

# Databases Illuminated

Second Edition

Catherine Ricardo, Iona College



JONES & BARTLETT  
LEARNING

*World Headquarters*

Jones & Bartlett Learning  
40 Tall Pine Drive  
Sudbury, MA 01776  
978-443-5000  
info@jblearning.com  
www.jblearning.com

Jones & Bartlett Learning Canada  
6339 Ormindale Way  
Mississauga, Ontario L5V 1J2  
Canada

Jones & Bartlett Learning International  
Barb House, Barb Mews  
London W6 7PA  
United Kingdom

Jones & Bartlett Learning books and products are available through most bookstores and online booksellers. To contact Jones & Bartlett Learning directly, call 800-832-0034, fax 978-443-8000, or visit our website, [www.jblearning.com](http://www.jblearning.com).

Substantial discounts on bulk quantities of Jones & Bartlett Learning publications are available to corporations, professional associations, and other qualified organizations. For details and specific discount information, contact the special sales department at Jones & Bartlett Learning via the above contact information or send an email to [specialsales@jblearning.com](mailto:specialsales@jblearning.com).

Copyright © 2012 by Jones & Bartlett Learning, LLC

All rights reserved. No part of the material protected by this copyright may be reproduced or utilized in any form, electronic or mechanical, including photocopying, recording, or by any information storage and retrieval system, without written permission from the copyright owner.

**Production Credits**

Publisher: Cathleen Sether  
Senior Acquisitions Editor: Timothy Anderson  
Senior Editorial Assistant: Stephanie Sguigna  
Production Director: Amy Rose  
Senior Marketing Manager: Andrea DeFronzo  
Associate Marketing Manager: Lindsay White  
V.P., Manufacturing and Inventory Control: Therese Connell  
Composition: Northeast Compositors, Inc.  
Cover and Title Page Design: Kristin E. Parker  
Cover Image: © ErickN/Shutterstock, Inc.  
Printing and Binding: Malloy, Inc.  
Cover Printing: Malloy, Inc.

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

Ricardo, Catherine M.  
Databases illuminated / Catherine M. Ricardo. — 2nd ed.  
p. cm.  
Includes bibliographical references and index.  
ISBN 978-1-4496-0600-8 (casebound)  
1. Database management. I. Title.  
QA76.9.D3R514 2011  
005.74—dc22

2010052808

6048

Printed in the United States of America

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

## **Dedication**

---

To my beloved husband, Henry



# Contents

Preface xix

Acknowledgments xxiv

## 1. Introductory Database Concepts 1

---

Chapter Objectives 1

1.1 Uses of Databases 2

1.2 A Sample Database 3

1.3 The Integrated Database Environment 7

1.4 Roles in the Integrated Database Environment 9

1.5 Advantages of the Integrated Database Approach 12

1.6 Historical Developments in Information Systems 15

1.7 Chapter Summary 21

Exercises 22

## 2. Database Planning and Database Architecture 25

---

Chapter Objectives 25

2.1 Data as a Resource 26

2.2 Characteristics of Data 27

2.2.1 Data and Information 27

2.2.2 Levels of Discussing Data 27

2.2.3 Data Sublanguages 29

<b>2.3</b>	<b>Stages in Database Design</b>	<b>31</b>
<b>2.4</b>	<b>Design Tools</b>	<b>34</b>
2.4.1	Data Dictionary	35
2.4.2	Project Management Software	37
<b>2.5</b>	<b>Database Administration</b>	<b>37</b>
2.5.1	Planning and Design	38
2.5.2	Developing the Database	39
2.5.3	Database Management	40
<b>2.6</b>	<b>The Three-Level Database Architecture</b>	<b>41</b>
2.6.1	External Views	42
2.6.2	Logical and Conceptual Models	44
2.6.3	Internal Model	46
2.6.4	Data Independence	49
<b>2.7</b>	<b>Overview of Data Models</b>	<b>50</b>
2.7.1	The Entity-Relationship Model	51
2.7.2	Relational Model	53
2.7.3	Object-Oriented Model	54
2.7.4	Object-Relational Model	55
2.7.5	Semistructured Data Model	55
<b>2.8</b>	<b>Chapter Summary</b>	<b>56</b>
	<b>Exercises</b>	<b>58</b>

## **3. The Entity-Relationship Model 61**

### **Chapter Objectives 61**

<b>3.1</b>	<b>Purpose of the E-R Model</b>	<b>62</b>
<b>3.2</b>	<b>Entities</b>	<b>63</b>
<b>3.3</b>	<b>Attributes</b>	<b>64</b>
3.3.1	Domains	65
3.3.2	Null Values	66

3.3.3	Multivalued Attributes	66
3.3.4	Composite Attributes	66
3.3.5	Derived Attributes	67
<b>3.4</b>	<b>Keys</b>	<b>68</b>
3.4.1	Superkeys	68
3.4.2	Candidate Keys	69
3.4.3	Primary Keys	69
<b>3.5</b>	<b>Relationships</b>	<b>70</b>
3.5.1	Degree of Relationships	71
3.5.2	Attributes of Relationship Sets	74
3.5.3	Cardinality of a Relationship	74
3.5.4	Showing Cardinalities on an E-R Diagram	76
3.5.5	Participation Constraints	78
<b>3.6</b>	<b>Roles</b>	<b>79</b>
<b>3.7</b>	<b>Existence Dependency and Weak Entities</b>	<b>80</b>
<b>3.8</b>	<b>A Sample E-R Diagram</b>	<b>82</b>
<b>3.9</b>	<b>Chapter Summary</b>	<b>85</b>
	<b>Exercises</b>	<b>87</b>

## **4. The Relational Model 93**

### **Chapter Objectives 93**

<b>4.1</b>	<b>Advantages of the Relational Model</b>	<b>94</b>
<b>4.2</b>	<b>Relational Data Structures</b>	<b>95</b>
4.2.1	Tables	95
4.2.2	Mathematical Relations	97
4.2.3	Database Relations and Tables	99
4.2.4	Properties of Relations	100
4.2.5	Degree and Cardinality	100
4.2.6	Relation Keys	101

- 4.3 Integrity Constraints: domain, key, foreign key, general constraints 103**
- 4.4 Representing Relational Database Schemas 104**
- 4.5 Relational Data Manipulation Languages 105**
  - 4.5.1 Categories of DMLs 105
  - 4.5.2 Relational Algebra 106
- 4.6 Views 115**
- 4.7 Mapping an E-R Model to a Relational Schema 117**
- 4.8 Chapter Summary 123**
  - Exercises 124**

## **5. Relational Database Management Systems and SQL 129**

### **Chapter Objectives 129**

- 5.1 Brief History of SQL in Relational Database Systems 130**
- 5.2 Architecture of a Relational Database Management System 131**
- 5.3 Defining the Database: SQL DDL 133**
  - 5.3.1 CREATE TABLE 133
  - 5.3.2 CREATE INDEX 139
  - 5.3.3 ALTER TABLE, RENAME TABLE 141
  - 5.3.4 DROP Statements 142
- 5.4 Manipulating the Database: SQL DML 143**
  - 5.4.1 Introduction to the SELECT Statement 143
  - 5.4.2 SELECT Using Multiple Tables 149
  - 5.4.3 SELECT with Other Operators 159
  - 5.4.4 Operators for Updating: UPDATE, INSERT, DELETE 165
  - 5.4.5 Creating and Using Views 173
- 5.5 Active Databases 178**
  - 5.5.1 Enabling and Disabling Constraints 178
  - 5.5.2 SQL Triggers 179

<b>5.6</b>	<b>Using COMMIT and ROLLBACK Statements</b>	<b>185</b>
<b>5.7</b>	<b>SQL Programming</b>	<b>186</b>
5.7.1	SQL PSMs	186
5.7.2	Embedded SQL	195
5.7.3	Dynamic SQL	199
5.7.4	APIs	200
<b>5.8</b>	<b>The System Catalog</b>	<b>201</b>
<b>5.9</b>	<b>Chapter Summary</b>	<b>203</b>
	<b>Exercises</b>	<b>205</b>

## **6. Normalization 211**

### **Chapter Objectives 211**

<b>6.1</b>	<b>Objectives of Normalization</b>	<b>212</b>
<b>6.2</b>	<b>Insertion, Update, and Deletion Anomalies</b>	<b>213</b>
<b>6.3</b>	<b>Functional Dependency</b>	<b>214</b>
<b>6.4</b>	<b>Superkeys, Candidate Keys, and Primary Keys</b>	<b>217</b>
<b>6.5</b>	<b>Normalization Using Primary Keys</b>	<b>218</b>
6.5.1	First Normal Form	219
6.5.2	Full Functional Dependency and Second Normal Form	222
6.5.3	Transitive Dependency and Third Normal Form	225
6.5.4	Boyce-Codd Normal Form	227
6.5.5	Comprehensive Example of Functional Dependencies	230
<b>6.6</b>	<b>Properties of Relational Decompositions</b>	<b>234</b>
6.6.1	Attribute Preservation	234
6.6.2	Dependency Preservation	235
6.6.3	Lossless Decomposition	235
6.6.4	Decomposition Algorithm for Boyce-Codd Normal Form with Lossless Join	238

- 6.7 The Normalization Process 238**
  - 6.7.1 Analysis 239
  - 6.7.2 Synthesis 239
  - 6.7.3 Normalization from an Entity-Relationship Diagram 239
- 6.8 When to Stop Normalizing 240**
- 6.9 Chapter Summary 241**
  - Exercises 242

## **7. The Object-Oriented Model 249**

### **Chapter Objectives 249**

- 7.1 Rationale for the Object-Oriented Data Model 250**
- 7.2 Object-Oriented Data Concepts 251**
  - 7.2.1 Objects and Literals 251
  - 7.2.2 Classes 253
  - 7.2.3 Class Hierarchies and Inheritance 255
  - 7.2.4 Object Identity 258
- 7.3 Object-Oriented Data Modeling Using UML 259**
- 7.4 The ODMG Model and ODL 263**
  - 7.4.1 Class Declarations 263
  - 7.4.2 Extent 263
  - 7.4.3 Attributes 263
  - 7.4.4 Relationships 268
  - 7.4.5 Methods 269
  - 7.4.6 Classes and Inheritance 269
  - 7.4.7 *N*-ary Relationships and *M*:*N* Relationships with Attributes 270
  - 7.4.8 Keys 270
- 7.5 Object Query Language 271**
- 7.6 Developing an OO Database 275**
- 7.7 Developing an OO Database in Caché 277**

- 7.8 Chapter Summary 281
- Exercises 282

## 8. The Enhanced Entity-Relationship Model and the Object-Relational Model 287

### Chapter Objectives 287

- 8.1 Rationale for Extending the E-R Model 288
- 8.2 Generalization and Specialization 288
  - 8.2.1 Specialization 289
  - 8.2.2 Generalization 290
  - 8.2.3 Generalization Constraints—Disjointness, Completeness, Definition Method 292
  - 8.2.4 Multiple Hierarchies and Inheritance 295
- 8.3 Union 297
- 8.4 Using (*min,max*) Notation for Cardinality and Participation 301
- 8.5 A Sample EE-R Diagram 302
- 8.6 Mapping the EE-R Model to a Relational Model 302
  - 8.6.1 Summary of E-R to Strictly Relational Mapping Concepts 304
  - 8.6.2 Mapping EE-R Class Hierarchies to Relational Tables 304
  - 8.6.3 Mapping Unions 306
- 8.7 Extending the Relational Model 306
  - 8.7.1 New Fundamental Data Types 307
  - 8.7.2 Collection Types 308
  - 8.7.3 User-Defined Data Types (UDT) 313
  - 8.7.4 Reference Types 318
  - 8.7.5 Type Hierarchies in Standard SQL 322
  - 8.7.6 Type Hierarchies in Oracle 323
  - 8.7.7 Nested Tables in Oracle 326
  - 8.7.8 Oracle Object Views 328
- 8.8 Converting an EE-R Diagram to an Object-Relational Database Model 330

**8.9 Chapter Summary 331****Exercises 333****9. Introduction to Database Security 337****Chapter Objectives 337****9.1 Issues in Database Security 338**

9.1.1 Accidental Security Threats 339

9.1.2 Deliberate Security Threats 339

**9.2 Physical Security and User Authentication 340****9.3 Authorization 341****9.4 Access Control 342****9.5 Using Views for Access Control 343****9.6 Security Logs and Audit Trails 343****9.7 Encryption 344****9.8 SQL Authorization Language 348****9.9 Security in Oracle 352**

9.9.1 The Oracle Enterprise Manager Database Control Facility 352

9.9.2 Security Control Using SQL\*Plus 353

9.9.3 Object Privileges 353

9.9.4 System Privileges 354

9.9.5 Roles 355

**9.10 Statistical Database Security 355****9.11 Database Security and the Internet 356**

9.11.1 Proxy Servers 356

9.11.2 Firewalls 357

9.11.3 Digital Signatures 357

9.11.4 Certification Authorities 357

**9.12 Chapter Summary 359****Exercises 360**

## **10. Transaction Management 363**

### **Chapter Objectives 363**

#### **10.1 Properties of Transactions 364**

#### **10.2 Need for Concurrency Control 367**

#### **10.3 Serializability 372**

#### **10.4 Locking 377**

##### 10.4.1 Deadlock 378

##### 10.4.2 Two-Phase Locking Protocol 381

##### 10.4.3 Levels of Locking 384

#### **10.5 Timestamping 386**

##### 10.5.1 Basic Timestamping Protocol 386

##### 10.5.2 Thomas' Write Rule 388

##### 10.5.3 Multiversion Timestamping 389

#### **10.6 Optimistic Techniques 391**

#### **10.7 Need for Recovery 393**

#### **10.8 Recovery Techniques 394**

##### 10.8.1 Deferred Update Protocol 395

##### 10.8.2 Checkpoints 397

##### 10.8.3 Immediate Update Protocol 398

##### 10.8.4 Shadow Paging 399

##### 10.8.5 Overview of ARIES Recovery Algorithm 399

#### **10.9 Transaction Management in Oracle 400**

#### **10.10 Chapter Summary 401**

##### **Exercises 404**

## **11. Relational Query Optimization 409**

### **Chapter Objectives 409**

#### **11.1 Query Processing and Query Optimization 410**

<b>11.2</b>	<b>Algebraic Techniques for Query Transformation</b>	<b>412</b>
11.2.1	The Query Tree	412
11.2.2	An SQL Query and Its Relational Algebra Translation	413
11.2.3	Performing SELECT Operations Early	414
11.2.4	Evaluating Conjunctive Conditions	415
11.2.5	Performing PROJECT Early	418
11.2.6	Equivalence of Algebraic Operations	419
11.2.7	Heuristics for Query Optimization	423
<b>11.3</b>	<b>Processing Techniques and Cost Estimation</b>	<b>424</b>
11.3.1	Cost Factors	424
11.3.2	Cost of Processing Selects	427
11.3.3	Processing Joins	430
11.3.4	Processing Other Operations	436
<b>11.4</b>	<b>Pipelining</b>	<b>440</b>
<b>11.5</b>	<b>Query Optimization in Oracle</b>	<b>440</b>
<b>11.6</b>	<b>Chapter Summary</b>	<b>441</b>
	<b>Exercises</b>	<b>442</b>

## **12. Distributed Databases 449**

### **Chapter Objectives 449**

<b>12.1</b>	<b>Rationale for Distribution</b>	<b>450</b>
<b>12.2</b>	<b>Architectures for a Distributed System</b>	<b>451</b>
12.2.1	Distributed Processing Using a Centralized Database	452
12.2.2	Client-Server Systems	452
12.2.3	Parallel Databases	454
12.2.4	Distributed Databases	458
<b>12.3</b>	<b>Components of a Distributed Database System</b>	<b>459</b>
<b>12.4</b>	<b>Data Placement</b>	<b>461</b>
<b>12.5</b>	<b>Transparency</b>	<b>467</b>

- 12.6 Transaction Control for Distributed Databases 469**
  - 12.6.1 Concurrency Control 469
  - 12.6.2 Recovery 475
- 12.7 Distributed Query Processing 479**
  - 12.7.1 Steps in Distributed Query Processing 480
  - 12.7.2 Estimating Data Communications Costs 483
  - 12.7.3 The Semijoin Operation 485
- 12.8 Chapter Summary 487**
  - Exercises 489**

## **13. Databases and the Internet 493**

---

### **Chapter Objectives 493**

- 13.1 Introduction 494**
- 13.2 Fundamental Concepts of the Internet and the World Wide Web 495**
  - 13.2.1 Origins of the World Wide Web 495
  - 13.2.2 Browsers, Links, and URIs 496
  - 13.2.3 HTTP 496
  - 13.2.4 HTML 498
  - 13.2.5 XML 499
- 13.3 Tiered Architectures 507**
  - 13.3.1 Single-Tier Architecture 507
  - 13.3.2 Two-Tier Architecture 510
  - 13.3.3 Three-Tier Architecture 512
- 13.4 Web Programming with Oracle 518**
  - 13.4.1 PL/SQL Web Toolkit 518
  - 13.4.2 PL/SQL Server Pages (PSP) 521
- 13.5 JDBC 523**

**13.6 The Semi-Structured Data Model 526**

13.6.1 Graphical Representation 526

13.6.2 XML Parsers 530

13.6.3 XML Data Manipulation 531

**13.7 XML and Relational Databases 537****13.8 XML in Oracle 538**

13.8.1 XML DB 538

13.8.2 Oracle XML Developer's Kits 545

**13.9 Chapter Summary 546****Exercises 549****14. Social and Ethical Issues 553****Chapter Objectives 553****14.1 Computerization and Society 554**

14.1.1 Ethical Issues 555

14.1.2 Ethical Standards for Computing Professionals 556

**14.2 Intellectual Property 556**

14.2.1 Definition of Intellectual Property 557

14.2.2 Legal Protections for Intellectual Property 559

14.2.3 Intellectual Property Protection for Software 571

**14.3 Databases and Privacy Issues 573**

14.3.1 Privacy and Security 574

14.3.2 Privacy as a Human Right 575

**14.4 Human Factors 584**

14.4.1 Human Factors in Software Development 584

14.4.2 The Human-Database Interface 586

14.4.3 Usability Testing for Database Applications 587

**14.5 Chapter Summary 588****Exercises 589**

## **15. Data Warehouses and Data Mining 593**

### **Chapter Objectives 593**

- 15.1 Origins of Data Warehouses 594**
- 15.2 Operational Databases and Data Warehouses 595**
- 15.3 Architecture of a Data Warehouse 596**
- 15.4 Developing a Data Warehouse 598**
  - 15.4.1 Top-Down Method 598
  - 15.4.2 Bottom-Up Method 604
  - 15.4.3 Comparison of Methods 607
- 15.5 Data Models for Data Warehouses 607**
  - 15.5.1 ROLAP and MOLAP 607
  - 15.5.2 Data Cubes and Hypercubes 608
  - 15.5.3 Star Schema 611
  - 15.5.4 Snowflake Schema 613
  - 15.5.5 Columnar Data 613
- 15.6 Data Warehouse Queries and SQL 613**
  - 15.6.1 Aggregate Functions 614
  - 15.6.2 RANK Functions 615
  - 15.6.3 CUBE and ROLLUP Functions 615
  - 15.6.4 Analytic Functions 616
- 15.7 Optimization and Index Techniques 616**
- 15.8 Views and View Materialization 619**
- 15.9 Data Warehouse Vendors 622**
- 15.10 Data Mining 623**
- 15.11 Purpose of Data Mining 624**
- 15.12 Types of Knowledge Discovered 625**
- 15.13 Models and Methods Used 627**
  - 15.13.1 Data Mining Process Model 627

15.13.2	Regression	629
15.13.3	Decision Trees	630
15.13.4	Artificial Neural Networks	631
15.13.5	Clustering	634
15.13.6	Genetic Algorithms	635
<b>15.14</b>	<b>Applications of Data Mining</b>	<b>635</b>
15.14.1	Retailing	635
15.14.2	Banking and Finance	636
15.14.3	Manufacturing	636
15.14.4	Science and Medicine	636
15.14.5	Homeland Security	638
15.14.6	Search Engines	638
<b>15.15</b>	<b>Data Mining Software Providers</b>	<b>638</b>
<b>15.16</b>	<b>Chapter Summary</b>	<b>639</b>
	<b>Exercises</b>	<b>641</b>
<b>Index</b>		<b>645</b>

# Preface

## Purpose of This Book

Since the publication of the first edition of this book in 2004, databases have become an even more central part of the framework of modern computer systems, and the study of database systems, design, and management is an essential part of the education of computer science and information science students. A database course should provide a strong theoretical background, practice in database design, and experience creating and developing a working database. This book is designed to help students integrate theoretical material with practical knowledge, using an approach that has a firm theoretical basis applied to practical database implementation.

## Structure

Theoretical foundations are presented early and the concepts are used repeatedly throughout the book, including the chapters that deal with implementation. Logical database design is given full consideration. The entity-relationship model is introduced early and then mapped to the relational model. Relational normalization is studied in detail, and many examples of the normalization process are discussed. The object-oriented model is presented using UML as a vehicle for logical design. The enhanced entity-relationship model is presented, and mapped to both relational and object-relational models. XML and the semi-structured data model are introduced in Chapter 13. A continuing example of a university database is incorporated throughout the text, to illustrate concepts and techniques and to provide both continuity and contrast. Other

examples are presented as needed. Purely relational, object-relational, and object-oriented database systems are described and used for implementation of the examples. Details of database management systems are described so that students can learn the specifics of these real-life systems, down to the physical implementation level. Microsoft® Access® is used initially in the examples, but Oracle® is introduced as the material is developed. However, the examples are suitable for use with any relational or object-relational DBMS. InterSystems Caché® DBMS is used to illustrate object-oriented databases.

### **Purpose of the Sample Project and Student Projects**

The various approaches are integrated in a sample project that is a unique feature of this book. Starting at the end of the first chapter, students can visit the Student Companion Website to see a sample project that is developed as the book progresses. This project is independent of the examples that are presented in the chapter text or in exercises. The design begins with a description of the application: a database needed by an art gallery. At the end of the first chapter, students see how the information needs are specified. After the students study planning techniques in Chapter 2, they see how to create a complete user-oriented data dictionary and how other planning tools could be used. When they learn about logical models, they see the step-by-step development of an E-R diagram at the end of Chapter 3. They see the steps involved in mapping the diagram to the relational model after Chapter 4, and normalizing the model after Chapter 6. Complete details of creating and manipulating a purely relational Oracle database using SQL are presented for the relational model after Chapter 5. The E-R diagram is expanded to a complete EE-R diagram and mapped to an object-relational database after Chapter 8, using the object-relational features of Oracle. The EE-R diagram is transformed into a UML diagram after Chapter 7, which includes the design and creation of an object-oriented database using ODL and also InterSystems Caché DBMS. The sample project is extended after Chapter 12 to show a distributed design that might be used in a client-server environment or a true distributed environment. Details of creating a simple website for the art gallery are provided after Chapter 13 using PL/SQL, JDBC, Caché, and XML. Thus, every important technique of planning, design, and implementation is illustrated using real-life systems. The sample pro-

ject section is always connected to continuing student projects, which require students to emulate the steps of the sample.

From my experiences assigning similar projects while teaching database courses, I have found that students learn most effectively by creating a working database and developing their own continuing projects to incorporate new concepts, and that they benefit by seeing samples as they progress. Such realistic experiences provide insights that no amount of lecturing can produce. Later chapters deal with topics of database security, concurrency control, recovery techniques, query optimization, distributed databases, social and ethical issues, data warehouses, and data mining. Online appendices cover physical data organization, the network model, the hierarchical model, and other specialized material.

## Learning Features

The writing style is conversational. Each chapter begins with a statement of learning objectives. Examples and applications are presented throughout the text. Illustrations are used both to clarify the material and to vary the presentation. Exercises, including lab exercises where appropriate, appear at the end of each chapter and on the Student Companion Website, with solutions provided in the Instructor's Guide. The sample project on the Student Companion Website is an important part of the text, providing an application of the material just presented. The student projects that follow the sample are introduced after the first chapter, and students are expected to choose one, or have one assigned, and to develop that project to parallel the sample as they progress. Student projects may be done individually or in a group. Access, Oracle, and Caché implementations of the example databases used in the text, and of the sample project, are available on the book's website; solutions for the student projects, since they are intended to be used as assignments, are not included there. Resources for student laboratory exercises are also available on the website. Chapter summaries are included in the text to provide a rapid review or preview of the material and to help students understand the relative importance of the concepts presented. The Instructor's Guide contains Microsoft® PowerPoint® presentations for each chapter, solutions for the student projects, copies of figures, full statements of objectives for each chapter, alternative student projects, quizzes, chapter tests, comprehensive examinations for multiple chapters, and solutions to exercises.

## Audience

The material is suitable for junior or senior computer science majors or information science majors with a good technical background. Students should have completed at least one year of programming, including data structures. The book could also be used as an introductory database text for graduate students or for self study.

## Mapping to Curriculum Guidelines

Although this book was based on the author's experiences, it fits the ACM-IEEE Curriculum Guidelines well. It covers the topics and supports the learning objectives listed in the IS 2010 Curriculum Guidelines for IS 2010.2, Data and Information Management. It also supports the guidelines of ACM Computer Science Curriculum 2008, An Interim Revision of CS2001. This revision retains the basic structure and guidelines of the 2001 recommendations, while including some new trends that have emerged since the first report. These include more attention to security, concurrency, and net-centric computing, all of which are addressed in this text. CC2001 provides a model database course, CS270T, for the topic-based curriculum. The course includes three areas, Human-Computer Interaction (HCI), Information Management (IM), and Social and Professional Issues (SP). The units included in CS270T are listed below, along with the corresponding chapters in the book.

IM1 Information models and systems (3 core hours)—Chapters 1, 2

IM2 Database systems (3 core hours)—Chapter 2

IM3 Data modeling (4 core hours)—Sections 2.7, 13.4; Chapters 3, 4, 7, 8, Appendices B, C

IM4 Relational databases (5 hours)—Chapter 4

IM5 Database query languages (4 hours)—Sections 4.5, 7.5, 8.7; Chapters 6, 11, 13; Appendices B, C

IM6 Relational database design (4 hours)—Chapter 6

IM7 Transaction processing (3 hours)—Chapter 10

IM8 Distributed databases (3 hours)—Chapters 12, 13

IM9 Physical database design (3 hours)—Appendix A

HCI Foundations of human–computer interaction (2 core hours)—  
Chapter 14

SP6 Intellectual property (3 core hours)—Chapter 14

SP7 Privacy and civil liberties (2 core hours)—Chapters 9, 14

Elective topics (1 hour)—Chapter 15

In addition to the text’s coverage, the sample project and student projects also provide practice in data modeling, database design, human–computer interaction, relational databases, database query languages, distributed databases, web-based databases, and physical database design. Some aspects of security, privacy, and civil liberties are discussed in the sample project, and similar issues will arise and should be treated in the student projects.

# Acknowledgments

Many people have offered useful comments, advice, and encouragement during the writing of this book. I am grateful to students and colleagues throughout the world who used the first edition and provided comments that have helped shape this book, especially Mary Courtney, Pace University and Stephen D'Alessio and Smiljana Petrovic, Iona College.

For his encouragement and support for this project, I am grateful to Tim Anderson, Senior Acquisitions Editor at Jones & Bartlett Learning. I would also like to thank the editorial and production staff at Jones & Bartlett Learning, especially Amy Rose.

Finally I would like to thank my family, especially my husband, Henry, for his love, patience, and understanding and for his active help in critiquing and proofreading the manuscript.