the broad context of parallelism. The authors begin by giving the reader a deeper understanding of the issues through a general examination of timing, data dependencies, and communication. These ideas are implemented with respect to shared memory, parallel and vector processing, and distributed memory cluster computing. Threads, OpenMP, and MPI are covered, along with code examples in Fortran, C, and Java.

The principles of parallel computation are applied throughout as the authors cover traditional topics in a first course in scientific computing. Building on the fundamentals of floating point representation and numerical error, a thorough treatment of numerical linear algebra and eigenvector/eigenvalue problems is provided. By studying how these algorithms parallelize, the reader is able to explore parallelism inherent in other computations, such as Monte Carlo methods.

Ronald W. Shonkwiler is a professor in the School of Mathematics at the Georgia Institute of Technology. He has authored or coauthored more than 50 research papers in areas of functional analysis, mathematical biology, image processing algorithms, fractal geometry, neural networks, and Monte Carlo optimization methods. His algorithm for monochrome image comparison is part of a U.S. patent for fractal image compression. He has coauthored two other books, *An Introduction to the Mathematics of Biology* and *The Handbook of Stochastic Analysis and Applications*.

Lew Lefton is the Director of Information Technology for the School of Mathematics and the College of Sciences at the Georgia Institute of Technology, where he has built and maintained several computing clusters that are used to implement parallel computations. Prior to that, he was a tenured faculty member in the Department of Mathematics at the University of New Orleans. He has authored or coauthored more than a dozen research papers in the areas of nonlinear differential equations and numerical analysis. His academic interests are in differential equations, applied mathematics, numerical analysis (in particular, finite element methods), and scientific computing.

Cambridge Texts in Applied Mathematics

All titles listed below can be obtained from good booksellers or from Cambridge University Press.
 For a complete series listing, visit <http://publishing.cambridge.org/stm/mathematics/ctam>

- Thinking about Ordinary Differential Equations*
 ROBERT E. O'MALLEY JR
- A Modern Introduction to the Mathematical Theory of Water Waves*
 ROBIN STANLEY JOHNSON
- Rarefied Gas Dynamics: From Basic Concepts to Actual Calculations*
 CARLO CERCIGNANI
- Symmetry Methods for Differential Equations: A Beginner's Guide*
 PETER E. HYDON
- High Speed Flow*
 C. J. CHAPMAN
- Wave Motion*
 J. BILLINGHAM, AND A. C. KING
- An Introduction to Magnetohydrodynamics*
 P. A. DAVIDSON
- Linear Elastic Waves*
 JOHN G. HARRIS
- Vorticity and Incompressible Flow*
 ANDREW J. MAJDA, AND ANDREA L. BERTOZZI
- Infinite-Dimensional Dynamical Systems*
 JAMES C. ROBINSON
- Introduction to Symmetry Analysis*
 BRAIN J. CANTWELL
- Bäcklund and Darboux Transformations*
 C. ROGERS, AND W. K. SCHIEF
- Finite Volume Methods for Hyperbolic Problems*
 RANDALL J. LEVEQUE
- Introduction to Hydrodynamic Stability*
 P. G. DRAZIN
- Theory of Vortex Sound*
 M. S. HOWE
- Scaling*
 GRIGORY ISAAKOVICH BARENBLATT
- Complex Variables: Introduction and Applications (2nd Edition)*
 MARK J. ABLOWITZ, AND ATHANASSIOS S. FOKAS
- A First Course in Combinatorial Optimization*
 JONE LEE
- Practical Applied Mathematics: Modelling, Analysis, Approximation*
 SAM HOWISON

Cambridge University Press
978-0-521-68337-1 — An Introduction to Parallel and Vector Scientific Computation
Ronald W. Shonkwiler , Lew Lefton
Frontmatter
[More Information](#)

An Introduction to Parallel and Vector Scientific Computing

RONALD W. SHONKWILER
Georgia Institute of Technology

LEW LEFTON
Georgia Institute of Technology



CAMBRIDGE
UNIVERSITY PRESS

Cambridge University Press
 978-0-521-68337-1 — An Introduction to Parallel and Vector Scientific Computation
 Ronald W. Shonkwiler, Lew Lefton
 Frontmatter
[More Information](#)

CAMBRIDGE UNIVERSITY PRESS

University Printing House, Cambridge CB2 8BS, United Kingdom
 One Liberty Plaza, 20th Floor, New York, NY 10006, USA
 477 Williamstown Road, Port Melbourne, VIC 3207, Australia
 314-321, 3rd Floor, Plot 3, Splendor Forum, Jasola District Centre, New Delhi - 110025, India
 103 Penang Road, #05-06/07, Visioncrest Commercial, Singapore 238467

Cambridge University Press is part of the University of Cambridge.
 It furthers the University's mission by disseminating knowledge in the pursuit of
 education, learning and research at the highest international levels of excellence.

www.cambridge.org
 Information on this title: www.cambridge.org/9780521683371

© Ronald Shonkwiler and Lew Lefton 2006

This publication is in copyright. Subject to statutory exception
 and to the provisions of relevant collective licensing agreements,
 no reproduction of any part may take place without the written
 permission of Cambridge University Press.

First published 2006

A catalogue record for this publication is available from the British Library

Library of Congress Cataloging in Publication data

Shonkwiler, Ronald W., 1942–

An introduction to parallel and vector scientific computing / by R. Shonkwiler and L. Lefton.
 p. cm. – (Cambridge texts in applied mathematics ; no. 41)

Includes bibliographical references and index.

ISBN-13: 978-0-521-86478-7

ISBN-10: 0-521-86478-X

ISBN-13: 978-0-521-68337-1 (pbk.)

ISBN-10: 0-521-68337-8 (pbk.)

1. Parallel processing (Electronic computers) 2. Vector processing (Computer science)

I. Lefton, L. (Lew), 1960– II. Title.

III. Series: Cambridge texts in applied mathematics ; 41.

QA76.58.S545 2006

004.3'5–dc22 2006007798

ISBN 978-0-521-86478-7 Hardback

ISBN 978-0-521-68337-1 Paperback

Cambridge University Press has no responsibility for the persistence or accuracy
 of URLs for external or third-party internet websites referred to in this publication,
 and does not guarantee that any content on such websites is, or will remain,
 accurate or appropriate. Information regarding prices, travel timetables, and other
 factual information given in this work is correct at the time of first printing but
 Cambridge University Press does not guarantee the accuracy of such information
 thereafter.

Contents

Preface	<i>page xi</i>
PART I MACHINES AND COMPUTATION	1
1 Introduction – The Nature of High-Performance Computation	3
1.1 Computing Hardware Has Undergone Vast Improvement	4
1.2 SIMD–Vector Computers	9
1.3 MIMD – True, Coarse-Grain Parallel Computers	14
1.4 Classification of Distributed Memory Computers	16
1.5 Amdahl’s Law and Profiling	20
Exercises	23
2 Theoretical Considerations – Complexity	27
2.1 Directed Acyclic Graph Representation	27
2.2 Some Basic Complexity Calculations	33
2.3 Data Dependencies	37
2.4 Programming Aids	41
Exercises	41
3 Machine Implementations	44
3.1 Early Underpinnings of Parallelization – Multiple Processes	45
3.2 Threads Programming	50
3.3 Java Threads	56
3.4 SGI Threads and OpenMP	63
3.5 MPI	67

viii	<i>Contents</i>	
3.6	Vector Parallelization on the Cray	80
3.7	Quantum Computation	89
	Exercises	97
PART II	LINEAR SYSTEMS	101
4	Building Blocks – Floating Point Numbers and Basic Linear Algebra	103
4.1	Floating Point Numbers and Numerical Error	104
4.2	Round-off Error Propagation	110
4.3	Basic Matrix Arithmetic	113
4.4	Operations with Banded Matrices	120
	Exercises	124
5	Direct Methods for Linear Systems and LU Decomposition	126
5.1	Triangular Systems	126
5.2	Gaussian Elimination	134
5.3	<i>ijk</i> -Forms for LU Decomposition	150
5.4	Bordering Algorithm for LU Decomposition	155
5.5	Algorithm for Matrix Inversion in $\log^2 n$ Time	156
	Exercises	158
6	Direct Methods for Systems with Special Structure	162
6.1	Tridiagonal Systems – Thompson’s Algorithm	162
6.2	Tridiagonal Systems – Odd–Even Reduction	163
6.3	Symmetric Systems – Cholesky Decomposition	166
	Exercises	170
7	Error Analysis and QR Decomposition	172
7.1	Error and Residual – Matrix Norms	172
7.2	Givens Rotations	180
	Exercises	184
8	Iterative Methods for Linear Systems	186
8.1	Jacobi Iteration or the Method of Simultaneous Displacements	186
8.2	Gauss–Seidel Iteration or the Method of Successive Displacements	189
8.3	Fixed-Point Iteration	191

<i>Contents</i>	ix
8.4 Relaxation Methods	193
8.5 Application to Poisson’s Equation	194
8.6 Parallelizing Gauss–Seidel Iteration	198
8.7 Conjugate Gradient Method	200
Exercises	204
9 Finding Eigenvalues and Eigenvectors	206
9.1 Eigenvalues and Eigenvectors	206
9.2 The Power Method	209
9.3 Jordan Canonical Form	210
9.4 Extensions of the Power Method	215
9.5 Parallelization of the Power Method	217
9.6 The QR Method for Eigenvalues	217
9.7 Householder Transformations	221
9.8 Hessenberg Form	226
Exercises	227
PART III MONTE CARLO METHODS	231
10 Monte Carlo Simulation	233
10.1 Quadrature (Numerical Integration)	233
Exercises	242
11 Monte Carlo Optimization	244
11.1 Monte Carlo Methods for Optimization	244
11.2 IIP Parallel Search	249
11.3 Simulated Annealing	251
11.4 Genetic Algorithms	255
11.5 Iterated Improvement Plus Random Restart	258
Exercises	262
APPENDIX: PROGRAMMING EXAMPLES	265
MPI Examples	267
Send a Message Sequentially to All Processors (C)	267
Send and Receive Messages (Fortran)	268
Fork Example	270
Polynomial Evaluation(C)	270

x	<i>Contents</i>	
Lan Example		275
Distributed Execution on a LAN(C)		275
Threads Example		280
Dynamic Scheduling(C)		280
SGI Example		282
Backsub.f(Fortran)		282
References		285
Index		286

Preface

Numerical computations are a fundamental tool for engineers and scientists. The current practice of science and engineering demands that nontrivial computations be performed with both great speed and great accuracy. More and more, one finds that scientific insight and technological breakthroughs are preceded by intense computational efforts such as modeling and simulation. It is clear that computing is, and will continue to be, central to the further development of science and technology.

As market forces and technological breakthroughs lowered the cost of computational power by several orders of magnitude, there was a natural migration from large-scale mainframes to powerful desktop workstations. Vector processing and parallelism became possible, and this parallelism gave rise to a new collection of algorithms. Parallel architectures matured, in part driven by the demand created by the algorithms. Large computational codes were modified to take advantage of these parallel supercomputers. Of course, the term *supercomputer* has referred, at various times, to radically different parallel architectures. This includes vector processors, various shared memory architectures, distributed memory clusters, and even computational grids. Although the landscape of scientific computing changes frequently, there is one constant; namely, that there will always be a demand in the research community for high-performance computing.

When computations are first introduced in beginning courses, they are often straightforward “vanilla” computations, which are well understood and easily done using standard techniques and/or commercial software packages on desktop computers. However, sooner or later, a working scientist or engineer will be faced with a problem that requires advanced techniques, more specialized software (perhaps coded from scratch), and/or more powerful hardware. This book is aimed at those individuals who are taking that step, from a novice to intermediate or even from intermediate to advanced user of tools that fall under

the broad heading of scientific computation. The text and exercises have been shown, over many years of classroom testing, to provide students with a solid foundation, which can be used to work on modern scientific computing problems. This book can be used as a guide for training the next generation of computational scientists.

This manuscript grew from a collection of lecture notes and exercises for a senior-level course entitled “Vector and Parallel Scientific Computing.” This course runs yearly at the Georgia Institute of Technology, and it is listed in both mathematics and computer science curricula. The students are a mix of math majors, computer scientists, all kinds of engineers (aerospace, mechanical, electrical, etc.), and all kinds of scientists (chemists, physicists, computational biologists, etc.). The students who used these notes came from widely varying backgrounds and varying levels of expertise with regard to mathematics and computer science.

Formally, the prerequisite for using this text is knowledge of basic linear algebra. We integrate many advanced matrix and linear algebra concepts into the text as the topics arise rather than offering them as a separate chapter. The material in Part II, Monte Carlo Methods, also assumes some familiarity with basic probability and statistics (e.g., mean, variance, t test, Markov chains).

The students should have some experience with computer programming. We do not teach nor emphasize a specific programming language. Instead, we illustrate algorithms through a pseudocode, which is very close to mathematics itself. For example, the mathematical expression $y = \sum_{i=1}^n x_i$ becomes

```
y=0;
loop i = 1 upto n
  y = y + xi;
end loop
```

We provide many example programs in Fortran, C, and Java. We also have examples of code that uses MPI libraries. When this course was originally taught, it took several weeks for the students to get accounts and access to the Cray system available at that time. As a result, the material in the first two chapters provides no programming exercises. If one wishes to start programming right away, then he or she should begin with Chapter 3.

The purpose of the course is to provide an introduction to important topics of scientific computing including the central algorithms for numerical linear algebra such as linear system solving and eigenvalue calculation. Moreover, we introduce this material from the very beginning in the context of vector and parallel computation. We emphasize a recognition of the sources and propagation of numerical error and techniques for its control. Numerical error starts with

the limitations inherent in the floating point representation of numbers leading to round-off error and continues with algorithmic sources of error.

The material has evolved over time along with the machines called supercomputers. At present, shared memory parallel computation has standardized on the threads model, and vector computation has moved from the machine level to the chip level. Of course, vendors provide parallelizing compilers that primarily automatically parallelize loops that the programmer has requested, sometimes referred to as the DOACROSS model. This is a convenient model for engineers and scientists as it allows them to take advantage of parallel and vector machines while making minimal demands on their programming time. For the purpose of familiarity, we include a section on the basic concepts of distributed memory computation, including topological connectivity and communication issues.

In teaching the course, we employ a hands-on approach, requiring the students to write and execute programs on a regular basis. Over the years, our students have had time on a wide variety of supercomputers, first at National Centers, and more recently at campus centers or even on departmental machines. Of course, even personal computers today can be multiprocessor with a vector processing chipset, and many compiled codes implement threads at the operating system level.

We base our approach to parallel computation on its representation by means of a directed acyclic graph. This cuts to the essence of the computation and clearly shows its parallel structure. From the graph it is easy to explain and calculate the complexity, speedup, efficiency, communication requirements, and scheduling of the computation. And, of course, the graph shows how the computation can be coded in parallel.

The text begins with an introduction and some basic terminology in Chapter 1. Chapter 2 gives a high-level view of the theoretical underpinnings of parallelism. Here we discuss data dependencies and complexity, using directed acyclic graphs to more carefully demonstrate a general way of thinking about parallelism. In Chapter 3, we have included a variety of machine implementations of parallelism. Although some of these architectures are not in widespread use any more (e.g., vector processors like the early Cray computers), there are still interesting and important ideas here. In fact, the Japanese Earth Simulator (the former world record holder for “fastest computer”) makes heavy use of vector processing and pipelining. Chapter 3 includes an introduction to low-level implementations of parallelism by including material on barriers, mutexes, and threads. Of course, not every scientific computing application will require thread programming, but as mentioned earlier, these objects provide many useful ideas about parallelization that can be generalized to many different parallel codes.

We have even included a short introduction to quantum computing because this technology may one day be the future of parallel scientific computation.

In the second half of the book, we start with basic mathematical and computational background, presented as building blocks in Chapter 4. This includes material on floating point numbers, round-off error, and basic matrix arithmetic. We proceed to cover mathematical algorithms, which we have found are most frequently used in scientific computing. Naturally, this includes a large measure of numerical linear algebra. Chapters 5, 6, and 7 discuss direct methods for solving linear systems. We begin with classical Gaussian elimination and then move on to matrices with special structure and more advanced topics such as Cholesky decomposition and Givens' rotation. Iterative methods are covered in Chapter 8. We study Jacobi and Gauss-Seidel as well as relaxation techniques. This chapter also includes a section on conjugate gradient methods. In Chapter 9, we examine eigenvalues and eigenvectors. This includes the power method and QR decomposition. We also cover the topics of Householder transformations and Hessenberg forms, since these can improve QR computations in practice.

Throughout all of Part II, our development of linear algebraic results relies heavily on the technique of partitioning matrices. This is introduced in Chapter 4 and continues through our presentation of Jordan form in Chapter 9.

The final section of the book is focused on Monte Carlo methods. We first develop classical quadrature techniques such as the Buffon Needle Problem in Chapter 10. We then advance in Chapter 11 to a presentation of Monte Carlo optimization, which touches on the ideas of simulated annealing, genetic algorithms, and iterated improvement with random restart.

Exercises are included at the end of every section. Some of these are meant to be done by hand, and some will require access to a computing environment that supports the necessary parallel architecture. This could be a vector machine, an SMP system supporting POSIX threads, a distributed memory cluster with MPI libraries and compilers, etc. We have attempted to isolate those exercises that require programming in a subsection of each exercise set. Exercises are followed by a number in parentheses, which is meant to be an indication of the level of difficulty.

Because scientific computing is often the result of significant research efforts by large distributed teams, it can be difficult to isolate meaningful self-contained exercises for a textbook such as this. We have found it very useful for students to work on and present a project as a substantial part of their course grade. A 10-minute oral presentation along with a written report (and/or a poster) is an excellent exercise for students at this level. One can ask them to submit a short project proposal in which they briefly describe the problem background, the

mathematical problem that requires computation, and how this computation may parallelize. Students do well when given the opportunity to perform a deeper study of a problem of interest to them.

Acknowledgments

Ron Shonkwiler would like to thank his coauthor, Dr. Lew Lefton, for bringing thoughtfulness and a fresh outlook to this project and for his “way with words,” and Lauren Cowles of Cambridge for her encouragement. Most of all, he would like to thank the many students who have taken the course on which this book is based. Their enthusiasm, creativity, energy, and industry have made teaching the course a challenge and a pleasure.

Lew Lefton would first and foremost like to express his deepest gratitude to his coauthor, Dr. Ron Shonkwiler. When they began collaborating on the project of getting the course notes into a book manuscript several years ago, Lefton naively thought it would be an interesting but straightforward project. He was half right; it has been very interesting. The notes have changed (improved!) significantly since that early draft and, having taught the course himself and gone through the process of publishing a book, he feels very fortunate to have had Ron, with his expertise and experience, as a guide. Lefton would also like to thank his wife, Dr. Enid Steinbart, and his daughters, Hannah, Monica, and Natalie, who put up with him spending more than his usual amount of “screen time” in front of the computer. He is truly grateful for their love and support.