Type Inheritance And Relational Theory Subtypes Supertypes And Substitutability

Expert MySQL is the leading reference for learning, understanding, and extending the MySQL server. It unlocks the full promise of open source by showing how to modify the code, create your own storage engine, build your own authentication plugins, and even add your own functions and commands to the SQL language. No other book provides the level of detail or the extensive examples of the inner workings of MySQL that have taken engineers years to master. Expert MySQL is a must have book for all systems integrators, engineers, and software developers working with the MySQL server code. Expert MySQL is also a wealth of information on key aspects of MySQL internals. You'll learn about internal query representation, how the optimizer creates execution plans, and how to exert control over those plans for optimal performance in your environment. You'll even learn to build your own query optimizer, giving insight that can help you understand and resolve tough performance problems. High-availability and replication are also covered, making Expert MySQL a must-have book for anyone doing high-end work involving MySQL. Shows how to customize MySQL and its storage and authentication engines Provides in-depth knowledge of internals for use in query tuning and performance troubleshooting Covers high-end features such as high-availability and replication

This book is based on the premise that knowledge of Information Technology (IT) is essential today for people in every walk of life and all types of profession. It is designed to impart a

unified body of knowledge and practice in IT to its readers. Readers can apply this knowledge in innovative ways for various strategic advantages such as increasing productivity, improving quality of products and services, problem solving, decision making, and improving their own and others living standards. The textbook takes a practical approach to introduce the various components of IT to its readers. While doing so, it demonstrates how IT is being used in modern enterprises by various departments to carry out their activities with greater ease, speed, and accuracy than before. It also introduces several new business models and practices made possible due to IT that enterprises are now using for better profitability. In the process, the book provides to its readers a sound foundation of various components and aspects of IT. It also introduces to its readers several latest concepts and technologies in IT such as Wearable computers, Green computing, Cloud computing, Speech recognition and voice response systems, 4G and 5G networks, Big data analytics, Data science, Web 3.0, IPv6, 3D printing, Enterprise 2.0 organization, etc.

A note from the authors: Dear Reader: "Database is boring." That sentiment is heard all too widely these days. But it's so wrong! The database field is full of important problems still to be solved and interesting issues still to be examined - and some of those problems and issues are explored in this book. Between us, we have nearly 80 years experience in this field, and we're still actively researching, exploring, and learning, as well as helping others do the same. The present book is the latest in a series devoted to these goals; using "The Third Manifesto" (a detailed proposal for the future of database technology) as a foundation, it reports on some of our most recent investigations in this field. Among many other things, it includes the most recent version of "The Third Manifesto" itself; specifications for a conforming language called

Tutorial D; and a detailed proposal for a model of type inheritance. Other significant features include: - Extending the foreign key concept - Simplifying queries using image relations - Closer looks at logic and relational algebra - Suggested approaches to "missing information" - Responses to certain "Manifesto" criticisms - Clarifying aspects of normalization The tone of the book overall is naturally somewhat serious, but there are moments of light relief as well. We hope you enjoy it. C.J. Date and Hugh Darwen

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.

A guide to building applications with Rails covers such topics as metaprogamming, Active Support library, advanced database functions, security principles, RESTful architecture, and optimizing performance.

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.

SQL is full of difficulties and traps for the unwary. You can avoid them if you understand relational theory, but only if you know how to put the theory into practice. In this insightful book, author C.J. Date explains relational theory in depth, and demonstrates through numerous examples and exercises how you can apply it directly to your use of SQL. This second edition includes new material on recursive queries, “missing information” without nulls, new update operators, and topics such as aggregate operators, grouping and ungrouping, and view updating. If you have a modest-to-advanced background in SQL, you’ll learn how to deal with a host of common SQL dilemmas. Why is proper column naming so important? Nulls in your database are causing you to get wrong answers. Why? What can you do about it? Is it possible to write an SQL query to find employees who have never been in the same department for more than six months at a time? SQL supports “quantified comparisons,” but they’re better avoided. Why? How do you avoid them? Constraints are crucially important, but most SQL products don’t support them properly. What can you do to resolve this situation? Database theory and practice have evolved since the relational model was developed more than 40 years ago. SQL and Relational Theory draws on decades of research to present the most up-

to-date treatment of SQL available. C.J. Date has a stature that is unique within the database industry. A prolific writer well known for the bestselling textbook An Introduction to Database Systems (Addison-Wesley), he has an exceptionally clear style when writing about complex principles and theory.

Type Inheritance and Relational TheorySubtypes, Supertypes, and Substitutability"O'Reilly Media, Inc."

Time and Relational Theory provides an in-depth description of temporal database systems, which provide special facilities for storing, querying, and updating historical and future data. Traditionally, database management systems provide little or no special support for temporal data at all. This situation is changing because: Cheap storage enables retention of large volumes of historical data in data warehouses Users are now faced with temporal data problems, and need solutions Temporal features have recently been incorporated into the SQL standard, and vendors have begun to add temporal support to their DBMS products Based on the groundbreaking text Temporal Data & the Relational Model (Morgan Kaufmann, 2002) and new research led by the authors, Time and Relational Theory is the only book to offer a complete overview of the functionality of a temporal DBMS. Expert authors Nikos Lorentzos, Hugh Darwen, and Chris Date describe an approach to temporal

database management that is firmly rooted in classical relational theory and will stand the test of time. This book covers the SQL:2011 temporal extensions in depth and identifies and discusses the temporal functionality still missing from SQL. Understand how the relational model provides an ideal basis for taming the complexities of temporal databases Learn how to analyze and evaluate commercial temporal products with this timely and important information Be able to use sound principles in designing and using temporal databases Understand the temporal support recently added to SQL with coverage of the new SQL features in this unique, accurate, and authoritative reference Appreciate the benefits of a truly relational approach to the problem with this clear, user friendly presentation

Create, develop and manage relational databases in real world applications using PostgreSQL About This Book Learn about the PostgreSQL development life cycle including its testing and refactoring Build productive database solutions and use them in Java applications A comprehensive guide to learn about SQL, PostgreSQL procedural language and PL/pgSQL Who This Book Is For If you are a student, database developer or an administrator, interested in developing and maintaining a PostgreSQL database, then this book is for you. No knowledge of database programming or administration is necessary. What You Will Learn

Learn concepts of data modelling and relation algebra Install and set up PostgreSQL database server and client software Implement data structures in PostgreSQL Manipulate data in the database using SQL Implement data processing logic in the database with stored functions, triggers and views Test database solutions and assess the performance Integrate database with Java applications Detailed knowledge of the main PostgreSQL building objects, most used extensions Practice database development life cycle including analysis, modelling, (documentation), testing, bug fixes and refactoring In Detail PostgreSQL is one of the most powerful and easy to use database management systems. It has strong support from the community and is being actively developed with a new release every year. PostgreSQL supports the most advanced features included in SQL standards. Also it provides NoSQL capabilities, and very rich data types and extensions. All that makes PostgreSQL a very attractive solution in various kinds of software systems. The book starts with the introduction of relational databases with PostegreSQL. It then moves on to covering data definition language (DDL) with emphasis on PostgreSQL and common DDL commands supported by ANSI SQL. You will then learn the data manipulation language (DML), and advanced topics like locking and multi version concurrency control (MVCC). This will give you a very robust background to tune

and troubleshoot your application. The book then covers the implementation of data models in the database such as creating tables, setting up integrity constraints, building indexes, defining views and other schema objects. Next, it will give you an overview about the NoSQL capabilities of PostgreSQL along with Hstore, XML, Json and arrays. Finally by the end of the book, you'll learn to use the JDBC driver and manipulate data objects in the Hibernate framework. Style and approach An easy-to-follow guide to learn programming build applications with PostgreSQL, and manage a PostgreSQL database instance.

-- Places object databases into perspective and shows how they fit into the relational continuum. -- Includes important new relational algebra and database programming ideas, and a complete new model for database subtyping and inheritance. -- Includes a detailed review of SQL:1999 (SQL3) and the proposals of the Object Data Management Group (ODMG). Foundation for Future Database Systems: The Third Manifesto offers a comprehensive, insightful proposal for the future of object/relational database management systems. Date and Darwen present a precise, formal definition of an abstract model of data that can be used as a blueprint for designing both databases and database languages -- and as a rock-solid foundation for integrating relational and object technologies. This new Second Edition has been revised extensively, with major extensions to its

inheritance model; new language proposals, and improved discussions of many key concepts. The book goes beyond formal specifications, with a detailed discussion of the rationale for each proposal. It will be essential reading for everyone with a serious interest in database technology.

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. There is an established interest in integrating databases and programming languages. This book on Data Types and Persistence evolved from the proceedings of a workshop held at the Appin in August 1985. The purpose of the Appin workshop was to focus on these two aspects: persistence and data types, and to bring together people from various disciplines who have thought about these problems. Particular topics of"interest include the design of type systems appropriate for database work, the representation of persistent objects such as data types and modules, and the provision of orthogonal persistence and certain aspects of transactions and concurrency. The programme was broken into three sessions: morning, late afternoon and evening to allow the participants to take advantage of two beautiful days in the Scottish Highlands. The financial assistance of the Science and Engineering Research Council, the National Science Foundation and International Computers Ltd. is gratefully acknowledged. We would also like to thank Isabel Graham, Anne Donnelly and Estelle Taylor for their help in organising the workshop. Finally our thanks to Pete Bailey, Ray

Carick and Dave Munro for the immense task they undertook in typesetting the book. The convergence of programming languages and databases to a coherent and consistent whole requires ideas from, and adjustment in, both intellectual camps. The first group of chapters in this book present ideas and adjustments coming from the programming language research community. This community frequently discusses types and uses them as a framework for other discussions.

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

A hands-on beginner's guide to designing relational databases and managing data using Microsoft Access Relational databases represent one of the most enduring and pervasive forms of information technology. Yet most texts covering relational database design assume an extensive, sophisticated computer science background. There are texts on relational database software tools like Microsoft

Access that assume less background, but they focus primarily on details of the user interface, with inadequate coverage of the underlying design issues of how to structure databases. Growing out of Professor Jonathan Eckstein's twenty years' experience teaching courses on management information systems (MIS) at Rutgers Business School, this book fills this gap in the literature by providing a rigorous introduction to relational databases for readers without prior computer science or programming experience. Relational Database Design for Business, with Microsoft Access helps readers to quickly develop a thorough, practical understanding of relational database design. It takes a step-by-step, real-world approach, using application examples from business and finance every step the way. As a result, readers learn to think concretely about database design and how to address issues that commonly arise when developing and manipulating relational databases. By the time they finish the final chapter, students will have the knowledge and skills needed to build relational databases with dozens of tables. They will also be able to build complete Microsoft Access applications around such databases. This text: Takes a hands-on approach using numerous real-world examples drawn from the worlds of business, finance, and more Gets readers up and running, fast, with the skills they need to use and develop relational databases with Microsoft Access Moves swiftly from conceptual

fundamentals to advanced design techniques Leads readers step-by-step through data management and design, relational database theory, multiple tables and the possible relationships between them, Microsoft Access features such as forms and navigation, formulating queries in SQL, and normalization Introductory Relational Database Design for Business, with MicrosoftAccess is the definitive guide for undergraduate and graduate students in business, finance, and data analysis without prior experience in database design. While Microsoft Access is its primary “hands-on” learning vehicle, most of the skills in this text are transferrable to other relational database software such as MySQL.

This book provides of a precise, formal definition of an abstract model of data, to be considered as a blueprint for the design of a DBMS and database language. In particular, it provides a rock-solid foundation for integrating relational and object technologies, a foundation conspicuously lacking in current approaches to such integration. The book is arranged into four parts plus appendixes: I. Preliminaries: Background and overview; objects and relations. II. Formal Specifications: The Manifesto proper; a new relational algebra; and a language called Tutorial D, a concrete realization of the ideas of the Manifesto. III. Informal Discussions and Explanations: A careful point-by-point examination and exposition of the Manifesto, with copious examples in Tutorial D. IV. Subtyping

and Inheritance: A detailed and comprehensive proposal for a model of type inheritance, also with numerous examples. Appendixes: Annotated references and bibliography; comparisons with SQL3 and ODMG; database design considerations; and many other topics.
--Book Jacket.

No matter what DBMS you are using—Oracle, DB2, SQL Server, MySQL, PostgreSQL—misunderstandings can always arise over the precise meanings of terms, misunderstandings that can have a serious effect on the success of your database projects. For example, here are some common database terms: attribute, BCNF, consistency, denormalization, predicate, repeating group, join dependency. Do you know what they all mean? Are you sure? The New Relational Database Dictionary defines all of these terms and many, many more. Carefully reviewed for clarity, accuracy, and completeness, this book is an authoritative and comprehensive resource for database professionals, with over 1700 entries (many with examples) dealing with issues and concepts arising from the relational model of data. DBAs, database designers, DBMS implementers, application developers, and database professors and students can find the information they need on a daily basis, information that isn't readily available anywhere else.

Bitemporal data has always been important. But it was not until 2011 that the ISO released a SQL standard that supported it. Currently, among major DBMS vendors,

Oracle, IBM and Teradata now provide at least some bitemporal functionality in their flagship products. But to use these products effectively, someone in your IT organization needs to know more than how to code bitemporal SQL statements. Perhaps, in your organization, that person is you. To correctly interpret business requests for temporal data, to correctly specify requirements to your IT development staff, and to correctly design bitemporal databases and applications, someone in your enterprise needs a deep understanding of both the theory and the practice of managing bitemporal data. Someone also needs to understand what the future may bring in the way of additional temporal functionality, so their enterprise can plan for it. Perhaps, in your organization, that person is you. This is the book that will show the do-it-yourself IT professional how to design and build bitemporal databases and how to write bitemporal transactions and queries, and will show those who will direct the use of vendor-provided bitemporal DBMSs exactly what is going on "under the covers" of that software. Explains the business value of bitemporal data in terms of the information that can be provided by bitemporal tables and not by any other form of temporal data, including history tables, version tables, snapshot tables, or slowly-changing dimensions. Provides an integrated account of the mathematics, logic, ontology and semantics of relational theory and relational databases, in terms of which current relational theory and practice can be seen as unnecessarily constrained to the management of nontemporal and incompletely temporal data. Explains how bitemporal tables can provide the time-

variance and nonvolatility hitherto lacking in Inmon historical data warehouses. Explains how bitemporal dimensions can replace slowly-changing dimensions in Kimball star schemas, and why they should do so. Describes several extensions to the current theory and practice of bitemporal data, including the use of episodes, "whenever" temporal transactions and queries, and future transaction time. Points out a basic error in the ISO's bitemporal SQL standard, and warns practitioners against the use of that faulty functionality. Recommends six extensions to the ISO standard which will increase the business value of bitemporal data. Points towards a tritemporal future for bitemporal data, in which an Aristotelian ontology and a speech-act semantics support the direct management of the statements inscribed in the rows of relational tables, and add the ability to track the provenance of database content to existing bitemporal databases. This book also provides the background needed to become a business ontologist, and explains why an IT data management person, deeply familiar with corporate databases, is best suited to play that role. Perhaps, in your organization, that person is you.

This series consists of collected volumes and monographs about specific issues dealing with interfaces among the subcomponents of linguistic structure: phonology-morphology, phonology-syntax, syntax-semantics, syntax-morphology, and syntax-lexicon. Recent linguistic research has recognized that the subcomponents of grammar interact in non-trivial ways. What is currently under debate is the actual range of such interactions and their most appropriate representation in grammar, and this is precisely

the focus of this series. Specifically, it provides a general overview of various topics by examining them through the interaction of grammatical components. The books function as a state-of- the-art report of research.

Proceedings of the 28th Annual International Conference on Very Large Data Bases held in Hong Kong, China on August 20-23, 2002. Organized by the VLDB Endowment, VLDB is the premier international conference on database technology.

This book constitutes the refereed proceedings of the 24th Seminar on Current Trends in Theory and Practice of Informatics, SOFSEM'97, held in Milovy, Czech Republic, in November 1997. SOFSEM is special in being a mix of a winter school, an international conference, and an advanced workshop meeting the demand for ongoing education in the area of computer science. The volume presents 22 invited contributions by leading experts together with 24 revised contributed papers selected from 63 submissions. The invited presentations are organized in topical sections on foundations, distributed and parallel computing, software engineering and methodology, and databases and information systems.

Databases, Types, and the Relational Model: The Third Manifesto is a proposal for the future direction of data and database management systems (DBMSs). It provide a precise, formal definition of an abstract model of data, to be considered as a foundation for the design of a DBMS and a database language.

Understanding SQL's underlying theory is the best way to guarantee that your SQL

code is correct and your database schema is robust and maintainable. On the other hand, if you're not well versed in the theory, you can fall into several traps. In SQL and Relational Theory, author C.J. Date demonstrates how you can apply relational theory directly to your use of SQL. With numerous examples and clear explanations of the reasoning behind them, you'll learn how to deal with common SQL dilemmas, such as: Should database access granted be through views instead of base tables? Nulls in your database are causing you to get wrong answers. Why? What can you do about it? Could you write an SQL query to find employees who have never been in the same department for more than six months at a time? SQL supports "quantified comparisons," but they're better avoided. Why? How do you avoid them? Constraints are crucially important, but most SQL products don't support them properly. What can you do to resolve this situation? Database theory and practice have evolved since Edgar Codd originally defined the relational model back in 1969. Independent of any SQL products, SQL and Relational Theory draws on decades of research to present the most up-to-date treatment of the material available anywhere. Anyone with a modest to advanced background in SQL will benefit from the many insights in this book.

This book is a comprehensive presentation of entity-relationship (ER) modeling with regard to an integrated development and modeling of database applications. It comprehensively surveys the achievements of research in this field and deals with the ER model and its extensions. In addition, the book presents techniques for the

translation of the ER model into classical database models and languages, such as relational, hierarchical, and network models and languages, as well as into object-oriented models.

E. F. Codd's relational model of data has been described as one of the three greatest inventions of all time (the other two being agriculture and the scientific method), and his receipt of the 1981 ACM Turing Award, the top award in computer science, for inventing it was thoroughly deserved. The papers in which Codd first described his model were staggering in their originality; they had, and continue to have, a huge impact on just about every aspect of the way we do business in the world today. And yet few people, even in the professional database community, are truly familiar with those papers. This book—a thorough overhaul and rewrite of an earlier book by the same name—is an attempt to remedy this sorry state of affairs. In it, well known author C. J. Date provides a detailed examination of all of Codd's major database publications, explaining the nature of his contribution in depth, and in particular highlighting not only the many things he got right but also some of the things he got wrong. Database theory and practice have evolved considerably since Codd first defined his relational model, back in 1969. This book draws on decades of experience to present the most up to date treatment of the material possible. Anyone with a professional interest in databases can benefit from the insights it contains. The book is product independent.

E. F. Codd's relational model of data has been described as one of the three greatest

inventions of all time (the other two being agriculture and the scientific method), and his receipt of the 1981 ACM Turing Award—the top award in computer science—for inventing it was thoroughly deserved. The papers in which Codd first described his model were staggering in their originality; they had, and continue to have, a huge impact on just about every aspect of the way we do business in the world today. And yet few people, even in the professional database community, are truly familiar with those papers. This book is an attempt to remedy this sorry state of affairs. In it, well known author C. J. Date provides a detailed examination of all of Codd's major technical publications, explaining the nature of his contribution in depth, and in particular highlighting not only the many things he got right but also some of the things he got wrong.

Fifty years of relational. It's hard to believe the relational model has been around now for over half a century! But it has—it was born on August 19th, 1969, when Codd's first database paper was published. And Chris Date has been involved with it for almost the whole of that time, working closely with Codd for many years and publishing the very first, and definitive, book on the subject in 1975. In this book's title essay, Chris offers his own unique perspective (two chapters) on those fifty years. No database professional can afford to miss this one of a kind history. But there's more to this book than just a little personal history. Another unique feature is an extensive and in depth discussion (nine chapters) of a variety of frequently asked questions on relational matters, covering such topics as mathematics and the relational model; relational

algebra; predicates; relation valued attributes; keys and normalization; missing information; and the SQL language. Another part of the book offers detailed responses to critics (four chapters). Finally, the book also contains the text of several recent interviews with Chris Date, covering such matters as RM/V2, XML, NoSQL, The Third Manifesto, and how SQL came to dominate the database landscape.

This book is a revised, upgraded, and hugely improved version of an earlier one called Logic and Databases. Although it's effectively a brand new book, therefore, the following remarks from that earlier book are still relevant here. First, logic and databases are inextricably intertwined. The relational model itself is essentially just elementary logic, tailored to database needs. Now, if you're a database professional, this won't be news to you—but you still might not realize just how much everything we do in the database world is (or should be!) affected by logic. Logic is fundamental, and everywhere. As a database professional, therefore, you owe it to yourself to understand the basics of formal logic, and you ought to be able to explain (and perhaps defend) the connections between formal logic and database technology. And that's what this book is about. What it does is show, through a series of partly independent, partly interrelated essays, just how various crucial aspects of database technology—some of them very familiar, others maybe less so—are solidly grounded in formal logic. Overall, the goal is to help you realize the importance of logic in everything you do, and also, I hope, to help you see that logic can be fun.

Views are virtual tables. That means they should be updatable, just as "real" or base tables are. In fact, view updatability isn't just desirable, it's crucial, for practical reasons as well as theoretical ones. But view updating has always been a controversial topic. Ever since the relational model first appeared, there has been widespread skepticism as to whether (in general) view updating is even possible. In stark contrast to this conventional wisdom, this book shows how views, just like base tables, can always be updated (so long as the updates don't violate any integrity constraints). More generally, it shows how updating always ought to work, regardless of whether the target is a base table or a view. The proposed scheme is 100% consistent with the relational model, but rather different from the way updating works in SQL products today. This book can: Help database products improve in the future Help with a "roll your own" implementation, absent such product improvements Make you aware of the crucial role of predicates and constraints Show you how relational products are really supposed to behave Anyone with a professional interest in the relational model, relational technology, or database systems in general can benefit from this book.
This volume was primarily intended to present selected papers from the workshop on Theory and Applications of Nested Relations and Complex Objects, held in Darmstadt, FRG, from April 6-8, 1987. Other papers were solicited in order to provide a picture of the field as general as possible. Research on nested relations and complex objects originates in the late seventies. The motivation was to obtain data models and systems

which would provide support for so-called complex objects or molecular structures, i.e., for hierarchically organized data, thereby overcoming severe shortcomings of the relational model. This theme of research is now maturing. Systems based on those ideas are beginning to be available. Languages of various natures (algebras, calculi, graphical, logic-oriented) have been designed and a theory is slowly emerging. Finally, new developments in database technology and research are incorporating features of models involving complex objects. A variety of approaches is represented in this volume. The first three papers give overviews of major pioneering implementation efforts. The fourth paper is devoted to the important issue of implementation of storage structures. The next three papers propose excursions in the foundations of nested relations and complex objects. The following six contributions are all devoted to modeling of complex objects. The area of database design is represented by the last four papers.

Some things seem so obvious that they don’t need to be spelled out in detail. Or do they? In computing, at least (and probably in any discipline where accuracy and precision are important), it can be quite dangerous just to assume that some given concept is “obvious,” and indeed universally understood. Serious mistakes can happen that way! The first part of this book discusses features of the database field—equality, assignment, naming—where just such an assumption seems to have been made, and it describes some of the unfortunate mistakes that have occurred as a consequence. It

also explains how and why the features in question aren't quite as obvious as they might seem, and it offers some advice on how to work around the problems caused by assumptions to the contrary. Other parts of the book also deal with database issues where devoting some preliminary effort to spelling out exactly what the issues in question entailed could have led to much better interfaces and much more carefully designed languages. The issues discussed include redundancy and indeterminacy; persistence, encapsulation, and decapsulation; the ACID properties of transactions; and types vs. units of measure. Finally, the book also contains a detailed deconstruction of, and response to, various recent pronouncements from the database literature, all of them having to do with relational technology. Once again, the opinions expressed in those pronouncements might seem "obvious" to some people (to the writers at least, presumably), but the fact remains that they're misleading at best, and in most cases just flat out wrong.

This book sheds light on the principles behind the relational model, which is fundamental to all database-backed applications--and, consequently, most of the work that goes on in the computing world today. Database in Depth: The Relational Model for Practitioners goes beyond the hype and gets to the heart of how relational databases actually work.Ideal for experienced database developers and designers, this concise guide gives you a clear view of the technology--a view that's not influenced by any vendor or product. Featuring an

extensive set of exercises, it will help you: understand why and how the relational model is still directly relevant to modern database technology (and will remain so for the foreseeable future) see why and how the SQL standard is seriously deficient use the best current theoretical knowledge in the design of their databases and database applications make informed decisions in their daily database professional activities Database in Depth will appeal not only to database developers and designers, but also to a diverse field of professionals and academics, including database administrators (DBAs), information modelers, database consultants, and more. Virtually everyone who deals with relational databases should have at least a passing understanding of the fundamentals of working with relational models.Author C.J. Date has been involved with the relational model from its earliest days. An exceptionally clear-thinking writer, Date lays out principle and theory in a manner that is easily understood. Few others can speak as authoritatively the topic of relational databases as Date can. Chris Date, one of the founders of the relational model, has updated and expanded his relational database dictionary to include more than 900 terms. Type inheritance is that phenomenon according to which we can say, for example, that every square is also a rectangle, and so properties that apply to rectangles in general apply to squares in particular. In other words, squares are a

subtype of rectangles, and rectangles are a supertype of squares. Recognizing and acting upon such subtype / supertype relationships provides numerous benefits: Certainly it can help in data modeling, and it can also provide for code reuse in applications. For these reasons, many languages, including the standard database language SQL, have long supported such relationships. However, there doesn't seem to be any consensus in the community at large on a formal, rigorous, and abstract model of inheritance. This book proposes such a model, one that enjoys several advantages over other approaches, not the least of which it is that it's fully compatible with the well known relational model of data. Topics the model covers include: Both single and multiple inheritance Scalar, tuple, and relation inheritance Type lattices and union and intersection types Polymorphism and substitutability Compile time and run time binding All of these topics are described in detail in the book, with numerous illustrative examples, exercises, and answers. The book also discusses several alternative approaches. In particular, it includes a detailed discussion and analysis of inheritance as supported in the SQL standard.

Craft the Right Design Using UML Whether building a relational, object-relational, or object-oriented database, database developers are increasingly relying on an object-oriented design approach as the best way to meet user needs and

performance criteria. This book teaches you how to use the Unified Modeling Language-the official standard of the Object Management Group-to develop and implement the best possible design for your database. Inside, the author leads you step by step through the design process, from requirements analysis to schema generation. You'll learn to express stakeholder needs in UML use cases and actor diagrams, to translate UML entities into database components, and to transform the resulting design into relational, object-relational, and object-oriented schemas for all major DBMS products. Features Teaches you everything you need to know to design, build, and test databases using an OO model. Shows you how to use UML, the accepted standard for database design according to OO principles. Explains how to transform your design into a conceptual schema for relational, object-relational, and object-oriented DBMSs. Offers practical examples of design for Oracle, SQL Server, Sybase, Informix, Object Design, POET, and other database management systems. Focuses heavily on re-using design patterns for maximum productivity and teaches you how to certify completed designs for re-use.