

Database Principles Programming And Performance Second Edition The Morgan Kaufmann Series In Data Management Systems

The most prominent Web applications in use today are data-intensive. Scores of database management systems across the Internet access and maintain large amounts of structured data for e-commerce, on-line trading, banking, digital libraries, and other high-volume sites. Developing and maintaining these data-intensive applications is an especially complex, multi-disciplinary activity, requiring all the tools and techniques that software engineering can provide. This book represents a breakthrough for Web application developers. Using hundreds of illustrations and an elegant intuitive modeling language, the authors—all internationally-known database researchers—present a methodology that fully exploits the conceptual modeling approach of software engineering, from idea to application. Readers will learn not only how to harness the design technologies of relational databases for use on the Web, but also how to transform their conceptual designs of data-intensive Web applications into effective software components. * A fully self-contained introduction and practitioner's guide suitable for both technical and non-technical members of staff, as well as students. * A methodology, development process, and notation (WebML) based on common practice but optimized for the unique challenges of high-volume Web applications. * Completely platform- and product-independent; even the use of WebML is optional. * Based on well-known industry standards such as UML and the Entity Relationship Model. * Enhanced by its own Web site (<http://www.webml.org>), containing additional examples, papers, teaching materials, developers' resources, and exercises with solutions.

Whether you are a software developer, systems architect, data analyst, or business analyst, if you want to take advantage of data mining in the development of advanced analytic applications, Java Data Mining, JDM, the new standard now implemented in core DBMS and data mining/analysis software, is a key solution component. This book is the essential guide to the usage of the JDM standard interface, written by contributors to the JDM standard. Data mining introduction - an overview of data mining and the problems it can address across industries; JDM's place in strategic solutions to data mining-related problems JDM essentials - concepts, design approach and design issues, with detailed code examples in Java; a Web Services interface to enable JDM functionality in an SOA environment; and illustration of JDM XML Schema for JDM objects JDM in practice - the use of JDM from vendor implementations and approaches to customer applications, integration, and usage; impact of data mining on IT infrastructure; a how-to guide for building applications that use the JDM API Free, downloadable KJDM source code referenced in the book available here

Dealing with the volume, complexity, and diversity of data currently being generated by scientific experiments and simulations often causes scientists to waste productive time. Scientific Data Management: Challenges, Technology, and Deployment describes cutting-edge technologies and solutions for managing and analyzing vast amounts of data, helping scientists focus on their scientific goals. The book begins with coverage of efficient storage systems, discussing how to write and read large volumes of data without slowing the simulation, analysis, or visualization processes. It then focuses on the efficient data movement and management of storage spaces and explores emerging database systems for scientific data. The book also addresses how to best organize data for analysis purposes, how to effectively conduct searches over large datasets, how to successfully automate multistep scientific process workflows, and how to automatically collect metadata and lineage information. This book provides a comprehensive understanding of the latest techniques for managing data during scientific exploration processes, from data generation to data analysis. Enhanced by numerous detailed color images, it includes real-world examples of applications drawn from biology, ecology, geology, climatology, and more. Check out Dr. Shoshani discuss the book during an interview with International Science Grid This Week (iSGTW): <http://www.isgtw.org/?pid=1002259>

Most of modern enterprises, institutions, and organizations rely on knowledge-based management systems. In these systems, knowledge is gained from data analysis. Today, knowledge-based management systems include data warehouses as their core components. Data integrated in a data warehouse are analyzed by the so-called On-Line Analytical Processing (OLAP) applications designed to discover trends, patterns of behavior, and anomalies as well as finding dependencies between data. Massive amounts of integrated data and the complexity of integrated data coming from many different sources make data integration and processing challenging. New Trends in Data Warehousing and Data Analysis brings together the most recent research and practical achievements in the DW and OLAP technologies. It provides an up-to-date bibliography of published works and the resource of research achievements. Finally, the book assists in the dissemination of knowledge in the field of advanced DW and OLAP.

Data Mining, Second Edition, describes data mining techniques and shows how they work. The book is a major revision of the first edition that appeared in 1999. While the basic core remains the same, it has been updated to reflect the changes that have taken place over five years, and now has nearly double the references. The highlights of this new edition include thirty new technique sections; an enhanced Weka machine learning workbench, which now features an interactive interface; comprehensive information on neural networks; a new section on Bayesian networks; and much more. This text is designed for information systems practitioners, programmers, consultants, developers, information technology managers, specification writers as well as professors and students of graduate-level data mining and machine learning courses. Algorithmic methods at the heart of successful data mining—including tried and true techniques as well as leading edge methods Performance improvement techniques that work by transforming the input or output Tuning your database for optimal performance means more than following a few short steps in a vendor-specific guide. For maximum improvement, you need a broad and deep knowledge of basic tuning principles, the ability to gather data in a systematic way, and the skill to make your system run faster. This is an art as well as a science, and Database

Tuning: Principles, Experiments, and Troubleshooting Techniques will help you develop portable skills that will allow you to tune a wide variety of database systems on a multitude of hardware and operating systems. Further, these skills, combined with the scripts provided for validating results, are exactly what you need to evaluate competing database products and to choose the right one. Forward by Jim Gray, with invited chapters by Joe Celko and Alberto Lerner Includes industrial contributions by Bill McKenna (RedBrick/Informix), Hany Saleeb (Oracle), Tim Shetler (TimesTen), Judy Smith (Deutsche Bank), and Ron Yorita (IBM) Covers the entire system environment: hardware, operating system, transactions, indexes, queries, table design, and application analysis Contains experiments (scripts available on the author's site) to help you verify a system's effectiveness in your own environment Presents special topics, including data warehousing, Web support, main memory databases, specialized databases, and financial time series Describes performance-monitoring techniques that will help you recognize and troubleshoot problems

Differing from other books on the subject, this one uses the framework of constraint databases to provide a natural and powerful generalization of relational databases. An important theme running through the text is showing how relational databases can smoothly develop into constraint databases, without sacrificing any of the benefits of relational databases whilst gaining new advantages. Peter Revesz begins by discussing data models and how queries may be addressed to them. From here, he develops the theory of relational and constraint databases, including Datalog and the relational calculus, concluding with three sample constraint database systems -- DISCO, DINGO, and RATHER. Advanced undergraduates and graduates in computer science will find this a clear introduction to the subject, while professionals and researchers will appreciate this novel perspective on their subject.

With the growth of Java and the rise of database-powered Web applications, the need to use Java with SQL is clear. Until now, authoritative coverage of the techniques available to meet these challenges and reap their benefits-both programming and career benefits-didn't exist. Understanding SQL and Java Together examines all the standards for combining SQL and Java. It shows you exactly how to use their features to write efficient and effective code supporting Java access to SQL data in a variety of ways. You'll gain a thorough understanding of the relationship between SQL and Java, which will allow you to write static and dynamic SQL programs in Java, merge Java code with SQL databases and SQL code, and use other data management techniques wherever appropriate. * Covers all the technologies for using SQL and Java together, including JDBC, Java Blend, and SQLJ Parts 0, 1, and 2 * Explains how to embed SQL code in Java and take advantage of Java's ability to compile that code for a specific DBMS * Explains how to store and invoke Java routines in an SQL database-and how to store Java objects in an SQL database for seamless interchange among application layers * Covers dynamic SQL access techniques using JDBC and advantageous ways to combine static and dynamic SQL * Comes with a CD-ROM containing Oracle's JDeveloper , Sybase's Adaptive Server Anywhere, Informix's Cloudscape, the complete database schema, and the complete text of most of the examples

An introduction to database management systems for anyone who is, or wishes to become, a database administrator, application programmer, or sophisticated SQL user. Starting with fundamental relational concepts, the author presents a full range of information about database standards and products.

DW 2.0: The Architecture for the Next Generation of Data Warehousing is the first book on the new generation of data warehouse architecture, DW 2.0, by the father of the data warehouse. The book describes the future of data warehousing that is technologically possible today, at both an architectural level and technology level. The perspective of the book is from the top down: looking at the overall architecture and then delving into the issues underlying the components. This allows people who are building or using a data warehouse to see what lies ahead and determine what new technology to buy, how to plan extensions to the data warehouse, what can be salvaged from the current system, and how to justify the expense at the most practical level. This book gives experienced data warehouse professionals everything they need in order to implement the new generation DW 2.0. It is designed for professionals in the IT organization, including data architects, DBAs, systems design and development professionals, as well as data warehouse and knowledge management professionals. * First book on the new generation of data warehouse architecture, DW 2.0. * Written by the "father of the data warehouse", Bill Inmon, a columnist and newsletter editor of The Bill Inmon Channel on the Business Intelligence Network. * Long overdue comprehensive coverage of the implementation of technology and tools that enable the new generation of the DW: metadata, temporal data, ETL, unstructured data, and data quality control.

SQL: 1999 is the best way to make the leap from SQL-92 to SQL:1999, but it is much more than just a simple bridge between the two. The latest from celebrated SQL experts Jim Melton and Alan Simon, SQL:1999 is a comprehensive, eminently practical account of SQL's latest incarnation and a potent distillation of the details required to put it to work. Written to accommodate both novice and experienced SQL users, SQL:1999 focuses on the language's capabilities, from the basic to the advanced, and the ways that real applications take advantage of them. Throughout, the authors illustrate features and techniques with clear and often entertaining references to their own custom database. Gives authoritative coverage from an expert team that includes the editor of the SQL-92 and SQL:1999 standards. Provides a general introduction to SQL that helps you understand its constituent parts, history, and place in the realm of computer languages. Explains SQL:1999's more sophisticated features, including advanced value expressions, predicates, advanced SQL query expressions, and support for active databases. Explores key issues for programmers linking applications to SQL databases. Provides guidance on troubleshooting, internationalization, and changes anticipated in the next version of SQL. Contains appendices devoted to database design, a complete SQL:1999 example, the standardization process, and more.

C. J. Date is one of the founding fathers of the relational database field. Many of today's seasoned database professionals "grew up" on Date's writings. Those same professionals, along with other serious database students and practitioners, form the core audience for Date's ongoing writing efforts. Date on Database: Writings 2000-2006 is a compilation of Date's most significant articles and papers over the past seven years. It gives readers a one-stop place in which to find Date's latest thinking on relational technology. Many papers are not easily found outside this

book.

An Overview of Multidatabase Systems: Past and Present / Athman Bouguettaya, Boualem Benatallah, Ahmed Elmagarmid / - Local Autonomy and Its Effects on Multidatabase Systems / Ahmed Elmagarmid, Weimin Du, Rafi Ahmed / - Semantic Similarities Between Objects in Multiple Databases / Vipul Kashyap, Amit Sheth / - Resolution of Representational Diversity in Multidatabase Systems / Joachim Hammer, Dennis McLeod / - Schema Integration: Past, Present, and Future / Sudha Ram, V. Ramesh / - Schema and Language Translation / Bogdan Czejdo, Le Gruenwald / - Multidatabase Languages / Paolo Missier, Marek Rusinkiewicz, W. Jin / - Interdependent Database Systems / George Karabatis, Marek Rusinkiewicz, Amit Sheth / - Correctness Criteria and Concurrency Control / Panos K. Chrysanthis, Krithi Ramamritham / - Transaction Management in Multidatabase Systems: Current Technologies and Formalisms / Ken Barker, Ahmed Elmagarmid / - Transaction-Based Recovery / Jari Veijalainen. ...

Joe Celkos SQL for Smarties: Advanced SQL Programming offers tips and techniques in advanced programming. This book is the fourth edition and it consists of 39 chapters, starting with a comparison between databases and file systems. It covers transactions and currency control, schema level objects, locating data and schema numbers, base tables, and auxiliary tables. Furthermore, procedural, semi-procedural, and declarative programming are explored in this book. The book also presents the different normal forms in database normalization, including the first, second, third, fourth, fifth, elementary key, domain-key, and Boyce-Codd normal forms. It also offers practical hints for normalization and denormalization. The book discusses different data types, such as the numeric, temporal and character data types; the different predicates; and the simple and advanced SELECT statements. In addition, the book presents virtual tables, and it discusses data partitions in queries; grouping operations; simple aggregate functions; and descriptive statistics, matrices and graphs in SQL. The book concludes with a discussion about optimizing SQL. It will be of great value to SQL programmers. Expert advice from a noted SQL authority and award-winning columnist who has given ten years service to the ANSI SQL standards committee Teaches scores of advanced techniques that can be used with any product, in any SQL environment, whether it is an SQL 92 or SQL 2008 environment Offers tips for working around deficiencies and gives insight into real-world challenges

DatabasePrinciples, Programming, and Performance, Second EditionMorgan Kaufmann

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.

The database field has experienced a rapid and incessant growth since the development of relational databases. The progress in database systems and applications has produced a diverse landscape of specialized technology areas that have often become the exclusive domain of research specialists. Examples include active databases, temporal databases, object-oriented databases, deductive databases, imprecise reasoning and queries, and multimedia information systems. This book provides a systematic introduction to and an in-depth treatment of these advanced database areas. It supplies practitioners and researchers with authoritative coverage of recent technological advances that are shaping the future of commercial database systems and intelligent information systems. Advanced Database Systems was written by a team of six leading specialists who have made significant contributions to the development of the technology areas covered in the book. Benefiting from the authors' long experience teaching graduate and professional courses, this book is designed to provide a gradual introduction to advanced research topics and includes many examples and exercises to support its use for individual study, desk reference, and graduate classroom teaching.

This book is a comprehensive, practical, and student-friendly textbook addressing fundamental concepts in database design and applications.

Moving Objects Databases is the first uniform treatment of moving objects databases, the technology that supports GPS and RFID. It focuses on the modeling and design of data from moving objects — such as people, animals, vehicles, hurricanes, forest fires, oil spills, armies, or other objects — as well as the storage, retrieval, and querying of that very voluminous data. It includes homework assignments at the end of each chapter, exercises throughout the text that students can complete as they read, and a solutions manual in the back of the book. This book is intended for graduate or advanced undergraduate students. It is also recommended for computer scientists and database systems engineers and programmers in government, industry and academia; professionals from other disciplines, e.g., geography, geology, soil science, hydrology, urban and regional planning, mobile computing, bioterrorism and homeland security, etc. Focuses on the modeling and design of data from moving objects--such as people, animals, vehicles, hurricanes, forest fires, oil spills, armies, or other objects--as well as the storage, retrieval, and querying of that very voluminous data. Demonstrates through many practical examples and illustrations how new concepts and techniques are used to integrate time and space in database applications. Provides exercises and solutions in each chapter to enable the reader to explore recent research results in practice.

Database: Principles Programming Performance provides an introduction to the fundamental principles of database systems. This book focuses on database programming and the relationships between principles, programming, and performance. Organized into 10 chapters, this book begins with an overview of database design principles and presents a comprehensive introduction to the concepts used by a DBA. This text then provides grounding in many abstract concepts of the relational model. Other chapters introduce SQL, describing its capabilities and covering the statements and functions of the programming language. This book provides as well an introduction to Embedded SQL and Dynamic SQL that is sufficiently detailed to enable students to immediately start writing database programs. The final chapter deals with some of the motivations for database systems spanning multiple CPUs, including client-server and distributed transactions. This book is a valuable resource for database administrators, application programmers, specialist users, and end users.

Are you a data mining analyst, who spends up to 80% of your time assuring data quality, then preparing that data for developing and deploying predictive models? And do you find lots of literature on data mining theory and concepts, but when it comes to practical advice on developing good mining views find little “how to information? And are you, like most analysts, preparing the data in SAS? This book is intended to fill this gap as your source of practical recipes. It introduces a framework for the process of data preparation for data mining, and presents the detailed implementation of each step in SAS. In addition, business applications of data mining modeling require you to deal with a large number of variables, typically hundreds if not thousands. Therefore, the book devotes several chapters to the methods of data transformation and variable selection. A complete framework for the data preparation process, including implementation details for each step. The complete SAS implementation code, which is readily usable by professional analysts and data miners. A unique and comprehensive approach for the treatment of missing values, optimal binning, and cardinality reduction. Assumes minimal proficiency in SAS and includes a quick-start chapter on writing SAS macros.

Component Database Systems is a collection of invited chapters by the researchers making the most influential contributions in the database industry's trend toward componentization. This book represents the sometimes-divergent, sometimes-convergent approaches taken by leading database vendors as they seek to establish commercially viable componentization strategies. Together, these contributions form the first book devoted entirely to the technical and architectural design of component-based database systems. In addition to detailing the current state of their research, the authors also take up many of the issues affecting the likely future directions of component databases. If you have a stake in the evolution of any of today's leading database systems, this book will make fascinating reading. It will also help prepare you for the technology that is likely to become widely available over the next several years. * Is comprised of contributions from the field's most highly respected researchers, including key figures at IBM, Oracle, Informix, Microsoft, and POET. * Represents the entire spectrum of approaches taken by leading software companies working on DBMS componentization strategies. * Covers component-focused architectures, methods for hooking components into an overall system, and support for component development. * Examines the component technologies that are most valuable to Web-based and multimedia databases. * Presents a thorough classification and overview of component database systems.

For programmers who prefer content to frills, this guide has succinct and straightforward information for putting Access to its full, individually tailored use.

The rapidly increasing volume of information contained in relational databases places a strain on databases, performance, and maintainability: DBAs are under greater pressure than ever to optimize database structure for system performance and administration. Physical Database Design discusses the concept of how physical structures of databases affect performance, including specific examples, guidelines, and best and worst practices for a variety of DBMSs and configurations. Something as simple as improving the table index design has a profound impact on performance. Every form of relational database, such as Online Transaction Processing (OLTP), Enterprise Resource Management (ERP), Data Mining (DM), or Management Resource Planning (MRP), can be improved using the methods provided in the book. The first complete treatment on physical database design, written by the authors of the seminal, Database Modeling and Design: Logical Design, Fourth Edition. Includes an introduction to the major concepts of physical database design as well as detailed examples, using methodologies and tools most popular for relational databases today: Oracle, DB2 (IBM), and SQL Server (Microsoft). Focuses on physical database design for exploiting B+tree indexing, clustered indexes, multidimensional clustering (MDC), range partitioning, shared nothing partitioning, shared disk data placement, materialized views, bitmap indexes, automated design tools, and more!

The study of relationship databases is a core component of virtually every undergraduate computer science degree course. This new edition of Theory and Practice of Relationship Databases retains all the features that made the previous edition such a success, and goes on to give even more comprehensive and informative coverage. Written in a tutorial style and containing a great many examples and exercises as well as extensively using illustrative and explanatory graphics, the author has produced an undergraduate textbook of great depth and clarity that is very easy to follow. The subject of relational databases is brought to life by the writing style and the inclusion of an homogenous case study that reinforces the issues dealt with in each chapter. The primary objective of the book is to present a comprehensive explanation of the process of development of database application systems within the framework of a set processing paradigm. Since the majority of these applications are built as relationship systems, a complete though reasonably concise account of that model is presented. Dr. Stanczyk has achieved this by concentrating on the issues that contribute significantly to the application development while de-emphasizing purely theoretical aspects of the subject. This has led to an imaginative and highly practical textbook that will be an excellent read for the undergraduate computer science student.

Joe Celko has looked deep into the code of SQL programmers and found a consistent and troubling pattern - a frightening lack of consistency between their individual encoding schemes and those of the industries in which they operate. This translates into a series of incompatible databases, each one an island unto itself that is unable to share information with others in an age of internationalization and business interdependence. Such incompatibility severely hinders information flow and the quality of company data. Data, Measurements and Standards in SQL reveals the shift these programmers need to make to overcome this deadlock. By collecting and detailing the diverse standards of myriad industries, and then giving a declaration for the units that can be used in an SQL schema, Celko enables readers to write and implement portable data that can interface to any number of external application systems! This book doesn't limit itself to one subject, but serves as a detailed synopsis of measurement scales and data standards for all industries, thereby giving RDBMS programmers and designers the knowledge and know-how they need to communicate effectively across business boundaries. * Collects and details the diverse data standards of myriad industries under one cover, thereby creating a definitive, one-stop-shopping opportunity for database programmers. * Enables readers to write and implement portable data that can interface to any number external application systems, allowing readers to cross business boundaries and move up the career ladder. * Expert advice from one of the most-read SQL authors in the world who is well known for his ten years of service on the ANSI SQL standards committee and Readers Choice Award winning column in Intelligent Enterprise.

XML has become the lingua franca for representing business data, for exchanging information between business partners and applications, and for adding structure—and sometimes meaning—to text-based documents. XML offers some special challenges and opportunities in the area of search: querying XML can produce very precise, fine-grained results, if you know how to express and execute those queries. For software developers and systems architects: this book teaches the most useful approaches to querying XML documents and repositories. This book will also help managers and project leaders grasp how “querying XML fits into the larger context of querying and XML. Querying XML provides a comprehensive background from fundamental concepts (What is XML?) to data models (the Infoset, PSVI, XQuery Data Model), to APIs (querying XML from SQL or Java) and more. * Presents the concepts clearly, and demonstrates them with illustrations and examples; offers a thorough mastery of the subject area in a single book. * Provides comprehensive coverage of XML query languages, and the concepts needed to understand them completely (such as the XQuery Data Model). * Shows how to query XML documents and data using: XPath (the XML Path Language); XQuery, soon to be the new W3C Recommendation for querying XML; XQuery's companion XQueryX; and

SQL, featuring the SQL/XML * Includes an extensive set of XQuery, XPath, SQL, Java, and other examples, with links to downloadable code and data samples.

"Michael Scott's book could have been entitled: *Why Programming Languages Work*. It takes a fresh look at programming languages by bringing together ideas and techniques usually covered in disparate language design, compiler, computer architecture, and operating system courses. Its comprehensive and integrated presentation of language design and implementation illustrates and explains admirably the many deep and profitable connections among these fields." - Jim Larus, Microsoft Research Programming Language Pragmatics addresses the fundamental principles at work in the most important contemporary languages, highlights the critical relationship between language design and language implementation, and devotes special attention to issues of importance to the expert programmer. Thanks to its rigorous but accessible teaching style, you'll emerge better prepared to choose the best language for particular projects, to make more effective use of languages you already know, and to learn new languages quickly and completely. Features Addresses the most recent developments in programming language design, spanning more than forty different languages, including Ada 95, C, C++, Fortran 95, Java, Lisp, Scheme, ML, Modula-3, Pascal, and Prolog. Places a special emphasis on implementation issues show the techniques used by compilers and related tools influence language design, and vice versa. Covers advanced topics in language design and implementation, such as iterators, coroutines, templates (generics), separate compilation, I/O, type inference, and exception handling. Reviews language-related topics in assembly-level architecture critical for understanding what a compiler does to a program. Offers in-depth coverage of object-oriented programming, including multiple inheritance and dynamic method binding. Devotes a special section to static and dynamic linking. Includes a comprehensive chapter on concurrency, with detailed coverage of both shared-memory and message-passing languages and libraries. Provides an accessible introduction to the formal foundations of compilation (automata theory), functional programming (lambda calculus), and logic programming (predicate calculus).

Our ability to generate and collect data has been increasing rapidly. Not only are all of our business, scientific, and government transactions now computerized, but the widespread use of digital cameras, publication tools, and bar codes also generate data. On the collection side, scanned text and image platforms, satellite remote sensing systems, and the World Wide Web have flooded us with a tremendous amount of data. This explosive growth has generated an even more urgent need for new techniques and automated tools that can help us transform this data into useful information and knowledge. Like the first edition, voted the most popular data mining book by KD Nuggets readers, this book explores concepts and techniques for the discovery of patterns hidden in large data sets, focusing on issues relating to their feasibility, usefulness, effectiveness, and scalability. However, since the publication of the first edition, great progress has been made in the development of new data mining methods, systems, and applications. This new edition substantially enhances the first edition, and new chapters have been added to address recent developments on mining complex types of data—including stream data, sequence data, graph structured data, social network data, and multi-relational data. A comprehensive, practical look at the concepts and techniques you need to know to get the most out of real business data Updates that incorporate input from readers, changes in the field, and more material on statistics and machine learning Dozens of algorithms and implementation examples, all in easily understood pseudo-code and suitable for use in real-world, large-scale data mining projects Complete classroom support for instructors at www.mkp.com/datamining2e companion site

ODMG is a widely accepted standard for object database modelling; every year more companies implement it. ODMG 3.0 integrates programming languages with databases and ensures the portability of applications across platforms and DBMS products

Active database systems enhance traditional database functionality with powerful rule-processing capabilities, providing a uniform and efficient mechanism for many database system applications. Among these applications are integrity constraints, views, authorization, statistics gathering, monitoring and alerting, knowledge-based systems, expert systems, and workflow management. This significant collection focuses on the most prominent research projects in active database systems. The project leaders for each prototype system provide detailed discussions of their projects and the relevance of their results to the future of active database systems. Features: A broad overview of current active database systems and how they can be extended and improved A comprehensive introduction to the core topics of the field, including its motivation and history Coverage of active database (trigger) capabilities in commercial products Discussion of forthcoming standards

Database Modeling and Design, Fifth Edition, focuses on techniques for database design in relational database systems. This extensively revised fifth edition features clear explanations, lots of terrific examples and an illustrative case, and practical advice, with design rules that are applicable to any SQL-based system. The common examples are based on real-life experiences and have been thoroughly class-tested. This book is immediately useful to anyone tasked with the creation of data models for the integration of large-scale enterprise data. It is ideal for a stand-alone data management course focused on logical database design, or a supplement to an introductory text for introductory database management. In-depth detail and plenty of real-world, practical examples throughout Loaded with design rules and illustrative case studies that are applicable to any SQL, UML, or XML-based system Immediately useful to anyone tasked with the creation of data models for the integration of large-scale enterprise data.

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.

There is no other manual for the over 200,000 database administrators using Visio.

"The chapter on object-relational database should be a great selling point for the book. No one else has the coverage on object relational that this chapter has; for example, the other new texts emphasize the purely object model. I think that the approach here is much more practical." --Betty Salzberg, Northeastern University "The coverage of this book is wonderful, especially the cutting-edge of object-relational systems . . . [and] this is the only text I have seen that is not by Jeffrey Ullman that treats the theoretical material appropriately. The chapter on dependencies and relational design is excellent. Examples abound, the explanations are crisp and clear, and the appropriate concepts are discussed. I cannot wait to use it." --Bill Grosky, Wayne State University "This book makes an excellent text for anyone just approaching database systems. It's both an accessible refresher for those of us who have not been paying careful attention to developments in this area and a useful reference for designers and implementers who need just-in-time education." --Jim Gray, Microsoft Research "This book is excellent!" --Mike Hartstein, Oracle Corporation, Senior Director of Oracle8i Product Management This second edition relies on the same successful approach that distinguished the first: it covers the principles of database theory with unmatched thoroughness, and it rigorously links theory to the real world of database programming and administration. A careful discussion of SQL standards and a multitude of examples drawn from actual databases-Oracle, DB2, and Informix-complements the authors' concept-oriented instruction, allowing you to develop product-specific understanding and to learn the important differences between the SQL dialects that will enable you to write portable applications. New Features Focuses extensively on the object-relational model that is rapidly gaining acceptance and revolutionizing the database industry. Collection types and UDF's are thoroughly covered. Introduces new relational features of SQL taken from the latest versions of today's most popular database products, Oracle, DB2, and Informix. Offers thorough coverage of the SQL-99 standard, including additions designed to help you take full advantage of the object-relational model. Provides expanded programming examples intended to improve your understanding of transaction processing and error handling. Explains clearly the principles of logical database design, including those relating to the E-R model and normalization, with a number of new illustrations and examples. Presents the latest indexing and query processing techniques, such as bitmap indexing, and shows how to use them to achieve significant performance improvements.

Information Modeling and Relational Databases, Second Edition, provides an introduction to ORM (Object-Role Modeling) and much more. In fact, it is 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. This book is intended for anyone with a stake in the accuracy and efficacy of databases: systems analysts, information modelers, database designers and administrators, and programmers. Terry Halpin, a pioneer in the development of ORM, 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. Presents the most in-depth coverage of Object-Role Modeling available anywhere, including a thorough update of the book for ORM2, as well as UML2 and E-R (Entity-Relationship) modeling. Includes clear coverage of relational database concepts, and the latest developments in SQL and XML, including a new chapter on the impact of XML on information modeling, exchange and transformation. New and improved case studies and exercises are provided for many topics.

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.

Publisher Description

[Copyright: ccccf2451d0df89e7b0311e99b1c3298](http://www.morgankaufmann.com/copyright/ccccf2451d0df89e7b0311e99b1c3298)