

# **Big Data Benchmarks Performance Optimization And Emerging Hardware 4th And 5th Workshops Bpoe 2014 Salt Lake City Usa March 1 2014 And Papers Lecture Notes In 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 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, MReBA 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 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.

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.

Big Data has been much in the news in recent years, and the advantages conferred by the collection and analysis of large datasets in fields such as marketing, medicine and finance have led to claims that almost any real world problem could be solved if sufficient data were available. This is of course a very simplistic view, and the usefulness of collecting, processing and storing large datasets must always be seen in terms of the communication, processing and storage capabilities of the computing platforms available. This book presents papers from the International Research Workshop, Advanced High Performance Computing Systems, held in Cetraro, Italy, in July 2014. The papers selected for publication here discuss fundamental aspects of the definition of Big Data, as well as considerations from practice where complex datasets are collected, processed and stored. The concepts, problems, methodologies and solutions presented are of much more general applicability than may be suggested by the particular application areas considered. As a result the book will be of interest to all those whose work involves the processing of very large data sets, exascale computing and the emerging fields of data science

This book constitutes the refereed proceedings of the 4th TPC Technology Conference, TPCTC 2012, held in Istanbul, Turkey, in August 2012. It contains 10 selected peer-reviewed papers, 2 invited talks, a report from the TPC Public Relations Committee, and a report from the workshop on Big Data Benchmarking, WBDB 2012. The papers present novel ideas and methodologies in performance evaluation, measurement, and characterization.

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.

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 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 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 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 thoroughly refereed post-conference proceedings of the 8th TPC Technology Conference, on Performance Evaluation and Benchmarking, TPCTC 2016, held in conjunction with the 41st International Conference on Very Large Databases (VLDB 2016) in New Delhi, India, in September 2016. The 9 papers presented were carefully reviewed and selected from 20 submissions. They reflect the rapid pace at which industry experts and researchers develop innovative techniques for evaluation, measurement and

characterization of complex systems.

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.

This book constitutes the thoroughly refereed post-conference proceedings of the 8th TPC Technology Conference, on Performance Evaluation and Benchmarking, TPCTC 2017, held in conjunction with the 43rd International Conference on Very Large Databases (VLDB 2017) in August/September 2017. The 12 papers presented were carefully reviewed and selected from numerous submissions. The TPC remains committed to developing new benchmark standards to keep pace with these rapid changes in technology. 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 transitions 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.

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.

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; ETL; transformation, extraction and archiving; modeling and ontologies; time series processing; performance and tuning; advanced query processing; approximation and skyline; confidentiality and trust.

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 contains the thoroughly refereed papers from the 9th International Ershov Informatics Conference, PSI 2014, held in St. Petersburg, Russia, in June 2014. The 17 revised full papers, 11 revised short papers, and 2 system and experimental papers presented in this book were carefully reviewed and selected from 80 submissions. The volume also contains 5 keynote talks which cover a range of hot topics in computer science and informatics. The papers cover various topics related to the foundations of program and system development and analysis, programming methodology and software engineering and information technologies.

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 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.

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.

This book constitutes the refereed proceedings of the 21st International Conference on Big Data Analytics and Knowledge Discovery, DaWaK 2019, held in Linz, Austria, in September 2019. The 12 full papers and 10 short papers presented were carefully reviewed and selected from 61 submissions. The papers are organized in the following topical sections: Applications; patterns; RDF and streams; big data systems; graphs and machine learning; databases.

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 DBMS performance benchmarking; benchmarking AI inference; a domain independent benchmark evolution model for the transaction processing performance council.

This book offers an introductory and intermediate-level treatment of all the essential ingredients for achieving high performance in numerical computations.

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 analyse 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.

This book constitutes the refereed post-conference proceedings of the 6th TPC Technology Conference, TPCTC 2014, held in Hangzhou, China, in September 2014. It contains 12 selected peer-reviewed papers, a report from the TPC Public Relations Committee. Many buyers use TPC benchmark results as points of comparison when purchasing new computing systems. The information technology landscape is evolving at a rapid pace, challenging industry experts and researchers to develop innovative techniques for evaluation, measurement and characterization of complex systems. The TPC remains committed to developing new benchmark standards to keep pace and one vehicle for achieving this objective is the sponsorship of the Technology Conference on Performance Evaluation and Benchmarking (TPCTC). Over the last five years TPCTC has been held successfully in conjunction with VLDB.

This book constitutes the revised selected papers of the 17th Smoky Mountains Computational

Sciences and Engineering Conference, SMC 2020, held in Oak Ridge, TN, USA\*, in August 2020. The 36 full papers and 1 short paper presented were carefully reviewed and selected from a total of 94 submissions. The papers are organized in topical sections of computational applications: converged HPC and artificial intelligence; system software: data infrastructure and life cycle; experimental/observational applications: use cases that drive requirements for AI and HPC convergence; deploying computation: on the road to a converged ecosystem; scientific data challenges. \*The conference was held virtually due to the COVID-19 pandemic. 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 Cambircon using AIBenc; 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.

Build production-ready machine learning and NLP systems using functional programming, development platforms, and cloud deployment. **KEY FEATURES** ? In-depth explanation and code samples highlighting the features of the Julia language. ? Extensive coverage of the Julia development ecosystem, package management, DevOps environment integration, and performance management tools. ? Exposure to the most important Julia packages that aid in Data and Text Analytics and Deep Learning. **DESCRIPTION** The Julia Programming language enables data scientists and programmers to create prototypes without sacrificing performance. Nonetheless, skeptics question its readiness for production deployments as a new platform with a 1.0 release in 2018. This book removes these doubts and offers a comprehensive glimpse at the language's use throughout developing and deploying production-ready applications. The first part of the book teaches experienced programmers and scientists about the Julia language features in great detail. The second part consists of gaining hands-on experience with the development environment, debugging, programming guidelines, package management, and cloud deployment strategies. In the final section, readers are introduced to a variety of third-party packages available in the Julia ecosystem for Data Processing, Text Analytics, and developing Deep Learning models. This book provides an extensive overview of the programming language and broadens understanding of the Julia ecosystem. As a result, it assists programmers, scientists, and information architects in selecting Julia for their next production deployments. **WHAT YOU WILL LEARN** ? Get to know the complete fundamentals of Julia programming. ? Explore Julia development frameworks and how to work with them. ? Dig deeper into the concepts and applications of functional programming. ? Uncover the Julia infrastructure for development, testing, and deployment. ? Learn to practice Julia libraries and the Julia package ecosystem. ? Processing Data, Deep Learning, and Natural Language Processing with Julia. **WHO THIS BOOK IS FOR** This book is for Data Scientists and application developers who want to learn about Julia application development. No prior Julia knowledge is required but knowing the basics of programming helps understand the objectives of this book. **TABLE OF CONTENTS** 1. Getting Started 2. Data Types 3. Conditions, Control Flow, and Iterations 4. Functions and Methods 5. Collections 6. Arrays 7. Strings 8. Metaprogramming 9. Standard Libraries Module 2. The Development Environment 10. Programming Guidelines in Julia 11. Performance Management 12. IDE and Debugging 13. Package Management 14. Deployment Module 3. Packages in Julia 15. Data Transformations 16. Text Analytics 17. Deep Learning

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 provides an insight into IoT intelligence in terms of applications and algorithmic challenges. The book is dedicated to addressing the major challenges in realizing the artificial intelligence in IoT-based applications including challenges that vary from cost and energy efficiency to availability to service quality in multidisciplinary fashion. The aim of this book is hence to focus on both the algorithmic and practical parts of the artificial intelligence approaches in IoT applications that are enabled and supported by wireless sensor networks and cellular networks. Targeted readers are from varying disciplines who are interested in implementing the smart planet/environments vision via intelligent wireless/wired enabling technologies. Includes the most up-to-date research and applications related to IoT artificial intelligence (AI); Provides new and innovative operational ideas regarding the IoT artificial intelligence that help advance the telecommunications industry; Presents AI challenges facing the IoT scientists and provides potential ways to solve them in critical daily life issues. 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 reader 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.

As metropolises continue to see a growth in population, planners are continually searching for trending methods 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 constitutes the thoroughly refereed post-conference proceedings of the 7th TPC Technology Conference on Performance Evaluation and Benchmarking, TPSTC 2015, held in conjunction with the 40th International Conference on Very Large Databases (VLDB 2015) in Kohala Coast, Hawaii, USA, in August/September 2015. The 8 papers presented together with 1 keynote, and 1 vision paper were carefully reviewed and selected from 24 submissions. Many buyers use TPC benchmark results

as points of comparison when purchasing new computing systems. The information technology landscape is evolving at a rapid pace, challenging industry experts and researchers to develop innovative techniques for evaluation, measurement and characterization of complex systems. The TPC remains committed to developing new benchmark standards to keep pace, and one vehicle for achieving this objective is the sponsorship of the Technology Conference on Performance Evaluation and Benchmarking (TPCTC).

This book focuses on the analytic principles of business practice and big data. Specifically, it provides an interface between the main disciplines of engineering/technology and the organizational and administrative aspects of management, serving as a complement to books in other disciplines such as economics, finance, marketing and risk analysis. The contributors present their areas of expertise, together with essential case studies that illustrate the successful application of engineering management theories in real-life examples.

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 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.

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.

