

Introduction To Parallel Processing Behrooz Parhami Solution

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.

Computer Architecture/Software Engineering

Sustainable Wireless Network-on-Chip Architectures focuses on developing novel Dynamic Thermal Management (DTM) and Dynamic Voltage and Frequency Scaling (DVFS) algorithms that exploit the advantages inherent in WiNoC architectures. The methodologies proposed—combined with extensive experimental validation—collectively represent efforts to create a sustainable NoC architecture for future many-core chips. Current research trends show a necessary paradigm shift towards green and sustainable computing. As implementing massively parallel energy-efficient CPUs and reducing resource consumption become standard, and their speed and power continuously increase, energy issues become a significant concern. The need for promoting research in sustainable computing is imperative. As hundreds of cores are integrated in a single chip, designing effective packages for dissipating maximum heat is infeasible. Moreover, technology scaling is pushing the limits of affordable cooling, thereby requiring suitable design techniques to reduce peak temperatures. Addressing thermal concerns at different design stages is critical to the success of future generation systems. DTM and DVFS appear as solutions to avoid high spatial and temporal temperature variations among NoC components, and thereby mitigate local network hotspots. Defines new complex, sustainable network-on-chip architectures to reduce network latency and energy Develops topology-agnostic dynamic thermal management and dynamic voltage and frequency scaling techniques Describes joint strategies for network- and core-level sustainability Discusses novel algorithms that exploit the advantages inherent in Wireless Network-on-Chip architectures

Foreword -- Foreword to the First Printing -- Preface -- Chapter 1 -- Introduction -- Chapter 2 -- Message Switching Layer -- Chapter 3 -- Deadlock, Livelock, and Starvation -- Chapter 4 -- Routing Algorithms -- Chapter 5 -- CollectiveCommunicationSupport -- Chapter 6 -- Fault-Tolerant Routing -- Chapter 7 -- Network Architectures -- Chapter 8 -- Messaging Layer Software -- Chapter 9 -- Performance Evaluation -- Appendix A -- Formal Definitions for Deadlock Avoidance -- Appendix B -- Acronyms -- References -- Index. Researches and developers of simulation models state that the Java programming language presents a unique and significant opportunity for important changes in the way we develop simulation models today. The most important characteristics of the Java language that are advantageous for simulation are its multi-threading capabilities, its facilities for executing programs across the Web, and its graphics facilities. It is feasible to develop compatible and reusable simulation components that will facilitate the construction of newer and more complex models. This is possible with Java development environments. Another important trend that begun very recently is web-based simulation, i.e., and the execution of simulation models using Internet browser software. This book introduces the application of the Java programming language in discrete-event simulation. In addition, the fundamental concepts and practical simulation techniques for modeling different types of systems to study their general behavior and their performance are introduced. The approaches applied are the process interaction approach to discrete-event simulation and object-oriented modeling. Java is used as the implementation language and UML as the modeling language. The first offers several advantages compared to C++, the most important being: thread handling, graphical user interfaces (GUI) and Web computing. The second language, UML (Unified Modeling Language) is the standard notation used today for modeling systems as a collection of classes, class relationships, objects, and object behavior.

Ideal for graduate and senior undergraduate courses in computer arithmetic and advanced digital design, Computer Arithmetic: Algorithms and Hardware Designs, Second Edition, provides a balanced, comprehensive treatment of computer arithmetic. It covers topics in arithmetic unit design and circuit implementation that complement the architectural and algorithmic speedup techniques used in high-performance computer architecture and parallel processing. Using a unified and consistent framework, the text begins with number representation and proceeds through basic arithmetic operations, floating-point arithmetic, and function evaluation methods. Later chapters cover broad

design and implementation topics-including techniques for high-throughput, low-power, fault-tolerant, and reconfigurable arithmetic. An appendix provides a historical view of the field and speculates on its future. An indispensable resource for instruction, professional development, and research, *Computer Arithmetic: Algorithms and Hardware Designs*, Second Edition, combines broad coverage of the underlying theories of computer arithmetic with numerous examples of practical designs, worked-out examples, and a large collection of meaningful problems. This second edition includes a new chapter on reconfigurable arithmetic, in order to address the fact that arithmetic functions are increasingly being implemented on field-programmable gate arrays (FPGAs) and FPGA-like configurable devices. Updated and thoroughly revised, the book offers new and expanded coverage of saturating adders and multipliers, truncated multipliers, fused multiply-add units, overlapped quotient digit selection, bipartite and multipartite tables, reversible logic, dot notation, modular arithmetic, Montgomery modular reduction, division by constants, IEEE floating-point standard formats, and interval arithmetic. Features: * Divided into 28 lecture-size chapters * Emphasizes both the underlying theories of computer arithmetic and actual hardware designs * Carefully links computer arithmetic to other subfields of computer engineering * Includes 717 end-of-chapter problems ranging in complexity from simple exercises to mini-projects * Incorporates many examples of practical designs * Uses consistent standardized notation throughout * Instructor's manual includes solutions to text problems * An author-maintained website http://www.ece.ucsb.edu/~parhami/text_comp_arit.htm contains instructor resources, including complete lecture slides

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.

The first contemporary comprehensive treatment of optimization without derivatives. This text explains how sampling and model techniques are used in derivative-free methods and how they are designed to solve optimization problems. It is designed to be readily accessible to both researchers and those with a modest background in computational mathematics.

This second edition of *Adaptive Filters: Theory and Applications* has been updated throughout to reflect the latest developments in this field; notably an increased coverage given to the practical applications of the theory to illustrate the much broader range of adaptive filters applications developed in recent years. The book offers an easy to understand approach to the theory and application of adaptive filters by clearly illustrating how the theory explained in the early chapters of the book is modified for the various applications discussed in detail in later chapters. This integrated approach makes the book a valuable resource for graduate students; and the inclusion of more advanced applications including antenna arrays and wireless communications makes it a suitable technical reference for engineers, practitioners and researchers. Key features: • Offers a thorough treatment of the theory of adaptive signal processing; incorporating new material on transform domain, frequency domain, subband adaptive filters, acoustic echo cancellation and active noise control. • Provides an in-depth study of applications which now includes extensive coverage of OFDM, MIMO and smart antennas. • Contains exercises and computer simulation problems at the end of each chapter. • Includes a new companion website hosting MATLAB® simulation programs which complement the theoretical analyses, enabling the reader to gain an in-depth understanding of the behaviours and properties of the various adaptive algorithms.

The authoritative reference on the theory and design practice of computer arithmetic.

This book constitutes the refereed proceedings of the 10th IFIP WG 11.11 International Conference on Trust Management, IFIPTM 2016, held in Darmstadt, Germany, in July 2016. The 7 revised full papers and 7 short papers presented together with an invited paper were carefully reviewed and selected from 26 submissions. The papers cover a wide range of topics including trust architecture, trust modeling, trust metrics and computation, reputation and privacy, security and trust, sociotechnical aspects of trust, and attacks on trust and reputation systems.

Solving the Immigrant Church Crisis: The Biblical Solution of Parallel Ministry (Acts 6:1-7) addresses the crisis of the immigrant church in which complex cultural and linguistic factors create a reticence on the part of immigrants to transfer financial and decision-making authority to succeeding generations, and this results in a culturally irrelevant ministry to those generations, an exodus of believers from the church, a spiritually immature remnant, and an inability to reach the lost. The thesis of this book is that parallel ministry, based on Acts 6:1-7, is the biblical solution to the crisis in the immigrant church. While there are at least two main aspects of this crisis, a spiritual-relational and an ecclesiastical aspect, this book focuses on the ecclesiastical aspect of defining the biblical structure of church government. Specifically, this book is for immigrant churches primarily in the United States and offers them a biblical and practical solution to the problem plaguing them for over two centuries of how to minister effectively to the succeeding generations.

Introduction to Parallel Algorithms and Architectures: Arrays Trees Hypercubes provides an introduction to the expanding field of parallel algorithms and architectures. This book focuses on parallel computation involving the most popular network architectures, namely, arrays, trees, hypercubes, and some closely related networks. Organized into three chapters, this book begins

with an overview of the simplest architectures of arrays and trees. This text then presents the structures and relationships between the dominant network architectures, as well as the most efficient parallel algorithms for a wide variety of problems. Other chapters focus on fundamental results and techniques and on rigorous analysis of algorithmic performance. This book discusses as well a hybrid of network architecture based on arrays and trees called the mesh of trees. The final chapter deals with the most important properties of hypercubes. This book is a valuable resource for readers with a general technical background.

This revised third edition presents the subject with the help of learning objectives (LO) guided by Bloom's Taxonomy and supports outcome-based learning. It discusses concepts from elementary to advanced levels with focus on mathematical preliminaries. Numerous solved examples, algorithms, illustrations & usage of fictitious characters make the text interesting and simple to read. Salient Features: Dedicated section on Elementary Mathematics Pseudo codes used to illustrate implementation of algorithm Includes new topics on Shannon's theory and Perfect Secrecy, Unicity Distance and Redundancy of Language Interesting elements introduced through QR codes - Solutions to select chapter-end problems (End of every chapter) - 19 Proofs of theorems (Appendix Q) - Secured Electronic Transaction (Appendix R) Enhanced Pedagogical Features: - Solved Examples: 260 - Exercises: 400 - Review Questions: 200 - Illustration: 400

This second edition of Distributed Systems, Principles & Paradigms, covers the principles, advanced concepts, and technologies of distributed systems in detail, including: communication, replication, fault tolerance, and security. Intended for use in a senior/graduate level distributed systems course or by professionals, this text systematically shows how distributed systems are designed and implemented in real systems.

This book provides an overview of the X Window System focusing on characteristics that have significant impact on the development of both application programs and widgets. We pay special attention to applications that go beyond graphical user interfaces (GUIs); therefore we discuss issues affecting video games, visualization and imaging programs, and designing widgets with a complex appearance. While the book does not assume previous knowledge of X, it is intended for experienced programmers, especially those who want to write programs that go beyond simple GUIs. X is the dominant window system under Unix, and X servers are available for Microsoft Windows, thus enabling graphics over a network in the PC world. While Java offers an apparently universal graphics library (the abstract window toolkit), the reality is quite different: For high-quality graphics and image display, we must program on the target platform itself (X or one of Microsoft's APIs) rather than rely on Java peer objects. X is a vast subject, so it is impossible to provide a complete coverage in a few hundred pages. Thus we selected topics that are fundamental to the system, so that the reader who masters them should be able to read the documentation of the numerous libraries and toolkits. Therefore we provide documentation on the most important Xlib and X toolkit functions only.

This three volume set (CCIS 1237-1239) constitutes the proceedings of the 18th International Conference on Information Processing and Management of Uncertainty in Knowledge-Based Systems, IPMU 2020, in June 2020. The conference was scheduled to take place in Lisbon, Portugal, at University of Lisbon, but due to COVID-19 pandemic it was held virtually. The 173 papers were carefully reviewed and selected from 213 submissions. The papers are organized in topical sections: homage to Enrique Ruspini; invited talks; foundations and mathematics; decision making, preferences and votes; optimization and uncertainty; games; real world applications; knowledge processing and creation; machine learning I; machine learning II; XAI; image processing; temporal data processing; text analysis and processing; fuzzy interval analysis; theoretical and applied aspects of imprecise probabilities; similarities in artificial intelligence; belief function theory and its applications; aggregation: theory and practice; aggregation: pre-aggregation functions and other generalizations of monotonicity; aggregation: aggregation of different data structures; fuzzy methods in data mining and knowledge discovery; computational intelligence for logistics and transportation problems; fuzzy implication functions; soft methods in statistics and data analysis; image understanding and explainable AI; fuzzy and generalized quantifier theory; mathematical methods towards dealing with uncertainty in applied sciences; statistical image processing and analysis, with applications in neuroimaging; interval uncertainty; discrete models and computational intelligence; current techniques to model, process and describe time series; mathematical fuzzy logic and graded reasoning models; formal concept analysis, rough sets, general operators and related topics; computational intelligence methods in information modelling, representation and processing.

The powertrain is at the heart of vehicle design; the engine – whether it is a conventional, hybrid or electric design – provides the motive power, which is then managed and controlled through the transmission and final drive components. The overall powertrain system therefore defines the dynamic performance and character of the vehicle. The design of the powertrain has conventionally been tackled by analyzing each of the subsystems individually and the individual components, for example, engine, transmission and driveline have received considerable attention in textbooks over the past decades. The key theme of this book is to take a systems approach – to look at the integration of the components so that the whole powertrain system meets the demands of overall energy efficiency and good drivability. Vehicle Powertrain Systems provides a thorough description and analysis of all the powertrain components and then treats them together so that the overall performance of the vehicle can be understood and calculated. The text is well supported by practical problems and worked examples. Extensive use is made of the MATLAB(R) software and many example programmes for vehicle calculations are provided in the text. Key features: Structured approach to explaining the fundamentals of powertrain engineering Integration of powertrain components into overall vehicle design Emphasis on practical vehicle design issues Extensive use of practical problems and worked examples Provision of MATLAB(R) programmes for the reader to use in vehicle performance calculations This comprehensive and integrated analysis of vehicle powertrain engineering provides an invaluable resource for undergraduate and postgraduate automotive engineering students and is a useful reference for practicing engineers in the vehicle industry

Don't engineer by coincidence-design it like you mean it! Filled with practical techniques, Design It! is the perfect introduction to software architecture for programmers who are ready to grow their design skills. Lead your team as a software architect, ask the right stakeholders the right questions, explore design options, and help your team implement a system that promotes the right -ilities. Share your design decisions, facilitate collaborative design workshops that are fast, effective, and fun-and develop more awesome software! With dozens of design methods, examples, and practical know-how, Design It! shows you how to become a software architect. Walk through the core concepts every architect must know, discover how to apply them, and learn a variety of skills that will make you a better programmer, leader, and designer. Uncover the big ideas behind software architecture and gain confidence working on projects big and small. Plan, design, implement, and evaluate software architectures and collaborate with your team, stakeholders, and other architects. Identify the right stakeholders and understand their needs, dig for architecturally significant requirements, write amazing quality attribute scenarios, and make confident decisions. Choose technologies based on their architectural impact, facilitate architecture-centric design workshops, and evaluate architectures using lightweight, effective methods. Write lean architecture descriptions people love to read. Run an architecture design studio, implement the architecture you've designed, and grow your team's architectural knowledge. Good design requires good communication. Talk about your software architecture with stakeholders using whiteboards, documents, and code, and apply architecture-focused design methods in your day-to-day practice. Hands-on exercises, real-world scenarios, and practical team-based decision-making tools will get everyone on board and give you the experience you need to become a confident software architect.

Neural Networks for Perception, Volume 1: Human and Machine Perception focuses on models for understanding human perception in terms of distributed computation and examples of PDP models for machine perception. This book addresses both theoretical and practical issues related to the feasibility of both explaining human perception and implementing machine perception in terms of neural network models. The book is organized into two parts. The first part focuses on human perception. Topics on network model of object recognition in human vision, the self-organization of functional architecture in the cerebral cortex, and the structure and interpretation of neuronal codes in the visual system are detailed under this part. Part two covers the relevance of neural networks for machine perception. Subjects considered under this section include the multi-dimensional linear lattice for Fourier and Gabor transforms, multiple- scale Gaussian filtering, and edge detection; aspects of invariant pattern and object recognition; and neural network for motion processing. Neuroscientists, computer scientists, engineers, and researchers in artificial intelligence will find the book useful.

Despite five decades of research, parallel computing remains an exotic, frontier technology on the fringes of mainstream computing. Its much-heralded triumph over sequential computing has yet to materialize. This is in spite of the fact that the processing needs of many signal processing applications continue to eclipse the capabilities of sequential computing. The culprit is largely the software development environment. Fundamental shortcomings in the development environment of many parallel computer architectures thwart the adoption of parallel computing. Foremost, parallel computing has no unifying model to accurately predict the execution time of algorithms on parallel architectures. Cost and scarce programming resources prohibit deploying multiple algorithms and partitioning strategies in an attempt to find the fastest solution. As a consequence, algorithm design is largely an intuitive art form dominated by practitioners who specialize in a particular computer architecture. This, coupled with the fact that parallel computer architectures rarely last more than a couple of years, makes for a complex and challenging design environment. To navigate this environment, algorithm designers need a road map, a detailed procedure they can use to efficiently develop high performance, portable parallel algorithms. The focus of this book is to draw such a road map. The Parallel Algorithm Synthesis Procedure can be used to design reusable building blocks of adaptable, scalable software modules from which high performance signal processing applications can be constructed. The hallmark of the procedure is a semi-systematic process for introducing parameters to control the partitioning and scheduling of computation and communication. This facilitates the tailoring of software modules to exploit different configurations of multiple processors, multiple floating-point units, and hierarchical memories. To showcase the efficacy of this procedure, the book presents three case studies requiring various degrees of optimization for parallel execution.

The Encyclopedia of Big Data Technologies provides researchers, educators, students and industry professionals with a comprehensive authority over the most relevant Big Data Technology concepts. With over 300 articles written by worldwide subject matter experts from both industry and academia, the encyclopedia covers topics such as big data storage systems, NoSQL database, cloud computing, distributed systems, data processing, data management, machine learning and social technologies, data science. Each peer-reviewed, highly structured entry provides the reader with basic terminology, subject overviews, key research results, application examples, future directions, cross references and a bibliography. The entries are expository and tutorial, making this reference a practical resource for students, academics, or professionals. In addition, the distinguished, international editorial board of the encyclopedia consists of well-respected scholars, each developing topics based upon their expertise.

Were the thirteen essays Michel Foucault wrote in 1978–1979 endorsing the Iranian Revolution an aberration of his earlier work or an inevitable pitfall of his stance on Enlightenment rationality, as critics have long alleged? Behrooz Ghamari-Tabrizi argues that the critics are wrong. He declares that Foucault recognized that Iranians were at a threshold and were considering if it were possible to think of dignity, justice, and liberty outside the cognitive maps and principles of the European Enlightenment. Foucault in Iran centers not only on the significance of the great thinker's writings on the revolution but also on the profound mark the event left on his later lectures on ethics, spirituality, and fearless speech. Contemporary events since 9/11, the War on Terror, and the Arab Uprisings have made Foucault's essays on the Iranian Revolution more relevant than ever. Ghamari-Tabrizi illustrates how Foucault saw in the revolution an instance of his antiteleological philosophy: here was an event that did not fit into the normative progressive

discourses of history. What attracted him to the Iranian Revolution was precisely its ambiguity. Theoretically sophisticated and empirically rich, this interdisciplinary work will spark a lively debate in its insistence that what informed Foucault's writing was not an effort to understand Islamism but, rather, his conviction that Enlightenment rationality has not closed the gate of unknown possibilities for human societies.

Designed as an introductory text for the students of computer science, computer applications, electronics engineering and information technology for their first course on the organization and architecture of computers, this accessible, student friendly text gives a clear and in-depth analysis of the basic principles underlying the subject. This self-contained text devotes one full chapter to the basics of digital logic. While the initial chapters describe in detail about computer organization, including CPU design, ALU design, memory design and I/O organization, the text also deals with Assembly Language Programming for Pentium using NASM assembler. What distinguishes the text is the special attention it pays to Cache and Virtual Memory organization, as well as to RISC architecture and the intricacies of pipelining. All these discussions are climaxed by an illuminating discussion on parallel computers which shows how processors are interconnected to create a variety of parallel computers. KEY FEATURES ? Self-contained presentation starting with data representation and ending with advanced parallel computer architecture. ? Systematic and logical organization of topics. ? Large number of worked-out examples and exercises. ? Contains basics of assembly language programming. ? Each chapter has learning objectives and a detailed summary to help students to quickly revise the material.

This book describes warehouse-scale computers (WSCs), the computing platforms that power cloud computing and all the great web services we use every day. It discusses how these new systems treat the datacenter itself as one massive computer designed at warehouse scale, with hardware and software working in concert to deliver good levels of internet service performance. The book details the architecture of WSCs and covers the main factors influencing their design, operation, and cost structure, and the characteristics of their software base. Each chapter contains multiple real-world examples, including detailed case studies and previously unpublished details of the infrastructure used to power Google's online services. Targeted at the architects and programmers of today's WSCs, this book provides a great foundation for those looking to innovate in this fascinating and important area, but the material will also be broadly interesting to those who just want to understand the infrastructure powering the internet. The third edition reflects four years of advancements since the previous edition and nearly doubles the number of pictures and figures. New topics range from additional workloads like video streaming, machine learning, and public cloud to specialized silicon accelerators, storage and network building blocks, and a revised discussion of data center power and cooling, and uptime. Further discussions of emerging trends and opportunities ensure that this revised edition will remain an essential resource for educators and professionals working on the next generation of WSCs.

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.

This book constitutes the refereed proceedings of the 6th International Conference on Algorithms and Architectures for Parallel Processing, ICA3PP 2005, held in Melbourne, Australia in October 2005. The 27 revised full papers and 25 revised short papers presented were carefully reviewed and selected from 95 submissions. The book covers new architectures of parallel and distributed systems, new system management facilities, and new application algorithms with special focus on two broad areas of parallel and distributed computing, i.e., architectures, algorithms and networks, and systems and applications.

Introduction to Parallel Processing Algorithms and Architectures Springer Science & Business Media

This book focuses on the future directions of the static scheduling and dynamic load balancing methods in parallel and distributed systems. It provides an overview and a detailed discussion of a wide range of topics from theoretical background to practical, state-of-the-art scheduling and load balancing techniques.

Multiple processor systems are an important class of parallel systems. Over the years, several architectures have been proposed to build such systems to satisfy the requirements of high performance computing. These architectures span a wide variety of system types. At the low end of the spectrum, we can build a small, shared-memory parallel system with tens of processors. These systems typically use a bus to interconnect the processors and memory. Such systems, for example, are becoming commonplace in high-performance graphics workstations. These systems are called uniform memory access (UMA) multiprocessors because they provide uniform access of memory to all processors. These systems provide a single address space, which is preferred by programmers. This architecture, however, cannot be extended even to medium systems with hundreds of processors due to bus bandwidth limitations. To scale systems to medium range i. e. , to hundreds of processors, non-bus interconnection networks have been proposed. These systems, for example, use a multistage dynamic interconnection network. Such systems also provide global, shared memory like the UMA systems. However, they introduce local and remote memories, which lead to non-uniform memory access (NUMA) architecture. Distributed-memory architecture is used for systems with thousands of processors. These systems differ from the shared-memory architectures in that there is no globally accessible shared memory. Instead, they use message passing to facilitate communication among the processors. As a result, they do not provide single address space.

You too can learn to design and develop classic arcade video games like Pong, Pac-Man, Space Invaders, and Scramble. Collision detection, extra lives, power-ups, and countless other essential design elements were invented by the mostly anonymous designers at the early pioneering companies that produced these great games. In this book you'll go step by step, using

modern, free software tools such as Unity3D to create five games in the classic style, inspired by these classics: Pong, Breakout, Space Invaders, Scramble, and Pac-Man. All the source code, art and sound sources for the projects are freely available on the companion DVD or at the book's Web site. You'll discover the fun of making your own games, putting in your own color graphics, adjusting the scoring, coding the AI, and creating the sound effects. You'll gain a deep understanding of the roots of modern video game design: the classics of the seventies and eighties. Features: * Uses seven Unity3D projects to allow for quick experimentation with classic game concepts * 4-color throughout with companion DVD that includes source code, art, and full projects * Includes historical anecdotes direct from one of the fabled Atari coin-op programmers * Detailed step-by-step instructions, dozens of exercises, and eight rules of classic game design * Contains unique insights on applying classic game design concepts to modern games

The book presents high-quality research papers presented at the first international conference, ICICCD 2016, organised by the Department of Electronics, Instrumentation and Control Engineering of University of Petroleum and Energy Studies, Dehradun on 2nd and 3rd April, 2016. The book is broadly divided into three sections: Intelligent Communication, Intelligent Control and Intelligent Devices. The areas covered under these sections are wireless communication and radio technologies, optical communication, communication hardware evolution, machine-to-machine communication networks, routing techniques, network analytics, network applications and services, satellite and space communications, technologies for e-communication, wireless Ad-Hoc and sensor networks, communications and information security, signal processing for communications, communication software, microwave informatics, robotics and automation, optimization techniques and algorithms, intelligent transport, mechatronics system, guidance and navigation, algorithms, linear/non-linear control, home automation, sensors, smart cities, control systems, high performance computing, cognition control, adaptive control, distributed control, prediction models, hybrid control system, control applications, power system, manufacturing, agriculture cyber physical system, network control system, genetic control based, wearable devices, nano devices, MEMS, bio-inspired computing, embedded and real-time software, VLSI and embedded systems, FPGA, digital system and logic design, image and video processing, machine vision, medical imaging, and reconfigurable computing systems.

An introduction to the principles and applications of passive seismic monitoring, providing an accessible overview of current research and technology.

Spinal disorders in very young children may be caused by a variety of conditions. The treatment of such conditions is often challenging due to the age of the patient and the progressive nature of the deformity. There also may be associated problems such as congenital anomalies, respiratory insufficiency, and neurological problems. Depending on the etiology of the deformity, these children are often cared for by multiple specialists including pediatricians, pediatric orthopaedists or orthopaedic spine surgeons, neurologists, pediatric surgeons, pediatric neurosurgeons, oncologists, and/or pulmonologists. Health professionals in all of the mentioned disciplines are involved in the management of these patients, which is why compiling a comprehensive textbook that is not limited to orthopedic specialists is essential. This textbook will effectively help to standardize the care of these patients. Furthermore, other professionals such as nurses, physical therapists and healthcare professionals in training are usually not familiar with these conditions and are in need of a reference book to consult when caring for children with spinal deformities.

This revised edition provides patient guidance in its clear and organized presentation of problems. It is rich in variety, large in number and provides very careful treatment of relativity. One outstanding feature is the inclusion of simple, standard examples demonstrated in different methods that will allow students to enhance and understand their calculating abilities. There are over 145 worked examples; virtually all of the standard problems are included.

It is our pleasure to welcome you to the proceedings of the 13th International Computer Society of Iran Computer Conference (CSICC-2008). The conference has been held annually since 1995, except for 1998, when it transitioned from a year-end to first-quarter schedule. It has been moving in the direction of greater selectivity (see Fig.1) and broader international participation. Holding it in Kish Island this year represents an effort to further facilitate and encourage international contributions. We feel privileged to participate in further advancing this strong technical tradition.

Year	Venue
1995	U of Tehran
1996	U of Tehran
1997	U of Tehran
1999	U of Tehran
2000	U of Tehran
2001	U of Tehran
2002	U of Tehran
2003	U of Tehran
2004	U of Tehran
2005	U of Tehran
2006	IPM, Tehran
2007	Sharif U, Tehran
2008	Sharif U, Tehran

With the advent of portable and autonomous computing systems, power consumption has emerged as a focal point in many research projects, commercial systems and DoD platforms. One current research initiative, which drew much attention to this area, is the Power Aware Computing and Communications (PAC/C) program sponsored by DARPA. Many of the chapters in this book include results from work that have been supported by the PACIC program. The performance of computer systems has been tremendously improving while the size and weight of such systems has been constantly shrinking. The capacities of batteries relative to their sizes and weights has been also improving but at a rate which is much slower than the rate of improvement in computer performance and the rate of shrinking in computer sizes. The relation between the power consumption of a computer system and its performance and size is a complex one which is very much dependent on the specific system and the technology used to build that system. We do not need a complex argument, however, to be convinced that energy and power, which is the rate of energy consumption, are becoming critical components in computer systems in general, and portable and autonomous systems, in particular. Most of the early research on power consumption in computer systems addressed the issue of minimizing power in a given platform, which usually translates into minimizing energy consumption, and thus, longer battery life.

[Copyright: 1aa4325f4bf21adbfe84039c086a2e76](http://www.1aa4325f4bf21adbfe84039c086a2e76)