

Mc
Graw
Hill Education

大学计算机教育国外著名教材系列(影印版)

PARALLEL PROGRAMMING IN C WITH MPI AND OPENMP

并行程序设计 C、MPI与OpenMP



Michael J. Quinn 著

清华大学出版社

PARALLEL PROGRAMMING IN C WITH MPI AND OPENMP

并行程序设计 C、MPI与OpenMP

内 容 简 介

本书是美国Oregon州立大学的Michael J. Quinn教授在多年讲授“并行程序设计”课程的基础上编写而成的，主要介绍用C语言，并结合使用MPI和OpenMP进行并行程序设计，内容包括并行体系结构、并行算法设计、消息传递编程、Eratosthenes筛法、Floyd算法、性能分析、矩阵向量乘法、文档分类、蒙特卡洛法、矩阵乘法、线性方程组求解、有限差分方法、排序、快速傅立叶变换、组合搜索、共享存储编程、融合OpenMP和MPI以及5个附录。

本书按授课方式安排章节，通过划分、通信、集聚和映射等四步的并行程序设计方法，来解决各种实际的并行性问题，使读者掌握系统化的并行程序设计方法，开发出高效的并行程序。

本书不仅是一本优秀的并行程序设计教材，对广大的相关专业人员也很有参考价值。

This edition is authorized for sale in the People's Republic of China only (excluding Hong Kong, Macao SAR and Taiwan).

此英文影印版仅限在中华人民共和国境内(不包括中国香港、澳门特别行政区及中国台湾地区)销售。

ISBN 7-302-11157-X



9 787302 111573 >

定价:33.00元

Mc
Graw
Hill Education

<http://www.mheducation.com>

Mc
Graw
Hill

影印版

PARALLEL PROGRAMMING
INC WITH MPI AND OPENMP

并行编程与
C、MPI 与 OpenMP

清华大学



大学计算机教育国外著名教材系列(影印版)

Parallel Programming
in C with MPI and OpenMP

并行程序设计 C、MPI 与 OpenMP

Michael J. Quinn

清华大学出版社
北 京

Michael J. Quinn

Parallel Programming in C with MPI and OpenMP

EISBN: 0-07-282256-2

Copyright © 2004 by The McGraw-Hill Companies, Inc.

Original language published by The McGraw-Hill Companies, Inc. All Rights reserved. No part of this publication may be reproduced or distributed by any means, or stored in a database or retrieval system, without the prior written permission of the publisher.

Authorized English language edition jointly published by McGraw-Hill Education (Asia) Co. and Tsinghua University Press. This edition is authorized for sale only to the educational and training institutions, and within the territory of the People's Republic of China (excluding Hong Kong, Macao 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.

本书英文影印版由清华大学出版社和美国麦格劳-希尔教育出版(亚洲)公司合作出版。此版本仅限在中华人民共和国境内(不包括中国香港、澳门特别行政区及中国台湾地区)针对教育及培训机构之销售。未经许可之出口,视为违反著作权法,将受法律之制裁。未经出版者预先书面许可,不得以任何方式复制或抄袭本书的任何部分。

北京市版权局著作权合同登记号 图字:01-2005-0654 号

版权所有,翻印必究。举报电话:010-62782989 13501256678 13801310933

本书封面贴有 McGraw-Hill 公司防伪标签,无标签者不得销售。

图书在版编目(CIP)数据

并行程序设计: C、MPI 与 OpenMP 英文/(美)奎因(Quinn, M. J.)著. —影印本. —北京:清华大学出版社, 2005. 8

(大学计算机教育国外著名教材系列)

ISBN 7-302-11157-X

I. 并… II. 奎… III. 并行程序—程序设计—高等学校—教材—英文 IV. TP311.11

中国版本图书馆 CIP 数据核字(2005)第 058944 号

出版者:清华大学出版社

<http://www.tup.com.cn>

社总机:010-62770175

地址:北京清华大学学研大厦

邮编:100084

客户服务:010-62776969

印刷者:北京牛山世兴印刷厂

装订者:三河市金元装订厂

发行者:新华书店总店北京发行所

开本:148×210 印张:16.875

版次:2005年8月第1版 2005年8月第1次印刷

书号:ISBN 7-302-11157-X/TP·7373

印数:1~3000

定价:33.00 元

出版说明

进入 21 世纪，世界各国的经济、科技以及综合国力的竞争将更加激烈。竞争的中心无疑是对人才的竞争。谁拥有大量高素质的人才，谁就能在竞争中取得优势。高等教育，作为培养高素质人才的事业，必然受到高度重视。目前我国高等教育的教材更新较慢，为了加快教材的更新频率，教育部正在大力促进我国高校采用国外原版教材。

清华大学出版社从 1996 年开始，与国外著名出版公司合作，影印出版了“大学计算机教育丛书(影印版)”等一系列引进图书，受到国内读者的欢迎和支持。跨入 21 世纪，我们本着为我国高等教育教材建设服务的初衷，在已有的基础上，进一步扩大选题内容，改变图书开本尺寸，一如既往地请有关专家挑选适用于我国高等本科及研究生计算机教育的国外经典教材或著名教材，组成本套“大学计算机教育国外著名教材系列(影印版)”，以飨读者。深切期盼读者及时将使用本系列教材的效果和意见反馈给我们。更希望国内专家、教授积极向我们推荐国外计算机教育的优秀教材，以利我们把“大学计算机教育国外著名教材系列(影印版)”做得更好，更适合高校师生的需要。

清华大学出版社

PREFACE

This book is a practical introduction to parallel programming in C using the MPI (Message Passing Interface) library and the OpenMP application programming interface. It is targeted to upper-division undergraduate students, beginning graduate students, and computer professionals learning this material on their own. It assumes the reader has a good background in C programming and has had an introductory class in the analysis of algorithms.

Fortran programmers interested in parallel programming can also benefit from this text. While the examples in the book are in C, the underlying concepts of parallel programming with MPI and OpenMP are essentially the same for both C and Fortran programmers.

In the past twenty years I have taught parallel programming to hundreds of undergraduate and graduate students. In the process I have learned a great deal about the sorts of problems people encounter when they begin “thinking in parallel” and writing parallel programs. Students benefit from seeing programs designed and implemented step by step. My philosophy is to introduce new functionality “just in time.” As much as possible, every new concept appears in the context of solving a design, implementation, or analysis problem. When you see the symbol



in a page margin, you’ll know I’m presenting a key concept.

The first two chapters explain when and why parallel computing began and gives a high-level overview of parallel architectures. Chapter 3 presents Foster’s parallel algorithm design methodology and shows how it is used through several case studies. Chapters 4, 5, 6, 8, and 9 demonstrate how to use the design methodology to develop MPI programs that solve a series of progressively more difficult programming problems. The 27 MPI functions presented in these chapters are a robust enough subset to implement parallel programs for a wide variety of applications. These chapters also introduce functions that simplify matrix and vector I/O. The source code for this I/O library appears in Appendix B.

The programs of Chapters 4, 5, 6, and 8 have been benchmarked on a commodity cluster of microprocessors, and these results appear in the text. Because new generations of microprocessors appear much faster than books can be produced, readers will observe that the processors are several generations old. The point of presenting the results is not to amaze the reader with the speed of the computations. Rather, the purpose of the benchmarking is to demonstrate that knowledge of the latency and bandwidth of the interconnection network, combined with information about the performance of a sequential program, are often sufficient to allow reasonably accurate predictions of the performance of a parallel program.

Chapter 7 focuses on four metrics for analyzing and predicting the performance of parallel systems: Amdahl's Law, Gustafson-Barsis' Law, the Karp-Flatt metric, and the isoefficiency metric.

Chapters 10–16 provide additional examples of how to analyze a problem and design a good parallel algorithm to solve it. At this point the development of MPI programs implementing the parallel algorithms is left to the reader. I present Monte Carlo methods and the challenges associated with parallel random number generation. Later chapters present a variety of key algorithms: matrix multiplication, Gaussian elimination, the conjugate gradient method, finite difference methods, sorting, the fast Fourier transform, backtrack search, branch-and-bound search, and alpha-beta search.

Chapters 17 and 18 are an introduction to the new shared-memory programming standard OpenMP. I present the features of OpenMP as needed to convert sequential code segments into parallel ones. I use two case studies to demonstrate the process of transforming MPI programs into hybrid MPI/OpenMP programs that can exhibit higher performance on multiprocessor clusters than programs based solely on MPI.

This book has more than enough material for a one-semester course in parallel programming. While parallel programming is more demanding than typical programming, it is also more rewarding. Even with a teacher's instruction and support, most students are unnerved at the prospect of harnessing multiple processors to perform a single task. However, this fear is transformed into a feeling of genuine accomplishment when they see their debugged programs run much faster than "ordinary" C programs. For this reason, programming assignments should play a central role in the course.

Fortunately, parallel computers are more accessible than ever. If a commercial parallel computer is not available, it is a straightforward task to build a small cluster out of a few PCs, networking equipment, and free software.

Figure P.1 illustrates the precedence relations among the chapters. A solid arrow from A to B indicates chapter B depends heavily upon material presented in chapter A. A dashed arrow from A to B indicates a weak dependence. If you cover the chapters in numerical order, you will satisfy all of these precedences. However, if you would like your students to start programming in C with MPI as quickly as possible, you may wish to skip Chapter 2 or only cover one or two sections of it. If you wish to focus on numerical algorithms, you may wish to skip Chapter 5 and introduce students to the function `MPI_Bcast` in another way. If you would like to start by having your students programming Monte Carlo algorithms, you can jump to Chapter 10 immediately after Chapter 4. If you want to cover OpenMP before MPI, you can jump to Chapter 17 after Chapter 3.

I thank everyone at McGraw-Hill who helped me create this book, especially Betsy Jones, Michelle Flomenhoft, and Kay Brimeyer. Thank you for your sponsorship, encouragement, and assistance. I also appreciate the help provided by Maggie Murphy and the rest of the composers at Interactive Composition Corporation.

CONTENTS

Preface xiv

CHAPTER 1

Motivation and History 1

- 1.1 Introduction 1
- 1.2 Modern Scientific Method 3
- 1.3 Evolution of Supercomputing 4
- 1.4 Modern Parallel Computers 5
 - 1.4.1 *The Cosmic Cube* 6
 - 1.4.2 *Commercial Parallel Computers* 6
 - 1.4.3 *Beowulf* 7
 - 1.4.4 *Advanced Strategic Computing Initiative* 8
- 1.5 Seeking Concurrency 9
 - 1.5.1 *Data Dependence Graphs* 9
 - 1.5.2 *Data Parallelism* 10
 - 1.5.3 *Functional Parallelism* 10
 - 1.5.4 *Pipelining* 12
 - 1.5.5 *Size Considerations* 13
- 1.6 Data Clustering 14
- 1.7 Programming Parallel Computers 17
 - 1.7.1 *Extend a Compiler* 17
 - 1.7.2 *Extend a Sequential Programming Language* 18
 - 1.7.3 *Add a Parallel Programming Layer* 19
 - 1.7.4 *Create a Parallel Language* 19
 - 1.7.5 *Current Status* 21
- 1.8 Summary 21
- 1.9 Key Terms 22
- 1.10 Bibliographic Notes 22
- 1.11 Exercises 23

CHAPTER 2

Parallel Architectures 27

- 2.1 Introduction 27
- 2.2 Interconnection Networks 28
 - 2.2.1 *Shared versus Switched Media* 28
 - 2.2.2 *Switch Network Topologies* 29
 - 2.2.3 *2-D Mesh Network* 29
 - 2.2.4 *Binary Tree Network* 30
 - 2.2.5 *Hypertree Network* 31
 - 2.2.6 *Butterfly Network* 32
 - 2.2.7 *Hypercube Network* 33
 - 2.2.8 *Shuffle-exchange Network* 35
 - 2.2.9 *Summary* 36
- 2.3 Processor Arrays 37
 - 2.3.1 *Architecture and Data-parallel Operations* 37
 - 2.3.2 *Processor Array Performance* 39
 - 2.3.3 *Processor Interconnection Network* 40
 - 2.3.4 *Enabling and Disabling Processors* 40
 - 2.3.5 *Additional Architectural Features* 42
 - 2.3.6 *Shortcomings of Processor Arrays* 42
- 2.4 Multiprocessors 43
 - 2.4.1 *Centralized Multiprocessors* 43
 - 2.4.2 *Distributed Multiprocessors* 45
- 2.5 Multicomputers 49
 - 2.5.1 *Asymmetrical Multicomputers* 49
 - 2.5.2 *Symmetrical Multicomputers* 51
 - 2.5.3 *Which Model Is Best for a Commodity Cluster?* 52
 - 2.5.4 *Differences between Clusters and Networks of Workstations* 53
- 2.6 Flynn's Taxonomy 54
 - 2.6.1 *SISD* 54
 - 2.6.2 *SIMD* 55

2.6.3	MISD	55
2.6.4	MIMD	56
2.7	Summary	58
2.8	Key Terms	59
2.9	Bibliographic Notes	59
2.10	Exercises	60
CHAPTER 3		
Parallel Algorithm Design 63		
3.1	Introduction	63
3.2	The Task/Channel Model	63
3.3	Foster's Design Methodology	64
3.3.1	Partitioning	65
3.3.2	Communication	67
3.3.3	Agglomeration	68
3.3.4	Mapping	70
3.4	Boundary Value Problem	73
3.4.1	Introduction	73
3.4.2	Partitioning	75
3.4.3	Communication	75
3.4.4	Agglomeration and Mapping	76
3.4.5	Analysis	76
3.5	Finding the Maximum	77
3.5.1	Introduction	77
3.5.2	Partitioning	77
3.5.3	Communication	77
3.5.4	Agglomeration and Mapping	81
3.5.5	Analysis	82
3.6	The n -Body Problem	82
3.6.1	Introduction	82
3.6.2	Partitioning	83
3.6.3	Communication	83
3.6.4	Agglomeration and Mapping	85
3.6.5	Analysis	85
3.7	Adding Data Input	86
3.7.1	Introduction	86
3.7.2	Communication	87
3.7.3	Analysis	88
3.8	Summary	89

3.9	Key Terms	90
3.10	Bibliographic Notes	90
3.11	Exercises	90

CHAPTER 4**Message-Passing Programming 93**

4.1	Introduction	93
4.2	The Message-Passing Model	94
4.3	The Message-Passing Interface	95
4.4	Circuit Satisfiability	96
4.4.1	Function <i>MPI_Init</i>	99
4.4.2	Functions <i>MPI_Comm_rank</i> and <i>MPI_Comm_size</i>	99
4.4.3	Function <i>MPI_Finalize</i>	101
4.4.4	Compiling MPI Programs	102
4.4.5	Running MPI Programs	102
4.5	Introducing Collective Communication	104
4.5.1	Function <i>MPI_Reduce</i>	105
4.6	Benchmarking Parallel Performance	108
4.6.1	Functions <i>MPI_Wtime</i> and <i>MPI_Wtick</i>	108
4.6.2	Function <i>MPI_Barrier</i>	108
4.7	Summary	110
4.8	Key Terms	110
4.9	Bibliographic Notes	110
4.10	Exercises	111

CHAPTER 5**The Sieve of Eratosthenes 115**

5.1	Introduction	115
5.2	Sequential Algorithm	115
5.3	Sources of Parallelism	117
5.4	Data Decomposition Options	117
5.4.1	Interleaved Data Decomposition	118
5.4.2	Block Data Decomposition	118
5.4.3	Block Decomposition Macros	120
5.4.4	Local Index versus Global Index	120

5.4.5 Ramifications of Block Decomposition 121

- 5.5 Developing the Parallel Algorithm 121
 - 5.5.1 Function `MPI_Bcast` 122
- 5.6 Analysis of Parallel Sieve Algorithm 122
- 5.7 Documenting the Parallel Program 123
- 5.8 Benchmarking 128
- 5.9 Improvements 129
 - 5.9.1 Delete Even Integers 129
 - 5.9.2 Eliminate Broadcast 130
 - 5.9.3 Reorganize Loops 131
 - 5.9.4 Benchmarking 131
- 5.10 Summary 133
- 5.11 Key Terms 134
- 5.12 Bibliographic Notes 134
- 5.13 Exercises 134

CHAPTER 6

Floyd's Algorithm 137

- 6.1 Introduction 137
- 6.2 The All-Pairs Shortest-Path Problem 137
- 6.3 Creating Arrays at Run Time 139
- 6.4 Designing the Parallel Algorithm 140
 - 6.4.1 Partitioning 140
 - 6.4.2 Communication 141
 - 6.4.3 Agglomeration and Mapping 142
 - 6.4.4 Matrix Input/Output 143
- 6.5 Point-to-Point Communication 145
 - 6.5.1 Function `MPI_Send` 146
 - 6.5.2 Function `MPI_Recv` 147
 - 6.5.3 Deadlock 148
- 6.6 Documenting the Parallel Program 149
- 6.7 Analysis and Benchmarking 151
- 6.8 Summary 154
- 6.9 Key Terms 154
- 6.10 Bibliographic Notes 154
- 6.11 Exercises 154

CHAPTER 7

Performance Analysis 159

- 7.1 Introduction 159
- 7.2 Speedup and Efficiency 159
- 7.3 Amdahl's Law 161
 - 7.3.1 Limitations of Amdahl's Law 164
 - 7.3.2 The Amdahl Effect 164
- 7.4 Gustafson-Barsis's Law 164
- 7.5 The Karp-Flatt Metric 167
- 7.6 The Isoefficiency Metric 170
- 7.7 Summary 174
- 7.8 Key Terms 175
- 7.9 Bibliographic Notes 175
- 7.10 Exercises 176

CHAPTER 8

Matrix-Vector Multiplication 178

- 8.1 Introduction 178
- 8.2 Sequential Algorithm 179
- 8.3 Data Decomposition Options 180
- 8.4 Rowwise Block-Striped Decomposition 181
 - 8.4.1 Design and Analysis 181
 - 8.4.2 Replicating a Block-Mapped Vector 183
 - 8.4.3 Function `MPI_Allgatherv` 184
 - 8.4.4 Replicated Vector Input/Output 186
 - 8.4.5 Documenting the Parallel Program 187
 - 8.4.6 Benchmarking 187
- 8.5 Columnwise Block-Striped Decomposition 189
 - 8.5.1 Design and Analysis 189
 - 8.5.2 Reading a Columnwise Block-Striped Matrix 191
 - 8.5.3 Function `MPI_Scatterv` 191
 - 8.5.4 Printing a Columnwise Block-Striped Matrix 193
 - 8.5.5 Function `MPI_Gatherv` 193
 - 8.5.6 Distributing Partial Results 195

8.5.7	<i>Function MPI_Alltoallv</i>	195
8.5.8	<i>Documenting the Parallel Program</i>	196
8.5.9	<i>Benchmarking</i>	198
8.6	<i>Checkerboard Block Decomposition</i>	199
8.6.1	<i>Design and Analysis</i>	199
8.6.2	<i>Creating a Communicator</i>	202
8.6.3	<i>Function MPI_Dims_create</i>	203
8.6.4	<i>Function MPI_Cart_create</i>	204
8.6.5	<i>Reading a Checkerboard Matrix</i>	205
8.6.6	<i>Function MPI_Cart_rank</i>	205
8.6.7	<i>Function MPI_Cart_coords</i>	207
8.6.8	<i>Function MPI_Comm_split</i>	207
8.6.9	<i>Benchmarking</i>	208
8.7	<i>Summary</i>	210
8.8	<i>Key Terms</i>	211
8.9	<i>Bibliographic Notes</i>	211
8.10	<i>Exercises</i>	211

CHAPTER 9

Document Classification 216

9.1	<i>Introduction</i>	216
9.2	<i>Parallel Algorithm Design</i>	217
9.2.1	<i>Partitioning and Communication</i>	217
9.2.2	<i>Agglomeration and Mapping</i>	217
9.2.3	<i>Manager/Worker Paradigm</i>	218
9.2.4	<i>Manager Process</i>	219
9.2.5	<i>Function MPI_Abort</i>	220
9.2.6	<i>Worker Process</i>	221
9.2.7	<i>Creating a Workers-only Communicator</i>	223
9.3	<i>Nonblocking Communications</i>	223
9.3.1	<i>Manager's Communication</i>	224
9.3.2	<i>Function MPI_Irecv</i>	224
9.3.3	<i>Function MPI_Wait</i>	225
9.3.4	<i>Workers' Communications</i>	225
9.3.5	<i>Function MPI_Isend</i>	225
9.3.6	<i>Function MPI_Probe</i>	225
9.3.7	<i>Function MPI_Get_count</i>	226
9.4	<i>Documenting the Parallel Program</i>	226
9.5	<i>Enhancements</i>	232

9.5.1	<i>Assigning Groups of Documents</i>	232
9.5.2	<i>Pipelining</i>	232
9.5.3	<i>Function MPI_Testsome</i>	234

9.6	<i>Summary</i>	235
9.7	<i>Key Terms</i>	236
9.8	<i>Bibliographic Notes</i>	236
9.9	<i>Exercises</i>	236

CHAPTER 10

Monte Carlo Methods 239

10.1	<i>Introduction</i>	239
10.1.1	<i>Why Monte Carlo Works</i>	240
10.1.2	<i>Monte Carlo and Parallel Computing</i>	243
10.2	<i>Sequential Random Number Generators</i>	243
10.2.1	<i>Linear Congruential</i>	244
10.2.2	<i>Lagged Fibonacci</i>	245
10.3	<i>Parallel Random Number Generators</i>	245
10.3.1	<i>Manager-Worker Method</i>	246
10.3.2	<i>Leapfrog Method</i>	246
10.3.3	<i>Sequence Splitting</i>	247
10.3.4	<i>Parameterization</i>	248
10.4	<i>Other Random Number Distributions</i>	248
10.4.1	<i>Inverse Cumulative Distribution Function Transformation</i>	249
10.4.2	<i>Box-Muller Transformation</i>	250
10.4.3	<i>The Rejection Method</i>	251
10.5	<i>Case Studies</i>	253
10.5.1	<i>Neutron Transport</i>	253
10.5.2	<i>Temperature at a Point Inside a 2-D Plate</i>	255
10.5.3	<i>Two-Dimensional Ising Model</i>	257
10.5.4	<i>Room Assignment Problem</i>	259
10.5.5	<i>Parking Garage</i>	262
10.5.6	<i>Traffic Circle</i>	264
10.6	<i>Summary</i>	268
10.7	<i>Key Terms</i>	269
10.8	<i>Bibliographic Notes</i>	269
10.9	<i>Exercises</i>	270

CHAPTER 11

Matrix Multiplication 273

- 11.1 Introduction 273
- 11.2 Sequential Matrix Multiplication 274
 - 11.2.1 *Iterative, Row-Oriented Algorithm* 274
 - 11.2.2 *Recursive, Block-Oriented Algorithm* 275
- 11.3 Rowwise Block-Striped Parallel Algorithm 277
 - 11.3.1 *Identifying Primitive Tasks* 277
 - 11.3.2 *Agglomeration* 278
 - 11.3.3 *Communication and Further Agglomeration* 279
 - 11.3.4 *Analysis* 279
- 11.4 Cannon's Algorithm 281
 - 11.4.1 *Agglomeration* 281
 - 11.4.2 *Communication* 283
 - 11.4.3 *Analysis* 284
- 11.5 Summary 286
- 11.6 Key Terms 287
- 11.7 Bibliographic Notes 287
- 11.8 Exercises 287

CHAPTER 12

Solving Linear Systems 290

- 12.1 Introduction 290
- 12.2 Terminology 291
- 12.3 Back Substitution 292
 - 12.3.1 *Sequential Algorithm* 292
 - 12.3.2 *Row-Oriented Parallel Algorithm* 293
 - 12.3.3 *Column-Oriented Parallel Algorithm* 295
 - 12.3.4 *Comparison* 295
- 12.4 Gaussian Elimination 296
 - 12.4.1 *Sequential Algorithm* 296
 - 12.4.2 *Parallel Algorithms* 298
 - 12.4.3 *Row-Oriented Algorithm* 299
 - 12.4.4 *Column-Oriented Algorithm* 303

12.4.5 *Comparison* 30312.4.6 *Pipelined, Row-Oriented Algorithm* 304

- 12.5 Iterative Methods 306
- 12.6 The Conjugate Gradient Method 309
 - 12.6.1 *Sequential Algorithm* 309
 - 12.6.2 *Parallel Implementation* 310
- 12.7 Summary 313
- 12.8 Key Terms 314
- 12.9 Bibliographic Notes 314
- 12.10 Exercises 314

CHAPTER 13

Finite Difference Methods 318

- 13.1 Introduction 318
- 13.2 Partial Differential Equations 320
 - 13.2.1 *Categorizing PDEs* 320
 - 13.2.2 *Difference Quotients* 321
- 13.3 Vibrating String 322
 - 13.3.1 *Deriving Equations* 322
 - 13.3.2 *Deriving the Sequential Program* 323
 - 13.3.3 *Parallel Program Design* 324
 - 13.3.4 *Isoefficiency Analysis* 327
 - 13.3.5 *Replicating Computations* 327
- 13.4 Steady-State Heat Distribution 329
 - 13.4.1 *Deriving Equations* 329
 - 13.4.2 *Deriving the Sequential Program* 330
 - 13.4.3 *Parallel Program Design* 332
 - 13.4.4 *Isoefficiency Analysis* 332
 - 13.4.5 *Implementation Details* 334
- 13.5 Summary 334
- 13.6 Key Terms 335
- 13.7 Bibliographic Notes 335
- 13.8 Exercises 335

CHAPTER 14

Sorting 338

- 14.1 Introduction 338
- 14.2 Quicksort 339

- 14.3 A Parallel Quicksort Algorithm 340
 - 14.3.1 Definition of Sorted 340
 - 14.3.2 Algorithm Development 341
 - 14.3.3 Analysis 341
- 14.4 Hyperquicksort 343
 - 14.4.1 Algorithm Description 343
 - 14.4.2 Isoefficiency Analysis 345
- 14.5 Parallel Sorting by Regular Sampling 346
 - 14.5.1 Algorithm Description 346
 - 14.5.2 Isoefficiency Analysis 347
- 14.6 Summary 349
- 14.7 Key Terms 349
- 14.8 Bibliographic Notes 350
- 14.9 Exercises 350

CHAPTER 15

The Fast Fourier Transform 353

- 15.1 Introduction 353
- 15.2 Fourier Analysis 353
- 15.3 The Discrete Fourier Transform 355
 - 15.3.1 Inverse Discrete Fourier Transform 357
 - 15.3.2 Sample Application: Polynomial Multiplication 357
- 15.4 The Fast Fourier Transform 360
- 15.5 Parallel Program Design 363
 - 15.5.1 Partitioning and Communication 363
 - 15.5.2 Agglomeration and Mapping 365
 - 15.5.3 Isoefficiency Analysis 365
- 15.6 Summary 367
- 15.7 Key Terms 367
- 15.8 Bibliographic Notes 367
- 15.9 Exercises 367

CHAPTER 16

Combinatorial Search 369

- 16.1 Introduction 369
- 16.2 Divide and Conquer 370

- 16.3 Backtrack Search 371
 - 16.3.1 Example 371
 - 16.3.2 Time and Space Complexity 374
- 16.4 Parallel Backtrack Search 374
- 16.5 Distributed Termination Detection 377
- 16.6 Branch and Bound 380
 - 16.6.1 Example 380
 - 16.6.2 Sequential Algorithm 382
 - 16.6.3 Analysis 385
- 16.7 Parallel Branch and Bound 385
 - 16.7.1 Storing and Sharing Unexamined Subproblems 386
 - 16.7.2 Efficiency 387
 - 16.7.3 Halting Conditions 387
- 16.8 Searching Game Trees 388
 - 16.8.1 Minimax Algorithm 388
 - 16.8.2 Alpha-Beta Pruning 392
 - 16.8.3 Enhancements to Alpha-Beta Pruning 395
- 16.9 Parallel Alpha-Beta Search 395
 - 16.9.1 Parallel Aspiration Search 396
 - 16.9.2 Parallel Subtree Evaluation 396
 - 16.9.3 Distributed Tree Search 397
- 16.10 Summary 399
- 16.11 Key Terms 400
- 16.12 Bibliographic Notes 400
- 16.13 Exercises 401

CHAPTER 17

Shared-Memory Programming 404

- 17.1 Introduction 404
- 17.2 The Shared-Memory Model 405
- 17.3 Parallel for Loops 407
 - 17.3.1 parallel for Pragma 408
 - 17.3.2 Function omp_get_num_procs 410
 - 17.3.3 Function omp_set_num_threads 410
- 17.4 Declaring Private Variables 410
 - 17.4.1 private Clause 411

17.4.2	<i>firstprivate Clause</i>	412
17.4.3	<i>lastprivate Clause</i>	412
17.5	Critical Sections	413
17.5.1	<i>critical Pragma</i>	415
17.6	Reductions	415
17.7	Performance Improvements	417
17.7.1	<i>Inverting Loops</i>	417
17.7.2	<i>Conditionally Executing Loops</i>	418
17.7.3	<i>Scheduling Loops</i>	419
17.8	More General Data Parallelism	421
17.8.1	<i>parallel Pragma</i>	422
17.8.2	<i>Function omp_get_thread_num</i>	423
17.8.3	<i>Function omp_get_num_threads</i>	425
17.8.4	<i>for Pragma</i>	425
17.8.5	<i>single Pragma</i>	427
17.8.6	<i>nowait Clause</i>	427
17.9	Functional Parallelism	428
17.9.1	<i>parallel sections Pragma</i>	429
17.9.2	<i>section Pragma</i>	429
17.9.3	<i>sections Pragma</i>	429
17.10	Summary	430
17.11	Key Terms	432
17.12	Bibliographic Notes	432
17.13	Exercises	433

CHAPTER 18 Combining MPI and OpenMP 436

18.1	Introduction	436
18.2	Conjugate Gradient Method	438
18.2.1	<i>MPI Program</i>	438
18.2.2	<i>Functional Profiling</i>	442
18.2.3	<i>Parallelizing Function matrix_vector_product</i>	442
18.2.4	<i>Benchmarking</i>	443
18.3	Jacobi Method	444
18.3.1	<i>Profiling MPI Program</i>	444

18.3.2	<i>Parallelizing Function find_steady_state</i>	444
18.3.3	<i>Benchmarking</i>	446

18.4	Summary	448
18.5	Exercises	448

APPENDIX A MPI Functions 450

APPENDIX B Utility Functions 485

B.1	Header File <i>MyMPI.h</i>	485
B.2	Source File <i>MyMPI.c</i>	486

APPENDIX C Debugging MPI Programs 505

C.1	Introduction	505
C.2	Typical Bugs in MPI Programs	505
C.2.1	<i>Bugs Resulting in Deadlock</i>	505
C.2.2	<i>Bugs Resulting in Incorrect Results</i>	506
C.2.3	<i>Advantages of Collective Communications</i>	507
C.3	Practical Debugging Strategies	507

APPENDIX D Review of Complex Numbers 509

APPENDIX E OpenMP Functions 513

Bibliography 515