

经 典 原 版 书 库

并行程序设计导论

(美) Peter S. Pacheco 著
旧金山大学

(英文版)

An Introduction to
**PARALLEL
PROGRAMMING**

Peter S. Pacheco

经


典

库

并行程序设计导论

(英文版)

*An Introduction to
Parallel Programming*



(美) Peter S. Pacheco 著
旧金山大学



机械工业出版社
China Machine Press

Peter S. Pacheco: An Introduction to Parallel Programming (ISBN 978-0-12-374260-5).

Original English language edition copyright © 2011 by Elsevier Inc. All rights reserved.

Authorized English language reprint edition published by the Proprietor.

Copyright © 2011 by Elsevier (Singapore) Pte Ltd.

Printed in China by China Machine Press under special arrangement with Elsevier (Singapore) Pte Ltd. This edition is authorized for sale in China only, excluding Hong Kong SAR and Taiwan.

Unauthorized export of this edition is a violation of the Copyright Act. Violation of this Law is subject to Civil and Criminal Penalties.

本书英文影印版由Elsevier (Singapore) Pte Ltd. 授权机械工业出版社在中国大陆境内独家发行。本版仅限在中国境内（不包括香港特别行政区及台湾地区）出版及标价销售。未经许可之出口，视为违反著作权法，将受法律之制裁。

封底无防伪标均为盗版

版权所有，侵权必究

本书法律顾问 北京市展达律师事务所

本书版权登记号：图字：01-2011-4800

图书在版编目（CIP）数据

并行程序设计导论（英文版）/（美）帕切克（Pacheco, P. S.）著. —北京：机械工业出版社，2011.9

（经典原版书库）

书名原文：An Introduction to Parallel Programming

ISBN 978-7-111-35828-2

I. 并… II. 帕… III. 并行程序—程序设计—英文 IV. TP311.11

中国版本图书馆CIP数据核字（2011）第181804号

机械工业出版社（北京市西城区百万庄大街22号 邮政编码 100037）

责任编辑：迟振春

北京京师印务有限公司印刷

2011年10月第1版第1次印刷

170mm × 242mm • 24.25印张

标准书号：ISBN 978-7-111-35828-2

定价：65.00元

凡购本书，如有缺页、倒页、脱页，由本社发行部调换

客服热线：(010) 88378991；88361066

购书热线：(010) 68326294；88379649；68995259

投稿热线：(010) 88379604

读者信箱：hzjsj@hzbook.com

出版者的话

机械工业出版社

文艺复兴以降，源远流长的科学精神和逐步形成的学术规范，使西方国家在自然科学的各个领域取得了垄断性的优势；也正是这样的传统，使美国在信息技术发展的六十多年间名家辈出、独领风骚。在商业化的进程中，美国的产业界与教育界越来越紧密地结合，计算机学科中的许多泰山北斗同时身处科研和教学的最前线，由此而产生的经典科学著作，不仅擘划了研究的范畴，还揭示了学术的源变，既遵循学术规范，又自有学者个性，其价值并不会因年月的流逝而减退。

近年，在全球信息化大潮的推动下，我国的计算机产业发展迅猛，对专业人才的需求日益迫切。这对计算机教育界和出版界都既是机遇，也是挑战；而专业教材的建设在教育战略上显得举足轻重。在我国信息技术发展时间较短的现状下，美国等发达国家在其计算机科学发展的几十年间积淀和发展的经典教材仍有许多值得借鉴之处。因此，引进一批国外优秀计算机教材将对我国计算机教育事业的发展起到积极的推动作用，也是与世界接轨、建设真正的世界一流大学的必由之路。

机械工业出版社华章公司较早意识到“出版要为教育服务”。自1998年开始，我们就将工作重点放在了遴选、移译国外优秀教材上。经过多年的不懈努力，我们与Pearson, McGraw-Hill, Elsevier, MIT, John Wiley & Sons, Cengage等世界著名出版公司建立了良好的合作关系，从他们现有的数百种教材中甄选出Andrew S. Tanenbaum, Bjarne Stroustrup, Brian W. Kernighan, Dennis Ritchie, Jim Gray, Alfred V. Aho, John E. Hopcroft, Jeffrey D. Ullman, Abraham Silberschatz, William Stallings, Donald E. Knuth, John L. Hennessy, Larry L. Peterson等大师名家的一批经典作品，以“计算机科学丛书”为总称出版，供读者学习、研究及珍藏。大理石纹理的封面，也正体现了这套丛书的品位和格调。

“计算机科学丛书”的出版工作得到了国内外学者的鼎力襄助，国内的专家不仅提供了中肯的选题指导，还不辞劳苦地担任了翻译和审校的工作；而原书的作者也相当关注其作品在中国的传播，有的还专程为其书的中译本作序。迄今，“计算机科学丛书”已经出版了近两年百个品种，这些书籍在读者中树立了良好的口碑，并被许多高校采用为正式教材和参考书籍。其影印版“经典原版书库”作为姊妹篇也被越来越多实施双语教学的学校所采用。

权威的作者、经典的教材、一流的译者、严格的审校、精细的编辑，这些因素使我们的图书有了质量的保证。随着计算机科学与技术专业学科建设的不断完善和教材改革的逐渐深化，教育界对国外计算机教材的需求和应用都将步入一个新的阶段，我们的目标是尽善尽美，而反馈的意见正是我们达到这一终极目标的重要帮助。华章公司欢迎老师和读者对我们的工作提出建议或给予指正，我们的联系方式如下：

华章网站：www.hzbook.com

电子邮件：hzjsj@hzbook.com

联系电话：(010) 88379604

联系地址：北京市西城区百万庄南街1号

邮政编码：100037



华章科技图书出版中心

In Praise of *An Introduction to Parallel Programming*

With the coming of multicore processors and the cloud, parallel computing is most certainly not a niche area off in a corner of the computing world. Parallelism has become central to the efficient use of resources, and this new textbook by Peter Pacheco will go a long way toward introducing students early in their academic careers to both the art and practice of parallel computing.

Duncan Buell

Department of Computer Science and Engineering
University of South Carolina

An Introduction to Parallel Programming illustrates fundamental programming principles in the increasingly important area of shared-memory programming using Pthreads and OpenMP and distributed-memory programming using MPI. More important, it emphasizes good programming practices by indicating potential performance pitfalls. These topics are presented in the context of a variety of disciplines, including computer science, physics, and mathematics. The chapters include numerous programming exercises that range from easy to very challenging. This is an ideal book for students or professionals looking to learn parallel programming skills or to refresh their knowledge.

Leigh Little

Department of Computational Science
The College at Brockport, The State University of New York

An Introduction to Parallel Programming is a well-written, comprehensive book on the field of parallel computing. Students and practitioners alike will appreciate the relevant, up-to-date information. Peter Pacheco's very accessible writing style, combined with numerous interesting examples, keeps the reader's attention. In a field that races forward at a dizzying pace, this book hangs on for the wild ride covering the ins and outs of parallel hardware and software.

Kathy J. Liszka

Department of Computer Science
University of Akron

Parallel computing is the future and this book really helps introduce this complicated subject with practical and useful examples.

Andrew N. Sloss, FBCS

Consultant Engineer, ARM
Author of *ARM System Developer's Guide*

Preface

Parallel hardware has been ubiquitous for some time now. It's difficult to find a laptop, desktop, or server that doesn't use a multicore processor. Beowulf clusters are nearly as common today as high-powered workstations were during the 1990s, and cloud computing could make distributed-memory systems as accessible as desktops. In spite of this, most computer science majors graduate with little or no experience in parallel programming. Many colleges and universities offer upper-division elective courses in parallel computing, but since most computer science majors have to take numerous required courses, many graduate without ever writing a multithreaded or multiprocess program.

It seems clear that this state of affairs needs to change. Although many programs can obtain satisfactory performance on a single core, computer scientists should be made aware of the potentially vast performance improvements that can be obtained with parallelism, and they should be able to exploit this potential when the need arises.

An Introduction to Parallel Programming was written to partially address this problem. It provides an introduction to writing parallel programs using MPI, Pthreads, and OpenMP—three of the most widely used application programming interfaces (APIs) for parallel programming. The intended audience is students and professionals who need to write parallel programs. The prerequisites are minimal: a college-level course in mathematics and the ability to write serial programs in C. They are minimal because we believe that students should be able to start programming parallel systems *as early as possible*.

At the University of San Francisco, computer science students can fulfill a requirement for the major by taking the course, on which this text is based, immediately after taking the "Introduction to Computer Science I" course that most majors take in the first semester of their freshman year. We've been offering this course in parallel computing for six years now, and it has been our experience that there really is no reason for students to defer writing parallel programs until their junior or senior year. To the contrary, the course is popular, and students have found that using concurrency in other courses is much easier after having taken the Introduction course.

If second-semester freshmen can learn to write parallel programs by taking a class, then motivated computing professionals should be able to learn to write parallel programs through self-study. We hope this book will prove to be a useful resource for them.

About This Book

As we noted earlier, the main purpose of the book is to teach parallel programming in MPI, Pthreads, and OpenMP to an audience with a limited background in computer science and no previous experience with parallelism. We also wanted to make it as

flexible as possible so that readers who have no interest in learning one or two of the APIs can still read the remaining material with little effort. Thus, the chapters on the three APIs are largely independent of each other: they can be read in any order, and one or two of these chapters can be bypass. This independence has a cost: It was necessary to repeat some of the material in these chapters. Of course, repeated material can be simply scanned or skipped.

Readers with no prior experience with parallel computing should read Chapter 1 first. It attempts to provide a relatively nontechnical explanation of why parallel systems have come to dominate the computer landscape. The chapter also provides a short introduction to parallel systems and parallel programming.

Chapter 2 provides some technical background in computer hardware and software. Much of the material on hardware can be scanned before proceeding to the API chapters. Chapters 3, 4, and 5 are the introductions to programming with MPI, Pthreads, and OpenMP, respectively.

In Chapter 6 we develop two longer programs: a parallel n -body solver and a parallel tree search. Both programs are developed using all three APIs. Chapter 7 provides a brief list of pointers to additional information on various aspects of parallel computing.

We use the C programming language for developing our programs because all three APIs have C-language interfaces, and, since C is such a small language, it is a relatively easy language to learn—especially for C++ and Java programmers, since they are already familiar with C’s control structures.

Classroom Use

This text grew out of a lower-division undergraduate course at the University of San Francisco. The course fulfills a requirement for the computer science major, and it also fulfills a prerequisite for the undergraduate operating systems course. The only prerequisites for the course are either a grade of “B” or better in a one-semester introduction to computer science or a “C” or better in a two-semester introduction to computer science. The course begins with a four-week introduction to C programming. Since most students have already written Java programs, the bulk of what is covered is devoted to the use pointers in C.¹ The remainder of the course provides introductions to programming in MPI, Pthreads, and OpenMP.

We cover most of the material in Chapters 1, 3, 4, and 5, and parts of the material in Chapters 2 and 6. The background in Chapter 2 is introduced as the need arises. For example, before discussing cache coherence issues in OpenMP (Chapter 5), we cover the material on caches in Chapter 2.

The coursework consists of weekly homework assignments, five programming assignments, a couple of midterms, and a final exam. The homework usually involves

¹Interestingly, a number of students have said that they found the use of C pointers more difficult than MPI programming.

writing a very short program or making a small modification to an existing program. Their purpose is to insure that students stay current with the course work and to give them hands-on experience with the ideas introduced in class. It seems likely that the existence of the assignments has been one of the principle reasons for the course's success. Most of the exercises in the text are suitable for these brief assignments.

The programming assignments are larger than the programs written for homework, but we typically give students a good deal of guidance: We'll frequently include pseudocode in the assignment and discuss some of the more difficult aspects in class. This extra guidance is often crucial: It's not difficult to give programming assignments that will take far too long for the students to complete. The results of the midterms and finals, and the enthusiastic reports of the professor who teaches operating systems, suggest that the course is actually very successful in teaching students how to write parallel programs.

For more advanced courses in parallel computing, the text and its online support materials can serve as a supplement so that much of the information on the syntax and semantics of the three APIs can be assigned as outside reading. The text can also be used as a supplement for project-based courses and courses outside of computer science that make use of parallel computation.

Support Materials

The book's website is located at <http://www.mkp.com/pacheco>. It will include errata and links to sites with related materials. Faculty will be able to download complete lecture notes, figures from the text, and solutions to the exercises and the programming assignments. All users will be able to download the longer programs discussed in the text.

We would greatly appreciate readers letting us know of any errors they find. Please send an email to peter@usfca.edu if you do find a mistake.

Acknowledgments

In the course of working on this book, I've received considerable help from many individuals. Among them I'd like to thank the reviewers who read and commented on the initial proposal: Fikret Ercal (Missouri University of Science and Technology), Dan Harvey (Southern Oregon University), Joel Hollingsworth (Elon University), Jens Mache (Lewis and Clark College), Don McLaughlin (West Virginia University), Manish Parashar (Rutgers University), Charlie Peck (Earlham College), Stephen C. Renk (North Central College), Rolfe Josef Sassenfeld (The University of Texas at El Paso), Joseph Sloan (Wofford College), Michela Taufer (University of Delaware), Pearl Wang (George Mason University), Bob Weems (University of Texas at Arlington), and Cheng-Zhong Xu (Wayne State University).

I'm also deeply grateful to the following individuals for their reviews of various chapters of the book: Duncan Buell (University of South Carolina), Matthias Gobbert (University of Maryland, Baltimore County), Krishna Kavi (University of North Texas), Hong Lin (University of Houston—Downtown), Kathy Liszka (University of Akron), Leigh Little (The State University of New York), Xinlian Liu (Hood College), Henry Tufo (University of Colorado at Boulder), Andrew Sloss (Consultant Engineer, ARM), and Gengbin Zheng (University of Illinois). Their comments and suggestions have made the book immeasurably better. Of course, I'm solely responsible for remaining errors and omissions.

Kathy Liszka is also preparing slides that can be used by faculty who adopt the text, and a former student, Jinyoung Choi, is working on preparing a solutions manual. Thanks to both of them.

The staff of Morgan Kaufmann has been very helpful throughout this project. I'm especially grateful to the developmental editor, Nate McFadden. He gave me much valuable advice, and he did a terrific job arranging for the reviews. He's also been tremendously patient with all the problems I've encountered over the past few years. Thanks also to Marilyn Rash and Megan Guiney, who have been very prompt and efficient during the production process.

My colleagues in the computer science and mathematics departments at USF have been extremely helpful during my work on the book. I'd like to single out Professor Gregory Benson for particular thanks: his understanding of parallel computing—especially Pthreads and semaphores—has been an invaluable resource for me. I'm also very grateful to our system administrators, Alexey Fedosov and Colin Bean. They've patiently and efficiently dealt with all of the “emergencies” that cropped up while I was working on programs for the book.

I would never have been able to finish this book without the encouragement and moral support of my friends Holly Cohn, John Dean, and Robert Miller. They helped me through some very difficult times, and I'll be eternally grateful to them.

My biggest debt is to my students. They showed me what was too easy, and what was far too difficult. In short, they taught me how to teach parallel computing. My deepest thanks to all of them.

About the Author

Peter Pacheco received a PhD in mathematics from Florida State University. After completing graduate school, he became one of the first professors in UCLA's "Program in Computing," which teaches basic computer science to students at the College of Letters and Sciences there. Since leaving UCLA, he has been on the faculty of the University of San Francisco. At USF Peter has served as chair of the computer science department and is currently chair of the mathematics department.

His research is in parallel scientific computing. He has worked on the development of parallel software for circuit simulation, speech recognition, and the simulation of large networks of biologically accurate neurons. Peter has been teaching parallel computing at both the undergraduate and graduate levels for nearly twenty years. He is the author of *Parallel Programming with MPI*, published by Morgan Kaufmann Publishers.

Contents

Preface	vi
Acknowledgments	ix
About the Author	x
CHAPTER 1 Why Parallel Computing?	1
1.1 Why We Need Ever-Increasing Performance	2
1.2 Why We're Building Parallel Systems	3
1.3 Why We Need to Write Parallel Programs	3
1.4 How Do We Write Parallel Programs?	6
1.5 What We'll Be Doing	8
1.6 Concurrent, Parallel, Distributed	9
1.7 The Rest of the Book	10
1.8 A Word of Warning	10
1.9 Typographical Conventions	11
1.10 Summary	12
1.11 Exercises	12
CHAPTER 2 Parallel Hardware and Parallel Software	15
2.1 Some Background	15
2.1.1 The von Neumann architecture	15
2.1.2 Processes, multitasking, and threads	17
2.2 Modifications to the von Neumann Model	18
2.2.1 The basics of caching	19
2.2.2 Cache mappings	20
2.2.3 Caches and programs: an example	22
2.2.4 Virtual memory	23
2.2.5 Instruction-level parallelism	25
2.2.6 Hardware multithreading	28
2.3 Parallel Hardware	29
2.3.1 SIMD systems	29
2.3.2 MIMD systems	32
2.3.3 Interconnection networks	35
2.3.4 Cache coherence	43
2.3.5 Shared-memory versus distributed-memory	46
2.4 Parallel Software	47
2.4.1 Caveats	47
2.4.2 Coordinating the processes/threads	48
2.4.3 Shared-memory	49

2.4.4	Distributed-memory	53
2.4.5	Programming hybrid systems	56
2.5	Input and Output	56
2.6	Performance	58
2.6.1	Speedup and efficiency	58
2.6.2	Amdahl's law	61
2.6.3	Scalability	62
2.6.4	Taking timings	63
2.7	Parallel Program Design	65
2.7.1	An example	66
2.8	Writing and Running Parallel Programs	70
2.9	Assumptions	70
2.10	Summary	71
2.10.1	Serial systems	71
2.10.2	Parallel hardware	73
2.10.3	Parallel software	74
2.10.4	Input and output	75
2.10.5	Performance	75
2.10.6	Parallel program design	76
2.10.7	Assumptions	76
2.11	Exercises	77
CHAPTER 3	Distributed-Memory Programming with MPI	83
3.1	Getting Started.....	84
3.1.1	Compilation and execution.....	84
3.1.2	MPI programs.....	86
3.1.3	MPI_Init and MPI_Finalize	86
3.1.4	Communicators, MPI_Comm_size and MPI_Comm_rank.....	87
3.1.5	SPMD programs	88
3.1.6	Communication	88
3.1.7	MPI_Send	88
3.1.8	MPI_Recv	90
3.1.9	Message matching	91
3.1.10	The status_p argument.....	92
3.1.11	Semantics of MPI_Send and MPI_Recv	93
3.1.12	Some potential pitfalls	94
3.2	The Trapezoidal Rule in MPI.....	94
3.2.1	The trapezoidal rule	94
3.2.2	Parallelizing the trapezoidal rule	96

3.3	Dealing with I/O	97
3.3.1	Output	97
3.3.2	Input	100
3.4	Collective Communication.....	101
3.4.1	Tree-structured communication.....	102
3.4.2	MPI_Reduce	103
3.4.3	Collective vs. point-to-point communications	105
3.4.4	MPI_Allreduce	106
3.4.5	Broadcast.....	106
3.4.6	Data distributions	109
3.4.7	Scatter	110
3.4.8	Gather	112
3.4.9	Allgather	113
3.5	MPI Derived Datatypes	116
3.6	Performance Evaluation of MPI Programs.....	119
3.6.1	Taking timings	119
3.6.2	Results.....	122
3.6.3	Speedup and efficiency.....	125
3.6.4	Scalability	126
3.7	A Parallel Sorting Algorithm	127
3.7.1	Some simple serial sorting algorithms	127
3.7.2	Parallel odd-even transposition sort	129
3.7.3	Safety in MPI programs	132
3.7.4	Final details of parallel odd-even sort	134
3.8	Summary	136
3.9	Exercises	140
3.10	Programming Assignments	147
CHAPTER 4	Shared-Memory Programming with Pthreads.....	151
4.1	Processes, Threads, and Pthreads	151
4.2	Hello, World	153
4.2.1	Execution.....	153
4.2.2	Preliminaries	155
4.2.3	Starting the threads	156
4.2.4	Running the threads	157
4.2.5	Stopping the threads	158
4.2.6	Error checking	158
4.2.7	Other approaches to thread startup	159
4.3	Matrix-Vector Multiplication	159
4.4	Critical Sections	162

4.5	Busy-Waiting	165
4.6	Mutexes	168
4.7	Producer-Consumer Synchronization and Semaphores.....	171
4.8	Barriers and Condition Variables.....	176
4.8.1	Busy-waiting and a mutex	177
4.8.2	Semaphores	177
4.8.3	Condition variables	179
4.8.4	Pthreads barriers	181
4.9	Read-Write Locks	181
4.9.1	Linked list functions.....	181
4.9.2	A multi-threaded linked list.....	183
4.9.3	Pthreads read-write locks	187
4.9.4	Performance of the various implementations	188
4.9.5	Implementing read-write locks	190
4.10	Caches, Cache Coherence, and False Sharing	190
4.11	Thread-Safety.....	195
4.11.1	Incorrect programs can produce correct output.....	198
4.12	Summary	198
4.13	Exercises	200
4.14	Programming Assignments	206
CHAPTER 5	Shared-Memory Programming with OpenMP	209
5.1	Getting Started.....	210
5.1.1	Compiling and running OpenMP programs.....	211
5.1.2	The program	212
5.1.3	Error checking	215
5.2	The Trapezoidal Rule.....	216
5.2.1	A first OpenMP version	216
5.3	Scope of Variables	220
5.4	The Reduction Clause	221
5.5	The <code>parallel for</code> Directive	224
5.5.1	Caveats	225
5.5.2	Data dependences.....	227
5.5.3	Finding loop-carried dependences.....	228
5.5.4	Estimating π	229
5.5.5	More on scope	231
5.6	More About Loops in OpenMP: Sorting	232
5.6.1	Bubble sort	232
5.6.2	Odd-even transposition sort.....	233
5.7	Scheduling Loops	236
5.7.1	The <code>schedule</code> clause.....	237

5.7.2	The static schedule type	238
5.7.3	The dynamic and guided schedule types.....	239
5.7.4	The runtime schedule type	239
5.7.5	Which schedule?.....	241
5.8	Producers and Consumers.....	241
5.8.1	Queues.....	241
5.8.2	Message-passing	242
5.8.3	Sending messages	243
5.8.4	Receiving messages	243
5.8.5	Termination detection	244
5.8.6	Startup	244
5.8.7	The atomic directive.....	245
5.8.8	Critical sections and locks	246
5.8.9	Using locks in the message-passing program.....	248
5.8.10	critical directives, atomic directives, or locks?.....	249
5.8.11	Some caveats.....	249
5.9	Caches, Cache Coherence, and False Sharing	251
5.10	Thread-Safety.....	256
5.10.1	Incorrect programs can produce correct output	258
5.11	Summary	259
5.12	Exercises	263
5.13	Programming Assignments	267
CHAPTER 6	Parallel Program Development	271
6.1	Two n-Body Solvers	271
6.1.1	The problem.....	271
6.1.2	Two serial programs	273
6.1.3	Parallelizing the n -body solvers	277
6.1.4	A word about I/O	280
6.1.5	Parallelizing the basic solver using OpenMP	281
6.1.6	Parallelizing the reduced solver using OpenMP	284
6.1.7	Evaluating the OpenMP codes	288
6.1.8	Parallelizing the solvers using pthreads	289
6.1.9	Parallelizing the basic solver using MPI	290
6.1.10	Parallelizing the reduced solver using MPI	292
6.1.11	Performance of the MPI solvers	297
6.2	Tree Search	299
6.2.1	Recursive depth-first search.....	302
6.2.2	Nonrecursive depth-first search.....	303
6.2.3	Data structures for the serial implementations.....	305

6.2.4	Performance of the serial implementations	306
6.2.5	Parallelizing tree search	306
6.2.6	A static parallelization of tree search using threads	309
6.2.7	A dynamic parallelization of tree search using threads	310
6.2.8	Evaluating the pthreads tree-search programs	315
6.2.9	Parallelizing the tree-search programs using OpenMP	316
6.2.10	Performance of the OpenMP implementations	318
6.2.11	Implementation of tree search using MPI and static partitioning	319
6.2.12	Implementation of tree search using MPI and dynamic partitioning	327
6.3	A Word of Caution	335
6.4	Which API?	335
6.5	Summary	336
6.5.1	Pthreads and OpenMP	337
6.5.2	MPI	338
6.6	Exercises	341
6.7	Programming Assignments	350
CHAPTER 7	Where to Go from Here	353
References	357
Index	361