

Where To Download Building A Columnar Database On Ramcloud Database Design For The Low Latency Enabled Data Center In Memory Data Management Research Free Download Pdf

Building a Columnar Database on RAMCloud Running a Transactional Database on Top of RamCloud Advanced Parallel Processing Technologies Non-Volatile Memory Database Management Systems The Use of a Low-latency Key-value Store for Implementing a Scalable and High-performance Graph Database Performance Evaluation and Benchmarking Encyclopedia of Business Analytics and Optimization Database Systems for Advanced Applications Main Memory Management on Relational Database Systems Database and Expert Systems Applications Database Systems for Advanced Applications Storage Systems Beyond Databases, Architectures and Structures. Paving the Road to Smart Data Processing and Analysis Beyond Databases, Architectures and Structures. Towards Efficient Solutions for Data Analysis and Knowledge Representation Web-Age Information Management Database Systems for Advanced Applications Handbook of Research on Cloud Infrastructures for Big Data Analytics **Computer Engineering and Technology Achieving Both Low Latency and Strong Consistency in**

Large-scale Systems *An Architecture for Fast and General Data Processing on Large Clusters* **Big Data Benchmarking** *Information Systems Development* **Data Management in the Cloud** **Middleware 2013** *International Joint Conference Designing Data-Intensive Applications* *Cloud Computing* **Intelligent Distributed Computing XIII** *Enabling Real-Time Business Intelligence* **Big Data Management and Processing** *Encyclopedia of Information Science and Technology, Fourth Edition* **In-Memory Data Management** **High-Performance Big Data Computing** *Information Security and Privacy* **Principles of Distributed Database Systems** **Transaction Processing on Modern Hardware** *Innovative Research and Applications in Next-Generation High Performance Computing* *Fast and Scalable Cloud Data Management* **Advanced Methodologies and Technologies in Network Architecture, Mobile Computing, and Data Analytics** *Big Data Computing and Communications*

Encyclopedia of Business Analytics and Optimization Aug 19 2022 As the age of Big Data emerges, it becomes necessary to take the five dimensions of Big Data- volume, variety, velocity, volatility, and veracity- and focus these dimensions towards one critical emphasis - value. The *Encyclopedia of Business Analytics and Optimization* confronts the challenges of information retrieval in the age of Big Data by exploring recent advances in the areas of knowledge management, data visualization, interdisciplinary communication, and others. Through its critical approach and practical application, this book will be a must-have reference for any professional, leader, analyst, or manager interested in making the most of the knowledge resources at their disposal.

Computer Engineering and Technology Sep 08 2021 This book constitutes the refereed proceedings of the 21st CCF Conference on Computer Engineering and Technology, NCCET 2017, held in Xiamen, China, in August 2017. The 13 full papers presented were carefully reviewed and selected from 108 submissions. They address topics such as processor architecture; application specific processors; computer application and software optimization; technology on the horizon. *Web-Age Information Management* Dec 11 2021 This two-volume set,

LNCS 9658 and 9659, constitutes the thoroughly refereed proceedings of the 17th International Conference on Web-Age Information Management, WAIM 2016, held in Nanchang, China, in June 2016. The 80 full research papers presented together with 8 demonstrations were carefully reviewed and selected from 266 submissions. The focus of the conference is on following topics: data mining, spatial and temporal databases, recommender systems, graph data management, information retrieval, privacy and trust, query processing and optimization, social media, big data analytics, and distributed and cloud computing.

Advanced Methodologies and Technologies in Network Architecture, Mobile Computing, and Data Analytics Nov 17 2019

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.

Intelligent Distributed Computing XIII Oct 29 2020 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

13th International Symposium on Intelligent Distributed Computing (IDC 2019), which was held in St. Petersburg, Russia, from October 7 to 9, 2019.

Cloud Computing Nov 29 2020 *Cloud Computing: Business Trends and Technologies* provides a broad introduction to Cloud computing technologies and their applications to IT and telecommunications businesses (i.e., the network function virtualization, NFV). To this end, the book is expected to serve as a textbook in a graduate course on Cloud computing. The book examines the business cases and then concentrates on the technologies necessary for supporting them. In the process, the book addresses the principles of – as well as the known problems with – the underlying technologies, such as virtualization, data communications, network and operations management, security and identity management. It introduces, through open-source case studies (based on OpenStack), an extensive illustration of lifecycle management. The book also looks at the existing and emerging standards, demonstrating their respective relation to each topic. Overall, this is an authoritative textbook on this emerging and still-developing discipline, which

- Guides the reader through basic concepts, to current practices, to state-of-the-art applications.
- Considers technical standards bodies involved in Cloud computing standardization.
- Is written by innovation experts in operating systems and data communications, each with over 20 years' experience in business, research, and teaching.

Achieving Both Low Latency and Strong Consistency in Large-scale Systems Aug 07 2021 Today's datacenter applications demand large-scale and low-latency systems. Unfortunately, consistency mechanisms were not designed in consideration of large-scale and low-latency settings. Most existing consistency mechanisms incur huge penalties on scalability or latency, so many datacenter systems have forgone consistency. As a result, application developers or end-users suffer from unpredictable system behaviors. This dissertation presents two new consistency mechanisms for large-scale and low-latency systems: Reusable Infrastructure for Linearizability (RIFL) and Consistent Unordered Replication Protocol (CURP). RIFL provides a general-purpose mechanism for converting at-least-once RPC semantics to exactly-once semantics, thereby making it easy to turn non-linearizable operations

into linearizable ones. RIFL is designed for large-scale systems and is lightweight enough to be used in low-latency environments. On the RAMCloud storage system, RIFL adds only 0.5 us to the 13.5 us base latency for durable writes, and it can support 1 million clients with minimal performance degradation (~5% latency increase). We also used RIFL to construct a new multi-server transaction mechanism (RIFL-TX) in RAMCloud; RIFL's facilities significantly simplified the transaction implementation. The transaction mechanism can commit simple distributed transactions in about 20 us and it outperforms the H-Store main-memory database system for the TPC-C benchmark. Replication is a must-have for large-scale systems for masking individual server failures and maintaining high availability. However, consistent replication incurs significant overhead, such as doubling the latency of operations. As a result, many large-scale systems have settled for weakly consistent replication. To address this dilemma, I present Consistent Unordered Replication Protocol (CURP), which removes most of the overhead of consistent replication. CURP avoids performance penalties by allowing clients to directly replicate their requests, as long as the requests are commutative. This strategy allows most operations to complete in 1 RTT (the same as an unreplicated system). On RAMCloud, CURP improved write latency by ~2x (14 us -> 7.1 us) and write throughput by 4x. Compared to unreplicated RAMCloud, CURP's latency overhead for 3-way replication is just 1 us (6.1 us vs. 7.1 us).

An Architecture for Fast and General Data Processing on Large Clusters

Jul 06 2021 The past few years have seen a major change in computing systems, as growing data volumes and stalling processor speeds require more and more applications to scale out to clusters. Today, a myriad data sources, from the Internet to business operations to scientific instruments, produce large and valuable data streams. However, the processing capabilities of single machines have not kept up with the size of data. As a result, organizations increasingly need to scale out their computations over clusters. At the same time, the speed and sophistication required of data processing have grown. In addition to simple queries, complex algorithms like machine learning and graph analysis are becoming common. And in addition to batch processing, streaming analysis of real-time data is required to let organizations take

timely action. Future computing platforms will need to not only scale out traditional workloads, but support these new applications too. This book, a revised version of the 2014 ACM Dissertation Award winning dissertation, proposes an architecture for cluster computing systems that can tackle emerging data processing workloads at scale. Whereas early cluster computing systems, like MapReduce, handled batch processing, our architecture also enables streaming and interactive queries, while keeping MapReduce's scalability and fault tolerance. And whereas most deployed systems only support simple one-pass computations (e.g., SQL queries), ours also extends to the multi-pass algorithms required for complex analytics like machine learning. Finally, unlike the specialized systems proposed for some of these workloads, our architecture allows these computations to be combined, enabling rich new applications that intermix, for example, streaming and batch processing. We achieve these results through a simple extension to MapReduce that adds primitives for data sharing, called Resilient Distributed Datasets (RDDs). We show that this is enough to capture a wide range of workloads. We implement RDDs in the open source Spark system, which we evaluate using synthetic and real workloads. Spark matches or exceeds the performance of specialized systems in many domains, while offering stronger fault tolerance properties and allowing these workloads to be combined. Finally, we examine the generality of RDDs from both a theoretical modeling perspective and a systems perspective. This version of the dissertation makes corrections throughout the text and adds a new section on the evolution of Apache Spark in industry since 2014. In addition, editing, formatting, and links for the references have been added.

Storage Systems Mar 14 2022 Storage Systems: Organization, Performance, Coding, Reliability and Their Data Processing was motivated by the 1988 Redundant Array of Inexpensive/Independent Disks proposal to replace large form factor mainframe disks with an array of commodity disks. Disk loads are balanced by striping data into strips—with one strip per disk—and storage reliability is enhanced via replication or erasure coding, which at best dedicates k strips per stripe to tolerate k disk failures. Flash memories have resulted in a paradigm shift with Solid State Drives (SSDs) replacing Hard Disk Drives (HDDs)

for high performance applications. RAID and Flash have resulted in the emergence of new storage companies, namely EMC, NetApp, SanDisk, and Purestorage, and a multibillion-dollar storage market. Key new conferences and publications are reviewed in this book. The goal of the book is to expose students, researchers, and IT professionals to the more important developments in storage systems, while covering the evolution of storage technologies, traditional and novel databases, and novel sources of data. We describe several prototypes: FAWN at CMU, RAMCloud at Stanford, and Lightstore at MIT; Oracle's Exadata, AWS' Aurora, Alibaba's PolarDB, Fungible Data Center; and author's paper designs for cloud storage, namely heterogeneous disk arrays and hierarchical RAID.

- Surveys storage technologies and lists sources of data: measurements, text, audio, images, and video
- Familiarizes with paradigms to improve performance: caching, prefetching, log-structured file systems, and merge-trees (LSMs)
- Describes RAID organizations and analyzes their performance and reliability
- Conserves storage via data compression, deduplication, compaction, and secures data via encryption
- Specifies implications of storage technologies on performance and power consumption
- Exemplifies database parallelism for big data, analytics, deep learning via multicore CPUs, GPUs, FPGAs, and ASICs, e.g., Google's Tensor Processing Units

Fast and Scalable Cloud Data Management Dec 19 2019

The unprecedented scale at which data is both produced and consumed today has generated a large demand for scalable data management solutions facilitating fast access from all over the world. As one consequence, a plethora of non-relational, distributed NoSQL database systems have risen in recent years and today's data management system landscape has thus become somewhat hard to overlook. As another consequence, complex polyglot designs and elaborate schemes for data distribution and delivery have become the norm for building applications that connect users and organizations across the globe – but choosing the right combination of systems for a given use case has become increasingly difficult as well. To help practitioners stay on top of that challenge, this book presents a comprehensive overview and classification of the current system landscape in cloud data management as well as a survey of the state-of-the-art approaches for efficient data distribution and

delivery to end-user devices. The topics covered thus range from NoSQL storage systems and polyglot architectures (backend) over distributed transactions and Web caching (network) to data access and rendering performance in the client (end-user). By distinguishing popular data management systems by data model, consistency guarantees, and other dimensions of interest, this book provides an abstract framework for reasoning about the overall design space and the individual positions claimed by each of the systems therein. Building on this classification, this book further presents an application-driven decision guidance tool that breaks the process of choosing a set of viable system candidates for a given application scenario down into a straightforward decision tree.

Database Systems for Advanced Applications Apr 15 2022 The 4 volume set LNCS 12112-12114 constitutes the papers of the 25th International Conference on Database Systems for Advanced Applications which will be held online in September 2020. The 119 full papers presented together with 19 short papers plus 15 demo papers and 4 industrial papers in this volume were carefully reviewed and selected from a total of 487 submissions. The conference program presents the state-of-the-art R&D activities in database systems and their applications. It provides a forum for technical presentations and discussions among database researchers, developers and users from academia, business and industry.

Information Systems Development May 04 2021 *Information Systems Development: Reflections, Challenges and New Directions*, is the collected proceedings of the 20th International Conference on Information Systems Development held in Edinburgh, Scotland, August 24 - 26, 2011. It follows in the tradition of previous conferences in the series in exploring the connections between industry, research and education. These proceedings represent ongoing reflections within the academic community on established information systems topics and emerging concepts, approaches and ideas. It is hoped that the papers herein contribute towards disseminating research and improving practice

Database and Expert Systems Applications May 16 2022 This two volume set of LNCS 11706 and LNCS 11707 constitutes the refereed proceedings of the 30th International Conference on Database and

Expert Systems Applications, DEXA 2019, held in Linz, Austria, in August 2019. The 32 full papers presented together with 34 short papers were carefully reviewed and selected from 157 submissions. The papers are organized in the following topical sections: Part I: Big data management and analytics; data structures and data management; management and processing of knowledge; authenticity, privacy, security and trust; consistency, integrity, quality of data; decision support systems; data mining and warehousing. Part II: Distributed, parallel, P2P, grid and cloud databases; information retrieval; Semantic Web and ontologies; information processing; temporal, spatial, and high dimensional databases; knowledge discovery; web services.

Enabling Real-Time Business Intelligence Sep 27 2020 This book constitutes the thoroughly refereed conference proceedings of the 4th International Workshop on Business Intelligence for the Real-Time Enterprise, BIRTE 2010, held in Singapore, in September 2010, in conjunction with VLDB 2010, the International Conference on Very Large Data Bases. The volume contains six research papers, which are carefully reviewed and selected out of 12 submissions, plus the three keynotes presented at the workshop. The topics cover all stages of the business intelligence cycle, including capturing of real-time data, handling of temporal or uncertain data, performance issues, event management, and the optimization of complex ETL workflows.

Encyclopedia of Information Science and Technology, Fourth Edition Jul 26 2020 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.

Handbook of Research on Cloud Infrastructures for Big Data Analytics

Oct 09 2021 Clouds are being positioned as the next-generation consolidated, centralized, yet federated IT infrastructure for hosting all kinds of IT platforms and for deploying, maintaining, and managing a wider variety of personal, as well as professional applications and services. Handbook of Research on Cloud Infrastructures for Big Data Analytics focuses exclusively on the topic of cloud-sponsored big data analytics for creating flexible and futuristic organizations. This book helps researchers and practitioners, as well as business entrepreneurs, to make informed decisions and consider appropriate action to simplify and streamline the arduous journey towards smarter enterprises.

Beyond Databases, Architectures and Structures. Paving the Road to Smart Data Processing and Analysis

Feb 13 2022 This book constitutes the refereed proceedings of the 15th International Conference entitled Beyond Databases, Architectures and Structures, BDAS 2019, held in Ustro?, Poland, in May 2019. It consists of 26 carefully reviewed papers selected from 69 submissions. The papers are organized in topical sections, namely big data and cloud computing; architectures, structures and algorithms for efficient data processing and analysis; artificial intelligence, data mining and knowledge discovery; image analysis and multimedia mining; bioinformatics and biomedical data analysis; industrial applications; networks and security.

Beyond Databases, Architectures and Structures. Towards Efficient

Solutions for Data Analysis and Knowledge Representation

Jan 12 2022 This book constitutes the refereed proceedings of the 13th International

Conference entitled Beyond Databases, Architectures and Structures, BDAS 2017, held in Ustro?, Poland, in May/June 2017. It consists of 44 carefully reviewed papers selected from 118 submissions. The papers are organized in topical sections, namely big data and cloud computing; artificial intelligence, data mining and knowledge discovery; architectures, structures and algorithms for efficient data processing; text mining, natural language processing, ontologies and semantic web; bioinformatics and biological data analysis; industrial applications; data mining tools, optimization and compression.

Database Systems for Advanced Applications Nov 10 2021 This book constitutes the workshop proceedings of the 17th International Conference on Database Systems for Advanced Applications, DASFAA 2012, held in Busan, South Korea, in April 2012. The volume contains five workshops, each focusing on specific area that contributes to the main themes of the DASFAA conference: The Second International Workshop on Flash-based Database Systems (FlashDB 2012), the First International Workshop on Information Technologies for Maritime and Logistics (ITEMS 2012), the Third International Workshop on Social Networks and Social Media Mining on the Web (SNSMW 2012), the Second International Workshop on Spatial Information Modeling, Management and Mining (SIM3 2012), and the Fifth International Workshop on Data Quality in Integration Systems (DQIS 2012).

Information Security and Privacy Apr 22 2020 The two-volume set LNCS 9722 and LNCS 9723 constitutes the refereed proceedings of the 21st Australasian Conference on Information Security and Privacy, ACISP 2016, held in Melbourne, VIC, Australia, in July 2016. The 52 revised full and 8 short papers presented together with 6 invited papers in this double volume were carefully revised and selected from 176 submissions. The papers of Part I (LNCS 9722) are organized in topical sections on National Security Infrastructure; Social Network Security; Bitcoin Security; Statistical Privacy; Network Security; Smart City Security; Digital Forensics; Lightweight Security; Secure Batch Processing; Pseudo Random/One-Way Function; Cloud Storage Security; Password/QR Code Security; and Functional Encryption and Attribute-Based Cryptosystem. Part II (LNCS 9723) comprises topics such as Signature and Key Management; Public Key and Identity-Based

Encryption; Searchable Encryption; Broadcast Encryption; Mathematical Primitives; Symmetric Cipher; Public Key and Identity-Based Encryption; Biometric Security; Digital Forensics; National Security Infrastructure; Mobile Security; Network Security; and Pseudo Random/One-Way Function.

Main Memory Management on Relational Database Systems Jun 17

2022 This book provides basic knowledge about main memory management in relational databases as it is needed to support large-scale applications processed completely in memory. In business operations, real-time predictability and high speed is a must. Hence every opportunity must be exploited to improve performance, including reducing dependency on the hard disk, adding more memory to make more data resident in the memory, and even deploying an in-memory system where all data can be kept in memory. The book provides one chapter for each of the main related topics, i.e. the memory system, memory management, virtual memory, and databases and their memory systems, and it is complemented by a short survey of six commercial systems: TimesTen, MySQL, VoltDB, Hekaton, HyPer/ScyPer, and SAP HANA.

The Use of a Low-latency Key-value Store for Implementing a Scalable and High-performance Graph Database Oct 21 2022

This dissertation presents the design, implementation, and evaluation of TorcDB, a high performance graph database system that inherits its ability to scale from a generic key-value store. The key-value store, RAMCloud, is an in memory and low latency distributed storage system that provides the necessary features for scaling, including fault-tolerance, automatic data distribution, and distributed transactions, without assuming anything about the type of data stored by clients. TorcDB achieves high performance using RAMCloud by carefully laying out data and translating traversals into a read and write schedule that minimizes query latency. Using this approach, TorcDB is able to achieve competitive and often better performance than the state of the art in graph database systems today. Evaluated on a large scale social network workload, TorcDB is able to achieve an overall 50th percentile latency of 246us for update queries, 139us for short read queries, and 4.6ms for complex analytical queries.

In-Memory Data Management Jun 24 2020 In the last fifty years the world has been completely transformed through the use of IT. We have now reached a new inflection point. This book presents, for the first time, how in-memory data management is changing the way businesses are run. Today, enterprise data is split into separate databases for performance reasons. Multi-core CPUs, large main memories, cloud computing and powerful mobile devices are serving as the foundation for the transition of enterprises away from this restrictive model. This book provides the technical foundation for processing combined transactional and analytical operations in the same database. In the year since we published the first edition of this book, the performance gains enabled by the use of in-memory technology in enterprise applications has truly marked an inflection point in the market. The new content in this second edition focuses on the development of these in-memory enterprise applications, showing how they leverage the capabilities of in-memory technology. The book is intended for university students, IT-professionals and IT-managers, but also for senior management who wish to create new business processes.

International Joint Conference Feb 01 2021 This volume of *Advances in Intelligent and Soft Computing* contains accepted papers presented at the 8th International Conference on Computational Intelligence in Security for Information Systems (CISIS 2015) and the 6th International Conference on European Transnational Education (ICEUTE 2015). These conferences were held in the beautiful and historic city of Burgos (Spain), in June 2015. The aim of the 8th CISIS conference is to offer a meeting opportunity for academic and industry-related researchers belonging to the various, vast communities of Computational Intelligence, Information Security, and Data Mining. The need for intelligent, flexible behaviour by large, complex systems, especially in mission-critical domains, is intended to be the catalyst and the aggregation stimulus for the overall event. After a thorough peer-review process, the CISIS 2015 International Program Committee selected 43 papers, written by authors from 16 different countries. In the case of 6th ICEUTE conference, the International Program Committee selected 12 papers (from 7 countries). These papers are published in present conference proceedings, achieving an acceptance rate of about 39%. The

selection of papers was extremely rigorous in order to maintain the high quality of the conference and we would like to thank the members of the Program Committees for their hard work in the reviewing process. This is a crucial process to the creation of a high standard conference and the CISIS and ICEUTE conferences would not exist without their help.

Performance Evaluation and Benchmarking Sep 20 2022 This book constitutes the proceedings of the Second Technology Conference on Performance Evaluation and Benchmarking, TPCTC 2010, held in conjunction with the 36th International Conference on Very Large Data Bases, VLDB 2010, in Singapore, September 13-17, 2010. The 14 full papers and two keynote papers were carefully selected and reviewed from numerous submissions. This book considers issues such as appliance; business intelligence; cloud computing; complex event processing; database optimizations; data compression; energy and space efficiency, green computing; hardware innovations; high speed data generation; hybrid workloads; very large memory systems; and virtualization.

Transaction Processing on Modern Hardware Feb 19 2020 The last decade has brought groundbreaking developments in transaction processing. This resurgence of an otherwise mature research area has spurred from the diminishing cost per GB of DRAM that allows many transaction processing workloads to be entirely memory-resident. This shift demanded a pause to fundamentally rethink the architecture of database systems. The data storage lexicon has now expanded beyond spinning disks and RAID levels to include the cache hierarchy, memory consistency models, cache coherence and write invalidation costs, NUMA regions, and coherence domains. New memory technologies promise fast non-volatile storage and expose uncharted trade-offs for transactional durability, such as exploiting byte-addressable hot and cold storage through persistent programming that promotes simpler recovery protocols. In the meantime, the plateauing single-threaded processor performance has brought massive concurrency within a single node, first in the form of multi-core, and now with many-core and heterogeneous processors. The exciting possibility to reshape the storage, transaction, logging, and recovery layers of next-generation systems on emerging hardware have prompted the database research community to vigorously

debate the trade-offs between specialized kernels that narrowly focus on transaction processing performance vs. designs that permit transactionally consistent data accesses from decision support and analytical workloads. In this book, we aim to classify and distill the new body of work on transaction processing that has surfaced in the last decade to navigate researchers and practitioners through this intricate research subject.

Big Data Computing and Communications Oct 17 2019 This book constitutes the proceedings of the Second International Conference on Big Data Computing and Communications, BigCom 2016, held in Shenyang, China, in July 2016. The 39 papers presented in this volume were carefully reviewed and selected from 90 submissions. BigCom is an international symposium dedicated to addressing the challenges emerging from big data related computing and networking. The conference is targeted to attract researchers and practitioners who are interested in Big Data analytics, management, security and privacy, communication and high performance computing in its broadest sense.

Big Data Management and Processing Aug 27 2020 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.

Innovative Research and Applications in Next-Generation High

Performance Computing Jan 20 2020 High-performance computing

(HPC) describes the use of connected computing units to perform complex tasks. It relies on parallelization techniques and algorithms to synchronize these disparate units in order to perform faster than a single processor could, alone. Used in industries from medicine and research to military and higher education, this method of computing allows for users to complete complex data-intensive tasks. This field has undergone many changes over the past decade, and will continue to grow in popularity in the coming years. Innovative Research Applications in Next-Generation High Performance Computing aims to address the future challenges, advances, and applications of HPC and related technologies. As the need for such processors increases, so does the importance of developing new ways to optimize the performance of these supercomputers. This timely publication provides comprehensive information for researchers, students in ICT, program developers, military and government organizations, and business professionals.

Data Management in the Cloud Apr 03 2021 Cloud computing has emerged as a successful paradigm of service-oriented computing and has revolutionized the way computing infrastructure is used. This success has seen a proliferation in the number of applications that are being deployed in various cloud platforms. There has also been an increase in

the scale of the data generated as well as consumed by such applications. Scalable database management systems form a critical part of the cloud infrastructure. The attempt to address the challenges posed by the management of big data has led to a plethora of systems. This book aims to clarify some of the important concepts in the design space of scalable data management in cloud computing infrastructures. Some of the questions that this book aims to answer are: the appropriate systems for a specific set of application requirements, the research challenges in data management for the cloud, and what is novel in the cloud for database researchers? We also aim to address one basic question: whether cloud computing poses new challenges in scalable data management or it is just a reincarnation of old problems? We provide a comprehensive background study of state-of-the-art systems for scalable data management and analysis. We also identify important aspects in the design of different systems and the applicability and scope of these systems. A thorough understanding of current solutions and a precise characterization of the design space are essential for clearing the "cloudy skies of data management" and ensuring the success of DBMSs in the cloud, thus emulating the success enjoyed by relational databases in traditional enterprise settings. Table of Contents: Introduction / Distributed Data Management / Cloud Data Management: Early Trends / Transactions on Co-located Data / Transactions on Distributed Data / Multi-tenant Database Systems / Concluding Remarks

Middleware 2013 Mar 02 2021 This book constitutes the refereed proceedings of the ACM/IFIP/USENIX 14th International Middleware Conference, held in Beijing, China, in December 2013. The 24 revised full papers presented were carefully reviewed and selected from 189 submissions. The papers cover a wide range of topics including design, implementation, deployment and evaluation of middleware for next-generation platforms such as cloud computing, social networks and large-scale storage and distributed systems. The middleware solutions introduced provide features such as availability, efficiency, scalability, fault-tolerance, trustworthy operation and support security and privacy needs.

Designing Data-Intensive Applications Dec 31 2020 Data is at the center of many challenges in system design today. Difficult issues need to be

figured out, such as scalability, consistency, reliability, efficiency, and maintainability. In addition, we have an overwhelming variety of tools, including relational databases, NoSQL datastores, stream or batch processors, and message brokers. What are the right choices for your application? How do you make sense of all these buzzwords? In this practical and comprehensive guide, author Martin Kleppmann helps you navigate this diverse landscape by examining the pros and cons of various technologies for processing and storing data. Software keeps changing, but the fundamental principles remain the same. With this book, software engineers and architects will learn how to apply those ideas in practice, and how to make full use of data in modern applications. Peer under the hood of the systems you already use, and learn how to use and operate them more effectively Make informed decisions by identifying the strengths and weaknesses of different tools Navigate the trade-offs around consistency, scalability, fault tolerance, and complexity Understand the distributed systems research upon which modern databases are built Peek behind the scenes of major online services, and learn from their architectures

Building a Columnar Database on RAMCloud Feb 25 2023 This book examines the field of parallel database management systems and illustrates the great variety of solutions based on a shared-storage or a shared-nothing architecture. Constantly dropping memory prices and the desire to operate with low-latency responses on large sets of data paved the way for main memory-based parallel database management systems. However, this area is currently dominated by the shared-nothing approach in order to preserve the in-memory performance advantage by processing data locally on each server. The main argument this book makes is that such an unilateral development will cease due to the combination of the following three trends: a) Today's network technology features remote direct memory access (RDMA) and narrows the performance gap between accessing main memory on a server and of a remote server to and even below a single order of magnitude. b) Modern storage systems scale gracefully, are elastic and provide high-availability. c) A modern storage system such as Stanford's RAM Cloud even keeps all data resident in the main memory. Exploiting these characteristics in the context of a main memory-based parallel database

management system is desirable. The book demonstrates that the advent of RDMA-enabled network technology makes the creation of a parallel main memory DBMS based on a shared-storage approach feasible. *Database Systems for Advanced Applications* Jul 18 2022 This two volume set LNCS 9642 and LNCS 9643 constitutes the refereed proceedings of the 21st International Conference on Database Systems for Advanced Applications, DASFAA 2016, held in Dallas, TX, USA, in April 2016. The 61 full papers presented were carefully reviewed and selected from a total of 183 submissions. The papers cover the following topics: crowdsourcing, data quality, entity identification, data mining and machine learning, recommendation, semantics computing and knowledge base, textual data, social networks, complex queries, similarity computing, graph databases, and miscellaneous, advanced applications.

Advanced Parallel Processing Technologies Dec 23 2022 This book constitutes the proceedings of the 11th International Symposium on Advanced Parallel Processing Technologies, APPT 2015, held in Jinan, China, in August 2015. The 8 papers presented in this volume were carefully reviewed and selected from 24 submissions. They deal with the recent advances in big data processing; parallel architectures and systems; parallel software; parallel algorithms and applications; and distributed and cloud computing.

High-Performance Big Data Computing May 24 2020 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.

Principles of Distributed Database Systems Mar 22 2020 This third edition of a classic textbook can be used to teach at the senior undergraduate and graduate levels. The material concentrates on fundamental theories as well as techniques and algorithms. The advent of the Internet and the World Wide Web, and, more recently, the emergence of cloud computing and streaming data applications, has forced a renewal of interest in distributed and parallel data management, while, at the same time, requiring a rethinking of some of the traditional techniques. This book covers the breadth and depth of this re-emerging field. The coverage consists of two parts. The first part discusses the fundamental principles of distributed data management and includes distribution design, data integration, distributed query processing and optimization, distributed transaction management, and replication. The second part focuses on more advanced topics and includes discussion of parallel database systems, distributed object management, peer-to-peer data management, web data management, data stream systems, and cloud computing. New in this Edition: • New chapters, covering database replication, database integration, multidatabase query processing, peer-to-peer data management, and web data management. • Coverage of emerging topics such as data streams and cloud computing • Extensive revisions and updates based on years of class testing and feedback Ancillary teaching materials are available.

Non-Volatile Memory Database Management Systems Nov 22 2022 This book explores the implications of non-volatile memory (NVM) for database management systems (DBMSs). The advent of NVM will fundamentally change the dichotomy between volatile memory and

durable storage in DBMSs. These new NVM devices are almost as fast as volatile memory, but all writes to them are persistent even after power loss. Existing DBMSs are unable to take full advantage of this technology because their internal architectures are predicated on the assumption that memory is volatile. With NVM, many of the components of legacy DBMSs are unnecessary and will degrade the performance of data-intensive applications. We present the design and implementation of DBMS architectures that are explicitly tailored for NVM. The book focuses on three aspects of a DBMS: (1) logging and recovery, (2) storage and buffer management, and (3) indexing. First, we present a logging and recovery protocol that enables the DBMS to support near-instantaneous recovery. Second, we propose a storage engine architecture and buffer management policy that leverages the durability and byte-addressability properties of NVM to reduce data duplication and data migration. Third, the book presents the design of a range index tailored for NVM that is latch-free yet simple to implement. All together, the work described in this book illustrates that rethinking the fundamental algorithms and data structures employed in a DBMS for NVM improves performance and availability, reduces operational cost, and simplifies software development.

Running a Transactional Database on Top of RamCloud Jan 24 2023

Big Data Benchmarking Jun 05 2021 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.

- [Building A Columnar Database On RAMCloud](#)
- [Running A Transactional Database On Top Of RamCloud](#)
- [Advanced Parallel Processing Technologies](#)
- [Non Volatile Memory Database Management Systems](#)

- [The Use Of A Low latency Key value Store For Implementing A Scalable And High performance Graph Database](#)
- [Performance Evaluation And Benchmarking](#)
- [Encyclopedia Of Business Analytics And Optimization](#)
- [Database Systems For Advanced Applications](#)
- [Main Memory Management On Relational Database Systems](#)
- [Database And Expert Systems Applications](#)
- [Database Systems For Advanced Applications](#)
- [Storage Systems](#)
- [Beyond Databases Architectures And Structures Paving The Road To Smart Data Processing And Analysis](#)
- [Beyond Databases Architectures And Structures Towards Efficient Solutions For Data Analysis And Knowledge Representation](#)
- [Web Age Information Management](#)
- [Database Systems For Advanced Applications](#)
- [Handbook Of Research On Cloud Infrastructures For Big Data Analytics](#)
- [Computer Engineering And Technology](#)
- [Achieving Both Low Latency And Strong Consistency In Large scale Systems](#)
- [An Architecture For Fast And General Data Processing On Large Clusters](#)
- [Big Data Benchmarking](#)
- [Information Systems Development](#)
- [Data Management In The Cloud](#)
- [Middleware 2013](#)
- [International Joint Conference](#)
- [Designing Data Intensive Applications](#)
- [Cloud Computing](#)
- [Intelligent Distributed Computing XIII](#)
- [Enabling Real Time Business Intelligence](#)
- [Big Data Management And Processing](#)
- [Encyclopedia Of Information Science And Technology Fourth Edition](#)
- [In Memory Data Management](#)

- [High Performance Big Data Computing](#)
- [Information Security And Privacy](#)
- [Principles Of Distributed Database Systems](#)
- [Transaction Processing On Modern Hardware](#)
- [Innovative Research And Applications In Next Generation High Performance Computing](#)
- [Fast And Scalable Cloud Data Management](#)
- [Advanced Methodologies And Technologies In Network Architecture Mobile Computing And Data Analytics](#)
- [Big Data Computing And Communications](#)