

Big Data Benchmarks Performance Optimization And Emerging Hardware 6th Workshop Epoe 2015 Kohala Hi Usa August 31 September 4 2015 Revised Papers Lecture Notes In Computer Science

This book constitutes the proceedings of the 10th International IFIP WG 8.9 Working Conference on Research and Practical Issues of Enterprise Information Systems, CONFENIS 2016, held in Vienna, Austria, in December 2016. The conference provided an international forum for the broader IFIP community to discuss the latest research findings in the area of EIS and specifically aimed at facilitating the exchange of ideas and advances on all aspects and developments of EIS. The 25 papers presented in this volume were carefully reviewed and selected from 63 submissions. They were organized in topical sections on: semantic concepts and open data; customer relationship management; security and privacy issues; advanced manufacturing and management aspects; business intelligence and big data; decision support in EIS; and EIS-practices.

This volume constitutes the proceedings of the 26th International Conference on Web Services, ICWS 2019, held as part of SCF 2019 in San Diego, CA, USA in June 2019. The 11 full papers together with 1 short paper published in this volume were carefully reviewed and selected from 31 submissions. ICWS has been a prime international forum for both researchers and industry practitioners to exchange the latest fundamental advances in the state of the art and practice of Web-based services, to identify emerging research topics, and to define the future of Web-based services. Topics include Internet services modeling, discovery, composition, testing, adaptation, delivery, as well as standards.

Big Data Benchmarks, Performance Optimization, and Emerging Hardware4th and 5th Workshops, BPOE 2014, Salt Lake City, USA, March 1, 2014 and Hangzhou, China, September 5, 2014, Revised Selected PapersSpringer

In order to carry out data analytics, we need powerful and flexible computing software. However the software available for data analytics is often proprietary and can be expensive. This book reviews Apache tools, which are open source and easy to use. After providing an overview of the background of data analytics, covering the different types of analysis and the basics of using Hadoop as a tool, it focuses on different Hadoop ecosystem tools, like Apache Flume, Apache Spark, Apache Storm, Apache Hive, R, and Python, which can be used for different types of analysis. It then examines the different machine learning techniques that are useful for data analytics, and how to visualize data with different graphs and charts. Presenting data analytics from a practice-oriented viewpoint, the book discusses useful tools and approaches for data analytics, supported by concrete code examples. The book is a valuable reference resource for graduate students and professionals in related fields, and is also of interest to general readers with an understanding of data analytics.

The digital age has presented an exponential growth in the amount of data available to individuals looking to draw conclusions based on given or collected information across industries. Challenges associated with the analysis, security, sharing, storage, and visualization of large and complex data sets continue to plague data scientists and analysts alike as traditional data processing applications struggle to adequately manage big data. The Handbook of Research on Big Data Storage and Visualization Techniques is a critical scholarly resource that explores big data analytics and technologies and their role in developing a broad understanding of issues pertaining to the use of big data in multidisciplinary fields. Featuring coverage on a broad range of topics, such as architecture patterns, programing systems, and computational energy, this publication is geared towards professionals, researchers, and students seeking current research and application topics on the subject.

Performance Evaluation and Benchmarking

First Workshop, WBDB 2012, San Jose, CA, USA, May 8–9, 2012 and Second Workshop, WBDB 2012, Pune, India, December 17–18, 2012, Revised Selected Papers

First Workshop, SDBA 2018, Beijing, China, June 12, 2018, Revised Selected Papers

Toward Solving Real-World Optimization Problems

Proceedings of ICBCCC18

Assured Cloud Computing

Research and Practical Issues of Enterprise Information Systems

This book gives detailed instructions on how to use, optimize, and troubleshoot mod_perl. It shows how to get this Apache module running quickly and easily.

Performance Optimization of Numerically Intensive Codes offers a comprehensive, tutorial-style, hands-on, introductory and intermediate-level treatment of all the essential ingredients for achieving high performance in numerical computations on modern computers. The authors explain computer architectures, data traffic and issues related to performance of serial and parallel code optimization exemplified by actual programs written for algorithms of wide interest. The unique hands-on style is achieved by extensive case studies using realistic computational problems. The performance gain obtained by applying the techniques described in this book can be very significant. The book bridges the gap between the literature in system architecture, the one in numerical methods and the occasional descriptions of optimization topics in computer vendors' literature. It also allows readers to better judge the suitability of certain computer architecture to their computational requirements. In contrast to standard textbooks on computer architecture and on programming techniques the book treats these topics together at the level necessary for writing high-performance programs. The book facilitates easy access to these topics for computational scientists and engineers mainly interested in practical issues related to efficient code development.

This six volume set LNCS 11063 – 11068 constitutes the thoroughly refereed conference proceedings of the 4th International Conference on Cloud Computing and Security, ICCCS 2018, held in Haikou, China, in June 2018. The 386 full papers of these six volumes were carefully reviewed and selected from 1743 submissions. The papers cover ideas and achievements in the theory and practice of all areas of inventive systems which includes control, artificial intelligence, automation systems, computing systems, electrical and informative systems. The six volumes are arranged according to the subject areas as follows: cloud computing, cloud security, encryption, information hiding, IoT security, multimedia forensics.

This book constitutes the refereed proceedings of workshops, held at the 33rd International Conference on Conceptual Modeling, ER 2014, in Atlanta, GA, USA in October 2014. The 24 revised full and 6 short papers were carefully reviewed and selected out of 59 submissions and are presented together with 4 demonstrations. The papers are organized in sections related to the individual workshops: the First International Workshop on Enterprise Modeling, ENMO 2014; the Second International Workshop on Modeling and Management of Big Data, MoBID 2014; the First International Workshop on Conceptual Modeling in Requirements and Business Analysis, MBReBA 2014; the First International Workshop on Quality of Models and Models of Quality, QMMQ 2014; the 8th International Workshop on Semantic and Conceptual Issues in GIS, SeCoGIS 2014; and the 11th International Workshop on Web Information Systems Modeling, WISM 2014. The contributions cover a variety of topics in conceptual modeling, including requirements and enterprise modeling, modeling of big data, spatial conceptual modeling, exploring the quality of models, and issues specific to the design of web information systems.

This book constitutes the refereed post-conference proceedings of the 11th TPC Technology Conference on Performance Evaluation and Benchmarking, TPCTC 2019, held in conjunction with the 45th International Conference on Very Large Databases (VLDB 2019) in August 2019. The 11 papers presented were carefully reviewed and focus on topics such as blockchain; big data and analytics; complex event processing; database Optimizations; data Integration; disaster tolerance and recovery; artificial Intelligence; emerging storage technologies (NVMe, 3D XPoint Memory etc.); hybrid workloads; energy and space efficiency; in-memory databases; internet of things; virtualization; enhancements to TPC workloads; lessons learned in practice using TPC workloads; collection and interpretation of performance data in public cloud environments.

6th International Workshop, WBDB 2015, Toronto, ON, Canada, June 16–17, 2015 and 7th International Workshop, WBDB 2015, New Delhi, India, December 14–15, 2015, Revised Selected Papers

10th TPC Technology Conference, TPCTC 2018, Rio de Janeiro, Brazil, August 25–31, 2018, Revised Selected Papers

High-Performance Big Data Computing

4th International Conference, ICCCS 2018, Haikou, China, June 8–10, 2018, Revised Selected Papers, Part II

Advancing Big Data Benchmarks

Systems Simulation and Modeling for Cloud Computing and Big Data Applications

Practical Mod_perl

Explores key challenges and solutions to assured cloud computing today and provides a provocative look at the face of cloud computing tomorrow This book offers readers a comprehensive suite of solutions for resolving many of the key challenges to achieving high levels of assurance in cloud computing. The distillation of critical research findings generated by the Assured Cloud Computing Center of Excellence (ACC-UCoE) of the University of Illinois, Urbana-Champaign, it provides unique insights into the current and future shape of robust, dependable, and secure cloud-based computing and data cyberinfrastructures. A survivable and distributed cloud-computing-based infrastructure can enable the configuration of any dynamic systems-of-systems that contain both trusted and partially trusted resources and services sourced from multiple organizations. To assure mission-critical computations and workflows that rely on such systems-of-systems it is necessary to ensure that a given configuration does not violate any security or reliability requirements. Furthermore, it is necessary to model the trustworthiness of a workflow or computation fulfillment to a high level of assurance. In presenting the substance of the work done by the ACC-UCoE, this book provides a vision for assured cloud computing illustrating how individual research contributions relate to each other and to the big picture of assured cloud computing. In addition, the book: Explores dominant themes in cloud-based systems, including design correctness, support for big data and analytics, monitoring and detection, network considerations, and performance Synthesizes heavily cited earlier work on topics such as DARE, trust mechanisms, and elastic graphs, as well as newer research findings on topics, including R-Storm, and RAMP transactions Addresses assured cloud computing concerns such as game theory, stream processing, storage, algorithms, workflow, scheduling, access control, formal analysis of safety, and streaming Bringing together the freshest thinking and applications in one of today's most important topics, Assured Cloud Computing is a must-read for researchers and professionals in the fields of computer science and engineering, especially those working within industrial, military, and governmental contexts. It is also a valuable reference for advanced students of computer science.

This two-volume set LNCS 11446 and LNCS 11447 constitutes the refereed proceedings of the 24th International Conference on Database Systems for Advanced Applications, DASFAA 2019, held in Chiang Mai, Thailand, in April 2019. The 92 full papers and 64 short papers were carefully selected from a total of 501 submissions. In addition, 13 demo papers and 6 tutorial papers are included. The full papers are organized in the following topics: big data; clustering and classification; crowdsourcing; data integration; embedding; graphs; knowledge graph; machine learning; privacy and graph; recommendation; social network; spatial; and spatio-temporal. The short papers, demo papers, and tutorial papers can be found in the volume LNCS 11448, which also includes the workshops of DASFAA 2019.

This book presents selected papers from the 3rd International Conference on Micro-Electronics and Telecommunication Engineering, held at SRM Institute of Science and Technology, Ghaziabad, India, on 30–31 August 2019. It covers a wide variety of topics in micro-electronics and telecommunication engineering, including micro-electronic engineering, computational remote sensing, computer science and intelligent systems, signal and image processing, and information and communication technology.

This book constitutes the thoroughly refereed post-workshop proceedings of the 6th International Workshop on Big Data Benchmarking, WBDB 2015, held in Toronto, ON, Canada, in June 2015 and the 7th International Workshop, WBDB 2015, held in New Delhi, India, in December 2015. The 8 full papers presented in this book were carefully reviewed and selected from 22 submissions. They deal with recent trends in big data and HPC convergence, new proposals for big data benchmarking, as well as tooling and performance results.

Systems Simulation and Modelling for Cloud Computing and Big Data Applications provides readers with the most current approaches to solving problems through the use of models and simulations, presenting SSM based approaches to performance testing and benchmarking that offer significant advantages. For example, multiple big data and cloud application developers and researchers can perform tests in a controllable and repeatable manner. Inspired by the need to analyze the performance of different big data processing and cloud frameworks, researchers have introduced several benchmarks, including BigDataBench, BigBench, HiBench, PigMix, CloudSuite and GridMix, which are all covered in this book. Despite the substantial progress, the research community still needs a holistic, comprehensive big data SSM to use in almost every scientific and engineering discipline involving multidisciplinary research. SSM develops frameworks that are applicable across disciplines to develop benchmarking tools that are useful in solutions development. Examines the methodology and requirements of benchmarking big data and cloud computing tools, advances in big data frameworks and benchmarks for large-scale data analytics, and frameworks for benchmarking and predictive analytics in big data deployment Discusses applications using big data benchmarks, such as BigDataBench, BigBench, HiBench, MapReduce, HPPC, ECL, HOBBIT, GridMix and PigMix, and applications using big data frameworks, such as Hadoop, Spark, Samza, Flink and SQL frameworks Covers development of big data benchmarks to evaluate workloads in state-of-the-practice heterogeneous hardware platforms, advances in modeling and simulation tools for performance evaluation, security problems and scalable cloud computing environments

Database Systems for Advanced Applications

Performance Optimization of Numerically Intensive Codes

Proceedings of the 2013 Workshop Series on Big Data Benchmarking, WBDB.cn, Xi'an, China, July16–17, 2013 and WBDB.us, San José, CA, USA, October 9–10, 2013, Revised Selected Papers

Big Data

Handbook of Research on Big Data Storage and Visualization Techniques

Understanding Big Data: Analytics for Enterprise Class Hadoop and Streaming Data

Big Scientific Data Benchmarks, Architecture, and Systems

This book constitutes the thoroughly refereed proceedings of the 19th East European Conference on Advances in Databases and Information Systems, ADBIS 2015, held in Poitiers, France, in September 2015. The 31 full papers and 18 short papers presented were carefully selected and reviewed from 135 submissions. The papers are organized in topical sections such as database theory and access methods; user requirements and database evolution; multidimensional modeling and OLAP; ETI.; transformation, extraction and archiving; modeling and ontologies; time series processing; performance and tuning; advanced queries data mining; approximation and sketch confidence and trust.

ICT technologies have contributed to the advances in wireless systems, which provide seamless connectivity for worldwide communication. The growth of interconnected devices and the need to store, manage, and process the data from them has led to increased research on the intersection of the internet of things and cloud computing. The Handbook of Research on the IoT, Cloud Computing, and Wireless Network Optimization is a pivotal reference source that provides the latest research findings and solutions for the design and augmentation of wireless systems and cloud computing. The content within this publication examines data mining, machine learning, and software engineering, and is designed for IT specialists, software engineers, researchers, academicians, industry professionals, and students.

Big data solutions enable us to change how we do business by exploiting previously unused sources of information in ways that were not possible just a few years ago. In IBM® Smarter Planet® terms, big data helps us to change the way that the world works. The purpose of this IBM Redpaper™ publication is to consider the performance and capacity implications of big data solutions, which must be taken into account for them to be viable. This paper describes the benefits that big data approaches can provide. We then cover performance and capacity considerations for creating big data solutions. We conclude with what this means for big data solutions, both now and in the future. Intended readers for this paper include decision-makers, consultants, and IT architects.

This book constitutes the refereed post-conference proceedings of the 12th TPC Technology Conference on Performance Evaluation and Benchmarking, TPCTC 2020, held in August 2020. The 8 papers presented were carefully reviewed and cover the following topics: testing ACID compliance in the LDBC social network benchmark; experimental performance evaluation of stream processing engines made easy; revisiting issues in benchmarking metric selection; performance evaluation for digital transformation; experimental comparison of relational and NoSQL document systems; a framework for supporting repetition and evaluation in the process of cloud-based benchmarking; benchmarking AI inference: a domain independent benchmark evaluation model for the transaction processing performance council.

This book constitutes the thoroughly revised selected papers of the 4th and 5th workshops on Big Data Benchmarks, Performance Optimization, and Emerging Hardware, BPOE 4 and BPOE 5, held respectively in Salt Lake City, in March 2014, and in Hangzhou, in September 2014. The 16 papers presented were carefully reviewed and selected from 30 submissions. Both workshops focus on architecture and system support for big data systems, such as benchmarking; workload characterization; performance optimization and evaluation; emerging hardware.

Performance and Capacity Implications for Big Data

Big Data Benchmarking

Second BenchCouncil International Symposium, Bench 2019, Denver, CO, USA, November 14–16, 2019, Revised Selected Papers

Integrated Platforms and Industry Use Cases

4th and 5th Workshops, BPOE 2014, Salt Lake City, USA, March 1, 2014 and Hangzhou, China, September 5, 2014, Revised Selected Papers

Benchmarking, Measuring, and Optimizing

From cloud computing to data analytics, society stores vast supplies of information through wireless networks and mobile computing. As organizations are becoming increasingly more wireless, ensuring the security and seamless function of electronic gadgets while creating a strong network is imperative. Advanced Methodologies and Technologies in Network Architecture, Mobile Computing, and Data Analytics highlights the challenges associated with creating a strong network architecture in a perpetually online society. Readers will learn various methods in building a seamless mobile computing option and the most effective means of analyzing big data. This book is an important resource for information technology professionals, software developers, data analysts, graduate-level students, researchers, computer engineers, and IT specialists seeking modern information on emerging methods in data mining, information technology, and wireless networks.

This book provides an overview of the resources and research projects that are bringing Big Data and High Performance Computing (HPC) on converging tracks. It demystifies Big Data and HPC for the leader by covering the primary resources, middleware, applications, and tools that enable the usage of HPC platforms for Big Data management and processing. Through interesting use-cases from traditional and non-traditional HPC domains, the book highlights the most critical challenges related to Big Data processing and management, and shows ways to mitigate them using HPC resources. Unlike most books on Big Data, it covers a variety of alternatives to Hadoop, and explains the differences between HPC platforms and Hadoop. Written by professionals and researchers in a range of departments and fields, this book is designed for anyone studying Big Data and its future directions. Those studying HPC will also find the content valuable.

This book constitutes the thoroughly refereed post-conference proceedings of the 10th TPC Technology Conference on Performance Evaluation and Benchmarking, TPCTC 2018, held in conjunction with the 44th International Conference on Very Large Databases (VLDB 2018) in August 2018. The 10 papers presented were carefully reviewed and selected from numerous submissions. The TPC encourages researchers and industry experts to present and debate novel ideas and methodologies in performance evaluation, measurement, and characterization.

Big Data: Principles and Paradigms captures the state-of-the-art research on the architectural aspects, technologies, and applications of Big Data. The book identifies potential future directions and technologies that facilitate insight into numerous scientific, business, and consumer applications. To help realize Big Data's full potential, the book addresses numerous challenges, offering the conceptual and technological solutions for tackling them. These challenges include life-cycle data management, large-scale storage, flexible processing infrastructure, data modeling, scalable machine learning, data analysis algorithms, sampling techniques, and privacy and ethical issues. Covers computational platforms supporting Big Data applications Addresses key principles underlying Big Data computing Examines key developments supporting next generation Big Data platforms Explores the challenges in Big Data computing and ways to overcome them Contains expert contributors from both academia and industry

This book comprehensively conveys the theoretical and practical aspects of IoT and big data analytics with the solid contributions from practitioners as well as academicians. This book examines and expounds the unique capabilities of the big data analytics platforms in capturing, cleansing and crunching IoT device/sensor data in order to extricate actionable insights. A number of experimental case studies and real-world scenarios are incorporated in this book in order to instigate our book readers. This book Analyzes current research and development in the domains of IoT and big data analytics Gives an overview of latest trends and transactions happening in the IoT data analytics space Illustrates the various platforms, processes, patterns, and practices for simplifying and streamlining IoT data analytics The Internet of Things and Big Data Analytics: Integrated Platforms and Industry Use Cases examines and accentuates how the multiple challenges at the cusp of IoT and big data can be fully met. The device ecosystem is growing steadily. It is forecast that there will be billions of connected devices in the years to come. When these IoT devices, resource-constrained as well as resource-intensive, interact with one another locally and remotely, the amount of multi-structured data generated, collected, and stored is bound to grow exponentially. Another prominent trend is the integration of IoT devices with cloud-based applications, services, infrastructures, middleware solutions, and databases. This book examines the pioneering technologies and tools emerging and evolving in order to collect, pre-process, store, process and analyze data heaps in order to disentangle actionable insights.

Proceedings of 3rd ICMETE 2019

5th International Workshop, WBDB 2014, Potsdam, Germany, August 5–6– 2014, Revised Selected Papers

Encyclopedia of Information Science and Technology, Fourth Edition

ER 2014 Workshops, ENMO, MoBID, MBReBA, QMMQ, SeCoGIS, WISM, and ER Demos, Atlanta, GA, USA, October 27–29, 2014, Proceedings

Performance Evaluation and Benchmarking for the Era of Artificial Intelligence

Advances in Conceptual Modeling

19th East European Conference, ADBIS 2015, Poitiers, France, September 8–11, 2015, Proceedings

This book constitutes the thoroughly refereed joint proceedings of the Third and Fourth Workshop on Big Data Benchmarking. The third WBDB was held in Xi'an, China, in July 2013 and the Fourth WBDB was held in San José, CA, USA, in October, 2013. The 15 papers presented in this book were carefully reviewed and selected from 33 presentations. They focus on big data benchmarks; applications and scenarios; tools, systems and surveys.

"This 10-volume compilation of authoritative, research-based articles contributed by thousands of researchers and experts from all over the world emphasized modern issues and the presentation of potential opportunities, prospective solutions, and future directions in the field of information science and technology"--Provided by publisher.

Big Data represents a new era in data exploration and utilization, and IBM is uniquely positioned to help clients navigate this transformation. This book reveals how IBM is leveraging open source Big Data technology, infused with IBM technologies, to deliver a robust, secure, highly available, enterprise-class Big Data platform. The three defining characteristics of Big Data--volume, variety, and velocity--are discussed. You'll get a primer on Hadoop and how IBM is hardening it for the enterprise, and learn when to leverage IBM InfoSphere BigInsights (Big Data at rest) and IBM InfoSphere Streams (Big Data in motion) technologies. Industry use cases are also included in this practical guide. Learn how IBM hardens Hadoop for enterprise-class scalability and reliability Gain insight into IBM's unique in-motion and at-rest Big Data analytics platform Learn tips and tricks for Big Data use cases and solutions Get a quick Hadoop primer

This book constitutes the thoroughly refereed post-workshop proceedings of the 5th International Workshop on Big Data Benchmarking, WBDB 2014, held in Potsdam, Germany, in August 2014. The 13 papers presented in this book were carefully reviewed and selected from numerous submissions and cover topics such as benchmarks specifications and proposals, Hadoop and MapReduce - in the different context such as virtualization and cloud - as well as in-memory, data generation, and graphs.

This book constitutes the refereed proceedings of the First Workshop on Big Scientific Data Benchmarks, Architecture, and Systems, SDBA 2018, held in Beijing, China, in June 2018. The 10 revised full papers presented were carefully reviewed and selected from 22 submissions. The papers are organized in topical sections on benchmarking; performance optimization; algorithms; big science data framework.

Micro-Electronics and Telecommunication Engineering

Performance Evaluation and Benchmarking for the Era of Cloud(s)

Big Data Optimization: Recent Developments and Challenges

Optimization in the Real World

Advances in Databases and Information Systems

Principles and Paradigms

Cloud Computing and Security

In recent years, our world has experienced a profound shift and progression in available computing and knowledge sharing innovations. These emerging advancements have developed at a rapid pace, disseminating into and affecting numerous aspects of contemporary society. This has created a pivotal need for an innovative compendium encompassing the latest trends, concepts, and issues surrounding this relevant discipline area. During the past 15 years, the Encyclopedia of Information Science and Technology has become recognized as one of the landmark sources of the latest knowledge and discoveries in this discipline. The Encyclopedia of Information Science and Technology, Fourth Edition is a 10-volume set which includes 705 original and previously unpublished research articles covering a full range of perspectives, applications, and techniques contributed by thousands of experts and researchers from around the globe. This authoritative encyclopedia is an all-encompassing, well-established reference source that is ideally designed to disseminate the most forward-thinking and diverse research findings. With critical perspectives on the impact of information science management and new technologies in modern settings, including but not limited to computer science, education, healthcare, government, engineering, business, and natural and physical sciences, it is a pivotal and relevant source of knowledge that will benefit every professional within the field of information science and technology and is an invaluable addition to every academic and corporate library.

As municipalities continue to see a growth in population, planners are continually searching for utilizing space and seeking the best geographical arrangements for these cities. Professionals have continually used geographic information systems (GIS) to solve these issues; however, limitations in this technology remain prevalent. Integrating multiple-criteria decision analysis and evolutionary computing tools with GIS has created an array of robust solutions for spatial optimization problems in densely populated areas. Interdisciplinary Approaches to Spatial Optimization Issues is a pivotal reference source that provides vital research on advancements within the field of GIS and evolutionary solutions for spatial optimization issues. While highlighting topics such as computing machinery, vehicular routing, and operational research, this publication is ideally designed for practitioners, technicians, developers, academicians, students, government officials, planners, and researchers seeking current research on applications and improvements within spatial optimization and GIS.

This book clearly shows the importance, usefulness, and powerfulness of current optimization technologies, in particular, mixed-integer programming and its remarkable applications. It is intended to be the definitive study of state-of-the-art optimization technologies for students, academic researchers, and non-professionals in industry. The chapters of this book are based on a collection of selected and extended papers from the "IMI Workshop on Optimization in the Real World" held in October 2014 in Japan.

An in-depth overview of an emerging field that brings together high-performance computing, big data processing, and deep learning. Over the last decade, the exponential explosion of data known as big data has changed the way we understand and harness the power of data. The emerging field of high-performance big data computing, which brings together high-performance computing (HPC), big data processing, and deep learning, aims to meet the challenges posed by large-scale data processing. This book offers an in-depth overview of high-performance big data computing and the associated technical issues, approaches, and solutions. The book covers basic concepts and necessary background knowledge, including data processing frameworks, storage systems, and hardware capabilities; offers a detailed discussion of technical issues in accelerating big data computing in terms of computation, communication, memory and storage, codesign, workload characterization and benchmarking, and system deployment and management; and surveys benchmarks and workloads for evaluating big data middleware systems. It presents a detailed discussion of big data computing systems and applications with high-performance networking, computing, and storage technologies, including state-of-the-art designs for data processing and storage systems. Finally, the book considers some advanced research topics in high-performance big data computing, including designing high-performance deep learning over big data (DLoBD) stacks and HPC cloud technologies.

The main objective of this book is to provide the necessary background to work with big data by introducing some novel optimization algorithms and codes capable of working in the big data setting as well as introducing some applications in big data optimization for both academics and practitioners interested, and to benefit society, industry, academia, and government. Presenting applications in a variety of industries, this book will be useful for the researchers aiming to analyses large scale data. Several optimization algorithms for big data including convergent parallel algorithms, limited memory bundle algorithm, diagonal bundle method, convergent parallel algorithms, network analytics, and many more have been explored in this book.

Interdisciplinary Approaches to Spatial Optimization Issues

6th Workshop, BPOE 2015, Kohala, HI, USA, August 31 - September 4, 2015, Revised Selected Papers

Advances in Big Data and Cloud Computing

Handbook of Research on the IoT, Cloud Computing, and Wireless Network Optimization

11th TPC Technology Conference, TPCTC 2019, Los Angeles, CA, USA, August 26, 2019, Revised Selected Papers

10th IFIP WG 8.9 Working Conference, CONFENIS 2016, Vienna, Austria, December 13–14, 2016, Proceedings

Supercomputing Frontiers

This book constitutes the thoroughly refereed revised selected papers of the First Workshop on Big Data Benchmarks, WBDB 2012, held in San Jose, CA, USA, in May 2012 and the Second Workshop on Big Data Benchmarks, WBDB 2012, held in Pune, India, in December 2012. The 14 revised papers presented were carefully reviewed and selected from 60 submissions. The papers are organized in topical sections on benchmarking, foundations and tools; domain specific benchmarking; benchmarking hardware and end-to-end big data benchmarks.

This book is a compendium of the proceedings of the International Conference on Big Data and Cloud Computing. It includes recent advances in the areas of big data analytics, cloud computing, internet of nano things, cloud security, data analytics in the cloud, smart cities and grids, etc. This volume primarily focuses on the application of the knowledge that promotes ideas for solving the problems of the society through cutting-edge technologies. The articles featured in this proceeding provide novel ideas that contribute to the growth of world class research and development. The contents of this volume will be of interest to researchers and practitioners alike.

This book constitutes the refereed proceedings of the First International Symposium on Benchmarking, Measuring, and Optimization, Bench 2018, held in Seattle, WA, USA, in December 2018. The 20 full papers presented were carefully reviewed and selected from 51 submissions. The papers are organized in topical sections named: AI Benchmarking; Cloud; Big Data; Modelling and Prediction; and Algorithm and Implementations.

This book constitutes the refereed proceedings of the Second International Symposium on Benchmarking, Measuring, and Optimization, Bench 2019, held in Denver, CO, USA, in November 2019. The 20 full papers and 11 short papers presented were carefully reviewed and selected from 79 submissions. The papers are organized in topical sections named: Best Paper Session; AI Challenges on Cambricon using AIBench; AI Challenges on RISC-V using AIBench; AI Challenges on X86 using AIBench; AI Challenges on 3D Face Recognition using AIBench; Benchmark; AI and Edge; Big Data; Datacenter; Performance Analysis; Scientific Computing.

It constitutes the refereed proceedings of the 4th Asian Supercomputing Conference, SCFA 2018, held in Singapore in March 2018. Supercomputing Frontiers will be rebranded as Supercomputing Frontiers Asia (SCFA), which serves as the technical programme for SCA18. The technical programme for SCA18 consists of four tracks: Application, Algorithms & Libraries Programming System Software Architecture, Network/Communications & Management Data, Storage & Visualisation The 20 papers presented in this volume were carefully reviewed nd selected from 60 submissions.

Advanced Methodologies and Technologies in Network Architecture, Mobile Computing, and Data Analytics

Encyclopedia of Information Science and Technology, Third Edition

Conquering Big Data with High Performance Computing

26th International Conference, Held as Part of the Services Conference Federation, SCF 2019, San Diego, CA, USA, June 25-30, 2019, Proceedings

First BenchCouncil International Symposium, Bench 2018, Seattle, WA, USA, December 10-13, 2018, Revised Selected Papers

Big Data Benchmarks, Performance Optimization, and Emerging Hardware

4th Asian Conference, SCFA 2018, Singapore, March 26-29, 2018, Proceedings

This book constitutes the thoroughly revised selected papers of the 6th workshop on Big Data Benchmarks, Performance Optimization, and Emerging Hardware, BPOE 2015, held in Kohala Coast, HI, USA, in August/September 2015 as satellite event of VLDB 2015, the 41st International Conference on Very Large Data Bases. The 8 papers presented were carefully reviewed and selected from 10

submissions. The workshop focuses on architecture and system support for big data systems, aiming at bringing researchers and practitioners from data management, architecture, and systems research communities together to discuss the research issues at the intersection of these areas. This book also invites three papers from several industrial partners, including two papers describing

tools used in system benchmarking and monitoring and one paper discussing principles and methodologies in existing big data benchmarks.

12th TPC Technology Conference, TPCTC 2020, Tokyo, Japan, August 31, 2020, Revised Selected Papers

24th International Conference, DASFAA 2019, Chiang Mai, Thailand, April 22-25, 2019, Proceedings, Part I

A Hands-On Approach for Application Development

The Internet of Things and Big Data Analytics

Specifying Big Data Benchmarks

Web Services - ICWS 2019

Network Data Analytics