

Time Series Databases: New Ways To Store And Access Data

The world is becoming more and more instrumented, interconnected, and intelligent in what IBM® terms a smarter planet, with more and more data being collected for analysis. In trade magazines, this trend is called big data. As part of this trend, the following types of time-based information are collected: Large data centers support a corporation or provide cloud services. These data centers need to collect temperature, humidity, and other types of information over time to optimize energy usage. Utility meters (referred to as smart meters) allow utility companies to collect information over a wireless network and to collect more data than ever before. IBM Informix® TimeSeries is optimized for the processing of time-based data and can provide the following benefits: Storage savings: Storage can be optimized when you know the characteristics of your time-based data. Informix TimeSeries offers one third of the storage space that is required by a standard relational database. Query performance: Informix TimeSeries takes into consideration the type of data to optimize its organization on disk and eliminates the need for some large indexes and additional sorting. For these reasons and more, some queries can easily have an order of magnitude performance improvement compared to standard relational. Simpler queries: Informix TimeSeries includes a large set of specialized functions that allow you to better express the processing that you want to execute. It even provides a toolkit so that you can add proprietary algorithms to the library. This IBM Redbooks® publication is for people who want to implement a solution that revolves around time-based data. It gives you the information that you need to get started and be productive with Informix TimeSeries.

"With an easy, step-by-step approach, this guide shows beginners how to install, use, and maintain the world's most popular open source database: MySQL. You'll learn through real-world examples and many practical tips, including information on how to improve database performance. Database systems such as MySQL help data handling for organizations large and small handle data, providing robust and efficient access in ways not offered by spreadsheets and other types of data stores. This book is also useful for web developers and programmers interested in adding MySQL to their skill sets. Topics include: Installation and basic administration; Introduction to relational and SQL; Functions, subqueries, and other query enhancements; Improving database performance; Accessing MySQL from popular languages" --

"Addresses the evolution of database management, technologies and applications along with the progress and endeavors of new research areas." --P. xiii.

Discover how graph databases can help you manage and query highly connected data. With this practical book, you'll learn how to design and implement a graph database that brings the power of graphs to bear on a broad range of problem domains. Whether you want to speed up your response to user queries or build a database that can adapt as your business evolves, this book shows you how to apply the schema-free graph model to real-world problems. Learn how different organizations are using graph databases to solve a variety of business and research issues. You'll learn about early decisions and pre-planning that can make the process easier and more productive. If you're already using these technologies, you'll discover ways to gain the full range of benefits possible with Hadoop. While you don't need a deep technical background to get started, this book does provide expert guidance to help managers, architects, and practitioners succeed with their Hadoop projects. Examine a day in the life of big data: India's ambitious Aadhaar project Review tools in the Hadoop ecosystem such as Apache's Spark, Storm, and Drill to learn how they can help you Pick up a collection of technical and strategic tips that have helped others succeed with Hadoop Learn from several prototypical Hadoop use cases, based on how organizations have actually applied the technology Explore real-world stories that reveal how MapR customers combine use cases when putting Hadoop and NoSQL to work, including in production

Information Modeling and Relational Databases provides an introduction to ORM (Object Role Modeling)-and much more. In fact, it's the only book to go beyond introductory coverage and provide all of the in-depth instruction you need to transform knowledge from domain experts into a sound database design. Inside, ORM authority Terry Halpin blends conceptual information with practical instruction that will let you begin using ORM effectively as soon as possible. Supported by examples, exercises, and useful background information, his step-by-step approach teaches you to develop a natural-language-based ORM model and then, where needed, abstract ER and UML models from it. This book will quickly make you proficient in the modeling technique that is proving vital to the development of accurate and efficient databases that best meet real business objectives. The most in-depth coverage of Object Role Modeling available anywhere-written by a pioneer in the development of ORM. Provides additional coverage of Entity Relationship (ER) modeling and the Unified Modeling Language-all from an ORM perspective. Intended for anyone with a stake in the accuracy and efficacy of databases: systems analysts, information modelers, database designers and administrators, instructors, managers, and programmers. Explains and illustrates required concepts from mathematics and set theory.

Advances in Security, Networks, and Internet of Things

Seven NoSQL Databases in a Week

Principles of Database Management

Research Methods and Case Studies

Heading in the Right Direction with MySQL and MariaDB

Real-World Hadoop

While traditional databases excel at complex queries over historical data, they are inherently pull-based and therefore ill-equipped to push new information to clients. Systems for data stream management and processing, on the other hand, are natively push-oriented and thus facilitate reactive behavior. However, they do not retain data indefinitely and are therefore not able to answer historical queries. The book provides an overview over the different (push-based) mechanisms for data retrieval in each system class and the semantic differences between them. It also provides a comprehensive overview over the current state of the art in real-time databases. It sfirst includes an in-depth system survey of today's real-time databases:

Discovery in Databases;Graphs;Graph Similarity;String Distance;Machine Learning;Segmentation;Change Detection
The two-volume set (LNAI 10448 and LNAI 10449) constitutes the refereed proceedings of the 9th International Conference on Collective Intelligence, ICCCI 2017, held in Nicosia, Cyprus, in September 2017. The 117 full papers presented were carefully reviewed and selected from 248 submissions. The conference focussed on the methodology and applications of computational collective intelligence, included: multi-agent systems, knowledge engineering and semantic web, social networks and recommender systems, text processing and information retrieval, data mining methods and applications, sensor networks and internet of things, decision support & control systems, and computer vision techniques.

When it comes to choosing, using, and maintaining a database, understanding its internals is essential. But with so many distributed databases and tools available today, it's often difficult to understand what each one offers and how they differ. With this practical guide, Alex Petrov guides developers through the concepts behind modern database and storage engine internals. Throughout the book, you'll explore relevant material gleaned from numerous books, papers, blog posts, and the source code of several open source databases. These resources are listed at the end of parts one and two. You'll discover that the most significant distinctions among many modern databases reside in subsystems that determine how storage is organized and how data is distributed. This book examines: Storage engines: Explore storage classification and taxonomy, and dive into B-Tree-based and immutable Log Structured storage engines, with differences and use-cases for each Storage building blocks: Learn how database files are organized to build efficient storage, using auxiliary data structures such as Page Cache, Buffer Pool and Write-Ahead Log Distributed systems: Learn step-by-step how nodes and processes connect and build complex communication patterns Database clusters: Which consistency models are commonly used by modern databases and how distributed storage systems achieve consistency

The need to support large-scale time series data is increasing rapidly. There are emerg- ing Time Series Databases built with conventional relational databases or newer NoSQL databases. The ScalaTion Time Series Database is built on top of its column-oriented in-memory database. ScalaTion is an open-source Scala based big data framework for simulation, optimization and analytics. This database provides support for large-scale stor- age, efficient query processing, pattern matching and a variety of forecasting techniques. Its design goals include the ability to scale up and scale out, and the ability to handle conven- tional multivariate time series. The database provides an easy way to transform a table into a matrix (or vector) which may be used as input for other data science/machine-learning models that are available in ScalaTion. The capabilities are illustrated via a case study of vehicle traffic forecasting. Multiple experiments are conducted to evaluate the performances of four databases: ScalaTion, MySQL, SQLite, and SparkSQL.

First European Symposium, PKDD '97, Trondheim, Norway, June 24-27, 1997 Proceedings
19th International Conference, IC3P'99 2019, Melbourne, VIC, Australia, December 9 - 11 2019, Proceedings, Part II
27th International Conference, ICONIP 2020, Bangkok, Thailand, November 23 - 27, 2020, Proceedings, Part III

High Performance Discovery In Time Series

Principles and Practice of Constraint Programming - CP '95

First International Conference, CP '95, Cassis, France, September 19 - 22, 1995, Proceedings

Solving Business Problems with Informix TimeSeries

The three-volume set of LNCS 12532, 12533, and 12534 constitutes the proceedings of the 27th International Conference on Neural Information Processing, ICONP 2020, held in Bangkok, Thailand, in November 2020. Due to COVID-19 pandemic the conference was held virtually. The 187 full papers presented were carefully reviewed and selected from 618 submissions. The papers address the emerging topics of theoretical research, empirical studies, and applications of neural information processing techniques across different domains. The third volume, LNCS 12534, is organized in topical sections on biomedical information; neural data science; network models; recommender systems; time series analysis.

This new edition updates Durbin & Koopman's important text on the state space approach to time series analysis. The distinguishing feature of state space time series models is that observations are regarded as made up of distinct components such as trend, seasonal, regression elements and disturbance terms, each of which is modelled separately. The techniques that emerge from this approach are very flexible and are capable of handling a much wider range of problems than the main analytical system currently in use for time series analysis, the Box-Jenkins ARIMA system. Additions to this second edition include the filtering of nonlinear and non-Gaussian series. Part I of the book obtains the mean and variance of the state, of a variable intended to measure the effect of an interaction and of regression coefficients, in terms of the observations. Part II extends the treatment to nonlinear and non-normal models. For these, analytical solutions are not available so methods are based on simulation.

This book focuses on different facets of flight data analysis, including the basic goals, methods, and implementation techniques. As mass flight data possesses the typical characteristics of time series, the time series analysis methods and their application for flight data have been illustrated from several different perspectives: flight data network models; recommender systems; time series analysis. A flight data analysis has been established to assist in aircraft condition monitoring, training evaluation and scientific maintenance. The book will serve as a reference resource for people working in aviation management and maintenance, as well as researchers and engineers in the fields of data analysis and data mining.

Foundations of Computational Intelligence Volume 6: Data Mining: Theoretical Foundations and Applications Finding information hidden in data is as theoretically difficult as it is practically important. With the objective of discovering unknown patterns from data, the methodologies of data mining were derived from statistics, machine learning, and artificial intelligence, and are being used successfully in application areas such as bioinformatics, business, health care, banking, retail, and many others. Advanced representation schemes and computational intelligence techniques such as rough sets, neural networks; decision trees; fuzzy logic; evolutionary algorithms; arti- cial immune systems; swarm intelligence; reinforcement learning, association rule mining, Web intelligence paradigms etc. have proved valuable when they are - plied to Data Mining problems. Computational tools or solutions based on intel- gent systems are being used to great success in Data Mining applications. It is also observed that strong scientific advances have been made when issues from different research areas are integrated. This Volume comprises of 15 chapters including an overview chapter providing an up-to-date and state-of-the research on the applications of Computational Intelligence techniques for Data Mining. The book is divided into 3 parts: Part-I: Data Click Streams and Temporal Data Mining Part-II: Text and Rule Mining Part-III: Applications Part I on Data Click Streams and Temporal Data Mining contains four chapters that describe several approaches in Data Click Streams and Temporal Data Mining.

Temporal Databases in the Relational Model and SQL

Cloud Native Java

Database Internals

A Deep Dive into How Distributed Data Systems Work

Web Database Applications with PHP and MySQL

Temporal Databases: Research and Practice

Information Modeling and Relational Databases

This book presents a selection of peer-reviewed contributions on the latest advances in time series analysis, presented at the International Conference on Time Series and Forecasting (ITISE 2019), held in Granada, Spain, on September 25-27, 2019. The first two parts of the book present theoretical contributions on statistical and advanced mathematical methods, and on econometric models, financial forecasting and risk analysis. The remaining four parts include practical contributions on time series analysis in energy; complex/big data time series and forecasting; time series analysis with computational intelligence; and time series analysis and prediction for other real-world problems. Given this mix of topics, readers will acquire a more comprehensive perspective on the field of time series analysis and forecasting. The ITISE conference series provides a forum for scientists, engineers, educators and students to discuss the latest advances and implementations in the foundations, theory, models and applications of time series analysis and forecasting. It focuses on interdisciplinary research encompassing computer science, mathematics, statistics and econometrics.

Describes the features and functions of Apache Hive, the data infrastructure for Hadoop.
"Many applications generate large datasets, like social networking and social influence programs, smart cities applications, smart house environments, Cloud applications, public web sites, scientific experiments and simulations, data warehouse, monitoring platforms, and e-government services. This book will primarily encompass practical approaches that advance research in all aspects of data processing, data analytics, data processing in Cloud/Edge/Fog systems, having a large variety of tools and software to manage them. The book focuses on focuses on topics covering algorithms, architectures, management models, high performance computing techniques and large-scale distributed systems"

The book introduces the topic of neurobiology, describing animal nervous systems, what they consist of, how they work, and how they are studied. Unlike many other neurobiology texts, considerable discussion is given to both human and non-human nervous systems. Written in an easy-to-read style, it will be useful for both biology and medical students. It provides the opportunity for self-testing at the end of each chapter, with objectives and questions. A CD-ROM entitled "The Human Brain" (ISBN 3-540-14666-0) has been produced to accompany this text, and can be purchased either separately or together with the book (ISBN 3-540-63778-8).

Graph Databases

Neural Information Processing

Data Science and Big Data Analytics in Smart Environments

Proceedings from SAM'20, ICWN'20, ICOMP'20, and ESCS'20

Programming Hive

Time Series Databases: New Ways to Store and Access Data

Seven Databases in Seven Weeks

Time series data analysis is increasingly important due to the massive production of such data through the internet of things, the digitalization of healthcare, and the rise of smart cities. As continuous monitoring and data collection become more common, the need for competent time series analysis with both statistical and machine learning techniques will increase. Covering innovations in time series data analysis and use cases from the real world, this practical guide will help you solve the most common data engineering and analysis challenges in time series, using both traditional statistical and modern machine learning techniques. Author Aileen Nielsen offers an accessible, well-rounded introduction to time series in both R and Python that will have data scientists, software engineers, and researchers up and running quickly. You'll get the guidance you need to confidently: Find and wrangle time series data Undertake exploratory time series data analysis Store temporal data Simulate time series data Generate and select features for a time series Measure error Forecast and classify time series with machine or deep learning Evaluate accuracy and performance

The book presents the proceedings of four conferences: The 19th International Conference on Security & Management (SAM'20), The 19th International Conference on Wireless Networks (ICWN'20), The 21st International Conference on Internet Computing & Internet of Things (ICOIMP'20), and The 18th International Conference on Embedded Systems, Cyber-physical Systems (ESCS'20). The conferences took place in Las Vegas, NV, USA, July 27-30, 2020. The conferences are part of the larger 2020 World Congress in Computer Science, Computer Engineering, & Applied Computing (CSCE'20), which features 20 major tracks. Authors include academics, researchers, professionals, and students. Presents the proceedings of four conferences as part of the 2020 World Congress in Computer Science, Computer Engineering, & Applied Computing (CSCE'20); Includes the tracks on security & management, wireless networks, internet computing and IoT, and embedded systems as well as cyber-physical systems; Features papers from SAM ' 20, ICWN ' 20, ICOMP ' 20 and ESCS ' 20.

' Big data ' is now readily available to economic historians, thanks to the digitisation of primary sources, collaborative research linking different data sets, and the publication of databases on the internet. Key economic indicators, such as the consumer price index, can be tracked over long periods, and qualitative information, such as land use, can be converted to a quantitative form. In order to fully exploit these innovations it is necessary to use sophisticated statistical techniques to reveal the patterns hidden in datasets, and this book shows how this can be done. A distinguished group of economic historians have teamed up with younger researchers to pilot the application of new techniques to ' big data '. Topics addressed in this volume include prices and the standard of living, money supply, credit markets, land values and land use, transport, technological innovation, and business networks. The research spans the medieval, early modern and modern periods. Research methods include simultaneous equation systems, stochastic trends and discrete choice modelling. This book is essential reading for doctoral and post-doctoral researchers in business, economic and social history. The case studies will also appeal to historical geographers and applied econometricians.

Adding the time dimension to real-world databases produces Time Series Databases (TSDB) and introduces new aspects and difficulties to data mining and knowledge discovery. This book covers the state-of-the-art methodology for mining time series databases. The novel data mining methods presented in the book include techniques for efficient segmentation, indexing, and classification of noisy and dynamic time series. A graph-based method for anomaly detection in time series is described and the book also studies the implications of a novel and potentially useful representation of time series as strings. The problem of detecting changes in data mining models that are induced from temporal databases is additionally discussed. Contents:Segmenting Time Series: A Survey and Novel Approach (E Keogh et al.)A Survey of Recent Methods for Efficient Retrieval of Similar Time Sequences (M L Hetland)Indexing of Compressed Time Series (E Fink & K Pratt)Indexing Time-Series under Conditions of Noise (M Vlachos et al.)Change Detection in Classification Models Induced from Time Series Data (G Zeira et al.)Classification and Detection of Abnormal Events in Time Series of Graphs (H Bunke & M Kraetzl)Boosting Interval-Based Literals: Variable Length and Early Classification (C Alonso Gonz ález & J J Rodr í guez Diez)Median Strings: A Review (X Jiang et al.) Readership: Graduate students, researchers and practitioners in the fields of data mining, machine learning, databases and statistics. Keywords:Times Series;Time Series Analysis;Data Mining, Knowledge Discovery in Databases;Graphs;Graph Similarity;String Distance;Machine Learning;Segmentation;Change Detection

Temporal Databases in the Relational Model and SQL
Cloud Native Java
Database Internals
A Deep Dive into How Distributed Data Systems Work
Web Database Applications with PHP and MySQL
Temporal Databases: Research and Practice
Information Modeling and Relational Databases

This book presents a selection of peer-reviewed contributions on the latest advances in time series analysis, presented at the International Conference on Time Series and Forecasting (ITISE 2019), held in Granada, Spain, on September 25-27, 2019. The first two parts of the book present theoretical contributions on statistical and advanced mathematical methods, and on econometric models, financial forecasting and risk analysis. The remaining four parts include practical contributions on time series analysis in energy; complex/big data time series and forecasting; time series analysis with computational intelligence; and time series analysis and prediction for other real-world problems. Given this mix of topics, readers will acquire a more comprehensive perspective on the field of time series analysis and forecasting. The ITISE conference series provides a forum for scientists, engineers, educators and students to discuss the latest advances and implementations in the foundations, theory, models and applications of time series analysis and forecasting. It focuses on interdisciplinary research encompassing computer science, mathematics, statistics and econometrics.

Describes the features and functions of Apache Hive, the data infrastructure for Hadoop.
"Many applications generate large datasets, like social networking and social influence programs, smart cities applications, smart house environments, Cloud applications, public web sites, scientific experiments and simulations, data warehouse, monitoring platforms, and e-government services. This book will primarily encompass practical approaches that advance research in all aspects of data processing, data analytics, data processing in Cloud/Edge/Fog systems, having a large variety of tools and software to manage them. The book focuses on focuses on topics covering algorithms, architectures, management models, high performance computing techniques and large-scale distributed systems"

The book introduces the topic of neurobiology, describing animal nervous systems, what they consist of, how they work, and how they are studied. Unlike many other neurobiology texts, considerable discussion is given to both human and non-human nervous systems. Written in an easy-to-read style, it will be useful for both biology and medical students. It provides the opportunity for self-testing at the end of each chapter, with objectives and questions. A CD-ROM entitled "The Human Brain" (ISBN 3-540-14666-0) has been produced to accompany this text, and can be purchased either separately or together with the book (ISBN 3-540-63778-8).

Graph Databases

Neural Information Processing

Data Science and Big Data Analytics in Smart Environments

Proceedings from SAM'20, ICWN'20, ICOMP'20, and ESCS'20

Programming Hive

Time Series Databases: New Ways to Store and Access Data

Seven Databases in Seven Weeks

Time series data analysis is increasingly important due to the massive production of such data through the internet of things, the digitalization of healthcare, and the rise of smart cities. As continuous monitoring and data collection become more common, the need for competent time series analysis with both statistical and machine learning techniques will increase. Covering innovations in time series data analysis and use cases from the real world, this practical guide will help you solve the most common data engineering and analysis challenges in time series, using both traditional statistical and modern machine learning techniques. Author Aileen Nielsen offers an accessible, well-rounded introduction to time series in both R and Python that will have data scientists, software engineers, and researchers up and running quickly. You'll get the guidance you need to confidently: Find and wrangle time series data Undertake exploratory time series data analysis Store temporal data Simulate time series data Generate and select features for a time series Measure error Forecast and classify time series with machine or deep learning Evaluate accuracy and performance

The book presents the proceedings of four conferences: The 19th International Conference on Security & Management (SAM'20), The 19th International Conference on Wireless Networks (ICWN'20), The 21st International Conference on Internet Computing & Internet of Things (ICOIMP'20), and The 18th International Conference on Embedded Systems, Cyber-physical Systems (ESCS'20). The conferences took place in Las Vegas, NV, USA, July 27-30, 2020. The conferences are part of the larger 2020 World Congress in Computer Science, Computer Engineering, & Applied Computing (CSCE'20), which features 20 major tracks. Authors include academics, researchers, professionals, and students. Presents the proceedings of four conferences as part of the 2020 World Congress in Computer Science, Computer Engineering, & Applied Computing (CSCE'20); Includes the tracks on security & management, wireless networks, internet computing and IoT, and embedded systems as well as cyber-physical systems; Features papers from SAM ' 20, ICWN ' 20, ICOMP ' 20 and ESCS ' 20.

' Big data ' is now readily available to economic historians, thanks to the digitisation of primary sources, collaborative research linking different data sets, and the publication of databases on the internet. Key economic indicators, such as the consumer price index, can be tracked over long periods, and qualitative information, such as land use, can be converted to a quantitative form. In order to fully exploit these innovations it is necessary to use sophisticated statistical techniques to reveal the patterns hidden in datasets, and this book shows how this can be done. A distinguished group of economic historians have teamed up with younger researchers to pilot the application of new techniques to ' big data '. Topics addressed in this volume include prices and the standard of living, money supply, credit markets, land values and land use, transport, technological innovation, and business networks. The research spans the medieval, early modern and modern periods. Research methods include simultaneous equation systems, stochastic trends and discrete choice modelling. This book is essential reading for doctoral and post-doctoral researchers in business, economic and social history. The case studies will also appeal to historical geographers and applied econometricians.

Adding the time dimension to real-world databases produces Time Series Databases (TSDB) and introduces new aspects and difficulties to data mining and knowledge discovery. This book covers the state-of-the-art methodology for mining time series databases. The novel data mining methods presented in the book include techniques for efficient segmentation, indexing, and classification of noisy and dynamic time series. A graph-based method for anomaly detection in time series is described and the book also studies the implications of a novel and potentially useful representation of time series as strings. The problem of detecting changes in data mining models that are induced from temporal databases is additionally discussed. Contents: A Survey of Recent Methods for Efficient Retrieval of Similar Time Sequences (E Fink & K Pratt); Boosting Interval-Based Literal: Variable Length and Early Classification (J J Rodríguez Diez); Segmenting Time Series: A Survey and Novel Approach (E Keogh et al.); Indexing Similar Time Series under Conditions of Noise (M Vlachos et al.); Classification of Events in Time Series of Graphs (H Bunke & M Kraetzl); Median Strings--A Review (X Jiang et al.); Change Detection in Classification Models of Data Mining (G Zeira et al.). Readership: Graduate students, researchers and practitioners in the fields of data mining, machine learning, databases and statistics.

Monitoring with Graphite

Volume 6: Data Mining

Learning in the 21st Century

7th European Conference on Principles and Practice of Knowledge Discovery in Databases, Cavtat-Dubrovnik, Croatia, September 22-26, 2003, Proceedings

Prediction with Statistics and Machine Learning

4th International Workshop, TMA 2012, Vienna, Austria, March 12, 2012, Proceedings

Tracking Dynamic Host and Application Metrics at Scale

This book constitutes the refereed proceedings of the First European Symposium on Principles of Data Mining and Knowledge Discovery, PKDD '97, held in Trondheim, Norway, in June 1997. The volume presents a total of 38 revised full papers together with abstracts of one invited talk and four tutorials. Among the topics covered are data and knowledge representation, statistical and probabilistic methods, logic-based approaches, man-machine interaction aspects, AI contributions, high performance computing support, machine learning, automated scientific discovery, quality assessment, and applications. Graphite has become one of the most powerful monitoring tools available today, due to its ease of use, rapid graph prototyping abilities, and a friendly rendering API. With this practical guide, system administrators and engineers will learn how to use this open source tool to track operational data you need to monitor your systems, as well as application-level metrics for profiling your services. Author Jason Dixon, member of the Graphite project, provides a thorough introduction of Graphite from the basics to the skills and tools you need for troubleshooting and scaling out its software components. If you want to learn more about monitoring systems, services, or applications, this is the book you need. Get an introduction to monitoring, including important concepts and terminology Examine the features and functionality of key Graphite components, including Carbon and Whisper Learn the typical user workflow necessary to create a basic line chart Build complex charts with chained functions and multiple axes that interact directly with the rendering API Understand how to use the native Graphite dashboard, as well as the more popular third-party dashboards Master the art of scaling and troubleshooting high-performance or highly available Graphite clusters

Time series data is of growing importance, especially with the rapid expansion of the Internet of Things. This concise guide shows you effective ways to collect, persist, and access large-scale time series data for analysis. You'll explore the theory behind time series databases and learn practical methods for implementing them. Authors Ted Dunning and Ellen Friedman provide a detailed examination of open source tools such as OpenTSDB and new modifications that greatly speed up data ingestion. You'll learn: A variety of time series use cases The advantages of NoSQL databases for large-scale time series data NoSQL table design for high-performance time series databases The benefits and limitations of OpenTSDB How to access data in OpenTSDB using R, Go, and Ruby How time series databases contribute to practical machine learning projects How to handle the added complexity of geo-temporal data For advice on analyzing time series data, check out Practical Machine Learning: A New Look at Anomaly Detection, also from Ted Dunning and Ellen Friedman.

What separates the traditional enterprise from the likes of Amazon, Netflix, and Etsy? Those companies have refined the art of cloud native development to maintain their competitive edge and stay well ahead of the competition. This practical guide shows Java/JVM developers how to build better software, faster, using Spring Boot, Spring Cloud, and Cloud Foundry. Many organizations have already waded into cloud computing, test-driven development, microservices, and continuous integration and delivery. Authors Josh Long and Kenny Bastani fully immerse you in the tools and methodologies that will help you transform your legacy application into one that is genuinely cloud native. In four sections, this book takes you through: The Basics: learn the motivations behind cloud native thinking; configure and test a Spring Boot application; and move your legacy application to the cloud Web Services: build HTTP and RESTful services with Spring; route requests in your distributed system; and build edge services closer to the data Data Integration: manage your data with Spring Data, and integrate distributed services with Spring's support for event-driven, messaging-centric architectures Production: make your system observable; use service brokers to connect stateful services; and understand the big ideas behind continuous delivery

Principles of Data Mining and Knowledge Discovery

The Scalation Time Series Database: Support for Big Data Analytics

Push-Based Data in Research & Practice

Get up and running with the fundamentals and functionalities of seven of the most popular NoSQL databases

From Conceptual Analysis to Logical Design

Foundations of Computational Intelligence

Computational Collective Intelligence

Data is getting bigger and more complex by the day, and so are your choices in handling it. Explore some of the most cutting-edge databases available - from a traditional relational database to newer NoSQL approaches - and make informed decisions about challenging data storage problems. This is the only comprehensive guide to the world of NoSQL databases, with in-depth practical and conceptual introductions to seven different technologies: Redis, Neo4J, CouchDB, MongoDB, HBase, Postgres, and DynamoDB. This second edition includes a new chapter on DynamoDB and updated content for each chapter. While relational databases such as MySQL remain as relevant as ever, the alternative, NoSQL paradigm has opened up new horizons in performance and scalability and changed the way we approach data-centric problems. This book presents the essential concepts behind each database alongside hands-on examples that make each technology come alive. With each database, tackle a real-world problem that highlights the concepts and features that make it shine. Along the way, explore five database models - relational, key/value, columnar, document, and graph - from the perspective of challenges faced by real applications. Learn how MongoDB and CouchDB are strikingly different, make your applications faster with Redis and more connected with Neo4J, build a cluster of HBase servers using cloud services such as Amazon's Elastic MapReduce, and more. This new edition brings a brand new chapter on DynamoDB, updated code samples and exercises, and a more up-to-date account of each database's feature set. Whether you're a programmer building the next big thing, a data scientist seeking solutions to thorny problems, or a technology enthusiast venturing into new territory, you will find something to inspire you in this book. What You Need: You'll need a *nix shell (Mac OS or Linux preferred, Windows users will need Cygwin), Java 6 (or greater), and Ruby 1.8.7 (or greater). Each chapter will list the downloads required for that database.

Time and Relational Theory provides an in-depth description of temporal database systems, which provide special facilities for storing, querying, and updating historical and future data. Traditionally, database management systems provide little or no special support for temporal data at all. This situation is changing because: Cheap storage enables retention of large volumes of historical data in data warehouses Users are now faced with temporal data problems, and need solutions Temporal features have recently been incorporated into the SQL standard, and vendors have begun to add temporal support to their DBMS products Based on the groundbreaking text Temporal Data & the Relational Model (Morgan Kaufmann, 2002) and new research led by the authors, Time and Relational Theory is the only book to offer a complete overview of the functionality of a temporal DBMS. Expert authors Nikos Lorentzos, Hugh Darwin, and Chris Date describe an approach to temporal database management that is firmly rooted in classical relational theory and will stand the test of time. This book covers the SQL:2011 temporal extensions in depth and identifies and discusses the temporal functionality still missing from SQL. Understand how the relational model provides an ideal basis for taming the complexities of temporal databases Learn how to analyze and evaluate commercial temporal products with this timely and important information Be able to use sound principles in designing and using temporal databases Understand the temporal support recently added to SQL with coverage of the new SQL features in this unique, accurate, and authoritative reference Appreciate the benefits of a truly relational approach to the problem with this clear, user friendly presentation

A beginner's guide to get you up and running with Cassandra, DynamoDB, HBase, InfluxDB, MongoDB, Neo4j, and Redis Key Features Covers the basics of 7 NoSQL databases and how they are used in the enterprises Quick introduction to MongoDB, DynamoDB, Redis, Cassandra, Neo4j, InfluxDB, and HBase Includes effective techniques for database querying and management Book Description This is the golden age of open source NoSQL databases. With enterprises having to work with large amounts of unstructured data and moving away from expensive monolithic architecture, the adoption of NoSQL databases is rapidly increasing. Being familiar with the popular NoSQL databases and knowing how to use them is a must for budding DBAs and developers. This book introduces you to the different types of NoSQL databases and gets you started with seven of the most popular NoSQL databases used by enterprises today. We start off with a brief overview of what NoSQL databases are, followed by an explanation of why and when to use them. The book then covers the seven most popular databases in each of these categories: MongoDB, Amazon DynamoDB, Redis, HBase, Cassandra, InfluxDB, and Neo4j. The book doesn't go into too much detail about each database but teaches you enough to get started with them. By the end of this book, you will have a thorough understanding of the different NoSQL databases and their functionalities, empowering you to select and use the right database according to your needs. What you will learn Understand how MongoDB provides high-performance, high-availability, and automatic scaling Interact with your Neo4j instances via database queries, Python scripts, and Java application code Get familiar with common querying and programming methods to interact with Redis Study the different types of problems Cassandra can solve Work with HBase components to support common operations such as creating tables and reading/writing data Discover data models and work with CRUD operations using DynamoDB Discover what makes InfluxDB a great choice for working with time-series data Who this book is for If you are a budding DBA or a developer who wants to get started with the fundamentals of NoSQL databases, this book is for you. Relational DBAs who want to get insights into the various offerings of popular NoSQL databases will also find this book to be very useful.

The authors set forth the theory and rationale behind adopting a Guided Inquiry approach to PreK-12 education, as well as the expertise, roles and responsibilities of each member of the instructional team.

Practical Time Series Analysis

Theory and Applications of Time Series Analysis

A Comparison of NoSQL Time Series Databases

9th International Conference, ICCCI 2017, Nicosia, Cyprus, September 27-29, 2017, Proceedings, Part II

Learning MySQL and MariaDB

A Guide to Modern Databases and the NoSQL Movement

Traffic Monitoring and Analysis

This monograph is a technical survey of concepts and techniques for describing and analyzing large-scale time-series data streams. Some topics covered are algorithms for query by humming, gamma-ray burst detection, pairs trading, and density detection. Included are self-contained descriptions of wavelets, fast Fourier transforms, and sketches as they apply to time-series analysis. Detailed applications are built on a solid scientific basis. The proceedings of ECLM/PKDD2003 are published in two volumes: the P- ceedings of the 14th European Conference on Machine Learning (LNAI 2837) and the Proceedings of the 7th European Conference on Principles and Practice of Knowledge Discovery in Databases (LNAI 2838). The two conferences were held on September 22-26, 2003 in Cavtat, a small tourist town in the vicinity of Dubrovnik, Croatia. As machine learning and knowledge discovery are two highly related ?elds, theco-locationofbothconferencesisbeneficialforbothresearchcommunities.InCavtat, ECLM and PKDD were co-located for the third time in a row, following the successful co-location of the two European conferences in Freiburg (2001) and Helsinki (2002). The co-location of ECLM2003 and PKDD2003 resulted in a joint program for the two conferences, including paper presentations, invited talks, tutorials, and workshops. Out of 332 submitted papers, 40 were accepted for publication in the ECLM2003proceedingsand40wereacceptedforpublicationinthePKDD2003 proceedings. All the submitted papers were reviewed by three referees. In ad- tion to submitted papers, the conference program consisted of four invited talks, four tutorials, seven workshops, two tutorials combined with a workshop, and a discovery challenge.

This book constitutes the proceedings of the First International Conference on Principles and Practice of Constraint Programming, CP '95, held in Cassis near Marseille, France in September 1995. The 33 refereed full papers included were selected out of 108 submissions and constitute the main part of the book; in addition there is a 60-p-ge documentation of the four invited papers and a section presenting industrial reports. Thus besides having a very strong research component, the volume will be attractive for practitioners. The papers are organized in sections on efficient constraint handling, constraint logic programming, concurrent constraint programming, computational logic, applications, and operations research.

Research Paper (undergraduate) from the year 2015 in the subject Engineering - Industrial Engineering and Management, grade: 1.0. Technical University of Berlin (Wirtschaftsinformatik - Information Systems Engineering (ISE)), course: Seminar: Hot Topics in Information Systems Engineering, language: English, abstract: During the last years NoSQL databases have been developed to ad-dress the needs of tremendous performance, reliability and horizontal scalability. NoSQL time series databases (TSDBs) have risen to combine valuable NoSQL properties with characteristics of time series data encountering many use-cases. Solutions offer the efficient handling of data volume and frequency related to time series. Developers and decision makers struggle with the choice of a TSDB among a large variety of solutions. Up to now no comparison exists focusing on the specific features and qualities of those heterogeneous applications. This paper aims to deliver two frameworks for the comparison of TSDBs, firstly with a focus on features and secondly on quality. Furthermore, we apply and evaluate the frameworks on up to seven open-source TSDBs such as InfluxDB and OpenTSDB. We come to the result that the investigated TSDBs differ mainly in support- and extension related points. They share performance-enhancing techniques, time-related query capabilities and data schemas optimized for the handling of time-series data.

Real-Time & Stream Data Management

Designing Resilient Systems with Spring Boot, Spring Cloud, and Cloud Foundry

Time and Relational Theory

Knowledge Discovery in Databases: PKDD 2003

Algorithms and Architectures for Parallel Processing

Encyclopedia of Database Technologies and Applications

Time Series Analysis Methods and Applications for Flight Data

Introductory, theory-practice balanced text teaching the fundamentals of databases to advanced undergraduates or graduate students in information systems or computer science.

This book constitutes the proceedings of the 4th International Workshop on Traffic Monitoring and Analysis, TMA 2012, held in Vienna, Austria, in March 2012. The thoroughly refereed 10 full papers and 8 short papers presented in this volume were carefully reviewed and selected from 31 submissions. The contributions are organized in topical sections on traffic analysis and characterization; new results and improved measurement techniques; measurement for QoS, security and service level agreements; and tools for network measurement and experimentation.

Firebase, Meteor, RethinkDB, Parse, Baqend, and others. Second, the high-level classification scheme illustrated above provides a gentle introduction into the system space of data management: Abstracting from the extreme system diversity in this field, it helps readers build a mental model of the available options. The two-volume set LNCS 11944-11945 constitutes the proceedings of the 19th International Conference on Algorithms and Architectures for Parallel Processing, ICA3PP 2019, held in Melbourne, Australia, in December 2019. The 73 full and 29 short papers presented were carefully reviewed and selected from 251 submissions. The papers are organized in topical sections on: Parallel and Distributed Architectures, Software Systems and Programming Models, Distributed and Parallel and Network-based Computing, Big Data and its Applications, Distributed and Parallel Algorithms, Applications of Distributed and Parallel Computing, Service Dependability and Security, IoT and CPS Computing, Performance Modelling and Evaluation. Adding the time dimension to real-world databases produces Time Series Databases (TSDB) and introduces new aspects and difficulties to data mining and knowledge discovery. This book covers the state-of-the-art methodology for mining time series databases. The novel data mining methods presented in the book include techniques for efficient segmentation, indexing, and classification of noisy and dynamic timeseries. A graph-based method for anomaly detection in time series is described and the book also studies the implications of a novel and potentially useful representation of time series as strings. The problem of detecting changes in data mining models that are induced from temporal databases is additionally discussed. Combines language tutorials with application design advice to cover the PHP server-side scripting language and the MySQL database engine.