教 育 部 高 等 教 育 司 推 荐
国外优秀信息科学与技术系列教学用书

并行程序设计
——技术与应用

（影印版）

PARALLEL PROGRAMMING
Techniques and Applications Using Networked Workstations and Parallel Computers

■ Barry Wilkinson
Michael Allen

高等教育出版社
Higher Education Press
Prentice Hall
Pearson Education
出 版 集 团

教育部高等教育司推荐
国外优秀信息科学与技术系列教学用书

并行程序设计

——技术与应用

（影印版）

PARALLEL PROGRAMMING

TECHNIQUES AND APPLICATIONS
Using Networked Workstations and Parallel Computers

Barry Wilkinson
Michael Allen

高等教育出版社
Pearson Education

图字：01-2001-5058 号

Parallel Programming: Techniques and Applications Using Networked Workstations and Parallel Computers, First Edition. Barry Wilkinson，Michael Allen

图书在版编目(CIP)数据

并行程序设计：技术与应用/(美)威尔金森(Wilkinson,B.)等著.—影印本.—北京:高等教育出版社,2002.5

ISBN 7-04-011162-4

Ⅰ.并… Ⅱ.威… 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(2002)第 031038 号

并行程序设计——技术与应用 (影印版)

Barry Wilkinson，Michael Allen

出版发行	高等教育出版社	购书热线	010-64054588
社　　址	北京市东城区沙滩后街 55 号	免费咨询	800-810-0598
邮政编码	100009	网　　址	http://www.hep.edu.cn
传　　真	010-64014048		http://www.hep.com.cn
经　　销	新华书店北京发行所		
印　　刷	北京铭成印刷有限公司		
开　　本	787×1092　1/16	版　　次	2002 年 5 月第 1 版
印　　张	28.25	印　　次	2002 年 5 月第 1 次印刷
字　　数	650 000	定　　价	29.50 元

本书如有缺页、倒页、脱页等质量问题，请到所购图书销售部门联系调换。

前　言

20 世纪末，以计算机和通信技术为代表的信息科学和技术对世界经济、科技、军事、教育和文化等产生了深刻影响。信息科学技术的迅速普及和应用，带动了世界范围信息产业的蓬勃发展，为许多国家带来了丰厚的回报。

进入 21 世纪，尤其随着我国加入 WTO，信息产业的国际竞争将更加激烈。我国信息产业虽然在 20 世纪末取得了迅猛发展，但与发达国家相比，甚至与印度、爱尔兰等国家相比，还有很大差距。国家信息化的发展速度和信息产业的国际竞争能力，最终都将取决于信息科学技术人才的质量和数量。引进国外信息科学和技术优秀教材，在有条件的学校推动开展英语授课或双语教学，是教育部为加快培养大批高质量的信息技术人才采取的一项重要举措。

为此，教育部要求由高等教育出版社首先开展信息科学和技术教材的引进试点工作。同时提出了两点要求，一是要高水平，二是要低价格。在高等教育出版社和信息科学技术引进教材专家组的努力下，经过比较短的时间，第一批引进的 20 多种教材已经陆续出版。这套教材出版后受到了广泛的好评，其中有不少是世界信息科学技术领域著名专家、教授的经典之作和反映信息科学技术最新进展的优秀作品，代表了目前世界信息科学技术教育的一流水平，而且价格也是最优惠的，与国内同类自编教材相当。

这项教材引进工作是在教育部高等教育司和高教社的共同组织下，由国内信息科学技术领域的专家、教授广泛参与，在对大量国外教材进行多次遴选的基础上，参考了国内和国外著名大学相关专业的课程设置进行系统引进的。其中，John Wiley 公司出版的贝尔实验室信息科学研究中心副总裁 Silberschatz 教授的经典著作《操作系统概念》，是我们经过反复谈判，做了很多努力才得以引进的。William Stallings 先生曾编写了在美国深受欢迎的信息科学技术系列教材，其中有多种教材获得过美国教材和学术著作者协会颁发的计算机科学与工程教材奖，这批引进教材中就有他的两本著作。留美中国学者 Jiawei Han 先生的《数据挖掘》是该领域中具有里程碑意义的著作。由达特茅斯学院 Thomas Cormen 和麻省理工学院、哥伦比亚大学的几

位学者共同编著的经典著作《算法导论》，在经历了11年的锤炼之后于2001年出版了第二版。目前任教于美国Massachusetts大学的James Kurose教授，曾在美国三所高校先后10次获得杰出教师或杰出教学奖，由他主编的《计算机网络》出版后，以其体系新颖、内容先进而倍受欢迎。在努力降低引进教材售价方面，高等教育出版社做了大量和细致的工作。这套引进的教材体现了权威性、系统性、先进性和经济性等特点。

教育部也希望国内和国外的出版商积极参与此项工作，共同促进中国信息技术教育和信息产业的发展。我们在与外商的谈判工作中，不仅要坚定不移地引进国外最优秀的教材，而且还要千方百计地将版权转让费降下来，要让引进教材的价格与国内自编教材相当，让广大教师和学生负担得起。中国的教育市场巨大，外国出版公司和国内出版社要通过扩大发行数量取得效益。

在引进教材的同时，我们还应做好消化吸收，注意学习国外先进的教学思想和教学方法，提高自编教材的水平，使我们的教学和教材在内容体系上，在理论与实践的结合上，在培养学生的动手能力上能有较大的突破和创新。

目前，教育部正在全国35所高校推动示范性软件学院的建设和实施，这也是加快培养信息科学技术人才的重要举措之一。示范性软件学院要立足于培养具有国际竞争力的实用性软件人才，与国外知名高校或著名企业合作办学，以国内外著名IT企业为实践教学基地，聘请国内外知名教授和软件专家授课，还要率先使用引进教材开展教学。

我们希望通过这些举措，能在较短的时间，为我国培养一大批高质量的信息技术人才，提高我国软件人才的国际竞争力，促进我国信息产业的快速发展，加快推动国家信息化进程，进而带动整个国民经济的跨越式发展。

教育部高等教育司

二〇〇二年三月

To my wife, Wendy
and my daughter, Johanna

Barry Wilkinson

To my wife, Bonnie

Michael Allen

Preface

The purpose of this text is to introduce parallel programming techniques. Parallel programming uses multiple computers, or computers with multiple internal processors, to solve a problem at a greater computational speed than using a single computer. It also offers the opportunity to tackle larger problems; that is, problems with more computational steps or more memory requirements, the latter because multiple computers and multiprocessor systems often have more total memory than a single computer. In this text, we concentrate upon the use of multiple computers that communicate between themselves by sending messages; hence the term *message-passing* parallel programming. The computers we use can be different types (PC, SUN, SGI, etc.) but must be interconnected by a network, and a software environment must be present for intercomputer message passing. Suitable networked computers are very widely available as the basic computing platform for students so that acquisition of specially designed multiprocessor systems can usually be avoided. Several software tools are available for message-passing parallel programming, including PVM and several implementations of MPI, which are all freely available. Such software can also be used on specially designed multiprocessor systems should these systems be available for use. So far as practicable, we discuss techniques and applications in a system-independent fashion.

The text is divided into two parts, Part I and Part II. In Part I, the basic techniques of parallel programming are developed. The chapters of Part I cover all the essential aspects, using simple problems to demonstrate techniques. The techniques themselves, however, can be applied to a wide range of problems. Sample code is given usually first as sequential code and then as realistic parallel *pseudocode*. Often, the underlying algorithm is already parallel in nature and the sequential version has "unnaturally" serialized it using loops. Of course, some algorithms have to be reformulated for efficient parallel solution, and this reformulation may not be immediately apparent. One chapter in Part I introduces a type of parallel programming not centered around message-passing multicomputers, but around specially designed *shared memory* multiprocessor systems. This chapter describes the use of Pthreads, an IEEE multiprocessor standard system that is widely available and can be used on a single computer.

The prerequisites for studying Part I are knowledge of sequential programming, such as from using the C language and associated data structures. Part I can be studied immediately after basic sequential programming has been mastered. Many assignments here can be attempted without specialized mathematical knowledge. If MPI or PVM is used for the assignments, programs are written in C with message-passing library calls. The descriptions of the specific library calls needed are given in the appendices.

Many parallel computing problems have specially developed algorithms, and in Part II problem-specific algorithms are studied in both non-numeric and numeric domains. For Part II, some mathematical concepts are needed such as matrices. Topics covered in Part II include sorting, matrix multiplication, linear equations, partial differential equations, image processing, and searching and optimization. Image processing is particularly suitable for parallelization and is included as an interesting application with significant potential for projects. The fast Fourier transform is discussed in the context of image processing. This important transform is also used in many other areas, including signal processing and voice recognition.

A large selection of "real-life" problems drawn from practical situations is presented at the end of each chapter. These problems require no specialized mathematical knowledge and are a unique aspect of this text. They develop skills in using parallel programming techniques rather than simply learning to solve specific problems such as sorting numbers or multiplying matrices.

Topics in Part I are suitable as additions to normal sequential programming classes. At the University of North Carolina at Charlotte (UNCC), we introduce our freshmen students to parallel programming in this way. In that context, the text is a supplement to a sequential programming course text. The sequential programming language is assumed to be C or C++. Part I and Part II together is suitable as a more advanced undergraduate parallel programming/computing course, and at UNCC we use the text in that manner.

Full details of the UNCC environment and site-specific details can be found at http://www.cs.uncc.edu/par_prog. Included at this site are extensive Web pages to help students learn how to compile and run parallel programs. Sample programs are provided. An Instructor's Manual is also available to instructors. Our work on teaching parallel programming is connected to that done by the Regional Training Center for Parallel Processing at North Carolina State University. Additional information about this center can be found at http://renoir.csc.ncsu.edu/RTCPP.

The text is a direct outcome of a National Science Foundation grant awarded to the authors at the University of North Carolina at Charlotte to introduce parallel programming in the freshman year.[1] It is a great pleasure to acknowledge Dr. M. Mulder, program director at the National Science Foundation, for supporting our project. Without his support, we would not be able to pursue the ideas presented in this text. We also wish to thank the graduate students that worked on this project, J. Alley, M. Antonious, M. Buchanan, and G.

[1]National Science Foundation grant "Introducing parallel programming techniques into the freshman curricula," ref. no. DUE 9554975.

Robins, and undergraduate students G. Feygin, W. Hasty, C. Beauregard, M. Moore, D. Lowery, K. Patel, Johns Cherian, and especially Uday Kamath. This team helped develop the material and assignments with us. We should like to record our thanks to James Robinson, the departmental system administrator who established our local workstation cluster, without which we would not have been able to conduct the work.

We should also like to thank the many students at UNCC who help us refine the material over the last few years, especially the "teleclasses," in which the materials were classroom tested in a unique setting. These teleclasses are broadcast to several North Carolina universities, including UNC-Asheville, UNC-Greensboro, UNC-Wilmington, and North Carolina State University, in addition to UNCC. We owe a debt of gratitude to many people, among which Professor Wayne Lang at UNC-Asheville and Professor Mladen Vouk of NC State University deserve special mention. Professor Lang truly contributed to the course development in the classroom and Professor Vouk, apart from presenting an expert guest lecture for us, set up an impressive Web page that included "real audio" of our lectures and "automatically turning" slides. (These lectures can be viewed at http://renoir.csc.ncsu.edu/CSC495A.) Professor John Board of Duke University and Professor Jan Prins of UNC Chapel Hill also kindly made expert guest presentations to classes. A parallel programming course based upon the material in this text was also given at the Universidad Nacional de San Luis in Argentina by kind invitation from Professor Raul Gallard — all these activities helped us in developing this text.

We would like to express our appreciation to Alan Apt and Laura Steele of Prentice Hall, who received our proposal for a textbook and supported us throughout its development. Reviewers provided us with very helpful advice.

Finally, may we ask that you please send comments and corrections to us at abw@uncc.edu (Barry Wilkinson) or cma@uncc.edu (Michael Allen).

Barry Wilkinson
Michael Allen
University of North Carolina
Charlotte

About the Authors

Barry Wilkinson is a Full Professor in the Department of Computer Science at the University of North Carolina at Charlotte. He has previously held faculty positions at Brighton Polytechnic, England (1984–87), State University of New York, College at New Paltz (1983–84), University College, Cardiff, Wales (1976–83), and the University of Aston, England (1973–76). From 1969 to 1970, he worked on process control computer systems at Ferranti Ltd. He is the author of *Computer Peripherals* (with D. Horrocks, Hodder and Stoughton, 1980, 2nd edition, 1987), *Digital System Design* (Prentice Hall, 1987, 2nd edition, 1992), *Computer Architecture Design and Performance* (Prentice Hall 1991, 2nd edition, 1996), and *The Essence of Digital Design* (Prentice Hall, 1997). In addition to these books, he has published many papers in major computer journals. He received a B.Sc. degree in Electrical Engineering (with first-class honors) from the University of Salford in 1969, and M.S. and Ph.D. degrees from the University of Manchester (Department of Computer Science), England in 1971 and 1974, respectively. He has been a senior member of the IEEE since 1983.

Michael Allen is a Full Professor in the Department of Computer Science at the University of North Carolina at Charlotte. He has previously held faculty positions as an Associate and Full Professor in the Electrical Engineering Department at the University of North Carolina at Charlotte (1974–85), and as an Instructor and an Assistant Professor in the Electrical Engineering Department at the State University of New York at Buffalo (1968–74). From 1985 to 1987, he was on leave from the University of North Carolina at Charlotte while serving as the President and Chairman of DataSpan, Inc. Additional industry experience includes electronics design and software systems development for Eastman Kodak, Sylvania Electronics, Bell of Pennsylvania, Wachovia Bank, and numerous other firms. He received B.S. and M.S. degrees in Electrical Engineering from Carnegie Mellon University in 1964 and 1965, respectively, and a Ph.D. from the State University of New York at Buffalo in 1968.

Contents