

High Performance Computing Hipc 2007 14th International Conference Goa India December 18 21 2007 Proceedings Lecture Notes In Computer Computer Science And General Issues

Dark Silicon and the Future of On-chip Systems, Volume 110, the latest release in the Advances in Computers series published since 1960, presents detailed coverage of innovations in computer hardware, software, theory, design and applications, with this release focusing on an Introduction to dark silicon and future processors, a Revisiting of processor allocation and application mapping in future CMPs in the dark silicon era, Multi-objectivism in the dark silicon age, Dark silicon aware resource management for many-core systems, Dynamic power management for dark silicon multi-core processors, Topology specialization for networks-on-chip in the dark silicon era, and Emerging SRAM-based FPGA architectures. Provides in-depth surveys and tutorials on new computer technology Covers well-known authors and researchers in the field Presents extensive bibliographies with most chapters Includes volumes that are devoted to single themes or subfields of computer science, with this release focusing on Dark Silicon and Future On-chip Systems

Identifies Recent Technological Developments Worldwide The field of grid computing has made rapid progress in the past few years, evolving and developing in almost all areas, including concepts, philosophy, methodology, and usages. Grid Computing: Infrastructure, Service, and Applications reflects the recent advances in this field, covering the research aspects that involve infrastructure, middleware, architecture, services, and applications. Grid Systems Across the Globe The first section of the book focuses on infrastructure and middleware and presents several national and international grid systems. The text highlights China Research and Development environment Over Wide-area Network (CROWN), several ongoing cyberinfrastructure efforts in New York State, and Enabling Grids for E-science (EGEE), which is co-funded by the European Commission and the world's largest multidisciplinary grid infrastructure today. The second part of the book discusses recent grid service advances. The authors examine the UK National Grid Service (NGS), the concept of resource allocation in a grid environment, OMIIBPEL, and the possibility of treating scientific workflow issues using techniques from the data stream community. The book describes an SLA model, reviews portal and workflow technologies, presents an overview of PKIs and their limitations, and introduces PIndex, a peer-to-peer model for grid information services. New Projects and Initiatives The third section includes an analysis of innovative grid applications. Topics covered include the WISDOM initiative, incorporating flow-level networking models into grid simulators, system-level virtualization, grid usage in the high-energy physics environment in the LHC project, and the Service Oriented HLA RTI (SOHR) framework. With a comprehensive summary of past advances, this text is a window into the future of this nascent technology, forging a path for the next generation of cyberinfrastructure developers.

This book constitutes the thoroughly refereed post-conference proceedings of the 8th International Conference on High Performance Computing for Computational Science, VECPAR 2008, held in Toulouse, France, in June 2008. The 51 revised full papers presented together with the abstract of a surveying and look-ahead talk were carefully reviewed and selected from 73 submissions. The papers are organized in topical sections on parallel and distributed computing, cluster and grid computing, problem solving environment and data centric, numerical methods, linear algebra, computing in geosciences and biosciences, imaging and graphics, computing for aerospace and engineering, and high-performance data management in grid environments. This book gathers a selection of peer-reviewed papers presented at the Tiangong-2 Data Utilization Conference, which was held in Beijing, China, in December 2018. As the first space laboratory in China, Tiangong-2 carries 3 new types of remote sensing payloads – the Wide-band Imaging Spectrometer (WIS), Three-dimensional Imaging Microwave Altimeter (TIMA), and Multi-band Ultraviolet Edge Imaging Spectrometer (MUEIS) – for observing the Earth. The spectrum of the WIS covers 18 bands, from visible to thermal infrared, with a swath of 300km. The TIMA is the first-ever system to use interferometric imaging radar altimeter (InIRA) technology to measure sea surface height and land topography at near-nadir angles with a wide swath. In turn, the MUEIS is the world's first large-field atmospheric detector capable of quasi-synchronously detecting the characteristics of ultraviolet limb radiation in the middle atmosphere. The Earth observation data obtained by Tiangong-2 has attracted many research groups and been applied in such diverse areas as land resources, water resources, climate change, environmental monitoring, agriculture, forestry, ecology, oceanography, meteorology and so on. The main subjects considered in this proceedings volume include: payload design, data processing, data service and application. It also provides a comprehensive introduction to the research results gleaned by engineers, researchers and scientists throughout the lifecycle of the Tiangong-2 Earth observation data, which will improve the payload development and enhance remote sensing data applications.

This book constitutes the refereed proceedings of the 7th International Conference on High Performance Computing, HiPC 2000, held in Bangalore, India in December 2000. The 46 revised papers presented together with five invited contributions were carefully reviewed and selected from a total of 127 submissions. The papers are organized in topical sections on system software, algorithms, high-performance middleware, applications, cluster computing, architecture, applied parallel processing, networks, wireless and mobile communication systems, and large scale data mining.

Euro-Par is an annual series of international conferences dedicated to the promotion and advancement of all aspects of parallel computing. The major themes can be divided into the broad categories of theory, high-performance, cluster and grid, and distributed and mobile computing. These four categories comprise 14 topics that focus on particular issues of the mentioned categories. The objective of Euro-Par is to provide a forum within which to promote the development of parallel computing both as an industrial technique and an academic discipline, extending the frontier of both the state of the art and the state of practice. The main audience for and participants in Euro-Par are seen as researchers in academic departments, government laboratories and industrial organizations. Previous Euro-Par conferences took place in Stockholm, Lyon, Passau, Southampton, Toulouse, Munich, Manchester, Paderborn, Klagenfurt, Pisa, Lisbon, Dresden and Rennes. Next year the conference will take place in Delft. Euro-Par 2008 was the 14th conference in the Euro-Par series. It was jointly organized by the Computer Architecture and Operating Systems Department of the Universitat Autònoma of Barcelona and the University Institute for Intelligent Systems and Numerical Applications in Engineering of the Universidad de Las Palmas de Gran Canaria, at the Tafira Campus of Universidad de Las Palmas de Gran Canaria.

"This book provides a compendium of terms, definitions, and explanations of concepts, issues, and trends in grid technology"--Provided by publisher.

Parallel Programming with Co-Arrays describes the basic techniques used to design parallel algorithms for high-performance, scientific computing. It is intended for upper-level undergraduate students and graduate students who need to develop parallel codes with little or no previous introduction to parallel computing. It is also intended as a reference manual for researchers active in the field of scientific computing. All the algorithms in the book are based on partition operators. These operators provide a unifying principle that fits seemingly disparate techniques into an overall framework for algorithm design. The book uses the co-array programming model to illustrate how to write code for concrete examples, but it emphasizes that the important concepts for algorithm design are independent of the programming model. With these concepts in mind, the reader can write algorithms in different programming models based on personal taste and comfort.

Today all kinds of ubiquitous systems, led by wireless sensor networks, can be seen as an unprecedented privacy risk given their ability to collect information on quantities and situations so far unsuspected. There is therefore an urgent need to develop mechanisms to ensure privacy in sensor networks. Location Privacy in Wireless Sensor Networks focuses on location privacy, by which an attacker might determine the source and destination of communications with simple techniques. This poses a serious threat as the attacker might use this information to reach the assets or individuals being monitored or even to destroy or compromise the whole network. This book will aid in the protection against this serious privacy threat.

This book introduces a new cyberphysical system that combines clinical and basic neuroscience research with advanced data analysis and medical management tools for developing novel applications for the management of epilepsy. The authors describe the algorithms and architectures needed to provide ambulatory, diagnostic and long-term monitoring services, through multi parametric data collection. Readers will see how to achieve in-hospital quality standards, addressing conventional "routine" clinic-based service purposes, at reduced cost, enhanced capability and increased geographical availability. The cyberphysical system described in this book is flexible, can be optimized for each patient and is demonstrated in several case studies.

This book constitutes thoroughly refereed post-conference proceedings of the workshops of the 17th International Conference on Parallel Computing, Euro-Par 2011, held in Bordeaux, France, in August 2011. The papers of these 12 workshops CCPI, CGWS, HeteroPar, HiBB, HPCVirt, HPPC, HPSS HPCF, PROPER, CCPI, and VHPC focus on promotion and advancement of all aspects of parallel and distributed computing.

The advances in sensor design have decreased the size, weight, and cost of sensors by orders of magnitude, yet with the increase of higher spatial and temporal resolution and accuracy. With the fast progress of sensors design and communications technique, sensor networks have also been quickly evolving in both research and practical domains in the last decade. More and more sensor networks have been deployed in real-world to gather information for our daily life. Applications of sensor networks can be found in battlefield surveillance, environmental monitoring, biological detection, smart spaces, industrial diagnostics, etc. Although the technique of sensor networks has a very promising future, many challenges are still deserving lots of research efforts for its successful applications.

This book is devoted to coverage control, one of the most fundamental and important research issues in sensor networks. The aim of the book is to provide tutorial-like and up-to-date reference resources on various coverage control problems in sensor networks, a hot topic that has been intensively researched in recent years. Due to some unique characteristics of sensor networks such as energy constraint and ad-hoc topology, the coverage problems in sensor networks have many new scenarios and features that entitle them an important research issue in recent years. I have done my best to include in the book the most recent advances, techniques, protocols, results, and findings in this field.

Programming multi-core and many-core computing systems Sabri Pillana, Linnaeus University, Sweden Fatos Xhafa, Technical University of Catalonia, Spain Provides state-of-the-art methods for programming multi-core and many-core systems The book comprises a selection of twenty two chapters covering: fundamental techniques and algorithms; programming approaches; methodologies and frameworks; scheduling and management; testing and evaluation methodologies; and case studies for programming multi-core and many-core systems. Program development for multi-core processors, especially for heterogeneous multi-core processors, is significantly more complex than for single-core processors. However, programmers have been traditionally trained for the development of sequential programs, and only a small percentage of them have experience with parallel programming. In the past, only a relatively small group of programmers interested in High Performance Computing (HPC) was concerned with the parallel programming issues, but the situation has changed dramatically with the appearance of multi-core processors on commonly used computing systems. It is expected that with the pervasiveness of multi-core processors, parallel programming will become mainstream. The pervasiveness of multi-core processors affects a large spectrum of systems, from embedded and general-purpose, to high-end computing systems. This book assists programmers in mastering the efficient programming of multi-core systems, which is of paramount importance for the software-intensive industry towards a more effective product-development cycle. Key features: Lessons, challenges, and roadmaps ahead. Contains real world examples and case studies. Helps programmers in mastering the efficient programming of multi-core and many-core systems. The book serves as a reference for a larger audience of practitioners, young researchers and graduate level students. A basic level of programming knowledge is required to use this book.

Energy Science and Applied Technology includes contributions on a wide range of topics:- Technologies in geology, mining, oil and gas exploration and exploitation of deposits- Energy transfer and conversion, materials and chemical technologies- Environmental engineering and sustainable development- Electrical and electronic technology, power system

This book constitutes the refereed proceedings of the 15th International Conference on High-Performance Computing, HiPC 2008, held in Bangalore, India, in December 2008. The 46 revised full papers presented together with the abstracts of 5 keynote talks were carefully reviewed and selected from 317 submissions. The papers are organized in topical sections on applications performance optimization, parallel algorithms and applications, scheduling and resource management, sensor networks, energy-aware computing, distributed algorithms, communication networks as well as architecture.

This book offers advanced parallel and distributed algorithms and experimental laboratory prototypes of unconventional shortest path solvers. In addition, it presents novel and unique algorithms of solving shortest problems in massively parallel cellular automaton machines. The shortest path problem is a fundamental and classical problem in graph theory and computer science and is frequently applied in the contexts of transport and logistics, telecommunication networks, virtual reality and gaming, geometry, and social networks analysis. Software implementations include distance-vector algorithms for distributed path computation in dynamics networks, parallel solutions of the constrained shortest path problem, and application of the shortest path solutions in gathering robotic swarms. Massively parallel algorithms utilise cellular automata, where a shortest path is computed either via matrix multiplication in automaton arrays, or via the representation of data graphs in automaton lattices and using the propagation of wave-like patterns. Unconventional shortest path solvers are presented in computer models of foraging behaviour and protoplasmic network optimisation by the slime mould *Physarum polycephalum* and fluidic devices, while experimental laboratory prototypes of path solvers using chemical media, flows and droplets, and electrical current are also highlighted. The book will be a pleasure to explore for readers from all walks of life, from undergraduate students to university professors, from mathematicians, computer scientists and engineers to chemists and biologists.

Peer-to-peer networking is a disruptive technology for large scale distributed applications that has recently gained wide interest

due to the successes of peer-to-peer (P2P) content sharing, media streaming, and telephony applications. There are a large range of other applications under development or being proposed. The underlying architectures share features such as decentralization, sharing of end system resources, autonomy, virtualization, and self-organization. These features constitute the P2P paradigm. This handbook broadly addresses a large cross-section of current research and state-of-the-art reports on the nature of this paradigm from a large number of experts in the field. Several trends in information and network technology such as increased performance and deployment of broadband networking, wireless networking, and mobile devices are synergistic with and reinforcing the capabilities of the P2P paradigm. There is general expectation in the technical community that P2P networking will continue to be an important tool for networked applications and impact the evolution of the Internet. A large amount of research activity has resulted in a relatively short time, and a growing community of researchers has developed. The Handbook of Peer-to-Peer Networking is dedicated to discussions on P2P networks and their applications. This is a comprehensive book on P2P computing.

This book constitutes the refereed proceedings of the 20th International Conference on Parallel and Distributed Computing, EuroPar 2014, held in Porto, Portugal, in August 2014. The 68 revised full papers presented were carefully reviewed and selected from 267 submissions. The papers are organized in 15 topical sections: support tools environments; performance prediction and evaluation; scheduling and load balancing; high-performance architectures and compilers; parallel and distributed data management; grid, cluster and cloud computing; green high performance computing; distributed systems and algorithms; parallel and distributed programming; parallel numerical algorithms; multicore and manycore programming; theory and algorithms for parallel computation; high performance networks and communication; high performance and scientific applications; and GPU and accelerator computing.

"This book offers historical perspectives on mobile computing, as well as new frameworks and methodologies for mobile networks, intelligent mobile applications, and mobile computing applications"--Provided by publisher.

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

This book constitutes the refereed proceedings of the 22nd International Conference on Parallel and Distributed Computing, EuroPar 2016, held in Grenoble, France, in August 2016. The 47 revised full papers presented together with 2 invited papers and one industrial paper were carefully reviewed and selected from 176 submissions. The papers are organized in 12 topical sections: Support Tools and Environments; Performance and Power Modeling, Prediction and Evaluation; Scheduling and Load Balancing; High Performance Architectures and Compilers; Parallel and Distributed Data Management and Analytics; Cluster and Cloud Computing; Distributed Systems and Algorithms; Parallel and Distributed Programming, Interfaces, Languages; Multicore and Manycore Parallelism; Theory and Algorithms for Parallel Computation and Networking; Parallel Numerical Methods and Applications; Accelerator Computing.

This book constitutes the refereed proceedings of the 13th International Conference on High-Performance Computing, HiPC 2006, held in Bangalore, India, December 2006. Coverage in this volume includes scheduling and load balancing, network and distributed algorithms, application software, network services, ad-hoc networks, systems software, sensor networks and performance evaluation, as well as routing and data management algorithms.

The edited volume contains original papers contributed to 1st International Conference on Smart System, Innovations and Computing (SSIC 2017) by researchers from different countries. The contributions focus on two main areas, i.e. Smart Systems Innovations which includes applications for smart cities, smart grid, social computing and privacy challenges with their theory, specification, design, performance, and system building. And second Computing of Complex Solutions which includes algorithms, security solutions, communication and networking approaches. The volume provides a snapshot of current progress in related areas and a glimpse of future possibilities. This volume is useful for researchers, Ph.D. students, and professionals working in the core areas of smart systems, innovations and computing.

The hybrid/heterogeneous nature of future microprocessors and large high-performance computing systems will result in a reliance on two major types of components: multicore/manycore central processing units and special purpose hardware/massively parallel accelerators. While these technologies have numerous benefits, they also pose substantial performance challenges for developers, including scalability, software tuning, and programming issues. Researchers at the Forefront Reveal Results from Their Own State-of-the-Art Work Edited by some of the top researchers in the field and with contributions from a variety of international experts, Scientific Computing with Multicore and Accelerators focuses on the architectural design and implementation of multicore and manycore processors and accelerators, including graphics processing units (GPUs) and the Sony Toshiba IBM (STI) Cell Broadband Engine (BE) currently used in the Sony PlayStation 3. The book explains how numerical libraries, such as LAPACK, help solve computational science problems; explores the emerging area of hardware-oriented numerics; and presents the design of a fast Fourier transform (FFT) and a parallel list ranking algorithm for the Cell BE. It covers stencil computations, auto-tuning, optimizations of a computational kernel, sequence alignment and homology, and pairwise computations. The book

also evaluates the portability of drug design applications to the Cell BE and illustrates how to successfully exploit the computational capabilities of GPUs for scientific applications. It concludes with chapters on dataflow frameworks, the Charm++ programming model, scan algorithms, and a portable intracore communication framework. Explores the New Computational Landscape of Hybrid Processors By offering insight into the process of constructing and effectively using the technology, this volume provides a thorough and practical introduction to the area of hybrid computing. It discusses introductory concepts and simple examples of parallel computing, logical and performance debugging for parallel computing, and advanced topics and issues related to the use and building of many applications.

The book is designed to serve as a textbook for courses offered to graduate and undergraduate students enrolled in electronics and electrical engineering and computer science. This book attempts to bridge the gap between electronics and computer science students, providing complementary knowledge that is essential for designing an embedded system. The book covers key concepts tailored for embedded system design in one place. The topics covered in this book are models and architectures, Executable Specific Languages – SystemC, Unified Modeling Language, real-time systems, real-time operating systems, networked embedded systems, Embedded Processor architectures, and platforms that are secured and energy-efficient. A major segment of embedded systems needs hard real-time requirements. This textbook includes real-time concepts including algorithms and real-time operating system standards like POSIX threads. Embedded systems are mostly distributed and networked for deterministic responses. The book covers how to design networked embedded systems with appropriate protocols for real-time requirements. Each chapter contains 2-3 solved case studies and 10 real-world problems as exercises to provide detailed coverage and essential pedagogical tools that make this an ideal textbook for students enrolled in electrical and electronics engineering and computer science programs.

Annual Reports in Computational Chemistry is a new periodical providing timely and critical reviews of important topics in computational chemistry as applied to all chemical disciplines. Topics covered include quantum chemistry, molecular mechanics, force fields, chemical education, and applications in academic and industrial settings. Each volume is organized into (thematic) sections with contributions written by experts. Focusing on the most recent literature and advances in the field, each article covers a specific topic of importance to computational chemists. Annual Reports in Computational Chemistry is a "must" for researchers and students wishing to stay up-to-date on current developments in computational chemistry. * Broad coverage of computational chemistry and up-to-date information * Topics covered include bioinformatics, drug discovery, protein NMR, simulation methodologies, and applications in academic and industrial settings * Each chapter reviews the most recent literature on a specific topic of interest to computational chemists

It is our great pleasure to present the proceedings of the symposia and workshops on parallel and distributed computing and applications associated with the ICA3PP 2010 conference. These symposia and workshops provide vibrant opportunities for researchers and industry practitioners to share their research experience, original research results and practical development experiences in the new challenging research areas of parallel and distributed computing technologies and applications. It was the first time that the ICA3PP conference series added symposia and workshops to its program in order to provide a wide range of topics that extend beyond the main conferences. The goal was to provide a better coverage of emerging research areas and also forums for focused and stimulating discussions. With this objective in mind, we selected three workshops to accompany the ICA3PP 2010 conference: • FPDC 2010, the 2010 International Symposium on Frontiers of Parallel and Distributed Computing • HPCTA 2010, the 2010 International Workshop on High-Performance Computing, Technologies and Applications • M2A 2010, the 2010 International Workshop on Multicore and Multi-threaded Architectures and Algorithms Each of the symposia / workshops focused on a particular theme and complemented the spectrum of the main conference. All papers published in the workshops proceedings were selected by the Program Committee on the basis of referee reports. Each paper was reviewed by independent referees who judged the papers for originality, quality, contribution, presentation and consistency with the theme of the workshops.

This book presents a range of cloud computing platforms for data-intensive scientific applications. It covers systems that deliver infrastructure as a service, including: HPC as a service; virtual networks as a service; scalable and reliable storage; algorithms that manage vast cloud resources and applications runtime; and programming models that enable pragmatic programming and implementation toolkits for eScience applications. Many scientific applications in clouds are also introduced, such as bioinformatics, biology, weather forecasting and social networks. Most chapters include case studies. Cloud Computing for Data-Intensive Applications targets advanced-level students and researchers studying computer science and electrical engineering. Professionals working in cloud computing, networks, databases and more will also find this book useful as a reference.

This two volume set LNCS 7016 and LNCS 7017 constitutes the refereed proceedings of the 11th International Conference on Algorithms and Architectures for Parallel Processing, ICA3PP 2011, held in Melbourne, Australia, in October 2011. The first volume presents 24 revised regular papers and 17 revised short papers together with the abstract of the keynote lecture - all carefully reviewed and selected from 85 initial submissions. The papers cover the many dimensions of parallel algorithms and architectures, encompassing fundamental theoretical approaches, practical experimental results, and commercial components and systems and focus on two broad areas of parallel and distributed computing, i.e., architectures, algorithms and networks, and systems and applications.

This book constitutes the refereed proceedings of the 20th International Conference on Architecture of Computing Systems, ARCS 2007, held in Zurich, Switzerland in March 2007. Coverage details a broad range of research topics related to basic technology, architecture, and application of computing systems with a strong focus on system aspects of pervasive computing and self organization techniques in both organic and autonomic computing.

An analytical overview of the state of the art, open problems, and future trends in heterogeneous parallel and distributed computing This book provides an overview of the ongoing academic research, development, and uses of heterogeneous parallel and distributed computing in the context of scientific computing. Presenting the state of the art in this challenging and rapidly evolving area, the book is organized in five distinct parts: Heterogeneous Platforms: Taxonomy, Typical Uses, and Programming Issues Performance Models of Heterogeneous Platforms and Design of Heterogeneous Algorithms Performance: Implementation and Software Applications Future Trends High Performance Heterogeneous Computing is a valuable reference for researchers and practitioners in the area of high performance heterogeneous computing. It also serves as an excellent supplemental text for graduate and postgraduate courses in related areas.

This tutorial book presents six carefully revised lectures given at the Spring School on Datatype-Generic Programming, SSDGP 2006. This was held in Nottingham, UK, in April 2006. It was colocated with the Symposium on Trends in Functional Programming (TFP 2006), and the Conference of the Types Project (TYPES 2006). All the lectures have been subjected to thorough internal review by the editors and contributors, supported by independent external reviews.

Advances in GPU Research and Practice focuses on research and practices in GPU based systems. The topics treated cover a range of issues, ranging from hardware and architectural issues, to high level issues, such as application systems, parallel programming, middleware, and power and energy issues. Divided into six parts, this edited volume provides the latest research on GPU computing. Part I: Architectural Solutions focuses on the architectural topics that improve on performance of GPUs, Part II: System Software discusses OS, compilers, libraries, programming

environment, languages, and paradigms that are proposed and analyzed to help and support GPU programmers. Part III: Power and Reliability Issues covers different aspects of energy, power, and reliability concerns in GPUs. Part IV: Performance Analysis illustrates mathematical and analytical techniques to predict different performance metrics in GPUs. Part V: Algorithms presents how to design efficient algorithms and analyze their complexity for GPUs. Part VI: Applications and Related Topics provides use cases and examples of how GPUs are used across many sectors. Discusses how to maximize power and obtain peak reliability when designing, building, and using GPUs Covers system software (OS, compilers), programming environments, languages, and paradigms proposed to help and support GPU programmers Explains how to use mathematical and analytical techniques to predict different performance metrics in GPUs Illustrates the design of efficient GPU algorithms in areas such as bioinformatics, complex systems, social networks, and cryptography Provides applications and use case scenarios in several different verticals, including medicine, social sciences, image processing, and telecommunications

"This book presents, discusses, shares ideas, results and experiences on the recent important advances and future challenges on enabling technologies for achieving higher performance"--Provided by publisher.

The two volume set LNCS 7133 and LNCS 7134 constitutes the thoroughly refereed post-conference proceedings of the 10th International Conference on Applied Parallel and Scientific Computing, PARA 2010, held in Reykjavík, Iceland, in June 2010. These volumes contain three keynote lectures, 29 revised papers and 45 minisymposia presentations arranged on the following topics: cloud computing, HPC algorithms, HPC programming tools, HPC in meteorology, parallel numerical algorithms, parallel computing in physics, scientific computing tools, HPC software engineering, simulations of atomic scale systems, tools and environments for accelerator based computational biomedicine, GPU computing, high performance computing interval methods, real-time access and processing of large data sets, linear algebra algorithms and software for multicore and hybrid architectures in honor of Fred Gustavson on his 75th birthday, memory and multicore issues in scientific computing - theory and praxis, multicore algorithms and implementations for application problems, fast PDE solvers and a posteriori error estimates, and scalable tools for high performance computing.

This book constitutes the refereed proceedings of the 14th International Conference on High-Performance Computing, HiPC 2007, held in Goa, India, in December 2007. The 53 revised full papers presented together with the abstracts of five keynote talks were carefully reviewed and selected from 253 submissions. The papers are organized in topical sections on a broad range of applications including I/O and FPGAs, and microarchitecture and multiprocessor architecture.

The LNCS journal Transactions on Large-Scale Data- and Knowledge-Centered Systems focuses on data management, knowledge discovery, and knowledge processing, which are core and hot topics in computer science. Since the 1990s, the Internet has become the main driving force behind application development in all domains. An increase in the demand for resource sharing across different sites connected through networks has led to an evolution of data- and knowledge-management systems from centralized systems to decentralized systems enabling large-scale distributed applications providing high scalability. Current decentralized systems still focus on data and knowledge as their main resource. Feasibility of these systems relies basically on P2P (peer-to-peer) techniques and the support of agent systems with scaling and decentralized control. Synergy between grids, P2P systems, and agent technologies is the key to data- and knowledge-centered systems in large-scale environments. This, the ninth issue of Transactions on Large-Scale Data- and Knowledge-Centered Systems, contains five revised selected regular papers focusing on the following topics: top-k query processing in P2P systems, self-stabilizing consensus average algorithms in distributed sensor networks, recoverable encryption schemes, xml data in a multi-system environment, and pairwise similarity for cluster ensemble problems.

This book provides basic and fundamental knowledge of various aspects of energy-aware computing at the component, software, and system level. It provides a broad range of topics dealing with power-, energy-, and temperature-related research areas for individuals from industry and academia.

Large Scale and Big Data: Processing and Management provides readers with a central source of reference on the data management techniques currently available for large-scale data processing. Presenting chapters written by leading researchers, academics, and practitioners, it addresses the fundamental challenges associated with Big Data processing

t
[Copyright: e5274daca2d609d910bdc1fa79e633fe](https://doi.org/10.1007/978-3-642-26099-9)