Advanced Computer Architecture Hwang Solution Manual

"The Encyclopedia of Microcomputers serves as the ideal companion reference to the popular Encyclopedia of Computer Science and Technology. Now in its 10th year of publication, this timely reference work details the broad spectrum of microcomputer technology, including microcomputer history; explains and illustrates the use of microcomputers throughout academe, business, government, and society in general; and assesses the future impact of this rapidly changing technology."

This book provides a non-technical introduction to High Performance Computing applications together with advice about how beginners can start to write parallel programs. The authors show what HPC can offer geographers and social scientists and how it can be used in GIS. They provide examples of where it has already been used and suggestions for other areas of application in geography and the social sciences. Case studies drawn from geography explain the key principles and help to understand the logic and thought processes that lie behind the parallel programming.

Content Description #Includes bibliographical references and index.

This concise text is designed to present the recent advances in parallel and distributed architectures and algorithms within an integrated framework. Beginning with an introduction to the basic concepts, the book goes on discussing the basic methods of parallelism exploitation in computation through vector processing, super scalar and VLIW processing, array processing, associative processing, systolic algorithms, and dataflow computation. After introducing interconnection networks, it discusses parallel algorithms for sorting, Fourier transform, matrix algebra, and graph theory. The second part focuses on basics and selected theoretical issues of distributed processing. Architectures and algorithms have been dealt in an integrated way throughout the book. The last chapter focuses on the different paradigms and issues of high performance computing making the reading more interesting. This book is meant for the senior level undergraduate and postgraduate students of computer science and engineering, and information technology. The book is also useful for the postgraduate students of computer science and computer application.

Intelligent/smart systems have become common practice in many engineering applications. On the other hand, current low cost standard CMOS technology (and future foreseeable developments) makes available enormous potentialities. The next breakthrough will be the design and development of "smart adaptive systems on silicon" i.e. very power and highly size efficient complete systems (i.e. sensing, computing and "actuating" actions) with intelligence on board on a single silicon die. Smart adaptive systems on silicon will be able to "adapt" autonomously to the changing environment and will be able to implement "intelligent" behaviour and both perceptual and cognitive tasks. At last, they will

communicate through wireless channels, they will be battery supplied or remote powered (via inductive coupling) and they will be ubiquitous in our every day life. Although many books deal with research and engineering topics (i.e. algorithms, technology, implementations, etc.) few of them try to bridge the gap between them and to address the issues related to feasibility, reliability and applications. Smart Adaptive Systems on Silicon, though not exhaustive, tries to fill this gap and to give answers mainly to the feasibility and reliability issues. Smart Adaptive Systems on Silicon mainly focuses on the analog and mixed mode implementation on silicon because this approach is amenable of achieving impressive energy and size efficiency. Moreover, analog systems can be more easily interfaced with sensing and actuating devices.

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 si e 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.

The salient features of the book are as follows: • Hybrid Elements including topics like Memory organization, Binary representation of data, Computer arithmetic Software for parallel programming, tagged across some chapters through Quick Response (QR) Codes • Learning objectives tagged across chapters: • Emphasis on parallelism, scalability and programmability aspects of computer architecture. It presents the analysis of scalability • Issues related to instruction level parallelism, processor clock speed, and power consumption defined according to the recent developments in processor design • Inclusion of important topics like processor design, control unit, input and output, parallelis • erial Bus, Real systems– IBM, Hitachi, Cray, Intel, UltraSparc, Blue Gene (from IBM), Cray XT series, XT5 and XMT, Fujitsu, DEC, MasPar, Tera, Stardent Topical inclusions include: • Pipelining hazards, data hazards and control

hazards • PCI Bus and PCI Express • Interconnection networks and cluster computers • MPI, openMP, PVM, Pthreads • Multicore processors • Impact of technology • Stream processing • Programming language Chapel • Updated coverage of recent processors and systems: Intel Pentium IV, Sun UltraSparc, Blue Gene (from IBM), Cray XT Series, XT5 and XMT Useful pedagogical features include the following: • Plenty of background material on OLC • Diagrams illustrating the basic concepts: 320 • A good number of case studies and: 6 • Solved problems: 114 • Exercise and review problems at the end of chapters: 251 • Tables: 40 • Solved Examples: 114 • Exercise Problems: 251

This comprehensive volume introduces educational units dealing with important topics in Mathematics, Modelling and Algorithms. Key Features: Illustrative examples and exercises Comprehensive bibliography

An interdisciplinary guide to enabling technologies for 3D ICs and 5G mobility, covering packaging, design to product life and reliability assessments Features an interdisciplinary approach to the enabling technologies and hardware for 3D ICs and 5G mobility Presents statistical treatments and examples with tools that are easily accessible, such as Microsoft’s Excel and Minitab Fundamental design topics such as electromagnetic design for logic and RF/passives centric circuits are explained in detail Provides chapter-wise review questions and powerpoint slides as teaching tools Mathematics of Computing -- Parallelism.

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 dev- opment 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.

This book constitutes the refereed proceedings of the 4th International Workshop on Distributed Computing, IWDC 2002, held in Calcutta, India, in December 2002. The 31 revised full papers and 3 student papers presented together with 3 keynote papers were carefully reviewed and selected from more than 90 submissions. The papers are organized in topical sections on Web caching, distributed computing, wireless networks, wireless mobile systems, VLSI and parallel systems, optical networks, and distributed

systems.

Over the past several decades, applications permeated by advances in digital signal processing have undergone unprecedented growth in capabilities. The editors and authors of High Performance Embedded Computing Handbook: A Systems Perspective have been significant contributors to this field, and the principles and techniques presented in the handbook are reinforced by examples drawn from their work. The chapters cover system components found in today's HPEC systems by addressing design trade-offs, implementation options, and techniques of the trade, then solidifying the concepts with specific HPEC system examples. This approach provides a more valuable learning tool, Because readers learn about these subject areas through factual implementation cases drawn from the contributing authors' own experiences. Discussions include: Key subsystems and components Computational characteristics of high performance embedded algorithms and applications Front-end real-time processor technologies such as analog-to-digital conversion, application-specific integrated circuits, field programmable gate arrays, and intellectual property–based design Programmable HPEC systems technology, including interconnection fabrics, parallel and distributed processing, performance metrics and software architecture, and automatic code parallelization and optimization Examples of complex HPEC systems representative of actual prototype developments Application examples, including radar, communications, electro-optical, and sonar applications The handbook is organized around a canonical framework that helps readers navigate through the chapters, and it concludes with a discussion of future trends in HPEC systems. The material is covered at a level suitable for practicing engineers and HPEC computational practitioners and is easily adaptable to their own implementation requirements.

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 Rünger 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.

This book constitutes the refereed proceedings of the 18th International Symposium on Computer and Information Sciences, ISCIS 2003, held in Antalya, Turkey in November 2003. The 135 revised papers presented together with 2 invited papers were carefully reviewed and selected from over 360 submissions. The papers are organized in topical sections on architectures and systems, theoretical computer science, databases and information retrieval, e-commerce, graphics and computer vision, intelligent systems and robotics, multimedia, networks and security, parallel and distributed computing, soft computing, and software engineering.

Storage Systems: Organization, Performance, Coding, Reliability and Their Data Processing was motivated by the 1988 Redundant Array of Inexpensive/Independent Disks proposal to replace large form factor mainframe disks with an array of commodity disks. Disk loads are balanced by striping data into strips—with one strip per disk— and storage reliability is enhanced via replication or erasure coding, which at best dedicates k strips per stripe to tolerate k disk failures. Flash memories have resulted in a paradigm shift with Solid State Drives (SSDs) replacing Hard Disk Drives (HDDs) for high performance applications. RAID and Flash have resulted in the emergence of new storage companies, namely EMC, NetApp, SanDisk, and Purestorage, and a multibillion-dollar storage market. Key new conferences and publications are reviewed in this book. The goal of the book is to expose students, researchers, and IT professionals to the more important developments in storage systems, while covering the evolution of storage technologies, traditional and novel databases, and novel sources of data. We describe several prototypes: FAWN at CMU, RAMCloud at Stanford, and Lightstore at MIT; Oracle's Exadata, AWS' Aurora, Alibaba's PolarDB, Fungible Data Center; and author's paper designs for cloud storage, namely heterogeneous disk arrays and hierarchical RAID. • Surveys storage technologies and lists sources of data: measurements, text, audio, images, and video • Familiarizes with paradigms to improve performance: caching, prefetching, log-structured file systems, and merge-trees (LSMs) • Describes RAID organizations and analyzes their performance and reliability • Conserves storage via data compression, deduplication, compaction, and secures data via encryption • Specifies implications of storage technologies on performance and power consumption • Exemplifies database parallelism for big data, analytics, deep learning via multicore CPUs, GPUs, FPGAs, and ASICs, e.g., Google's Tensor Processing Units

Parallel computers have started to completely revolutionize scientific computation. Articles in this volume represent applied mathematics, computer science, and application aspects of parallel scientific computing. Major advances are discussed dealing with multiprocessor architectures, parallel algorithm development and analysis, parallel systems and programming languages. The optimization of the application of massively parallel architectures to real world problems will provide the impetus for the development of entirely new approaches to these technical situations.

This book addresses the issues of nonlinearity and disorder. It covers mathematical and numerical techniques as well as applications of nonlinearity and disorder. The analysis of continuous and discrete systems is also shown.

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 Suitable for a one- or two-semester undergraduate or beginning graduate course in computer science and computer engineering, Computer Organization, Design, and Architecture, Fourth Edition presents the operating principles, capabilities, and limitations of digital computers to enable development of complex yet efficient systems. With 40% updated material and four new chapters, this edition takes students through a solid, up-to-date exploration of single- and multiple-processor systems, embedded architectures, and performance evaluation. New to the Fourth Edition Additional material that covers the ACM/IEEE computer science and engineering curricula More coverage on computer organization, embedded systems, networks, and performance evaluation Expanded discussions of RISC, CISC, VLIW, and parallel/pipelined architectures The latest information on integrated circuit technologies and devices, memory hierarchy, and storage Updated examples, references, and problems Supplying appendices with relevant details of integrated circuits reprinted from vendors' manuals, this book provides all of the necessary information to program and design a computer system.

Today's highly parameterized large-scale distributed computing systems may be composed of a large number of various components (computers, databases, etc) and must provide a wide range of services. The users of such systems, located at different (geographical or managerial) network cluster may have a limited access to the system's services and resources, and different, often conflicting, expectations and requirements. Moreover, the information and data processed in such dynamic environments may be incomplete, imprecise, fragmentary, and overloading. All of the above mentioned issues require some intelligent scalable methodologies for the management of the whole complex structure, which unfortunately may increase the energy consumption of such systems. An optimal energy utilization has reached to a point that many information technology (IT) managers and corporate executives are all up in arms to identify scalable solution that can reduce electricity consumption (so that the total cost of operation is minimized) of their respective large-scale computing systems and simultaneously improve upon or maintain the current throughput of the system. This book in its eight chapters, addresses the fundamental issues related to the energy usage and the optimal low-cost system design in high performance ``green computing'' systems. The recent evolutionary and general metaheuristic-based solutions for energy optimization in data processing, scheduling, resource allocation, and communication in modern computational grids, could and network computing are presented along with several important conventional technologies to cover the hot topics from the fundamental theory of the ''green computing'' concept and to describe the basic architectures of systems. This book points out the potential application areas and provides detailed examples of application case studies in low-energy computational systems.

The development trends and open research issues are also outlined. All of those technologies have formed the foundation for the green computing that we know of today.

Entity Identification to Virtual Reality in Driving Simulation

This highly acclaimed work, first published by Prentice Hall in 1989, is a comprehensive and theoretically sound treatment of parallel and distributed numerical methods. It focuses on algorithms that are naturally suited for massive parallelization, and it explores the fundamental convergence, rate of convergence, communication, and synchronization issues associated with such algorithms. This is an extensive book, which aside from its focus on parallel and distributed algorithms, contains a wealth of material on a broad variety of computation and optimization topics. It is an excellent supplement to several of our other books, including Convex Optimization Algorithms (Athena Scientific, 2015), Nonlinear Programming (Athena Scientific, 1999), Dynamic Programming and Optimal Control (Athena Scientific, 2012), Neuro-Dynamic Programming (Athena Scientific, 1996), and Network Optimization (Athena Scientific, 1998). The on-line edition of the book contains a 95-page solutions manual.

This volume reviews, in the context of partial differential equations, algorithm development that has been specifically aimed at computers that exhibit some form of parallelism. Emphasis is on the solution of PDEs because these are typically the problems that generate high computational demands. The authors discuss architectural features of these computers insomuch as they influence algorithm performance, and provide insight into algorithm characteristics that allow effective use of hardware.

This volume is the first part of a four-volume set (CCIS 190, CCIS 191, CCIS 192, CCIS 193), which constitutes the refereed proceedings of the First International Conference on Computing and Communications, ACC 2011, held in Kochi, India, in July 2011. The 68 revised full papers presented in this volume were carefully reviewed and selected from a large number of submissions. The papers are organized in topical sections on ad hoc networks; advanced micro architecture techniques; autonomic and context-aware computing; bioinformatics and bio-computing; cloud, cluster, grid and P2P computing; cognitive radio and cognitive networks; cyber forensics; database and information systems.

Computer Architecture/Software Engineering

Describes the introduction of advanced computer architecture and parallel processing. Covers the paradigms of computing like synchronous and asynchronous. Detailed explanation of the Flynn's classification, kung's taxonomy and reduction paradigm. provides a detailed treatment of abstract parallel computational models like combination circuits, sorting network, PRAM models, interconnection RAMs. Covers the parallelism in uni processor systems. Provides an extensive treatment of parallel computer structures like pipeline computers, array computers and multiprocessor systems. Covers the concepts of pipeline and classification of pipeline processors. Give description of super scalar, super pipeline design and VLIW processors. Explains the design structures and

algorithms for array processors.

Computer Organization and Design, Fifth Edition, is the latest update to the classic introduction to computer organization. The text now contains new examples and material highlighting the emergence of mobile computing and the cloud. It explores this generational change with updated content featuring tablet computers, cloud infrastructure, and the ARM (mobile computing devices) and x86 (cloud computing) architectures. The book uses a MIPS processor core to present the fundamentals of hardware technologies, assembly language, computer arithmetic, pipelining, memory hierarchies and I/O.Because an understanding of modern hardware is essential to achieving good performance and energy efficiency, this edition adds a new concrete example, Going Faster, used throughout the text to demonstrate extremely effective optimization techniques. There is also a new discussion of the Eight Great Ideas of computer architecture. Parallelism is examined in depth with examples and content highlighting parallel hardware and software topics. The book features the Intel Core i7, ARM Cortex-A8 and NVIDIA Fermi GPU as real-world examples, along with a full set of updated and improved exercises. This new edition is an ideal resource for professional digital system designers, programmers, application developers, and system software developers. It will also be of interest to undergraduate students in Computer Science, Computer Engineering and Electrical Engineering courses in Computer Organization, Computer Design, ranging from Sophomore required courses to Senior Electives. Winner of a 2014 Texty Award from the Text and Academic Authors Association Includes new examples, exercises, and material highlighting the emergence of mobile computing and the cloud Covers parallelism in depth with examples and content highlighting parallel hardware and software topics Features the Intel Core i7, ARM Cortex-A8 and NVIDIA Fermi GPU as real-world examples throughout the book Adds a new concrete example, "Going Faster," to demonstrate how understanding hardware can inspire software optimizations that improve performance by 200 times Discusses and highlights the "Eight Great Ideas" of computer architecture: Performance via Parallelism; Performance via Pipelining; Performance via Prediction; Design for Moore's Law; Hierarchy of Memories; Abstraction to Simplify Design; Make the Common Case Fast; and Dependability via Redundancy Includes a full set of updated and improved exercises

This volume gives an overview of the state-of-the-art with respect to the development of all types of parallel computers and their application to a wide range of problem areas. The international conference on parallel computing ParCo97 (Parallel Computing 97) was held in Bonn, Germany from 19 to 22 September 1997. The first conference in this biannual series was held in 1983 in Berlin. Further conferences were held in Leiden (The Netherlands), London (UK), Grenoble (France) and Gent (Belgium). From the outset the aim with the ParCo (Parallel Computing) conferences was to promote the application of parallel computers to solve real life problems. In the case of ParCo97 a new milestone

was reached in that more than half of the papers and posters presented were concerned with application aspects. This fact reflects the coming of age of parallel computing. Some 200 papers were submitted to the Program Committee by authors from all over the world. The final programme consisted of four invited papers, 71 contributed scientific/industrial papers and 45 posters. In addition a panel discussion on Parallel Computing and the Evolution of Cyberspace was held. During and after the conference all final contributions were refereed. Only those papers and posters accepted during this final screening process are included in this volume. The practical emphasis of the conference was accentuated by an industrial exhibition where companies demonstrated the newest developments in parallel processing equipment and software. Speakers from participating companies presented papers in industrial sessions in which new developments in parallel computing were reported.

Assessing the most valuable technology for an organization is becoming a growing challenge for business professionals confronted with an expanding array of options. This 2007 book is an A-Z compendium of technological terms written for the non-technical executive, allowing quick identification of what the term is and why it is significant. This is more than a dictionary - it is a concise review of the most important aspects of information technology from a business perspective: the major advantages, disadvantages and business value propositions of each term are discussed, as well as sources for further reading, and cross-referencing with other terms where applicable. The essential elements of each concept are covered in a succinct manner so the reader can quickly obtain the required knowledge without wading through exhaustive descriptions. With over 200 terms, this is a valuable reference for non- and semi-technical managers, executives and graduate students in business and technology management.

Scalable Coherent Interface (SCI) is an innovative interconnect standard (ANSI/IEEE Std 1596-1992) addressing the high-performance computing and networking domain. This book describes in depth one specific application of SCI: its use as a high-speed interconnection network (often called a system area network, SAN) for compute clusters built from commodity workstation nodes. The editors and authors, coming from both academia and industry, have been instrumental in the SCI standardization process, the development and deployment of SCI adapter cards, switches, fully integrated clusters, and software systems, and are closely involved in various research projects on this important interconnect. This thoroughly cross-reviewed state-of-the-art survey covers the complete hardware/software spectrum of SCI clusters, from the major concepts of SCI, through SCI hardware, networking, and low-level software issues, various programming models and environments, up to tools and application experiences.

Internet usage has become a facet of everyday life, especially as more technological advances have made it easier to connect to the web from virtually

anywhere in the developed world. However, with this increased usage comes heightened threats to security within digital environments. The Handbook of Research on Modern Cryptographic Solutions for Computer and Cyber Security identifies emergent research and techniques being utilized in the field of cryptology and cyber threat prevention. Featuring theoretical perspectives, best practices, and future research directions, this handbook of research is a vital resource for professionals, researchers, faculty members, scientists, graduate students, scholars, and software developers interested in threat identification and prevention.

This is an introductory book on supercomputer applications written by a researcher who is working on solving scientific and engineering application problems on parallel computers. The book is intended to quickly bring researchers and graduate students working on numerical solutions of partial differential equations with various applications into the area of parallel processing.The book starts from the basic concepts of parallel processing, like speedup, efficiency and different parallel architectures, then introduces the most frequently used algorithms for solving PDEs on parallel computers, with practical examples. Finally, it discusses more advanced topics, including different scalability metrics, parallel time stepping algorithms and new architectures and heterogeneous computing networks which have emerged in the last few years of high performance computing. Hundreds of references are also included in the book to direct interested readers to more detailed and in-depth discussions of specific topics.

Future requirements for computing speed, system reliability, and cost-effectiveness entail the development of alternative computers to replace the traditional von Neumann organization. As computing networks come into being, one of the latest dreams is now possible - distributed computing. Distributed computing brings transparent access to as much computer power and data as the user needs for accomplishing any given task - simultaneously achieving high performance and reliability. The subject of distributed computing is diverse, and many researchers are investigating various issues concerning the structure of hardware and the design of distributed software. Distributed System Design defines a distributed system as one that looks to its users like an ordinary system, but runs on a set of autonomous processing elements (PEs) where each PE has a separate physical memory space and the message transmission delay is not negligible. With close cooperation among these PEs, the system supports an arbitrary number of processes and dynamic extensions. Distributed System Design outlines the main motivations for building a distributed system, including: inherently distributed applications performance/cost resource sharing flexibility and extendibility availability and fault tolerance scalability Presenting basic concepts, problems, and possible solutions, this reference serves graduate students in distributed system design as well as computer professionals analyzing and designing distributed/open/parallel systems. Chapters discuss: the

scope of distributed computing systems general distributed programming languages and a CSP-like distributed control description language (DCDL) expressing parallelism, interprocess communication and synchronization, and fault-tolerant design two approaches describing a distributed system: the time-space view and the interleaving view mutual exclusion and related issues, including election, bidding, and self-stabilization prevention and detection of deadlock reliability, safety, and security as well as various methods of handling node, communication, Byzantine, and software faults efficient interprocessor communication mechanisms as well as these mechanisms without specific constraints, such as adaptiveness, deadlock-freedom, and fault-tolerance virtual channels and virtual networks load distribution problems synchronization of access to shared data while supporting a high degree of concurrency