

## Computer Architecture And Parallel Processing Kai Hwang

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 searchers 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, Amdahls 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

**THE CONTEXT OF PARALLEL PROCESSING** The field of digital computer architecture has grown explosively in the past two decades. Through a steady stream of experimental research, tool-building efforts, and theoretical studies, the design of an instruction-set architecture, once considered an art, has been transformed into one of the most quantitative branches of computer technology. At the same time, better understanding of various forms of concurrency, from standard pipelining to massive parallelism, and invention of architectural structures to support a reasonably efficient and user-friendly programming model for such systems, has allowed hardware performance to continue its exponential growth. This trend is expected to continue in the near future. This explosive growth, linked with the expectation that performance will continue its exponential rise with each new generation of hardware and that (in stark contrast to software) computer hardware will function correctly as soon as it comes off the assembly line, has its down side. It has led to unprecedented hardware complexity and almost intolerable development costs. The challenge facing current and future computer designers is to institute simplicity where we now have complexity; to use fundamental theories being developed in this area to gain performance and ease-of-use benefits from simpler circuits; to understand the interplay between technological capabilities and limitations, on the one hand, and design decisions based on user and application requirements on the other.

A complete source of information on almost all aspects of parallel computing from introduction, to architectures, to programming paradigms, to algorithms, to programming standards. It covers traditional Computer Science algorithms, scientific computing algorithms and data intensive algorithms.

The computing world today is in the middle of a revolution: mobile clients and cloud computing have emerged as the dominant paradigms driving programming and hardware innovation today. The Fifth Edition of Computer Architecture focuses on this dramatic shift, exploring the ways in which software and technology in the cloud are accessed by cell phones, tablets, laptops, and other mobile computing devices. Each chapter includes two real-world examples, one mobile and one datacenter, to illustrate this revolutionary change. Updated to cover the mobile computing revolution Emphasizes the two most important topics in architecture today: memory hierarchy and parallelism in all its forms. Develops common themes throughout each chapter: power, performance, cost, dependability, protection, programming models, and emerging trends ("What's Next") Includes three review appendices in the printed text. Additional reference appendices are available online. Includes updated Case Studies and completely new exercises.

Computer Architecture

Parallel Computer Architecture

Designing Embedded Hardware

Architecture, Programming and Algorithms

Computer Architecture and Organization

***Technological improvements continue to push back the frontier of processor speed in modern computers. Unfortunately, the computational intensity demanded by modern research problems grows even faster. Parallel computing has emerged as the most successful bridge to this computational gap, and many popular solutions have emerged based on its concepts***

***This book covers the syllabus of GGSIPU, DU, UPTU, PTU, MDU, Pune University and many other universities. •It is useful for B.Tech(CSE/IT), M.Tech(CSE), MCA(SE) students. •Many solved problems have been added to make this book more fresh. •It has been divided in three parts :Parallel Algorithms, Parallel Programming and Super Computers.***

***Since the publication of the first edition, parallel computing technology has gained considerable momentum. A large proportion of this has come from the improvement in VLSI techniques, offering one to two orders of magnitude more devices than previously possible. A second contributing factor in the fast development of the subject is commercialization. The supercomputer is no longer restricted to a few well-established research institutions and large companies. A new computer breed combining the architectural advantages of the supercomputer with the advance of VLSI technology is now available at very attractive prices. A pioneering device in this development is the transputer, a VLSI processor specifically designed to operate in large concurrent systems. Parallel Computers 2: Architecture, Programming and Algorithms reflects the shift in emphasis of parallel computing and tracks the development of supercomputers in the years since the first edition was published. It looks at large-scale parallelism as found in transputer ensembles. This extensively rewritten second edition includes major new sections on the transputer and the OCCAM language. The book contains specific information on the various types of machines available, details of computer architecture and technologies, and descriptions of programming languages and algorithms. Aimed at an advanced undergraduate and postgraduate level, this handbook is also useful for research workers, machine designers, and programmers concerned with parallel computers. In addition, it will serve as a guide for potential parallel computer users, especially in disciplines where large amounts of computer time are***

*regularly used.*

***Parallel Computing Deals With The Topics Of Current Interests In Parallel Processing Architectures (Synchronous Parallel Architectures). The Synchronous Model Of Parallel Processing Is Based On Two Orthogonal Fundamental Ideas, Viz.,1. Temporal Parallelism (Pipeline Processing), And2. Spatial Parallelism (Simd Parallel Processing).This Book Is Devoted To An Indepth Treatment Of Both Of The Above Ideas. The Primary Goal Here Is To Provide A Deeper Understanding Of The Ideas And Principles Involved And Not The Description Of Machines Which Could Be Found Elsewhere. The Material Presented In This Book Has Evolved Through The Advanced Courses Taught By The Author In Architecture And Parallel Processing. A One Semester Advanced Course Can Be Planned Employing The Material From This Book, Supplemented By The Papers Of Current Interests From Current Technical Literature.***

***Parallel Processing, 1980 to 2020***

***Technology, Architecture, Programming***

***Parallel Programming***

***Fundamentals and Architecture Security***

***Parallel Computer Organization and Design***

Intelligent readers who want to build their own embedded computer systems-- installed in everything from cell phones to cars to handheld organizers to refrigerators-- will find this book to be the most in-depth, practical, and up-to-date guide on the market. Designing Embedded Hardware carefully steers between the practical and philosophical aspects, so developers can both create their own devices and gadgets and customize and extend off-the-shelf systems. There are hundreds of books to choose from if you need to learn programming, but only a few are available if you want to learn to create hardware. Designing Embedded Hardware provides software and hardware engineers with no prior experience in embedded systems with the necessary conceptual and design building blocks to understand the architectures of embedded systems. Written to provide the depth of coverage and real-world examples developers need, Designing Embedded Hardware also provides a road-map to the pitfalls and traps to avoid in designing embedded systems. Designing Embedded Hardware covers such essential topics as: The principles of developing computer hardware Core hardware designs Assembly language concepts Parallel I/O Analog-digital conversion Timers (internal and external) UART Serial Peripheral Interface Inter-Integrated Circuit Bus Controller Area Network (CAN) Data Converter Interface (DCI) Low-power operation This invaluable and eminently useful book gives you the practical tools and skills to develop, build, and program your own application-specific computers.

Computer architecture deals with the physical configuration, logical structure, formats, protocols, and operational sequences for processing data, controlling the configuration, and controlling the operations over a computer. It also encompasses word lengths, instruction codes, and the interrelationships among the main parts of a computer or group of computers. This two-volume set offers a comprehensive coverage of the field of computer organization and architecture.

Computer Architecture/Software Engineering

This two-volume set provides comprehensive coverage of the field of computer organization and architecture. The first book in the set gives complete coverage of the subjects pertaining to introductory courses in computer organization and architecture, including: Instruction set architecture and design Assembly language programming Computer arithmetic Processing unit design Memory system design Input-output design and organization Pipelining design techniques Reduced Instruction Set Computers (RISCs) The second volume provides advanced coverage of the field. Both books benefit from the authors' many years' experience in teaching this field as they offer real world applications, examples of machines, case studies and practical experiences in each chapter.

Introduction to Parallel Computing

Advanced Computer Architecture

for Multicore and Cluster Systems

From Algorithms to Programming on State-of-the-Art Platforms

A Survey of Available Parallel Computer Systems

**Programming Massively Parallel Processors: A Hands-on Approach, Second Edition, teaches students how to program massively parallel processors. It offers a detailed discussion of various techniques for constructing parallel programs. Case studies are used to demonstrate the development process, which begins with computational thinking and ends with effective and efficient parallel programs. This guide shows both student and professional alike the basic concepts of parallel programming and GPU architecture. Topics of performance, floating-point format, parallel patterns, and dynamic parallelism are covered in depth. This revised edition contains more parallel programming examples, commonly-used libraries such as Thrust, and explanations of the latest tools. It also provides new coverage of CUDA 5.0, improved performance, enhanced development tools, increased hardware support, and more; increased coverage of related technology, OpenCL and new material on algorithm patterns, GPU clusters, host programming, and data parallelism; and two new case studies (on MRI reconstruction and molecular visualization) that explore the latest applications of CUDA and GPUs for scientific research and high-performance computing. This book should be a valuable resource for advanced students, software engineers, programmers, and hardware engineers. New coverage of CUDA 5.0, improved performance, enhanced development tools, increased hardware support, and more Increased coverage of related technology,**

OpenCL and new material on algorithm patterns, GPU clusters, host programming, and data parallelism Two new case studies (on MRI reconstruction and molecular visualization) explore the latest applications of CUDA and GPUs for scientific research and high-performance computing

Today all computers, from tablet/desktop computers to super computers, work in parallel. A basic knowledge of the architecture of parallel computers and how to program them, is thus, essential for students of computer science and IT professionals. In its second edition, the book retains the lucidity of the first edition and has added new material to reflect the advances in parallel computers. It is designed as text for the final year undergraduate students of computer science and engineering and information technology. It describes the principles of designing parallel computers and how to program them. This second edition, while retaining the general structure of the earlier book, has added two new chapters, 'Core Level Parallel Processing' and 'Grid and Cloud Computing' based on the emergence of parallel computers on a single silicon chip popularly known as multicore processors and the rapid developments in Cloud Computing. All chapters have been revised and some chapters are re-written to reflect the emergence of multicore processors and the use of MapReduce in processing vast amounts of data. The new edition begins with an introduction to how to solve problems in parallel and describes how parallelism is used in improving the performance of computers. The topics discussed include instruction level parallel processing, architecture of parallel computers, multicore processors, grid and cloud computing, parallel algorithms, parallel programming, compiler transformations, operating systems for parallel computers, and performance evaluation of parallel computers.

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.

A clear illustration of how parallel computers can be successfully applied to large-scale scientific computations. This book demonstrates how a variety of applications in physics, biology, mathematics and other sciences were implemented on real parallel computers to produce new scientific results. It investigates issues of fine-grained parallelism relevant for future supercomputers with particular emphasis on hypercube architecture. The authors describe how they used an experimental approach to configure different massively parallel machines, design and implement basic system software, and develop algorithms for frequently used mathematical computations. They also devise performance models, measure the performance characteristics of several computers, and create a high-performance computing facility based exclusively on parallel computers. By addressing all issues involved in scientific problem solving, Parallel Computing Works! provides valuable insight into computational science for large-scale parallel architectures. For those in the sciences, the findings reveal the usefulness of an important experimental tool. Anyone in supercomputing and related computational fields will gain a new perspective on the potential contributions of parallelism. Includes over 30 full-color illustrations.

Parallel Computing

Encyclopedia of Parallel Computing

Computer Organization and Design

Parallelism, Scalability, Programmability

A Hands-on Approach

Motivation It is now possible to build powerful single-processor and multiprocessor systems and use them efficiently for data processing, which has seen an explosive expansion in many areas of computer science and engineering. One approach to meeting the performance requirements of the applications has been to utilize the most powerful single-processor system that is available. When such a system does not provide the performance requirements, pipelined and parallel processing structures can be employed. The concept of parallel processing is a departure from sequential processing. In sequential computation one processor is involved and performs one operation at a time. On the other hand, in parallel computation several processors cooperate to solve a problem, which reduces computing time because several operations can be carried out simultaneously. Using several processors that work together on a given computation illustrates a new paradigm in computer problem solving which is completely different from sequential processing. From the practical point of view, this provides sufficient justification to investigate the concept of parallel processing and related issues, such as parallel algorithms. Parallel processing involves utilizing several factors, such as parallel architectures, parallel algorithms, parallel programming languages and performance analysis, which are strongly interrelated. In general, four steps are involved in performing a computational problem in parallel. The first step is to understand the nature of computations in the specific application domain.

In today's workplace, computer and cybersecurity professionals must understand both hardware and software to deploy effective security solutions. This book introduces readers to the fundamentals of computer architecture and organization for security, and provides them with both theoretical and practical solutions to design and implement secure computer systems. Offering an in-depth and innovative introduction to modern computer systems and patent-pending technologies in computer security, the text integrates design considerations with hands-on lessons learned to help practitioners design computer systems that are immune from attacks. Studying computer architecture and organization from a security perspective is a new area. There are many books on computer architectures and many others on computer security. However, books introducing computer architecture and organization with security as the main focus are still rare. This book addresses not only how to secure computer components (CPU, Memory, I/O, and network) but also how to secure data and the computer system as a whole. It also incorporates experiences from the author's recent award-winning teaching and research. The

book also introduces the latest technologies, such as trusted computing, RISC-V, QEMU, cache security, virtualization, cloud computing, IoT, and quantum computing, as well as other advanced computing topics into the classroom in order to close the gap in workforce development. The book is chiefly intended for undergraduate and graduate students in computer architecture and computer organization, as well as engineers, researchers, cybersecurity professionals, and middleware designers.

This historical survey of parallel processing from 1980 to 2020 is a follow-up to the authors' 1981 Tutorial on Parallel Processing, which covered the state of the art in hardware, programming languages, and applications. Here, we cover the evolution of the field since 1980 in: parallel computers, ranging from the Cyber 205 to clusters now approaching an exaflop, to multicore microprocessors, and Graphic Processing Units (GPUs) in commodity personal devices; parallel programming notations such as OpenMP, MPI message passing, and CUDA streaming notation; and seven parallel applications, such as finite element analysis and computer vision. Some things that looked like they would be major trends in 1981, such as big Single Instruction Multiple Data arrays disappeared for some time but have been revived recently in deep neural network processors. There are now major trends that did not exist in 1980, such as GPUs, distributed memory machines, and parallel processing in nearly every commodity device. This book is intended for those that already have some knowledge of parallel processing today and want to learn about the history of the three areas. In parallel hardware, every major parallel architecture type from 1980 has scaled-up in performance and scaled-out into commodity microprocessors and GPUs, so that every personal and embedded device is a parallel processor. There has been a confluence of parallel architecture types into hybrid parallel systems. Much of the impetus for change has been Moore's Law, but as clock speed increases have stopped and feature size decreases have slowed down, there has been increased demand on parallel processing to continue performance gains. In programming notations and compilers, we observe that the roots of today's programming notations existed before 1980. And that, through a great deal of research, the most widely used programming notations today, although the result of much broadening of these roots, remain close to target system architectures allowing the programmer to almost explicitly use the target's parallelism to the best of their ability. The parallel versions of applications directly or indirectly impact nearly everyone, computer expert or not, and parallelism has brought about major breakthroughs in numerous application areas. Seven parallel applications are studied in this book.

A variety of programming models relevant to scientists explained, with an emphasis on how programming constructs map to parts of the computer. What makes computer programs fast or slow? To answer this question, we have to get behind the abstractions of programming languages and look at how a computer really works. This book examines and explains a variety of scientific programming models (programming models relevant to scientists) with an emphasis on how programming constructs map to different parts of the computer's architecture. Two themes emerge: program speed and program modularity. Throughout this book, the premise is to "get under the hood," and the discussion is tied to specific programs. The book digs into linkers, compilers, operating systems, and computer architecture to understand how the different parts of the computer interact with programs. It begins with a review of C/C++ and explanations of how libraries, linkers, and Makefiles work. Programming models covered include Pthreads, OpenMP, MPI, TCP/IP, and CUDA. The emphasis on how computers work leads the reader into computer architecture and occasionally into the operating system kernel. The operating system studied is Linux, the preferred platform for scientific computing. Linux is also open source, which allows users to peer into its inner workings. A brief appendix provides a useful table of machines used to time programs. The book's website (<https://github.com/divakarvi/bk-spc>) has all the programs described in the book as well as a link to the html text.

Algorithms and Architectures

Architectures, Algorithms, and Applications

Computer Architecture and Parallel Processing

Fundamentals of Computer Organization and Architecture & Advanced Computer Architecture and Parallel Processing, 2 Volume Set

Algorithms and Parallel Computing

**Computer Architecture and Parallel Processing** McGraw-Hill College

**Innovations in hardware architecture, like hyper-threading or multicore processors, mean that parallel computing resources are available for inexpensive desktop computers. In only a few years, many standard software products will be based on concepts of parallel programming implemented on such hardware, and the range of applications will be much broader than that of scientific computing, up to now the main application area for parallel computing. Rauber and Runger take up these recent developments in processor architecture by giving detailed descriptions of parallel programming techniques that are necessary for developing efficient programs for multicore processors as well as for parallel cluster systems and supercomputers. Their book is structured in three main parts, covering all areas of parallel computing: the architecture of parallel systems, parallel programming models and environments, and the implementation of efficient application algorithms. The emphasis lies on parallel programming techniques needed for different architectures. For this second edition, all chapters have been carefully revised. The chapter on architecture of parallel systems has been updated considerably, with a greater emphasis on the architecture of multicore systems and adding new material on the latest developments in computer architecture. Lastly, a completely new chapter on general-purpose GPUs and the corresponding programming techniques has been added. The main goal of the book is to present parallel programming techniques that can be used in many situations for a broad range of application areas and which enable the reader to develop correct and efficient parallel programs. Many examples and exercises are provided to show how to apply the techniques. The book can be used as both a textbook for students and a reference book for professionals. The material presented has been used for courses in parallel programming at different universities for many years.**

**Past, Present, Parallel is a survey of the current state of the parallel processing industry. In the early 1980s, parallel computers were generally regarded as academic curiosities whose natural environment was the research laboratory. Today, parallelism is being used by every major computer manufacturer, although in very different ways, to produce increasingly powerful and cost-**

effective machines. The first chapter introduces the basic concepts of parallel computing; the subsequent chapters cover different forms of parallelism, including descriptions of vector supercomputers, SIMD computers, shared memory multiprocessors, hypercubes, and transputer-based machines. Each section concentrates on a different manufacturer, detailing its history and company profile, the machines it currently produces, the software environments it supports, the market segment it is targeting, and its future plans. Supplementary chapters describe some of the companies which have been unsuccessful, and discuss a number of the common software systems which have been developed to make parallel computers more usable. The appendices describe the technologies which underpin parallelism. *Past, Present, Parallel* is an invaluable reference work, providing up-to-date material for commercial computer users and manufacturers, and for researchers and postgraduate students with an interest in parallel computing.

There is a software gap between the hardware potential and the performance that can be attained using today's software parallel program development tools. The tools need manual intervention by the programmer to parallelize the code. Programming a parallel computer requires closely studying the target algorithm or application, more so than in the traditional sequential programming we have all learned. The programmer must be aware of the communication and data dependencies of the algorithm or application. This book provides the techniques to explore the possible ways to program a parallel computer for a given application.

**Solutions Manual to Accompany: Hwang Advanced Computer Architecture**

**A Flexible Parallel Processing Computer Architecture Design for Internetworked Computer Systems**

**Advanced Computer Architecture and Parallel Processing**

**Scientific Programming and Computer Architecture**

**Software development for parallel processing in a distributed computer architecture**

Teaching fundamental design concepts and the challenges of emerging technology, this textbook prepares students for a career designing the computer systems of the future. In-depth coverage of complexity, power, reliability and performance, coupled with treatment of parallelism at all levels, including ILP and TLP, provides the state-of-the-art training that students need. The whole gamut of parallel architecture design options is explained, from core microarchitecture to chip multiprocessors to large-scale multiprocessor systems. All the chapters are self-contained, yet concise enough that the material can be taught in a single semester, making it perfect for use in senior undergraduate and graduate computer architecture courses. The book is also teeming with practical examples to aid the learning process, showing concrete applications of definitions. With simple models and codes used throughout, all material is made open to a broad range of computer engineering/science students with only a basic knowledge of hardware and software.

This is the instructor's manual to a text which presents the latest technologies for parallel processing and high performance computing. The main text deals with advanced computer architecture and parallel processing systems and techniques, providing an integrated study of computer hardware and software systems, and the material is suitable for use on courses found in computer science, computer engineering, or electrical engineering departments. This material is only available to lecturers.

Parallel computing has been the enabling technology of high-end machines for many years. Now, it has finally become the ubiquitous key to the efficient use of any kind of multi-processor computer architecture, from smart phones, tablets, embedded systems and cloud computing up to exascale computers. **ParCo2013** – the latest edition of the biennial International Conference on Parallel Computing – held from 10 to 13 September 2013, in Garching, Germany. The conference focused on several key parallel computing areas. Themes included parallel programming models for multi- and manycore CPUs, GPUs, FPGAs and heterogeneous platforms, the performance engineering processes that must be adapted to efficiently use these new and innovative platforms, novel numerical algorithms and approaches to large-scale simulations of problems in science and engineering. The conference programme also included twelve mini-symposia (including an industry session and a special PhD Symposium), which comprehensively represented and intensified the discussion of current hot topics in high performance and parallel computing. These special sessions covered large-scale supercomputing, novel challenges arising from parallel architectures (multi-/manycore, heterogeneous platforms, FPGAs), multi-level algorithms as well as multi-scale, multi-physics and multi-dimensional problems. It is clear that parallel computing – including the processing of large data sets (“Big Data”) – will remain a persistent driver of research in all fields of innovative computing, which makes this book relevant to all those with an interest in this field.

**Computer Systems Organization -- Parallel architecture.**

**Parallel Computing Works!**

**High-performance Computer Architecture**

**A Quantitative Approach**

**Past, Present, Parallel**

**Theory and Computation**

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

This is the first book in the two-volume set offering comprehensive coverage of the field of computer organization and architecture. This book provides complete coverage of the subjects pertaining to introductory courses in computer organization and architecture, including: \* Instruction set architecture and design \* Assembly language programming \* Computer arithmetic \* Processing unit design \* Memory system design \* Input-output design and organization \* Pipelining design techniques \* Reduced Instruction Set Computers (RISCs) The authors, who share over 15 years of undergraduate and graduate level instruction in computer architecture, provide real world applications, examples of machines, case studies and practical experiences in each chapter.

This comprehensive new text from author Kai Hwang covers four important aspects of parallel and distributed computing -- principles, technology, architecture, and programming -- and can be used for several upper-level courses.

The performance of software systems is dramatically affected by how well software designers understand the basic hardware technologies at work in a system. Similarly, hardware designers must understand the far-reaching effects their design decisions have on software applications. For readers in either category, this classic introduction to the field provides a look deep into the computer. It demonstrates the relationships between the software and hardware and focuses on the foundational concepts that are the basis for current computer

design.

Parallel Computing: Accelerating Computational Science and Engineering (CSE)

Handbook of Parallel Computing and Statistics

Programming Massively Parallel Processors

Fundamentals of Computer Organization and Architecture

Introduction to Parallel Processing

Advancements in microprocessor architecture, interconnection technology, and software development have fueled rapid growth in parallel and distributed computing. However, this development is only of practical benefit if it is accompanied by progress in the design, analysis and programming of parallel algorithms. This concise textbook provides, in one place, three mainstream parallelization approaches, Open MPP, MPI and OpenCL, for multicore computers, interconnected computers and graphical processing units. An overview of practical parallel computing and principles will enable the reader to design efficient parallel programs for solving various computational problems on state-of-the-art personal computers and computing clusters. Topics covered range from parallel algorithms, programming tools, OpenMP, MPI and OpenCL, followed by experimental measurements of parallel programs' run-times, and by engineering analysis of obtained results for improved parallel execution performances. Many examples and exercises support the exposition.

Parallel Computers 2

Scalable Parallel Computing

Pipelined and Parallel Processor Design

Parallelism

Parallel Processing and Parallel Algorithms