

GLOBAL
EDITION



Database Systems

*A Practical Approach to Design, Implementation,
and Management*

SIXTH EDITION

Thomas Connolly • Carolyn Begg

ALWAYS LEARNING

PEARSON

ONLINE ACCESS

Thank you for purchasing a new copy of Database Systems, Sixth Edition. Your textbook includes one year of prepaid access to the book's Companion Website. This prepaid subscription provides you with full access to the following student support areas:

- online appendices
- tutorials on selected chapters
- DreamHome web implementation

**Use a coin to scratch off the coating and reveal your student access code.
Do not use a knife or other sharp object as it may damage the code.**

To access the Database Systems, Sixth Edition., Companion Website for the first time, you will need to register online using a computer with an Internet connection and a web browser. The process takes just a couple of minutes and only needs to be completed once.

- 1.** Go to **www.pearsonglobaleditions.com/connolly**
- 2.** Click on **Companion Website**.
- 3.** Click on the **Register** button.
- 4.** On the registration page, enter your student access code* found beneath the scratch-off panel. Do not type the dashes. You can use lower- or uppercase.
- 5.** Follow the on-screen instructions. If you need help at any time during the online registration process, simply click the **Need Help?** icon.
- 6.** Once your personal Login Name and Password are confirmed, you can begin using the Database Systems Companion Website!

To log in after you have registered:

You only need to register for this Companion Website once. After that, you can log in any time at **www.pearsonglobaleditions.com/connolly** by providing your Login Name and Password when prompted.

*Important: The access code can only be used once. This subscription is valid for one year upon activation and is not transferable. If this access code has already been revealed, it may no longer be valid. If this is the case, you can purchase a subscription by going to **www.pearsonglobaleditions.com/connolly** and following the on-screen instructions.

DATABASE SYSTEMS

A Practical Approach to Design, Implementation, and Management

SIXTH EDITION
GLOBAL EDITION

DATABASE SYSTEMS

A Practical Approach to Design, Implementation, and Management

SIXTH EDITION
GLOBAL EDITION

THOMAS M. CONNOLLY | CAROLYN E. BEGG
UNIVERSITY OF THE WEST OF SCOTLAND

PEARSON

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

Editorial Director, ECS:	Marcia Horton	Senior Production Manufacturing	
Head of Learning Asset		Controller, Global Editions:	Trudy Kimber
Acquisition, Global Editions:	Laura Dent	Manufacturing Buyer:	Linda Sager
Acquisitions Editor:	Matt Goldstein	Art Director:	Jayne Conte
Acquisitions Editor,		Cover Designer:	Lumina Datamatics
Global Editions:	Subhasree Patra	Text Designer:	Susan Raymond
Program Manager:	Kayla Smith-Tarbox	Manager, Text Permissions:	Tim Nicholls
Director of Marketing:	Christy Lesko	Text Permission Project Manager:	Jenell Forschler
Marketing Manager:	Yezan Alayan	Cover Image:	© Africa Studio/ Shutterstock
Marketing Assistant:	Jon Bryant	Media Project Manager:	Renata Butera
Director of Production:	Erin Gregg	Full-Service Project Management:	Vasundhara Sawhney/ Cenveo® Publisher Services
Senior Managing Editor:	Scott Disanno		
Senior Project Manager:	Marilyn Lloyd		
Media Producer, Global Editions:	M Vikram Kumar		
Project Editor, Global Editions:	K.K. Neelakantan		

Credits and acknowledgments borrowed from other sources and reproduced, with permission, in this textbook appear on the Credits page at the end of the book.

Pearson Education Limited
Edinburgh Gate
Harlow
Essex CM20 2JE
England

and Associated Companies throughout the world

Visit us on the World Wide Web at:
www.pearsonglobaleditions.com

© Pearson Education Limited 2015

The rights of Thomas Connolly and Carolyn Begg to be identified as the authors of this work have been asserted by them in accordance with the Copyright, Designs and Patents Act 1988.

Authorized adaptation from the United States edition, entitled Database Systems: A Practical Approach to Design, Implementation, and Management, 6th edition, ISBN 978-0-13-294326-0, by Thomas Connolly and Carolyn Begg, published by Pearson Education © 2015.

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, without either the prior written permission of the publisher or a license permitting restricted copying in the United Kingdom issued by the Copyright Licensing Agency Ltd, Saffron House, 6–10 Kirby Street, London EC1N 8TS.

All trademarks used herein are the property of their respective owners. The use of any trademark in this text does not vest in the author or publisher any trademark ownership rights in such trademarks, nor does the use of such trademarks imply any affiliation with or endorsement of this book by such owners.

British Library Cataloguing-in-Publication Data

A catalogue record for this book is available from the British Library

10 9 8 7 6 5 4 3 2 1
15 14 13 12 11

ISBN 10: 1-292-06118-9

ISBN 13: 978-1-292-06118-4

Typeset by Cenveo Publishing Services.

Printed and bound by Courier Westford in The United States of America.

To Sheena, for her patience, understanding, and love.

To our beautiful children Kathryn, Michael and Stephen with all our love.

And to my brother, Francis, who died during the writing of this book.

Thomas M. Connolly

To my past, present, and future students at UWS.

Carolyn E. Begg

Brief Contents

<i>Preface</i>	35
Part I Background	49
Chapter 1 Introduction to Databases	51
Chapter 2 Database Environment	83
Chapter 3 Database Architectures and the Web	105
Part 2 The Relational Model and Languages	147
Chapter 4 The Relational Model	149
Chapter 5 Relational Algebra and Relational Calculus	167
Chapter 6 SQL: Data Manipulation	191
Chapter 7 SQL: Data Definition	233
Chapter 8 Advanced SQL	271
Chapter 9 Object-Relational DBMSs	291
Part 3 Database Analysis and Design	343
Chapter 10 Database System Development Lifecycle	345
Chapter 11 Database Analysis and the <i>DreamHome</i> Case Study	375
Chapter 12 Entity–Relationship Modeling	405
Chapter 13 Enhanced Entity–Relationship Modeling	433
Chapter 14 Normalization	451
Chapter 15 Advanced Normalization	481
Part 4 Methodology	501
Chapter 16 Methodology—Conceptual Database Design	503
Chapter 17 Methodology—Logical Database Design for the Relational Model	527

Chapter 18	Methodology—Physical Database Design for Relational Databases	561
Chapter 19	Methodology—Monitoring and Tuning the Operational System	585
Part 5	Selected Database Issues	605
Chapter 20	Security and Administration	607
Chapter 21	Professional, Legal, and Ethical Issues in Data Management	641
Chapter 22	Transaction Management	667
Chapter 23	Query Processing	727
Part 6	Distributed DBMSs and Replication	783
Chapter 24	Distributed DBMSs—Concepts and Design	785
Chapter 25	Distributed DBMSs—Advanced Concepts	831
Chapter 26	Replication and Mobile Databases	875
Part 7	Object DBMSs	939
Chapter 27	Object-Oriented DBMSs—Concepts and Design	941
Chapter 28	Object-Oriented DBMSs—Standards and Systems	995
Part 8	The Web and DBMSs	1045
Chapter 29	Web Technology and DBMSs	1047
Chapter 30	Semistructured Data and XML	1129
Part 9	Business Intelligence	1221
Chapter 31	Data Warehousing Concepts	1223
Chapter 32	Data Warehousing Design	1257
Chapter 33	OLAP	1285
Chapter 34	Data Mining	1315

Appendices

A Users' Requirements Specification for <i>DreamHome</i> Case Study	A-1
B Other Case Studies	B-1
C Alternative ER Modeling Notations	C-1
D Summary of the Database Design Methodology for Relational Databases	D-1
E Introduction to Pyrrho: A Lightweight RDBMS	E-1
F File Organizations and Indexes (Online)	F-1
G When Is a DBMS Relational? (Online)	G-1
H Commercial DBMSs: Access® and Oracle® (Online)	H-1
I Programmatic SQL (Online)	I-1
J Estimating Disk Space Requirements (Online)	J-1
K Introduction to Object-Oriented Concepts (Online)	K-1
L Example Web Scripts (Online)	L-1
M Query-By-Example (QBE) (Online)	M-1
N Third Generation Manifestos (Online)	N-1
O Postgres—An Early ORDBMS (Online)	O-1

References**R-I****Further Reading****FR-I****Index****IN-I****I329**

Contents

<i>Preface</i>	35
Part I Background	49
Chapter 1 Introduction to Databases	51
1.1 Introduction	52
1.2 Traditional File-Based Systems	55
1.2.1 File-Based Approach	55
1.2.2 Limitations of the File-Based Approach	60
1.3 Database Approach	62
1.3.1 The Database	63
1.3.2 The Database Management System (DBMS)	64
1.3.3 (Database) Application Programs	65
1.3.4 Components of the DBMS Environment	66
1.3.5 Database Design: The Paradigm Shift	69
1.4 Roles in the Database Environment	69
1.4.1 Data and Database Administrators	69
1.4.2 Database Designers	70
1.4.3 Application Developers	71
1.4.4 End-Users	71
1.5 History of Database Management Systems	71
1.6 Advantages and Disadvantages of DBMSs	75
<i>Chapter Summary</i>	79
<i>Review Questions</i>	80
<i>Exercises</i>	80
Chapter 2 Database Environment	83
2.1 The Three-Level ANSI-SPARC Architecture	84
2.1.1 External Level	85
2.1.2 Conceptual Level	86
2.1.3 Internal Level	86
2.1.4 Schemas, Mappings, and Instances	87
2.1.5 Data Independence	88
2.2 Database Languages	89
2.2.1 The Data Definition Language (DDL)	90

2.2.2	The Data Manipulation Language (DML)	90
2.2.3	Fourth-Generation Languages (4GLs)	92
2.3	Data Models and Conceptual Modeling	93
2.3.1	Object-Based Data Models	94
2.3.2	Record-Based Data Models	94
2.3.3	Physical Data Models	97
2.3.4	Conceptual Modeling	97
2.4	Functions of a DBMS	97
	<i>Chapter Summary</i>	102
	<i>Review Questions</i>	103
	<i>Exercises</i>	104
Chapter 3	Database Architectures and the Web	105
3.1	Multi-user DBMS Architectures	106
3.1.1	Teleprocessing	106
3.1.2	File-Server Architecture	107
3.1.3	Traditional Two-Tier Client–Server Architecture	108
3.1.4	Three-Tier Client–Server Architecture	111
3.1.5	<i>N</i> -Tier Architectures	112
3.1.6	Middleware	113
3.1.7	Transaction Processing Monitors	115
3.2	Web Services and Service-Oriented Architectures	117
3.2.1	Web Services	117
3.2.2	Service-Oriented Architectures (SOA)	119
3.3	Distributed DBMSs	120
3.4	Data Warehousing	123
3.5	Cloud Computing	125
3.5.1	Benefits and Risks of Cloud Computing	127
3.5.2	Cloud-based database solutions	130
3.6	Components of a DBMS	134
3.7	Oracle Architecture	137
3.7.1	Oracle’s Logical Database Structure	137
3.7.2	Oracle’s Physical Database Structure	140
	<i>Chapter Summary</i>	144
	<i>Review Questions</i>	145
	<i>Exercises</i>	145
Part 2	The Relational Model and Languages	147
Chapter 4	The Relational Model	149
4.1	Brief History of the Relational Model	150
4.2	Terminology	152
4.2.1	Relational Data Structure	152

4.2.2	Mathematical Relations	155
4.2.3	Database Relations	156
4.2.4	Properties of Relations	156
4.2.5	Relational Keys	158
4.2.6	Representing Relational Database Schemas	159
4.3	Integrity Constraints	161
4.3.1	Nulls	161
4.3.2	Entity Integrity	162
4.3.3	Referential Integrity	162
4.3.4	General Constraints	163
4.4	Views	163
4.4.1	Terminology	163
4.4.2	Purpose of Views	164
4.4.3	Updating Views	165
	<i>Chapter Summary</i>	165
	<i>Review Questions</i>	166
	<i>Exercises</i>	166
Chapter 5	Relational Algebra and Relational Calculus	167
5.1	The Relational Algebra	168
5.1.1	Unary Operations	168
5.1.2	Set Operations	171
5.1.3	Join Operations	174
5.1.4	Division Operation	177
5.1.5	Aggregation and Grouping Operations	178
5.1.6	Summary of the Relational Algebra Operations	180
5.2	The Relational Calculus	181
5.2.1	Tuple Relational Calculus	181
5.2.2	Domain Relational Calculus	184
5.3	Other Languages	186
	<i>Chapter Summary</i>	187
	<i>Review Questions</i>	187
	<i>Exercises</i>	188
Chapter 6	SQL: Data Manipulation	191
6.1	Introduction to SQL	192
6.1.1	Objectives of SQL	192
6.1.2	History of SQL	193
6.1.3	Importance of SQL	195
6.1.4	Terminology	195
6.2	Writing SQL Commands	195

6.3	Data Manipulation	196
6.3.1	Simple Queries	197
6.3.2	Sorting Results (ORDER BY Clause)	205
6.3.3	Using the SQL Aggregate Functions	207
6.3.4	Grouping Results (GROUP BY Clause)	209
6.3.5	Subqueries	212
6.3.6	ANY and ALL	214
6.3.7	Multi-table Queries	216
6.3.8	EXISTS and NOT EXISTS	222
6.3.9	Combining Result Tables (UNION, INTERSECT, EXCEPT)	223
6.3.10	Database Updates	225
	<i>Chapter Summary</i>	229
	<i>Review Questions</i>	230
	<i>Exercises</i>	230
Chapter 7	SQL: Data Definition	233
7.1	The ISO SQL Data Types	234
7.1.1	SQL Identifiers	234
7.1.2	SQL Scalar Data Types	235
7.2	Integrity Enhancement Feature	240
7.2.1	Required Data	240
7.2.2	Domain Constraints	240
7.2.3	Entity Integrity	241
7.2.4	Referential Integrity	242
7.2.5	General Constraints	243
7.3	Data Definition	244
7.3.1	Creating a Database	244
7.3.2	Creating a Table (CREATE TABLE)	245
7.3.3	Changing a Table Definition (ALTER TABLE)	248
7.3.4	Removing a Table (DROP TABLE)	249
7.3.5	Creating an Index (CREATE INDEX)	250
7.3.6	Removing an Index (DROP INDEX)	250
7.4	Views	251
7.4.1	Creating a View (CREATE VIEW)	251
7.4.2	Removing a View (DROP VIEW)	253
7.4.3	View Resolution	254
7.4.4	Restrictions on Views	255
7.4.5	View Updatability	255
7.4.6	WITH CHECK OPTION	256
7.4.7	Advantages and Disadvantages of Views	258
7.4.8	View Materialization	260

7.5	Transactions	261
7.5.1	Immediate and Deferred Integrity Constraints	262
7.6	Discretionary Access Control	262
7.6.1	Granting Privileges to Other Users (GRANT)	264
7.6.2	Revoking Privileges from Users (REVOKE)	265
	<i>Chapter Summary</i>	267
	<i>Review Questions</i>	268
	<i>Exercises</i>	268
Chapter 8	Advanced SQL	271
8.1	The SQL Programming Language	272
8.1.1	Declarations	272
8.1.2	Assignments	273
8.1.3	Control Statements	274
8.1.4	Exceptions in PL/SQL	276
8.1.5	Cursors in PL/SQL	277
8.2	Subprograms, Stored Procedures, Functions, and Packages	280
8.3	Triggers	281
8.4	Recursion	287
	<i>Chapter Summary</i>	288
	<i>Review Questions</i>	289
	<i>Exercises</i>	289
Chapter 9	Object-Relational DBMSs	291
9.1	Advanced Database Applications	292
9.2	Weaknesses of RDBMSs	297
9.3	Storing Objects in a Relational Database	302
9.3.1	Mapping Classes to Relations	303
9.3.2	Accessing Objects in the Relational Database	304
9.4	Introduction to Object-Relational Database Systems	305
9.5	SQL:2011	308
9.5.1	Row Types	309
9.5.2	User-Defined Types	310
9.5.3	Subtypes and Supertypes	313
9.5.4	User-Defined Routines	314
9.5.5	Polymorphism	317
9.5.6	Reference Types and Object Identity	318
9.5.7	Creating Tables	318
9.5.8	Querying Data	321

9.5.9	Collection Types	323
9.5.10	Typed Views	326
9.5.11	Persistent Stored Modules	327
9.5.12	Triggers	327
9.5.13	Large Objects	330
9.5.14	Recursion	331
9.6	Object-Oriented Extensions in Oracle	331
9.6.1	User-Defined Data Types	332
9.6.2	Manipulating Object Tables	337
9.6.3	Object Views	338
9.6.4	Privileges	339
	<i>Chapter Summary</i>	340
	<i>Review Questions</i>	340
	<i>Exercises</i>	341

Part 3 Database Analysis and Design 343

Chapter 10	Database System Development Lifecycle	345
10.1	The Information Systems Lifecycle	346
10.2	The Database System Development Lifecycle	347
10.3	Database Planning	347
10.4	System Definition	350
10.4.1	User Views	350
10.5	Requirements Collection and Analysis	350
10.5.1	Centralized Approach	352
10.5.2	View Integration Approach	352
10.6	Database Design	354
10.6.1	Approaches to Database Design	355
10.6.2	Data Modeling	355
10.6.3	Phases of Database Design	356
10.7	DBMS Selection	359
10.7.1	Selecting the DBMS	359
10.8	Application Design	363
10.8.1	Transaction Design	364
10.8.2	User Interface Design Guidelines	365
10.9	Prototyping	367
10.10	Implementation	367
10.11	Data Conversion and Loading	368

10.12	Testing	368
10.13	Operational Maintenance	369
10.14	CASE Tools	370
	<i>Chapter Summary</i>	372
	<i>Review Questions</i>	373
	<i>Exercises</i>	374
Chapter 11	Database Analysis and the <i>DreamHome</i> Case Study	375
11.1	When Are Fact-Finding Techniques Used?	376
11.2	What Facts Are Collected?	377
11.3	Fact-Finding Techniques	378
	11.3.1 Examining Documentation	378
	11.3.2 Interviewing	378
	11.3.3 Observing the Enterprise in Operation	379
	11.3.4 Research	380
	11.3.5 Questionnaires	380
11.4	Using Fact-Finding Techniques: A Worked -Example	381
	11.4.1 The <i>DreamHome</i> Case Study—An Overview of the Current System	382
	11.4.2 The <i>DreamHome</i> Case Study—Database Planning	386
	11.4.3 The <i>DreamHome</i> Case Study—System Definition	392
	11.4.4 The <i>DreamHome</i> Case Study—Requirements Collection and Analysis	393
	11.4.5 The <i>DreamHome</i> Case Study—Database Design	401
	<i>Chapter Summary</i>	402
	<i>Review Questions</i>	402
	<i>Exercises</i>	402
Chapter 12	Entity–Relationship Modeling	405
12.1	Entity Types	406
12.2	Relationship Types	408
	12.2.1 Degree of Relationship Type	410
	12.2.2 Recursive Relationship	412
12.3	Attributes	413
	12.3.1 Simple and Composite Attributes	413
	12.3.2 Single-valued and Multi-valued Attributes	414
	12.3.3 Derived Attributes	414
	12.3.4 Keys	415
12.4	Strong and Weak Entity Types	417
12.5	Attributes on Relationships	418

12.6	Structural Constraints	419
12.6.1	One-to-One (1:1) Relationships	420
12.6.2	One-to-Many (1:*) Relationships	421
12.6.3	Many-to-Many (*:*) Relationships	422
12.6.4	Multiplicity for Complex Relationships	423
12.6.5	Cardinality and Participation Constraints	424
12.7	Problems with ER Models	426
12.7.1	Fan Traps	426
12.7.2	Chasm Traps	428
	<i>Chapter Summary</i>	430
	<i>Review Questions</i>	430
	<i>Exercises</i>	431
Chapter 13	Enhanced Entity–Relationship Modeling	433
13.1	Specialization/Generalization	434
13.1.1	Superclasses and Subclasses	434
13.1.2	Superclass/Subclass Relationships	435
13.1.3	Attribute Inheritance	436
13.1.4	Specialization Process	436
13.1.5	Generalization Process	437
13.1.6	Constraints on Specialization/Generalization	440
13.1.7	Worked Example of using Specialization/ Generalization to Model the Branch View of the <i>DreamHome</i> Case Study	441
13.2	Aggregation	445
13.3	Composition	446
	<i>Chapter Summary</i>	447
	<i>Review Questions</i>	448
	<i>Exercises</i>	448
Chapter 14	Normalization	451
14.1	The Purpose of Normalization	452
14.2	How Normalization Supports Database Design	453
14.3	Data Redundancy and Update Anomalies	454
14.3.1	Insertion Anomalies	455
14.3.2	Deletion Anomalies	455
14.3.3	Modification Anomalies	456
14.4	Functional Dependencies	456
14.4.1	Characteristics of Functional Dependencies	456
14.4.2	Identifying Functional Dependencies	460
14.4.3	Identifying the Primary Key for a Relation Using Functional Dependencies	463
14.5	The Process of Normalization	464

14.6	First Normal Form (1NF)	466
14.7	Second Normal Form (2NF)	470
14.8	Third Normal Form (3NF)	471
14.9	General Definitions of 2NF and 3NF	473
	<i>Chapter Summary</i>	475
	<i>Review Questions</i>	475
	<i>Exercises</i>	476
Chapter 15	Advanced Normalization	481
15.1	More on Functional Dependencies	482
	15.1.1 Inference Rules for Functional Dependencies	482
	15.1.2 Minimal Sets of Functional Dependencies	484
15.2	Boyce–Codd Normal Form (BCNF)	485
	15.2.1 Definition of BCNF	485
15.3	Review of Normalization Up to BCNF440	
15.4	Fourth Normal Form (4NF)	493
	15.4.1 Multi-Valued Dependency	494
	15.4.2 Definition of Fourth Normal Form	495
15.5	Fifth Normal Form (5NF)	495
	15.5.1 Lossless-Join Dependency	496
	15.5.2 Definition of Fifth Normal Form	496
	<i>Chapter Summary</i>	498
	<i>Review Questions</i>	498
	<i>Exercises</i>	499
Part 4	Methodology	501
Chapter 16	Methodology—Conceptual Database Design	503
16.1	Introduction to the Database Design Methodology	504
	16.1.1 What Is a Design Methodology?	504
	16.1.2 Conceptual, Logical, and Physical Database Design	505
	16.1.3 Critical Success Factors in Database Design	505
16.2	Overview of the Database Design Methodology	506
16.3	Conceptual Database Design Methodology	508
	Step 1: Build Conceptual Data Model	508
	<i>Chapter Summary</i>	524
	<i>Review Questions</i>	524
	<i>Exercises</i>	525

Chapter 17	Methodology—Logical Database Design for the Relational Model	527
17.1	Logical Database Design Methodology for the Relational Model	528
	Step 2: Build Logical Data Model	528
	<i>Chapter Summary</i>	556
	<i>Review Questions</i>	557
	<i>Exercises</i>	557
Chapter 18	Methodology—Physical Database Design for Relational Databases	561
18.1	Comparison of Logical and Physical Database Design	562
18.2	Overview of the Physical Database Design Methodology	563
18.3	The Physical Database Design Methodology for Relational Databases	564
	Step 3: Translate Logical Data Model for Target DBMS	564
	Step 4: Design File Organizations and Indexes	569
	Step 5: Design User Views	582
	Step 6: Design Security Mechanisms	582
	<i>Chapter Summary</i>	583
	<i>Review Questions</i>	584
	<i>Exercises</i>	584
Chapter 19	Methodology—Monitoring and Tuning the Operational System	585
19.1	Denormalizing and Introducing Controlled Redundancy	585
	Step 7: Consider the Introduction of Controlled Redundancy	585
19.2	Monitoring the System to Improve Performance	598
	Step 8: Monitor and Tune the Operational System	598
	<i>Chapter Summary</i>	602
	<i>Review Questions</i>	603
	<i>Exercises</i>	603
Part 5	Selected Database Issues	605
Chapter 20	Security and Administration	607
20.1	Database Security	608
	20.1.1 Threats	609

20.2	Countermeasures—Computer-Based Controls	611
20.2.1	Authorization	612
20.2.2	Access Controls	613
20.2.3	Views	616
20.2.4	Backup and Recovery	616
20.2.5	Integrity	617
20.2.6	Encryption	617
20.2.7	RAID (Redundant Array of Independent Disks)	618
20.3	Security in Microsoft Office Access DBMS	621
20.4	Security in Oracle DBMS	623
20.5	DBMSs and Web Security	627
20.5.1	Proxy Servers	628
20.5.2	Firewalls	628
20.5.3	Message Digest Algorithms and Digital Signatures	629
20.5.4	Digital Certificates	629
20.5.5	Kerberos	630
20.5.6	Secure Sockets Layer and Secure HTTP	630
20.5.7	Secure Electronic Transactions and Secure Transaction Technology	631
20.5.8	Java Security	632
20.5.9	ActiveX Security	634
20.6	Data Administration and Database Administration	634
20.6.1	Data Administration	635
20.6.2	Database Administration	636
20.6.3	Comparison of Data and Database Administration	636
	<i>Chapter Summary</i>	637
	<i>Review Questions</i>	638
	<i>Exercises</i>	638
Chapter 21	Professional, Legal, and Ethical Issues in Data Management	641
21.1	Defining Legal and Ethical Issues in IT	642
21.1.1	Defining Ethics in the Context of IT	642
21.1.2	The Difference Between Ethical and Legal Behavior	643
21.1.3	Ethical Behavior in IT	644
21.2	Legislation and Its Impact on the IT Function	645
21.2.1	Securities and Exchange Commission (SEC) Regulation National Market System (NMS)	645
21.2.2	The Sarbanes-Oxley Act, COBIT, and COSO	646
21.2.3	The Health Insurance Portability and Accountability Act	649
21.2.4	The European Union (EU) Directive on Data Protection of 1995	650
21.2.5	The United Kingdom's Data Protection Act of 1998	651

21.2.6	Access to Information Laws	652
21.2.7	International Banking—Basel II Accords	654
21.3	Establishing a Culture of Legal and Ethical	
	Data Stewardship	655
21.3.1	Developing an Organization-Wide Policy for Legal and Ethical Behavior	655
21.3.2	Professional Organizations and Codes of Ethics	656
21.3.3	Developing an Organization-Wide Policy for Legal and Ethical Behavior for <i>DreamHome</i>	659
21.4	Intellectual Property	660
21.4.1	Patent	661
21.4.2	Copyright	661
21.4.3	Trademark	662
21.4.4	Intellectual Property Rights Issues for Software	662
21.4.5	Intellectual Property Rights Issues for Data	664
	<i>Chapter Summary</i>	664
	<i>Review Questions</i>	665
	<i>Exercises</i>	666
Chapter 22	Transaction Management	667
22.1	Transaction Support	668
22.1.1	Properties of Transactions	671
22.1.2	Database Architecture	671
22.2	Concurrency Control	672
22.2.1	The Need for Concurrency Control	672
22.2.2	Serializability and Recoverability	675
22.2.3	Locking Methods	683
22.2.4	Deadlock	689
22.2.5	Timestamping Methods	692
22.2.6	Multiversion Timestamp Ordering	695
22.2.7	Optimistic Techniques	696
22.2.8	Granularity of Data Items	697
22.3	Database Recovery	700
22.3.1	The Need for Recovery	700
22.3.2	Transactions and Recovery	701
22.3.3	Recovery Facilities	704
22.3.4	Recovery Techniques	707
22.3.5	Recovery in a Distributed DBMS	709
22.4	Advanced Transaction Models	709
22.4.1	Nested Transaction Model	711
22.4.2	Sagas	712
22.4.3	Multilevel Transaction Model	713
22.4.4	Dynamic Restructuring	714
22.4.5	Workflow Models	715

22.5	Concurrency Control and Recovery in Oracle	716
22.5.1	Oracle's Isolation Levels	717
22.5.2	Multiversion Read Consistency	717
22.5.3	Deadlock Detection	719
22.5.4	Backup and Recovery	719
	<i>Chapter Summary</i>	722
	<i>Review Questions</i>	723
	<i>Exercises</i>	724
Chapter 23	Query Processing	727
23.1	Overview of Query Processing	729
23.2	Query Decomposition	732
23.3	Heuristical Approach to Query Optimization	736
23.3.1	Transformation Rules for the Relational Algebra Operations	736
23.3.2	Heuristical Processing Strategies	741
23.4	Cost Estimation for the Relational Algebra Operations	742
23.4.1	Database Statistics	742
23.4.2	Selection Operation ($S = \sigma_p(R)$)	743
23.4.3	Join Operation ($T = (R \bowtie_F S)$)	750
23.4.4	Projection Operation ($S = \Pi_{A_1, A_2, \dots, A_m}(R)$)	757
23.4.5	The Relational Algebra Set Operations ($T = R \cup S$, $T = R \cap S$, $T = R - S$)	759
23.5	Enumeration of Alternative Execution Strategies	760
23.5.1	Pipelining	761
23.5.2	Linear Trees	761
23.5.3	Physical Operators and Execution Strategies	762
23.5.4	Reducing the Search Space	764
23.5.5	Enumerating Left-Deep Trees	765
23.5.6	Semantic Query Optimization	766
23.5.7	Alternative Approaches to Query Optimization	767
23.5.8	Distributed Query Optimization	768
23.6	Query Processing and Optimization	768
23.6.1	New Index Types	771
23.7	Query Optimization in Oracle	772
23.7.1	Rule-Based and Cost-Based Optimization	772
23.7.2	Histograms	776
23.7.3	Viewing the Execution Plan	778
	<i>Chapter Summary</i>	779
	<i>Review Questions</i>	780
	<i>Exercises</i>	781

Part 6 Distributed DBMSs and Replication **783**

Chapter 24	Distributed DBMSs—Concepts and Design	785
24.1	Introduction	786
	24.1.1 Concepts	787
	24.1.2 Advantages and Disadvantages of DDBMSs	791
	24.1.3 Homogeneous and Heterogeneous DDBMSs	794
24.2	Overview of Networking	797
24.3	Functions and Architectures of a DDBMS	801
	24.3.1 Functions of a DDBMS	801
	24.3.2 Reference Architecture for a DDBMS	801
	24.3.3 Reference Architecture for a Federated MDBS	803
	24.3.4 Component Architecture for a DDBMS	804
24.4	Distributed Relational Database Design	805
	24.4.1 Data Allocation	806
	24.4.2 Fragmentation	807
24.5	Transparencies in a DDBMS	816
	24.5.1 Distribution Transparency	816
	24.5.2 Transaction Transparency	819
	24.5.3 Performance Transparency	822
	24.5.4 DBMS Transparency	824
	24.5.5 Summary of Transparencies in a DDBMS	824
24.6	Date's Twelve Rules for a DDBMS	825
	<i>Chapter Summary</i>	827
	<i>Review Questions</i>	828
	<i>Exercises</i>	828
Chapter 25	Distributed DBMSs—Advanced Concepts	831
25.1	Distributed Transaction Management	832
25.2	Distributed Concurrency Control	833
	25.2.1 Objectives	833
	25.2.2 Distributed Serializability	834
	25.2.3 Locking Protocols	834
25.3	Distributed Deadlock Management	837
25.4	Distributed Database Recovery	840
	25.4.1 Failures in a Distributed Environment	841
	25.4.2 How Failures Affect Recovery	842
	25.4.3 Two-Phase Commit (2PC)	842
	25.4.4 Three-Phase Commit (3PC)	849
	25.4.5 Network Partitioning	852
25.5	The X/Open Distributed Transaction Processing Model	854
25.6	Distributed Query Optimization	856

25.6.1 Data Localization	858
25.6.2 Distributed Joins	861
25.6.3 Global Optimization	862
25.7 Distribution in Oracle	866
25.7.1 Oracle's DDBMS Functionality	866
Chapter Summary	872
Review Questions	872
Exercises	873
Chapter 26 Replication and Mobile Databases	875
26.1 Introduction to Data Replication	876
26.1.1 Applications of Replication	877
26.1.2 Replication Model	878
26.1.3 Functional Model of Replication Protocols	879
26.1.4 Consistency	880
26.2 Replication Architecture	880
26.2.1 Kernel-Based Replication	880
26.2.2 Middleware-Based Replication	881
26.2.3 Processing of Updates	882
26.2.4 Propagation of Updates	884
26.2.5 Update Location (Data Ownership)	884
26.2.6 Termination Protocols	888
26.3 Replication Schemes	888
26.3.1 Eager Primary Copy	889
26.3.2 Lazy Primary Copy	894
26.3.3 Eager Update Anywhere	898
26.3.4 Lazy Update Anywhere	899
26.3.5 Update Anywhere with Uniform Total Order Broadcast	903
26.3.6 SI and Uniform Total Order Broadcast Replication	907
26.4 Introduction to Mobile Databases	913
26.4.1 Mobile DBMSs	915
26.4.2 Issues with Mobile DBMSs	916
26.5 Oracle Replication	929
26.5.1 Oracle's Replication Functionality	929
Chapter Summary	936
Review Questions	937
Exercises	937
Part 7 Object DBMSs	939
Chapter 27 Object-Oriented DBMSs—Concepts and Design	941
27.1 Next-Generation Database Systems	943

27.2	Introduction to OODBMSs	945
	27.2.1 Definition of Object-Oriented DBMSs	945
	27.2.2 Functional Data Models	946
	27.2.3 Persistent Programming Languages	951
	27.2.4 Alternative Strategies for Developing an OODBMS	953
27.3	Persistence in OODBMSs	954
	27.3.1 Pointer Swizzling Techniques	956
	27.3.2 Accessing an Object	959
	27.3.3 Persistence Schemes	961
	27.3.4 Orthogonal Persistence	962
27.4	Issues in OODBMSs	964
	27.4.1 Transactions	964
	27.4.2 Versions	965
	27.4.3 Schema Evolution	966
	27.4.4 Architecture	969
	27.4.5 Benchmarking	971
27.5	Advantages and Disadvantages of OODBMSs	974
	27.5.1 Advantages	974
	27.5.2 Disadvantages	976
27.6	Comparison of ORDBMS and OODBMS	978
27.7	Object-Oriented Database Design	979
	27.7.1 Comparison of Object-Oriented Data Modeling and Conceptual Data Modeling	979
	27.7.2 Relationships and Referential Integrity	980
	27.7.3 Behavioral Design	982
27.8	Object-Oriented Analysis and Design with UML	984
	27.8.1 UML Diagrams	985
	27.8.2 Usage of UML in the Methodology for Database Design	990
	<i>Chapter Summary</i>	992
	<i>Review Questions</i>	993
	<i>Exercises</i>	993
Chapter 28	Object-Oriented DBMSs—Standards and Systems	995
28.1	Object Management Group	996
	28.1.1 Background	996
	28.1.2 The Common Object Request Broker Architecture	999
	28.1.3 Other OMG Specifications	1004
	28.1.4 Model-Driven Architecture	1007
28.2	Object Data Standard ODMG 3.0, 1999	1007
	28.2.1 Object Data Management Group	1009
	28.2.2 The Object Model	1010

28.2.3	The Object Definition Language	1018
28.2.4	The Object Query Language	1021
28.2.5	Other Parts of the ODMG Standard	1027
28.2.6	Mapping the Conceptual Design to a Logical (Object-Oriented) Design	1030
28.3	ObjectStore	1031
28.3.1	Architecture	1031
28.3.2	Building an ObjectStore Application	1034
28.3.3	Data Definition in ObjectStore	1035
28.3.4	Data Manipulation in ObjectStore	1039
	<i>Chapter Summary</i>	1042
	<i>Review Questions</i>	1043
	<i>Exercises</i>	1043

Part 8 The Web and DBMSs **I045**

Chapter 29 Web Technology and DBMSs **I047**

29.1	Introduction to the Internet and the Web	1048
29.1.1	Intranets and Extranets	1050
29.1.2	e-Commerce and e-Business	1051
29.2	The Web	1052
29.2.1	HyperText Transfer Protocol	1053
29.2.2	HyperText Markup Language	1055
29.2.3	Uniform Resource Locators	1057
29.2.4	Static and Dynamic Web Pages	1058
29.2.5	Web Services	1058
29.2.6	Requirements for Web–DBMS Integration	1059
29.2.7	Advantages and Disadvantages of the Web–DBMS Approach	1060
29.2.8	Approaches to Integrating the Web and DBMSs	1064
29.3	Scripting Languages	1065
29.3.1	JavaScript and JScript	1065
29.3.2	VBScript	1066
29.3.3	Perl and PHP	1067
29.4	Common Gateway Interface (CGI)	1067
29.4.1	Passing Information to a CGI Script	1069
29.4.2	Advantages and Disadvantages of CGI	1071
29.5	HTTP Cookies	1072
29.6	Extending the Web Server	1073
29.6.1	Comparison of CGI and API	1074
29.7	Java	1074
29.7.1	JDBC	1078
29.7.2	SQLJ	1084

29.7.3	Comparison of JDBC and SQLJ	1084
29.7.4	Container-Managed Persistence (CMP)	1085
29.7.5	Java Data Objects (JDO)	1089
29.7.6	JPA (Java Persistence API)	1096
29.7.7	Java Servlets	1104
29.7.8	JavaServer Pages	1104
29.7.9	Java Web Services	1105
29.8	Microsoft's Web Platform	1107
29.8.1	Universal Data Access	1108
29.8.2	Active Server Pages and ActiveX Data Objects	1109
29.8.3	Remote Data Services	1110
29.8.4	Comparison of ASP and JSP	1113
29.8.5	Microsoft .NET	1113
29.8.6	Microsoft Web Services	1118
29.9	Oracle Internet Platform	1119
29.9.1	Oracle WebLogic Server	1120
29.9.2	Oracle Metadata Repository	1121
29.9.3	Oracle Identity Management	1121
29.9.4	Oracle Portal	1122
29.9.5	Oracle WebCenter	1122
29.9.6	Oracle Business Intelligence (BI) Discoverer	1122
29.9.7	Oracle SOA (Service-Oriented Architecture) Suite	1123
	<i>Chapter Summary</i>	1126
	<i>Review Questions</i>	1127
	<i>Exercises</i>	1127
Chapter 30	Semistructured Data and XML	1129
30.1	Semistructured Data	1130
30.1.1	Object Exchange Model (OEM)	1132
30.1.2	Lore and Lorel	1133
30.2	Introduction to XML	1137
30.2.1	Overview of XML	1140
30.2.2	Document Type Definitions (DTDs)	1142
30.3	XML-Related Technologies	1145
30.3.1	DOM and SAX Interfaces	1146
30.3.2	Namespaces	1147
30.3.3	XSL and XSLT	1147
30.3.4	XPath (XML Path Language)	1148
30.3.5	XPointer (XML Pointer Language)	1149
30.3.6	XLink (XML Linking Language)	1150
30.3.7	XHTML	1150
30.3.8	Simple Object Access Protocol (SOAP)	1151
30.3.9	Web Services Description Language (WSDL)	1152

30.3.10 Universal Discovery, Description, and Integration (UDDI)	1152
30.3.11 JSON (JavaScript Object Notation)	1154
30.4 XML Schema	1156
30.4.1 Resource Description Framework (RDF)	1162
30.5 XML Query Languages	1166
30.5.1 Extending Lore and Lorel to Handle XML	1167
30.5.2 XML Query Working Group	1168
30.5.3 XQuery—A Query Language for XML	1169
30.5.4 XML Information Set	1179
30.5.5 XQuery 1.0 and XPath 2.0 Data Model (XDM)	1180
30.5.6 XQuery Update Facility 1.0	1186
30.5.7 Formal Semantics	1188
30.6 XML and Databases	1196
30.6.1 Storing XML in Databases	1196
30.6.2 XML and SQL	1199
30.6.3 Native XML Databases	1213
30.7 XML in Oracle	1214
<i>Chapter Summary</i>	<i>1217</i>
<i>Review Questions</i>	<i>1219</i>
<i>Exercises</i>	<i>1220</i>

Part 9 Business Intelligence

1221

Chapter 31 Data Warehousing Concepts

1223

31.1 Introduction to Data Warehousing	1224
31.1.1 The Evolution of Data Warehousing	1224
31.1.2 Data Warehousing Concepts	1225
31.1.3 Benefits of Data Warehousing	1226
31.1.4 Comparison of OLTP Systems and Data Warehousing	1226
31.1.5 Problems of Data Warehousing	1228
31.1.6 Real-Time Data Warehouse	1230
31.2 Data Warehouse Architecture	1231
31.2.1 Operational Data	1231
31.2.2 Operational Data Store	1231
31.2.3 ETL Manager	1232
31.2.4 Warehouse Manager	1232
31.2.5 Query Manager	1233
31.2.6 Detailed Data	1233
31.2.7 Lightly and Highly Summarized Data	1233
31.2.8 Archive/Backup Data	1233
31.2.9 Metadata	1234
31.2.10 End-User Access Tools	1234

31.3	Data Warehousing Tools and Technologies	1235
31.3.1	Extraction, Transformation, and Loading (ETL)	1236
31.3.2	Data Warehouse DBMS	1237
31.3.3	Data Warehouse Metadata	1240
31.3.4	Administration and Management Tools	1242
31.4	Data Mart	1242
31.4.1	Reasons for Creating a Data Mart	1243
31.5	Data Warehousing and Temporal Databases	1243
31.5.1	Temporal Extensions to the SQL Standard	1246
31.6	Data Warehousing Using Oracle	1248
31.6.1	Warehouse Features in Oracle 11g	1251
31.6.2	Oracle Support for Temporal Data	1252
	<i>Chapter Summary</i>	1253
	<i>Review Questions</i>	1254
	<i>Exercises</i>	1255
Chapter 32	Data Warehousing Design	1257
32.1	Designing a Data Warehouse Database	1258
32.2	Data Warehouse Development Methodologies	1258
32.3	Kimball's Business Dimensional Lifecycle	1260
32.4	Dimensionality Modeling	1261
32.4.1	Comparison of DM and ER models	1264
32.5	The Dimensional Modeling Stage of Kimball's Business Dimensional Lifecycle	1265
32.5.1	Create a High-Level Dimensional Model (Phase I)	1265
32.5.2	Identify All Dimension Attributes for the Dimensional Model (Phase II)	1270
32.6	Data Warehouse Development Issues	1273
32.7	Data Warehousing Design Using Oracle	1274
32.7.1	Oracle Warehouse Builder Components	1274
32.7.2	Using Oracle Warehouse Builder	1275
32.7.3	Warehouse Builder Features in Oracle 11g	1279
	<i>Chapter Summary</i>	1280
	<i>Review Questions</i>	1281
	<i>Exercises</i>	1282
Chapter 33	OLAP	1285
33.1	Online Analytical Processing	1286
33.1.1	OLAP Benchmarks	1287