

Joe Celkos Trees And Hierarchies In Sql For Smarties The Morgan Kaufmann Series In Data Management Systems

The authors explore and explain current techniques for handling the specialised data that describes geographical phenomena in a study that will be of great value to computer scientists and geographers working with spatial databases.

While the REST design philosophy has captured the imagination of web and enterprise developers alike, using this approach to develop real web services is no picnic. This cookbook includes more than 100 recipes to help you take advantage of REST, HTTP, and the infrastructure of the Web. You'll learn ways to design RESTful web services for client and server applications that meet performance, scalability, reliability, and security goals, no matter what programming language and development framework you use. Each recipe includes one or two problem statements, with easy-to-follow, step-by-step instructions for solving them, as well as examples using HTTP requests and responses, and XML, JSON, and Atom snippets. You'll also get implementation guidelines, and a discussion of the pros, cons, and trade-offs that come with each solution. Learn how to design resources to meet various application scenarios Successfully design representations and URIs Implement the hypertext constraint using links and link headers Understand when and how to use Atom and AtomPub Know what and what not to do to support caching Learn how to implement concurrency control Deal with advanced use cases involving copying, merging, transactions, batch processing, and partial updates Secure web services and support OAuth

An industry consultant shares his most useful tips and tricks for advanced SQL programming to help the working programmer gain performance and work around system deficiencies.

Are you an SQL programmer that, like many, came to SQL after learning and writing procedural or object-oriented code? Or have switched jobs to where a different brand of SQL is being used, or maybe even been told to learn SQL yourself? If even one answer is yes, then you need this book. A "Manual of Style" for the SQL programmer, this book is a collection of heuristics and rules, tips, and tricks that will help you improve SQL programming style and proficiency, and for formatting and writing portable, readable, maintainable SQL code. Based on many years of experience consulting in SQL shops, and gathering questions and resolving his students' SQL style issues, Joe Celko can help you become an even better SQL programmer. Help you write Standard SQL without an accent or a dialect that is used in another programming language or a specific flavor of SQL, code that can be maintained and used by other people. Enable you to give your group a coding standard for internal use, to enable programmers to use a consistent style. Give you the mental tools to approach a new problem with SQL as your tool, rather than another programming language — one that someone else might not know!

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

Read PDF Joe Celkos Trees And Hierarchies In Sql For Smarties The Morgan Kaufmann Series In Data Management Systems

from both academia and industry

This guide documents SQL: 1999's advanced features in the same practical, "programmercentric" way that the first volume documented the language's basic features. This is no mere representation of the standard, but rather authoritative guidance on making an application conform to it, both formally and effectively.

If you are a typical Oracle professional, you don't have the luxury of time to keep up with new technology and read all the new manuals to understand each new feature of the latest release from Oracle. You need a comprehensive source of information and in-depth tips and techniques for using the new technology. You need Oracle Internals: Tips, Trick

Joe Celko's SQL Puzzles and Answers, Second Edition, challenges you with his trickiest puzzles and then helps solve them with a variety of solutions and explanations. Author Joe Celko demonstrates the thought processes that are involved in attacking a problem from an SQL perspective to help advanced database programmers solve the puzzles you frequently face. These techniques not only help with the puzzle at hand, but also help develop the mindset needed to solve the many difficult SQL puzzles you face every day. This updated edition features many new puzzles; dozens of new solutions to puzzles; and new chapters on temporal query puzzles and common misconceptions about SQL and RDBMS that leads to problems. This book is recommended for database programmers with a good knowledge of SQL. A great collection of tricky SQL puzzles with a variety of solutions and explanations Uses the proven format of puzzles and solutions to provide a user-friendly, practical look into SQL programming problems - many of which will help users solve their own problems New edition features: Many new puzzles added!, Dozens of new solutions to puzzles, and using features in SQL-99, Code is edited to conform to SQL STYLE rules, New chapter on temporal query puzzles, New chapter on common misconceptions about SQL and RDBMS that leads to problems

Demonstrates important concepts and offers working Transact-SQL code, covering data filtering, DDL, DML, statistical functions, runs and sequences, transactions, stored procedures and triggers, and performance tuning.

Tackle the toughest set-based querying and query tuning problems—guided by an author team with in-depth, inside knowledge of T-SQL. Deepen your understanding of architecture and internals—and gain practical approaches and advanced techniques to optimize your code's performance. Discover how to: Move from procedural programming to the language of sets and logic Optimize query tuning with a top-down methodology Assess algorithmic complexity to predict performance Compare data-aggregation techniques, including new grouping sets Manage data modification—insert, delete, update, merge—for performance Write more efficient queries against partitioned tables Work with graphs, trees, hierarchies, and recursive queries Plus—Use pure-logic puzzles to sharpen your problem-solving skills

This indispensable SQL reference book is the first of its kind to leverage the benefits of design patterns to relational database SQL queries; all common SQL structures and design patterns are clearly categorized and described. Emphasizing the theoretical foundation for almost every type of SQL query problem, accompanying figures are included to help visualize the problem. Because SQL is a declarative language there are many ways to write any SQL query and professional database programmers must understand the correct way to write SQL for complicated database queries, and managers must institute formal SQL coding standards to improve productivity and maintainability. The SQL design patterns in this resource greatly improve the quality and productivity of systems development projects by forming a "best practices" foundation for all relational database queries.

Provides information on developing database applications in SQL, covering such topics as adjacency list model, nested sets, binary trees, data modeling, graphs, and hierarchical database systems.

Read PDF Joe Celkos Trees And Hierarchies In Sql For Smarties The Morgan Kaufmann Series In Data Management Systems

The only book you'll ever need on SQL. The authors detail the changes in the new standard and provide a thorough guide to programming with SQL 2 for both newcomers and experienced programmers. The book is one that novice programmers should read cover to cover and experienced DBMS professionals should have as a definitive reference book for the new SQL 2 standard.

Joe Celko's Trees and Hierarchies in SQL for SmartiesElsevier

Django is the leading Python web application development framework. Learn how to leverage the Django web framework to its full potential in this advanced tutorial and reference. Endorsed by Django, Pro Django more or less picks up where The Definitive Guide to Django left off and examines in greater detail the unusual and complex problems that Python web application developers can face and how to solve them. Provides in-depth information about advanced tools and techniques available in every Django installation Runs the gamut from the theory of Django's internal operations to actual code that solves real-world problems for high-volume environments Goes above and beyond other books, leaving the basics behind Shows how Django can do things even its core developers never dreamed possible

A deceptively simple process called Scrum, a new approach to systems development projects, is described here. Application development managers will learn how to simplify the implementation of Agile processes with Scrum, how to simplify XP implementation through a Scrum wrapper, how to understand the theoretical underpinnings of Agile processes, and why Agile processes work and how to manage them.

Material is of interest to the software development community at large. Schwaber is president of a software development consultancy.

Annotation copyrighted by Book News Inc., Portland, OR.

Perfectly intelligent programmers often struggle when forced to work with SQL. Why? Joe Celko believes the problem lies with their procedural programming mindset, which keeps them from taking full advantage of the power of declarative languages. The result is overly complex and inefficient code, not to mention lost productivity. This book will change the way you think about the problems you solve with SQL programs.. Focusing on three key table-based techniques, Celko reveals their power through detailed examples and clear explanations. As you master these techniques, you'll find you are able to conceptualize problems as rooted in sets and solvable through declarative programming. Before long, you'll be coding more quickly, writing more efficient code, and applying the full power of SQL • Filled with the insights of one of the world's leading SQL authorities - noted for his knowledge and his ability to teach what he knows. • Focuses on auxiliary tables (for computing functions and other values by joins), temporal tables (for temporal queries, historical data, and audit information), and virtual tables (for improved performance). • Presents clear guidance for selecting and correctly applying the right table technique.

SQL for Smarties was hailed as the first book devoted explicitly to the advanced techniques needed to transform an experienced SQL programmer into an expert. Now, 10 years later and in the third edition, this classic still reigns supreme as the book written by an SQL master that teaches future SQL masters. These are not just tips and techniques; Joe also offers the best solutions to old and new challenges and conveys the way you need to think in order to get the most out of SQL programming efforts for both correctness and performance. In the third edition, Joe features new examples and

updates to SQL-99, expanded sections of Query techniques, and a new section on schema design, with the same war-story teaching style that made the first and second editions of this book classics. Expert advice from a noted SQL authority and award-winning columnist, who has given ten years of service to the ANSI SQL standards committee and many more years of dependable help to readers of online forums. Teaches scores of advanced techniques that can be used with any product, in any SQL environment, whether it is an SQL-92 or SQL-99 environment. Offers tips for working around system deficiencies. Continues to use war stories--updated!--that give insights into real-world SQL programming challenges.

The implementation of stored procedures in MySQL 5.0 a huge milestone -- one that is expected to lead to widespread enterprise adoption of the already extremely popular MySQL database. If you are serious about building the web-based database applications of the future, you need to get up to speed quickly on how stored procedures work -- and how to build them the right way. This book, destined to be the bible of stored procedure development, is a resource that no real MySQL programmer can afford to do without. In the decade since MySQL burst on the scene, it has become the dominant open source database, with capabilities and performance rivaling those of commercial RDBMS offerings like Oracle and SQL Server. Along with Linux and PHP, MySQL is at the heart of millions of applications. And now, with support for stored procedures, functions, and triggers in MySQL 5.0, MySQL offers the programming power needed for true enterprise use. MySQL's new procedural language has a straightforward syntax, making it easy to write simple programs. But it's not so easy to write secure, easily maintained, high-performance, and bug-free programs. Few in the MySQL world have substantial experience yet with stored procedures, but Guy Harrison and Steven Feuerstein have decades of combined expertise. In MySQL Stored Procedure Programming, they put that hard-won experience to good use. Packed with code examples and covering everything from language basics to application building to advanced tuning and best practices, this highly readable book is the one-stop guide to MySQL development. It consists of four major sections: MySQL stored programming fundamentals -- tutorial, basic statements, SQL in stored programs, and error handling Building MySQL stored programs -- transaction handling, built-in functions, stored functions, and triggers MySQL stored programs in applications -- using stored programs with PHP, Java, Perl, Python, and .NET (C# and VB.NET) Optimizing MySQL stored programs -- security, basic and advanced SQL tuning, optimizing stored program code, and programming best practices A companion web site contains many thousands of lines of code, that you can put to use immediately. Guy Harrison is Chief Architect of Database Solutions at Quest Software and a frequent speaker and writer on MySQL topics. Steven Feuerstein is the author of Oracle PL/SQL Programming, the classic reference for Oracle stored programming for more than ten years. Both have decades of experience as database developers, and between them they have authored a

dozen books.

T-SQL insiders help you tackle your toughest queries and query-tuning problems Squeeze maximum performance and efficiency from every T-SQL query you write or tune. Four leading experts take an in-depth look at T-SQL's internal architecture and offer advanced practical techniques for optimizing response time and resource usage. Emphasizing a correct understanding of the language and its foundations, the authors present unique solutions they have spent years developing and refining. All code and techniques are fully updated to reflect new T-SQL enhancements in Microsoft SQL Server 2014 and SQL Server 2012. Write faster, more efficient T-SQL code: Move from procedural programming to the language of sets and logic Master an efficient top-down tuning methodology Assess algorithmic complexity to predict performance Compare data aggregation techniques, including new grouping sets Efficiently perform data-analysis calculations Make the most of T-SQL's optimized bulk import tools Avoid date/time pitfalls that lead to buggy, poorly performing code Create optimized BI statistical queries without additional software Use programmable objects to accelerate queries Unlock major performance improvements with In-Memory OLTP Master useful and elegant approaches to manipulating graphs About This Book For experienced T-SQL practitioners Includes coverage updated from Inside Microsoft SQL Server 2008 T-SQL Querying and Inside Microsoft SQL Server 2008 T-SQL Programming Valuable to developers, DBAs, BI professionals, and data scientists Covers many MCSE 70-464 and MCSA/MCSE 70-461 exam topics

Joe Celko's Analytics and OLAP in SQL is the first book that teaches what SQL programmers need in order to successfully make the transition from On-Line Transaction Processing (OLTP) systems into the world of On-Line Analytical Processing (OLAP). This book is not an in-depth look at particular subjects, but an overview of many subjects that will give the working RDBMS programmers a map of the terra incognita they will face — if they want to grow. It contains expert advice from a noted SQL authority and award-winning columnist, who has given ten years of service to the ANSI SQL standards committee and many more years of dependable help to readers of online forums. It offers real-world insights and lots of practical examples. It covers the OLAP extensions in SQL-99; ETL tools, OLAP features supported in DBMSs, other query tools, simple reports, and statistical software. This book is ideal for experienced SQL programmers who have worked with OLTP systems who need to learn techniques—and even some tricks—that they can use in an OLAP situation. Expert advice from a noted SQL authority and award-winning columnist, who has given ten years of service to the ANSI SQL standards committee and many more years of dependable help to readers of online forums First book that teaches what SQL programmers need in order to successfully make the transition from transactional systems (OLTP) into the world of data warehouse data and OLAP Offers real-world insights and lots of

practical examples Covers the OLAP extensions in SQL-99; ETL tools, OLAP features supported in DBMSs, other query tools, simple reports, and statistical software

Data Mining: Concepts and Techniques provides the concepts and techniques in processing gathered data or information, which will be used in various applications. Specifically, it explains data mining and the tools used in discovering knowledge from the collected data. This book is referred as the knowledge discovery from data (KDD). It focuses on the feasibility, usefulness, effectiveness, and scalability of techniques of large data sets. After describing data mining, this edition explains the methods of knowing, preprocessing, processing, and warehousing data. It then presents information about data warehouses, online analytical processing (OLAP), and data cube technology. Then, the methods involved in mining frequent patterns, associations, and correlations for large data sets are described. The book details the methods for data classification and introduces the concepts and methods for data clustering. The remaining chapters discuss the outlier detection and the trends, applications, and research frontiers in data mining. This book is intended for Computer Science students, application developers, business professionals, and researchers who seek information on data mining. Presents dozens of algorithms and implementation examples, all in pseudo-code and suitable for use in real-world, large-scale data mining projects Addresses advanced topics such as mining object-relational databases, spatial databases, multimedia databases, time-series databases, text databases, the World Wide Web, and applications in several fields Provides a comprehensive, practical look at the concepts and techniques you need to get the most out of your data

Data Mining: Practical Machine Learning Tools and Techniques, Third Edition, offers a thorough grounding in machine learning concepts as well as practical advice on applying machine learning tools and techniques in real-world data mining situations. This highly anticipated third edition of the most acclaimed work on data mining and machine learning will teach you everything you need to know about preparing inputs, interpreting outputs, evaluating results, and the algorithmic methods at the heart of successful data mining. Thorough updates reflect the technical changes and modernizations that have taken place in the field since the last edition, including new material on Data Transformations, Ensemble Learning, Massive Data Sets, Multi-instance Learning, plus a new version of the popular Weka machine learning software developed by the authors. Witten, Frank, and Hall include both tried-and-true techniques of today as well as methods at the leading edge of contemporary research. The book is targeted at information systems practitioners, programmers, consultants, developers, information technology managers, specification writers, data analysts, data modelers, database R&D professionals, data warehouse engineers, data mining professionals. The book will also be useful for professors and students of upper-level undergraduate and graduate-level data mining and machine learning courses who want to

incorporate data mining as part of their data management knowledge base and expertise. Provides a thorough grounding in machine learning concepts as well as practical advice on applying the tools and techniques to your data mining projects Offers concrete tips and techniques for performance improvement that work by transforming the input or output in machine learning methods Includes downloadable Weka software toolkit, a collection of machine learning algorithms for data mining tasks—in an updated, interactive interface. Algorithms in toolkit cover: data pre-processing, classification, regression, clustering, association rules, visualization

"XQuery Kick Start" delivers a concise introduction to the XQuery standard, and useful implementation advice for developers needing to put it into practice. The book starts by explaining the role of XQuery in the XML family of specifications, and its relationship with XPath. The authors then explain the specification in detail, describing the semantics and data model, before moving to examples using XQuery to manipulate XML databases and document storage systems. Later chapters discuss Java implementations of XQuery and development tools that facilitate the development of Web sites with XQuery. This book is up to date with the latest XQuery specifications, and includes coverage of new features for extending the XQuery language.

This book brings all of the elements of database design together in a single volume, saving the reader the time and expense of making multiple purchases. It consolidates both introductory and advanced topics, thereby covering the gamut of database design methodology ? from ER and UML techniques, to conceptual data modeling and table transformation, to storing XML and querying moving objects databases. The proposed book expertly combines the finest database design material from the Morgan Kaufmann portfolio. Individual chapters are derived from a select group of MK books authored by the best and brightest in the field. These chapters are combined into one comprehensive volume in a way that allows it to be used as a reference work for those interested in new and developing aspects of database design. This book represents a quick and efficient way to unite valuable content from leading database design experts, thereby creating a definitive, one-stop-shopping opportunity for customers to receive the information they would otherwise need to round up from separate sources. Chapters contributed by various recognized experts in the field let the reader remain up to date and fully informed from multiple viewpoints. Details multiple relational models and modeling languages, enhancing the reader's technical expertise and familiarity with design-related requirements specification. Coverage of both theory and practice brings all of the elements of database design together in a single volume, saving the reader the time and expense of making multiple purchases.

Learn Business Intelligence Markup Language (Biml) for automating much of the repetitive, manual labor involved in data integration. We teach you how to build frameworks and use advanced Biml features to get more out of SQL Server

Integration Services (SSIS), Transact-SQL (T-SQL), and SQL Server Analysis Services (SSAS) than you ever thought possible. The first part of the book starts with the basics—getting your development environment configured, Biml syntax, and scripting essentials. Whether a beginner or a seasoned Biml expert, the next part of the book guides you through the process of using Biml to build a framework that captures both your design patterns and execution management. Design patterns are reusable code blocks that standardize the approach you use to perform certain types of data integration, logging, and other key data functions. Design patterns solve common problems encountered when developing data integration solutions. Because you do not have to build the code from scratch each time, design patterns improve your efficiency as a Biml developer. In addition to leveraging design patterns in your framework, you will learn how to build a robust metadata store and how to package your framework into Biml bundles for deployment within your enterprise. In the last part of the book, we teach you more advanced Biml features and capabilities, such as SSAS development, T-SQL recipes, documentation autogeneration, and Biml troubleshooting. The Biml Book: Provides practical and applicable examples Teaches you how to use Biml to reduce development time while improving quality Takes you through solutions to common data integration and BI challenges What You'll Learn Master the basics of Business Intelligence Markup Language (Biml) Study patterns for automating SSIS package generation Build a Biml Framework Import and transform database schemas Automate generation of scripts and projects Who This Book Is For BI developers wishing to quickly locate previously tested solutions, Microsoft BI specialists, those seeking more information about solution automation and code generation, and practitioners of Data Integration Lifecycle Management (DILM) in the DevOps enterprise

The demand for SQL information and training continues to grow with the need for a database behind every website capable of offering web-based information queries. SQL is the de facto standard for database retrieval, and if you need to access, update, or utilize data in a modern database management system, you will need SQL to do it. The Second Edition of Joe Celko's Trees and Hierarchies in SQL for Smarties covers two new sets of extensions over three entirely new chapters and expounds upon the changes that have occurred in SQL standards since the previous edition's publication. Benefit from mastering the challenging aspects of these database applications in SQL as taught by Joe Celko, one of the most-read SQL authors in the world. Expert advice from a noted SQL authority and award-winning columnist who has given 10 years of service to the ANSI SQL standards committee Teaches scores of advanced techniques that can be used with any product, in any SQL environment Offers graph theory and programming techniques for working around deficiencies and gives insight into real-world challenges

This text covers basic database concepts to provide a conceptual understanding of data and databases necessary for database design and development.

200 practice questions to prepare & pass exam Exam 70-461: Querying Microsoft SQL Server 2012! Study this book to perform to excellence as a database developer! The exam text book follows the official Microsoft exam guidelines. The book features over 200 practice exam questions & answers. Once you learn the book, very likely you pass the exam effortlessly. SQL (Structured Query Language) programming & relational database design teach-by-practical-diagrams-&-examples book for developers, programmers, systems analysts and project managers who are new to relational database and client/server technologies. Also for database developers, database designers and database administrators (DBA), who know some SQL programming and database design, and who wish to refresh & expand their RDBMS design & development technology horizons. The book has special orientation for passing Exam 70-461. Familiarity with at least one computer programming language, Windows file system & Excel is assumed. Since the book is career advancement oriented, it has a great number of 3NF database design examples with metadata explanations along with practical SQL queries (over 1,400 SELECT queries) and T-SQL scripts, plenty to learn indeed. Great emphasis is placed on explaining the FOREIGN KEY - PRIMARY KEY constraints among tables, the connections which make the collection of individual tables a database. The database diagrams and queries are based on historic and current SQL Server sample databases: pubs (PRIMARY KEYs 9, FOREIGN KEYs 10) , Northwind (PRIMARY KEYs 13, FOREIGN KEYs 13) and the latest AdventureWorks series. Among them: AdventureWorks, AdventureWorks2008, AdventureWorks2012 (PRIMARY KEYs 71, FOREIGN KEYs 90), & AdventureWorksDW2012 (PRIMARY KEYs 27, FOREIGN KEYs 44). The last one is a data warehouse database which is the basis for multi-dimensional OLAP cubes. Sample databases installation instructions are included. The book teaches through vivid database diagrams and T-SQL queries & stored procedures; how to think in terms of sets at a very high level, focusing on set-based operations instead of loops like in procedural programming languages. There is a chapter dedicated to the new programming features of SQL Server 2012 and XML. The best way to master T-SQL programming is to type the query in your own SQL Server Management Studio Query Editor, test it, examine it, change it and study it. Wouldn't it be easier just to copy & paste it? It would, but the learning value would diminish rapidly. You need to feel relational database design and the SQL language in your DNA. SQL queries must "pour" out from your fingers into the keyboard. Why is knowing SQL queries by heart so important? After all everything can be found on the web so why not just copy & paste? Well not exactly. If you want to be an database designer & development expert, it has to be in your head not on the web. Second, when your supervisor is looking over your shoulder, "Charlie, can you tell me what is the total revenue for March?", you have to be able to type the query without documentation or SQL forum search and provide the results to your superior promptly. The book was designed to be readable in any environment, even on the beach laptop around or no laptop in sight at all. All

queries are followed by results row count and /or full/partial results listing in tabular (grid) format. Screenshots are used when dealing with GUI tools such as SQL Server Management Studio. SQL Server 2012 installation instructions with screenshots are included. Mastery of the database design & SQL programming book likely to be sufficient for career advancement as a database designer and database developer.

The key to client/server computing. Transaction processing techniques are deeply ingrained in the fields of databases and operating systems and are used to monitor, control and update information in modern computer systems. This book will show you how large, distributed, heterogeneous computer systems can be made to work reliably. Using transactions as a unifying conceptual framework, the authors show how to build high-performance distributed systems and high-availability applications with finite budgets and risk. The authors provide detailed explanations of why various problems occur as well as practical, usable techniques for their solution. Throughout the book, examples and techniques are drawn from the most successful commercial and research systems. Extensive use of compilable C code fragments demonstrates the many transaction processing algorithms presented in the book. The book will be valuable to anyone interested in implementing distributed systems or client/server architectures.

This guide contains a wealth of solutions to problems that SQL Server programmers face. The recipes in the book range from those that show how to perform simple tasks to ones that are more complicated.

The central challenge to imperial powers entering the modern era was the schooling of their peoples. How could they insure the literacy that modernity required without providing a foundation for nationalism among the colonised? In Russia's eastern empire in the late nineteenth century, Orthodox Christianity vied with Islam for people's souls; Russian language competed with Tatar and local vernaculars in market squares, peasant cottages, and schoolrooms; Arabic and Cyrillic alphabets clashed in school textbooks; and western secularism undermined traditional religious authority among both Muslim and Orthodox faithful. Russian nationalism peaked in the early twentieth century and public support for policies of the russification of non-Russian minorities increased. The inevitable clash with local languages shook the stability of the empire. Classroom and Empire tells the story of the politics of alphabets, languages, and schooling in the eastern empire of Russia from 1860 to 1917. Wayne Dowler presents an intriguing cast of characters, including Nikolai Il'minskii, whose method of schooling non-Russian children lay at the heart of nationalist controversy; Ismail Bey Gaspirali, whose new method schools attempted to reconcile Islam with modern secular philosophy and science; Konstantin Pobedonostsev, procurator of the Holy Synod and minence grise of the reigns of Alexander III and his son Nicholas II; and Sophia Chicherina, feisty defender of the Il'minskii school. Dowler shows us that the problem of schooling non-Russians was unresolved by the fall of the Romanovs in 1917, smouldered through much of the Soviet

period, and has re-emerged today as a major source of divisiveness in the Russian Federation. Wayne Dowler is professor of history at University of Toronto at Scarborough.

A poorly performing database application not only costs users time, but also has an impact on other applications running on the same computer or the same network. SQL Tuning provides an essential next step for SQL developers and database administrators who want to extend their SQL tuning expertise and get the most from their database applications. There are two basic issues to focus on when tuning SQL: how to find and interpret the execution plan of an SQL statement and how to change SQL to get a specific alternate execution plan. SQL Tuning provides answers to these questions and addresses a third issue that's even more important: how to find the optimal execution plan for the query to use. Author Dan Tow outlines a timesaving method he's developed for finding the optimum execution plan--rapidly and systematically--regardless of the complexity of the SQL or the database platform being used. You'll learn how to understand and control SQL execution plans and how to diagram SQL queries to deduce the best execution plan for a query. Key chapters in the book include exercises to reinforce the concepts you've learned. SQL Tuning concludes by addressing special concerns and unique solutions to "unsolvable problems." Whether you are a programmer who develops SQL-based applications or a database administrator or other who troubleshoots poorly tuned applications, SQL Tuning will arm you with a reliable and deterministic method for tuning your SQL queries to gain optimal performance.

Three of CouchDB's creators show you how to use this document-oriented database as a standalone application framework or with high-volume, distributed applications. With its simple model for storing, processing, and accessing data, CouchDB is ideal for web applications that handle huge amounts of loosely structured data. That alone would stretch the limits of a relational database, yet CouchDB offers an open source solution that's reliable, scales easily, and responds quickly. CouchDB works with self-contained data that has loose or ad-hoc connections. It's a model that fits many real-world items, such as contacts, invoices, and receipts, but you'll discover that this database can easily handle data of any kind. With this book, you'll learn how to work with CouchDB through its RESTful web interface, and become familiar with key features such as simple document CRUD (create, read, update, delete), advanced MapReduce, deployment tuning, and more. Understand the basics of document-oriented storage and manipulation Interact with CouchDB entirely through HTTP using its RESTful interface Model data as self-contained JSON documents Handle evolving data schemas naturally Query and aggregate data in CouchDB using MapReduce views Replicate data between nodes Tune CouchDB for increased performance and reliability

Offers tips for improving the performance of any SQL database, no matter what the platform. Written for experienced database administrators familiar with SQL, the book identifies the similarities and differences of eight DBMSs, including Oracle 9i, IBM DB2 7.2, and Microsoft SQL server 2000. It provides strategies for refining sorts, subqueries, columns, tables, indexes, constraints, and locks. Annotation copyrighted by Book News, Inc., Portland, OR

Joe Celko's Trees and Hierarchies in SQL is an intermediate to advanced-level practitioner's guide to mastering the two most

Read PDF Joe Celkos Trees And Hierarchies In Sql For Smarties The Morgan Kaufmann Series In Data Management Systems

challenging aspects of developing database applications in SQL. In this book, Celko illustrates several major approaches to representing trees and hierarchies and related topics that should be of interest to the working database programmer. These topics include hierarchical encoding schemes, graphs, IMS, binary trees, and more. This book covers SQL-92 and SQL:1999. · Includes graph theory and programming techniques. · Running examples throughout the book help illustrate and tie concepts together. · Loads of code, available for download from www.mkp.com.

This is the eBook of the printed book and may not include any media, website access codes, or print supplements that may come packaged with the bound book. Database Systems: The Complete Book is ideal for Database Systems and Database Design and Application courses offered at the junior, senior and graduate levels in Computer Science departments. A basic understanding of algebraic expressions and laws, logic, basic data structure, OOP concepts, and programming environments is implied. Written by well-known computer scientists, this introduction to database systems offers a comprehensive approach, focusing on database design, database use, and implementation of database applications and database management systems. The first half of the book provides in-depth coverage of databases from the point of view of the database designer, user, and application programmer. It covers the latest database standards SQL:1999, SQL/PSM, SQL/CLI, JDBC, ODL, and XML, with broader coverage of SQL than most other texts. The second half of the book provides in-depth coverage of databases from the point of view of the DBMS implementor. It focuses on storage structures, query processing, and transaction management. The book covers the main techniques in these areas with broader coverage of query optimization than most other texts, along with advanced topics including multidimensional and bitmap indexes, distributed transactions, and information integration techniques.

In Advanced Transact-SQL for SQL Server 2000, authors Itzik Ben-Gan and Thomas Moreau explore the powerful capabilities of Transact-SQL (T-SQL). Ben-Gan and Moreau offer solutions to common problems encountered using all versions of SQL Server, with a focus on the latest version, SQL Server 2000. Expert tips and real code examples teach advanced database programmers to write more efficient and better-performing code that takes full advantage of T-SQL. The authors offer practical solutions to the everyday problems programmers face and include in-depth information on advanced T-SQL topics such as joins, subqueries, stored procedures, triggers, user-defined functions (UDFs), indexed views, cascading actions, federated views, hierarchical structures, cursors, and more.

* This is the "official" book on FirebirdSQL—it's being written with the support of the development and management team. * Includes an extensive set of working, real-world examples, a troubleshooting guide, and a guide to migrating existing databases to FirebirdSQL. * This is the only book on the topic; appeals to users worldwide, especially in Europe; Most people deploy Firebird on Windows.

Joe Celko's Complete Guide to NoSQL provides a complete overview of non-relational technologies so that you can become more nimble to meet the needs of your organization. As data continues to explode and grow more complex, SQL is becoming less useful for querying data and extracting meaning. In this new world of bigger and faster data, you will need to leverage non-

Read PDF Joe Celkos Trees And Hierarchies In Sql For Smarties The Morgan Kaufmann Series In Data Management Systems

relational technologies to get the most out of the information you have. Learn where, when, and why the benefits of NoSQL outweigh those of SQL with Joe Celko's Complete Guide to NoSQL. This book covers three areas that make today's new data different from the data of the past: velocity, volume and variety. When information is changing faster than you can collect and query it, it simply cannot be treated the same as static data. Celko will help you understand velocity, to equip you with the tools to drink from a fire hose. Old storage and access models do not work for big data. Celko will help you understand volume, as well as different ways to store and access data such as petabytes and exabytes. Not all data can fit into a relational model, including genetic data, semantic data, and data generated by social networks. Celko will help you understand variety, as well as the alternative storage, query, and management frameworks needed by certain kinds of data. Gain a complete understanding of the situations in which SQL has more drawbacks than benefits so that you can better determine when to utilize NoSQL technologies for maximum benefit Recognize the pros and cons of columnar, streaming, and graph databases Make the transition to NoSQL with the expert guidance of best-selling SQL expert Joe Celko

[Copyright: 6507b3c8aab7431d7d3c87af67fbe49b](#)