


数据库系统基础教程

(英文版·第3版)

A First Course in DATABASE SYSTEMS

THIRD EDITION

Jeffrey D. Ullman
Jennifer Widom



Jeffrey D. Ullman
斯坦福大学
(美) Jennifer Widom 著
斯坦福大学



机械工业出版社
China Machine Press

数据库系统基础教程

(英文版·第3版)

A First Course in Database Systems

(Third Edition)

本书由美国斯坦福大学知名计算机科学家Jeffrey Ullman和Jennifer Widom合作编写。书中介绍了核心DBMS概念、理论和模型，描述了如何使用抽象语言和SQL查询与更新DBMS。在介绍了SQL扩展内容(包括嵌入式SQL程序设计和对象关系特征)后，又介绍了使用XML的系统。设计语言包括XML模式，查询语言包括XPath和XQuery。

本书特色

- 全面改版的组织结构。
- UML数据库模型的新内容。
- 包括3NF综合算法在内的操作依赖新算法的引入。
- 更多的3NF，包括3NF综合算法。
- 扩展的SQL触发讨论。
- 新增的索引选择和物化视图。
- 新增的三层体系结构。
- 新增的PHP。
- 新增的OLAP和SQL立方体操作符介绍。
- 扩展的XML内容，包括XML模式、XPath、XQuery和XSLT。


作者简介

Jeffrey D. Ullman 斯坦福大学计算机科学系Stanford W. Ascherman教授，数据库技术专家。他独立或与入合作出版了15本著作，发表了170多篇技术论文。他的研究兴趣包括数据库理论、数据库集成、数据挖掘和利用信息基础设施进行教育。他是美国国家工程院院士成员，曾获得Knuth奖、SIGMOD贡献奖、Karlstrom杰出教育家奖和Edgar F. Codd发明奖。



Jennifer Widom 美国康奈尔大学计算机科学博士，现为斯坦福大学计算机科学与工程系教授，研究非传统的数据管理。她是ACM会士、Guggenheim会士和美国国家工程院院士成员，并且是多个编辑委员会、程序委员会和顾问委员会的成员。她在2007年获得了ACM SIGMOD Edgar F. Codd发明奖。



 www.PearsonEd.com
PRENTICE HALL




For sale and distribution in the People's Republic of China exclusively (except Taiwan, Hong Kong SAR and Macau SAR).

仅限于中华人民共和国境内(不包括中国香港、澳门特别行政区和中国台湾地区)销售发行。

投稿热线: (010) 88379604
购书热线: (010) 68995259, 68995264
读者信箱: hzjsj@hzbook.com

华章网站 <http://www.hzbook.com>

 网上购书: www.china-pub.com



上架指导: 计算机/数据库

ISBN 978-7-111-24733-3



9 787111 247333

定价: 45.00元



数据库系统基础教程

(美)

Jeffrey D. Ullman


Jennifer Widom

著

PRENTICE
HALL

A First Course in Database Systems (Third Edition)

英文版
第3版



机械工业出版社

China Machine Press

经典原版书库

数据库系统基础教程

(英文版·第3版)

A First Course in Database Systems

(Third Edition)

Jeffrey D. Ullman

斯坦福大学

Jennifer Widom

斯坦福大学

著



机械工业出版社
China Machine Press

English reprint edition copyright © 2008 by Pearson Education Asia Limited and China Machine Press.

Original English language title: *A First Course in Database Systems, Third Edition* (ISBN 978-0-13-502176-7) by Jeffrey D. Ullman and Jennifer Widom, Copyright © 2008, 2002, 1997 Pearson Education, Inc.

All rights reserved.

Published by arrangement with the original publisher, Pearson Education, Inc., publishing as Prentice Hall.

For sale and distribution in the People's Republic of China exclusively (except Taiwan, Hong Kong SAR and Macau SAR).

本书英文影印版由Pearson Education Asia Ltd.授权机械工业出版社独家出版。未经出版者书面许可，不得以任何方式复制或抄袭本书内容。

仅限于中华人民共和国境内（不包括中国香港、澳门特别行政区和中国台湾地区）销售发行。

本书封面贴有Pearson Education（培生教育出版集团）激光防伪标签，无标签者不得销售。

版权所有，侵权必究。

本书法律顾问 北京市展达律师事务所

本书版权登记号：图字：01-2008-3272

图书在版编目（CIP）数据

数据库系统基础教程（英文版·第3版）/（美）厄尔曼（Ullman, J. D.），（美）怀德姆（Widom, J.）著.—北京：机械工业出版社，2008.8
（经典原版书库）

书名原文：A First Course in Database Systems, Third Edition

ISBN 978-7-111-24733-3

I. 数… II. ①厄… ②怀… III. 数据库系统—教材—英文 IV. TP311.13

中国版本图书馆CIP数据核字（2008）第109993号

机械工业出版社（北京市西城区百万庄大街22号 邮政编码 100037）

责任编辑：迟振春

北京京北制版厂印刷 · 新华书店北京发行所发行

2008年8月第1版第1次印刷

145mm × 210mm · 18.375印张

标准书号：ISBN 978-7-111-24733-3

定价：45.00元

凡购本书，如有倒页、脱页、缺页，由本社发行部调换
本社购书热线：（010）68326294

试读结束：需要全本请在线购买：www.ertongbook.com

出版者的话

文艺复兴以降，源远流长的科学精神和逐步形成的学术规范，使西方国家在自然科学的各个领域取得了垄断性的优势；也正是这样的传统，使美国在信息技术发展的六十多年间名家辈出、独领风骚。在商业化的进程中，美国的产业界与教育界越来越紧密地结合，计算机学科中的许多泰山北斗同时身处科研和教学的最前线，由此而产生的经典科学著作，不仅擘划了研究的范畴，还揭示了学术的源变，既遵循学术规范，又自有学者个性，其价值并不会因年月的流逝而减退。

近年，在全球信息化大潮的推动下，我国的计算机产业发展迅猛，对专业人才的需求日益迫切。这对计算机教育界和出版界都既是机遇，也是挑战；而专业教材的建设在教育战略上显得举足轻重。在我国信息技术发展时间较短的现状下，美国等发达国家在其计算机科学发展的几十年间积淀和发展的经典教材仍有许多值得借鉴之处。因此，引进一批国外优秀计算机教材将对我国计算机教育事业的发展起到积极的推动作用，也是与世界接轨、建设真正的世界一流大学的必由之路。

机械工业出版社华章分社较早意识到“出版要为教育服务”。自1998年开始，华章分社就将工作重点放在了遴选、移译国外优秀教材上。经过多年的不懈努力，我们与Pearson, McGraw-Hill, Elsevier, MIT, John Wiley & Sons, Cengage等世界著名出版公司建立了良好的合作关系，从他们现有的数百种教材中甄选出Andrew S. Tanenbaum, Bjarne Stroustrup, Brian W. Kernighan, Dennis Ritchie, Jim Gray, Alfred V. Aho, John E. Hopcroft, Jeffrey D. Ullman, Abraham Silberschatz, William Stallings, Donald E. Knuth, John L. Hennessy, Larry L. Peterson等大师名家的一批经典作品，以“计算机科学丛书”为总称出版，供读者学习、研究及珍藏。大理石纹理的封面，也正体现了这套丛书的品位和格调。

“计算机科学丛书”的出版工作得到了国内外学者的鼎力襄助，国内的专家不仅提供了中肯的选题指导，还不辞劳苦地担任了翻译和审校的工作；而原书的作者也相当关注其作品在中国的传播，有的还专程为其书的中译本作序。迄今，“计算机科学丛书”已经出版了近两百个品种，这些书籍在读者中树立了良好的口碑，并被许多高校采用为正式教材和参考书籍。其影印版“经典原版书库”作为姊妹篇也被越来越多实施双语教学的学校所采用。

权威的作者、经典的教材、一流的译者、严格的审校、精细的编辑，这些因素使我们的图书有了质量的保证。随着计算机科学与技术专业学科建设的不断完善和教材改革的逐渐深化，教育界对国外计算机教材的需求和应用都将步入一个新的阶段，我们的目标是尽善尽美，而反馈的意见正是我们达到这一终极目标的重要帮助。华章分社欢迎老师和读者对我们的工作提出建议或给予指正，我们的联系方式如下：

华章网站：www.hzbook.com

电子邮件：hzedu@hzbook.com

联系电话：(010) 68995264

联系地址：北京市西城区百万庄南街1号

邮政编码：100037



华章教育

华章科技图书出版中心

Preface

At Stanford, we are on the quarter system, and as a result, our introductory database instruction is divided into two courses. The first, CS145, is designed for students who will use database systems but not necessarily take a job implementing a DBMS. It is a prerequisite for CS245, which is the introduction to DBMS implementation. Students wishing to go further in the database field then take CS345 (special topics), CS346 (DBMS implementation project), and CS347 (transaction processing and distributed databases).

Starting in 1997, we published a pair of books. *A First Course in Database Systems* was designed for CS145, and *Database System Implementation* was for CS245 and parts of CS346. Because many schools are on the semester system or combine the two kinds of database instruction into one introductory course, we felt that there was a need to produce the two books as a single volume, which we call *Database Systems: The Complete Book*.

- If you are a student considering buying this book, and anticipate a later study of implementation, you should consider buying *Database Systems: The Complete Book* instead.

However, because many more students need to know how to use database systems than to implement them, we have continued to package the first half of the “complete book” as *A First Course in Database Systems*. In the third edition, we have introduced many new topics and altered the overall viewpoint somewhat. Today, there are two important models for database systems: relational and semistructured (XML). We have decided therefore to downplay object-oriented databases, except in the contexts of design and object-relational systems.

What’s New in the Third Edition

After a brief introduction in Chapter 1, we cover relational modeling in Chapters 2–4. Chapter 4 is devoted to high-level modeling. There, in addition to the E/R model, we now cover UML (Unified Modeling Language). We also have moved to Chapter 4 a shorter version of the material on ODL, treating it as

a design language for relational database schemas. The earlier, more extensive treatment of ODL and OQL is available on the book's Web site.

The material on functional and multivalued dependencies has been modified and remains in Chapter 3. We have changed our viewpoint, so that a functional dependency is assumed to have a set of attributes on the right. We have also given explicitly certain algorithms, including the "chase," that allow us to manipulate dependencies. We have augmented our discussion of third normal form to include the 3NF synthesis algorithm and to make clear what the tradeoff between 3NF and BCNF is.

Chapter 5 contains the coverage of relational algebra from the previous edition, and is joined by (part of) the treatment of Datalog from the old Chapter 10. The discussion of recursion in Datalog is either moved to the book's Web site or combined with the treatment of recursive SQL in Chapter 10 of this edition.

Chapters 6–10 are devoted to aspects of SQL programming, and they represent a reorganization and augmentation of the earlier book's Chapters 6, 7, 8, and parts of 10. The material on indexes and views has been moved to its own chapter, number 8, and this material has been augmented with a discussion of important new topics, including materialized views, and automatic selection of indexes.

The new Chapter 9 is based on the old Chapter 8 (embedded SQL). It is introduced by a new section on 3-tiered architecture. It also includes an expanded discussion of JDBC and new coverage of PHP.

Chapter 10 collects a number of advanced SQL topics. The discussion of authorization from the old Chapter 8 has been moved here, as has the discussion of recursive SQL from the old Chapter 10. Most of the chapter is devoted to the nested-relation model (from the old Chapter 4) and object-relational features of SQL (from the old Chapter 9).

Then, Chapters 11 and 12 cover XML and systems based on XML. Except for material at the end of the old Chapter 4, which has been moved to Chapter 11, this material is all new. Chapter 11 covers modeling; it includes expanded coverage of DTD's, along with new material on XML Schema. Chapter 12 is devoted to programming, and it includes sections on XPath, XQuery, and XSLT.

Use of the Book

There is adequate material in this volume for a one-semester course on database modeling and programming. For a one-quarter course, you will probably have to omit some of the topics. We regard Chapters 2–7 as the core of the course. The remaining five chapters contain material from which it is safe to select at will, although we believe that every student should get some exposure to the issues of embedding SQL in standard host languages from one of the sections in Chapter 9.

If, as we do in CS145, you give students a substantial database-application design and implementation project, then you may have to reorder the material somewhat, so that SQL instruction occurs earlier. You may wish to defer material such as dependencies, although students need normalization for design.

Prerequisites

We have used the book at the “mezzanine” level, in a course taken both by undergraduates and beginning graduate students. The formal prerequisites for the course are Sophomore-level treatments of: (1) Data structures, algorithms, and discrete math, and (2) Software systems, software engineering, and programming languages. Of this material, it is important that students have at least a rudimentary understanding of such topics as: algebraic expressions and laws, logic, basic data structures, object-oriented programming concepts, and programming environments. However, we believe that adequate background is acquired by the Junior year of a typical computer science program.

Exercises

The book contains extensive exercises, with some for almost every section. We indicate harder exercises or parts of exercises with an exclamation point. The hardest exercises have a double exclamation point.

Gradiance On-Line Exercises

There is an accompanying set of on-line homeworks using a technology developed by Gradiance Corp. Instructors may assign these homeworks to their class, or students not enrolled in a class may enroll in an “omnibus class” that allows them to do the homeworks as a tutorial (without an instructor-created class). Gradiance questions look like ordinary questions, but your solutions are sampled. If you make an incorrect choice you are given specific advice or feedback to help you correct your solution. If your instructor permits, you are allowed to try again, until you get a perfect score.

In addition, the Gradiance package for the book includes programming exercises in SQL and XQuery. Submitted queries are tested for correctness, and incorrect results lead to examples of where the query goes wrong. Students can try as many times as they like, but writing queries that only respond correctly to the examples is not sufficient to get credit for the problem.

Gradiance service can be purchased at

<http://www.prenhall.com/goal>

Instructors who want to use the system in their classes should contact their Prentice-Hall representative.

Support on the World Wide Web

The book's home page is

<http://www-db.stanford.edu/~ullman/fcdb.html>

You will find errata as we learn of them, and backup materials. We are making available the notes for each offering of CS145 as we teach it, including homeworks, projects and exams. We shall also make available there the sections from the second edition that have been removed from the third.

Mapping the Second Edition to the Third

Here is a table giving the correspondence between old sections and new.

Old	New	Old	New	Old	New	Old	New	Old	New
1.1	1.1	1.2	1.2	1.3	1.3	2.1	4.1	2.2	4.2
2.3	4.3	2.4	4.4	3.1	2.2	3.2	4.5	3.3	4.6
3.4	3.1	3.5	3.2	3.6	3.3-5	3.7	3.6-7	4.1	Web
4.2	4.9	4.3	4.9	4.4	4.10	4.5	10.3	4.6	11.1
4.7	11.2	5.1	2.2	5.2	2.4	5.3	5.1	5.4	5.2
5.5	2.5	6.1	6.1	6.2	6.2	6.3	6.3	6.4	6.4
6.5	6.5	6.6	2.3	6.7	8.1-2	7.1	7.1	7.2	7.2
7.3	7.3	7.4	7.4-5	8.1	9.3	8.2	9.4	8.3	9.2
8.4	9.5	8.5	9.6	8.6	6.6	8.7	10.1	9.1	Web
9.2	Web	9.3	Web	9.4	10.4	9.5	10.5	10.1	5.3
10.2	5.4	10.3	Web	10.4	10.2				

Acknowledgements

We would like to thank Donald Kossmann for helpful discussions, especially concerning XML and its associated programming systems. Also, Bobbie Cochrane assisted us in understanding trigger semantics for a earlier edition.

A large number of people have helped us, either with the development of this book or its predecessors, or by contacting us with errata in the books and/or other Web-based materials. It is our pleasure to acknowledge them all here.

Marc Abromowitz, Joseph H. Adamski, Brad Adelberg, Gleb Ashimov, Donald Aingworth, Teresa Almeida, Brian Babcock, Bruce Baker, Yunfan Bao, Jonathan Becker, Margaret Benitez, Eberhard Bertsch, Larry Bonham, Phillip Bonnet, David Brokaw, Ed Burns, Alex Butler, Karen Butler, Mike Carey, Christopher Chan, Sudarshan Chawathe.

Also Per Christensen, Ed Chang, Surajit Chaudhuri, Ken Chen, Rada Chirkova, Nitin Chopra, Lewis Church, Jr., Bobbie Cochrane, Michael Cole, Alissa Cooper, Arturo Crespo, Linda DeMichiel, Matthew F. Dennis, Tom Diestbier, Pearl D'Souza, Oliver Duschka, Xavier Faz, Greg Fichtenholtz, Bart Fisher, Simon Frettloeh, Jarl Friis.

Also John Fry, Chiping Fu, Tracy Fujieda, Prasanna Ganesan, Suzanne Garcia, Mark Gjol, Manish Godara, Seth Goldberg, Jeff Goldblat, Meredith Goldsmith, Luis Gravano, Gerard Guillemette, Himanshu Gupta, Petri Gynther, Jon Heggland, Rafael Hernandez, Masanori Higashihara, Antti Hjelt, Ben Holtzman, Steve Huntsberry.

Also Sajid Hussain, Leonard Jacobson, Thulasiraman Jeyaraman, Dwight Joe, Brian Jorgensen, Mathew P. Johnson, Sameh Kamel, Seth Katz, Pedram Keyani, Victor Kimeli, Ed Knorr, Yeong-Ping Koh, David Koller, Gyorgy Kovacs, Phillip Koza, Brian Kulman, Bill Labiosa, Sang Ho Lee, Younghan Lee, Miguel Licona.

Also Olivier Lobry, Chao-Jun Lu, Waynn Lue, John Manz, Arun Marathe, Philip Minami, Le-Wei Mo, Fabian Modoux, Peter Mork, Mark Mortensen, Ramprakash Narayanaswami, Hankyung Na, Mor Naaman, Mayur Naik, Marie Nilsson, Torbjorn Norbye, Chang-Min Oh, Mehul Patel, Soren Peen, Jian Pei.

Also Xiaobo Peng, Bert Porter, Limbek Reka, Prahash Ramanan, Nisheeth Ranjan, Suzanne Rivoire, Ken Ross, Tim Roughgarten, Mema Roussopoulos, Richard Scherl, Loren Shevitz, June Yoshiko Sison, Man Cho A. So, Elizabeth Stinson, Qi Su, Ed Swierk, Catherine Tornabene, Anders Uhl, Jonathan Ullman, Mayank Upadhyay.

Also Anatoly Varakin, Vassilis Vassalos, Krishna Venuturimilli, Vikram Vijayaraghavan, Terje Viken, Qiang Wang, Mike Wiacek, Kristian Widjaja, Janet Wu, Sundar Yamunachari, Takeshi Yokukawa, Bing Yu, Min-Sig Yun, Torben Zahle, Sandy Zhang. The remaining errors are ours, of course.

J. D. U.

J. W.

Stanford, CA

July, 2007

Table of Contents

1	The Worlds of Database Systems	1
1.1	The Evolution of Database Systems	1
1.1.1	Early Database Management Systems	2
1.1.2	Relational Database Systems	3
1.1.3	Smaller and Smaller Systems	3
1.1.4	Bigger and Bigger Systems	4
1.1.5	Information Integration	4
1.2	Overview of a Database Management System	5
1.2.1	Data-Definition Language Commands	5
1.2.2	Overview of Query Processing	5
1.2.3	Storage and Buffer Management	7
1.2.4	Transaction Processing	8
1.2.5	The Query Processor	9
1.3	Outline of Database-System Studies	10
1.4	References for Chapter 1	12
I	Relational Database Modeling	15
2	The Relational Model of Data	17
2.1	An Overview of Data Models	17
2.1.1	What is a Data Model?	17
2.1.2	Important Data Models	18
2.1.3	The Relational Model in Brief	18
2.1.4	The Semistructured Model in Brief	19
2.1.5	Other Data Models	20
2.1.6	Comparison of Modeling Approaches	21
2.2	Basics of the Relational Model	21
2.2.1	Attributes	22
2.2.2	Schemas	22
2.2.3	Tuples	22
2.2.4	Domains	23
2.2.5	Equivalent Representations of a Relation	23

2.2.6	Relation Instances	24
2.2.7	Keys of Relations	25
2.2.8	An Example Database Schema	26
2.2.9	Exercises for Section 2.2	28
2.3	Defining a Relation Schema in SQL	29
2.3.1	Relations in SQL	29
2.3.2	Data Types	30
2.3.3	Simple Table Declarations	31
2.3.4	Modifying Relation Schemas	33
2.3.5	Default Values	34
2.3.6	Declaring Keys	34
2.3.7	Exercises for Section 2.3	36
2.4	An Algebraic Query Language	38
2.4.1	Why Do We Need a Special Query Language?	38
2.4.2	What is an Algebra?	38
2.4.3	Overview of <u>Relational Algebra</u>	39
2.4.4	Set Operations on Relations	39
2.4.5	Projection	41
2.4.6	Selection	42
2.4.7	Cartesian Product	43
2.4.8	Natural Joins	43
2.4.9	Theta-Joins	45
2.4.10	Combining Operations to Form Queries	47
2.4.11	Naming and Renaming	49
2.4.12	Relationships Among Operations	50
2.4.13	A Linear Notation for Algebraic Expressions	51
2.4.14	Exercises for Section 2.4	52
2.5	Constraints on Relations	58
2.5.1	Relational Algebra as a Constraint Language	59
2.5.2	Referential Integrity Constraints	59
2.5.3	Key Constraints	60
2.5.4	Additional Constraint Examples	61
2.5.5	Exercises for Section 2.5	62
2.6	Summary of Chapter 2	63
2.7	References for Chapter 2	65
3	Design Theory for Relational Databases	67
3.1	Functional Dependencies	67
3.1.1	Definition of Functional Dependency	68
3.1.2	Keys of Relations	70
3.1.3	Superkeys	71
3.1.4	Exercises for Section 3.1	71
3.2	Rules About Functional Dependencies	72
3.2.1	Reasoning About Functional Dependencies	72
3.2.2	The Splitting/Combining Rule	73

3.2.3	Trivial Functional Dependencies	74
3.2.4	Computing the Closure of Attributes	75
3.2.5	Why the Closure Algorithm Works	77
3.2.6	The Transitive Rule	79
3.2.7	Closing Sets of Functional Dependencies	80
3.2.8	Projecting Functional Dependencies	81
3.2.9	Exercises for Section 3.2	83
3.3	Design of Relational Database Schemas	85
3.3.1	Anomalies	86
3.3.2	Decomposing Relations	86
3.3.3	Boyce-Codd Normal Form	88
3.3.4	Decomposition into BCNF	89
3.3.5	Exercises for Section 3.3	92
3.4	Decomposition: The Good, Bad, and Ugly	93
3.4.1	Recovering Information from a Decomposition	94
3.4.2	The Chase Test for Lossless Join	96
3.4.3	Why the Chase Works	99
3.4.4	Dependency Preservation	100
3.4.5	Exercises for Section 3.4	102
3.5	Third Normal Form	102
3.5.1	Definition of Third Normal Form	102
3.5.2	The Synthesis Algorithm for 3NF Schemas	103
3.5.3	Why the 3NF Synthesis Algorithm Works	104
3.5.4	Exercises for Section 3.5	105
3.6	Multivalued Dependencies	105
3.6.1	Attribute Independence and Its Consequent Redundancy	106
3.6.2	Definition of Multivalued Dependencies	107
3.6.3	Reasoning About Multivalued Dependencies	108
3.6.4	Fourth Normal Form	110
3.6.5	Decomposition into Fourth Normal Form	111
3.6.6	Relationships Among Normal Forms	113
3.6.7	Exercises for Section 3.6	113
3.7	An Algorithm for Discovering MVD's	115
3.7.1	The Closure and the Chase	115
3.7.2	Extending the Chase to MVD's	116
3.7.3	Why the Chase Works for MVD's	118
3.7.4	Projecting MVD's	119
3.7.5	Exercises for Section 3.7	120
3.8	Summary of Chapter 3	121
3.9	References for Chapter 3	122

4	High-Level Database Models	125
4.1	The Entity/Relationship Model	126
4.1.1	Entity Sets	126
4.1.2	Attributes	126
4.1.3	Relationships	127
4.1.4	Entity-Relationship Diagrams	127
4.1.5	Instances of an E/R Diagram	128
4.1.6	Multiplicity of Binary E/R Relationships	129
4.1.7	Multiway Relationships	130
4.1.8	Roles in Relationships	131
4.1.9	Attributes on Relationships	134
4.1.10	Converting Multiway Relationships to Binary	134
4.1.11	Subclasses in the E/R Model	135
4.1.12	Exercises for Section 4.1	138
4.2	Design Principles	140
4.2.1	Faithfulness	140
4.2.2	Avoiding Redundancy	141
4.2.3	Simplicity Counts	142
4.2.4	Choosing the Right Relationships	142
4.2.5	Picking the Right Kind of Element	144
4.2.6	Exercises for Section 4.2	145
4.3	Constraints in the E/R Model	148
4.3.1	Keys in the E/R Model	148
4.3.2	Representing Keys in the E/R Model	149
4.3.3	Referential Integrity	150
4.3.4	Degree Constraints	151
4.3.5	Exercises for Section 4.3	151
4.4	Weak Entity Sets	152
4.4.1	Causes of Weak Entity Sets	152
4.4.2	Requirements for Weak Entity Sets	153
4.4.3	Weak Entity Set Notation	155
4.4.4	Exercises for Section 4.4	156
4.5	From E/R Diagrams to Relational Designs	157
4.5.1	From Entity Sets to Relations	157
4.5.2	From E/R Relationships to Relations	158
4.5.3	Combining Relations	160
4.5.4	Handling Weak Entity Sets	161
4.5.5	Exercises for Section 4.5	163
4.6	Converting Subclass Structures to Relations	165
4.6.1	E/R-Style Conversion	166
4.6.2	An Object-Oriented Approach	167
4.6.3	Using Null Values to Combine Relations	168
4.6.4	Comparison of Approaches	169
4.6.5	Exercises for Section 4.6	171
4.7	Unified Modeling Language	171

4.7.1	UML Classes	172
4.7.2	Keys for UML classes	173
4.7.3	Associations	173
4.7.4	Self-Associations	175
4.7.5	Association Classes	175
4.7.6	Subclasses in UML	176
4.7.7	Aggregations and Compositions	177
4.7.8	Exercises for Section 4.7	179
4.8	From UML Diagrams to Relations	179
4.8.1	UML-to-Relations Basics	179
4.8.2	From UML Subclasses to Relations	180
4.8.3	From Aggregations and Compositions to Relations	181
4.8.4	The UML Analog of Weak Entity Sets	181
4.8.5	Exercises for Section 4.8	183
4.9	Object Definition Language	183
4.9.1	Class Declarations	184
4.9.2	Attributes in ODL	184
4.9.3	Relationships in ODL	185
4.9.4	Inverse Relationships	186
4.9.5	Multiplicity of Relationships	186
4.9.6	Types in ODL	188
4.9.7	Subclasses in ODL	190
4.9.8	Declaring Keys in ODL	191
4.9.9	Exercises for Section 4.9	192
4.10	From ODL Designs to Relational Designs	193
4.10.1	From ODL Classes to Relations	193
4.10.2	Complex Attributes in Classes	194
4.10.3	Representing Set-Valued Attributes	195
4.10.4	Representing Other Type Constructors	196
4.10.5	Representing ODL Relationships	198
4.10.6	Exercises for Section 4.10	198
4.11	Summary of Chapter 4	200
4.12	References for Chapter 4	202

II Relational Database Programming 203

5	Algebraic and Logical Query Languages	205
5.1	Relational Operations on Bags	205
5.1.1	Why Bags?	206
5.1.2	Union, Intersection, and Difference of Bags	207
5.1.3	Projection of Bags	208
5.1.4	Selection on Bags	209
5.1.5	Product of Bags	210
5.1.6	Joins of Bags	210