

## Distributed And Cloud Computing Kai Hwang Geoffrey Free

This book provides an overview of Cloud Computing in an enterprise environment, describes the benefits and challenges, and then leads the reader through the process of assessing the suitability of a cloud-based approach for a given situation, calculating and justifying the investment that is required to transform the process or application, and then developing a solid design that considers the implementation as well as the ongoing operations and governance required to maintain the solution in a partially outsourced delivery model.

In this text, Smith and Nair take a new approach by examining virtual machines as a unified discipline and pulling together cross-cutting technologies. Topics include instruction set emulation, dynamic program translation and optimization, high level virtual machines (including Java and CLI), and system virtual machines for both single-user systems and servers.

The salient features of the book are as follows:

- Hybrid Elements including topics like Memory organization, Binary representation of data, Computer arithmetic Software for parallel programming, tagged across some chapters through Quick Response (QR) Codes
- Learning objectives tagged across chapters:
- Emphasis on parallelism, scalability and programmability aspects of computer architecture. It presents the analysis of scalability
- Issues related to instruction level parallelism, processor clock speed, and power consumption defined according to the recent developments in processor design
- Inclusion of important topics like processor design, control unit, input and output, parallelis
- erial Bus, Real systems— IBM, Hitachi, Cray,

## Get Free Distributed And Cloud Computing Kai Hwang Geoffrey Free

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

Mastering Cloud Computing is designed for undergraduate students learning to develop cloud computing applications. Tomorrow's applications won't live on a single computer but will be deployed from and reside on a virtual server, accessible anywhere, any time. Tomorrow's application developers need to understand the requirements of building apps for these virtual systems, including concurrent programming, high-performance computing, and data-intensive systems. The book introduces the principles of distributed and parallel computing underlying cloud architectures and specifically focuses on virtualization, thread programming, task programming, and map-reduce programming. There are examples demonstrating all of these and more, with exercises and labs throughout. Explains how to make design choices and tradeoffs to consider when building applications to run in a virtual cloud environment Real-world case studies include scientific, business, and energy-efficiency considerations Component Oriented Programming offers a unique programming-centered approach to

## Get Free Distributed And Cloud Computing Kai Hwang Geoffrey Free

component-based software development that delivers the well-developed training and practices you need to successfully apply this cost-effective method. Following an overview of basic theories and methodologies, the authors provide a unified component infrastructure for building component software using JavaBeans, EJB, OSGi, CORBA, CCM, .NET, and Web services. You'll learn how to develop reusable software components; build a software system of pre-built software components; design and implement a component-based software system using various component-based approaches. Clear organization and self-testing features make Component Oriented Programming an ideal textbook for graduate and undergraduate courses in computer science, software engineering, or information technology as well as a valuable reference for industry professionals.

This book provides the users with quick and easy data acquisition, processing, storage and product generation services. It describes the entire life cycle of remote sensing data and builds an entire high performance remote sensing data processing system framework. It also develops a series of remote sensing data management and processing standards. Features:

- Covers remote sensing cloud computing
- Covers remote sensing data integration across distributed data centers
- Covers cloud storage based remote sensing data share service
- Covers high performance remote sensing data processing
- Covers distributed remote sensing products analysis

The Cloud Computing Bible is a complete reference to cloud computing that presents the technologies, protocols, platforms and infrastructure that make cloud computing possible and desirable. Many of the cloud computing books on the market today are small books of 300 pages or less and the larger books tend to be programming books or security titles. A longer

## Get Free Distributed And Cloud Computing Kai Hwang Geoffrey Free

format book such as Cloud Computing Bible allows a complete definition of the topic as well as in-depth introductions to essential technologies and platforms. Additionally it allows significant technologies to be presented in a form that provides enough detail for the reader to determine if it is something that they are interested in learning more about. It is important to stress platform and technologies as the main subject and intersperse that with products in order to provide an extended life span, but have current appeal. The book will be divided into five parts: The Value Proposition, Platforms, Infrastructure, Services and Applications, and The Mobile Cloud.

The first textbook to teach students how to build data analytic solutions on large data sets using cloud-based technologies. This is 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

## Get Free Distributed And Cloud Computing Kai Hwang Geoffrey Free

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, IBKai HwangM SyNapse, Bluemix 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.

This book gathers research contributions on recent advances in intelligent and distributed computing. A major focus is placed on new techniques and applications for several highlydemanded research directions: Internet of Things, Cloud Computing and Big Data, Data Mining and Machine Learning, Multi-agent and Service-Based Distributed Systems, Distributed Algorithms and Optimization, Modeling Operational Processes, Social Network Analysis and Inappropriate Content Counteraction, Cyber-Physical Security and Safety, Intelligent Distributed Decision Support Systems, Intelligent Human-Machine Interfaces, VisualAnalytics and others. The book represents the peer-reviewed proceedings of the 13thInternational Symposium on Intelligent Distributed Computing (IDC 2019), which was held in St. Petersburg, Russia, from October 7 to 9, 2019. Service Level Agreements for Cloud Computing provides a unique combination of business-driven application scenarios and advanced research in the area of

service-level agreements for Clouds and service-oriented infrastructures. Current state-of-the-art research findings are presented in this book, as well as business-ready solutions applicable to Cloud infrastructures or ERP (Enterprise Resource Planning) environments. Service Level Agreements for Cloud Computing contributes to the various levels of service-level management from the infrastructure over the software to the business layer, including horizontal aspects like service monitoring. This book provides readers with essential information on how to deploy and manage Cloud infrastructures. Case studies are presented at the end of most chapters. Service Level Agreements for Cloud Computing is designed as a reference book for high-end practitioners working in cloud computing, distributed systems and IT services. Advanced-level students focused on computer science will also find this book valuable as a secondary text book or reference.

Distributed and Cloud Computing From Parallel Processing to the Internet of Things Morgan Kaufmann

If you're involved in planning IT infrastructure as a network or system architect, system administrator, or developer, this book will help you adapt your skills to work with these highly scalable, highly redundant infrastructure services. While analysts hotly debate the advantages and risks of cloud computing, IT staff and

programmers are left to determine whether and how to put their applications into these virtualized services. Cloud Application Architectures provides answers -- and critical guidance -- on issues of cost, availability, performance, scaling, privacy, and security. With Cloud Application Architectures, you will: Understand the differences between traditional deployment and cloud computing Determine whether moving existing applications to the cloud makes technical and business sense Analyze and compare the long-term costs of cloud services, traditional hosting, and owning dedicated servers Learn how to build a transactional web application for the cloud or migrate one to it Understand how the cloud helps you better prepare for disaster recovery Change your perspective on application scaling To provide realistic examples of the book's principles in action, the author delves into some of the choices and operations available on Amazon Web Services, and includes high-level summaries of several of the other services available on the market today. Cloud Application Architectures provides best practices that apply to every available cloud service. Learn how to make the transition to the cloud and prepare your web applications to succeed. The efficient management of a consistent and integrated database is a central task in modern IT and highly relevant for science and industry. Hardly any critical enterprise solution comes without any functionality for managing data in its

different forms. Web-Scale Data Management for the Cloud addresses fundamental challenges posed by the need and desire to provide database functionality in the context of the Database as a Service (DBaaS) paradigm for database outsourcing. This book also discusses the motivation of the new paradigm of cloud computing, and its impact to data outsourcing and service-oriented computing in data-intensive applications. Techniques with respect to the support in the current cloud environments, major challenges, and future trends are covered in the last section of this book. A survey addressing the techniques and special requirements for building database services are provided in this book as well.

This book presents a collection of contributions addressing recent advances and research in synergistic combinations of topics in the joint fields of intelligent computing and distributed computing. It focuses on the following specific topics: distributed data mining and machine learning, reasoning and decision-making in distributed environments, distributed evolutionary algorithms, trust and reputation models for distributed systems, scheduling and resource allocation in distributed systems, intelligent multi-agent systems, advanced agent-based and service-based architectures, and Smart Cloud and Internet of Things (IoT) environments. The book represents the combined peer-reviewed proceedings of the 11th

International Symposium on Intelligent Distributed Computing (IDC 2017) and the 7th International Workshop on Applications of Software Agents (WASA 2017), both of which were held in Belgrade, Serbia from October 11 to 13, 2017.

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

"This book compiles recent research trends and practical issues in the fields of distributed computing and Internet technologies, providing advancements on emerging technologies that aim to support the effective design and implementation of service-oriented networks, future Internet environments and building management frameworks"--

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. The energy consumption issue in distributed computing systems raises various monetary, environmental and system performance concerns. Electricity consumption in the US doubled from 2000 to 2005. From a financial and environmental

standpoint, reducing the consumption of electricity is important, yet these reforms must not lead to performance degradation of the computing systems. These contradicting constraints create a suite of complex problems that need to be resolved in order to lead to 'greener' distributed computing systems. This book brings together a group of outstanding researchers that investigate the different facets of green and energy efficient distributed computing. Key features: One of the first books of its kind Features latest research findings on emerging topics by well-known scientists Valuable research for grad students, postdocs, and researchers Research will greatly feed into other technologies and application domains

This important text provides a single point of reference for state-of-the-art cloud computing design and implementation techniques. The book examines cloud computing from the perspective of enterprise architecture, asking the question; how do we realize new business potential with our existing enterprises? Topics and features: with a Foreword by Thomas Erl; contains contributions from an international selection of preeminent experts; presents the state-of-the-art in enterprise architecture approaches with respect to cloud computing models, frameworks, technologies, and applications; discusses potential research directions, and technologies to facilitate the realization of emerging business models through enterprise architecture approaches; provides relevant theoretical frameworks, and the latest empirical research findings.

Distributed and Cloud Computing: From Parallel Processing to the Internet of Things

## Get Free Distributed And Cloud Computing Kai Hwang Geoffrey Free

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

This book constitutes the refereed proceedings of the First International Conference on Big Data Analytics, BDA 2012, held in New Delhi, India, in December 2012. The 5 regular papers and 5 short papers presented were carefully reviewed and selected from 42 submissions. The volume also contains two tutorial papers in the section perspectives on big data analytics. The regular contributions are organized in topical sections on: data analytics applications; knowledge discovery through information extraction; and data models in analytics.

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

## Get Free Distributed And Cloud Computing Kai Hwang Geoffrey Free

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](http://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. About the Book Recent industry surveys expect the cloud computing services market to be in excess of \$20 billion and cloud computing jobs to be in excess of 10 million

worldwide in 2014 alone. In addition, since a majority of existing information technology (IT) jobs is focused on maintaining legacy in-house systems, the demand for these kinds of jobs is likely to drop rapidly if cloud computing continues to take hold of the industry. However, there are very few educational options available in the area of cloud computing beyond vendor-specific training by cloud providers themselves. Cloud computing courses have not found their way (yet) into mainstream college curricula. This book is written as a textbook on cloud computing for educational programs at colleges. It can also be used by cloud service providers who may be interested in offering a broader perspective of cloud computing to accompany their own customer and employee training programs. The typical reader is expected to have completed a couple of courses in programming using traditional high-level languages at the college-level, and is either a senior or a beginning graduate student in one of the science, technology, engineering or mathematics (STEM) fields. We have tried to write a comprehensive book that transfers knowledge through an immersive "hands-on approach", where the reader is provided the necessary guidance and knowledge to develop working code for real-world cloud applications. Additional support is available at the book's website: [www.cloudcomputingbook.info](http://www.cloudcomputingbook.info) Organization The book is organized into three main parts. Part I covers technologies that form the foundations of cloud computing. These include topics such as virtualization, load balancing, scalability & elasticity, deployment, and replication. Part II introduces the reader to the design &

programming aspects of cloud computing. Case studies on design and implementation of several cloud applications in the areas such as image processing, live streaming and social networks analytics are provided. Part III introduces the reader to specialized aspects of cloud computing including cloud application benchmarking, cloud security, multimedia applications and big data analytics. Case studies in areas such as IT, healthcare, transportation, networking and education are provided.

From the Foreword: "Big Data Management and Processing is [a] state-of-the-art book that deals with a wide range of topical themes in the field of Big Data. The book, which probes many issues related to this exciting and rapidly growing field, covers processing, management, analytics, and applications... [It] is a very valuable addition to the literature. It will serve as a source of up-to-date research in this continuously developing area. The book also provides an opportunity for researchers to explore the use of advanced computing technologies and their impact on enhancing our capabilities to conduct more sophisticated studies." ---Sartaj Sahni, University of Florida, USA "Big Data Management and Processing covers the latest Big Data research results in processing, analytics, management and applications. Both fundamental insights and representative applications are provided. This book is a timely and valuable resource for students, researchers and seasoned practitioners in Big Data fields. --Hai Jin, Huazhong University of Science and Technology, China Big Data Management and Processing explores a range of big data related issues and their impact on the design

of new computing systems. The twenty-one chapters were carefully selected and feature contributions from several outstanding researchers. The book endeavors to strike a balance between theoretical and practical coverage of innovative problem solving techniques for a range of platforms. It serves as a repository of paradigms, technologies, and applications that target different facets of big data computing systems. The first part of the book explores energy and resource management issues, as well as legal compliance and quality management for Big Data. It covers In-Memory computing and In-Memory data grids, as well as co-scheduling for high performance computing applications. The second part of the book includes comprehensive coverage of Hadoop and Spark, along with security, privacy, and trust challenges and solutions. The latter part of the book covers mining and clustering in Big Data, and includes applications in genomics, hospital big data processing, and vehicular cloud computing. The book also analyzes funding for Big Data projects.

Big data technologies are used to achieve any type of analytics in a fast and predictable way, thus enabling better human and machine level decision making. Principles of distributed computing are the keys to big data technologies and analytics. The mechanisms related to data storage, data access, data transfer, visualization and predictive modeling using distributed processing in multiple low cost machines are the key considerations that make big data analytics possible within stipulated cost and time practical for consumption by human and

machines. However, the current literature available in big data analytics needs a holistic perspective to highlight the relation between big data analytics and distributed processing for ease of understanding and practitioner use. This book fills the literature gap by addressing key aspects of distributed processing in big data analytics. The chapters tackle the essential concepts and patterns of distributed computing widely used in big data analytics. This book discusses also covers the main technologies which support distributed processing. Finally, this book provides insight into applications of big data analytics, highlighting how principles of distributed computing are used in those situations. Practitioners and researchers alike will find this book a valuable tool for their work, helping them to select the appropriate technologies, while understanding the inherent strengths and drawbacks of those technologies.

### Cloud Computing and Distributed Systems

From its first appearance in 1995, this book has been consistently well received by tutors and students alike. Now with a revised and updated 3rd edition the authors have updated the original text to better reflect the latest developments in Software Project Management.

This book introduces the reader to the fundamentals of contemporary, emerging and future technologies and services in Internet computing. It covers essential

concepts such as distributed systems architectures and web technologies, contemporary paradigms such as cloud computing and the Internet of things, and emerging technologies like distributed ledger technologies and fog computing. The book also highlights the interconnection and recombination of these Internet-based technologies, which together form a critical information infrastructure with major impacts on individuals, organizations, governments, economies, and society as a whole. Intended as a textbook for upper undergraduate and graduate classes, it features a wealth of examples, learning goals and summaries for every chapter, numerous recommendations for further reading, and questions for checking students' comprehension. A dedicated author website offers additional teaching material and more elaborate examples. Accordingly, the book enables students and young professionals in IT-related fields to familiarize themselves with the Internet's basic mechanisms, and with the most promising Internet-based technologies of our time. This book provides readers with a thorough understanding of various research areas within the field of data science. The book introduces readers to various techniques for data acquisition, extraction, and cleaning, data summarizing and modeling, data analysis and communication techniques, data science tools, deep learning, and various data science applications. Researchers can extract and

conclude various future ideas and topics that could result in potential publications or thesis. Furthermore, this book contributes to Data Scientists' preparation and to enhancing their knowledge of the field. The book provides a rich collection of manuscripts in highly regarded data science topics, edited by professors with long experience in the field of data science. Introduces various techniques, methods, and algorithms adopted by Data Science experts Provides a detailed explanation of data science perceptions, reinforced by practical examples Presents a road map of future trends suitable for innovative data science research and practice

The development of smart cities is one of the most important challenges over the next few decades. Governments and companies are leveraging billions of dollars in public and private funds for smart cities. Next generation smart cities are heavily dependent on distributed smart sensing systems and devices to monitor the urban infrastructure. The smart sensor networks serve as autonomous intelligent nodes to measure a variety of physical or environmental parameters. They should react in time, establish automated control, and collect information for intelligent decision-making. In this context, one of the major tasks is to develop advanced frameworks for the interpretation of the huge amount of information provided by the emerging testing and monitoring systems. Data Analytics for

Smart Cities brings together some of the most exciting new developments in the area of integrating advanced data analytics systems into smart cities along with complementary technological paradigms such as cloud computing and Internet of Things (IoT). The book serves as a reference for researchers and engineers in domains of advanced computation, optimization, and data mining for smart civil infrastructure condition assessment, dynamic visualization, intelligent transportation systems (ITS), cyber-physical systems, and smart construction technologies. The chapters are presented in a hands-on manner to facilitate researchers in tackling applications. Arguably, data analytics technologies play a key role in tackling the challenge of creating smart cities. Data analytics applications involve collecting, integrating, and preparing time- and space-dependent data produced by sensors, complex engineered systems, and physical assets, followed by developing and testing analytical models to verify the accuracy of results. This book covers this multidisciplinary field and examines multiple paradigms such as machine learning, pattern recognition, statistics, intelligent databases, knowledge acquisition, data visualization, high performance computing, and expert systems. The book explores new territory by discussing the cutting-edge concept of Big Data analytics for interpreting massive amounts of data in smart city applications.

Innovations in cloud and service-oriented architectures continue to attract attention by offering interesting opportunities for research in scientific communities. Although advancements such as computational power, storage, networking, and infrastructure have aided in making major progress in the implementation and realization of cloud-based systems, there are still significant concerns that need to be taken into account. Principles, Methodologies, and Service-Oriented Approaches for Cloud Computing aims to present insight into Cloud principles, examine associated methods and technologies, and investigate the use of service-oriented computing technologies. In addressing supporting infrastructure of the Cloud, including associated challenges and pressing issues, this reference source aims to present researchers, engineers, and IT professionals with various approaches in Cloud computing.

Minimize Power Consumption and Enhance User Experience Essential for high-speed fifth-generation mobile networks, mobile cloud computing (MCC) integrates the power of cloud data centers with the portability of mobile computing devices. Mobile Cloud Computing: Architectures, Algorithms and Applications covers the latest technological and architectural

An Introduction to Parallel Programming, Second Edition presents a tried-and-true tutorial approach that shows students how to develop effective parallel

programs with MPI, Pthreads and OpenMP. As the first undergraduate text to directly address compiling and running parallel programs on multi-core and cluster architecture, this second edition carries forward its clear explanations for designing, debugging and evaluating the performance of distributed and shared-memory programs while adding coverage of accelerators via new content on GPU programming and heterogeneous programming. New and improved user-friendly exercises teach students how to compile, run and modify example programs. Takes a tutorial approach, starting with small programming examples and building progressively to more challenging examples Explains how to develop parallel programs using MPI, Pthreads and OpenMP programming models A robust package of online ancillaries for instructors and students includes lecture slides, solutions manual, downloadable source code, and an image bank New to this edition: New chapters on GPU programming and heterogeneous programming New examples and exercises related to parallel algorithms

This volume contains the proceedings of CloudCom 2009, the First International Conference on Cloud Computing. The conference was held in Beijing, China, during December 1–4, 2009, and was the first in a series initiated by the Cloud Computing Association ([www.cloudcom.org](http://www.cloudcom.org)). The Cloud Computing Association

was founded in 2009 by Chunming Rong, Martin Gilje Jaatun, and Frode Eika Sandnes. This first conference was organized by the Beijing Ji-tong University, Chinese Institute of Electronics, and Wuhan University, and co-organized by Huazhong University of Science and Technology, South China Normal University, and Sun Yat-sen University. Ever since the inception of the Internet, a “Cloud” has been used as a metaphor for a network-accessible infrastructure (e.g., data storage, computing hardware, or entire networks) which is hidden from users. To some, the concept of cloud computing may seem like a throwback to the days of big mainframe computers, but we believe that cloud computing makes data truly mobile, allowing a user to access services anywhere, anytime, with any Internet browser. In cloud computing, IT-related capabilities are provided as services, accessible without requiring control of, or even knowledge of, the underlying technology. Cloud computing provides dynamic scalability of services and computing power, and although many mature technologies are used as components in cloud computing, there are still many unresolved and open problems.

Cloud computing promises to revolutionize IT and business by making computing available as a utility over the internet. This book is intended primarily for practising software architects who need to assess the impact of such a

transformation. It explains the evolution of the internet into a cloud computing platform, describes emerging development paradigms and technologies, and discusses how these will change the way enterprise applications should be architected for cloud deployment. Gautam Shroff provides a technical description of cloud computing technologies, covering cloud infrastructure and platform services, programming paradigms such as MapReduce, as well as 'do-it-yourself' hosted development tools. He also describes emerging technologies critical to cloud computing. The book also covers the fundamentals of enterprise computing, including a technical introduction to enterprise architecture, so it will interest programmers aspiring to become software architects and serve as a reference for a graduate-level course in software architecture or software engineering.

This volume offers readers various perspectives and visions for cutting-edge research in ubiquitous healthcare. The topics emphasize large-scale architectures and high performance solutions for smart healthcare, healthcare monitoring using large-scale computing techniques, Internet of Things (IoT) and big data analytics for healthcare, Fog Computing, mobile health, large-scale medical data mining, advanced machine learning methods for mining multidimensional sensor data, smart homes, and resource allocation methods for

the BANs. The book contains high quality chapters contributed by leading international researchers working in domains, such as e-Health, pervasive and context-aware computing, cloud, grid, cluster, and big-data computing. We are optimistic that the topics included in this book will provide a multidisciplinary research platform to the researchers, practitioners, and students from biomedical engineering, health informatics, computer science, and computer engineering. "This book provide relevant theoretical frameworks covering the latest empirical research findings in the area of grid computing, with a critical perspective bridging the gap between academia and the latest achievements of the computer industry"--Provided by publisher.

[Copyright: 4faf5e85b0194968f42e3aa07f88bf15](#)