

Database Systems Second Edition Thomas Connolly

Business Database Systems arms you with the knowledge to analyse, design and implement effective, robust and successful databases. This book is ideal for students of Business/Management Information Systems, or Computer Science, who will be expected to take a course in database systems for their degree programme. It is also excellently suited to any practitioner who needs to learn, or refresh their knowledge of, the essentials of database management systems.

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.

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.

Programming Language Pragmatics, Third Edition, is the most comprehensive programming language book available today. Taking the perspective that language design and implementation are tightly interconnected and that neither can be fully understood in isolation, this critically acclaimed and bestselling book has been thoroughly updated to cover the most recent developments in programming language design, including Java 6 and 7, C++0X, C# 3.0, F#, Fortran 2003 and 2008, Ada 2005, and Scheme R6RS. A new chapter on run-time program management covers virtual machines, managed code, just-in-time and dynamic compilation, reflection, binary translation and rewriting, mobile code, sandboxing, and debugging and program analysis tools. Over 800 numbered examples are provided to help the reader quickly cross-reference and access content. This text is designed for undergraduate Computer Science students, programmers, and systems and software engineers. Classic programming foundations text now updated to familiarize students with the languages they are most likely to encounter in the workforce, including including Java 7, C++, C# 3.0, F#, Fortran 2008, Ada 2005, Scheme R6RS, and Perl 6. New and expanded coverage of concurrency and run-time systems ensures students and professionals understand the most important advances driving software today. Includes over 800 numbered examples to help the reader quickly cross-reference and access content.

Multimedia Database Systems

Abstraction in Ontology-based Data Management

Querying XML

Programming Language Pragmatics

Advanced Relational Programming

Data Mining, Southeast Asia Edition

This volume aims to present recent advances in database technology from the viewpoint of the novel database paradigms proposed in the last decade. It focuses on the theory of the extended relational model and an example of an extended relational database programming language. Algres complements this work, and is available on the Internet. Audience: This work will be of interest to graduate students following advanced database courses, advanced data-oriented applications developers, and researchers in the field of database programming languages and prototyping platform for the development of software tools.

Due to market forces and technological evolution, Big Data computing is developing at an increasing rate. A wide variety of novel approaches and tools have emerged to tackle the challenges of Big Data, creating both more opportunities and more challenges for students and practitioners and analysis. Presenting a mix of industry cases and theory, Big Data Computing discusses the technical and practical issues related to Big Data in intelligent information management. Emphasizing the adoption and diffusion of Big Data tools and technologies in industry, the book covers concepts, tools, and techniques. It covers a wide range of research, and provides comparisons between state-of-the-art approaches. Comprised of five sections, the book focuses on: What Big Data is and why it is important Semantic technologies Tools and methods Business and across industries

The second edition of this bestselling title is a perfect blend of theoretical knowledge and practical application. It progresses gradually from basic to advance concepts in database management systems, with numerous solved exercises to make learning easier and interesting. New commercial database management systems.

This volume presents the results of approximately 15 years of work from researchers around the world on the use of fuzzy set theory to represent imprecision in databases. The maturity of the research in the discipline and the recent developments in commercial/industrial fuzzy set theory. In this introduction we will describe briefly how fuzzy databases fit into the overall design of database systems and then overview the organization of the text. FUZZY DATABASE LANDSCAPE The last five years have been witness to a revolution in the database research community. have changed and the consensus on what constitutes worthwhile research is in flux. Also, at this time, it is possible to gain a perspective on what has been accomplished in the area of fuzzy databases. Therefore, now is an opportune time to take stock of the past and establish evaluating future research through a better understanding of the different aspects of imprecision that a database can model [1].

Data Preparation for Data Mining Using SAS

Concepts, Design and Applications

Database Modeling and Design

Achieving Buzzword Compliance

A Practitioner’s Reference, Second Edition

With the rapid growth in the use of computers to manipulate, process, and reason about multimedia data, the problem of how to store and retrieve such data is becoming increasingly important. Thus, although the field of multimedia database systems is only about 5 years old, it is rapidly becoming a focus for much excitement and research effort. Multimedia database systems are intended to provide unified frameworks for requesting and integrating information in a wide variety of formats, such as audio and video data, document data, and image data. Such data often have special storage requirements that are closely coupled to the various kinds of devices that are used for recording and presenting the data, and for each form of data there are often multiple representations and multiple standards - all of which make the database integration task quite complex. Some of the problems include: - what a multimedia database query means - what kinds of languages to use for posing queries - how to develop compilers for such languages - how to develop indexing structures for storing media on ancillary devices - data compression techniques - how to present and author presentations based on user queries. Although approaches are being developed for a number of these problems, they have often been ad hoc in nature, and there is a need to provide a principled theoretical foundation.

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.

This book is a result of the Seventh International Conference on Information Systems Development-Methods and Tools, Theory and Practice held in Bled, Slovenia, September 21-23, 1998. The purpose of the conference was to address issues facing academia and industry when specifying, developing, managing, and improving information computerized systems. During the past few years, many new concepts and approaches emerged in the Information Systems Development (ISD) field. The various theories, methods, and tools available to system developers also bring problems such as choosing the most effective approach for a specific task. This conference provides a meeting place for IS researchers and practitioners from Eastern and Western Europe as well as from other parts of the world. An objective of the conference is not only to share scientific knowledge and in interests but to establish strong professional ties among the participants. The Seventh International Conference on Information Systems Development-ISD’98 continues the concepts of the first Polish-Scandinavian Seminar on Current Trends in Information Systems Development Methodologies held in Gdansk, Poland in 1988. Through the years, the Seminar developed into the International Conference on Information Systems Development. ISD’99 will be held in Boise, Idaho. The selection of papers was carried out by the International Program Committee. All papers were reviewed in advance by three people. Papers were judged according to their originality, relevance, and presentation quality. All papers were judged only on their own merits, independent of other submissions.

Database SystemsThe Complete BookPearson Higher Ed

Java Data Mining: Strategy, Standard, and Practice

Logical Design

Valuepack

Large Scale and Big Data

Database Systems

Principles of Transaction Processing

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 Publisher Description

The book has an introductory chapter that gets the reader started quickly with programming in Perl. The initial part of the book discusses Perl expressions, statements, control flow, built-in data types such as arrays and hashes, and complex data structures built using references. On Perl has several chapters covering specialized topics. The chapter on socket-based network programming deals with forking and using fork to write complex interactive client-server programs. There is a chapter with in-depth discussion of CGI programming including error-handling and security issues that arise. The chapter on web-client programming deals with writing programs that access Web pages, fill up GET and POST forms, handle cookies and redirected Web pages. The book has several unique chapters not found in any other book on Perl in the market. The chapter on security discusses hashes such as MD5, message authentication codes (MACs), digital signature schemes, and encryption techniques such as DES, Rijndael, and RSA. Other chapters deal with writing recursive programs that work with files and directories; this chapter also discusses predefined modules that deal with portability in file names and paths across operating systems, recursive traversal of file hierarchies and tarring and untarring of files. The chapter on functional programming illustrates that Perl functions are first-class, can be used to write closures and can be composed to form more complex functions. In particular, this can be useful for programming in artificial intelligence.

Are you responsible for designing and creating the databases that keep your business running? Or are you studying for a module in database design? If so, Database Solutions is for you! This fully revised and updated edition will make the database design and build process smoother, quicker and more reliable. Recipe for database success Take one RDMS of any of the major commercial products will do: Oracle, Informix, SQL Server, Access, Paradox Add one thorough reading of Database Solutions if you are an inexperienced database designer, or one recap of the methodology if you are an old hand Use the design and implementation frameworks to plan your timetable, use a common data model that fits your requirements and adapt as necessary

The Complete Book

Principles and Applications

Information Modeling and Relational Databases

A Practical Approach to Design, Implementation, and Management

Physical Database Design

Database Solutions

Information Security Policies and Procedures: A Practitioner’s Reference, Second Edition illustrates how policies and procedures support the efficient running of an organization. This book is divided into two parts, an

overview of security policies and procedures, and an information security reference guide. This volume points out how security

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

There are more than one billion documents on the Web, with the count continually rising at a pace of over one million new documents per day. As information increases, the motivation and interest in data warehousing and

mining research and practice remains high in organizational interest. The Encyclopedia of Data Warehousing and Mining, Second Edition, offers thorough exposure to the issues of importance in the rapidly changing field of

data warehousing and mining. This essential reference source informs decision makers, problem solvers, and data mining specialists in business, academia, government, and other settings with over 300 entries on theories,

methodologies, functionalities, and applications.

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

DW 2.0: The Architecture for the Next Generation of Data Warehousing

Business Database Systems

Information Security Policies and Procedures

A Practical Guide for Architecture, Design, and Implementation

XQuery, XPath, and SQL/XML in context

Perl for Students and Professionals

The latest edition of a popular text and reference on database research, with substantial new material and revision; covers classical literature and recent hot topics. Lessons from database research have been applied in academic fields ranging from bioinformatics to next-generation Internet architecture and in industrial uses including Web-based e-commerce and search engines. The core ideas in the field have become increasingly influential. This text provides both students and professionals with a grounding in database research and a technical context for understanding recent innovations in the field. The readings included treat the most important issues in the database area--the basic material for any DBMS professional. This fourth edition has been substantially updated and revised, with 21 of the 48 papers new to the edition, four of them published for the first time. Many of the sections have been newly organized, and each section includes a new or substantially revised introduction that discusses the context, motivation, and controversies in a particular area, placing it in the broader perspective of database research. Two introductory articles, never before published, provide an organized, current introduction to basic knowledge of the field; one discusses the history of data models and query languages and the other offers an architectural overview of a database system. The remaining articles range from the classical literature on database research to treatments of current hot topics, including a paper on search engine architecture and a paper on application servers, both written expressly for this edition. The result is a collection of papers that are seminal and also accessible to a reader who has a basic familiarity with database systems.

Principles of Transaction Processing is a comprehensive guide to developing applications, designing systems, and evaluating engineering products. The book provides detailed discussions of the internal workings of transaction processing systems, and it discusses how these systems work and how best to utilize them. It covers the architecture of Web Application Servers and transactional communication paradigms. The book is divided into 11 chapters, which cover the following: Overview of transaction processing application and system structure Software abstractions found in transaction processing systems Architecture of multitier applications and the functions of transactional middleware and database servers Queued transaction processing and its internals, with IBM’s Websphere MQ and Oracle’s Stream AQ as examples Business process management and its mechanisms Description of the two-phase locking function, B-tree locking and multigranularity locking used in SQL database systems and nested transaction locking System recovery and its failures Two-phase commit protocol Comparison between the tradeoffs of replicating servers versus replication resources Transactional middleware products and standards Future trends, such as cloud computing platforms, composing scalable systems using distributed computing components, the use of flash storage to replace disks and data streams from sensor devices as a source of transaction requests. The text meets the needs of systems professionals, such as IT application programmers who construct TP applications, application analysts, and product developers. The book will also be invaluable to students and novices in application programming. Complete revision of the classic “non mathematical” transaction processing reference for systems professionals. Updated to focus on the needs of transaction processing via the Internet-- the main focus of business data processing investments, via web application servers, SOA, and important new TP standards. Retains the practical, non-mathematical, but thorough conceptual basis of the first edition.

Database Systems is ideal for a one- or two-term course in database management or database design in an undergraduate or graduate level course. With its comprehensive coverage, this book can also be used as a reference for IT professionals. This best-selling text introduces the theory behind databases in a concise yet comprehensive manner, providing database design methodology that can be used by both technical and non-technical readers. The methodology for relational Database Management Systems is presented in simple, step-by-step instructions in conjunction with a realistic worked example using three explicit phases—conceptual, logical, and physical database design. Teaching and Learning Experience This program presents a better teaching and learning experience—for you and your students. It provides: Database Design Methodology that can be Used by Both Technical and Non-technical Readers A Comprehensive

Introduction to the Theory behind Databases A Clear Presentation that Supports Learning

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 Fuzzy Databases

Encyclopedia of Data Warehousing and Mining, Second Edition

Evolution and Challenges in System Development

On Perl

Database Systems:A Practical Approach to Design, Implementation and Management with Corporate Computer and Network Security:(International Edition) and Making the Team (International Edition) with Success in Your Project Unix and C Programming

This undergraduate textbook first introduces basic electronic circuitry before explaining more advanced elements such as the Arithmetic Logic Unit, sequential circuits, and finally microprocessors. In keeping with this

integrated and graduated approach, the authors then explain the relationship to first assembly programming, then higher-level languages, and finally computer organisation. Authors use the Raspberry Pi and ARM

microprocessors for their explanations The material has been extensively class tested at TU Eindhoven by an experienced team of lecturers and researchers. This is a modern, holistic treatment of well-established topics, valuable for undergraduate students of computer science and electronics engineering and for self-study. The authors use the Raspberry Pi and ARM microprocessors for their explanations.

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 indepth 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.

Best-selling author and database expert with more than 25 years of experience modeling application and enterprise data, Dr. Michael Blaha provides tried and tested data model patterns, to help readers avoid common modeling mistakes and unnecessary frustration on their way to building effective data models. Unlike the typical methodology book, Patterns of Data Modeling provides advanced techniques for those who have mastered the basics. Recognizing that database representation sets the path for software, determines its flexibility, affects its quality, and influences whether it succeeds or fails, the text focuses on databases rather than programming. It is one of the first books to apply the popular patterns perspective to database systems and data models. It offers practical advice on the core aspects of applications and provides authoritative coverage of mathematical templates, antipatterns, archetypes, identity, canonical models, and relational database design.

Understand the language and vocabulary of Data Architecture. The Data Architecture field is rife with terms that have become "fashionable". Some of the terms began with very specific, specialized, meanings – but as their

use spread, they lost the precision of their technical definitions and become, well, "buzzwords". A buzzword is "a word or expression from a particular subject area that has become fashionable because it has been used a

lot". Compliance is "the obeying of an accepted principle or instruction that states the way things are or should be done." The assignment is to take buzzwords and follow rules to use them correctly. We cut through the

hype to arrive at buzzword compliance – the state where you fully understand the words that in fact have real meaning in the data architecture industry. This book will rationalize the various ways all these terms are

defined. Of necessity, the book must address all aspects of describing an enterprise and its data management technologies. This includes a wide range of subjects, from entity/relationship modeling, through the semantic web, to database issues like relational and "beyond relational" ("NoSQL") approaches. In each case, the definitions for the subject are meant to be detailed enough to make it possible to understand basic principles—while

recognizing that a full understanding will require consulting the sources where they are more completely described. The book's Glossary contains a catalogue of definitions and its Bibliography contains a comprehensive set

of references.

Patterns of Data Modeling

Psychology, Pedagogy, and Assessment in Serious Games

Foundations of Multidimensional and Metric Data Structures

The Database Professional's Guide to Exploiting Indexes, Views, Storage, and More

Joe Celko's SQL for Smarties

Processing and Management

Database systems and database design technology have undergone significant evolution in recent years. The relational data model and relational database systems dominate business applications; in turn, they are extended by other technologies like data warehousing,

OLAP, and data mining. How do you model and design your database application in consideration of new technology or new business needs? In the extensively revised fifth edition, you'll get clear explanations, lots of terrific examples and an illustrative case, and the really practical advice you have come to count on--with design rules that are applicable to any SQL-based system. But you'll also get plenty to help you grow from a new database designer to an experienced designer developing industrial-sized systems. 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.

Effectively documenting data services is a crucial issue in any organization, not only for governing data but also for interoperation purposes. Indeed, in order to fully realize the promises and benefits of a data-driven society, data-driven approaches need to be resilient, transparent, and fully accountable. This book, Abstraction in Ontology-based Data Management, proposes a new approach to automatically associating formal semantic description to data services, thus bringing them into compliance with the FAIR (Findable, Accessible, Interoperable, and Reusable) guiding principles. The approach is founded on the Ontology-based Data Management (OBDM) paradigm, in which a domain ontology is used to provide a high-level semantic layer mapped to the source schema of an organization containing

data, thus abstracting from the technical details of the data layer implementation. A formal framework for a novel reasoning task in OBDM, called Abstraction, is introduced in which a data service is assumed to be expressed as a query over the source schema, and the aim is to derive a query over the ontology that semantically describes the given data service best with respect to the underlying OBDM specification. In a general scenario that uses the most popular languages in the OBDM literature, an in-depth complexity analysis of two computational problems associated with the framework is carried out. Also investigated is the problem of expressing abstractions in a non-monotonic query language as well as the impact of adding inequalities. Regarding the latter, the problem of answering queries

with inequalities over lightweight ontologies is first studied. Lastly, the author illustrates how the achieved results contribute to new results in the Semantic Web context and in the Relational Database theory. The book will be of interest to all those engaged in Artificial Intelligence and Data Management.

"This book addresses issues the potential of games to support learning and change behaviour offering empirical evidence pertaining to the effectiveness of Serious Games in the key areas of psychology, pedagogy, and assessment"--

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.

Logic Gates, Circuits, Processors, Compilers and Computers

A Deep Dive into How Distributed Data Systems Work

Issues and Research Directions

Data Architecture Language and Vocabulary

Joe Celko's Data, Measurements and Standards in SQL

Database Internals

A major technological trend for large database systems has been the introduction of ever-larger mass storage systems. This allows computing centers and business data processing installations to maintain on line their program libraries, less frequently used data files, transaction logs and backup copies under unified system control. Tapes, disks and drums are classical examples of mass storage media. The more recent IBM 3851 Mass Storage Facility, part of the IBM 3850 Mass Storage System, represents a new direction in mass storage development, namely, it is two-dimensional. With the maturity of

magnetic bubble technology, more sophisticated, massive, multi-trillion-bit storage systems are not far in the future. While large in capacity, mass storage systems have in general relatively long access times. Since record access probabilities are usually not uniform, various algorithms have been devised to position the records to decrease the average access time. The first two chapters of this book are devoted mainly to such algorithmic studies in linear and two-dimensional mass storage systems. In the third chapter, we view the bubble memory as more than a storage medium. In fact, we discuss different

structures where routine operations, such as data rearrangement, sorting, searching, etc., can be done in the memory itself, freeing the CPU for more complicated tasks. The problems discussed in this book are combinatorial in nature. Recent years have seen an explosive growth in the use of new database applications such as CAD/CAM systems, spatial information systems, and multimedia information systems. The needs of these applications are far more complex than traditional business applications. They call for support of objects

with complex data types, such as images and spatial objects, and for support of objects with wildly varying numbers of index terms, such as documents. Traditional indexing techniques such as the B-tree and its variants do not efficiently support these applications, and so new indexing mechanisms have been developed. As a result of the demand for database support for new applications, there has been a proliferation of new indexing techniques. The need for a book addressing indexing problems in advanced applications is evident. For practitioners and database and application developers, this book

explains best practice, guiding the selection of appropriate indexes for each application. For researchers, this book provides a foundation for the development of new and more robust indexes. For newcomers, this book is an overview of the wide range of advanced indexing techniques. Indexing Techniques for Advanced Database Systems is suitable as a secondary text for a graduate level course on indexing techniques, and as a reference for researchers and practitioners in industry.

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. 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. Joe Celko's Thinking in Sets: Auxiliary, Temporal, and Virtual Tables in SQL Readings in Database Systems Algorithmic Studies in Mass Storage Systems Advanced SQL Programming Indexing Techniques for Advanced Database Systems Big Data Computing

Large Scale and Big Data: Processing and Management provides readers with a central source of reference on the data management techniques currently available for large-scale data processing. Presenting chapters written by leading researchers, academics, and practitioners, it addresses the fundamental challenges associated with Big Data processing tools and techniques across a range of computing environments. The book begins by discussing the basic concepts and tools of large-scale Big Data processing and cloud computing. It also provides an overview of different

programming models and cloud-based deployment models. The book ' s second section examines the usage of advanced Big Data processing techniques in different domains, including semantic web, graph processing, and stream processing. The third section discusses advanced topics of Big Data processing such as consistency management, privacy, and security. Supplying a comprehensive summary from both the research and applied perspectives, the book covers recent research discoveries and applications, making it an ideal reference for a wide range of audiences, including researchers and academics working on databases, data mining, and web scale data processing. After reading this book, you will gain a fundamental understanding of how to use Big Data-processing tools and techniques effectively across application domains. Coverage includes cloud data

management architectures, big data analytics visualization, data management, analytics for vast amounts of unstructured data, clustering, classification, link analysis of big data, scalable data mining, and machine learning techniques. 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 book is intended to provide an insight into the DBMS concepts. An effort has been made to familiarize the readers with the concepts of database normalization, concurrency control, deadlock handling and recovery etc., which are extremely vital for a clear understanding of DBMS. To familiarize the readers with the equivalence amongst Relational Algebra, Tuple Relational Calculus, and SQL, a large number of equivalent queries have been provided. The concepts of normalization have been elaborated very systematically by fully covering the underlying concepts of

functional dependencies, multi-valued dependencies, join dependencies, loss-less-join decomposition, dependency-preserving decomposition etc. It is hoped that with the help of the information provided in the text, a reader will be able to design a flawless database. Also, the concepts of serializability, concurrency control, deadlock handling and log-based recovery have been covered in full detail. An overview has also been provided of the issues related to distributed-databases. Database Management Systems