

Kai Hwang Parallel Processing Solution Tianqior

The scope of the present book is to offer the most efficient tools for the vectorization of serial computer programs. Here, by vectorization we understand the adaptation of computer programs to the special architecture of modern available vector computers to exploit fully their potential, which will often result in remarkable performance improvements. The book is written primarily for users working in the various fields of computational physics, for scientists as well as for programmers running their jobs on a vector computer. The text may, however, also be of value to those who are interested in numerical algorithms. Although the examples discussed in chapter 9 have been taken from Computational Fluid Dynamics, the numerical methods are well-known, and are applied in many fields of Computational Physics. The book is divided into four parts. After a short introduction which outlines the limits of conventional serial computers in contrast to the possibilities offered by the new vector machines, the second part is addressed to the discussion of some main features of existing computer architectures. We restrict ourselves to the vector computers CRAY-IS and CDC-CYBER 205, although, in the meantime, many vector and parallel computers and array processors are available such as BBN/COR's Heterogeneous Element Processor (HEP), TCI's Distributed Array Processor (DAP), SPERRY UNIVAC's Array Processor (APS), STAR TECHNOLOGIES SP-100, FLOATING POINT SYSTEMS' Array Processor (FPS), FUJITSU's FACOM VP-100 and VP-200, HITACHI's Integrated Array Processor (IAP), HITACHI's S 810/10 and S 810/20 and others.

The book, in addition, includes a balanced survey of the fundamentals of artificial intelligence, emphasizing the relationship between symbolic and numeric processing. The text is structured around an innovative, interactive combination of LISP programming and AI; it uses the constructs of the programming language to help readers understand the array of artificial intelligence concepts presented. After an overview of the field of artificial intelligence, the text presents the fundamentals of LISP, explaining the language's features in more detail than any other AI text. Common Lisp is then used consistently, in both programming exercises and plentiful examples of actual AI code. - Back cover This text is intended to provide an introduction to both AI and LISP for those having a background in computer science and mathematics. -Pref. This book, first published in 1990, analyses how to train end-users to search with both natural language and controlled vocabularies in the sciences, describes a planning assessment for implementing end-user searching in a sci-tech organization, examines how the scientists at a major industrial research organization have begun to do more online searching with the encouragement of the information center, and explores the proactive role that medical libraries have taken in training health care professionals to search MEDLINE.

Proceedings of the . . . International Conference on Parallel and Distributed Information Systems

Parallel Processing for Supercomputers and Artificial Intelligence

Advances in Petri Nets 1987

August 23-26, 1983

End-User Training for Sci-Tech Databases

Proceedings of the 1993 International Conference on Parallel Processing

Proceedings - Miscellaneous

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.

Over the years, the promise of artificial intelligence has inspired many researchers and many schemes, only to have incipient hopes thwarted by its complexity. With each generation of computational engines, a new wave of enthusiasm sweeps the community as solutions to a few problems come within reach. However, intractability and undecidability continue to frustrate the unwary practitioner, while unsubstantiated methodologies offer ingenious solutions that hold more promise than potential. Despite its undulate past and variegated present, AI has made solid contributions to a growing information technology. Expert systems and allied tools have become a mainstay of industrial and business organizations. Intelligent interfaces have increased accessibility of computational resources, and robotic innovations have redefined the manufacturing industries. Meanwhile, research in evolutionary algorithms, neural networks, fuzzy reasoning, and other exciting approaches promise continued progress in surprising new directions. These proceedings record the latest results of industrial, commercial, military, and academic artificial intelligence exploration. Seventy-seven papers divided into twenty different areas document a significant slice of this broad and exciting field. Although dozens of themes are treated in the papers, the topical divisions of this volume comprise: The Software Engineering/AI Interface, Knowledge-Based Systems, Temporal Reasoning, Machine Learning, Robotics, Intelligent Databases, Planning, Expert Systems Applications, Search Techniques, Genetic and Evolutionary Methods, Design, Qualitative Reasoning, Neural Networks, Knowledge Representation, Application Paradigms, Fuzzy and Pattern Recognition, Reasoning about Physical Systems, Parallel and Distributed AI, and Diagnostic Systems.

Proceedings of the International Conference, Leiden, 29 August-1 September 1989

Parallelism operational research and its applications

Big-Data Analytics for Cloud, IoT and Cognitive Computing

Distributed and Cloud Computing

Vectorization of Computer Programs with Applications to Computational Fluid Dynamics

This book is the result of the first conference organized under the auspices of the Parallel Computing Society, the only independent international forum where researchers, manufacturers and users of parallel computers can exchange expertise and experience in the development and use of all types of systems. The emphasis is on applications of parallel computing. In addition to papers discussing the development of new algorithms for, and applications of, various types of parallel computers, an overview of the newest developments regarding parallel computer hardware and software is given.

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 datflow computation. After introducing interconnection networks, it discusses parallel algorithms for sorting, Fourier transform, matrix algebra, and graph theory. The second part focuses on 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.

"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."

International Aerospace Abstracts

Networks on Chip

6th International Conference, Calcutta, India, December 17-20, 1999 Proceedings

Encyclopedia of Microcomputers

Volume 13 - Optical Disks to Production Scheduling

Advanced Computer Architecture

What does Google's management of billions of Web pages have in common with analysis of a genome with billions of nucleotides? Both apply methods that coordinate many processors to accomplish a single task. From mining genomes to the World Wide Web, from modeling financial markets to global weather patterns, parallel computing enables computations that would otherwise be impractical if not impossible with sequential approaches alone. Its fundamental role as an enabler of simulations and data analysis continues an advance in a wide range of application areas. Scientific Parallel Computing is the first textbook to integrate all the fundamentals of parallel computing in a single volume while also providing a basis for a deeper understanding of the subject. Designed for graduate and advanced undergraduate courses in the sciences and in engineering, computer science, and mathematics, it focuses on the three key areas of algorithms, architecture, languages, and their crucial synthesis in performance. The book's computational examples, whose math prerequisites are not beyond the level of advanced calculus, derive from a breadth of topics in scientific and engineering simulation and data analysis. The programming exercises presented early in the book are designed to bring students up to speed quickly, while the book later develops projects challenging enough to guide students toward research questions in the field. The new paradigm of cluster computing is fully addressed. A supporting web site provides access to all the codes and software mentioned in the book, and offers topical information on popular parallel computing systems. Integrates all the fundamentals of parallel computing essential for today's high-performance requirements Ideal for graduate and advanced undergraduate students in the sciences and in engineering, computer science, and mathematics Extensive programming and theoretical exercises enable students to write parallel codes quickly More challenging projects later in the book introduce research questions New paradigm of cluster computing fully addressed Supporting web site provides access to all the codes and software mentioned in the book

The four-volume set LNCS 3991-3994 constitutes the refereed proceedings of the 6th International Conference on Computational Science, ICCS 2006, held in Reading, UK, in May 2006. The main conference and its 32 topical workshops attracted over 1400 submissions. The 98 revised full papers and 29 revised poster papers of the main track presented together with 500 accepted workshop papers were carefully reviewed and selected for inclusion in the four volumes. The papers span the whole range of computational science with focus on the following major themes: tackling grand challenges problems, modelling and simulations of complex systems, scalable algorithms and tools and environments for computational science. Of particular interest were the following major recent developments in novel methods and modelling of complex systems for diverse areas of science, scalable scientific algorithms, advanced software tools, computational grids, advanced numerical methods, and novel application areas where the above novel models, algorithms and tools were efficiently applied such as physical systems, computational and systems biology, environmental and systems biology, financial systems, finance, and others.

As the number of processor cores and IP blocks integrated on a single chip is steadily growing, a systematic approach to design the communication infrastructure becomes necessary. Different variants of packed switched on-chip networks have been proposed by several groups during the past two years. This book summarizes the state of the art of these efforts and discusses the major issues from the physical integration to architecture to operating systems and application interfaces. It also provides a guideline and vision about the direction this field is moving to. Moreover, the book outlines the consequences of adopting design platforms based on packet switched network. The consequences may in fact be far reaching because many of the topics of distributed systems, distributed real-time systems, fault tolerant systems, parallel computer architecture, parallel programming as well as traditional system-on-chip issues will appear relevant but within the constraints of a single chip VLSI implementation.

Parallel Architectures for Data/Knowledge-based Systems

Artificial Intelligence with Common Lisp

Euro-Par 2002. Parallel Processing

Computational Science - ICCS 2006

Parallel Computing 89

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

By the end of the 1960s, a new discipline named computer science had come into being. A new scientific paradigm—the ‘computational paradigm’—was in place, suggesting that computer science had reached a certain level of maturity. Yet as a science it was still precociously young. New forces, some technological, some socio-economic, some cognitive impinged upon it, the outcome of which was that new kinds of computational problems arose over the next two decades. Indeed, by the beginning of the 1990’s the structure of the computational paradigm looked markedly different in many important respects from how it was at the end of the 1960s. Author Subrata Dasgupta named the two decades from 1970 to 1990 as the second age of computer science to distinguish it from the preceding genesis of the science and the age of the Internet/World Wide Web that followed. This book describes the evolution of computer science in this second age in the form of seven overlapping, intermingling, parallel histories that unfold concurrently in the course of the two decades. Certain themes characteristic of this second age thread through this narrative: the desire for a genuine science of computing; the realization that computing is as much a human experience as it is a technological one; the search for a unified theory of intelligence spanning machines and mind; the desire to liberate the computational mind from the shackles of sequentiality; and, most ambitiously, a quest to subvert the very core of the computational paradigm itself. We see how the computer scientists of the second age address these desires and challenges, in what manner they succeed or fail and how, along the way, the shape of computational paradigm was altered. And to complete this history, the author asks and seeks to answer the question of how computer science shows evidence of progress over the course of its second age.

This book contains four review articles in the area of scalable computing. Two of the articles discuss methods and tools for the parallel solution of irregular problems, which have been satisfactorily worked out in heterogeneous systems. One surveys the technology and applications of multimedia server clusters, which are playing an increasing role in the current networked environment. An additional article discusses SilkRoad, which adds distributed shared memory capabilities to the Cilk parallel programming system. Once again, the book represents a new set of steps forward in parallel systems. Contents:Parallel Computing Strategies for Irregular AlgorithmsA Runtime Support for Large-Scale Irregular Computing on Clusters and GridsMemory Model Support for Mixed Programming Paradigm in SilkRoadClustered Multimedia Servers: Architectures and Storage Systems Readership: Graduate students, academics and researchers in supercomputing and computer engineering. Keywords:Clusters;Distributed Shared Memory;Heterogeneous Systems;Irregular Problems;Multimedia Servers;Parallel Processing;Scalable Computing

25 Years of the International Symposia on Computer Architecture

Annual Review of Scalable Computing

UPC

Instructor's Solutions Manual to Accompany Scaladle Parallel Computing, Technology, Architecture and Programming [by] Kai Hwang, Zhiwei Xu

Scientific Parallel Computing

PARALLEL AND DISTRIBUTED COMPUTING : ARCHITECTURES AND ALGORITHMS

The first textbook to teach students how to build data analytic solutions on large data sets (specifically in Internet of Things applications) using cloud-based technologies for data storage, transmission and mashup, and AI techniques to analyze this data. This textbook is designed to train college students to master modern cloud computing systems in operating principles, architecture design, machine learning algorithms, programming models and software tools for big data mining, analytics, and cognitive applications. The book will be suitable for use in one-semester computer science or electrical engineering courses on cloud computing, machine learning, cloud programming, cognitive computing, or big data science. The book will also be very useful as a reference for professionals who want to work in cloud computing and data science. Cloud and Cognitive Computing begins with two introductory chapters on fundamentals of cloud computing, data science, and adaptive computing that lay the foundation for the rest of the book. Subsequent chapters cover topics including cloud architecture, mashup services, virtual machines, Docker containers, mobile clouds, IoT and AI, inter-cloud mashups, and cloud performance and benchmarks, with a focus on Google's Brain Project, DeepMind, and X-Lab programs, IBM Kai Hwang/M SyNapSe, Bluerim programs, cognitive initiatives, and neurocomputers. The book then covers machine learning algorithms and cloud programming software tools and application development, applying the tools in machine learning, social media, deep learning, and cognitive applications. All cloud systems are illustrated with big data and cognitive application examples.

These are the proceedings of the Sixth International Conference on High Performance Computing (HPC'99) held December 17-20 in Calcutta, India. The meeting serves as a forum for presenting current work by researchers from around the world, as well as highlighting activities in Asia in the high performance computing area. The meeting emphasizes both the design and the analysis of high performance computing systems and their scientific, engineering, and commercial applications. Topics covered in the meeting series include: Parallel Algorithms; Scientific Computation; Parallel Architectures; Visualization; Parallel Languages & Compilers; Network and Cluster Based Computing; Distributed Systems; Signal & Image Processing; Systems; Programming; Environments; Supercomputing; Applications; Memory; Systems; Internet and High Speed Networks; Scalable Servers. We would like to thank Alfred Hofmann and Ruth Abraham of Springer-Verlag for their excellent support in bringing out the proceedings. The detailed messages from the steering committee chair, general co-chair and program chair pay tribute to numerous volunteers who helped us in organizing the meeting. October 1999 Viktor K. Prasanna Bhabani Sinha Prithviraj Banerjee Message from the Steering Chair It is my pleasure to welcome you to the Sixth International Conference on High Performance Computing. I hope you enjoy the meeting, the rich cultural heritage of Calcutta, as well as the mother Ganges, "the river of life". The aim of these series of volumes "Advances in Petri Nets" is to present to the general computer science community the most significant recent results with regard to the development in the area. The main source of the papers are the annual European Workshops on Applications and Theory of Petri Nets"; the highest ranked papers from the past workshops are considered for the series, i.e., they are again reviewed and accordingly revised or extended. In addition to the workshop papers, the "Advances" also present invited papers. The present volume Advances in Petri Nets 1987 covers the 7th "European Workshop on Applications and Theory of Petri Nets" held in Oxford, Great Britain, in June 1986. It also contains a survey on complexity of problems related to Petri nets written by R.R. Howell and L.E. Rosier. A special feature of this volume is a bibliography on Petri nets, containing more than 2000 entries.

Analysis and Design of Scalable Parallel Algorithms for Scientific Computing

7th Int. Conf. Industrial & En

From Algal Genes to Neural Nets

Selected Papers

Computers for Artificial Intelligence Processing

Partitioning Strategy for Efficient Nonlinear Finite Element Dynamic Analysis on Multiprocessor Computers

Euro-Par – the European Conference on Parallel Computing – is an international conference series dedicated to the promotion and advancement of all aspects of parallel computing. The major themes can be divided into the broad categories of hardware, software, algorithms, and applications for parallel computing. The objective of Euro-Par is to provide a forum within which to promote the dev- opment of parallel computing both as an industrial technique and an academic discipline, extending the frontiers of both the state of the art and the state of the practice. This is particularly important at a time when parallel computing is undergoing strong and sustained development and experiencing real industrial take-up. The main audience for and participants in Euro-Par are researchers in academic departments, government laboratories, and industrial organizations. Euro-Par aims to become the primary choice of such professionals for the p- sentation of new results in their speci?c areas. Euro-Par is also interested in applications that demonstrate the e?ctiveness of the main Euro-Par themes. Euro-Par has its own Internet domain with a permanent website where the history of the conference series is described: http://www. euro-par. org. The Euro-Par conference series is sponsored by the Association of Computer - chinery and the International Federation of Information Processing. Euro-Par 2002 at Paderborn, Germany Euro-Par 2002 was organized by the Paderborn Center for Parallel Comput- 2.2 Ing (PC) and was held at the Heinz Nixdorf MuseumsForum (HNF).

The definitive guide to successfully integrating social, mobile, Big-Data analytics, cloud and IoT principles and technologies The main goal of this book is to spur the development of effective big-data computing operations on smart clouds that are fully supported by IoT sensing, machine learning and analytics systems. To that end, the authors draw upon their original research and proven track record in the field to describe a practical approach integrating big-data theories, cloud design principles, Internet of Things (IoT) sensing, machine learning, data analytics and Hadoop and Spark programming. Part 1 focuses on data science, the roles of clouds and IoT devices and frameworks for big-data computing. Big data analytics and cognitive machine learning, as well as cloud architecture, IoT and cognitive systems are explored, and mobile cloud-IoT-interaction frameworks are illustrated with concrete system design examples. Part 2 is devoted to the principles of and algorithms for machine learning, data analytics and deep learning in big data applications. Part 3 concentrates on cloud programming software libraries from MapReduce to Hadoop, Spark and TensorFlow and describes business, educational, healthcare and social media applications for those tools. The first book describing a practical approach to integrating social, mobile, analytics, cloud and IoT (SMACT) principles and technologies Covers theory and computing techniques and technologies, making it suitable for use in both computer science and electrical engineering programs Offers an extremely well-informed vision of future intelligent and cognitive computing environments Integrating SMACT technologies Fully illustrated throughout with examples, figures and approximately 150 problems to support and reinforce learning Features a companion website with an instructor manual and PowerPoint slides www.wiley.com/go/hwangIoT Big-Data Analytics for Cloud, IoT and Cognitive Computing satisfies the demand among university faculty and students for cutting-edge information on emerging intelligent and cognitive computing systems and technologies. Professionals working in data science, cloud computing and IoT applications will also find this book to be an extremely useful working resource.

The present book supports the increasing complexity and the growing need for computational power of artificial intelligence (AI) by providing comprehensive treatments of new hardware and software engineering met in AI language design and applications. The book is a collection of 16 substantial papers (chapters), the contributors being 51 well-known researchers in the AI fields. The papers are grouped into the following five sections: Section 1 represents a well documented survey on symbolic processing computers. Section 2 (Language-based AI Architectures) supports the design and implementation of AI language-oriented computers. Three (2-4) chapters are devoted to (computer architecture concerning) sequential Lisp processing; architectural features of Lisp computers, Symbolics' Lisp computer architecture, memory management and performance evaluation of a Lisp machine system. Other three (5-7) chapters discuss multiprocessing and parallel processing of Lisp (and general functional) programs. The last two chapters of section 2 are presenting architectures supporting object-oriented programming (Smalltalk) and production systems. Section 3 (Multiprocessor AI Architecture) contains two (10-11) chapters, dealing with Connection Machine architecture and its applications, and with the design of data/knowledge base machines for AI processing. Section 4 (Connectionist Architectures and Applications) include two (12-13) chapters, illustrating the connectionist model architecture design and learning. Section 5 (Software Architectures for AI Applications) is made up of three (14-16) chapters, analysing the relationship between AI and software engineering, the development tools for AI programs, and the problem of AI hardware and software reliability. This book addresses a wide range of AI readers, from beginners to professionals. It carries forth doubtless qualities: compact and well-dimensioned chapters, comprehensively written by AI remarkable professionals, covering up-to-date AI topics and trends.

Dissertation Abstracts International

From Parallel Processing to the Internet of Things

Cloud Computing for Machine Learning and Cognitive Applications

NASA Technical Paper

Proceedings: Industrial systems, supercomputer projects, and prototype architectures

Proceedings of the 1993 International Conference on Parallel Processing

Instructor's Solutions Manual to Accompany Scaladle Parallel Computing, Technology, Architecture and Programming [by] Kai Hwang, Zhiwei XuDistributed and Cloud ComputingFrom Parallel Processing to the Internet of ThingsMorgan Kaufmann

This is the first book to explain the language Unified Parallel Cand its use. Authors El-Ghazawi, Carlson, and Sterling are among the developers of UPC, with close links with the industrial members of the UPC consortium. Their text covers background material onparallel architectures and algorithms, and includes UPC programmingcase studies. This book represents an invaluable resource for thegrowing number of UPC users and applications developers. Moreinformation about UPC can be found at: http://upc.gwu.edu/ An Instructor Support FTP site is available from the Wileyeditorial department.

This three-volume work presents a compendium of current and seminal papers on parallel/distributed processing offered at the 22nd International Conference on Parallel Processing, held August 16-20, 1993 in Chicago, Illinois. Topics include processor architectures; mapping algorithms to parallel systems, performance evaluations; fault diagnosis, recovery, and tolerance; cube networks; portable software; synchronization; compilers; hypercube computing; and image processing and graphics. Computer professionals in parallel processing, distributed systems, and software engineering will find this book essential to complete their computer reference library.

Fundamentals of Symbolic and Numeric Processing

High Performance Computing - HIPC'99

The Second Age of Computer Science

Parallel Computing: Fundamentals, Applications and New Directions

The sciences and engineering. B

Parallelism, Scalability, Programmability

This book contains four review articles in the area of scalable computing. Two of the articles discuss methods and tools for the parallel solution of irregular problems, which have been satisfactorily worked out in heterogeneous systems. One surveys the technology and applications of multimedia server clusters, which are playing an increasing role in the current networked environment. An additional article discusses SilkRoad, which adds distributed shared memory capabilities to the Cilk parallel programming system. Once again, the book represents a new set of steps forward in parallel systems. Graduate students, academics and researchers in supercomputing and computer engineering.

8th International Euro-Par Conference Paderborn, Germany, August 27-30, 2002 Proceedings

6th International Conference, Reading, UK, May 28-31, 2006, Proceedings

Distributed Shared Memory Programming