HZ BOOKS 华章教育

经 典 原 版 书 库

PEARSON Addison Wesley

并行程序设计原理

（英文版）

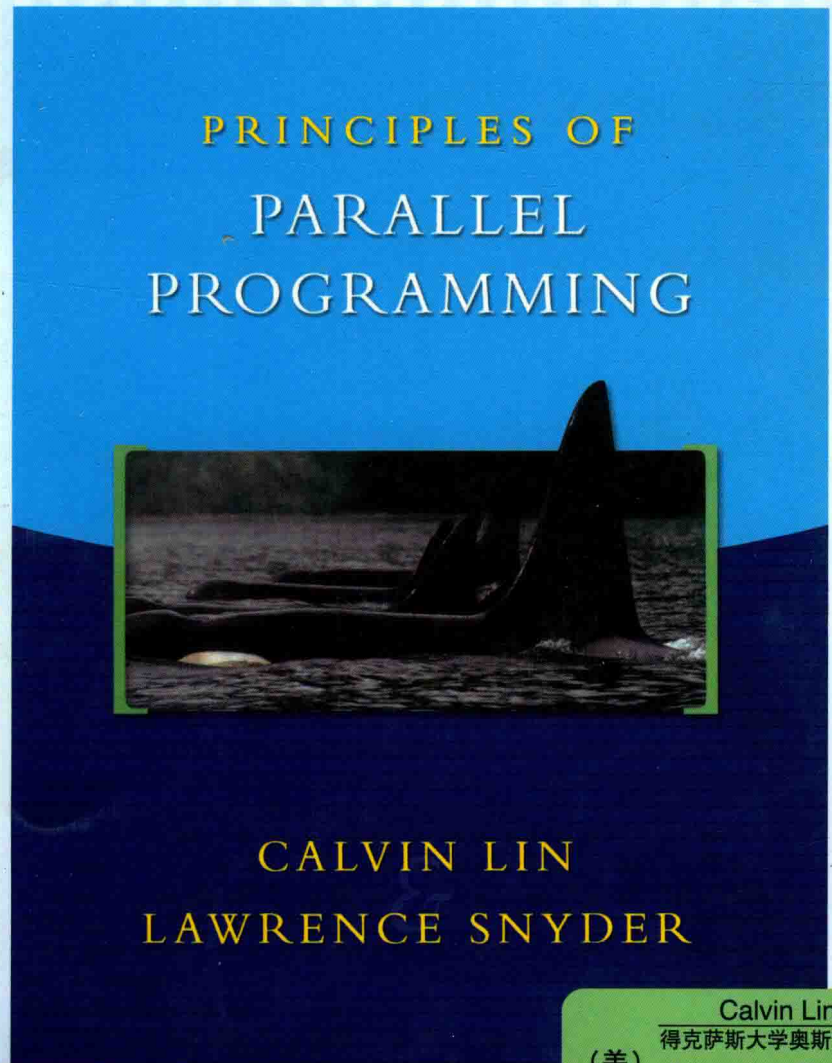

（美） Calvin Lin 得克萨斯大学奥斯汀分校
Lawrence Snyder 华盛顿大学西雅图分校 著

机械工业出版社
China Machine Press

经 典 原 版 书 库

并行程序设计原理

（英文版）

Principles of Parallel Programming

（美）Calvin Lin 得克萨斯大学奥斯汀分校
Lawrence Snyder 华盛顿大学西雅图分校 著

机械工业出版社
China Machine Press

本书版权登记号：图字：01-2008-3270

图书在版编目（CIP）数据

并行程序设计原理（英文版）/（美）林（Lin, C.），（美）斯奈德（Snyder, L.）著. —北京：机械工业出版社，2008.8

（经典原版书库）

书名原文：Principles of Parallel Programming

ISBN 978-7-111-24734-0

Ⅰ.并…　Ⅱ.①林…　②斯…　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（2008）第109994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

北京瑞德印刷有限公司印刷·新华书店北京发行所发行

2008年8月第1版第1次印刷

170mm×242mm　·　22.25印张

标准书号：ISBN 978-7-111-24734-0

定价：49.00元

凡购本书，如有倒页、脱页、缺页，由本社发行部调换

本社购书热线：（010）68326294

出版者的话

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，还揭橥了学术的源变，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分社较早意识到“出版要为教育服务”。自1998年开始，华章分社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，Elsevien，MIT，John Wiley & Sons Wiley，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，Brain W. Kernighan，Dennis Ritchie Jim Gray，Afred V. Aho，John E. Hopcroft，Jeffrey D. Ullman，Abraham Silberschatz，William Stallings，Donald E. Knuth，John L. Hennessy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、研究及庋藏。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。华章分社欢迎老师和读者对我们的工作提出建议或给予指正，我们的联系方法如下：

华章网站：www.hzbook.com

电子邮件：hzedu@hzbook.com

联系电话：（010）68995264

联系地址：北京市西城区百万庄南街1号

邮政编码：100037

华章教育

华章科技图书出版中心

To Mom and Dad
(Josette and Min Shuey)

To Julie, Dave, and Dan

Preface

Welcome!

For readers who are motivated by the advent of multi-core chips to learn parallel programming, you've come to the right place. This book is written for a world in which parallel computers are everywhere, ranging from laptops with two-core chips to supercomputers to huge data-center clusters that index the Internet.

This book focuses on scalable parallelism, that is, the ability of a parallel program to run well on any number of processors. This notion is critical for two reasons: (1) Most of the techniques needed to create scalable parallel computations are the same techniques that produce efficient solutions on a multi-core chip, and (2) while multi-core chips currently have a modest number of processors, typically 2–8, the number of cores per chip promises to increase dramatically in the coming years, making the notion of scalable parallelism directly relevant. Thus, while today's multi-core chips offer opportunities for low latency communication among cores, this characteristic is likely a short-term advantage, as on-chip delays to different parts of the chip will become increasingly apparent as the number of cores grows. So, we focus not on exploiting such short-term advantages, but on emphasizing approaches that work well now and in the future. Of course, multi-core chips present their own challenges, particularly with their limited bandwidth to off-chip memory and their limited aggregate on-chip cache. This book discusses these issues as well.

First, we discuss the principles that underlie effective and efficient parallel programs. Learning the principles is essential to acquiring any capability as sophisticated as programming, of course, but principles are perhaps even more important for parallel programming because the state of the art changes rapidly. Training that is tied too closely to a specific computer or language will not have the staying power needed to keep pace with advancing technology. But the principles—concepts that apply to any parallel computing system and ideas that exploit these features—lead to an understanding that is timeless and knowledge that will always be applicable.

But we do more than discuss abstract concepts. We also apply those principles to everyday computations, which makes the book very practical. We introduce several parallel programming systems, and we describe how to apply the principles in those

programming systems. On completion, we expect readers to be able to write parallel programs. Indeed, the final chapter is devoted to parallel programming techniques and the development of a term-long parallel programming capstone project.

Audience

Our intended audience is anyone—students or professionals—who has written successful programs in C or similar languages and who describes himself as a programmer. It is helpful to have a basic idea of how a computer executes sequential programs, including knowledge of the fetch/execute cycle and basics of caching. This book was originally targeted to upper level undergraduate computer science majors or first year graduate students with a CS undergraduate degree, and it continues to be appropriate for that level. However, as the book evolved, we reduced the assumed knowledge and emphasized pedagogy in the belief that if some explanations cover knowledge the reader already has, it's easy to skip forward.

Organization

Because parallel programming is not a direct extension of sequential programming with which the reader is doubtless familiar, we have organized this book into four parts:

Foundations: Chapters 1–3

Abstractions: Chapters 4–5

Languages: Chapters 6–9

Looking Forward: Chapters 10–11

To enable you to select intelligently from these parts, we now explain their goals and content.

Foundations. In Chapter 1 we discover the many issues that parallel programmers must address by showing how difficult it is to implement a computation that is trivial when written for sequential computers. The example focuses our attention on issues that concern us throughout the entire book, but it also emphasizes the importance of understanding how a parallel computer operates. Chapter 2 introduces five different types of parallel computers, giving a few details about their architecture and their ability to scale to a larger size. There are two key conclusions from the chapter: First, unlike sequential computing, there is no standard architecture. Second, to be successful at spanning this architectural diversity we need an abstract machine model to guide our programming. And we give one. With the architectures in mind, Chapter 3 covers basic ideas of concurrency, including threads and processes, latency, bandwidth, speedup, and so forth, with an emphasis on issues related to performance. These foundations of Part 1 prepare us for an exploration of algorithms and abstractions.

Abstractions. As an aid to designing and discussing parallel algorithms, Chapter 4 introduces an informal pseuodcode notation for writing parallel programs in a language-independent way. The notation has a variety of features that span various programming models and approaches, allowing us to discuss algorithms without bias toward any particular language or machine. To bootstrap your thinking about parallel algorithms, Chapter 5 covers a series of basic algorithmic techniques. By the end of Part 2, you should be able to conceptualize ways to solve a problem in parallel, bringing us to the final issue of encoding your algorithms in a concrete parallel programming language.

Languages. There is no single parallel programming language that fulfills the role that, say, C or Java plays in sequential programming, that is, a language widely known and accepted as a baseline medium to encode algorithms. As a result, Part 3 introduces three kinds of parallel programming languages: thread-based (Chapter 6), message-passing (Chapter 7), and high-level (Chapter 8). We cover each language well enough for you to write small exercises; serious computations require a more complete language introduction that is available through online resources. In addition to introducing a language, each chapter includes a brief overview of related languages that have a following in the parallel programming community. Chapter 9 briefly compares and contrasts all of the languages presented, noting their strengths and weaknesses. There is benefit to reading all three chapters, but we realize that many readers will focus on one approach, so these chapters are independent of one another.

Onward. Part 4 looks to the future. Chapter 10 covers a series of new, promising parallel technologies that will doubtless impact future research and practice. In our view, they are not quite "ready for prime time," but they are important and worth becoming familiar with even before they are fully deployed. Finally, Chapter 11 focuses on hands-on techniques for programming parallel machines. The first two sections of the chapter can be read early in your study of parallel programming, perhaps together with your study of abstractions in Chapters 4 and 5. But the main goal of the chapter is to assist you in writing a substantial program as a capstone design project. In this capacity we assume that you will return to Chapter 11 repeatedly.

Using This Book

Although the content is presented in a logical order, it is not necessary to read this book front to back. Indeed, in a one term course, it may be sensible to begin programming exercises before all of the topics have been introduced. We see the following as a sensible general plan:

- Chapters 1, 2
- Chapter 11 first section, Chapter 3 through Performance Tradeoffs; begin programming exercises
- Chapters 4, 5

- One of Chapters 6–8, programming language chapters
- Complete Chapter 3 and 11, begin term project
- Complete remaining chapters in order: language chapters, Chapters 9, 10

There is, of course, no harm in reading the book straight through, but the advantage of this approach is that the reading and programming can proceed in parallel.

Acknowledgments

Sincere thanks are due to E Christopher Lewis and Robert van de Geijn, who critiqued an early draft of this book. Thanks also to the following reviewers for their valuable feedback and suggestions:

David Bader, Georgia Institute of Technology

Purushotham Bangalore, University of Alabama, Birmingham

John Cavazos, University of Delaware

Sandhya Dwarkadas, University of Rochester

John Gilbert, UC Santa Barbara

Robert Henry, Cray Inc.

E Christopher Lewis, VMWare

Kai Li, Princeton

Glenn Reinman, UCLA

Darko Stefanovic, University of New Mexico

We thank Karthik Murthy and Brandon Plost for their assistance in writing and running parallel programs and for finding bugs in the text, and we are grateful to Bobby Blumofe, whose early collaborations on a multi-threaded programming course are evident in many places in the book. We recognize and thank the students of the Parallel Programming Environments Seminar (CSE590o) at the University of Washington in autumn quarter, 2006 for their contributions to the text: Ivan Beschastnikh, Alex Colburn, Roxana Geambasu, Sangyun Hahn, Ethan Katz-Bassett, Nathan Kuchta, Harsha Madhyastha, Marianne Shaw, Brian Van Essen, and Benjamin Ylvisaker. Other contributors are Sonja Keserovic, Kate Moore, Brad Chamberlain, Steven Deitz, Dan Grossman, Jeff Diamond, Don Fussell, Bill Mark, and David Mohr.

We would like to thank our editor, Matt Goldstein, and the Addison Wesley team: Sarah Milmore, Marilyn Lloyd, Barbara Atkinson, Joyce Wells, and Chris Kelly. Thanks to Gillian Hall who has been especially tolerant of our antics.

Finally, we thank our families for their patience through the writing of this book.

Calvin Lin
Lawrence Snyder
February 2008

Any sufficiently advanced technology
is indistinguishable from magic.

—Arthur C. Clarke
Profiles of the Future, 1961

造就龙脉神韵 成就品质传奇

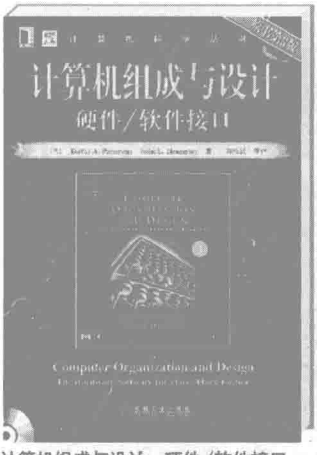

计算机组成与设计：硬件/软件接口

作　者：[美] John L.Hennessy,David A.Patterson
译　者：郑纬民 等
中文版：7-111-20214-1 定 价：75.00元
英文版：7-111-19339-3 定 价：85.00元
■本书采用MIPS处理器作为展示计算机硬件技术基本核心

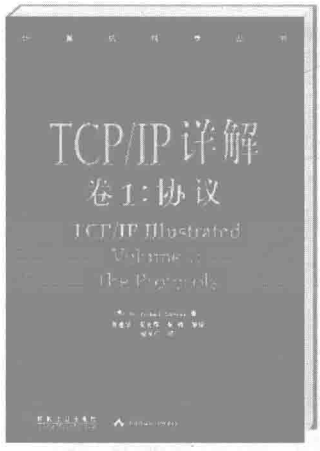

TCP／IP详解 卷1：协议

作　者：[美] W. Richard Stevens
译　者：范建华 胥光辉 张涛 等
中文版：7-111-09505-7 定 价：39.00元
英文版：7-111-07566-8 定 价：45.00元
■国际知名Unix和网络专家杰作 ■中文版畅销八年，持续热销中

数据挖掘：概念与技术，原书第2版

作　者：[加] Jiawei Han 等
译　者：范明 孟小峰
中文版：ISBN 7-111-20538-3 定 价：55.00元
英文版：ISBN 7-111-18828-4 定 价：79.00元
■KDnuggests读者评选为最受欢迎的数据挖掘专著

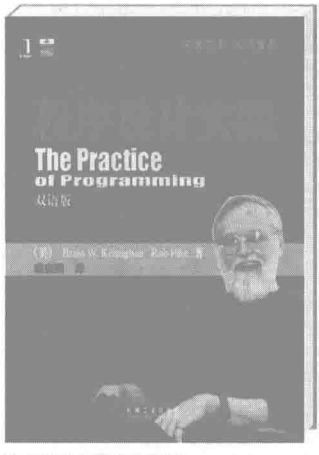

程序设计实践(双语版)

作　者：[美] Brian W Kernighan
书　号：978-7-111-21127-3 书 号：59.00
■著名计算机专家Brian W. Kernighan的畅销作品，曾在国内外受到广泛赞誉
■译者裘宗燕，名著名译，双语重温

华章图书 服务中国教育

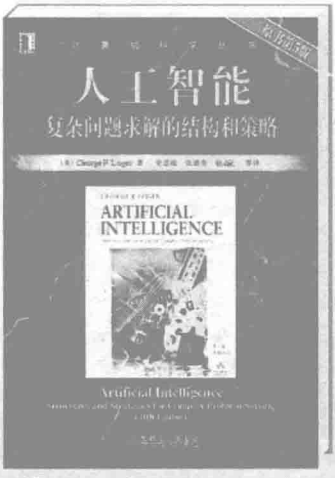

人工智能：复杂问题求解的结构和策略，原书第5版

作　者：[美] George F.Luger 译 者：史忠植 等
中文版：7-111-15916-0 定 价：75.00元
英文版：7-111-15916-0 定 价：75.00元
■加州伯克利分校教材，注重该领域的产业实践及计算机应用

高性能嵌入式计算

作　者：[美] Wayne Wolf 译 者：汪东升
中文版：2008年
英文版：7-111-20416-6 定 价：65.00元
■包含大量现实世界中嵌入式计算应用和体系结构的实例

计算机集成制造，原书第3版

作　者：[美] James A.Rehg 译 者：夏链 等
中文版：978-7-111-21017-7 定 价：55.00元
英文版：7-111-14687-5　　定 价：59.00元
■描述了世界各地企业所采用的不同类型的制造系统或生产策略，并展示了如何运用技术来解决实际的工业问题

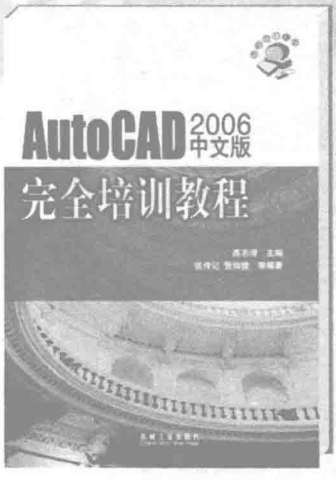

Auto CAD 2006中文版完全培训教程

作 者：高志清
书 号：7-111-19057-2
定 价：46.00元
■本书实例典型、图文并茂，系统讲解了AutoCad 中文版的二维绘图功能和三维建模功能

教师服务登记表

尊敬的老师：

您好！感谢您购买我们出版的 _____ 教材。

机械工业出版社华章公司本着为服务高等教育的出版原则，为进一步加强与高校教师的联系与沟通，更好地为高校教师服务，特制此表，请您填妥后发回给我们，我们将定期向您寄送华章公司最新的图书出版信息。为您的教材、论著或译著的出版提供可能的帮助。欢迎您对我们的教材和服务提出宝贵的意见，感谢您的大力支持与帮助！

个人资料（请用正楷完整填写）

<table>
<tr><td>教师姓名</td><td></td><td>□先生
□女士</td><td>出生年月</td><td></td><td>职务</td><td></td><td colspan="2">职称：□教授 □副教授
□讲师 □助教 □其他</td></tr>
<tr><td>学校</td><td colspan="2"></td><td>学院</td><td colspan="2"></td><td>系别</td><td colspan="2"></td></tr>
<tr><td rowspan="2">联系电话</td><td colspan="3" rowspan="2">办公：
宅电：
移动：</td><td>联系地址及邮编</td><td colspan="3"></td></tr>
<tr><td>E-mail</td><td colspan="3"></td></tr>
<tr><td>学历</td><td></td><td>毕业院校</td><td colspan="2"></td><td>国外进修及讲学经历</td><td colspan="3"></td></tr>
<tr><td>研究领域</td><td colspan="8"></td></tr>
<tr><td colspan="2">主讲课程</td><td colspan="2">现用教材名</td><td>作者及出版社</td><td>共同授课教师</td><td colspan="3">教材满意度</td></tr>
<tr><td colspan="2">课程：
□专 □本 □研
人数： 学期：□春□秋</td><td colspan="2"></td><td></td><td></td><td colspan="3">□满意 □一般
□不满意 □希望更换</td></tr>
<tr><td colspan="2">课程：
□专 □本 □研
人数： 学期：□春□秋</td><td colspan="2"></td><td></td><td></td><td colspan="3">□满意 □一般
□不满意 □希望更换</td></tr>
<tr><td colspan="9">样书申请</td></tr>
<tr><td colspan="3">已出版著作</td><td colspan="2"></td><td colspan="2">已出版译作</td><td colspan="2"></td></tr>
<tr><td colspan="4">是否愿意从事翻译/著作工作 □是 □否</td><td>方向</td><td colspan="4"></td></tr>
<tr><td>意见和建议</td><td colspan="8"></td></tr>
</table>

填妥后请选择以下任何一种方式将此表返回：（如方便请赐名片）

地 址：北京市西城区百万庄南街1号 华章公司营销中心 邮编：100037

电 话：(010) 68353079 88378995 传真：(010)68995260

E-mail:hzedu@hzbook.com markerting@hzbook.com 图书详情可登录http://www.hzbook.com网站查询

Contents