

The Sourcebook Of Parallel Computing The Morgan Kaufmann Series In Computer Architecture And Design

This comprehensive tutorial for building 3D virtual worlds on the Internet using VRML is now updated to reflect the significant upgrade to VRML 2.0. The second edition includes extensive coverage of Behaviors and new introductions to Perl, Java and JavaScript. The VRML 2.0 Sourcebook tells where to get VRML browsers and how to make them work with your current Web browser. The CD contains code for all VRML worlds.

Euro-Par 2005 was the eleventh conference in the Euro-Par series. It was organized by the Centre for Informatics and Information Technology (CITI) and the Department of Informatics of the Faculty of Science and Technology of Universidade Nova de Lisboa, at the Campus of Monte de Caparica.

Parallel processing has been an enabling technology in scientific computing for more than 20 years. This book is the first in-depth discussion of parallel computing in 10 years; it reflects the mix of topics that mathematicians, computer scientists, and computational scientists focus on to make parallel processing effective for scientific problems.

Presently, the impact of parallel processing on scientific computing varies greatly across disciplines, but it plays a vital role in most problem domains and is absolutely essential in many of them. Parallel Processing for Scientific Computing is divided into four parts: The first concerns performance modeling, analysis, and optimization; the second focuses on parallel algorithms and software for an array of problems common to many modeling and simulation applications; the third emphasizes tools and environments that can ease and enhance the process of application development; and the fourth provides a sampling of applications that require parallel computing for scaling to solve larger and realistic models that can advance science and engineering.

The Parallel Programming Guide for Every Software Developer From grids and clusters to next-generation game consoles, parallel computing is going mainstream. Innovations such as Hyper-Threading Technology, HyperTransport Technology, and multicore microprocessors from IBM, Intel, and Sun are accelerating the movement's growth. Only one thing is missing: programmers with the skills to meet the soaring demand for parallel software. That's where Patterns for Parallel Programming comes in. It's the first parallel programming guide written specifically to serve working software developers, not just computer scientists. The authors introduce a complete, highly accessible pattern language that will help any experienced developer "think parallel"-and start writing effective parallel code almost immediately. Instead of formal theory, they deliver proven solutions to the challenges faced by parallel programmers, and pragmatic guidance for using today's parallel APIs in the real world. Coverage includes: Understanding the parallel computing landscape and the challenges faced by parallel developers Finding the concurrency in a software design problem and decomposing it into concurrent tasks Managing the use of data across tasks Creating an algorithm structure that effectively exploits the concurrency you've identified Connecting your algorithmic structures to the APIs needed to implement them Specific software constructs for implementing parallel programs Working with today's leading parallel programming environments: OpenMP, MPI, and Java Patterns have helped thousands of programmers master object-oriented development and other complex programming technologies. With this book, you will learn that they're the best way to master parallel programming too.

Scientific Computing

Parallel I/O for High Performance Computing

From Parallel Processing to the Internet of Things

Living Modern

Going Beyond Words

Patterns for Parallel Programming

Level of Detail for 3D Graphics

Compelling and practical view of computer security in a multinational environment – for everyone who does business in more than one country.

In the early days of computer networking IBM mainframes could only connect to other IBM mainframes, Burroughs only to other Burroughs, etc. Beginning in 1967 the US Defense Department's Advanced Research Projects Agency (ARPA) office sponsored development of a "heterogeneous" network compatible with computers from any manufacturer. That R&D effort, one of the most successful in history, resulted in the on-time, on-budget construction of the revolutionary ARPANET, the immediate predecessor of today's Internet. The ARPANET Sourcebook: The Unpublished Foundations of the Internet reproduces the seminal papers, reports, and RFCs that led to the birth of modern network computing. Most appear here in book form for the first time. Part A, Imagining the ARPANET, covers the initial studies of network feasibility and includes: the introductory and concluding chapters of Paul Baran's seminal but little-known RAND research report On Distributed Communications in which packet switching was first conceptualized, the classic 1968 paper The Computer as a Communication Device by J.C.R. Licklider and Robert Taylor, respectively the ARPANET's earliest proponent and the ARPA administrator who pushed the development project. Part B, Planning the ARPANET includes: scans of the earliest RFCs ("Requests for Comments"), some publicly available here for the first time. RFCs were in effect the design documents for the ARPANET and later the Internet. the 1968 ARPA-commissioned SRI study that modeled a heterogeneous network and concluded that it was indeed feasible, forewords by Steve Crocker (author of RFC #1) and Leonard Kleinrock (noted author and head of the UCLA computing lab that hosted the first ARPANET node). Part C, Building the ARPANET, reproduces the quarterly technical reports from the government's contractor Bolt Beranek and Newman contemporaneously describing the development group's progress, difficulties encountered, and final success. Dave Walden, former BBN VP and a key member of the ARPANET team, has contributed a retrospective Foreword. Other noteworthy material: historical perspectives from Peter Salus, Robert Taylor, Willis Ware, Michael Padlipsky, and Les Earnest, and a long-forgotten RFC which anticipated JAVA by more than 20 years.

Containing over 300 entries in an A-Z format, the Encyclopedia of Parallel Computing provides easy, intuitive access to relevant information for professionals and researchers seeking access to any aspect within the broad field of parallel computing. Topics for this comprehensive reference were selected, written, and peer-reviewed by an international pool of distinguished researchers in the field. The Encyclopedia is broad in scope, covering machine organization, programming languages, algorithms, and applications. Within each area, concepts, designs, and specific implementations are presented. The highly-structured essays in this work comprise synonyms, a definition and discussion of the topic, bibliographies, and links to related literature. Extensive cross-references to other entries within the Encyclopedia support efficient, user-friendly searches for immediate access to useful information. Key concepts presented in the Encyclopedia of Parallel Computing include; laws and metrics; specific numerical and non-numerical algorithms; asynchronous algorithms; libraries of subroutines; benchmark suites; applications; sequential consistency and cache coherency; machine classes such as clusters, shared-memory multiprocessors, special-purpose machines and dataflow machines; specific machines such as Cray supercomputers, IBM's cell processor and Intel's multicore machines; race detection and auto parallelization; parallel programming languages, synchronization primitives, collective operations, message passing libraries, checkpointing, and operating systems. Topics covered: Speedup, Efficiency, Isoefficiency, Redundancy, Amdahl's law, Computer Architecture Concepts, Parallel Machine Designs, Benmarks, Parallel Programming concepts & design, Algorithms, Parallel applications. This authoritative reference will be published in two formats: print and online. The online edition features hyperlinks to cross-references and to additional significant research. Related Subjects: supercomputing, high-performance computing, distributed computing

"I enjoyed reading this book immensely. The author was uncommonly careful in his explanations. I'd recommend this book to anyone writing scientific application codes." -Peter S. Pacheco, University of San Francisco "This text provides a useful overview of an area that is currently not addressed in any book. The presentation of parallel I/O issues across all levels of abstraction is this book's greatest strength." -Alan Sussman, University of Maryland Scientific and technical programmers can no longer afford to treat I/O as an afterthought. The speed, memory size, and disk capacity of parallel computers continue to grow rapidly, but the rate at which disk drives can read and write data is improving far less quickly. As a result, the performance of carefully tuned parallel programs can slow dramatically when they read or write files-and the problem is likely to get far worse. Parallel input and output techniques can help solve this problem by creating multiple data paths between memory and disks. However, simply adding disk drives to an I/O system without considering the overall software design will not significantly improve performance. To reap the full benefits of a parallel I/O system, application programmers must understand how parallel I/O systems work and where the performance pitfalls lie. Parallel I/O for High Performance Computing directly addresses this critical need by examining parallel I/O from the bottom up. This important new book is recommended to anyone writing scientific application codes as the best single source on I/O techniques and to computer scientists as a solid up-to-date introduction to parallel I/O research. Features: An overview of key I/O issues at all levels of abstraction-including hardware, through the OS and file systems, up to very high-level scientific libraries. Describes the important features of MPI-IO, netCDF, and HDF-5 and presents numerous examples illustrating how to use each of these I/O interfaces. Addresses the basic question of how to read and write data efficiently in HPC applications. An explanation of various layers of storage - and techniques for using disks (and sometimes tapes) effectively in HPC applications.

UNIX Administration

Numerical Solution of Partial Differential Equations on Parallel Computers

A Sourcebook for Fast 32-bit Software Development

Parallel Processing for Scientific Computing

Euro-Par 2005 Parallel Processing

Parallel Processing and Applied Mathematics

This book constitutes the refereed joint proceedings of ten international workshops held in conjunction with the 4th International Symposium on Parallel and Distributed Processing and Applications, ISPA 2006, held in Sorrento, Italy in December 2006. It contains 116 papers that contribute to enlarging the spectrum of the more general topics treated in the ISPA 2006 main conference.

Since the dawn of computing, the quest for a better understanding of Nature has been a driving force for technological development. Groundbreaking achievements by great scientists have paved the way from the abacus to the supercomputing power of today. When trying to replicate Nature in the computer's silicon test tube, there is need for precise and computable process descriptions. The scienti?c ?elds of Ma- ematics and Physics provide a powerful vehicle for such descriptions in terms of Partial Differential Equations (PDEs). Formulated as such equations, physical laws can become subject to computational and analytical studies. In the computational setting, the equations can be discreti ed for ef?cient solution on a computer, leading to valuable tools for simulation of natural and man-made processes. Numerical so- tion of PDE-based mathematical models has been an important research topic over centuries, and will remain so for centuries to come. In the context of computer-based simulations, the quality of the computed results is directly connected to the model's complexity and the number of data points used for the computations. Therefore, computational scientists tend to ?ll even the largest and most powerful computers they can get access to, either by increasing the size of the data sets, or by introducing new model terms that make the simulations more realistic, or a combination of both. Today, many important simulation problems can not be solved by one single computer, but calls for parallel computing.

This book constitutes the refereed proceedings of the 14th International Conference on Parallel Computing, Euro-Par 2008, held in Las Palmas de Gran Canaria, Spain, in August 2008. The 86 revised papers presented were carefully reviewed and selected from 264 submissions. The papers are organized in topical sections on support tools and environments; performance prediction and evaluation; scheduling and load balancing; high performance architectures and compilers; parallel and distributed databases; grid and cluster computing; peer-to-peer computing; distributed systems and algorithms; parallel and distributed programming; parallel numerical algorithms; distributed and high-performance multimedia; theory and algorithms for parallel computation; and high performance networks.

This book is the first to introduce a mesoscale polymer simulation system called OCTA. With its name derived from "Open Computational Tool for Advanced material technology," OCTA is a unique software product, available without charge, that was developed in a project funded by Japanese government. OCTA contains a series of simulation programs focused on mesoscale simulation of the soft matter COGNAC, SUSHI, PASTA, NAPLES, MUFFIN, and KAPSEL. When mesoscale polymer simulation is performed, one may encounter many difficulties that this book will help to overcome. The book not only introduces the theoretical background and functions of each simulation engine, it also provides many examples of the practical applications of the OCTA system. Those examples include predicting mechanical properties of plastic and rubber, morphology formation of polymer blends and composites, the micelle structure of surfactants, and optical properties of polymer films. This volume is strongly recommended as a valuable resource for both academic and industrial researchers who work in polymer simulation.

Introduction to Parallel Computing

A Practical Approach to High-Performance Computing

Computer Simulation of Polymeric Materials

Frontiers of High Performance Computing and Networking

A Comprehensive Sourcebook for Effective Systems & Network Management

7th International Conference, PaCT 2003, Novosibirsk, Russia, September 15-19, 2003, Proceedings

7th International Conference, PPAM 2007, Gdansk, Poland, September 9-12, 2007, Revised Selected Papers

Teach yourself accounting while using your software

This book constitutes the thoroughly refereed post-workshop proceedings of the First and the Second International Workshop on OpenMP, IWOMP 2005 and IWOMP 2006, held in Eugene, OR, USA, and in Reims, France, in June 2005 and 2006 respectively. The first part of the book presents 16 revised full papers carefully reviewed and selected from the IWOMP 2005 program and organized in topical sections on performance tools, compiler technology, run-time environment, applications, as well as the OpenMP language and its evaluation. In the second part there are 19 papers of IWOMP 2006, fully revised and grouped thematically in sections on advanced performance tuning aspects of code development applications, and proposed extensions to OpenMP.

This book constitutes the thoroughly refereed post-conference proceedings of the 7th International Conference on Parallel Processing and Applied Mathematics, PPAM 2007, held in Gdansk, Poland, in September 2007. The 63 revised full papers of the main conference presented together with 85 revised workshop papers were carefully reviewed and selected from over 250 initial submissions. The papers are organized in topical sections on parallel/distributed architectures and mobile computing, numerical algorithms and parallel numerics, parallel and distributed non-numerical algorithms, environments and tools for as well as applications of parallel/distributed/grid computing, evolutionary computing, meta-heuristics and neural networks. The volume proceeds with the outcome of 11 workshops and minisymposia dealing with novel data formats and algorithms for dense linear algebra computations, combinatorial tools for parallel sparse matrix computations, grid applications and middleware, large scale computations on grids, models, algorithms and methodologies for grid-enabled computing environments, scheduling for parallel computing, language-based parallel programming models, performance evaluation of parallel applications on large-scale systems, parallel computational biology, high performance computing for engineering applications, and the minisymposium on interval analysis.

A comprehensive guide for students and practitioners to parallel computing models, processes, metrics, and implementation in MPI and OpenMP.

An Introduction to Parallel Programming

Parallel Programming in OpenMP

The Corporate Security Sourcebook for Today's Global Economy

Lunar Sourcebook

Inner Loops

Parallel Computing Technologies

The Foundations of Artificial Intelligence

Mathematics of Computing -- Parallelism.

Software -- Programming Techniques.

An Introduction to Parallel Programming, Second Edition presents a tried-and-true tutorial approach that shows students how to develop effective parallel programs with MPI, Pthreads and OpenMP. As the first undergraduate text to directly address compiling and running parallel programs on multi-core and cluster architecture, this second designing, debugging and evaluating the performance of distributed and shared-memory programs while adding coverage of accelerators via new content on GPU programming and heterogeneous programming. New and improved user-friendly exercises teach students how to compile, run and modify example programs. Takes a tutorial approach building progressively to more challenging examples Explains how to develop parallel programs using MPI, Pthreads and OpenMP programming models A robust package of online ancillaries for instructors and students includes lecture slides, solutions manual, downloadable source code, and an image bank New to this edition: New chapters on programming New examples and exercises related to parallel algorithms

Understanding modern more as lifestyle than as style, it is all about clean lines, elegant colour combinations, maximising indoor-outdoor relationships, artfully collecting and displaying design objects and open areas for lounging, cooking and dining.

VRML 2.0 Sourcebook

The Unpublished Foundations of the Internet

The Sourcebook of Nonverbal Measures

Encyclopedia of Parallel Computing

The ARPANET Sourcebook

19th International Conference, Aachen, Germany, August 26-30, 2013, Proceedings

Distributed and Cloud Computing

This outstanding collection is designed to address the fundamental issues and principles underlying the task of Artificial Intelligence.

With practical advice and a broad sampling of important algorithms, Inner Loops shows how to design programs that extend the edge of the envelope of PC performance. It provides a thorough review of 32-bit code optimization for the 486, Pentium and Pentium Pro, as well as design tips for writing fast 32-bit software.

Distributed and Cloud Computing: From Parallel Processing to the Internet of Things offers complete coverage of modern distributed computing technology including clusters, the grid, service-oriented architecture, massively parallel processors, peer-to-peer networking, and cloud computing. It is the first modern, up-to-date distributed systems textbook; it explains how to create high-performance, scalable, reliable systems, exposing the design principles, architecture, and innovative applications of parallel, distributed, and cloud computing systems. Topics covered by this book include: facilitating management, debugging, migration, and disaster recovery through virtualization; clustered systems for research or ecommerce applications; designing systems as web services; and social networking systems using peer-to-peer computing. The principles of cloud computing are discussed using examples from open-source and commercial applications, along with case studies from the leading distributed computing vendors such as Amazon, Microsoft, and Google. Each chapter includes

exercises and further reading, with lecture slides and more available online. This book will be ideal for students taking a distributed systems or distributed computing class, as well as for professional system designers and engineers looking for a reference to the latest distributed technologies including cloud, P2P and grid computing. Complete coverage of modern distributed computing technology including clusters, the grid, service-oriented architecture, massively parallel processors, peer-to-peer networking, and cloud computing Includes case studies from the leading distributed computing vendors: Amazon, Microsoft, Google, and more Explains how to use virtualization to facilitate management, debugging, migration, and disaster recovery Designed for undergraduate or graduate students taking a distributed systems course—each chapter includes exercises and further reading, with lecture slides and more available online

Sourcebook of Parallel ComputingMorgan Kaufmann Pub

OpenMP Shared Memory Parallel Programming

An Introductory Survey, Revised Second Edition

GARF Assessment Sourcebook

The Mathematics of Egypt, Mesopotamia, China, India, and Islam

International Workshop, IWOMP 2005 and IWOMP 2006, Eugene, OR, USA, June 1-4, 2005, and Reims, France, June 12-15, 2006, Proceedings

A Hardware/software Approach

A User's Guide to the Moon

Parallel Computing is a compelling vision of how computation can seamlessly scale from a single processor to virtually limitless computing power. Unfortunately, the scaling of application performance has not matched peak speed, and the programming burden for these machines remains heavy. This book represents the collected knowledge and experience of over 30 leading parallel computing researchers. The parallel computing hardware, programming considerations, algorithms, software and enabling technologies, as well as several parallel application case studies. (Midwest).

In recent decades it has become obvious that mathematics has always been a worldwide activity. But this is the first book to provide a substantial collection of English translations of key mathematical texts from the five most important ancient and medieval non-Western mathematical cultures, and to put them into full historical and mathematical context. The Mathematics of Egypt, Mesopotamia, China, India, and Islam: A Sourcebook of Parallel Computing is an appreciation of these cultures' important contributions to world mathematics. The five section authors--Annette Imhausen (Egypt), Eleanor Robson (Mesopotamia), Joseph Dauben (China), Kim Plofker (India), and J. Lennart Berggren (Islam)--are experts in their fields. Each author has selected key texts and in many cases provided new translations. The authors have also written substantial sections of explanatory notes that put each selection into context. This authoritative commentary allows readers to understand the sometimes unfamiliar mathematics of these civilizations and the purpose and significance of each text. Addressing a critical gap in the mathematics literature in English, this book is an essential resource for anyone with at least an undergraduate degree in mathematics who wishes to appreciate the role that mathematics has played in the development of modern science and technology. The book is also an indispensable guide for mathematics teachers who want to use non-Western mathematical ideas in the classroom.

Designed for introductory parallel computing courses at the advanced undergraduate or beginning graduate level, Elements of Parallel Computing presents the fundamental concepts of parallel computing not from the point of view of hardware, but from a more abstract view of algorithmic and implementation patterns. The aim is to facilitate the teaching of parallel programming by surveying some of the most important abstract representation of the underlying hardware. The presentation is friendly and informal. The content of the book is language neutral, using pseudocode that represents common programming language models. The first five chapters present core concepts in parallel computing. SIMD, shared memory, and distributed memory machine models are covered, along with a brief discussion of what the fundamental activity in parallel algorithmic design, starting with a naive example, and continuing with a discussion of some key algorithmic structures. Important programming models are presented in depth, as well as important concepts of performance analysis, including work-depth analysis of task graphs, communication analysis of distributed memory algorithms, key performance metrics, and analysis of the book presents three case studies that reinforce the concepts of the earlier chapters. One feature of these chapters is to contrast different solutions to the same problem, using select problems that aren't discussed frequently in parallel computing textbooks. They include the Single Source Shortest Path Problem, the Eikonal equation, and a classical computational geometry problem: computing the shortest path in a planar graph. Sequential algorithms, each chapter first discusses the sources of parallelism then surveys parallel algorithms.

The PaCT-2003 (Parallel Computing Technologies) conference was a four-day conference held in Nizhni Novgorod on September 15–19, 2003. This was the 7th international conference of the PaCT series, organized in Russia every odd year. The first conference, PaCT-91, was held in Novosibirsk (Academgorodok), September 7–11, 1991. The next PaCT conferences were held in: Obninsk (near Moscow) in 1993, Yaroslavl, September 9–12, 1997; Pushkin (near St. Petersburg) September 6– 10, 1999; and Akademgorodok (Novosibirsk), September 3–7, 2001. The PaCT proceedings are published by Springer-Verlag in the LNCS series. PaCT-2003 was jointly organized by the Institute of Computational Mathematics and Mathematical Geophysics of the Russian Academy of Sciences (Novosibirsk) and the State University of Nizhni Novgorod. The conference attracted about 100 participants from around the world. Authors from 23 countries submitted 78 papers. Of those submitted, 38 papers were selected for the conference as regular ones; there were also 40 posters presented. All the papers were internationally reviewed by at least three referees. As usual a demo session was organized for the participants. Many thanks to our sponsors: the Russian Academy of Sciences, the Russian Fund for Basic Research, the Russian State Committee of Higher Education, IBM and Intel (Intel laboratory in Nizhni Novgorod) for their financial support.

Architectures, Algorithms, and Applications

Mapping Security

A Sourcebook

Parallel Computing

Euro-Par 2008 Parallel Processing

Parallel Computer Architecture

Parallel Programming with MPI

The Sourcebook of Nonverbal Measures provides a comprehensive discussion of research choices for investigating nonverbal phenomena. The volume presents many of the primary means by which researchers assess nonverbal cues. Editor Valerie Manusov has collected both well-established and new measures used in researching nonverbal behaviors, illustrating the broad spectrum of measures appropriate for use in research, and providing a critical resource for future studies. With chapters written by the creators of the research measures, this volume represents work across disciplines, and provides first-hand experience and thoughtful guidance on the use of nonverbal measures. It also offers research strategies researchers can use to answer their research questions; discussions of larger research paradigms into which a measure may be placed; and analysis tools to help researchers think through the research choices available to them. With its thorough and pragmatic approach, this Sourcebook will be an invaluable resource for studying nonverbal behavior. Researchers in interpersonal communication, psychology, personal relationships, and related areas will find it to be an essential research tool.

The only work to date to collect data gathered during the American and Soviet missions in an accessible and complete reference of current scientific and technical information about the Moon.

To configure and maintain an operating system is serious business. With UNIX and its wide variety of "flavors," it can be especially difficult and frustrating, and networking with UNIX adds still more challenges. UNIX Administration: A Comprehensive Sourcebook for Effective Systems & Network Management is a one-stop handbook for the administration and maintenance of UNIX systems and networks. With an outstanding balance of concepts and practical matters, it covers the entire range of administrative tasks, from the most basic to the advanced, from system startup and shutdown to network security and kernel reconfiguration. While focusing on the primary UNIX platforms, the author discusses all of the most common UNIX "flavors," including Solaris, Linux, HP-UX, AIX and SGI IRIX. Three chapters of case studies offer a practical look at UNIX implementation issues: UNIX installation, disk space upgrade, and several emergency situations that every administrator must expect to face at some point. Diverse yet detailed, filled with examples and specific procedures, this is the one book that both the novice and the seasoned professional need to learn UNIX administration and effectively perform their daily system and network-related duties.

First published in 1998. Routledge is an imprint of Taylor & Francis, an informa company.

Sourcebook of Parallel Computing

Euro-Par 2013: Parallel Processing

Understanding the Buzzwords and Basics of Accounting Using Today's Most Popular Software

Optimizing Supercompilers for Supercomputers

14th International Euro-Par Conference, Las Palmas de Gran Canaria, Spain, August 26-29, 2008, Proceedings

The Sourcebook of Contemporary Interiors

Applications of the OCTA System

A breakthrough sourcebook to the challenges and solutions for mobile database systems This text enables readers to effectively manage mobile database systems (MDS) and data dissemination via wireless channels. The author explores the mobile communication platform and analyzes its use in the development of a distributed database management system. Workable solutions for key challenges in wireless information management are presented throughout the text. Following an introductory chapter that includes important milestones in the history and development of mobile data processing, the text provides the information, tools, and resources needed for MDS management, including:

- * Fundamentals of wireless communication**
- * Location and handoff management**
- * Fundamentals of conventional database management systems and why existing approaches are not adequate for mobile databases**
- * Concurrency control mechanism schemes**
- * Data processing and mobility**
- * Management of transactions**
- * Mobile database recovery schemes**
- * Data dissemination via wireless channels**

Case studies and examples are used liberally to aid in the understanding and visualization of complex concepts. Various exercises enable readers to test their grasp of each topic before advancing in the text. Each chapter also concludes with a summary of key concepts as well as references for further study. Professionals in the mobile computing industry, particularly e-commerce, will find this text indispensable. With its extensive use of case studies, examples, and exercises, it is also highly recommended as a graduate-level textbook.

This book differs from traditional numerical analysis texts in that it focuses on the motivation and ideas behind the algorithms presented rather than on detailed analyses of them. It presents a broad overview of methods and software for solving mathematical problems arising in computational modeling and data analysis, including proper problem formulation, selection of effective solution algorithms, and interpretation of results. In the 20 years since its original publication, the modern, fundamental perspective of this book has aged well, and it continues to be used in the classroom. This Classics edition has been updated to include pointers to Python software and the Chebfun package, expansions on barycentric formulation for Lagrange polynomial interpretation and stochastic methods, and the availability of about 100 interactive educational modules that dynamically illustrate the concepts and algorithms in the book. Scientific Computing: An Introductory Survey, Second Edition is intended as both a textbook and a reference for computationally oriented disciplines that need to solve mathematical problems.

This book outlines a set of issues that are critical to all of parallel architecture--communication latency, communication bandwidth, and coordination of cooperative work (across modern designs). It describes the set of techniques available in hardware and in software to address each issues and explore how the various techniques interact.

An Introduction to Parallel Programming is the first undergraduate text to directly address compiling and running parallel programs on the new multi-core and cluster architecture. It explains how to design, debug, and evaluate the performance of distributed and shared-memory programs. The author Peter Pacheco uses a tutorial approach to show students how to develop effective parallel programs with MPI, Pthreads, and OpenMP, starting with small programming examples and building progressively to more challenging ones. The text is written for students in undergraduate parallel programming or parallel computing courses designed for the computer science major or as a service course to other departments; professionals with no background in parallel computing. Takes a tutorial approach, starting with small programming examples and building progressively to more challenging examples Focuses on designing, debugging and evaluating the performance of distributed and shared-memory programs Explains how to develop parallel programs using MPI, Pthreads, and OpenMP programming models

ISPA 2006 Workshops : ISPA 2006 International Workshops, FHPCN, XHPC, S-GRACE, GridGIS, HPC-GTP, PDCE, ParDMCom, WOMP, ISDF, and UPWN, Sorrento, Italy, December 4-7, 2006 : Proceedings

Accounting Made Easy with Your Computer

Elements of Parallel Computing

Mobile Database Systems

11th International Euro-Par Conference, Lisbon, Portugal, August 30 – September 2, 2005, Proceedings

The book discusses the fundamentals of high-performance computing. The authors combine visualization, comprehensibility, and strictness in their material presentation, and thus influence the reader towards practical application and learning how to solve real computing problems. They address both key approaches to programming modern computing systems: multithreading-based parallelizing in shared memory systems, and applying message-passing technologies in distributed systems. The book is suitable for undergraduate and graduate students, and for researchers and practitioners engaged with high-performance computing systems. Each chapter begins with a theoretical part, where the relevant terminology is introduced along with the basic theoretical results and methods of parallel programming, and concludes with a list of test questions and problems of varying difficulty. The authors include many solutions and hints, and often sample code.

Written by recognized LOD leaders, this is a coherent, state-of-the-art account of cutting-edge LOD research and development. This complete resource enables programmers to incorporate LOD technology into their own systems.

This book constitutes the refereed proceedings of the 19th International Conference on Parallel and Distributed Computing, Euro-Par 2013, held in Aachen, Germany, in August 2013. The 70 revised full papers presented were carefully reviewed and selected from 261 submissions. The papers are organized in 16 topical sections: support tools and environments; performance prediction and evaluation; scheduling and load balancing; high-performance architectures and compilers; parallel and distributed data management; grid, cluster and cloud computing; peer-to-peer computing; distributed systems and algorithms; parallel and distributed programming; parallel numerical algorithms; multicore and manycore programming; theory and algorithms for parallel computation; high performance networks and communication; high performance and scientific applications; GPU and accelerator computing; and extreme-scale computing.