

国外优秀信息科学与技术系列教学用书

数据库系统

——面向应用的方法


(第二版 影印版)

DATABASE SYSTEMS

An Application-Oriented Approach

(Second Edition)

■ Michael Kifer
Arthur Bernstein
Philip M. Lewis



高等教育出版社
Higher Education Press

国外优秀信息科学与技术系列教学用书

数据库系统 ——面向应用的方法

(第二版 影印版)

DATABASE SYSTEMS
An Application-Oriented Approach
(Second Edition)

Michael Kifer
Arthur Bernstein
Philip M. Lewis



高等教育出版社

图字:01-2005-0965 号

Database Systems: An Application-Oriented Approach, Introductory Version, Second Edition

Michael Kifer, Arthur Bernstein, Philip M. Lewis

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。

English reprint edition copyright © 2005 by PEARSON EDUCATION ASIA LIMITED and HIGHER EDUCATION PRESS. (Database Systems: An Application-Oriented Approach, Introductory Version, 2e from Pearson Education's edition of the Work)

Database Systems: An Application-Oriented Approach, Introductory Version, 2e by Michael Kifer, Arthur Bernstein, Philip M. Lewis, Copyright © 2005.

All Rights Reserved.

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

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

原版 ISBN: 0-321-22838-3

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

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

图书在版编目(CIP)数据

数据库系统:面向应用的方法 = Database Systems:

An Application - Oriented Approach: 第2版 / (美)

基弗(Kifer, M.), (美)伯恩斯坦(Bernstein, A.),

(美)刘易斯(Lewis, P. M.). —影印本. —北京:

高等教育出版社, 2005. 12

ISBN 7-04-017817-6

I. 数... II. ①基... ②伯... ③刘... III. 数据库系统 - 高等学校 - 教材 - 英文 IV. TP311.13

中国版本图书馆 CIP 数据核字(2005)第 134986 号

出版发行 高等教育出版社
社 址 北京市西城区德外大街 4 号
邮政编码 100011
总 机 010-58581000

经 销 蓝色畅想图书发行有限公司
印 刷 北京民族印刷厂

开 本 787 × 1092 1/16
印 张 43.25
字 数 830 000

购书热线 010-58581118
免费咨询 800-810-0598
网 址 <http://www.hep.edu.cn>
<http://www.hep.com.cn>
网上订购 <http://www.landaco.com>
<http://www.landaco.com.cn>
畅想教育 <http://www.widedu.com>

版 次 2005 年 12 月第 1 版
印 次 2005 年 12 月第 1 次印刷
定 价 49.50 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

物料号 17817-00

出版说明

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。引进国外信息科学与技术优秀教材,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与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,第一批由教育部高等教育司推荐的 20 多种引进教材已经陆续出版。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。这套教材基本覆盖了计算机科学与技术专业的课程体系,体现了权威性、系统性、先进性和经济性等特点。

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。

希望这些教学用书的引进出版,对于提高我国高等学校信息科学技术的教学水平,缩小与国际先进水平的差距,加快培养一大批具有国际竞争力的高质量信息技术人才,起到积极的推动作用。同时也欢迎广大教师和专家们对教材引进工作提出宝贵的意见和建议。联系方式:hep.cs@263.net。

高等教育出版社

二〇〇四年十月

前言

此教材的第二版将会发行两个版本：

- 本书的版本,主要是入门性内容,适合本科生或研究生的数据库初级课程。
- 另一个版本则是完全版,适合以下三方面的课程。

■ 数据库入门课程,面向本科生或研究生。

■ 事务处理方面的课程,面向学习过数据库入门性课程的本科生或研究生。

■ 关于数据库的本科生高级课程或者研究生初级课程,面向学习过数据库入门性课程的学生。

我们的目标之一就是减少入门版本的篇幅,使学生更容易接受。另一个目标是在使用本书第一版所获经验的基础上,尽量使本版本更加完善。

本书的章节不仅仅是完全版的子集。我们认为讲授数据库入门课程的教师应该选择增加对象数据库和 XML(完全版中许多章节详细叙述的主题),以丰富这门课程的内容。因此在本书中,我们新加了两章,第 16 章“对象数据库绪论”和第 17 章“XML 与 Web 数据绪论”,都是从本书完全版中经过适当挑选出来的内容。

为了使本书能与迅速发展的技术保持同步,我们在本书两个版本的许多章节中加入了大量关于 UML 的内容,并且新加了一章关于数据库优化的内容(第 12 章)。

正如第一版,我们主要关注如何利用数据库来创建应用程序,而不是如何创建数据库管理系统本身。我们相信,开发应用程序的学生比创建 DBMS 的学生要多得多。因此,本书加入了大量关于通过事务访问数据库的语言(如嵌入式 SQL、ODBC 和 JDBC)和 API 的内容。

尽管本书涵盖了许多关于数据库和事务处理应用程序的实用内容,但是我们更关心这些主题的基本概念,而不是特定的商用系统或应用程序的细节。因此,我们把重点集中在关系和对象数据模型的概念上。即使在 SQL 过时很久之后,这些概念仍然是数据库处理的基础。

为了加强学生对技术内容的理解,我们增加了一个事务处理应用程序的案例,即贯彻全书的“学生注册系统”。尽管一个学生注册系统可能并不吸引人,但它有一个独特的优势,即所有学生都曾作为用户和这样的系统进行过交互。更重要的是,“麻雀虽小,五脏俱全”,因此我们可以用它来阐述在数据库设计、查询处理和事务处理等方面的诸多问题。

本书的一个重要特点就是将“学生注册系统”作为例子,介绍了用来实现事务处理应用程序所需的软件工程概念。由于许多信息系统的实现因为项目管理不善和软件工程的欠缺而最终夭折,因此我们认为这些内容应成为学生教育中一个重要组成部分。由于许多学生还需要在这方面另修专门的课程,因此本书对软件工程问题的处理是简明扼要的。然而,我们相信,如果学生们看到这些内容在信息系统实现中会涉及,他们能更好地理解并应用它。由于在 Stony Brook 使用本教材的课程不是软件工程课程,因此我们上课时不涉及这些内容。事实上,我们要求学生课后阅读并在他们的课程项目中很好地利用软件工程实践。我们在课堂上会介绍“学生注册系统”的特点,以说明关于数据库和事务处理的重要问题。

第二版中的变化

数据库和事务处理系统的基本技术变化如此之快,使我们对本书第一版的内容进行了大量的更改和增补。一种快速发展的技术就是统一建模语言(UML)。在第4章“数据库设计”原有的E-R图的基础上,增加了大量关于UML的内容。同时还在第2章、第14章和第15章关于软件工程的内容中添加了UML的内容。

由于在实际应用中,大量投资被用在日益增加的数据库和事务处理应用程序的吞吐量方面,我们增加了第12章关于“数据库优化”的内容。

除此之外,几乎所有章节都进行了增补和更新。一个重要的例子就是融入了SQL/XML和RAID技术。

本书虽然没有包括Web业务,但它仍是一个很重要的内容。因为这是一个快速发展的、面向应用的有意义的技术,我们对本书的完全版做了大量修改以加入这方面的内容。除了通过更新关于XML和Web数据的章节并加入关于SQL/XML的小节来加强本书在XML技术方面的内容之外,我们也加入了新的关于Web业务的章节,包括SOAP、WSDL、BPEL、UDDI和利用WS-Coordination与WS-Transaction处理基于XML的事务等方面的内容。在“安全与网络商务”这一章,加入了基于XML的加密、利用XML加密、XML签名、WS安全和SAML等方面的内容。而在“事务处理系统体系结构”这一章,加入了关于网络应用程序服务器和J2EE方面的内容,它们被用来实现对许多Web业务的支持。

本书的组织

第1章~第7章应该以它们在本书中出现的顺序来进行讲授。第8章包含了许多信息,学生需要利用这些信息将前几章所学到的知识运用到实践中。但是,后续的章节并不是很依赖于第8章。第3部分中的第9章~第12章也应当按照顺序进行讲授。第4部分中软件工程方面的章节利用了第2部分和第3部分中的章节内容,但是关于软件工程的章节可以和数据库方面的内容并行讲授。本书高级主题中的第16章和第17章依赖于第2部分中的前7章。

最后,需要指出的是,书中打了“*”的小节是可选的,因此如果教师愿意,可以将其省略。目录中前面有C_s图标的小节是关于案例研究的。另外,带有星号的习题会比其他题目难一些,而带有两个星号的题目更难。

补充材料

除了本教材以外,下列补充材料可以协助教师教学:

- 所有章节的在线 PowerPoint 讲稿。
- 所有图表的在线 PowerPoint。
- 一份在线的习题解答手册,包括所有练习的答案。
- 附加的参考文献、笔记、勘误表,作业和测试题。

请访问本书原版的伴随网站(www.aw-bc.com/kifer)来取得关于如何获取这些补充材料的详细信息。答案手册和 PowerPoint 讲稿只能通过 Addison-Wesley 公司的销售代表向教师提供。请访问 www.aw.com 来联系你所在区域的销售代表。

致谢

我们要感谢下列审阅者,他们的意见和建议极大地促进了本书第二版的修改:

Tran Cao Son, New Mexico State University

Frantisek Franek, McMaster University

Junping Sun, Nova Southeastern University

Philip Cannata, Sun Microsystems

Dehu Qi, Lamar University

Nematollah Shiri, Concordia University

Jian Pei, State University of New York at Buffalo

Jack Wileden, University of Massachusetts Amherst

Sibel Adali, Rensselaer Polytechnic Institute

Roger King, University of Colorado at Boulder

Markus Schneider, University of Florida

Yaron Y. Goland, BEA

Dennis Shasha, New York University

Christelle Scharff, Pace University

Zhiwei Wang, Graduate Programs in Software Engineering, IT, and IS University of St. Thomas

我们还要感谢本书第一版的审阅者:

Suad Alagic, Wichita University

Catriel Beerli, The Hebrew University

Rick Cattel, Sun Microsystems

Jan Chomicki, SUNY Buffalo

Henry A. Etlinger, Rochester Institute of Technology

Leonidas Fegaras, University of Texas at Arlington

Alan Fekete, University of Sidney

Johannes Gehrke, Cornell University

Hershel Gottesman, Consultant

Jiawei Han, Simon Fraser University

Peter Honeyman, University of Michigan

Vijay Kumar, University of Missouri-Kansas City

Jonathan Lazar, Towson University

Dennis McLeod, University of Southern California

Rokia Missaoui, University of Quebec in Montreal

Clifford Neuman, University of Southern California

Fabian Pascal, Consultant

Sudha Ram, University of Arizona

Krithi Ramamritham, University of Massachusetts-Amherst, and IIT Bombay

Andreas Reuter, International University in Germany, Bruchsal

Arijit Sengupta, Georgia State University

Munindar P. Singh, North Carolina State University

Greg Speegle, Baylor University

Junping Sun, Nova Southeastern University

Joe Trubisz, Consultant

Vassilis J. Tsotras, University of California, Riverside

Emilia E. Villarreal, California Polytechnic State University

同时还要感谢以下人员,他们对我们的问题提供了补充信息和解答:Don Chamberlin、Daniela Florescu、Jim Gray、Pankaj Gupta、Rob Kelly 和 C. Mohan。

下列两位使用本书的测试版进行教学,并且提出了宝贵的意见和建议:David S. Warren 和 Radu Grosu。Joe Trubicz 不仅是手稿完成后的审阅者,同时对本书许多章节的早期版本提出了宝贵的意见。

许多学生对本书各部分的阅读和勘误提供了很大的帮助:Ziyang Duan、Shiyong Lu、Swapnil Patil、Guizhen Yang 和 Yan Zhang。

非常感谢 Stony Brook 计算机科学系的教职员工,尤其是 Kathy Germana,她的工作使得一切都成为可能。

我们要特别感谢 Matt Goldstein 和 Maite Suarez-Rivas——Addison-Wesley 出版公司的编辑,他们在本书早期内容和方法的策划以及全书的编写过程中都起着非常重要的作用。我们也要感谢 Addison-Wesley 公司和 Windfall 软件公司的员工,感谢他们在本书的编辑和生产过程中的出色工作:Jeffrey Holcomb、Paul Anagnostopoulos、Elisabeth Beller、Jennifer McClain 和 Joe Snowden。

最后,我们需要感谢的是我们的妻子 Lora、Edie 和 Rhoda,感谢她们在我们编写本书的过程中所给予的大力支持和鼓励。

Michael Kifer
Arthur Bernstein
Philip M. Lewis

Preface

We are publishing the second edition of our textbook in two versions:

- This version, which consists of introductory material, is appropriate for a first undergraduate or graduate course in databases.
- The second version, which is the complete book, is appropriate for three courses:
 - An introductory undergraduate or graduate course in databases
 - An undergraduate or graduate course in transaction processing for students who have had an introductory course in databases
 - An advanced undergraduate or a first graduate course in databases for students who have had an introductory course in databases

One of our goals was to reduce the size and make this introductory version more affordable to students. Another was to capitalize on our experience in using the first edition of the book to make an even better introductory text.

The chapters in this book are not just a subset of those in the complete book. We believe that instructors of an introductory database course should have the option of enriching an introductory course by including material on object databases and XML—topics that are covered in great detail in several chapters in the complete book. Therefore we have added to the introductory book two new chapters, Chapter 16, Introduction to Object Databases, and Chapter 17, Introduction to XML and Web Data, which contain an appropriately chosen subset of the material in the full version of this book.

To keep the book up-to-date with the rapidly changing technology, we have added a substantial amount of material on UML to a number of chapters and have included a new chapter on Database Tuning, Chapter 12, in both the introductory and complete books.

As with the first edition, our focus is on how to build applications using databases rather than on how to build the database management system itself. We believe that many more students will be implementing applications than will be building DBMSs. Thus, we include substantial material describing the languages and APIs used by transactions to access a database, such as embedded SQL, ODBC, and JDBC.

Although we cover many practical aspects of database and transaction processing applications, we are primarily concerned with the concepts that underlie these topics rather than with the details of particular commercial systems or applications.

Thus we concentrate on the concepts behind the relational and object data models. These concepts will remain the foundation of database processing long after SQL is obsolete.

To enhance students' understanding of the technical material, we have included a case study of a transaction processing application, the Student Registration System, which is carried through the book. While a student registration system can hardly be considered glamorous, it has the unique advantage that all students have interacted with such a system as users. More importantly, it turns out to be a surprisingly rich application, so we can use it to illustrate many of the issues in database design, query processing, and transaction processing.

A unique aspect of the book is a presentation of the software engineering concepts required to implement transaction processing applications, using the Student Registration System as an example. Since the implementations of many information systems fail because of poor project management and inadequate software engineering, we feel that these topics should be an important part of the student's education. Our treatment of software engineering issues is brief, since many students will take a separate course in this subject. However, we believe that they will be better able to understand and apply that material when they see it presented in the context of an information system implementation. Since the courses that use this text at Stony Brook are not software engineering courses, we do not cover this material in class. Instead, we ask the students to read it and require that they use good software engineering practice in their class projects. We do cover in class those aspects of the Student Registration System that illustrate important issues in databases and transaction processing.

Changes in the Second Edition

The technology underlying database and transaction processing systems is changing so rapidly that we have made a large number of changes and additions to the material of the first edition. One rapidly advancing technology is the Unified Modeling Language, UML. We added substantial amount of material on UML in Chapter 4 on database design, in addition to the material on E-R diagrams that was already there. We also added UML to the material on software engineering in Chapters 2, 14, and 15.

A new chapter on Database Tuning, Chapter 12, was added because so much effort in the real world is spent increasing the throughput of database and transaction processing applications.

In addition, material has been added and updated in almost all the chapters. Significant examples of this are the coverage of SQL/XML and RAID technology.

One important area that is *not* included in this volume is Web Services. Since this is a rapidly developing and interesting application-oriented subject we have significantly revised the compete version of this text to include material on this topic. In addition to strengthening the book on the subject of XML Technology by updating the chapter on XML and Web Data and adding a section on SQL/XML, we have added a new chapter on Web Services that contains material on SOAP,

WSDL, BPEL, UDDI, and XML-based transaction processing using WS-Coordination and WS-Transaction. In the chapter on Security and Internet Commerce, we added a section on XML-based encryption, using XML-Encryption, XML-Signature, WS-Security, and SAML. And in the chapter on Architecture of Transaction Processing Systems, we added material on Web Application Servers and J2EE, which are used to implement the back-end of many Web services.

Organization of the Book

Chapters 1 through 7 should be taught in the order in which they appear in the book. Chapter 8 contains much of the information that students need in order to put the knowledge they acquired in the preceding chapters into practice. However, subsequent chapters do not significantly depend on Chapter 8. Chapters 9 through 12 in Part 3 should be taught sequentially. Chapter 13 in the same part is largely independent. The software engineering chapters in Part 4 utilize the material of the chapters in Parts 2 and 3, but the software engineering chapters can be read in parallel with the database material. Chapters 16 and 17 in the advanced part of the book depend on the first seven chapters in Part 2.

Finally we note that the sections in this book that are marked with an asterisk (*) are optional and can be omitted, if the instructor prefers to do so. Sections marked with the © icon in the table of contents deal with the case study. Also, exercises that are marked with an asterisk are slightly harder than the rest, and exercises that are marked with two asterisks are even harder.

Supplements

In addition to the text, the following supplementary materials are available to assist instructors:

- Online PowerPoint presentations for all chapters
- Online PowerPoint slides of all figures
- An online solution manual containing solutions for the exercises
- Additional references, notes, errata, homeworks, and exams.

For more information on obtaining these supplements, please visit this book's Companion Website at www.aw-bc.com/kifer. The solutions manual and PowerPoint presentations are available only to instructors through your Addison-Wesley sales representative. To contact your representative, please visit www.aw.com.

Acknowledgments

We would like to thank the reviewers, whose comments and suggestions significantly improved the second edition of the book:

Tran Cao Son, New Mexico State University

Frantisek Franek, McMaster University

Junping Sun, Nova Southeastern University
Philip Cannata, Sun Microsystems
Dehu Qi, Lamar University
Nematollah Shiri, Concordia University
Jian Pei, State University of New York at Buffalo
Jack Wileden, University of Massachusetts Amherst
Sibel Adali, Rensselaer Polytechnic Institute
Roger King, University of Colorado at Boulder
Markus Schneider, University of Florida
Yaron Y. Goland, BEA
Dennis Shasha, New York University
Christelle Scharff, Pace University
Zhiwei Wang, Graduate Programs in Software Engineering, IT, and IS University of St. Thomas.

We would also like to thank the reviewers of the first edition of the book:

Suad Alagic, Wichita University
Catriel Beer, The Hebrew University
Rick Cattell, Sun Microsystems
Jan Chomicki, SUNY Buffalo
Henry A. Etlinger, Rochester Institute of Technology
Leonidas Fegaras, University of Texas at Arlington
Alan Fekete, University of Sydney
Johannes Gehrke, Cornell University
Hershel Gottesman, Consultant
Jiawei Han, Simon Fraser University
Peter Honeyman, University of Michigan
Vijay Kumar, University of Missouri–Kansas City
Jonathan Lazar, Towson University
Dennis McLeod, University of Southern California
Rokia Missaoui, University of Quebec in Montreal
Clifford Neuman, University of Southern California
Fabian Pascal, Consultant
Sudha Ram, University of Arizona
Krithi Ramamritham, University of Massachusetts–Amherst, and IIT Bombay
Andreas Reuter, International University in Germany, Bruchsal
Arijit Sengupta, Georgia State University

Munindar P. Singh, North Carolina State University

Greg Speegle, Baylor University

Junping Sun, Nova Southeastern University

Joe Trubisz, Consultant

Vassilis J. Tsotras, University of California, Riverside

Emilia E. Villarreal, California Polytechnic State University

We would also like to thank the following people who were kind enough to provide us with additional information and answers to our questions: Don Chamberlin, Daniela Florescu, Jim Gray, Pankaj Gupta, Rob Kelly, and C. Mohan.

Two people taught out of beta versions of the book and made useful comments and suggestions: David S. Warren and Radu Grosu. Joe Trubicz served not only as a reviewer when the manuscript was complete, but provided critical comments on early versions of many of the chapters.

A number of students were very helpful in reading and checking the correctness of various parts of the book: Ziyang Duan, Shiyong Lu, Swapnil Patil, Guizhen Yang, and Yan Zhang.

Many thanks to the staff of the Computer Science Department at Stony Brook, and in particular Kathy Germana, who helped make things happen at work.

We would particularly like to thank Matt Goldstein and Maite Suarez-Rivas, our editors at Addison-Wesley, who played an important role in shaping the contents and approach of the book in its early stages and throughout the time we were writing it. We would also like to thank the various staff members of Addison-Wesley and Windfall Software, who did an excellent job of editing and producing the book: Jeffrey Holcomb, Paul Anagnostopoulos, Elisabeth Beller, Jennifer McClain, and Joe Snowden.

Last, but not least, we would like to thank our wives, Lora, Edie, and Rhoda, who provided much needed support and encouragement while we were writing the book.

Contents

Preface	xvii
PART ONE Introduction	1
1 Overview of Databases and Transactions	3
1.1 What Are Databases and Transactions?	3
1.2 Features of Modern Database and Transaction Processing Systems	6
1.3 Major Players in the Implementation and Support of Database and Transaction Processing Systems	7
1.4 Decision Support Systems—OLAP and OLTP	9
2 The Big Picture	13
2.1 Case Study: A Student Registration System	13
2.2 Introduction to Relational Databases	14
2.3 What Makes a Program a Transaction—The ACID Properties	20
Bibliographic Notes	25
Exercises	25
PART TWO Database Management	29
3 The Relational Data Model	31
3.1 What Is a Data Model?	31
3.2 The Relational Model	35
3.2.1 Basic Concepts	35
3.2.2 Integrity Constraints	38
3.3 SQL—Data Definition Sublanguage	46
3.3.1 Specifying the Relation Type	46
3.3.2 The System Catalog	46
3.3.3 Key Constraints	47
3.3.4 Dealing with Missing Information	48

3.3.5	Semantic Constraints	49
3.3.6	User-Defined Domains	53
3.3.7	Foreign-Key Constraints	53
3.3.8	Reactive Constraints	56
3.3.9	Database Views	59
3.3.10	Modifying Existing Definitions	60
3.3.11	SQL-Schemas	62
3.3.12	Access Control	63
	Bibliographic Notes	65
	Exercises	66

4	Conceptual Modeling of Databases with Entity-Relationship Diagrams and the Unified Modeling Language	69
4.1	Conceptual Modeling with the E-R Approach	70
4.2	Entities and Entity Types	70
4.3	Relationships and Relationship Types	73
4.4	Advanced Features in Conceptual Data Modeling	78
4.4.1	Entity Type Hierarchies	78
4.4.2	Participation Constraints	81
4.4.3	The Part-of Relationship	83
4.5	From E-R Diagrams to Relational Database Schemas	86
4.5.1	Representation of Entities	86
4.5.2	Representation of Relationships	88
4.5.3	Representing IsA Hierarchies in the Relational Model	90
4.5.4	Representation of Participation Constraints	92
4.5.5	Representation of the Part-of Relationship	94
4.6	UML: A New Kid on the Block*	95
4.6.1	Representing Entities in UML	96
4.6.2	Representing Relationships in UML	97
4.6.3	Advanced Modeling Concepts in UML	101
4.6.4	Translation to SQL	105
4.7	A Brokerage Firm Example	106
4.7.1	An Entity-Relationship Design	106
4.7.2	A UML Design*	110
cs 4.8	Case Study: A Database Design for the Student Registration System	111
4.8.1	The Database Part of the Requirements Document	112
4.8.2	The Database Design	113
4.9	Limitations of Data Modeling Methodologies	119
	Bibliographic Notes	123
	Exercises	123

5 Relational Algebra and SQL 127

- 5.1 Relational Algebra: Under the Hood of SQL 128
 - 5.1.1 Basic Operators 128
 - 5.1.2 Derived Operators 137
- 5.2 The Query Sublanguage of SQL 147
 - 5.2.1 Simple SQL Queries 148
 - 5.2.2 Set Operations 154
 - 5.2.3 Nested Queries 157
 - 5.2.4 Quantified Predicates 163
 - 5.2.5 Aggregation over Data 164
 - 5.2.6 Join Expressions in the FROM Clause 170
 - 5.2.7 A Simple Query Evaluation Algorithm 171
 - 5.2.8 More on Views in SQL 174
 - 5.2.9 Materialized Views 177
 - 5.2.10 The Null Value Quandary 181
- 5.3 Modifying Relation Instances in SQL 182
 - 5.3.1 Inserting Data 182
 - 5.3.2 Deleting Data 184
 - 5.3.3 Updating Existing Data 185
 - 5.3.4 Updates on Views 185
- Bibliographic Notes 187
- Exercises 188

6 Database Design with the Relational Normalization Theory 193

- 6.1 The Problem of Redundancy 193
- 6.2 Decompositions 195
- 6.3 Functional Dependencies 198
- 6.4 Properties of Functional Dependencies 200
- 6.5 Normal Forms 207
 - 6.5.1 The Boyce-Codd Normal Form 208
 - 6.5.2 The Third Normal Form 210
- 6.6 Properties of Decompositions 211
 - 6.6.1 Lossless and Lossy Decompositions 212
 - 6.6.2 Dependency-Preserving Decompositions 215
- 6.7 An Algorithm for BCNF Decomposition 219
- 6.8 Synthesis of 3NF Schemas 221
 - 6.8.1 Minimal Cover 222
 - 6.8.2 3NF Decomposition through Schema Synthesis 224
 - 6.8.3 BCNF Decomposition through 3NF Synthesis 226
- 6.9 The Fourth Normal Form 228

- 6.10 Advanced 4NF Design* 233
 - 6.10.1 MVDs and Their Properties 234
 - 6.10.2 The Difficulty of Designing for 4NF 235
 - 6.10.3 A 4NF Decomposition How-To 238
- 6.11 Summary of Normal Form Decomposition 240
- ⑧ 6.12 Case Study: Schema Refinement for the Student Registration System 241
- 6.13 Tuning Issues: To Decompose or Not to Decompose? 244
 - Bibliographic Notes 245
 - Exercises 246

7 Triggers and Active Databases 251

- 7.1 What Is a Trigger? 251
- 7.2 Semantic Issues in Trigger Handling 252
- 7.3 Triggers in SQL:1999 256
- 7.4 Avoiding a Chain Reaction 264
 - Bibliographic Notes 265
 - Exercises 265

8 Using SQL in an Application 267

- 8.1 What Are the Issues Involved? 267
- 8.2 Embedded SQL 268
 - 8.2.1 Status Processing 271
 - 8.2.2 Sessions, Connections, and Transactions 273
 - 8.2.3 Executing Transactions 274
 - 8.2.4 Cursors 276
 - 8.2.5 Stored Procedures on the Server 282
- 8.3 More on Integrity Constraints 285
- 8.4 Dynamic SQL 286
 - 8.4.1 Statement Preparation in Dynamic SQL 287
 - 8.4.2 Prepared Statements and the Descriptor Area* 290
 - 8.4.3 Cursors 293
 - 8.4.4 Stored Procedures on the Server 293
- 8.5 JDBC and SQLJ 294
 - 8.5.1 JDBC Basics 294
 - 8.5.2 Prepared Statements 297
 - 8.5.3 Result Sets and Cursors 297
 - 8.5.4 Obtaining Information about a Result Set 300
 - 8.5.5 Status Processing 300