
Access Free Relational Database Practices Bridging The Gap Between The Theory Of Database Design And Real World Practices

As recognized, adventure as competently as experience very nearly lesson, amusement, as well as understanding can be gotten by just checking out a book **Relational Database Practices Bridging The Gap Between The Theory Of Database Design And Real World Practices** then it is not directly done, you could believe even more just about this life, all but the world.

We provide you this proper as with ease as simple pretentiousness to get those all. We have the funds for Relational Database Practices Bridging The Gap Between The Theory Of Database Design And Real World Practices and numerous book collections from fictions to scientific research in any way. along with them is this Relational Database Practices Bridging The Gap Between The Theory Of Database Design And Real World Practices that can be your partner.

BF4 - CONRAD CAYDEN

This book constitutes the refereed proceedings of the 20th International Conference on Scientific and Statistical Database Management, SSDBM 2008, held in Hong Kong, China, in July 2008. The 28 revised full papers, 7 revised short papers and 8 poster and demo papers presented together with 3 invited talks were carefully reviewed and selected from 84 submissions. The papers are organized in topical sections on query optimization in scientific databases, privacy, searching and mining

graphs, data streams, scientific database applications, advanced indexing methods, data mining, as well as advanced queries and uncertain data.

This book, gathering the Proceedings of the 2018 Computing Conference, offers a remarkable collection of chapters covering a wide range of topics in intelligent systems, computing and their real-world applications. The Conference attracted a total of 568 submissions from pioneering researchers, scientists, industrial engineers, and students from all around the world. These submissions under-

went a double-blind peer review process. Of those 568 submissions, 192 submissions (including 14 poster papers) were selected for inclusion in these proceedings. Despite computer science's comparatively brief history as a formal academic discipline, it has made a number of fundamental contributions to science and society—in fact, along with electronics, it is a founding science of the current epoch of human history ('the Information Age') and a main driver of the Information Revolution. The goal of this conference is to provide a plat-

form for researchers to present fundamental contributions, and to be a premier venue for academic and industry practitioners to share new ideas and development experiences. This book collects state of the art chapters on all aspects of Computer Science, from classical to intelligent. It covers both the theory and applications of the latest computer technologies and methodologies. Providing the state of the art in intelligent methods and techniques for solving real-world problems, along with a vision of future research, the book will be interesting and valuable for a broad readership.

The two-volume set LNCS 7031 and 7032 constitutes the proceedings of the 10th International Semantic Web Conference, ISWC 2011, held in Bonn, Germany, in October 2011. Part I, LNCS 7031, contains 50 research papers which were carefully reviewed and selected from 264 submissions. The 17 semantic web in-use track papers contained in part II, LNCS 7032, were selected from 75 submissions. This volume also contains 15 doctoral consortium papers, selected from 31 submissions. The topics covered are: ontologies and seman-

tics; database, IR, and AI technologies for the semantic web; management of semantic web data; reasoning over semantic web data; search, query, integration, and analysis on the semantic web; robust and scalable knowledge management and reasoning on the web; interacting with semantic web data; ontology modularity, mapping, merging and alignment; languages, tools, and methodologies for representing and managing semantic web data; ontology, methodology, evaluation, reuse, extraction and evolution; evaluation of semantic web technologies or data; specific ontologies and ontology pattern for the semantic web; new formalisms for semantic web; user interfaces to the semantic web; cleaning, assurance, and provenance of semantic web data; services, and processes; social semantic web, evaluation of semantic web technology; semantic web population from the human web.

Here are selected papers of the First International Workshop on Model-Based Software and Data Integration 2008, part of the Berlin Software Integration Week 2008. The 9 revised full papers and 3 invited lectures were carefully picked from numerous

submissions.

Six-Step Relational Database DesignTM bridges the gaps between database theory, database modeling, and database implementation by outlining a simple but reliable six-step process for accurately modeling user data on a Crow's Foot Relational Model Diagram, and then demonstrating how to implement this model on any relational database management system. The second edition contains a new chapter on implementation that goes through the steps necessary to implement each of the case studies on a relational database management system, clearly relating the design to implementation and database theory. In addition, questions are also included at the end of each of the six steps and one of the previous case studies has been replaced, making the case study selection more diverse. Six-Step Relational Database DesignTM uses three case studies and starts with a statement of the problem by the client and then goes through the six steps necessary to create a reliable and accurate data model of the client's business requirements. This model can

then be used to implement the database on any relational database management system. Six-Step Relational Database Design™ should be used as a handbook for students and professionals in the software-development field. The technique described in this book can be used by students for quickly developing relational databases for their applications, and by professionals for developing sturdy, reliable, and accurate relational database models for their software applications.

All of today's mainstream database products support the SQL language, and relational theory is what SQL is supposed to be based on. But are those products truly relational? Sadly, the answer is no. This book shows you what a real relational product would be like, and how and why it would be so much better than what's currently available. With this unique book, you will: Learn how to see database systems as programming systems Get a careful, precise, and detailed definition of the relational model Explore a detailed analysis of SQL from a relational point of view There are literally hundreds of books on relational theory or the SQL

language or both. But this one is different. First, nobody is more qualified than Chris Date to write such a book. He and Ted Codd, inventor of the relational model, were colleagues for many years, and Chris's involvement with the technology goes back to the time of Codd's first papers in 1969 and 1970. Second, most books try to use SQL as a vehicle for teaching relational theory, but this book deliberately takes the opposite approach. Its primary aim is to teach relational theory as such. Then it uses that theory as a vehicle for teaching SQL, showing in particular how that theory can help with the practical problem of using SQL correctly and productively. Any computer professional who wants to understand what relational systems are all about can benefit from this book. No prior knowledge of databases is assumed. This book constitutes the refereed proceedings of the 11th International Conference on Model Transformation, ICMT 2018, held as part of STAF 2018, in Toulouse, France, in June 2018. The 9 full papers were carefully reviewed and selected from 24 submissions. This book also presents 1 invited paper. The papers include re-

search, application, and tool demonstration papers presented in the context of four sessions on verification of model transformations, model transformation tools, transformation reuse and graph transformations.

Architecture of a Database System presents an architectural discussion of DBMS design principles, including process models, parallel architecture, storage system design, transaction system implementation, query processor and optimizer architectures, and typical shared components and utilities.

This book discusses how model-based approaches can improve the daily practice of software professionals. This is known as Model-Driven Software Engineering (MDSE) or, simply, Model-Driven Engineering (MDE). MDSE practices have proved to increase efficiency and effectiveness in software development, as demonstrated by various quantitative and qualitative studies. MDSE adoption in the software industry is foreseen to grow exponentially in the near future, e.g., due to the convergence of software development and business analysis. The aim of this book is

to provide you with an agile and flexible tool to introduce you to the MDSE world, thus allowing you to quickly understand its basic principles and techniques and to choose the right set of MDSE instruments for your needs so that you can start to benefit from MDSE right away. The book is organized into two main parts. The first part discusses the foundations of MDSE in terms of basic concepts (i.e., models and transformations), driving principles, application scenarios, and current standards, like the well-known MDA initiative proposed by OMG (Object Management Group) as well as the practices on how to integrate MDSE in existing development processes. The second part deals with the technical aspects of MDSE, spanning from the basics on when and how to build a domain-specific modeling language, to the description of Model-to-Text and Model-to-Model transformations, and the tools that support the management of MDSE projects. The second edition of the book features: a set of completely new topics, including: full example of the creation of a new modeling language (IFML), discussion of modeling issues and approach-

es in specific domains, like business process modeling, user interaction modeling, and enterprise architecture complete revision of examples, figures, and text, for improving readability, understandability, and coherence better formulation of definitions, dependencies between concepts and ideas addition of a complete index of book content In addition to the contents of the book, more resources are provided on the book's website <http://www.mdse-book.com>, including the examples presented in the book.

Refactoring has proven its value in a wide range of development projects—helping software professionals improve system designs, maintainability, extensibility, and performance. Now, for the first time, leading agile methodologist Scott Ambler and renowned consultant Pramodkumar Sadalage introduce powerful refactoring techniques specifically designed for database systems. Ambler and Sadalage demonstrate how small changes to table structures, data, stored procedures, and triggers can significantly enhance virtually any database design—without changing semantics. You'll learn how to evolve database sche-

mas in step with source code—and become far more effective in projects relying on iterative, agile methodologies. This comprehensive guide and reference helps you overcome the practical obstacles to refactoring real-world databases by covering every fundamental concept underlying database refactoring. Using start-to-finish examples, the authors walk you through refactoring simple standalone database applications as well as sophisticated multi-application scenarios. You'll master every task involved in refactoring database schemas, and discover best practices for deploying refactorings in even the most complex production environments. The second half of this book systematically covers five major categories of database refactorings. You'll learn how to use refactoring to enhance database structure, data quality, and referential integrity; and how to refactor both architectures and methods. This book provides an extensive set of examples built with Oracle and Java and easily adaptable for other languages, such as C#, C++, or VB.NET, and other databases, such as DB2, SQL Server, MySQL, and Sybase. Using this

book's techniques and examples, you can reduce waste, rework, risk, and cost—and build database systems capable of evolving smoothly, far into the future.

Create database designs that scale, meet business requirements, and inherently work toward keeping your data structured and usable in the face of changing business models and software systems. This book is about database design theory. Design theory is the scientific foundation for database design, just as the relational model is the scientific foundation for database technology in general. Databases lie at the heart of so much of what we do in the computing world that negative impacts of poor design can be extraordinarily widespread. This second edition includes greatly expanded coverage of exotic and little understood normal forms such as: essential tuple normal form (ETNF), redundancy free normal form (RFNF), superkey normal form (SKNF), sixth normal form (6NF), and domain key normal form (DKNF). Also included are new appendices, including one that provides an in-depth look into the crucial notion of data consistency. Sequencing

of topics has been improved, and many explanations and examples have been rewritten and clarified based upon the author's teaching of the content in instructor-led courses. This book aims to be different from other books on design by bridging the gap between the theory of design and the practice of design. The book explains theory in a way that practitioners should be able to understand, and it explains why that theory is of considerable practical importance. Reading this book provides you with an important theoretical grounding on which to do the practical work of database design. Reading the book also helps you in going to and understanding the more academic texts as you build your base of knowledge and expertise. Anyone with a professional interest in database design can benefit from using this book as a stepping-stone toward a more rigorous design approach and more lasting database models. What You Will Learn Understand what design theory is and is not Be aware of the two different goals of normalization Know which normal forms are truly significant Apply design theory in practice Be familiar

with techniques for dealing with redundancy Understand what consistency is and why it is crucially important Who This Book Is For Those having a professional interest in database design, including data and database administrators; educators and students specializing in database matters; information modelers and database designers; DBMS designers, implementers, and other database vendor personnel; and database consultants. The book is product independent.

Leverage the power of PostgreSQL 11 to build powerful database and data warehousing applications Key Features Monitor, secure, and fine-tune your PostgreSQL 11 database Learn client-side and server-side programming using SQL and PL/pgSQL Discover tips on implementing efficient database solutions Book Description PostgreSQL is one of the most popular open source database management systems in the world, and it supports advanced features included in SQL standards. This book will familiarize you with the latest features in PostgreSQL 11, and get you up and running with building efficient PostgreSQL database solu-

tions from scratch. Learning PostgreSQL, 11 begins by covering the concepts of relational databases and their core principles. You'll explore the Data Definition Language (DDL) and commonly used DDL commands supported by ANSI SQL. You'll also learn how to create tables, define integrity constraints, build indexes, and set up views and other schema objects. As you advance, you'll come to understand Data Manipulation Language (DML) and server-side programming capabilities using PL/pgSQL, giving you a robust background to develop, tune, test, and troubleshoot your database application. The book will guide you in exploring NoSQL capabilities and connecting to your database to manipulate data objects. You'll get to grips with using data warehousing in analytical solutions and reports, and scaling the database for high availability and performance. By the end of this book, you'll have gained a thorough understanding of PostgreSQL 11 and developed the necessary skills to build efficient database solutions. What you will learn Understand the basics of relational databases, relational algebra, and data modeling

tall a PostgreSQL server, create a database, and implement your data model. Create tables and views, define indexes and stored procedures, and implement triggers. Make use of advanced data types such as Arrays, hstore, and JSONB. Connect your Python applications to PostgreSQL and work with data efficiently. Identify bottlenecks to enhance reliability and performance of database applications. Who this book is for This book is for you if you're interested in learning about PostgreSQL from scratch. Those looking to build solid database or data warehousing applications or wanting to get up to speed with the latest features of PostgreSQL 11 will also find this book useful. No prior knowledge of database programming or administration is required to get started.

Medical Imaging in Clinical Practice is a compendium of the various applications of imaging modalities in specific clinical conditions. It captures in an easy to read manner, the experiences of various experts drawn from across the globe. It explores the conventional techniques, advanced modalities and on going research efforts in the ever widening horizon

of medical imaging. The various topics would be relevant to residents, radiologists and specialists who order and interpret various medical imaging procedures. It is an essential for the inquisitive mind, seeking to understand the scope of medical imaging in clinical practice.

Because databases often stay in production for decades, careful design is critical to making the database serve the needs of your users over years, and to avoid subtle errors or performance problems. In this book, C.J. Date, a leading exponent of relational databases, lays out the principles of good database design.

This series equips the student with clear understanding of the concepts of informatics. Based on the latest recommendation of CBSE, this series provides in-depth knowledge to students on Informatics Practices under one cover. This series is prepared with extensive practice papers, assignments, chapter-wise solved and unsolved examples including CBSE sample paper questions and previous year's questions.

"It's not easy to find such a generous book on big data and databases. Fortunately, this book is the

one." Feng Yu. Computing Reviews. June 28, 2016. This is a book for enterprise architects, database administrators, and developers who need to understand the latest developments in database technologies. It is the book to help you choose the correct database technology at a time when concepts such as Big Data, NoSQL and NewSQL are making what used to be an easy choice into a complex decision with significant implications. The relational database (RDBMS) model completely dominated database technology for over 20 years. Today this "one size fits all" stability has been disrupted by a relatively recent explosion of new database technologies. These paradigm-busting technologies are powering the "Big Data" and "NoSQL" revolutions, as well as forcing fundamental changes in databases across the board. Deciding to use a relational database was once truly a no-brainer, and the various commercial relational databases competed on price, performance, reliability, and ease of use rather than on fundamental architectures. Today we are faced with choices between radically different database technologies. Choosing the right

database today is a complex undertaking, with serious economic and technological consequences. Next Generation Databases demystifies today's new database technologies. The book describes what each technology was designed to solve. It shows how each technology can be used to solve real world application and business problems. Most importantly, this book highlights the architectural differences between technologies that are the critical factors to consider when choosing a database platform for new and upcoming projects. Introduces the new technologies that have revolutionized the database landscape Describes how each technology can be used to solve specific application or business challenges Reviews the most popular new wave databases and how they use these new database technologies

The need to handle increasingly larger data volumes is one factor driving the adoption of a new class of nonrelational "NoSQL" databases. Advocates of NoSQL databases claim they can be used to build systems that are more performant, scale better, and are easier to program. NoSQL Distilled

is a concise but thorough introduction to this rapidly emerging technology. Pramod J. Sadalage and Martin Fowler explain how NoSQL databases work and the ways that they may be a superior alternative to a traditional RDBMS. The authors provide a fast-paced guide to the concepts you need to know in order to evaluate whether NoSQL databases are right for your needs and, if so, which technologies you should explore further. The first part of the book concentrates on core concepts, including schemaless data models, aggregates, new distribution models, the CAP theorem, and map-reduce. In the second part, the authors explore architectural and design issues associated with implementing NoSQL. They also present realistic use cases that demonstrate NoSQL databases at work and feature representative examples using Riak, MongoDB, Cassandra, and Neo4j. In addition, by drawing on Pramod Sadalage's pioneering work, NoSQL Distilled shows how to implement evolutionary design with schema migration: an essential technique for applying NoSQL databases. The book concludes by describing how NoSQL is ushering in a new age of Poly-

glot Persistence, where multiple data-storage worlds coexist, and architects can choose the technology best optimized for each type of data access. This book is a guide to designing and building knowledge graphs from enterprise relational databases in practice. It presents a principled framework centered on mapping patterns to connect relational databases with knowledge graphs, the roles within an organization responsible for the knowledge graph, and the process that combines data and people. The content of this book is applicable to knowledge graphs being built either with property graph or RDF graph technologies. Knowledge graphs are fulfilling the vision of creating intelligent systems that integrate knowledge and data at large scale. Tech giants have adopted knowledge graphs for the foundation of next-generation enterprise data and metadata management, search, recommendation, analytics, intelligent agents, and more. We are now observing an increasing number of enterprises that seek to adopt knowledge graphs to develop a competitive edge. In order for enterprises to design and build knowledge graphs, they

need to understand the critical data stored in relational databases. How can enterprises successfully adopt knowledge graphs to integrate data and knowledge, without boiling the ocean? This book provides the answers.

Information Modeling and Relational Databases provides an introduction to ORM (Object Role Modeling)-and much more. In fact, it's the only book to go beyond introductory coverage and provide all of the in-depth instruction you need to transform knowledge from domain experts into a sound database design. Inside, ORM authority Terry Halpin blends conceptual information with practical instruction that will let you begin using ORM effectively as soon as possible. Supported by examples, exercises, and useful background information, his step-by-step approach teaches you to develop a natural-language-based ORM model and then, where needed, abstract ER and UML models from it. This book will quickly make you proficient in the modeling technique that is proving vital to the development of accurate and efficient databases that best meet real business objectives. The most in-depth coverage of Ob-

ject Role Modeling available anywhere-written by a pioneer in the development of ORM. Provides additional coverage of Entity Relationship (ER) modeling and the Unified Modeling Language-all from an ORM perspective. Intended for anyone with a stake in the accuracy and efficacy of databases: systems analysts, information modelers, database designers and administrators, instructors, managers, and programmers. Explains and illustrates required concepts from mathematics and set theory.

"This book takes the somewhat daunting process of database design and breaks it into completely manageable and understandable components. Mike's approach whilst simple is completely professional, and I can recommend this book to any novice database designer." --Sandra Barker, Lecturer, University of South Australia, Australia
 "Databases are a critical infrastructure technology for information systems and today's business. Mike Hernandez has written a literate explanation of database technology--a topic that is intricate and often obscure. If you design databases yourself, this book will educate you

about pitfalls and show you what to do. If you purchase products that use a database, the book explains the technology so that you can understand what the vendor is doing and assess their products better." --Michael Blaha, consultant and trainer, author of *A Manager's Guide to Database Technology* "If you told me that Mike Hernandez could improve on the first edition of *Database Design for Mere Mortals* I wouldn't have believed you, but he did! The second edition is packed with more real-world examples, detailed explanations, and even includes database-design tools on the CD-ROM! This is a must-read for anyone who is even remotely interested in relational database design, from the individual who is called upon occasionally to create a useful tool at work, to the seasoned professional who wants to brush up on the fundamentals. Simply put, if you want to do it right, read this book!" ---Matt Greer, Process Control Development, The Dow Chemical Company "Mike's approach to database design is totally common-sense based, yet he's adhered to all the rules of good relational database design. I use

Mike's books in my starter database-design class, and I recommend his books to anyone who's interested in learning how to design databases or how to write SQL queries." --Michelle Poolet, President, MVDS, Inc. "Slapping together sophisticated applications with poorly designed data will hurt you just as much now as when Mike wrote his first edition, perhaps even more. Whether you're just getting started developing with data or are a seasoned pro; whether you've read Mike's previous book or this is your first; whether you're happier letting someone else design your data or you love doing it yourself--this is the book for you. Mike's ability to explain these concepts in a way that's not only clear, but fun, continues to amaze me." - -From the Foreword by Ken Getz, MCW Technologies, coauthor *ASP.NET Developer's JumpStart* "The first edition of Mike Hernandez's book *Database Design for Mere Mortals* was one of the few books that survived the cut when I moved my office to smaller quarters. The second edition expands and improves on the original in so many ways. It is not only a good, clear read, but con-

tains a remarkable quantity of clear, concise thinking on a very complex subject. It's a must for anyone interested in the subject of database design." - -Malcolm C. Rubel, Performance Dynamics Associates "Mike's excellent guide to relational database design deserves a second edition. His book is an essential tool for fledgling Microsoft Access and other desktop database developers, as well as for client/server pros. I recommend it highly to all my readers." --Roger Jennings, author of *Special Edition Using Access 2002* "There are no silver bullets! Database technology has advanced dramatically, the newest crop of database servers perform operations faster than anyone could have imagined six years ago, but none of these technological advances will help fix a bad database design, or capture data that you forgot to include! *Database Design for Mere Mortals(TM), Second Edition*, helps you design your database right in the first place!" --Matt Nunn, Product Manager, SQL Server, Microsoft Corporation "When my brother started his professional career as a developer, I gave him Mike's book to help him understand

database concepts and make real-world application of database technology. When I need a refresher on the finer points of database design, this is the book I pick up. I do not think that there is a better testimony to the value of a book than that it gets used. For this reason I have wholeheartedly recommended to my peers and students that they utilize this book in their day-to-day development tasks." --Chris Kunic-ki, Senior Consultant, OfficeZealot.com "Mike has always had an incredible knack for taking the most complex topics, breaking them down, and explaining them so that anyone can 'get it.' He has honed and polished his first very, very good edition and made it even better. If you're just starting out building database applications, this book is a must-read cover to cover. Expert designers will find Mike's approach fresh and enlightening and a source of great material for training others." --John Viescas, President, Viescas Consulting, Inc., author of *Running Microsoft Access 2000* and coauthor of *SQL Queries for Mere Mortals* "Whether you need to learn about relational database design in general, design a relational

database, understand relational database terminology, or learn best practices for implementing a relational database, *Database Design for Mere Mortals*(T-M), Second Edition, is an indispensable book that you'll refer to often. With his many years of real-world experience designing relational databases, Michael shows you how to analyze and improve existing databases, implement keys, define table relationships and business rules, and create data views, resulting in data integrity, uniform access to data, and reduced data-entry errors." --Paul Cornell, Site Editor, MSDN Office Developer Center Sound database design can save hours of development time and ensure functionality and reliability. *Database Design for Mere Mortals*(TM), Second Edition, is a straightforward, platform-independent tutorial on the basic principles of relational database design. It provides a commonsense design methodology for developing databases that work. Database design expert Michael J. Hernandez has expanded his best-selling first edition, maintaining its hands-on approach and accessibility while updating its coverage and including even more exam-

ples and illustrations. This edition features a CD-ROM that includes diagrams of sample databases, as well as design guidelines, documentation forms, and examples of the database design process. This book will give you the knowledge and tools you need to create efficient and effective relational databases.

Java developers typically go through four "stages" in mastering Java. In the first stage, they learn the language itself. In the second stage, they study the APIs. In the third stage, they become proficient in the environment. It is in the fourth stage -- "the expert stage"-- where things really get interesting, and *Java Enterprise Best Practices* is the tangible compendium of experience that developers need to breeze through this fourth and final stage of Enterprise Java mastery. Crammed with tips and tricks, *Java Enterprise Best Practices* distills years of solid experience from eleven experts in the J2EE environment into a practical, to-the-point guide to J2EE. *Java Enterprise Best Practices* gives developers the unvarnished, expert-tested advice that the man pages don't provide--what areas

of the APIs should be used frequently (and which are better avoided); elegant solutions to problems you face that other developers have already discovered; what things you should always do, what things you should consider doing, and what things you should never do—even if the documentation says it's ok. Until Java Enterprise Best Practices, Java developers in the fourth stage of mastery relied on the advice of a loose-knit community of fellow developers, time-consuming online searches for examples or suggestions for the immediate problem they faced, and tedious trial-and-error. But Java has grown to include a huge number of APIs, classes, and methods. Now it is simply too large for even the most intrepid developer to know it all. The need for a written compendium of J2EE Best Practices has never been greater. Java Enterprise Best Practices focuses on the Java 2 Enterprise Edition (J2EE) APIs. The J2EE APIs include such alphabet soup acronyms as EJB, JDBC, RMI, XML, and JMX. 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!

This book constitutes the thoroughly refereed post-proceedings of the 11th International Symposium on Database Programming Languages, DBPL 2007, held in conjunction with VLDB 2007. The 16 revised full papers presented together with one invited lecture were carefully selected during two rounds of reviewing. The papers are organized in topical sections on algorithms, XML query languages, inconsistency handling, data provenance, emerging data models, and type checking.

Open source has had a profound effect on the Java community. Many Java open source projects have even become de-facto standards. The principal purpose of Enterprise Java Development on a Budget is to guide you through the development of a real enterprise Java application using nothing but open source Java tools, projects, and frameworks. This book is organized by activities and by particular open source projects

that can help you take on the challenges of building the different tiers of your applications. The authors also present a realistic example application that covers most areas of enterprise application development. You'll find information on how to use and configure JBoss, Ant, XDoclet, Struts, ArgoUML, OJB, Hibernate, JUnit, SWT/JFace, and others. Not only will you learn how to use each individual tool, but you'll also understand how to use them in synergy to create robust enterprise Java applications within your budget. Enterprise Java Development on a Budget combines coverage of best practices with information on the right open source Java tools and technologies, all of which will help support your Java development budget and goals.

Describes Agile Modeling Driven Design (AMDD) and Test-Driven Design (TDD) approaches, database refactoring, database encapsulation strategies, and tools that support evolutionary techniques Agile software developers often use object and relational database (RDB) technology together and as a result must overcome the impedance mismatch The author covers techniques

for mapping objects to RDBs and for implementing concurrency control, referential integrity, shared business logic, security access control, reports, and XML An agile foundation describes fundamental skills that all agile software developers require, particularly Agile DBAs Includes object modeling, UML data modeling, data normalization, class normalization, and how to deal with legacy databases Scott W. Ambler is author of Agile Modeling (0471202827), a contributing editor with Software Development (www.sdmagazine.com), and a featured speaker at software conferences worldwide

We increasingly view the world around us as a product of science and technology. Accordingly, we have begun to appreciate that science does not take its problems only from nature and then produces technological applications, but that the very problems of scientific research themselves are generated by science and technology. Simultaneously, problems like global warming, the toxicology of nanoparticles, or the use of renewable energies are constituted by many factors that interact with great complexity. Science in the context of applica-

tion is challenged to gain new understanding and control of such complexity—it cannot seek shelter in the ivory tower or simply pursue its internal quest for understanding and gradual improvement of grand theories. Science in the Context of Application will identify, explore and assess these changes. Part I considers the "Changing Conditions of Scientific Research" and part II "Science, Values, and Society". Examples are drawn from pharmaceutical research, the information sciences, simulation modelling, nanotechnology, cancer research, the effects of commercialization, and many other fields. The book assembles papers from well-known European and American Science Studies scholars like Bernadette Bensaude-Vincent, Janet Kourany, Michael Mahoney, Margaret Morrison, Hans-Jörg Rheinberger, Arie Rip, Dan Sarewitz, Peter Weingart, and others. The individual chapters are written to address anyone who is concerned about the role of contemporary science in society, including scientists, philosophers, and policy makers.

Information Management: Gaining a Competitive Advantage with Data is

about making smart decisions to make the most of company information. Expert author William McKnight develops the value proposition for information in the enterprise and succinctly outlines the numerous forms of data storage. Information Management will enlighten you, challenge your preconceived notions, and help activate information in the enterprise. Get the big picture on managing data so that your team can make smart decisions by understanding how everything from workload allocation to data stores fits together. The practical, hands-on guidance in this book includes: Part 1: The importance of information management and analytics to business, and how data warehouses are used Part 2: The technologies and data that advance an organization, and extend data warehouses and related functionality Part 3: Big Data and NoSQL, and how technologies like Hadoop enable management of new forms of data Part 4: Pulls it all together, while addressing topics of agile development, modern business intelligence, and organizational change management Read the book cover-to-cover, or keep it within reach for a quick and useful resource.

Either way, this book will enable you to master all of the possibilities for data or the broadest view across the enterprise. Balances business and technology, with non-product-specific technical detail Shows how to leverage data to deliver ROI for a business Engaging and approachable, with practical advice on the pros and cons of each domain, so that you learn how information fits together into a complete architecture Provides a path for the data warehouse professional into the new normal of heterogeneity, including NoSQL solutions

Did you ever try getting Businesspeople and IT to agree on the project scope for a new application? Or try getting Marketing and Sales to agree on the target audience? Or try bringing new team members up to speed on the hundreds of tables in your data warehouse — without them dozing off? Whether you are a businessperson or an IT professional, you can be the hero in each of these and hundreds of other scenarios by building a High-Level Data Model. The High-Level Data Model is a simplified view of our complex environment. It can be a powerful communication tool of the key

concepts within our application development projects, business intelligence and master data management programs, and all enterprise and industry initiatives. Learn about the High-Level Data Model and master the techniques for building one, including a comprehensive ten-step approach and hands-on exercises to help you practice topics on your own. In this book, we review data modeling basics and explain why the core concepts stored in a high-level data model can have significant business impact on an organization. We explain the technical notation used for a data model and walk through some simple examples of building a high-level data model. We also describe how data models relate to other key initiatives you may have heard of or may be implementing in your organization. This book contains best practices for implementing a high-level data model, along with some easy-to-use templates and guidelines for a step-by-step approach. Each step will be illustrated using many examples based on actual projects we have worked on. Names have been changed to protect the innocent, but the pain

points and lessons have been preserved. One example spans an entire chapter and will allow you to practice building a high-level data model from beginning to end, and then compare your results to ours. Building a high-level data model following the ten step approach you'll read about is a great way to ensure you will retain the new skills you learn in this book. As is the case in many disciplines, using the right tool for the right job is critical to the overall success of your high-level data model implementation. To help you in your tool selection process, there are several chapters dedicated to discussing what to look for in a high-level data modeling tool and a framework for choosing a data modeling tool, in general. This book concludes with a real-world case study that shows how an international energy company successfully used a high-level data model to streamline their information management practices and increase communication throughout the organization—between both businesspeople and IT. Data modeling is one of the under-exploited, and potentially very valuable, business capabilities that are often hidden away in an

organization's Information Technology department. Data Modeling for the Business highlights both the resulting damage to business value, and the opportunities to make things better. As an easy-to follow and comprehensive guide on the 'why' and 'how' of data modeling, it also reminds us that a successful strategy for exploiting IT depends at least as much on the information as the technology. Chris Potts, Corporate IT Strategist and Author of *frulTion: Creating the Ultimate Corporate Strategy for Information Technology* One of the most critical systems issues is aligning business with IT and fulfilling business needs using data models. The authors of *Data Modeling for the Business* do a masterful job at simply and clearly describing the art of using data models to communicate with business representatives and meet business needs. The book provides many valuable tools, analogies, and step-by-step methods for effective data modeling and is an important contribution in bridging the much needed connection between data modeling and realizing business requirements. Len Silverston, author of *The Data Model Resource Book* series

With the newly introduced 2 Term Examination Pattern, CBSE has eased out the pressure of preparation of subjects and cope up with lengthy syllabus. Introducing Arihant's CBSE TERM II - 2022 Series, the first of its kind that gives complete emphasis on the rationalized syllabus of Class 9th to 12th. The all new "CBSE Term II 2022 - Informatics Practices" of Class 11th provides explanation and guidance to the syllabus required to study efficiently and succeed in the exams. The book provides topical coverage of all the chapters in a complete and comprehensive manner. Covering the 50% of syllabus as per Latest Term wise pattern 2021-22, this book consists of: 1. Complete Theory in each Chapter covering all topics 2. Case-Based, Short and Long Answer Type Question in each chapter 3. Coverage of NCERT, NCERT Exemplar & Board Exams' Questions 4. Complete and Detailed explanations for each question 5. 3 Practice papers based on the entire Term II Syllabus. Table of Content Database Concepts, Introduction to MySQL and SQL, Queries in SQL, Emerging Trends, Practice Papers (1-3).

This guide for Web developers and database programmers shows how to build robust XML applications backed by SQL databases. After an overview of advantages of XML and SQL, stages of application development are detailed step-by-step, illustrated with examples of when and how each technology is most effective. Coverage includes project definition, data modeling, database schema design, and Java programming with XML and SQL. The book is intended for software developers managing small- to medium-scale projects. Appelquist is a technology consultant in content management and e-business strategy. Annotation copyrighted by Book News, Inc., Portland, OR.

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

How can managers increase their ability to calculate price and risk data for financial instruments while decreasing their dependence on a myriad of specific instrument variants? Wolfgang Schwerdt and Marcelle von Wendland created a sim-

ple and consistent way to handle and process large amounts of complex financial data. By means of a practical framework, their approach analyzes market and credit risk exposure of financial instruments and portfolios and calculates risk adjusted performance measures. Its emphasis on standardization yields significant improvements in speed and accuracy. Schwerdt and von Wendland's focus on practical implementation directly addresses limitations imposed by the complex and costly processing time required for advanced risk management models and pricing hundreds of thousands of securities each day. Their many examples and programming codes demonstrate how to use standards to build financial instruments, how to price them, and how to measure the risk and performance of the portfolios that include them. Feature: The authors have designed and implemented a standard for the description of financial instruments Benefit: The reader can rely on accurate and valid information about describing financial instruments Feature: The authors have developed an approach for pricing and analyzing any financial in-

strument using a limited set of atomic instruments Benefit: The reader can use these instruments to define and set up even very large numbers of financial instruments. Feature: The book builds a practical framework for analysing the market and credit risk exposure of financial instruments and portfolios Benefit: Readers can use this framework today in their work and identify and measure market and credit risk using a reliable method.

This book constitutes the thoroughly refereed proceedings of five international workshops held in Ljubljana, Slovenia, in conjunction with the 28th International Conference on Advanced Information Systems Engineering, CAiSE 2016, in June 2016. The 16 full and 9 short papers were carefully selected from 51 submissions. The associated workshops were the Third International Workshop on Advances in Services Design based on the Notion of Capabiliy (ASDENCA) co-arranged with the First International Workshop on Business Model Dynamics and Information Systems Engineering (BumDISE), the Fourth International Workshop on Cognitive Aspects of Information Systems Engineer-

ing (COGNISE), the First International Workshop on Energy-awareness and Big Data Management in Information Systems (EnBIS), the Second International Workshop on Enterprise Modeling (EM), and the Sixth International Workshop on Information Systems Security Engineering (WISSE).

Master the most agile and resilient design for building analytics applications: the Unified Star Schema (USS) approach. The USS has many benefits over traditional dimensional modeling. Witness the power of the USS as a single star schema that serves as a foundation for all present and future business requirements of your organization. Data warehouse legend Bill Inmon and business intelligence innovator, Francesco Pupini, explain step-by-step why the Unified Star Schema is the recommended approach for business intelligence designs today, and show through many examples how to build and use this new solution. This book contains two parts. Part I, Architecture, explains the benefits of data marts and data warehouses, covering how organizations progressed to their current state of analytics, and to the challenges that result from

current business intelligence architectures. Chapter 1 covers the drivers behind and the characteristics of the data warehouse and data mart. Chapter 2 introduces dimensional modeling concepts, including fact tables, dimensions, star joins, and snowflakes. Chapter 3 recalls the evolution of the data mart. Chapter 4 explains Extract, Transform, and Load (ETL), and the value ETL brings to reporting. Chapter 5 explores the Integrated Data Mart Approach, and Chapter 6 explains how to monitor this environment. Chapter 7 describes the different types of metadata within the data warehouse environment. Chapter 8 progresses through the evolution to our current modern data warehouse environment. Part II, the Unified Star Schema, covers the Unified Star Schema (USS) approach and how it solves the challenges introduced in Part I. There are eight chapters within Part II: · Chapter 9, Introduction to the Unified Star Schema: Learn about its architecture and use cases, as well as how the USS approach differs from the traditional approach. · Chapter 10, Loss of Data: Learn about the loss of data and the USS Bridge. Un-

derstand that the USS approach does not create any join, and for this reason, it has no loss of data. · Chapter 11, The Fan Trap: Get introduced to the Oriented Data Model convention, and learn the dangers of a fan trap through an example. Differentiate join and association, and realize that an “in-memory association” is the preferred solution to the fan trap. · Chapter 12, The Chasm Trap: Become familiar with the Cartesian product, and then follow along with an example based on LinkedIn, which illustrates that a chasm trap produces unwanted duplicates. See that the USS Bridge is based on a union, which does not create any duplicates. · Chapter 13, Multi-Fact Queries: Distinguish between multiple facts “with direct connection” versus multiple facts “with no direct connection”. Explore how BI tools are capable of building aggregated virtual rows. · Chapter 14, Loops: Learn more about loops and five traditional techniques to solve them. Follow along with an implementation, which will illustrate the solution based on the USS approach. · Chapter 15, Non-Conformed Granularities: Learn about non-con-

formed granularities, and learn that the Unified Star Schema introduces a solution called “re-normalization”. · Chapter 16, Northwind Case Study. Witness how easy it is to detect the pitfalls of Northwind using the ODM convention. Follow along with an implementation of the USS approach on the Northwind database with various BI tools.

Go beyond the basics with Spring Boot! This practical guide presents dozens of relevant scenarios in a convenient problem-solution-discussion format. Spring Boot in Practice covers dozens of handy Spring Boot development techniques, from basic functions to hidden features you probably didn't even know existed. Each recipe is built around a real-world problem, complete with a full solution and thoughtful discussion. You'll work your way from fundamentals to advanced functionalities as you take deep dives into auto-configuration, security implementation, and even support for reactive application development. Learn how to work with Spring Boot and Kotlin, handling connections for multiple platforms, and how Spring Boot can simplify building microservices and APIs. You're

sure to keep this practical book on hand as a reference as you build your next Spring Boot apps. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications.

Information that is crucial to your case can be stored just about anywhere in Blackberries, on home computers, in cell-phones, in voicemail transcription programs, on flash drives, in native files, in metadata... Knowing what you're looking for is essential, but understanding technology and data storage systems can literally make or break your discovery efforts and your case. If you can't write targeted discovery requests, you won't get all the information you need. With *Electronic Discovery: Law and Practice*, Third Edition, you'll have the first single-source guide to the emerging law of electronic discovery and delivering reliable guidance on such topics as: Duty to Preserve Electronic Evidence Spoliation Document Retention Policies and Electronic Information Cost Shifting in Electronic Discovery Evidentiary Issues Inadvertent Waiver Table of State eDiscovery rules Litigation Hold Notices Application of the

Work Product Doctrine to Litigation Support Systems Collection, Culling and Coding of ESI Inspection of Hard Disks in Civil Litigation Privacy Concerns Disclosure under FOIA Fully grasp the complexities of data sources and IT systems as they relate to electronic discovery, including cutting-edge software tools that facilitate discovery and litigation. Achieve a cooperative and efficient approach to conducting cost-effective ESI discovery. Employ sophisticated and effective discovery tools, including concept and contextual searching, statistical sampling, relationship mapping, and artificial intelligence that help automate the discovery process, reduce costs and enhance process and information integrity Written by Adam Cohen of Ernst & Young and David Lender of Weil, Gotshal & Manges LLP, *Electronic Discovery: Law and Practice*, Third Edition offers detailed analysis and guidance on the legal aspects of electronic discovery never before collected in such a comprehensive guide. You'll save time on research while benefiting from the knowledge and experience of the leading experts. Note: Online subscriptions are for three--

month periods. Previous Edition: Electronic Discovery: Law & Practice, Second Edition, ISBN 9781454815600

This book constitutes the refereed proceedings of the 46th International Conference on Current Trends in Theory and Practice of Informatics, SOFSEM 2020, held in Limassol, Cyprus, in January 2020. The 40 full papers presented together with 17 short papers and 3 invited papers were carefully reviewed and selected from

125 submissions. They presented new research results in the theory and practice of computer science in the each sub-area of SOFSEM 2020: foundations of computer science, foundations of data science and engineering, foundations of software engineering, and foundations of algorithmic computational biology.

This book constitutes the refereed proceedings of the 8th International Conference on Model Transformation, ICMT 2015, held

in L'Aquila, Italy, in July 2015, as Part of STAF 2015, the federation of a number of the leading conferences on software technologies. The 16 revised papers were carefully selected from 34 submissions. The papers are organized in topical sections on change management; reuse and industrial applications; new paradigms for model transformation; transformation validation and verification; and foundations of model transformation.