



# DATABASE SYSTEM CONCEPTS

4TH EDITION

THE MCGRAW-HILL COMPANIES  
FOR DONATION ONLY  
NOT FOR RESALE

MCGRAW-HILL EDUCATION  
09-ASA-005

THE MCGRAW-HILL COMPANIES  
FOR DONATION ONLY  
NOT FOR RESALE



MARE DATUM



SILBERSCHATZ · KORTH · SUDARSHAN



# Database System Concepts

FOURTH EDITION

**Abraham Silberschatz**

*Bell Laboratories*

**Henry F. Korth**

*Bell Laboratories*

**S. Sudarshan**

*Indian Institute of Technology, Bombay*



Boston Burr Ridge, IL Dubuque, IA Madison, WI New York San Francisco St. Louis  
Bangkok Bogotá Caracas Kuala Lumpur Lisbon London Madrid Mexico City  
Milan Montreal New Delhi Santiago Seoul Singapore Sydney Taipei Toronto

# McGraw-Hill Higher Education

A Division of The McGraw-Hill Companies

## DATABASE SYSTEM CONCEPTS, FOURTH EDITION

Published by McGraw-Hill, a business unit of The McGraw-Hill Companies, Inc., 1221 Avenue of the Americas, New York, NY 10020. Copyright © 2002, 1999, 1991, 1986 by The McGraw-Hill Companies, Inc. All rights reserved. No part of this publication may be reproduced or distributed in any form or by any means, or stored in a database or retrieval system, without the prior written consent of The McGraw-Hill Companies, Inc., including, but not limited to, in any network or other electronic storage or transmission, or broadcast for distance learning.

Some ancillaries, including electronic and print components, may not be available to customers outside the United States.

This book is printed on acid-free paper.

International 6 7 8 9 0 DOC/DOC 0 9 8 7 6 5 4 3 2  
Domestic 3 4 5 6 7 8 9 0 DOC/DOC 0 9 8 7 6 5 4 3

ISBN 0-07-228363-7

ISBN 0-07-112268-0 (ISE)

General manager: *Thomas E. Casson*

Publisher: *Elizabeth A. Jones*

Senior developmental editor: *Kelley Butcher*

Executive marketing manager: *John Wannemacher*

Project manager: *Jill R. Peter*

Production supervisor: *Kara Kudronowicz*

Coordinator of freelance design: *Rick D. Noel*

Cover designer: *JoAnne Schopler*

Cover illustration: *Paul D. Turnbaugh*

Supplement producer: *Jodi K. Banowetz*

Media technology senior producer: *Phillip Meek*

Compositor: *Abraham Silberschatz/Techsetters*

Typeface: *10/12 Palatino*

Printer: *R. R. Donnelley & Sons Company/Crawfordsville, IN*

### Library of Congress Cataloging-in-Publication Data

Silberschatz, Abraham.

Database system concepts / Abraham Silberschatz, Henry F. Korth, S. Sudarshan. — 4th ed.  
p. cm.

Includes bibliographical references and index.

ISBN 0-07-228363-7

I. Database management. I. Korth, Henry F. II. Sudarshan, S. III. Title.

QA76.9.D3 S5637 2002

005.74—dc21

2001030845

CIP

INTERNATIONAL EDITION ISBN 0-07-112268-0

Copyright © 2002. Exclusive rights by The McGraw-Hill Companies, Inc., for manufacture and export. This book cannot be re-exported from the country to which it is sold by McGraw-Hill. The International Edition is not available in North America.

# Database System Concepts





*In memory of my father Joseph Silberschatz,  
and my grandparents Stepha and Aaron Rosenblum*

*Avi Silberschatz*

*To my wife, Joan  
my children, Abigail and Joseph  
and my parents, Henry and Frances*

*Hank Korth*

*To my wife, Sita  
my son, Madhur  
and my mother, Indira*

*S. Sudarshan*

# Preface

Database management has evolved from a specialized computer application to a central component of a modern computing environment, and, as a result, knowledge about database systems has become an essential part of an education in computer science. In this text, we present the fundamental concepts of database management. These concepts include aspects of database design, database languages, and database-system implementation.

This text is intended for a first course in databases at the junior or senior undergraduate, or first-year graduate, level. In addition to basic material for a first course, the text contains advanced material that can be used for course supplements, or as introductory material for an advanced course.

We assume only a familiarity with basic data structures, computer organization, and a high-level programming language such as Java, C, or Pascal. We present concepts as intuitive descriptions, many of which are based on our running example of a bank enterprise. Important theoretical results are covered, but formal proofs are omitted. The bibliographical notes contain pointers to research papers in which results were first presented and proved, as well as references to material for further reading. In place of proofs, figures and examples are used to suggest why a result is true.

The fundamental concepts and algorithms covered in the book are often based on those used in existing commercial or experimental database systems. Our aim is to present these concepts and algorithms in a general setting that is not tied to one particular database system. Details of particular commercial database systems are discussed in Part 8, “Case Studies.”

In this fourth edition of *Database System Concepts*, we have retained the overall style of the first three editions, while addressing the evolution of database management. Several new chapters have been added to cover new technologies. Every chapter has been edited, and most have been modified extensively. We shall describe the changes in detail shortly.



## Organization

The text is organized in eight major parts, plus three appendices:

- **Overview** (Chapter 1). Chapter 1 provides a general overview of the nature and purpose of database systems. We explain how the concept of a database system has developed, what the common features of database systems are, what a database system does for the user, and how a database system interfaces with operating systems. We also introduce an example database application: a banking enterprise consisting of multiple bank branches. This example is used as a running example throughout the book. This chapter is motivational, historical, and explanatory in nature.
- **Data models** (Chapters 2 and 3). Chapter 2 presents the entity-relationship model. This model provides a high-level view of the issues in database design, and of the problems that we encounter in capturing the semantics of realistic applications within the constraints of a data model. Chapter 3 focuses on the relational data model, covering the relevant relational algebra and relational calculus.
- **Relational databases** (Chapters 4 through 7). Chapter 4 focuses on the most influential of the user-oriented relational languages: SQL. Chapter 5 covers two other relational languages, QBE and Datalog. These two chapters describe data manipulation: queries, updates, insertions, and deletions. Algorithms and design issues are deferred to later chapters. Thus, these chapters are suitable for introductory courses or those individuals who want to learn the basics of database systems, without getting into the details of the internal algorithms and structure.

Chapter 6 presents constraints from the standpoint of database integrity and security; Chapter 7 shows how constraints can be used in the design of a relational database. Referential integrity; mechanisms for integrity maintenance, such as triggers and assertions; and authorization mechanisms are presented in Chapter 6. The theme of this chapter is the protection of the database from accidental and intentional damage.

Chapter 7 introduces the theory of relational database design. The theory of functional dependencies and normalization is covered, with emphasis on the motivation and intuitive understanding of each normal form. The overall process of database design is also described in detail.
- **Object-based databases and XML** (Chapters 8 through 10). Chapter 8 covers object-oriented databases. It introduces the concepts of object-oriented programming, and shows how these concepts form the basis for a data model. No prior knowledge of object-oriented languages is assumed. Chapter 9 covers object-relational databases, and shows how the SQL:1999 standard extends the relational data model to include object-oriented features, such as inheritance, complex types, and object identity.

Chapter 10 covers the XML standard for data representation, which is seeing increasing use in data communication and in the storage of complex data types. The chapter also describes query languages for XML.

- **Data storage and querying** (Chapters 11 through 14). Chapter 11 deals with disk, file, and file-system structure, and with the mapping of relational and object data to a file system. A variety of data-access techniques are presented in Chapter 12, including hashing, B<sup>+</sup>-tree indices, and grid file indices. Chapters 13 and 14 address query-evaluation algorithms, and query optimization based on equivalence-preserving query transformations.

These chapters provide an understanding of the internals of the storage and retrieval components of a database.

- **Transaction management** (Chapters 15 through 17). Chapter 15 focuses on the fundamentals of a transaction-processing system, including transaction atomicity, consistency, isolation, and durability, as well as the notion of serializability.

Chapter 16 focuses on concurrency control and presents several techniques for ensuring serializability, including locking, timestamping, and optimistic (validation) techniques. The chapter also covers deadlock issues. Chapter 17 covers the primary techniques for ensuring correct transaction execution despite system crashes and disk failures. These techniques include logs, shadow pages, checkpoints, and database dumps.

- **Database system architecture** (Chapters 18 through 20). Chapter 18 covers computer-system architecture, and describes the influence of the underlying computer system on the database system. We discuss centralized systems, client-server systems, parallel and distributed architectures, and network types in this chapter. Chapter 19 covers distributed database systems, revisiting the issues of database design, transaction management, and query evaluation and optimization, in the context of distributed databases. The chapter also covers issues of system availability during failures and describes the LDAP directory system.

Chapter 20, on parallel databases explores a variety of parallelization techniques, including I/O parallelism, interquery and intraquery parallelism, and interoperation and intraoperation parallelism. The chapter also describes parallel-system design.

- **Other topics** (Chapters 21 through 24). Chapter 21 covers database application development and administration. Topics include database interfaces, particularly Web interfaces, performance tuning, performance benchmarks, standardization, and database issues in e-commerce. Chapter 22 covers querying techniques, including decision support systems, and information retrieval. Topics covered in the area of decision support include online analytical processing (OLAP) techniques, SQL:1999 support for OLAP, data mining, and data warehousing. The chapter also describes information retrieval techniques for



querying textual data, including hyperlink-based techniques used in Web search engines.

Chapter 23 covers advanced data types and new applications, including temporal data, spatial and geographic data, multimedia data, and issues in the management of mobile and personal databases. Finally, Chapter 24 deals with advanced transaction processing. We discuss transaction-processing monitors, high-performance transaction systems, real-time transaction systems, and transactional workflows.

- **Case studies** (Chapters 25 through 27). In this part we present case studies of three leading commercial database systems, including Oracle, IBM DB2, and Microsoft SQL Server. These chapters outline unique features of each of these products, and describe their internal structure. They provide a wealth of interesting information about the respective products, and help you see how the various implementation techniques described in earlier parts are used in real systems. They also cover several interesting practical aspects in the design of real systems.
- **Online appendices.** Although most new database applications use either the relational model or the object-oriented model, the network and hierarchical data models are still in use. For the benefit of readers who wish to learn about these data models, we provide appendices describing the network and hierarchical data models, in Appendices A and B respectively; the appendices are available only online (<http://www.bell-labs.com/topic/books/db-book>).

Appendix C describes advanced relational database design, including the theory of multivalued dependencies, join dependencies, and the project-join and domain-key normal forms. This appendix is for the benefit of individuals who wish to cover the theory of relational database design in more detail, and instructors who wish to do so in their courses. This appendix, too, is available only online, on the Web page of the book.

## The Fourth Edition

The production of this fourth edition has been guided by the many comments and suggestions we received concerning the earlier editions, by our own observations while teaching at IIT Bombay, and by our analysis of the directions in which database technology is evolving.

Our basic procedure was to rewrite the material in each chapter, bringing the older material up to date, adding discussions on recent developments in database technology, and improving descriptions of topics that students found difficult to understand. Each chapter now has a list of review terms, which can help you review key topics covered in the chapter. We have also added a tools section at the end of most chapters, which provide information on software tools related to the topic of the chapter. We have also added new exercises, and updated references.

We have added a new chapter covering XML, and three case study chapters covering the leading commercial database systems, including Oracle, IBM DB2, and Microsoft SQL Server.

We have organized the chapters into several parts, and reorganized the contents of several chapters. For the benefit of those readers familiar with the third edition, we explain the main changes here:

- **Entity-relationship model.** We have improved our coverage of the entity-relationship (E-R) model. More examples have been added, and some changed, to give better intuition to the reader. A summary of alternative E-R notations has been added, along with a new section on UML.
- **Relational databases.** Our coverage of SQL in Chapter 4 now references the SQL:1999 standard, which was approved after publication of the third edition. SQL coverage has been significantly expanded to include the **with** clause, expanded coverage of embedded SQL, and coverage of ODBC and JDBC whose usage has increased greatly in the past few years. Coverage of Quel has been dropped from Chapter 5, since it is no longer in wide use. Coverage of QBE has been revised to remove some ambiguities and to add coverage of the QBE version used in the Microsoft Access database.

Chapter 6 now covers integrity constraints and security. Coverage of security has been moved to Chapter 6 from its third-edition position of Chapter 19. Chapter 6 also covers triggers. Chapter 7 covers relational-database design and normal forms. Discussion of functional dependencies has been moved into Chapter 7 from its third-edition position of Chapter 6. Chapter 7 has been significantly rewritten, providing several short-cut algorithms for dealing with functional dependencies and extended coverage of the overall database design process. Axioms for multivalued dependency inference, PJNF and DKNF, have been moved into an appendix.

- **Object-based databases.** Coverage of object orientation in Chapter 8 has been improved, and the discussion of ODMG updated. Object-relational coverage in Chapter 9 has been updated, and in particular the SQL:1999 standard replaces the extended SQL used in the third edition.
- **XML.** Chapter 10, covering XML, is a new chapter in the fourth edition.
- **Storage, indexing, and query processing.** Coverage of storage and file structures, in Chapter 11, has been updated; this chapter was Chapter 10 in the third edition. Many characteristics of disk drives and other storage mechanisms have changed greatly in the past few years, and our coverage has been correspondingly updated. Coverage of RAID has been updated to reflect technology trends. Coverage of data dictionaries (catalogs) has been extended.

Chapter 12, on indexing, now includes coverage of bitmap indices; this chapter was Chapter 11 in the third edition. The B<sup>+</sup>-tree insertion algorithm has been simplified, and pseudocode has been provided for search. Partitioned hashing has been dropped, since it is not in significant use.

Our treatment of query processing has been reorganized, with the earlier chapter (Chapter 12 in the third edition) split into two chapters, one on query processing (Chapter 13) and another on query optimization (Chapter 14). All details regarding cost estimation and query optimization have been moved



to Chapter 14, allowing Chapter 13 to concentrate on query processing algorithms. We have dropped several detailed (and tedious) formulae for calculating the exact number of I/O operations for different operations. Chapter 14 now has pseudocode for optimization algorithms, and new sections on optimization of nested subqueries and on materialized views.

- **Transaction processing.** Chapter 15, which provides an introduction to transactions, has been updated; this chapter was numbered Chapter 13 in the third edition. Tests for view serializability have been dropped.

Chapter 16, on concurrency control, includes a new section on implementation of lock managers, and a section on weak levels of consistency, which was in Chapter 20 of the third edition. Concurrency control of index structures has been expanded, providing details of the crabbing protocol, which is a simpler alternative to the B-link protocol, and next-key locking to avoid the phantom problem. Chapter 17, on recovery, now includes coverage of the ARIES recovery algorithm. This chapter also covers remote backup systems for providing high availability despite failures, an increasingly important feature in “ $24 \times 7$ ” applications.

As in the third edition, instructors can choose between just introducing transaction-processing concepts (by covering only Chapter 15), or offering detailed coverage (based on Chapters 15 through 17).

- **Database system architectures.** Chapter 18, which provides an overview of database system architectures, has been updated to cover current technology; this was Chapter 16 in the third edition. The order of the parallel database chapter and the distributed database chapters has been flipped. While the coverage of parallel database query processing techniques in Chapter 20 (which was Chapter 16 in the third edition) is mainly of interest to those who wish to learn about database internals, distributed databases, now covered in Chapter 19, is a topic that is more fundamental; it is one that anyone dealing with databases should be familiar with.

Chapter 19 on distributed databases has been significantly rewritten, to reduce the emphasis on naming and transparency and to increase coverage of operation during failures, including concurrency control techniques to provide high availability. Coverage of three-phase commit protocol has been abbreviated, as has distributed detection of global deadlocks, since neither is used much in practice. Coverage of query processing issues in heterogeneous databases has been moved up from Chapter 20 of the third edition. There is a new section on directory systems, in particular LDAP, since these are quite widely used as a mechanism for making information available in a distributed setting.

- **Other topics.** Although we have modified and updated the entire text, we concentrated our presentation of material pertaining to ongoing database research and new database applications in four new chapters, from Chapter 21 to Chapter 24.

Chapter 21 is new in the fourth edition and covers application development and administration. The description of how to build Web interfaces to databases, including servlets and other mechanisms for server-side scripting, is new. The section on performance tuning, which was earlier in Chapter 19, has new material on the famous 5-minute rule and the 1-minute rule, as well as some new examples. Coverage of materialized view selection is also new. Coverage of benchmarks and standards has been updated. There is a new section on e-commerce, focusing on database issues in e-commerce, and a new section on dealing with legacy systems.

Chapter 22, which covers advanced querying and information retrieval, includes new material on OLAP, particularly on SQL:1999 extensions for data analysis. Coverage of data warehousing and data mining has also been extended greatly. Coverage of information retrieval has been significantly extended, particularly in the area of Web searching. Earlier versions of this material were in Chapter 21 of the third edition.

Chapter 23, which covers advanced data types and new applications, has material on temporal data, spatial data, multimedia data, and mobile databases. This material is an updated version of material that was in Chapter 21 of the third edition. Chapter 24, which covers advanced transaction processing, contains updated versions of sections on TP monitors, workflow systems, main-memory and real-time databases, long-duration transactions, and transaction management in multidatabases, which appeared in Chapter 20 of the third edition.

- **Case studies.** The case studies covering Oracle, IBM DB2 and Microsoft SQL Server are new to the fourth edition. These chapters outline unique features of each of these products, and describe their internal structure.

## Instructor's Note

The book contains both basic and advanced material, which might not be covered in a single semester. We have marked several sections as advanced, using the symbol “\*\*”. These sections may be omitted if so desired, without a loss of continuity.

It is possible to design courses by using various subsets of the chapters. We outline some of the possibilities here:

- Chapter 5 can be omitted if students will not be using QBE or Datalog as part of the course.
- If object orientation is to be covered in a separate advanced course, Chapters 8 and 9, and Section 11.9, can be omitted. Alternatively, they could constitute the foundation of an advanced course in object databases.
- Chapter 10 (XML) and Chapter 14 (query optimization) can be omitted from an introductory course.
- Both our coverage of transaction processing (Chapters 15 through 17) and our coverage of database-system architecture (Chapters 18 through 20) consist of

an overview chapter (Chapters 15 and 18, respectively), followed by chapters with details. You might choose to use Chapters 15 and 18, while omitting Chapters 16, 17, 19, and 20, if you defer these latter chapters to an advanced course.

- Chapters 21 through 24 are suitable for an advanced course or for self-study by students, although Section 21.1 may be covered in a first database course.

Model course syllabi, based on the text, can be found on the Web home page of the book (see the following section).

## Web Page and Teaching Supplements

A Web home page for the book is available at the URL:

<http://www.bell-labs.com/topic/books/db-book>

The Web page contains:

- Slides covering all the chapters of the book
- Answers to selected exercises
- The three appendices
- An up-to-date errata list
- Supplementary material contributed by users of the book

A complete solution manual will be made available only to faculty. For more information about how to get a copy of the solution manual, please send electronic mail to [customer.service@mcgraw-hill.com](mailto:customer.service@mcgraw-hill.com). In the United States, you may call 800-338-3987. The McGraw-Hill Web page for this book is

<http://www.mhhe.com/silberschatz>

## Contacting Us and Other Users

We provide a mailing list through which users of our book can communicate among themselves and with us. If you wish to be on the list, please send a message to [db-book@research.bell-labs.com](mailto:db-book@research.bell-labs.com), include your name, affiliation, title, and electronic mail address.

We have endeavored to eliminate typos, bugs, and the like from the text. But, as in new releases of software, bugs probably remain; an up-to-date errata list is accessible from the book's home page. We would appreciate it if you would notify us of any errors or omissions in the book that are not on the current list of errata.

We would be glad to receive suggestions on improvements to the books. We also welcome any contributions to the book Web page that could be of use to other read-

ers, such as programming exercises, project suggestions, online labs and tutorials, and teaching tips.

E-mail should be addressed to [db-book@research.bell-labs.com](mailto:db-book@research.bell-labs.com). Any other correspondence should be sent to Avi Silberschatz, Bell Laboratories, Room 2T-310, 600 Mountain Avenue, Murray Hill, NJ 07974, USA.

## Acknowledgments

This edition has benefited from the many useful comments provided to us by the numerous students who have used the third edition. In addition, many people have written or spoken to us about the book, and have offered suggestions and comments. Although we cannot mention all these people here, we especially thank the following:

- Phil Bernhard, Florida Institute of Technology; Eitan M. Gurari, The Ohio State University; Irwin Levinstein, Old Dominion University; Ling Liu, Georgia Institute of Technology; Ami Motro, George Mason University; Bhagirath Narahari, Meral Ozsoyoglu, Case Western Reserve University; and Odinaldo Rodriguez, King's College London; who served as reviewers of the book and whose comments helped us greatly in formulating this fourth edition.
- Soumen Chakrabarti, Sharad Mehrotra, Krithi Ramamritham, Mike Reiter, Sunita Sarawagi, N. L. Sarda, and Dilys Thomas, for extensive and invaluable feedback on several chapters of the book.
- Phil Bohannon, for writing the first draft of Chapter 10 describing XML.
- Hakan Jakobsson (Oracle), Sriram Padmanabhan (IBM), and César Galindo-Legaria, Goetz Graefe, José A. Blakeley, Kalen Delaney, Michael Rys, Michael Zwillling, Sameet Agarwal, Thomas Casey (all of Microsoft) for writing the appendices describing the Oracle, IBM DB2, and Microsoft SQL Server database systems.
- Yuri Breitbart, for help with the distributed database chapter; Mike Reiter, for help with the security sections; and Jim Melton, for clarifications on SQL:1999.
- Marilyn Turnamian and Nandprasad Joshi, whose excellent secretarial assistance was essential for timely completion of this fourth edition.

The publisher was Betsy Jones. The senior developmental editor was Kelley Butcher. The project manager was Jill Peter. The executive marketing manager was John Wannemacher. The cover illustrator was Paul Tumbaugh while the cover designer was JoAnne Schopler. The freelance copyeditor was George Watson. The freelance proofreader was Marie Zartman. The supplement producer was Jodi Banowetz. The designer was Rick Noel. The freelance indexer was Tobiah Waldron.

This edition is based on the three previous editions, so we thank once again the many people who helped us with the first three editions, including R. B. Abhyankar, Don Batory, Haran Boral, Paul Bourgeois, Robert Brazile, Michael Carey, J. Edwards, Christos Faloutsos, Homma Farian, Alan Fekete, Shashi Gadia, Jim Gray, Le Gruen-



wald, Ron Hitchens, Yannis Ioannidis, Hyoung-Joo Kim, Won Kim, Henry Korth (father of Henry F.), Carol Kroll, Gary Lindstrom, Dave Maier, Keith Marzullo, Fletcher Mattox, Alberto Mendelzon, Hector Garcia-Molina, Ami Motro, Anil Nigam, Cyril Orji, Bruce Porter, Jim Peterson, K. V. Raghavan, Mark Roth, Marek Rusinkiewicz, S. Seshadri, Shashi Shekhar, Amit Sheth, Nandit Soparkar, Greg Speegle, and Marianne Winslett. Lyn Dupré copyedited the third edition and Sara Strandtman edited the text of the third edition. Greg Speegle, Dawn Bezviner, and K. V. Raghavan helped us to prepare the instructor's manual for earlier editions. The new cover is an evolution of the covers of the first three editions; Marilyn Turnamian created an early draft of the cover design for this edition. The idea of using ships as part of the cover concept was originally suggested to us by Bruce Stephan.

Finally, Sudarshan would like to acknowledge his wife, Sita, for her love and support, two-year old son Madhur for his love, and mother, Indira, for her support. Hank would like to acknowledge his wife, Joan, and his children, Abby and Joe, for their love and understanding. Avi would like to acknowledge his wife Haya, and his son, Aaron, for their patience and support during the revision of this book.

A. S.  
H. F. K.  
S. S.

# Contents

## Preface xv

## Chapter 1 Introduction

- |  |    |                                  |    |
|--|----|----------------------------------|----|
| 1.1 Database System Applications         | 1  | 1.7 Transaction Management       | 15 |
| 1.2 Database Systems versus File Systems | 3  | 1.8 Database System Structure    | 16 |
| 1.3 View of Data                         | 5  | 1.9 Application Architectures    | 18 |
| 1.4 Data Models                          | 7  | 1.10 History of Database Systems | 18 |
| 1.5 Database Languages                   | 11 | 1.11 Summary                     | 21 |
| 1.6 Database Users and Administrators    | 13 | Exercises                        | 23 |
|  |    | Bibliographical Notes            | 24 |

## PART 1 ■ DATA MODELS

## Chapter 2 Entity-Relationship Model

- |                                      |    |  |    |
|--------------------------------------|----|--|----|
| 2.1 Basic Concepts                   | 27 | 2.9 Reduction of an E-R Schema to Tables | 62 |
| 2.2 Constraints                      | 33 | 2.10 The Unified Modeling Language UML   | 68 |
| 2.3 Keys                             | 35 | 2.11 Summary                             | 70 |
| 2.4 Design Issues                    | 37 | Exercises                                | 72 |
| 2.5 Entity-Relationship Diagram      | 42 | Bibliographical Notes                    | 77 |
| 2.6 Weak Entity Sets                 | 47 |  |    |
| 2.7 Extended E-R Features            | 49 |  |    |
| 2.8 Design of an E-R Database Schema | 56 |  |    |

## Chapter 3 Relational Model

- 3.1 Structure of Relational Databases 79
- 3.2 The Relational Algebra 89
- 3.3 Extended Relational-Algebra Operations 103
- 3.4 Modification of the Database 111
- 3.5 Views 113
- 3.6 The Tuple Relational Calculus 118
- 3.7 The Domain Relational Calculus 122
- 3.8 Summary 126  
Exercises 127  
Bibliographical Notes 131

## PART 2 ■ RELATIONAL DATABASES

### Chapter 4 SQL

- 4.1 Background 135
- 4.2 Basic Structure 137
- 4.3 Set Operations 144
- 4.4 Aggregate Functions 146
- 4.5 Null Values 148
- 4.6 Nested Subqueries 149
- 4.7 Views 154
- 4.8 Complex Queries 155
- 4.9 Modification of the Database 157
- 4.10 Joined Relations 163
- 4.11 Data-Definition Language 168
- 4.12 Embedded SQL 172
- 4.13 Dynamic SQL 175
- 4.14 Other SQL Features 180
- 4.15 Summary 182  
Exercises 183  
Bibliographical Notes 186

### Chapter 5 Other Relational Languages

- 5.1 Query-by-Example 189
- 5.2 Datalog 203
- 5.3 User Interfaces and Tools 217
- 5.4 Summary 219  
Exercises 220  
Bibliographical Notes 223

### Chapter 6 Integrity and Security

- 6.1 Domain Constraints 225
- 6.2 Referential Integrity 227
- 6.3 Assertions 232
- 6.4 Triggers 233
- 6.5 Security and Authorization 238
- 6.6 Authorization in SQL 244
- 6.7 Encryption and Authentication 248
- 6.8 Summary 250  
Exercises 252  
Bibliographical Notes 254