

PEARSON

MANAGEMENT  
SCIENCE AND  
ENGINEERING CLASSICS

管理科学与工程经典教材

Science and Engineering Classics

# 现代数据库管理

(英文版·第10版)

## Modern Database Management

(Tenth Edition)

杰弗里·A·霍弗 (Jeffrey A. Hoffer)

V·拉梅什 (V. Ramesh)

著

海基·托皮 (Heikki Topi)

PEARSON

中国人民大学出版社



管理科学与工程经典教材

Science and Engineering Classics

# 现代数据库管理

(英文版·第10版)

## Modern Database Management

(Tenth Edition)

杰弗里·A·霍弗 (Jeffrey A. Hoffer)

V·拉梅什 (V. Ramesh)

著

海基·托皮 (Heikki Topi)

中国人民大学出版社

· 北京 ·

图书在版编目 ( CIP ) 数据

现代数据库管理：第10版：英文 / 霍弗等著. —北京：中国人民大学出版社，2012.6

管理科学与工程经典教材

ISBN 978-7-300-15981-2

I. ①现… II. ①霍… III. ①数据库管理系统-高等学校-教材-英文 IV. ①TP311.13

中国版本图书馆 CIP 数据核字 (2012) 第128192号

管理科学与工程经典教材

现代数据库管理 ( 英文版 · 第10版 )

杰弗里 · A · 霍弗

V · 拉梅什 著

海基 · 托皮

---

出版发行 中国人民大学出版社

社 址 北京中关村大街31号

邮政编码 100080

电 话 010-62511242 ( 总编室 )

010-62511398 ( 质管部 )

010-82501766 ( 邮购部 )

010-62514148 ( 门市部 )

010-62515195 ( 发行公司 )

010-62515275 ( 盗版举报 )

网 址 [http:// www. crup. com. cn](http://www.crup.com.cn)

[http:// www. ttrnet. com](http://www.ttrnet.com) ( 人大教研网 )

经 销 新华书店

印 刷 北京东君印刷有限公司

规 格 215 mm × 275 mm 16开本

版 次 2012 年 6 月第 1 版

印 张 37 插页1

印 次 2012 年 6 月第 1 次印刷

字 数 990 000

定 价 75.00 元

---

版权所有

侵权必究

印装差错

负责调换

# 总 序

随着我国加入 WTO,越来越多的国内企业参与到国际竞争中来,用国际上通用的语言思考、工作、交流的能力也越来越受到重视。这样一种能力也成为我国各类人才参与竞争的一种有效工具。国家教育机构、各类院校以及一些主要的教材出版单位一直在思考,如何顺应这一发展潮流,推动各层次人员通过学习来获取这种能力。双语教学就是这种背景下的一种尝试。

双语教学在我国主要指汉语和国际通用的英语教学。事实上,双语教学在我国教育界已经不是一个陌生的词汇了,以双语教学为主的科研课题也已列入国家“十五”规划的重点课题。但从另一方面来看,双语教学从其诞生的那天起就被包围在人们的赞成与反对声中。如今,依然是有人赞成有人反对,但不论是赞成居多还是反对占上,双语教学的规模 and 影响都在原有的基础上不断扩大,且呈大发展之势。一些率先进行双语教学的院校在实践中积累了经验,不断加以改进;一些待进入者也在模仿中学习,并静待时机成熟时加入这一行列。由于我国长期缺乏讲第二语言(包括英语)的环境,开展双语教学面临特殊的困难,因此,选用合适的教材就成为双语教学成功与否的一个重要问题。我们认为,双语教学从一开始就应该使用原版的各类学科的教材,而不是由本土教师自编的教材,从而可以避免中国式英语问题,保证语言的原汁原味。各院校除应执行国家颁布的教学大纲和课程标准外,还应根据双语教学的特点和需要,适当调整教学课时的设置,合理选择优秀的、合适的双语教材。

顺应这样一种大的教育发展趋势,中国人民大学出版社同众多国际知名的大出版公司,如麦格劳-希尔出版公司、培生教育出版公司等合作,面向大学本科生层次,遴选了一批国外最优秀的管理类原版教材,涉及专业基础课,人力资源管理、市场营销及国际化管理等专业方向课,并广泛听取有着丰富的双语一线教学经验的教师的建议和意见,对原版教材进行了适当的改编,删减了一些不适合我国国情和不适合教学的内容;另一方面,根据教育部对双语教学教材篇幅合理、定价低的要求,我们更是努力区别于目前市场上形形色色的各类英文版、英文影印版的大部头,将目标受众锁定在大学本科生层次。本套教材尤其突出了以下一些特点:

- 保持英文原版教材的特色。本套双语教材根据国内教学实际需要,对原书进行了一定的改编,主要是删减了一些不适合教学以及不符合我国国情的内容,但在体系结构和内容特色方面都保持了原版教材的风貌。专家们的认真改编和审定,使本套教材既保持了学术上的完整性,又贴近中国实际;既方便教师教学,又方便学生理解和掌握。



● 突出管理类专业教材的实用性。本套教材既强调学术的基础性，又兼顾应用的广泛性；既侧重让学生掌握基本的理论知识、专业术语和专业表达方式，又考虑到教材和管理实践的紧密结合，有助于学生形成专业的思维能力，培养实际的管理技能。

● 体系经过精心组织。本套教材在体系架构上充分考虑到当前我国在本科教育阶段推广双语教学的进度安排，首先针对那些课程内容国际化程度较高的学科进行双语教材开发，在其专业模块内精心选择各专业教材。这种安排既有利于我国教师摸索双语教学的经验，使得双语教学贴近现实教学的需要；也有利于我们收集关于双语教学教材的建议，更好地推出后续的双语教材及教辅材料。

● 篇幅合理，价格相对较低。为适应国内双语教学内容和课时上的实际需要，本套教材进行了一定的删减和改编，使总体篇幅更为合理；而采取低定价，则充分考虑到了学生实际的购买能力，从而使本套教材得以真正走近广大读者。

● 提供强大的教学支持。依托国际大出版公司的力量，本套教材为教师提供了配套的教辅材料，如教师手册、PowerPoint 讲义、试题库等，并配有内容极为丰富的网络资源，从而使教学更为便利。

本套教材是在双语教学教材出版方面的一种尝试。我们在选书、改编及出版的过程中得到了国内许多高校的专家、教师的支持和指导，在此深表谢意。同时，为使我们后续推出的教材更适于教学，我们也真诚地期待广大读者提出宝贵的意见和建议。需要说明的是，尽管我们在改编的过程中已加以注意，但由于各教材的作者所处的政治、经济和文化背景不同，书中内容仍可能有不妥之处，望读者在阅读时注意比较和甄别。

徐二明  
中国人民大学商学院

# PREFACE

This text is designed to be used with an introductory course in database management. Such a course is usually required as part of an information systems curriculum in business schools, computer technology programs, and applied computer science departments. The Association for Information Systems (AIS), the Association for Computing Machinery (ACM), and the International Federation of Information Processing Societies (IFIPS) curriculum guidelines (e.g., IS 2010) all outline this type of database management course. Previous editions of this text have been used successfully for more than 27 years at both the undergraduate and graduate levels, as well as in management and professional development programs.

## WHAT'S NEW IN THIS EDITION?

This 10th edition of *Modern Database Management* updates and expands materials in areas undergoing rapid change due to improved managerial practices, database design tools and methodologies, and database technology. Later we detail changes to each chapter. The themes of this 10th edition reflect the major trends in the information systems field and the skills required of modern information systems graduates:

- Data quality and database processing accuracy, which are extremely important with the national and international regulations such as the Sarbanes-Oxley Act, Basel II, COSI, and HIPAA that now require organizations to comply with standards for reporting accurate financial data and ensuring data privacy. Material on data quality and master data management has been updated with a stronger coverage of the people, process, and technology aspects and internationally accepted best practices for information systems development and management (specifically, ITIL).
- Integration of data from multiple internal and external databases and data sources, which is now common for building data warehouses and other types of enterprise systems, and dealing with the rapid organizational changes in information systems brought on by corporate reorganizations, mergers, and acquisitions. These first two bullets are implemented with the revised Chapter 10 on data quality and integration, which updates and improves the focus of the material and introduces the latest principles in these areas.
- Demonstrating knowledge of how to use databases in the context of developing database applications in two and three-tier client/server environments. In this 10th edition (in Chapters 8 and 14), we provide examples of how to connect to databases from popular programming languages such as Java and VB.NET as well Web development languages such as Java Server Pages (JSP), ASP.NET, and PHP. Coverage of XML has also been revised to emphasize the role of XML in data storage and retrieval.
- Linking object-oriented information systems development environments (such as Java Technology and Microsoft .NET) with mainstream technology for maintaining organizational data—relational databases—and in the process dealing with significant paradigm differences between object-oriented and relational frameworks. This major change that was introduced for the ninth edition and has been updated for the 10th edition reflects what is a rapidly changing environment for database processing.

Also, we are very excited to now provide on the student Companion Web site several new, custom-developed short videos that address key concepts and skills from different sections of the book. These videos, produced using Camtasia by the textbook authors, help students to learn difficult material by using both the printed text and a mini lecture or tutorial. Videos have been developed to support Chapters 1 (introduction to database), 2 and 3 (conceptual data modeling), 4 (normalization), and 6 and 7 (SQL).



More will be produced with future editions. Look for special icons on the opening page of these chapters to call attention to these videos, and go to [www.pearsonhighered.com/hoffer](http://www.pearsonhighered.com/hoffer) to find these videos.

Specific improvements to the textbook have been made in the following areas:

- Arranged the Problems and Exercises into roughly increasing order of difficulty to make it easier for instructors and students to select problems and exercises for practice and assignments.
- Applied standard data naming conventions throughout the book to make it easier for students to distinguish data elements from conceptual to physical forms.
- Clarified system requirements through systems modeling and design and outlined a process to use the increasingly popular industry and business function commercial data models to speed up the systems development process. The new material focuses on changes to the database development process when an organization uses packaged data models. Students are now better prepared to understand why these data models are important and how to read and work with (tailor) them.
- Expanded coverage of SQL, with a few more frequently used components of the language. We have also created new figures to graphically depict the set processing logic of SQL queries, which gives students, especially visual learners, new tools to use when writing queries.
- Included new screen captures to reflect the latest database technologies and an updated Web Resources section in each chapter that lists Web sites that can provide the student with information on the latest database trends and expanded background details on important topics covered in the text.
- Reduced the length of the printed book, which we began doing with the eighth edition. The reduced length is more consistent with what our reviewers say can be covered in a database course today, given the need for depth of coverage in the most important topics. Specifically, for the 10th edition, we combined the first two chapters from the ninth edition into one, so that students can more quickly cover/review background topics and then dig into the material central to database management. We have also combined the two chapters from the ninth edition on client/server and Internet databases into one chapter addressing database issues in a multitier computing environment. We continue to update the chapters on distributed databases, the object-oriented data model, and using relational databases to provide object persistence, including an overview in the printed textbook and full versions on the textbook's Web site. Care has been given to the layout of figures and tables to also reduce the length of the book, while adding some new figures and figure elements to better link the text narrative with the figures. The reduced length should encourage more students to purchase and read the text, without any loss of coverage and learning. The book is also now available through CourseSmart, an innovative e-book delivery system.

## **MODERN DATABASE MANAGEMENT: A RETRO AND FUTURE PERSPECTIVE**

This 10th edition is a humbling milestone. We are extremely grateful for the support of adopters, reviewers, students, colleagues, editors, and publisher staff who have been with us for some or, in a few cases, all of the past 27 years. Database technology has "grown up" over these years, from a resource for only the most sophisticated organizations to being a mainstay of almost any computing environment. Some topics, such as relational databases, have been a central part of the text from the beginning; other topics, such as data warehousing, business intelligence, object-oriented databases, and databases on the Internet, are newer topics. Database management used to be able to be explained in 531 pages that were about 80 percent the size of current pages, and now it takes 624 larger pages (really, we aren't just wordier). One of the original authors of this text is still a co-author, while a newer generation of database academic

experts now contributes to these pages with zest and creativity. The original book authors were educated in fields other than business information systems, whereas today our newer authors are experienced and educated in this rich field central to the success of modern organizations.

As a book that we believe has succeeded in leading the database management textbook market, this book is positioned to continue (in some printed or electronic form) for at least another 27 years. Writing this book has been and remains an awesome responsibility. We authors realize that the course that this text supports will be the foundation for student careers with databases. Over the years, we've seen students reading our book on airplanes while traveling on business, and, believe it or not, reading it on a Florida beach during spring break. The authors remain committed to presenting material with sound pedagogy, including topics (both easy and difficult, traditional and emerging) that are critical for the practical success of database professionals, and being informed by research that reveals what will be the "next big thing" in database management. It is in this spirit that we celebrate our milestone edition, and lay the foundation for many more editions to come.

## FOR THOSE NEW TO MODERN DATABASE MANAGEMENT

*Modern Database Management* has been a leading text since its first edition in 1983. In spite of this market leadership position, some instructors have used other good database management texts. Why might you want to switch at this time? There are several good reasons to switch to *Modern Database Management*, including:

- One of our goals, in every edition, has been to lead other books in coverage of the latest principles, concepts, and technologies. See what we have added for the 10th edition in "What's New in This Edition." In the past, we have led in coverage of object-oriented data modeling and UML, Internet databases, data warehousing, and the use of CASE tools in support of data modeling. For the 10th edition, we are taking the lead on database development for Internet-based applications, data quality and integration, the linking of object-oriented development environments with relational databases, and the increasingly important role of packaged database model as a component of agile, rapid development of information systems. We also have for the first time Camtasia-produced tutorial videos to accompany the book, with more to come for future editions.
- While remaining current, this text focuses on what leading practitioners say is most important for database developers. We work with many practitioners, including the professionals of the Data Management Association (DAMA) and The Data Warehousing Institute (TDWI), leading consultants, technology leaders, and authors of articles in the most widely read professional publications. We draw on these experts to ensure that what the book includes is important and covers not only important entry-level knowledge and skills, but also those fundamentals and mindsets that lead to long-term career success.
- In this highly successful book in its 10th edition, material is presented in a way that has been viewed as very accessible to students. Our methods have been refined through continuous market feedback for over 27 years, as well as through our own teaching. Overall, the pedagogy of the book is sound. We use many illustrations that help to make important concepts and techniques clear. We use the most modern notations. The organization of the book is flexible, so you can use chapters in whatever sequence makes sense for your students. We supplement the book with data sets to facilitate hands-on, practical learning, and with new media resources to make some of the more challenging topics more engaging.
- You may have particular interest in introducing SQL early in your course. Our text can accommodate this. First, we cover SQL in depth, devoting two full chapters to this core technology of the database field. Second, we include many SQL examples in early chapters. Third, many instructors have successfully used the two SQL chapters early in their course. Although logically appearing in the life cycle of systems development as Chapters 6 and 7, part of the implementation section of the text, many instructors have used these chapters immediately after



Chapter 1 or in parallel with other early chapters. Finally, we use SQL throughout the book, for example, to illustrate Web application connections to relational databases in Chapter 8, online analytical processing in Chapter 9, and accessing relational databases from object-oriented development environments in Chapter 14.

- We have the latest in supplements and Web site support for the text. See the supplement package for details on all the resources available to you and your students.
- This text is written to be part of a modern information systems curriculum with a strong business systems development focus. Topics are included and addressed so as to reinforce principles from other typical courses, such as systems analysis and design, networking, Web site design and development, MIS principles, and computer programming. Emphasis is on the development of the database component of modern information systems and on the management of the data resource. Thus, the text is practical, supports projects and other hands-on class activities, and encourages linking database concepts to concepts being learned throughout the curriculum the student is taking.

## SUMMARY OF ENHANCEMENTS TO EACH CHAPTER

The following sections present a chapter-by-chapter description of the major changes in this edition. Each chapter description presents a statement of the purpose of that chapter, followed by a description of the changes and revisions that have been made for the 10th edition. Each paragraph concludes with a description of the strengths that have been retained from prior editions.

### Part I: The Context of Database Management

**CHAPTER 1: THE DATABASE ENVIRONMENT AND DEVELOPMENT PROCESS** This chapter discusses the role of databases in organizations and previews the major topics in the remainder of the text. This chapter has undergone extensive reorganization for the 10th edition because it is a consolidation of two previous chapters, allowing students to more quickly cover material that previews the rest of the book. After presenting a brief introduction to the basic terminology associated with storing and retrieving data, the chapter presents a well organized comparison of traditional file-processing systems and modern database technology. The chapter then introduces the core components of a database environment and the range of database applications that are currently in use within organizations—personal, two-tier, multitier, and enterprise applications. The explanation of enterprise databases includes databases that are part of enterprise resource planning systems and data warehouses. A brief history of the evolution of database technology, from pre-database files to modern object-relational technologies, is also presented. The chapter then goes on to explain the process of database development in the context of structured life cycle, prototyping, and agile methodologies. The presentation remains consistent with the companion systems analysis text by Hoffer, George, and Valacich. The chapter also discusses important issues in database development, including management of the diverse group of people involved in database development and frameworks for understanding database architectures and technologies (e.g., the three-schema architecture). Reviewers frequently note the compatibility of this chapter with what students learn in systems analysis and design classes.

### Part II: Database Analysis

**CHAPTER 2: MODELING DATA IN THE ORGANIZATION** This chapter presents a thorough introduction to conceptual data modeling with the entity-relationship (E-R) model. The chapter title emphasizes the reason for the entity-relationship model: to unambiguously document the rules of the business that influence database design. Specific subsections explain in detail how to name and define elements of a data model, which are essential

in developing an unambiguous E-R diagram. In the 10th edition, we have provided some new problems and exercises, improved an example from Pine Valley Furniture to show the relationship of conceptual database design with implementation, and provided more annotations in figures to better highlight key elements and better link text to figures. The chapter continues to proceed from simple to more complex examples, and it concludes with a comprehensive E-R diagram for the Pine Valley Furniture Company.

**CHAPTER 3: THE ENHANCED E-R MODEL** This chapter presents a discussion of several advanced E-R data model constructs, primarily supertype/subtype relationships. A major change in this chapter is the elimination of the section on business rules, which many adopters and reviewers said they did not have time to cover in class. The most significant addition is a more thorough description of how to conduct a data modeling project when using a packaged data model; this new material better prepares students for working with commercial off-the-shelf (COTS) software and purchased data models, which support major efforts to implement patterns and reusability for application deployment in organizations. As in Chapter 2, figures have been improved with more annotations to clarify important data modeling structures. The chapter continues to present a thorough coverage of supertype/subtype relationships and includes a comprehensive example of an extended E-R data model for the Pine Valley Furniture Company.

### Part III: Database Design

**CHAPTER 4: LOGICAL DATABASE DESIGN AND THE RELATIONAL MODEL** This chapter describes the process of converting a conceptual data model to the relational data model, as well as how to merge new relations into an existing normalized database. It provides a conceptually sound and practically relevant introduction to normalization, emphasizing the importance of the use of functional dependencies and determinants as the basis for normalization. Concepts of normalization and normal forms are extended in Appendix B. The chapter features a discussion of the characteristics of foreign keys and introduces the important concept of a nonintelligent enterprise key. Enterprise keys (also called surrogate keys for data warehouses) are being emphasized as some concepts of object-orientation migrate into the relational technology world. A number of new review questions and problems and exercises are included, and revision also has further clarified the presentation of some of the key concepts. The chapter continues to emphasize the basic concepts of the relational data model and the role of the database designer in the logical design process.

**CHAPTER 5: PHYSICAL DATABASE DESIGN AND PERFORMANCE** This chapter describes the steps that are essential in achieving an efficient database design, with a strong focus on those aspects of database design and implementation that are typically within the control of a database professional in a modern database environment. The revised chapter is significantly shorter than the previous one, but we believe the reduction in length has been achieved without loss of significant content. The cuts are in areas that are either not relevant anymore because of changes in technology or not directly related to database design (e.g., storage technologies). Consequently, the chapter has a stronger and clearer focus on the core concepts. Several new review questions and problems and exercises are included. The chapter contains an emphasis on ways to improve database performance, with references to specific techniques available in Oracle and other DBMSs to improve database processing performance. The discussion of indexes includes descriptions of the types of indexes (primary and secondary indexes, join index, hash index table) that are widely available in database technologies as techniques to improve query processing speed. Appendix C provides excellent background on fundamental data structures for programs of study that need coverage of this topic. The chapter continues to emphasize the physical design process and the goals of that process.

## Part IV: Implementation

**CHAPTER 6: INTRODUCTION TO SQL** This chapter presents a thorough introduction to the SQL used by most DBMSs (SQL:1999) and introduces the changes that are included in the latest proposed standard (SQL:200n). The coverage of SQL is extensive and divided into this and the next chapter. This chapter includes examples of SQL code, using mostly SQL:1999 and SQL:200n syntax, as well as some Oracle 11g and Microsoft SQL Server syntax. Some unique features of MySQL are mentioned. Views, both dynamic and materialized, are also covered. Chapter 6 explains the SQL commands needed to create and maintain a database and to program single-table queries. The history and SQL technology environment sections have been streamlined for the 10th edition. Coverage of dual-table, IS NULL/IS NOT NULL, more built-in functions, derived tables, and rules for aggregate functions and the GROUP BY clause are included or improved. New problems and exercises have been added to the chapter. It continues to use the Pine Valley Furniture Company case to illustrate a wide variety of practical queries and query results.

**CHAPTER 7: ADVANCED SQL** This chapter continues the description of SQL, with a careful explanation of multiple-table queries, transaction integrity, data dictionaries, triggers and stored procedures (the differences between which are now more clearly explained), and embedded SQL in other programming language programs. All forms of the OUTER JOIN command are covered. Standard SQL is also used in Chapter 7. This chapter illustrates how to store the results of a query in a derived table, the CAST command to convert data between different data types, and the CASE command for doing conditional processing in SQL. The chapter reduces its coverage of online analytical processing (OLAP) features of SQL:200n, which are also covered in Chapter 9. A new section on self-joins and an explanation of when to use EXISTS (NOT EXISTS) versus IN (NOT IN) are now included. The explanation of cursors with embedded SQL is enhanced. Emphasis continues on the set-processing style of SQL compared with the record-processing of programming languages with which the student may be familiar. New problems and exercises have been added to the chapter. The chapter continues to contain a clear explanation of subqueries and correlated subqueries, two of the most complex and powerful constructs in SQL.

**CHAPTER 8: DATABASE APPLICATION DEVELOPMENT** This chapter provides a modern discussion of the concepts of client/server architecture and applications, middleware, and database access in contemporary database environments. Technologies that are commonly used to create two- and three-tier applications are presented. Many figures are included to show the options in multitiered networks, including application and database servers, database processing distribution alternatives among network tiers, and browser (thin) clients. New to this edition is the presentation of sample application programs that demonstrate how to access databases from popular programming languages such as, Java, VB.NET, ASP.NET, JSP, and PHP. This chapter lays the technology groundwork for the Internet topics presented in the remainder of the text and highlights some of the key considerations in creating three-tier Internet-based applications. The chapter also presents expanded coverage of the role of Extensible Markup Language (XML) and related technologies in data storage and retrieval. Topics covered include basics of XML schemas, XQuery, and XSLT. The chapter concludes with an overview of Web services, associated standards and technologies, and their role in seamless, secure movement of data in Web-based applications. A brief introduction to service-oriented architecture (SOA) is also presented. Security topics, including Web security, are covered in Chapter 11.

**CHAPTER 9: DATA WAREHOUSING** This chapter describes the basic concepts of data warehousing, the reasons data warehousing is regarded as critical to competitive advantage in many organizations, and the database design activities and structures unique to data warehousing. An updated section reviews best practices for determining requirements for a dimensional model. A short new section introduces the emerging column databases technology, which has been developed especially for data warehousing applications.

New exercises provide hands-on practice with a data mart, using SQL and a BI tool called MicroStrategy that is supported on Teradata University Network. Topics include alternative data warehouse architectures and the dimensional data model (or star schema) for data warehouses. Coverage of architectures has been streamlined consistent with trends in data warehousing, and a deep explanation of how to handle slowly changing dimensional data is provided. Operational data store; independent, dependent, and logical data marts; and various forms of online analytical processing (OLAP) are defined (including the SAMPLE SQL command, which is useful for analyzing data from market research activities). User interfaces, including OLAP, data visualization, business performance management and dashboards, and data mining are also described.

## **Part V: Advanced Database Topics**

**CHAPTER 10: DATA QUALITY AND INTEGRATION** This chapter, first introduced in the ninth edition, has been reorganized to better reflect the nature of enterprise data management (EDM) activities in organizations. The principles of data governance, which are at the core of EDM activities, are introduced first. This is followed by coverage of data quality. This chapter describes the need for an active program to manage data quality in organizations and outlines the steps that are considered today to be best practices for data quality management. Quality data are defined, and reasons for poor-quality data are identified. Methods for data quality improvement, such as data auditing, improving data capturing (a key part of database design), data stewardship and governance, TQM principles, modern data management technologies, and high-quality data models are all discussed. The current hot topic of master data management, one approach to integrating key business data, is motivated and explained. Different approaches to data integration are overviewed, and the reasons for each are outlined. The ETL process for data warehousing is discussed in detail. The authors believe that the material covered in this chapter continues to represent a major step forward in database management textbooks.

**CHAPTER 11: DATA AND DATABASE ADMINISTRATION** This chapter presents a thorough discussion of the importance and roles of data and database administration and describes a number of the key issues that arise when these functions are performed. This chapter emphasizes the changing roles and approaches of data and database administration, with emphasis on data quality and high performance. It contains a thorough discussion of database backup procedures, as well as extensively expanded and consolidated coverage of data security threats and responses, and data availability. The data security topics include database security policies, procedures, and technologies (including encryption and smart cards). New to this edition is expanded coverage of the role of databases in Sarbanes-Oxley compliance. We have again added to our discussion of open source DBMS to cover more on the benefits and hazards of this technology and how to choose an open source DBMS. In addition, the topic of heartbeat queries is included in the coverage of database performance improvements. The chapter continues to emphasize the critical importance of data and database management in managing data as a corporate asset.

**CHAPTER 12: DISTRIBUTED DATABASES** This chapter reviews the role, technologies, and unique database design opportunities of distributed databases. The objectives and trade-offs for distributed databases, data replication alternatives, factors in selecting a data distribution strategy, and distributed database vendors and products are covered. This chapter provides thorough coverage of database concurrency access controls. The revision of the chapter introduces several technical updates that are related to the significant advancements in both data management and networking technologies, which form the context for distributed database. An overview of this chapter is included in the printed textbook, and the full version of this chapter has been moved to the textbook's Web site. Many reviewers indicated that they seldom are able to cover this chapter in an introductory course, but having the material available is critical for advanced students or special topics. Having an overview in the printed text with the full chapter available to students provides the greatest flexibility and economy.



**CHAPTER 13: OBJECT-ORIENTED DATA MODELING** This chapter presents an introduction to object-oriented modeling using Object Management Group's Unified Modeling Language (UML). This chapter has been carefully reviewed to ensure consistency with the latest UML notation and best industry practices. UML provides an industry-standard notation for representing classes and objects. The chapter continues to emphasize basic object-oriented concepts, such as inheritance, encapsulation, composition, and polymorphism. The revised version of the chapter also includes several brand-new modeling exercises. As with Chapters 12 and 14, the full version of this chapter is available on the textbook's Web site, with a brief overview included in the printed text.

**CHAPTER 14: USING RELATIONAL DATABASES TO PROVIDE OBJECT PERSISTENCE** This chapter presents an up-to-date approach to how relational databases are used with object-oriented development environments, such as Java EE and Microsoft .NET. Object-oriented and relational approaches have critical design mismatches, which are outlined in the chapter, along with ways database and application developers can deal with these issues. The chapter reviews call-level application program interfaces, SQL query mapping frameworks, and object-relational mapping frameworks as approaches to providing object persistence, which is an essential need in modern development environments that integrate object-oriented development and relational databases. The chapter has been revised to take into account the changing landscape of object-relational mapping (ORM) technologies and the strengthening of the Java Persistence API (JPA) standard. Object-relational mapping is illustrated using the XML mapping files of Hibernate, the most popular ORM framework and the most widely used implementation of the JPA standard. As with Chapters 12 and 13, the full version of this chapter is available on the textbook's Web site, with a brief overview included in the printed text.

## Appendices

The 10th edition contains three appendices intended for those who wish to explore certain topics in greater depth.

**APPENDIX A: DATA MODELING TOOLS AND NOTATION** This appendix addresses a need raised by many readers—how to translate the E-R notation in the text into the form used by the CASE tool or DBMS used in class. Specifically, this appendix compares the notations of CA ERwin Data Modeler r7.3, Oracle Designer 10g, Sybase PowerDesigner 15, and Microsoft Visio Pro 2003. Tables and illustrations show the notations used for the same constructs in each of these popular software packages.

**APPENDIX B: ADVANCED NORMAL FORMS** This appendix presents a description (with examples) of Boyce-Codd and fourth normal forms, including an example of BCNF to show how to handle overlapping candidate keys. Other normal forms are briefly introduced. The Web Resources section includes a reference for information on many advanced normal form topics.

**APPENDIX C: DATA STRUCTURES** This appendix describes several data structures that often underlie database implementations. Topics include the use of pointers, stacks, queues, sorted lists, inverted lists, and trees.

## PEDAGOGY

A number of additions and improvements have been made to chapter-end materials to provide a wider and richer range of choices for the user. The most important of these improvements are the following:

1. **Review Questions** Questions have been updated to support new and enhanced chapter material.
2. **Problems and Exercises** This section has been reviewed in every chapter, and many chapters contain new problems and exercises to support updated chapter material. Of special interest are questions in many chapters that give students

opportunities to use the data sets provided for the text. Also, Problems and Exercises have been re-sequenced into roughly increasing order of difficulty, which should help instructors and students to find exercises appropriate for what they want to accomplish.

3. **Field Exercises** This section provides a set of “hands-on” minicases that can be assigned to individual students or to small teams of students. Field exercises range from directed field trips to Internet searches and other types of research exercises.
4. **Case** The Mountain View Community Hospital (MVCH) case was updated for the 10th edition only to the extent that chapters have been combined from the ninth edition. In each chapter, the case begins with a description of a realistic, modern hospital situation as it relates to that chapter. The case then presents a series of case questions and exercises that focus on specific aspects of the case. The final section includes project assignments, which tie together some issues and activities across chapters. These project assignments can be completed by individual students or by small project teams. This case provides an excellent means for students to gain hands-on experience with the concepts and tools they have studied.
5. **Web Resources** Each chapter contains a list of updated and validated URLs for Web sites that contain information that supplements the chapter. These Web sites cover online publication archives, vendors, electronic publications, industry standards organizations, and many other sources. These sites allow students and instructors to find updated product information, innovations that have appeared since the printing of the book, background information to explore topics in greater depth, and resources for writing research papers.



We have also updated the pedagogical features that helped make the 10th edition widely accessible to instructors and students. These features include the following:

1. **Learning objectives** appear at the beginning of each chapter, as a preview of the major concepts and skills students will learn from that chapter. The learning objectives also provide a great study review aid for students as they prepare for assignments and examinations.
2. **Chapter introductions and summaries** both encapsulate the main concepts of each chapter and link material to related chapters, providing students with a comprehensive conceptual framework for the course.
3. **The chapter review** includes the Review Questions, Problems and Exercises, and Field Exercises discussed earlier, also contains a Key Terms list to test the student’s grasp of important concepts, basic facts, and significant issues.
4. **A running glossary** defines key terms in the page margins as they are discussed in the text. These terms are also defined at the end of the text, in the Glossary of Terms. Also included is the end-of-book Glossary of Acronyms for abbreviations commonly used in database management.

## ORGANIZATION

We encourage instructors to customize their use of this book to meet the needs of both their curriculum and student career paths. The modular nature of the text, its broad coverage, extensive illustrations, and its inclusion of advanced topics and emerging issues make customization easy. The many references to current publications and Web sites can help instructors develop supplemental reading lists or expand classroom discussion beyond material presented in the text. The use of appendices for several advanced topics allows instructors to easily include or omit these topics.

The modular nature of the text allows the instructor to omit certain chapters or to cover chapters in a different sequence. For example, an instructor who wishes to emphasize data modeling may cover Chapter 13 on object-oriented data modeling along with or instead of Chapters 2 and 3. An instructor who wishes to cover only basic entity-relationship concepts (but not the enhanced E-R model) may skip Chapter 3 or cover it after Chapter 4 on the relational model. Three of the advanced topics

chapters—Chapters 12 through 14—are provided in overview form in the printed text and in full version on the book's Companion Web site; this gives the instructor added flexibility to cover these advanced topics at different levels.

We have contacted many adopters of *Modern Database Management* and asked them to share with us their syllabi. Most adopters cover the chapters in sequence, but several alternative sequences have also been successful. These alternatives include:

- Some instructors cover Chapter 11 on data and database administration immediately after Chapter 5 on physical database design and the relational model.
- To cover SQL as early as possible, instructors have effectively covered Chapters 6 and 7 immediately after Chapter 4; some have even covered Chapter 6 immediately after Chapter 1.
- Many instructors have students read appendices along with chapters, such as reading Appendix A on data modeling notations with Chapters 2 or 3 on E-R modeling, Appendix B on advanced normal forms with Chapter 4 on the relational model, and Appendix C on data structures with Chapter 5.

## CASE TOOLS

*Modern Database Management*, 10th edition, offers adopters the option of acquiring outstanding CASE tools software packages from Microsoft and Oracle. Students can purchase this book packaged with the full editions of Microsoft Visio Pro and Oracle 11g at a greatly reduced fee. We are proud to offer such highly valued, powerful software packages to students at such a low cost. These packages can be used to draw data models, generate normalized relations from conceptual data models, and generate database definition code, among other tasks. These tools are also useful in other courses on information systems development.

## THE SUPPLEMENT PACKAGE: [WWW.PEARSONHIGHERED.COM/HOFFER](http://WWW.PEARSONHIGHERED.COM/HOFFER)

A comprehensive and flexible technology support package is available to enhance the teaching and learning experience. All instructor and student supplements are available on the text Web site: [www.pearsonhighered.com/hoffer](http://www.pearsonhighered.com/hoffer).

**FOR STUDENTS** The following online resources are available to students:

- The *Web Resources* module includes the Web links referenced at the end of each chapter in the text to help students further explore database management topics on the Web.
- A full *glossary* is available, along with a glossary of acronyms.
- *Links to sites where students can use our data sets* are provided. Although our data sets are provided in formats that are easily loaded on computers at your university or on student PCs, some instructors will not want the responsibility of supporting local data sets. The application service providers with whom we have developed arrangements (e.g., [www.teradatastudentnetwork.com](http://www.teradatastudentnetwork.com)) provide thin-client interfaces to SQL coding environments. See the text's Web site and the inside front cover for more details.
- *Complete chapters on distributed databases, object-oriented data modeling, and object-oriented development with relational databases* allow you to learn in depth about topics that are overviewed in Chapters 12 through 14 of the textbook.
- *Accompanying databases* are also provided. Two versions of the Pine Valley Furniture Company case have been created and populated for the 10th edition. One version is scoped to match the textbook examples. A second version is fleshed out with more data and tables, as well as sample forms, reports, and modules coded in Visual Basic. This version is not complete, however, so that students can create missing tables and additional forms, reports, and modules. Databases are provided in several formats (ASCII tables, Oracle script, and Microsoft Access), but formats vary for the two versions. Some documentation of the databases is also provided. Both versions of the PVFC database are also provided on Teradata University Network.

- Several new, custom-developed short videos that address key concepts and skills from different sections of the book help students to learn material that may be more difficult to understand by using both the printed text and a mini lecture.



**FOR INSTRUCTORS** The following online resources are available to instructors:

- The *Instructor's Resource Manual* by Chelley Vician, University of St. Thomas, provides chapter-by-chapter instructor objectives, classroom ideas, and answers to Review Questions, Problems and Exercises, Field Exercises, and Project Case Questions. The Instructor's Resource Manual is available for download on the instructor area of the text's Web site.
- The *Test Item File and TestGen* by John P. Russo, Wentworth Institute of Technology, includes a comprehensive set of test questions in multiple-choice, true/false, and short-answer format, ranked according to level of difficulty and referenced with page numbers and topic headings from the text. The Test Item File is available in Microsoft Word and as the computerized TestGen. TestGen is a comprehensive suite of tools for testing and assessment. It allows instructors to easily create and distribute tests for their courses, either by printing and distributing through traditional methods or by online delivery via a local area network (LAN) server. Test Manager features Screen Wizards to assist you as you move through the program, and the software is backed with full technical support.
- *PowerPoint presentation slides* by Michel Mitri, James Madison University, feature lecture notes that highlight key terms and concepts. Instructors can customize the presentation by adding their own slides or editing existing ones.
- The *Image Library* is a collection of the text art organized by chapter. It includes all figures, tables, and screenshots (as permission allows) and can be used to enhance class lectures and PowerPoint slides.
- *Accompanying databases* are also provided. Two versions of the Pine Valley Furniture Company case have been created and populated for the 10th edition. One version is scoped to match the textbook examples. A second version is fleshed out with more data and tables, and sample forms, reports, and modules coded in Visual Basic. This version is not complete, however, so that students can create missing tables and additional forms, reports, and modules. Databases are provided in several formats (ASCII tables, Oracle script, and Microsoft Access), but formats vary for the two versions. Some documentation of the databases is also provided. Both versions of the PVFC database are also provided on Teradata University Network.
- A white paper by Willard Baird, Progress Telecom, *Achieving Optimal Database Performance*, provides supplemental reading for students interested in tuning an Oracle database. This paper provides a real-world perspective from a very experienced database administrator. It offers students and instructors an opportunity to consider the differences between material taught in classrooms and the hands-on experience gained through professional database administration.

## MATERIALS FOR YOUR ONLINE COURSE

Pearson Prentice Hall supports our adopters using online courses by providing files ready for upload into Blackboard course management systems for our testing, quizzing, and other supplements. Please contact your local Pearson Prentice Hall representative for further information on your particular course.

## COURSESMART eTEXTBOOK

**CourseSmart** is an exciting new choice for students looking to save money. As an alternative to purchasing the print textbook, students can purchase an electronic version of the same content and save up to 50% off the suggested list price of the print text. With a CourseSmart eTextbook, students can search the text, make notes online, print out reading assignments that incorporate lecture notes, and bookmark important passages for later review. For more information, or to purchase access to the CourseSmart eTextbook, visit [www.coursesmart.com](http://www.coursesmart.com).



## ACKNOWLEDGMENTS

We are grateful to numerous individuals who contributed to the preparation of *Modern Database Management*, 10th edition. First, we wish to thank our reviewers for their detailed suggestions and insights, characteristic of their thoughtful teaching style. As always, analysis of topics and depth of coverage provided by the reviewers was crucial. Our reviewers and others who gave us many useful comments to improve the text include Tamara Babaian, Bentley University; Gary Baram, Temple University; Timothy Bridges, University of Central Oklahoma; Bijoy Bordoloi, Southern Illinois University, Edwardsville; Traci Carte, University of Oklahoma; Wingyan Chung, Santa Clara University; Jon Gant, Syracuse University; Jinzhu Gao, University of the Pacific; Monica Garfield, Bentley University; Rick Gibson, American University; William H. Hochstetler III, Franklin University; Weiling Ke, Clarkson University; Dongwon Lee, Pennsylvania State University; Ingyu Lee, Troy University; Brian Mennecke, Iowa State University; Dat-Dao Nguyen, California State University, Northridge; Lara Preiser-Houy, California State Polytechnic University, Pomona; John Russo, Wentworth Institute of Technology; Ioulia Rytikova, George Mason University; Richard Segall, Arkansas State University; Chelley Vician, University of St. Thomas; and Daniel S. Weaver, Messiah College.

We received excellent input from people in industry, including Todd Walter, Carrie Ballinger, Rob Armstrong, and Dave Schoeff (all of Teradata Corp); Chad Gronbach and Philip DesAutels (Microsoft Corp.); Peter Gauvin (Ball Aerospace); Mridvika Raisinghani (Ernst & Young); Nisha Subramanian (Cook Medical, Inc.); Paul Longhurst (Overstock.com); Derek Strauss (Gavroshe International); Richard Hackathorn (Bolder Technology); and Michael Alexander (Open Access Technology, International).

We also thank Klara Nelson at the University of Tampa, who authored the Mountain View Community Hospital case study. This extensive real-world situation is a notable addition to the text. Linda Jayne, formerly operations manager of Suncoast Hospital, Largo, Florida, provided many relevant stories and validation of the situations faced by Mountain View Community Hospital, introduced in the eighth edition. We appreciate her taking the time to meet with us and to review the existing case.

Thanks are also due Willard Baird, Progress Telecom, a highly experienced Oracle database administrator, who has co-authored the Web site piece *Achieving Optimal Database Performance* with us. Willard is able to convey trends in database administration to us very clearly and in a timely fashion, and his experience with the trend toward more holistic approaches to database optimization that have occurred over the past few years have led to our decision to put up a piece on the Web site for students and instructors who desire a deeper understanding of this area.

We have special admiration for and gratitude to Chelley Vician of University of St. Thomas, author of the *Instructor's Resource Manual*. Chelley has been extremely careful in preparing the *Instructor's Resource Manual* and in the process has helped us clarify and fix various parts of the text. Chelley has added great value to this book. Christina Cooper, Indiana University, has been of considerable assistance to Chelley Vician, to ensure notational changes in the *Instructor's Resource Manual*. We also thank Sven Aelterman, Troy University, for his many excellent suggestions for improvements and clarifications throughout the text. And we thank Pragya Seth, an MS in Information Systems (MSIS) student at Indiana University, for help in reviewing material in several chapters.

We are also grateful to the staff and associates of Pearson Prentice Hall for their support and guidance throughout this project. In particular, we wish to thank Executive Editor Bob Horan, who coordinated the planning for the text, Editorial Project Manager Kelly Loftus, who kept us on track and made sure everything was complete, Production Project Manager Becca Richter, Senior Marketing Manager Anne Fahlgren, Media Project Manager Denise Vaughn, and Marketing Assistant Melinda Jensen. We extend special thanks to Jen Carley at PMG, whose supervision of the production process was excellent.

Finally, we give immeasurable thanks to our spouses, who endured many evenings and weekends of solitude for the thrill of seeing a book cover hang on a den