教育部高等教育司推荐
国外优秀信息科学与技术系列教学用书

数据库系统概念

（第四版　影印版）

DATABASE SYSTEM CONCEPTS

(Fourth Edition)

Abraham Silberschatz
Henry F. Korth
S. Sudarshan

高等教育出版社

McGraw-Hill Companies

需要全本请访问 www.ertongbook.com

图字：01-2001-3752 号

Database System Concepts, Fourth Edition

Abraham Silberschatz，Henry F. Korth，S. Sudarshan

Copyright ©2002 by **The McGraw-Hill Companies, Inc.**

All rights reserved.

This edition is authorized for sale only in the People's Republic of China (excluding the Special Administrative Regions of Hong Kong and Macau).

图书在版编目(CIP)数据

数据库系统概念：第 4 版 /（美）西尔伯沙茨（Silberschatz,A.）等著. —影印版. —北京：高等教育出版社，2002.5

计算机专业本科、研究生

ISBN 7-04-011049-0

Ⅰ.数... Ⅱ.西.. Ⅲ.数据库系统-英文

Ⅳ.TP311.13

中国版本图书馆 CIP 数据核字（2002）第 030749 号

责任编辑 康兆华　**封面设计** 张　楠　**责任印制** 陈伟光

数据库系统概念（第四版　影印版）

Abraham Silberschatz, Henry F. Korth, S. Sudarshan

出版发行	高等教育出版社	**购书热线**	010-64054588
社　　址	北京市东城区沙滩后街 55 号	**免费咨询**	800-810-0598
邮政编码	100009	**网　　址**	http://www.hep.edu.cn
传　　真	010-64014048		http://www.hep.com.cn
经　　销	新华书店北京发行所		
印　　刷	北京外文印刷厂		
开　　本	787×1092　1/16	**版　　次**	2002 年 5 月第 4 版
印　　张	68	**印　　次**	2002 年 5 月第 1 次印刷
字　　数	1 480 000	**定　　价**	59.50 元

本书如有缺页、倒页、脱页等质量问题，请到所购图书销售部门联系调换。

版权所有　侵权必究

前　言

20世纪末，以计算机和通信技术为代表的信息科学和技术对世界经济、科技、军事、教育和文化等产生了深刻影响。信息科学技术的迅速普及和应用，带动了世界范围信息产业的蓬勃发展，为许多国家带来了丰厚的回报。

进入21世纪，尤其随着我国加入WTO，信息产业的国际竞争将更加激烈。我国信息产业虽然在20世纪末取得了迅猛发展，但与发达国家相比，甚至与印度、爱尔兰等国家相比，还有很大差距。国家信息化的发展速度和信息产业的国际竞争能力，最终都将取决于信息科学技术人才的质量和数量。引进国外信息科学和技术优秀教材，在有条件的学校推动开展英语授课或双语教学，是教育部为加快培养大批高质量的信息技术人才采取的一项重要举措。

为此，教育部要求由高等教育出版社首先开展信息科学和技术教材的引进试点工作。同时提出了两点要求，一是要高水平，二是要低价格。在高等教育出版社和信息科学技术引进教材专家组的努力下，经过比较短的时间，第一批引进的20多种教材已经陆续出版。这套教材出版后受到了广泛的好评，其中有不少是世界信息科学技术领域著名专家、教授的经典之作和反映信息科学技术最新进展的优秀作品，代表了目前世界信息科学技术教育的一流水平，而且价格也是最优惠的，与国内同类自编教材相当。

这项教材引进工作是在教育部高等教育司和高教社的共同组织下，由国内信息科学技术领域的专家、教授广泛参与，在对大量国外教材进行多次遴选的基础上，参考了国内和国外著名大学相关专业的课程设置进行系统引进的。其中，John Wiley公司出版的贝尔实验室信息科学研究中心副总裁Silberschatz教授的经典著作《操作系统概念》，是我们经过反复谈判，做了很多努力才得以引进的。William Stallings先生曾编写了在美国深受欢迎的信息科学技术系列教材，其中有多种教材获得过美国教材和学术著作者协会颁发的计算机科学与工程教材奖，这批引进教材中就有他的两本著作。留美中国学者Jiawei Han先生的《数据挖掘》是该领域中具有里程碑意义的著作。由达特茅斯学院Thomas Cormen和麻省理工学院、哥伦比亚大学的几

位学者共同编著的经典著作《算法导论》，在经历了 11 年的锤炼之后于 2001 年出版了第二版。目前任教于美国 Massachusetts 大学的 James Kurose 教授，曾在美国三所高校先后 10 次获得杰出教师或杰出教学奖，由他主编的《计算机网络》出版后，以其体系新颖、内容先进而倍受欢迎。在努力降低引进教材售价方面，高等教育出版社做了大量和细致的工作。这套引进的教材体现了权威性、系统性、先进性和经济性等特点。

教育部也希望国内和国外的出版商积极参与此项工作，共同促进中国信息技术教育和信息产业的发展。我们在与外商的谈判工作中，不仅要坚定不移地引进国外最优秀的教材，而且还要千方百计地将版权转让费降下来，要让引进教材的价格与国内自编教材相当，让广大教师和学生负担得起。中国的教育市场巨大，外国出版公司和国内出版社要通过扩大发行数量取得效益。

在引进教材的同时，我们还应做好消化吸收，注意学习国外先进的教学思想和教学方法，提高自编教材的水平，使我们的教学和教材在内容体系上，在理论与实践的结合上，在培养学生的动手能力上能有较大的突破和创新。

目前，教育部正在全国 35 所高校推动示范性软件学院的建设和实施，这也是加快培养信息科学技术人才的重要举措之一。示范性软件学院要立足于培养具有国际竞争力的实用性软件人才，与国外知名高校或著名企业合作办学，以国内外著名 IT 企业为实践教学基地，聘请国内外知名教授和软件专家授课，还要率先使用引进教材开展教学。

我们希望通过这些举措，能在较短的时间，为我国培养一大批高质量的信息技术人才，提高我国软件人才的国际竞争力，促进我国信息产业的快速发展，加快推动国家信息化进程，进而带动整个国民经济的跨越式发展。

教育部高等教育司

二〇〇二年三月

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

需要全本请访问 www.ertongbook.com

Contents

Preface xv

Chapter 1 Introduction

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

PART 1 ■ DATA MODELS

Chapter 2 Entity-Relationship Model

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

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

Chapter 7 Relational-Database Design

7.1 First Normal Form 257
7.2 Pitfalls in Relational-Database Design 258
7.3 Functional Dependencies 260
7.4 Decomposition 271
7.5 Desirable Properties of Decomposition 275
7.6 Boyce–Codd Normal Form 279
7.7 Third Normal Form 284
7.8 Fourth Normal Form 289
7.9 More Normal Forms 293
7.10 Overall Database Design Process 293
7.11 Summary 297
Exercises 299
Bibliographical Notes 303

PART 3 ■ OBJECT-BASED DATABASES AND XML

Chapter 8 Object-Oriented Databases

8.1 Need for Complex Data Types 307
8.2 The Object-Oriented Data Model 308
8.3 Object-Oriented Languages 318
8.4 Persistent Programming Languages 318
8.5 Persistent C++ Systems 322
8.6 Persistent Java Systems 330
8.7 Summary 331
Exercises 332
Bibliographical Notes 333

Chapter 9 Object-Relational Databases

9.1 Nested Relations 335
9.2 Complex Types 337
9.3 Inheritance 342
9.4 Reference Types 346
9.5 Querying with Complex Types 348
9.6 Functions and Procedures 351
9.7 Object-Oriented versus Object-Relational 356
9.8 Summary 357
Exercises 358
Bibliographical Notes 360

Chapter 10 XML

10.1 Background 361
10.2 Structure of XML Data 364
10.3 XML Document Schema 367
10.4 Querying and Transformation 372
10.5 The Application Program Interface 380
10.6 Storage of XML Data 381
10.7 XML Applications 384
10.8 Summary 386
Exercises 388
Bibliographical Notes 390

PART 4 ■ DATA STORAGE AND QUERYING

Chapter 11 Storage and File Structure

11.1 Overview of Physical Storage Media 393
11.2 Magnetic Disks 396
11.3 RAID 402
11.4 Tertiary Storage 410
11.5 Storage Access 412
11.6 File Organization 415
11.7 Organization of Records in Files 422
11.8 Data-Dictionary Storage 426
11.9 Storage for Object-Oriented Databases 428
11.10 Summary 438
Exercises 440
Bibliographical Notes 443

Chapter 12 Indexing and Hashing

12.1 Basic Concepts 445
12.2 Ordered Indices 446
12.3 B^+-Tree Index Files 453
12.4 B-Tree Index Files 464
12.5 Static Hashing 465
12.6 Dynamic Hashing 471
12.7 Comparison of Ordered Indexing and Hashing 477
12.8 Index Definition in SQL 479
12.9 Multiple-Key Access 480
12.10 Summary 487
Exercises 489
Bibliographical Notes 491

Chapter 13 Query Processing

13.1 Overview 493
13.2 Measures of Query Cost 495
13.3 Selection Operation 496
13.4 Sorting 501
13.5 Join Operation 503
13.6 Other Operations 514
13.7 Evaluation of Expressions 518
13.8 Summary 523
Exercises 525
Bibliographical Notes 526

Chapter 14 Query Optimization

14.1 Overview 529
14.2 Estimating Statistics of Expression Results 531
14.3 Transformation of Relational Expressions 537
14.4 Choice of Evaluation Plans 544
14.5 Materialized Views 553
14.6 Summary 557
Exercises 559
Bibliographical Notes 561

PART 5 ■ TRANSACTION MANAGEMENT

Chapter 15 Transactions

15.1 Transaction Concept 565
15.2 Transaction State 568
15.3 Implementation of Atomicity and Durability 571
15.4 Concurrent Executions 573
15.5 Serializability 576
15.6 Recoverability 582
15.7 Implementation of Isolation 583
15.8 Transaction Definition in SQL 584
15.9 Testing for Serializability 584
15.10 Summary 586
Exercises 588
Bibliographical Notes 590

Chapter 16 Concurrency Control

16.1 Lock-Based Protocols 591
16.2 Timestamp-Based Protocols 604
16.3 Validation-Based Protocols 607
16.4 Multiple Granularity 609
16.5 Multiversion Schemes 612
16.6 Deadlock Handling 615
16.7 Insert and Delete Operations 620
16.8 Weak Levels of Consistency 623
16.9 Concurrency in Index Structures 625
16.10 Summary 629
Exercises 632
Bibliographical Notes 636

Chapter 17 Recovery System

17.1 Failure Classification 639
17.2 Storage Structure 640
17.3 Recovery and Atomicity 644
17.4 Log-Based Recovery 645
17.5 Shadow Paging 653
17.6 Recovery with Concurrent Transactions 657
17.7 Buffer Management 660
17.8 Failure with Loss of Nonvolatile Storage 663
17.9 Advanced Recovery Techniques 664
17.10 Remote Backup Systems 672
17.11 Summary 674
Exercises 677
Bibliographical Notes 679

PART 6 ■ DATABASE SYSTEM ARCHITECTURE

Chapter 18 Database System Architectures

18.1 Centralized and Client–Server Architectures 683
18.2 Server System Architectures 687
18.3 Parallel Systems 691
18.4 Distributed Systems 697
18.5 Network Types 701
18.6 Summary 703
Exercises 705
Bibliographical Notes 707

Chapter 19 Distributed Databases

19.1 Homogeneous and Heterogeneous Databases 709
19.2 Distributed Data Storage 710
19.3 Distributed Transactions 713
19.4 Commit Protocols 716
19.5 Concurrency Control in Distributed Databases 722
19.6 Availability 730
19.7 Distributed Query Processing 735
19.8 Heterogeneous Distributed Databases 738
19.9 Directory Systems 741
19.10 Summary 746
Exercises 749
Bibliographical Notes 752

Chapter 20 Parallel Databases

20.1 Introduction 755
20.2 I/O Parallelism 756
20.3 Interquery Parallelism 760
20.4 Intraquery Parallelism 761
20.5 Intraoperation Parallelism 762
20.6 Interoperation Parallelism 770
20.7 Design of Parallel Systems 772
20.8 Summary 773
Exercises 775
Bibliographical Notes 777

PART 7 ■ OTHER TOPICS

Chapter 21 Application Development and Administration

21.1 Web Interfaces to Databases 781
21.2 Performance Tuning 790
21.3 Performance Benchmarks 798
21.4 Standardization 802
21.5 E-Commerce 806
21.6 Legacy Systems 809
21.7 Summary 810
Exercises 812
Bibliographical Notes 815

Chapter 22 Advanced Querying and Information Retrieval

22.1 Decision-Support Systems 817
22.2 Data Analysis and OLAP 819
22.3 Data Mining 830
22.4 Data Warehousing 842
22.5 Information-Retrieval Systems 846
22.6 Summary 856
Exercises 859
Bibliographical Notes 861

Chapter 23 Advanced Data Types and New Applications

23.1 Motivation 863
23.2 Time in Databases 864
23.3 Spatial and Geographic Data 866
23.4 Multimedia Databases 877
23.5 Mobility and Personal Databases 880
23.6 Summary 885
Exercises 887
Bibliographical Notes 889

需要全本请访问 www.ertongbook.com

Chapter 24 Advanced Transaction Processing

24.1 Transaction-Processing Monitors 891
24.2 Transactional Workflows 895
24.3 Main-Memory Databases 901
24.4 Real-Time Transaction Systems 903
24.5 Long-Duration Transactions 904
24.6 Transaction Management in Multidatabases 910
24.7 Summary 914
Exercises 916
Bibliographical Notes 917

PART 8 ■ CASE STUDIES

Chapter 25 Oracle

25.1 Database Design and Querying Tools 921
25.2 SQL Variations and Extensions 923
25.3 Storage and Indexing 925
25.4 Query Processing and Optimization 934
25.5 Concurrency Control and Recovery 940
25.6 System Architecture 942
25.7 Replication, Distribution, and External Data 945
25.8 Database Administration Tools 946
Bibliographical Notes 947

Chapter 26 IBM DB2 Universal Database

26.1 Database Design and Querying Tools 950
26.2 SQL Variations and Extensions 951
26.3 Storage and Indexing 953
26.4 Query Processing and Optimization 956
26.5 Concurrency Control and Recovery 960
26.6 System Architecture 963
26.7 Replication, Distribution, and External Data 965
26.8 Database Administration Tools 965
26.9 Summary 967
Bibliographical Notes 967

Chapter 27 Microsoft SQL Server

27.1 Management, Design, and Querying Tools 969
27.2 SQL Variations and Extensions 975
27.3 Storage and Indexing 980
27.4 Query Processing and Optimization 982
27.5 Concurrency and Recovery 986
27.6 System Architecture 991
27.7 Data Access 992
27.8 Distribution and Replication 994
27.9 Full-Text Queries on Relational Data 998
27.10 Data Warehousing and Analysis Services 999
27.11 XML and Web Support 1002
27.12 Summary 1005
Bibliographical Notes 1005

PART 9 ■ APPENDICES

Appendix A Network Model (contents online)

A.1 Basic Concepts A1
A.2 Data-Structure Diagrams A2
A.3 The DBTG CODASYL Model A7
A.4 DBTG Data-Retrieval Facility A13
A.5 DBTG Update Facility A20
A.6 DBTG Set-Processing Facility A22
A.7 Mapping of Networks to Files A27
A.8 Summary A31
Exercises A32
Bibliographical Notes A35

Appendix B Hierarchical Model (contents online)

B.1 Basic Concepts B1
B.2 Tree-Structure Diagrams B2
B.3 Data-Retrieval Facility B13
B.4 Update Facility B18
B.5 Virtual Records B21
B.6 Mapping of Hierarchies to Files B22
B.7 The IMS Database System B24
B.8 Summary B25
Exercises B26
Bibliographical Notes B29

Appendix C Advanced Relational Design (contents online)

C.1 Multivalued Dependencies C1
C.2 Join Dependencies C5
C.3 Domain-Key Normal Form C8
C.4 Summary C10
Exercises C10
Bibliographical Notes C11

Bibliography 1007

Index 1043

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^+-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.

需要全本请访问 www.ertongbook.com