


GLOBAL  
EDITION



# Fundamentals of Database Systems

SEVENTH EDITION

Ramez Elmasri • Shamkant B. Navathe

ALWAYS LEARNING

PEARSON



FUNDAMENTALS OF

# Database Systems

SEVENTH EDITION

GLOBAL EDITION



**Ramez Elmasri**

*Department of Computer Science and Engineering  
The University of Texas at Arlington*

**Shamkant B. Navathe**

*College of Computing  
Georgia Institute of Technology*

**PEARSON**

Boston Columbus Indianapolis New York San Francisco Hoboken  
Amsterdam Cape Town Dubai London Madrid Milan Munich Paris Montreal Toronto  
Delhi Mexico City São Paulo Sydney Hong Kong Seoul Singapore Taipei Tokyo



# Preface

This book introduces the fundamental concepts necessary for designing, using, and implementing database systems and database applications. Our presentation stresses the fundamentals of database modeling and design, the languages and models provided by the database management systems, and database system implementation techniques. The book is meant to be used as a textbook for a one- or two-semester course in database systems at the junior, senior, or graduate level, and as a reference book. Our goal is to provide an in-depth and up-to-date presentation of the most important aspects of database systems and applications, and related technologies. We assume that readers are familiar with elementary programming and data-structuring concepts and that they have had some exposure to the basics of computer organization.

## New to This Edition

The following key features have been added in the seventh edition:

- A reorganization of the chapter ordering (this was based on a survey of the instructors who use the textbook); however, the book is still organized so that the individual instructor can choose to follow the new chapter ordering or *choose a different ordering of chapters* (for example, follow the chapter order from the sixth edition) when presenting the materials.
- There are two new chapters on recent advances in database systems and big data processing; one new chapter (Chapter 24) covers an introduction to the newer class of database systems known as **NOSQL databases**, and the other new chapter (Chapter 25) covers technologies for processing **big data**, including **MapReduce** and **Hadoop**.
- The chapter on query processing and optimization has been expanded and reorganized into two chapters; Chapter 18 focuses on strategies and algorithms for query processing whereas Chapter 19 focuses on query optimization techniques.
- A second UNIVERSITY database example has been added to the early chapters (Chapters 3 through 8) in addition to our COMPANY database example from the previous editions.
- Many of the individual chapters have been updated to varying degrees to include newer techniques and methods; rather than discuss these enhancements here,



we will describe them later in the preface when we discuss the organization of the seventh edition.

The following are key features of the book:

- A self-contained, flexible organization that can be tailored to individual needs; in particular, *the chapters can be used in different orders* depending on the instructor's preference.
- A companion website ([www.pearsonglobaleditions.com/Elmasri](http://www.pearsonglobaleditions.com/Elmasri)) includes data to be loaded into various types of relational databases for more realistic student laboratory exercises.
- A dependency chart (shown later in this preface) to show which chapters depend on other earlier chapters; this can guide the instructor who wants to tailor the *order of presentation of the chapters*.
- A collection of supplements, including a robust set of materials for instructors and students such as PowerPoint slides, figures from the text, and an instructor's guide with solutions.

## Organization and Contents of the Seventh Edition

There are some organizational changes in the seventh edition as well as improvement to the individual chapters. The book is now divided into 12 parts as follows:

- Part 1 (Chapters 1 and 2) describes the basic introductory concepts necessary for a good understanding of database models, systems, and languages. Chapters 1 and 2 introduce databases, typical users, and DBMS concepts, terminology, and architecture, as well as a discussion of the progression of database technologies over time and a brief history of data models. These chapters have been updated to introduce some of the newer technologies such as NOSQL systems.
- Part 2 (Chapters 3 and 4) includes the presentation on entity-relationship modeling and database design; however, it is *important to note* that instructors can cover the relational model chapters (Chapters 5 through 8) *before Chapters 3 and 4* if that is their preferred order of presenting the course materials. In Chapter 3, the concepts of the Entity-Relationship (ER) model and ER diagrams are presented and used to illustrate conceptual database design. Chapter 4 shows how the basic ER model can be extended to incorporate additional modeling concepts such as subclasses, specialization, generalization, union types (categories) and inheritance, leading to the enhanced-ER (EER) data model and EER diagrams. The notation for the class diagrams of UML are also introduced in Chapters 7 and 8 as an alternative model and diagrammatic notation for ER/EER diagrams.
- Part 3 (Chapters 5 through 8) includes a detailed presentation on relational databases and SQL with some additional new material in the SQL chapters to cover a few SQL constructs that were not in the previous edition. Chapter 5



describes the basic relational model, its integrity constraints, and update operations. Chapter 6 describes some of the basic parts of the SQL standard for relational databases, including data definition, data modification operations, and simple SQL queries. Chapter 7 presents more complex SQL queries, as well as the SQL concepts of triggers, assertions, views, and schema modification. Chapter 8 describes the formal operations of the relational algebra and introduces the relational calculus. The material on SQL (Chapters 6 and 7) is presented before our presentation on relational algebra and calculus in Chapter 8 to allow instructors to start SQL projects early in a course if they wish (it is possible to cover Chapter 8 before Chapters 6 and 7 if the instructor desires this order). The final chapter in Part 2, Chapter 9, covers ER- and EER-to-relational mapping, which are algorithms that can be used for designing a relational database schema from a conceptual ER/EER schema design.

- Part 4 (Chapters 10 and 11) are the chapters on database programming techniques; these chapters can be assigned as reading materials and augmented with materials on the particular language used in the course for programming projects (much of this documentation is readily available on the Web). Chapter 10 covers traditional SQL programming topics, such as embedded SQL, dynamic SQL, ODBC, SQLJ, JDBC, and SQL/CLI. Chapter 11 introduces Web database programming, using the PHP scripting language in our examples, and includes new material that discusses Java technologies for Web database programming.
- Part 5 (Chapters 12 and 13) covers the updated material on object-relational and object-oriented databases (Chapter 12) and XML (Chapter 13); both of these chapters now include a presentation of how the SQL standard incorporates object concepts and XML concepts into more recent versions of the SQL standard. Chapter 12 first introduces the concepts for object databases, and then shows how they have been incorporated into the SQL standard in order to add object capabilities to relational database systems. It then covers the ODMG object model standard, and its object definition and query languages. Chapter 13 covers the XML (eXtensible Markup Language) model and languages, and discusses how XML is related to database systems. It presents XML concepts and languages, and compares the XML model to traditional database models. We also show how data can be converted between the XML and relational representations, and the SQL commands for extracting XML documents from relational tables.
- Part 6 (Chapters 14 and 15) are the normalization and relational design theory chapters (we moved all the formal aspects of normalization algorithms to Chapter 15). Chapter 14 defines functional dependencies, and the normal forms that are based on functional dependencies. Chapter 14 also develops a step-by-step intuitive normalization approach, and includes the definitions of multivalued dependencies and join dependencies. Chapter 15 covers normalization theory, and the formalisms, theories,



and algorithms developed for relational database design by normalization, including the relational decomposition algorithms and the relational synthesis algorithms.

- Part 7 (Chapters 16 and 17) contains the chapters on file organizations on disk (Chapter 16) and indexing of database files (Chapter 17). Chapter 16 describes primary methods of organizing files of records on disk, including ordered (sorted), unordered (heap), and hashed files; both static and dynamic hashing techniques for disk files are covered. Chapter 16 has been updated to include materials on buffer management strategies for DBMSs as well as an overview of new storage devices and standards for files and modern storage architectures. Chapter 17 describes indexing techniques for files, including B-tree and B<sup>+</sup>-tree data structures and grid files, and has been updated with new examples and an enhanced discussion on indexing, including how to choose appropriate indexes and index creation during physical design.
- Part 8 (Chapters 18 and 19) includes the chapters on query processing algorithms (Chapter 18) and optimization techniques (Chapter 19); these two chapters have been updated and reorganized from the single chapter that covered both topics in the previous editions and include some of the newer techniques that are used in commercial DBMSs. Chapter 18 presents algorithms for searching for records on disk files, and for joining records from two files (tables), as well as for other relational operations. Chapter 18 contains new material, including a discussion of the semi-join and anti-join operations with examples of how they are used in query processing, as well as a discussion of techniques for selectivity estimation. Chapter 19 covers techniques for query optimization using cost estimation and heuristic rules; it includes new material on nested subquery optimization, use of histograms, physical optimization, and join ordering methods and optimization of typical queries in data warehouses.
- Part 9 (Chapters 20, 21, and 22) covers transaction processing concepts; concurrency control; and database recovery from failures. These chapters have been updated to include some of the newer techniques that are used in some commercial and open source DBMSs. Chapter 20 introduces the techniques needed for transaction processing systems, and defines the concepts of recoverability and serializability of schedules; it has a new section on buffer replacement policies for DBMSs and a new discussion on the concept of snapshot isolation. Chapter 21 gives an overview of the various types of concurrency control protocols, with a focus on two-phase locking. We also discuss timestamp ordering and optimistic concurrency control techniques, as well as multiple-granularity locking. Chapter 21 includes a new presentation of concurrency control methods that are based on the snapshot isolation concept. Finally, Chapter 23 focuses on database recovery protocols, and gives an overview of the concepts and techniques that are used in recovery.



- Part 10 (Chapters 23, 24, and 25) includes the chapter on distributed databases (Chapter 23), plus the two new chapters on NOSQL storage systems for big data (Chapter 24) and big data technologies based on Hadoop and MapReduce (Chapter 25). Chapter 23 introduces distributed database concepts, including availability and scalability, replication and fragmentation of data, maintaining data consistency among replicas, and many other concepts and techniques. In Chapter 24, NOSQL systems are categorized into four general categories with an example system in each category used for our examples, and the data models, operations, as well as the replication/distribution/scalability strategies of each type of NOSQL system are discussed and compared. In Chapter 25, the MapReduce programming model for distributed processing of big data is introduced, and then we have presentations of the Hadoop system and HDFS (Hadoop Distributed File System), as well as the Pig and Hive high-level interfaces, and the YARN architecture.
- Part 11 (Chapters 26 through 29) is entitled Advanced Database Models, Systems, and Applications and includes the following materials: Chapter 26 introduces several advanced data models including active databases/triggers (Section 26.1), temporal databases (Section 26.2), spatial databases (Section 26.3), multimedia databases (Section 26.4), and deductive databases (Section 26.5). Chapter 27 discusses information retrieval (IR) and Web search, and includes topics such as IR and keyword-based search, comparing DB with IR, retrieval models, search evaluation, and ranking algorithms. Chapter 28 is an introduction to data mining including overviews of various data mining methods such as associate rule mining, clustering, classification, and sequential pattern discovery. Chapter 29 is an overview of data warehousing including topics such as data warehousing models and operations, and the process of building a data warehouse.
- Part 12 (Chapter 30) includes one chapter on database security, which includes a discussion of SQL commands for discretionary access control (GRANT, REVOKE), as well as mandatory security levels and models for including mandatory access control in relational databases, and a discussion of threats such as SQL injection attacks, as well as other techniques and methods related to data security and privacy.

Appendix A gives a number of alternative diagrammatic notations for displaying a conceptual ER or EER schema. These may be substituted for the notation we use, if the instructor prefers. Appendix B gives some important physical parameters of disks. Appendix C gives an overview of the QBE graphical query language, and Appendixes D and E (available on the book's Companion Website located at [www.pearsonglobaleditions.com/Elmasri](http://www.pearsonglobaleditions.com/Elmasri)) cover legacy database systems, based on the hierarchical and network database models. They have been used for more than thirty years as a basis for many commercial database applications and transaction-processing systems.



## Guidelines for Using This Book

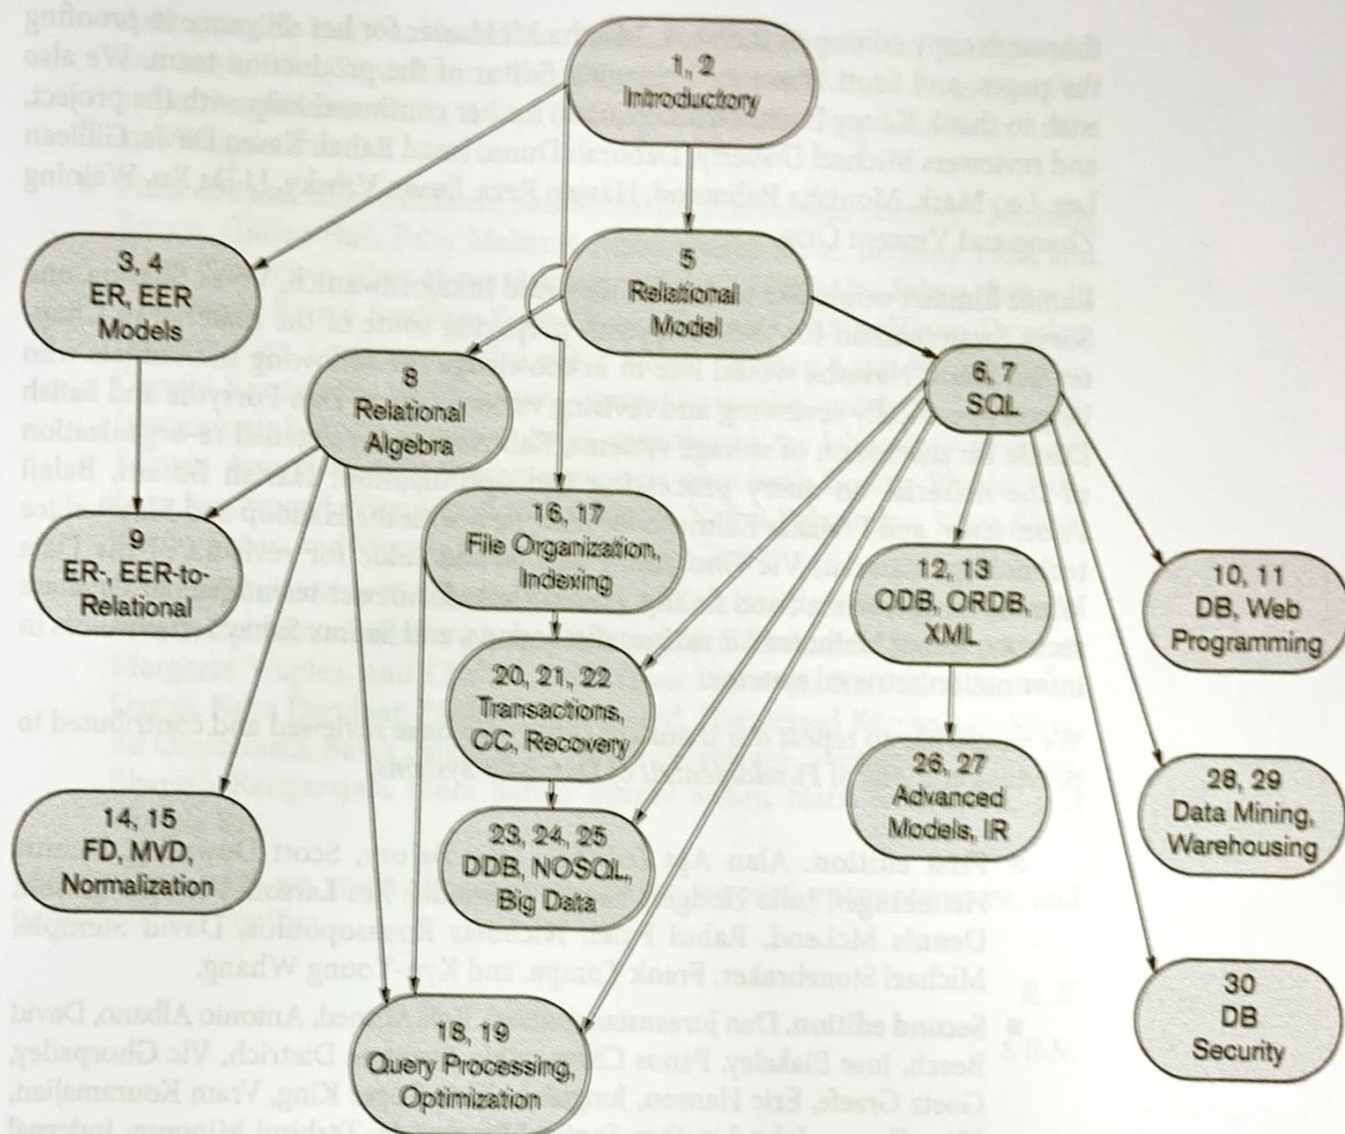
There are many different ways to teach a database course. The chapters in Parts 1 through 7 can be used in an introductory course on database systems in the order that they are given or in the preferred order of individual instructors. Selected chapters and sections may be left out and the instructor can add other chapters from the rest of the book, depending on the emphasis of the course. At the end of the opening section of some of the book's chapters, we list sections that are candidates for being left out whenever a less-detailed discussion of the topic is desired. We suggest covering up to Chapter 15 in an introductory database course and including selected parts of other chapters, depending on the background of the students and the desired coverage. For an emphasis on system implementation techniques, chapters from Parts 7, 8, and 9 should replace some of the earlier chapters.

Chapters 3 and 4, which cover conceptual modeling using the ER and EER models, are important for a good conceptual understanding of databases. However, they may be partially covered, covered later in a course, or even left out if the emphasis is on DBMS implementation. Chapters 16 and 17 on file organizations and indexing may also be covered early, later, or even left out if the emphasis is on database models and languages. For students who have completed a course on file organization, parts of these chapters can be assigned as reading material or some exercises can be assigned as a review for these concepts.

If the emphasis of a course is on database design, then the instructor should cover Chapters 3 and 4 early on, followed by the presentation of relational databases. A total life-cycle database design and implementation project would cover conceptual design (Chapters 3 and 4), relational databases (Chapters 5, 6, and 7), data model mapping (Chapter 9), normalization (Chapter 14), and application programs implementation with SQL (Chapter 10). Chapter 11 also should be covered if the emphasis is on Web database programming and applications. Additional documentation on the specific programming languages and RDBMS used would be required. The book is written so that it is possible to cover topics in various sequences. The following chapter dependency chart shows the major dependencies among chapters. As the diagram illustrates, it is possible to start with several different topics following the first two introductory chapters. Although the chart may seem complex, it is important to note that if the chapters are covered in order, the dependencies are not lost. The chart can be consulted by instructors wishing to use an alternative order of presentation.

For a one-semester course based on this book, selected chapters can be assigned as reading material. The book also can be used for a two-semester course sequence. The first course, *Introduction to Database Design and Database Systems*, at the sophomore, junior, or senior level, can cover most of Chapters 1 through 15. The second course, *Database Models and Implementation Techniques*, at the senior or first-year graduate level, can cover most of Chapters 16 through 30. The two-semester sequence can also be designed in various other ways, depending on the preferences of the instructors.





## Supplemental Materials

Support material is available to qualified instructors at Pearson's instructor resource center ([www.pearsonglobaleditions.com/elmasri](http://www.pearsonglobaleditions.com/elmasri)). For access, contact your local Pearson representative.

- PowerPoint lecture notes and figures.
- A solutions manual.

## Acknowledgments

It is a great pleasure to acknowledge the assistance and contributions of many individuals to this effort. First, we would like to thank our editor, Matt Goldstein, for his guidance, encouragement, and support. We would like to acknowledge the excellent work of Rose Kernan for production management, Patricia Daly for a



thorough copy editing of the book, Martha McMaster for her diligence in proofing the pages, and Scott Disanno, Managing Editor of the production team. We also wish to thank Kelsey Loanes from Pearson for her continued help with the project, and reviewers Michael Doherty, Deborah Dunn, Imad Rahal, Karen Davis, Gilliean Lee, Leo Mark, Monisha Pulimood, Hassan Reza, Susan Vrbsky, Li Da Xu, Weining Zhang and Vincent Oria.

Ramez Elmasri would like to thank Kulsawasd Jitkajornwanich, Vivek Sharma, and Surya Swaminathan for their help with preparing some of the material in Chapter 24. Sham Navathe would like to acknowledge the following individuals who helped in critically reviewing and revising various topics. Dan Forsythe and Satish Damle for discussion of storage systems; Rafi Ahmed for detailed re-organization of the material on query processing and optimization; Harish Butani, Balaji Palanisamy, and Prajakta Kalmegh for their help with the Hadoop and MapReduce technology material; Vic Ghorpadey and Nenad Jukic for revision of the Data Warehousing material; and finally, Frank Rietta for newer techniques in database security, Kunal Malhotra for various discussions, and Saurav Sahay for advances in information retrieval systems.

We would like to repeat our thanks to those who have reviewed and contributed to previous editions of *Fundamentals of Database Systems*.

- **First edition.** Alan Apt (editor), Don Batory, Scott Downing, Dennis Heimbinger, Julia Hodges, Yannis Ioannidis, Jim Larson, Per-Ake Larson, Dennis McLeod, Rahul Patel, Nicholas Roussopoulos, David Stemple, Michael Stonebraker, Frank Tompa, and Kyu-Young Whang.
- **Second edition.** Dan Joraanstad (editor), Rafi Ahmed, Antonio Albano, David Beech, Jose Blakeley, Panos Chrysanthis, Suzanne Dietrich, Vic Ghorpadey, Goetz Graefe, Eric Hanson, Junguk L. Kim, Roger King, Vram Kouramajian, Vijay Kumar, John Lowther, Sanjay Manchanda, Toshimi Minoura, Inderpal Mumick, Ed Omiecinski, Girish Pathak, Raghu Ramakrishnan, Ed Robertson, Eugene Sheng, David Stotts, Marianne Winslett, and Stan Zdonick.
- **Third edition.** Maite Suarez-Rivas and Katherine Harutunian (editors); Suzanne Dietrich, Ed Omiecinski, Rafi Ahmed, Francois Bancilhon, Jose Blakeley, Rick Cattell, Ann Chervenak, David W. Embley, Henry A. Etlinger, Leonidas Fegaras, Dan Forsyth, Farshad Fotouhi, Michael Franklin, Sreejith Gopinath, Goetz Craefe, Richard Hull, Sushil Jajodia, Ramesh K. Karne, Harish Kotbagi, Vijay Kumar, Tarcisio Lima, Ramon A. Mata-Toledo, Jack McCaw, Dennis McLeod, Rokia Missaoui, Magdi Morsi, M. Narayanaswamy, Carlos Ordonez, Joan Peckham, Betty Salzberg, Ming-Chien Shan, Junping Sun, Rajshekhar Sunderraman, Aravindan Veerasamy, and Emilia E. Villareal.
- **Fourth edition.** Maite Suarez-Rivas, Katherine Harutunian, Daniel Rausch, and Juliet Silveri (editors); Phil Bernhard, Zhengxin Chen, Jan Chomicki, Hakan Ferhatosmanoglu, Len Fisk, William Hankley, Ali R. Hurson, Vijay Kumar, Peretz Shoval, Jason T. L. Wang (reviewers); Ed Omiecinski (who contributed to Chapter 27). Contributors from the University of Texas at



# Contents

Preface	7
About the Authors	30



## ■ part 1

### Introduction to Databases ■

#### chapter 1 Databases and Database Users 33

1.1 Introduction	34
1.2 An Example	36
1.3 Characteristics of the Database Approach	40
1.4 Actors on the Scene	45
1.5 Workers behind the Scene	47
1.6 Advantages of Using the DBMS Approach	47
1.7 A Brief History of Database Applications	53
1.8 When Not to Use a DBMS	57
1.9 Summary	57
Review Questions	58
Exercises	58
Selected Bibliography	59

#### chapter 2 Database System Concepts and Architecture 61

2.1 Data Models, Schemas, and Instances	62
2.2 Three-Schema Architecture and Data Independence	66
2.3 Database Languages and Interfaces	68
2.4 The Database System Environment	72
2.5 Centralized and Client/Server Architectures for DBMSs	76
2.6 Classification of Database Management Systems	81
2.7 Summary	84
Review Questions	85
Exercises	85
Selected Bibliography	86



## ■ part 2

### Conceptual Data Modeling and Database Design ■

#### chapter 3 Data Modeling Using the Entity–Relationship (ER) Model 89

- 3.1 Using High-Level Conceptual Data Models for Database Design 90
- 3.2 A Sample Database Application 92
- 3.3 Entity Types, Entity Sets, Attributes, and Keys 93
- 3.4 Relationship Types, Relationship Sets, Roles, and Structural Constraints 102
- 3.5 Weak Entity Types 109
- 3.6 Refining the ER Design for the COMPANY Database 110
- 3.7 ER Diagrams, Naming Conventions, and Design Issues 111
- 3.8 Example of Other Notation: UML Class Diagrams 115
- 3.9 Relationship Types of Degree Higher than Two 118
- 3.10 Another Example: A UNIVERSITY Database 122
- 3.11 Summary 124
- Review Questions 126
- Exercises 126
- Laboratory Exercises 133
- Selected Bibliography 134

#### chapter 4 The Enhanced Entity–Relationship (EER) Model 137

- 4.1 Subclasses, Superclasses, and Inheritance 138
- 4.2 Specialization and Generalization 140
- 4.3 Constraints and Characteristics of Specialization and Generalization Hierarchies 143
- 4.4 Modeling of UNION Types Using Categories 150
- 4.5 A Sample UNIVERSITY EER Schema, Design Choices, and Formal Definitions 152
- 4.6 Example of Other Notation: Representing Specialization and Generalization in UML Class Diagrams 157
- 4.7 Data Abstraction, Knowledge Representation, and Ontology Concepts 158
- 4.8 Summary 165
- Review Questions 165
- Exercises 166
- Laboratory Exercises 173
- Selected Bibliography 176



## ■ part 3

### The Relational Data Model and SQL ■

#### chapter 5 The Relational Data Model and Relational Database Constraints 179

- 5.1 Relational Model Concepts 180
- 5.2 Relational Model Constraints and Relational Database Schemas 187
- 5.3 Update Operations, Transactions, and Dealing with Constraint Violations 195
- 5.4 Summary 199
- Review Questions 200
- Exercises 200
- Selected Bibliography 205

#### chapter 6 Basic SQL 207

- 6.1 SQL Data Definition and Data Types 209
- 6.2 Specifying Constraints in SQL 214
- 6.3 Basic Retrieval Queries in SQL 217
- 6.4 INSERT, DELETE, and UPDATE Statements in SQL 228
- 6.5 Additional Features of SQL 231
- 6.6 Summary 232
- Review Questions 233
- Exercises 233
- Selected Bibliography 235

#### chapter 7 More SQL: Complex Queries, Triggers, Views, and Schema Modification 237

- 7.1 More Complex SQL Retrieval Queries 237
- 7.2 Specifying Constraints as Assertions and Actions as Triggers 255
- 7.3 Views (Virtual Tables) in SQL 258
- 7.4 Schema Change Statements in SQL 262
- 7.5 Summary 264
- Review Questions 266
- Exercises 266
- Selected Bibliography 268

#### chapter 8 The Relational Algebra and Relational Calculus 269

- 8.1 Unary Relational Operations: SELECT and PROJECT 271
- 8.2 Relational Algebra Operations from Set Theory 276



8.3 Binary Relational Operations: JOIN and DIVISION	281
8.4 Additional Relational Operations	289
8.5 Examples of Queries in Relational Algebra	295
8.6 The Tuple Relational Calculus	298
8.7 The Domain Relational Calculus	307
8.8 Summary	309
Review Questions	310
Exercises	311
Laboratory Exercises	316
Selected Bibliography	318

## chapter 9 Relational Database Design by ER- and EER-to-Relational Mapping 319

9.1 Relational Database Design Using ER-to-Relational Mapping	320
9.2 Mapping EER Model Constructs to Relations	328
9.3 Summary	333
Review Questions	333
Exercises	333
Laboratory Exercises	335
Selected Bibliography	336

## ■ part 4

### Database Programming Techniques ■

## chapter 10 Introduction to SQL Programming Techniques 339

10.1 Overview of Database Programming Techniques and Issues	340
10.2 Embedded SQL, Dynamic SQL, and SQLJ	344
10.3 Database Programming with Function Calls and Class Libraries: SQL/CLI and JDBC	356
10.4 Database Stored Procedures and SQL/PSM	365
10.5 Comparing the Three Approaches	368
10.6 Summary	369
Review Questions	370
Exercises	370
Selected Bibliography	371

## chapter 11 Web Database Programming Using PHP 373

11.1 A Simple PHP Example	374
11.2 Overview of Basic Features of PHP	376



11.3 Overview of PHP Database Programming	383
11.4 Brief Overview of Java Technologies for Database Web Programming	388
11.5 Summary	388
Review Questions	389
Exercises	389
Selected Bibliography	389

## ■ part 5

### Object, Object-Relational, and XML: Concepts, Models, Languages, and Standards ■

#### chapter 12 Object and Object-Relational Databases 393

12.1 Overview of Object Database Concepts	395
12.2 Object Database Extensions to SQL	409
12.3 The ODMG Object Model and the Object Definition Language ODL	416
12.4 Object Database Conceptual Design	435
12.5 The Object Query Language OQL	438
12.6 Overview of the C++ Language Binding in the ODMG Standard	447
12.7 Summary	448
Review Questions	450
Exercises	451
Selected Bibliography	452

#### chapter 13 XML: Extensible Markup Language 455

13.1 Structured, Semistructured, and Unstructured Data	456
13.2 XML Hierarchical (Tree) Data Model	460
13.3 XML Documents, DTD, and XML Schema	463
13.4 Storing and Extracting XML Documents from Databases	472
13.5 XML Languages	473
13.6 Extracting XML Documents from Relational Databases	477
13.7 XML/SQL: SQL Functions for Creating XML Data	483
13.8 Summary	485
Review Questions	486
Exercises	486
Selected Bibliography	486



## ■ part 6

### Database Design Theory and Normalization ■

#### chapter 14 Basics of Functional Dependencies and Normalization for Relational Databases 489

- 14.1 Informal Design Guidelines for Relation Schemas 491
- 14.2 Functional Dependencies 501
- 14.3 Normal Forms Based on Primary Keys 504
- 14.4 General Definitions of Second and Third Normal Forms 513
- 14.5 Boyce-Codd Normal Form 517
- 14.6 Multivalued Dependency and Fourth Normal Form 521
- 14.7 Join Dependencies and Fifth Normal Form 524
- 14.8 Summary 525
- Review Questions 526
- Exercises 527
- Laboratory Exercises 531
- Selected Bibliography 532

#### chapter 15 Relational Database Design Algorithms and Further Dependencies 533

- 15.1 Further Topics in Functional Dependencies: Inference Rules, Equivalence, and Minimal Cover 535
- 15.2 Properties of Relational Decompositions 543
- 15.3 Algorithms for Relational Database Schema Design 549
- 15.4 About Nulls, Dangling Tuples, and Alternative Relational Designs 553
- 15.5 Further Discussion of Multivalued Dependencies and 4NF 557
- 15.6 Other Dependencies and Normal Forms 560
- 15.7 Summary 563
- Review Questions 564
- Exercises 565
- Laboratory Exercises 566
- Selected Bibliography 567



## ■ part 7

### File Structures, Hashing, Indexing, and Physical Database Design ■

#### chapter 16 Disk Storage, Basic File Structures, Hashing, and Modern Storage Architectures 571

- 16.1 Introduction 572
- 16.2 Secondary Storage Devices 577
- 16.3 Buffering of Blocks 586
- 16.4 Placing File Records on Disk 590
- 16.5 Operations on Files 594
- 16.6 Files of Unordered Records (Heap Files) 597
- 16.7 Files of Ordered Records (Sorted Files) 598
- 16.8 Hashing Techniques 602
- 16.9 Other Primary File Organizations 612
- 16.10 Parallelizing Disk Access Using RAID Technology 614
- 16.11 Modern Storage Architectures 618
- 16.12 Summary 622
- Review Questions 623
- Exercises 625
- Selected Bibliography 628

#### chapter 17 Indexing Structures for Files and Physical Database Design 631

- 17.1 Types of Single-Level Ordered Indexes 632
- 17.2 Multilevel Indexes 643
- 17.3 Dynamic Multilevel Indexes Using B-Trees and B<sup>+</sup>-Trees 647
- 17.4 Indexes on Multiple Keys 661
- 17.5 Other Types of Indexes 663
- 17.6 Some General Issues Concerning Indexing 668
- 17.7 Physical Database Design in Relational Databases 673
- 17.8 Summary 676
- Review Questions 677
- Exercises 678
- Selected Bibliography 680



## ■ part 8

### Query Processing and Optimization ■

#### chapter **18** Strategies for Query Processing 685

- 18.1 Translating SQL Queries into Relational Algebra and Other Operators 687
- 18.2 Algorithms for External Sorting 690
- 18.3 Algorithms for SELECT Operation 693
- 18.4 Implementing the JOIN Operation 698
- 18.5 Algorithms for PROJECT and Set Operations 706
- 18.6 Implementing Aggregate Operations and Different Types of JOINS 708
- 18.7 Combining Operations Using Pipelining 711
- 18.8 Parallel Algorithms for Query Processing 713
- 18.9 Summary 718
- Review Questions 718
- Exercises 719
- Selected Bibliography 719

#### chapter **19** Query Optimization 721

- 19.1 Query Trees and Heuristics for Query Optimization 722
- 19.2 Choice of Query Execution Plans 731
- 19.3 Use of Selectivities in Cost-Based Optimization 740
- 19.4 Cost Functions for SELECT Operation 744
- 19.5 Cost Functions for the JOIN Operation 747
- 19.6 Example to Illustrate Cost-Based Query Optimization 756
- 19.7 Additional Issues Related to Query Optimization 758
- 19.8 An Example of Query Optimization in Data Warehouses 761
- 19.9 Overview of Query Optimization in Oracle 763
- 19.10 Semantic Query Optimization 767
- 19.11 Summary 768
- Review Questions 769
- Exercises 770
- Selected Bibliography 770



## ■ part 9

### Transaction Processing, Concurrency Control, and Recovery ■

#### chapter **20** Introduction to Transaction Processing Concepts and Theory 775

- 20.1 Introduction to Transaction Processing 776
- 20.2 Transaction and System Concepts 783
- 20.3 Desirable Properties of Transactions 787
- 20.4 Characterizing Schedules Based on Recoverability 789
- 20.5 Characterizing Schedules Based on Serializability 793
- 20.6 Transaction Support in SQL 803
- 20.7 Summary 806
- Review Questions 807
- Exercises 807
- Selected Bibliography 809

#### chapter **21** Concurrency Control Techniques 811

- 21.1 Two-Phase Locking Techniques for Concurrency Control 812
- 21.2 Concurrency Control Based on Timestamp Ordering 822
- 21.3 Multiversion Concurrency Control Techniques 825
- 21.4 Validation (Optimistic) Techniques and Snapshot Isolation Concurrency Control 828
- 21.5 Granularity of Data Items and Multiple Granularity Locking 830
- 21.6 Using Locks for Concurrency Control in Indexes 835
- 21.7 Other Concurrency Control Issues 836
- 21.8 Summary 837
- Review Questions 838
- Exercises 839
- Selected Bibliography 840

#### chapter **22** Database Recovery Techniques 843

- 22.1 Recovery Concepts 844
- 22.2 NO-UNDO/REDO Recovery Based on Deferred Update 851
- 22.3 Recovery Techniques Based on Immediate Update 853



22.4 Shadow Paging	856
22.5 The ARIES Recovery Algorithm	857
22.6 Recovery in Multidatabase Systems	861
22.7 Database Backup and Recovery from Catastrophic Failures	862
22.8 Summary	863
Review Questions	864
Exercises	865
Selected Bibliography	868

## ■ part 10

### Distributed Databases, NOSQL Systems, and Big Data ■

#### chapter 23 Distributed Database Concepts 871

23.1 Distributed Database Concepts	872
23.2 Data Fragmentation, Replication, and Allocation Techniques for Distributed Database Design	877
23.3 Overview of Concurrency Control and Recovery in Distributed Databases	884
23.4 Overview of Transaction Management in Distributed Databases	887
23.5 Query Processing and Optimization in Distributed Databases	889
23.6 Types of Distributed Database Systems	895
23.7 Distributed Database Architectures	898
23.8 Distributed Catalog Management	905
23.9 Summary	906
Review Questions	907
Exercises	908
Selected Bibliography	910

#### chapter 24 NOSQL Databases and Big Data Storage Systems 913

24.1 Introduction to NOSQL Systems	914
24.2 The CAP Theorem	918
24.3 Document-Based NOSQL Systems and MongoDB	920
24.4 NOSQL Key-Value Stores	925
24.5 Column-Based or Wide Column NOSQL Systems	930
24.6 NOSQL Graph Databases and Neo4j	933
24.7 Summary	939
Review Questions	939
Selected Bibliography	940



## chapter **25** Big Data Technologies Based on MapReduce and Hadoop 941

- 25.1 What Is Big Data? 944
- 25.2 Introduction to MapReduce and Hadoop 946
- 25.3 Hadoop Distributed File System (HDFS) 951
- 25.4 MapReduce: Additional Details 956
- 25.5 Hadoop v2 alias YARN 966
- 25.6 General Discussion 974
- 25.7 Summary 983
- Review Questions 984
- Selected Bibliography 986

## ■ part **11**

### Advanced Database Models, Systems, and Applications ■

## chapter **26** Enhanced Data Models: Introduction to Active, Temporal, Spatial, Multimedia, and Deductive Databases 991

- 26.1 Active Database Concepts and Triggers 993
- 26.2 Temporal Database Concepts 1004
- 26.3 Spatial Database Concepts 1017
- 26.4 Multimedia Database Concepts 1024
- 26.5 Introduction to Deductive Databases 1029
- 26.6 Summary 1042
- Review Questions 1044
- Exercises 1045
- Selected Bibliography 1048

## chapter **27** Introduction to Information Retrieval and Web Search 1051

- 27.1 Information Retrieval (IR) Concepts 1052
- 27.2 Retrieval Models 1059
- 27.3 Types of Queries in IR Systems 1065
- 27.4 Text Preprocessing 1067
- 27.5 Inverted Indexing 1070
- 27.6 Evaluation Measures of Search Relevance 1074
- 27.7 Web Search and Analysis 1077



27.8 Trends in Information Retrieval	1087
27.9 Summary	1093
Review Questions	1094
Selected Bibliography	1096

## chapter **28** Data Mining Concepts 1099

28.1 Overview of Data Mining Technology	1100
28.2 Association Rules	1103
28.3 Classification	1115
28.4 Clustering	1118
28.5 Approaches to Other Data Mining Problems	1121
28.6 Applications of Data Mining	1124
28.7 Commercial Data Mining Tools	1124
28.8 Summary	1127
Review Questions	1127
Exercises	1128
Selected Bibliography	1129

## chapter **29** Overview of Data Warehousing and OLAP 1131

29.1 Introduction, Definitions, and Terminology	1132
29.2 Characteristics of Data Warehouses	1133
29.3 Data Modeling for Data Warehouses	1135
29.4 Building a Data Warehouse	1141
29.5 Typical Functionality of a Data Warehouse	1144
29.6 Data Warehouse versus Views	1145
29.7 Difficulties of Implementing Data Warehouses	1146
29.8 Summary	1147
Review Questions	1147
Selected Bibliography	1148

# ■ part **12**

## Additional Database Topics: Security ■

## chapter **30** Database Security 1151

30.1 Introduction to Database Security Issues	1152
30.2 Discretionary Access Control Based on Granting and Revoking Privileges	1159
30.3 Mandatory Access Control and Role-Based Access Control for Multilevel Security	1164



30.4 SQL Injection	1173
30.5 Introduction to Statistical Database Security	1176
30.6 Introduction to Flow Control	1177
30.7 Encryption and Public Key Infrastructures	1179
30.8 Privacy Issues and Preservation	1183
30.9 Challenges to Maintaining Database Security	1184
30.10 Oracle Label-Based Security	1185
30.11 Summary	1188
Review Questions	1189
Exercises	1190
Selected Bibliography	1191
<b>appendix A Alternative Diagrammatic Notations for ER Models</b>	<b>1193</b>
<b>appendix B Parameters of Disks</b>	<b>1197</b>
<b>appendix C Overview of the QBE Language</b>	<b>1201</b>
C.1 Basic Retrievals in QBE	1201
C.2 Grouping, Aggregation, and Database Modification in QBE	1205
<b>appendix D Overview of the Hierarchical Data Model</b> (located on the Companion Website at <a href="http://www.pearsonglobaleditions.com/Elmasri">www.pearsonglobaleditions.com/Elmasri</a> )	
<b>appendix E Overview of the Network Data Model</b> (located on the Companion Website at <a href="http://www.pearsonglobaleditions.com/Elmasri">www.pearsonglobaleditions.com/Elmasri</a> )	
<b>Bibliography</b>	<b>1209</b>
<b>Index</b>	<b>1245</b>