

MODERN DATABASE MANAGEMENT

SEVENTH EDITION



Jeffrey A. Hoffer • Mary B. Prescott • Fred R. McFadden

MODERN DATABASE MANAGEMENT

SEVENTH

JEFFREY A. HOFFER

University of Dayton

MARY B. PRESCOTT

University of Tampa

FRED R. McFADDEN

University of Colorado—Colorado Springs

PEARSON

Prentice
Hall

Upper Saddle River, New Jersey 07458

Library of Congress Cataloging-in-Publication Data

Hoffer, Jeffrey A.

Modern database management / Jeffrey A. Hoffer, Mary B. Prescott, Fred R. McFadden. —7th ed.

p. cm.

Includes bibliographical references and index.

ISBN 0-13-145320-3 (alk. paper)

I. Database management. I. Prescott, Mary B. II. McFadden, Fred R., 1933 – III. Title.

QA76.9.D3M395 2004

005.74—dc22

2003068969

AVP/Executive Editor: David Alexander

AVP/Executive Editor: Bob Horan

VP/Publisher: Natalie E. Anderson

Project Manager (Editorial): Lori Cerreto

Editorial Assistant: Robyn Goldenberg

Media Project Manager: Joan Waxman

Senior Marketing Manager: Sharon M. Koch

Marketing Assistant: Danielle Torio

Managing Editor (Production): John Roberts

Production Editor: Suzanne Grappi

Production Assistant: Joe DeProspero

Associate Director, Manufacturing: Vincent Scelta

Production Manager: Arnold Vila

Manufacturing Buyer: Michelle Klein

Design Manager: Maria Lange

Cover Design: Jayne Conte

Cover Illustration/Photo: Photodisc/George Dannatt

Manager, Multimedia Production: Christy Mahon

Composition: Laserwords

Full-Service Project Management: Jennifer Welsch, BookMasters, Inc.

Printer/Binder: Courier/Westford

Credits and acknowledgments borrowed from other sources and reproduced, with permission, in this textbook appear on page 688.

Microsoft® and Windows® are registered trademarks of the Microsoft Corporation in the U.S.A. and other countries. Screen shots and icons reprinted with permission from the Microsoft Corporation. This book is not sponsored or endorsed by or affiliated with the Microsoft Corporation.

Copyright © 2005, 2002 by Pearson Education, Inc., Upper Saddle River, New Jersey 07458.

Pearson Prentice Hall. All rights reserved. Printed in the United States of America. This publication is protected by Copyright and permission should be obtained from the publisher prior to any prohibited reproduction, storage in a retrieval system, or transmission in any form or by any means, electronic, mechanical, photocopying, recording, or likewise. For information regarding permission(s), write to: Rights and Permissions Department.

Pearson Prentice Hall™ is a trademark of Pearson Education, Inc.

Pearson® is a registered trademark of Pearson plc

Prentice Hall® is a registered trademark of Pearson Education, Inc.

Pearson Education LTD.

Pearson Education Singapore, Pte. Ltd

Pearson Education, Canada, Ltd

Pearson Education-Japan

Pearson Education Australia PTY, Limited

Pearson Education North Asia Ltd

Pearson Educación de México, S.A. de C.V.

Pearson Education Malaysia, Pte. Ltd



10 9 8 7 6 5 4 3 2 1
ISBN 0-13-145320-3

*To Patty, for her sacrifices, encouragement, and support.
To my students and colleagues for being receptive and critical and for
challenging me to be a better teacher.*

—J.A.H.

*To Larry, for his love, patience, and support of my determination to manage
an administrative position and the writing of this text. To Fred, who
willingly helped me with my challenge by agreeing to revise some of my
assigned chapters. And to Jeff, who inspires and coordinates us so well as we
work to keep this text current.*

—M.B.P.

*To Evelyn, for her patience, love, and support. And to my colleagues and
students, from whom I continue to learn so much.*

—F.R.M.

Preface

This text is designed for 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 of Information Technology Professionals (AITP), Association for Computing Machinery (ACM), and the International Federation of Information Processing Societies (IFIPS) curriculum guidelines (e.g., IS 2002) all outline this type of database management course. Previous editions of our text have been used successfully for 20 years at both the undergraduate and graduate levels, as well as in management and professional development programs.

This text represents a meaningful updating of the sixth edition of *Modern Database Management*. These revisions are necessary to accommodate the technical, managerial, and methodological changes occurring at an ever-increasing pace in this field. However, we have endeavored to retain the best features of our previous editions. We have made every effort to justify the title *Modern Database Management*, which was introduced in the fourth edition.

For Those New to *Modern Database Management*

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

- This text has in every edition led other books in coverage of the latest principles, concepts, and technologies. See below in “New to This Edition” what we have added for the seventh 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.
- While being 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), 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.
- As a highly successful book in its seventh 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 20 years. Overall, the pedagogy of the book is sound. We use many illustrations that help to make important concepts and techniques clear. In this edition, we have put special

emphasis on updating all the end-of-chapter questions and exercises to increase learning for students. The organization of the book is flexible, so you can use chapters in whatever sequence makes sense for your students.

- 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. Finally, many faculty have successfully used the two SQL chapters early in their course. Although logically appearing in the life cycle of systems development as Chapters 7 and 8, part of the implementation section of the text, many faculty have used these chapters immediately after Chapter 2 or in parallel with other early chapters.
- We have the latest in supplements and Website 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, Website 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.

New to This Edition

The seventh 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. The themes of the seventh edition reflect the major trends in the information systems field and the skills required of modern information systems graduates:

- Data security, quality, and availability, which are increasingly important with Web-enabled databases and threats of computer viruses and worms
- Web-enabled systems design and programming within an overall client/server architecture for systems, including Web services
- Large-scale databases and data warehouses
- Clarifying system requirements through thorough systems modeling and design and using industry and business function commercial data models to speed up the systems development process
- Criticality of database performance in an increasingly online environment
- SQL as a standard for database querying

In all the chapters, new screen captures are included to reflect the latest database technologies and an extensively updated Web Resources section lists Websites that can provide the student with information on the latest database trends and expanded background details on important topics covered in the text. Major changes to the text include:

- The two object-oriented database chapters have been updated to show the latest notation and now provide more thorough examples of UML data modeling for the Pine Valley Furniture Company case study.
- The client/server section of the book has been extensively updated to describe the latest technologies and concepts for databases in multitier systems and on the Internet, including Web services and XML.

- The client/server, Internet, and data and database administration chapters have been thoroughly updated to include many new ideas on database security and data quality and availability management.
- More data model illustrations are included in the data modeling chapters and in the appendix on data modeling tools, showing the variety of notations found in practice. The result is that the student should be able to more easily draw data models using whatever tools they have available to them.
- More examples of the extended entity-relationship notation are included, with emphasis on the concepts embodied in the increasingly popular universal data models provided by industry experts for specific industries and business functions.

The following presents 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 from the sixth edition. Each paragraph concludes with a description of the strengths that have been retained from the sixth edition.

Part I: The Context of Database Management

Chapter 1: The Database Environment This chapter discusses the role of databases in organizations and previews the major topics in the remainder of the text. The chapter introduces a revised classification scheme that now recognizes five types of databases—personal, work group, departmental, enterprise, and Internet/intranet/extranet. The explanation of enterprise databases includes databases that are part of enterprise resource planning systems and data warehouses. The chapter updates the discussion of the evolution of database technologies from predatabase files to modern object-relational and Web-enabled systems, including mobile databases; also updated are the discussions of unstructured data and databases in support of business intelligence. The chapter continues to present a well-organized comparison of database technology and conventional file-processing systems.

Chapter 2: The Database Development Process This chapter presents a detailed discussion of the role of database development within the broader context of information systems development. The chapter explains the process of database development for both structured life cycle and prototyping methodologies. The chapter 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). The chapter also continues to emphasize the information engineering methodology in database development, including the role of the enterprise data model. New to this chapter is an introduction to the increasingly popular and standard-setting packaged data models (so-called universal data models) for industries and business functional areas. Reviewers frequently note the compatibility of this chapter with what students learn in systems analysis and design classes.

Part II: Database Analysis

Chapter 3: 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. New to the chapter are coverage of other popular data modeling notations, so students can more easily use whatever tools are available for drawing data models, and an improved illustration

of a bill-of-materials E-R model. 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 4: The Enhanced E-R Model and Business Rules This chapter presents a discussion of several advanced E-R data model constructs. New to the chapter is coverage of other popular data modeling notations so students can more easily use whatever tools are available for drawing data models (in particular, Visio and Oracle Designer-like notations are used throughout). Also new is an in-depth explanation of the use of supertype/subtype relationships and the “party-role” constructs used in universal data models (which were introduced in Chapter 2). The chapter extensively updates coverage of the GUIDE business rules methodology based on the latest guidelines and now shows the structure of these guidelines, which will facilitate student understanding. 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 5: Logical Database Design and the Relational Model This chapter describes the process of converting a conceptual data model to the relational data model. It features an improved 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. The discussion of functional dependencies and normalization has been improved with an extensively revised example from the Pine Valley Furniture Company that more clearly illustrates the role of functional dependencies. 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 6: Physical Database Design and Performance This chapter describes the steps that are essential in achieving an efficient database design. The chapter contains a new 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. New is an in-depth discussion of the merits and hazards of denormalization. The chapter continues to emphasize the physical design process and the goals of that process.

Part IV: Implementation

Chapter 7: Introduction to SQL This chapter presents a thorough introduction to the SQL used by most DBMSs (SQL-92) and introduces the changes that are included in the latest standard (SQL-99). The coverage of SQL is extensive and divided into this and the next chapter. This chapter includes examples of SQL code, using mostly SQL-99 syntax and some Oracle 9i syntax. New are examples that compare SQL-99 and Oracle 9i syntax with the syntax of MS SQL Server and Microsoft Access 2002 SQL Views, both dynamic and materialized, are also covered. Chapter 7 explains the SQL commands needed to create and maintain a database and to program single-table queries. The chapter continues to use the Pine Valley Furniture Company case to illustrate a wide variety of practical queries and query results.

Chapter 8: 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, and embedded SQL in other programming language

programs. All forms of the OUTER JOIN command are covered. Standard SQL is compared with Oracle 9i, Microsoft SQL Server, and Microsoft Access SQL. 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 also outlines the new on-line analytical processing (OLAP) features of SQL-99, which are necessary for SQL to be useful as a data access tool for data warehouses. As in Chapter 7, most SQL code illustrations are written using SQL-99 syntax and some Oracle 9i syntax. New is an emphasis on the set-processing style of SQL compared with the record-processing of programming languages with which the student may be familiar. The chapter continues to contain a clear explanation of subqueries and correlated subqueries, two of the most complex and powerful constructs in SQL.

Chapter 9: The Client/Server Database Environment This chapter provides a thoroughly modern discussion of the client/server architecture, applications, middleware, and client database access in contemporary database environments. This chapter lays the technology groundwork for the Internet topics in the remainder of the text. 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 an expanded coverage of security for Web-enabled databases. The chapter continues with coverage of the three-tier client/server architecture, application partitioning, role of the mainframe, use of parallel computer architectures, middleware, and Microsoft Access 2002 Query-by-Example. Symmetric multiprocessing (SMP) and massively parallel processing (MPP) architectures are described and compared.

Chapter 10: The Internet Database Environment The purpose of this chapter is to describe how databases are connected with Web-based applications. This chapter includes a discussion of scripting languages and embedded SQL in scripts, with examples from ASP and ColdFusion for a shopping cart application (all of the code for these examples appears on the book's Website). The chapter also includes a review of the Internet-related terminology and concepts (such as firewall, proxy server, static and dynamic Web pages, HTML/SGML/XML/XHTML languages, Cascading Style Sheets, Common Gateway Interface, and servlets) necessary to understand how to connect a database to a Web page. The role of Web servers and server-side extensions for database connectivity is addressed. Web security and privacy issues are also covered. New to this edition is an overview of Web services and its associated standards and technologies and its role in seamless, secure movement of data in Web-based applications. There is also a new section on XML, including an illustration of XML and a discussion of its future in database management. This chapter continues to present the networking concepts that are important for connecting databases to Web-based applications.

Chapter 11: 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. Topics include alternative data warehouse architectures, techniques for data transformation and reconciliation, and the dimensional data model (or star schema) for data warehouses. Operational data store; independent, dependent, and logical data marts; and various forms of on-line analytical processing (OLAP) are defined. User interfaces, including OLAP and data mining, are also described. New to this edition is a summary of award-winning best practices of data warehousing that illustrate how leading organizations have drastically reduced the chance of data warehousing project failure. We have also updated the explanations of @ctive (or real-time) data warehousing, which is becoming increasingly common as more Web-based services are provided to customers and suppliers. Also new

is a data warehouse dataset and associated problems and exercises for the Pine Valley Furniture Company.

Part V: Advanced Database Topics

Chapter 12: Data and Database Administration This extensively revised 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 being performed. This chapter emphasizes the changing roles and approaches of data and database administration, with emphasis on tuning the database and queries for improved performance. It contains a thorough discussion of database backup procedures, an extensively expanded coverage of data security threats and responses, and an updated description of managing data quality and availability. The data security topics now include considerable new material on database security policies, procedures, and technologies (including encryption and smart cards). Also new is an introduction to open-source DBMS, as well as an introduction to data and database administration duties for databases related to Web-based and mobile applications. The chapter continues to emphasize the critical importance of data and database management in managing data as a corporate asset.

Chapter 13: 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, along with Chapter 12, provides thorough coverage of database concurrency access controls.

Chapter 14: Object-Oriented Data Modeling This chapter presents an introduction to object-oriented modeling using the Unified Modeling Language (UML) of Booch, Jacobson, and Rumbaugh. This chapter has been updated to illustrate the latest UML notations, including the notation for stereotypes. UML provides an industry-standard notation for representing classes and objects. New to this edition is a UML diagram in Microsoft Visio, representing the Pine Valley Furniture Company database shown in Chapters 3 and 4, that features E-R and EER notation. The chapter continues to emphasize basic object-oriented concepts, such as inheritance and aggregation.

Chapter 15: Object-Oriented Database Development The purpose of this chapter is to show how to translate object-oriented models (explained in Chapter 14) into class, object, relationship, and operation definitions for an object-oriented DBMS. The chapter also introduces the latest format for object definition language (ODL) and object query language (OQL), the standard language for ODBMSs. The chapter includes an object-oriented database definition using ODL for the Pine Valley Furniture database design of the previous chapter. The chapter concludes with an updated survey of ODBMSs—both vendors and products.

Appendices

The seventh edition contains four appendices intended for persons who wish to explore certain topics in greater depth.

Appendix A: E-R Modeling Tools and Notation This extensively updated 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 Visible Analyst 7.5.5, Computer Associates' AllFusion ERwin Data Modeler 4.1, Microsoft Access 2002, Oracle Designer 6i, and Microsoft Visio Pro 2002. 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. The Web Resources section has been expanded to include 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.

Appendix D: Object-Relational Databases This appendix presents a description of object-relational database management systems (ORDBMSs). Topics include features of an ORDBMS, enhanced SQL, advantages of the object-relational approach, and a summary of ORDBMS vendors and products. New examples are included to show how relational DBMSs such as Oracle 9i include object-oriented data specifications.

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* Many new questions have been added to support new chapter material.
2. *Problems and Exercises* This section has been expanded in every chapter and contains many new problems and exercises to support updated chapter material. New sets of questions have been added in many chapters so that students can use the datasets provided for the text.
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. *Project Case* The Mountain View Community Hospital (MVCH) case continues to be included as a student project. In each chapter, the case begins with a brief description of the project as it relates to that chapter. The case then presents a series of project questions and exercises to 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. Some cases include questions for students to use the MVCH dataset provided for the text.
5. *Web Resources* Each chapter contains a list of updated and validated URLs for Websites that contain information that supplements the chapter. These Websites cover online publication archives, vendors, electronic publications, industry standards organizations, and many other sources. These sites allow students and faculty 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 seventh edition widely accessible to instructors and students. These features include the following:

1. *Learning objectives* appear at the beginning of each chapter to preview 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*, which includes the Review Questions, Problems and Exercises, and Field Exercises discussed earlier, also contains *key terms* 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 an 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 inclusion of advanced topics and emerging issues make customization easy. The many references to current publications and Websites 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 14 on object-oriented data modeling along with or instead of Chapters 3 and 4. An instructor who wishes to cover only basic entity-relationship concepts (but not the enhanced E-R model or business rules) may skip Chapter 4 or cover it after Chapter 5 on the relational model.

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:

- Covering Chapter 12 on data and database administration immediately after Chapter 6 on physical database design and the relational model.
- In order to cover SQL as early as possible, instructors have effectively covered Chapters 7 and 8 immediately after Chapter 5; some have even covered Chapter 7 immediately after Chapter 2.
- Others have had students read the material on Query-by-Example in Chapter 9 while covering Chapter 2.
- Many faculty have students read appendices along with chapters, such as reading Appendix A on data modeling notations with Chapters 3 or 4 on E-R modeling, Appendix B on advanced normal forms can be read with Chapter 5 on the relational model, and Appendix C on data structures can be read with Chapter 6.

Case Tools

Modern Database Management, seventh edition, offers adopters the option of acquiring outstanding CASE tools software packages from Microsoft, Oracle, and Visible Systems. Students can purchase this book packaged with the full editions of Microsoft Visio Pro, Oracle Designer, Oracle Forms and Reports (Developer), Personal Oracle, or Visible Analyst 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 also are useful in other courses on information systems development.

The Supplement Package: www.prenhall.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 Website: www.prenhall.com/hoffer.



For Students The following online resources are available to students:

- An *Interactive Study Guide* that contains multiple choice, true/false, and essay questions. Students receive automatic feedback to their answers. Responses to the essay questions and results from the multiple choice and true/false questions can be e-mailed to the instructor after a student finishes a quiz.
- 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.
- *PowerPoint Presentation Slides* feature lecture notes that highlight key terms and concepts.
- A full *glossary* is available both alphabetically and by chapter, along with a glossary of acronyms.
- *Links to two sites where students can use our datasets* are provided, either on servers running Oracle or Teradata DBMSs. Although our datasets are provided to you in formats that are easily loaded on computers at your university or on student PCs, some faculty will not want the responsibility of supporting local datasets. The application service providers with whom we have developed arrangements provide thin-client interfaces to SQL coding environments. See the text's Website for more details.

For Instructors The following online resources are available to instructors:

- The *Instructor's Resource Manual* by John P. Russo, Wentworth Institute of Technology, 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 also available in print and from the faculty area of the text's Website.
- 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 Prentice-Hall 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. Professors 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.

PINE VALLEY FURNITURE



- *Accompanying databases* are also provided. Three versions of the Pine Valley Furniture case have been created and populated for the seventh edition. One version is scoped to match the textbook examples. A second version is fleshed out with sample forms, reports, and modules coded in Visual Basic. This version is not complete, however, so that students may create missing tables and additional forms, reports, and modules. The third version is a data warehousing environment, and special exercises for using this dataset are located on the text's Website. A version of the Mountain View Community Hospital case is also included. Oracle scripts are included to create the tables and insert sample data for both Pine Valley Furniture and Mountain View Community Hospital. Robert Lewis of the University of South Florida has created these datasets and applications for us.
- The *shopping cart application* illustrated in Chapter 10 provides the starting point for some interesting student exercises. The code and documentation for a "bare bones" shopping cart, with either an ASP or a ColdFusion connection, are included on the book's Website. The shopping carts can be run on the student's personal PC without having to establish a live Website. Microsoft's Personal Web Server can be used for the ASP connection, and a free trial download of ColdFusion makes it possible to try out the ColdFusion shopping cart. Most students should be able to read the code and then build enhancements to their shopping cart as they desire. This experiential learning opportunity, aided by simple code that gives the student the ability to examine the connection of a database to the Web and to enhance that application, is a significant strength of the book that addresses the learning styles of many current students.

Materials for Your Online Course

Prentice-Hall supports our adopters using online courses by providing files ready for upload into both WebCT and Blackboard course management systems for our testing, quizzing, and other supplements. Please contact your local Prentice-Hall representative or mis_service@prenhall.com for further information on your particular course.

Acknowledgments

We are grateful to numerous individuals who contributed to the preparation of *Modern Database Management*, seventh 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 include the following:

Kirk Arnett, Mississippi State University
 Rajeev Kaula, Southwest Missouri State University
 Ruth C. King, University of Illinois at Urbana–Champaign
 Alexis Koster, San Diego State University
 Lynn M. Mason, Lubbock Christian University
 Richard S. Segall, Arkansas State University
 Charles Watkins, Villa Julie College
 Saeed Yazdani, Lane College

We received excellent input from people in industry, including: Todd Walter, Carrie Ballinger, Rob Armstrong, and Dave Schoeff (all of Teradata Division of NCR); Patty Melanson, Kelly Carrigan, and Joanna Frankenberry (all of Catalina

Marketing); Don Berndt (University of South Florida); Bernadette Lynch (Oracle Corporation); and Michael Alexander (Open Access Technology, Int'l.).

We also thank Joe Valacich at Washington State University and Joey George at Florida State University for their great insights on the relationship of database development to the overall development of information systems. Their careful attention helps make our book compatible with books used in systems analysis and design courses, such as *Modern Systems Analysis and Design* by Hoffer, George, and Valacich and *Essentials of Systems Analysis and Design* by Valacich, George, and Hoffer (Prentice Hall).

We also thank Atish Sinha at the University of Wisconsin, Milwaukee, who authored the original versions of Chapters 14 and 15 on object-oriented database modeling and implementation. We sincerely appreciate his efforts to integrate these chapters into the whole text by using the Pine Valley Furniture Company and Mountain View Community Hospital cases, to create parallel exercises that contrast object-oriented approaches with entity-relationship and relational approaches, and to provide suggestions concerning all chapters.

This edition of the textbook will be accompanied by more sample database applications that faculty and students can explore than any previous edition. Our thanks for creating these supplements go to Robert Lewis, University of South Florida, for the Pine Valley Furniture Company and Mountain View Community Hospital cases. We also want to thank Michael Alexander, Open Access Technology, Int'l., for creating the ASP and ColdFusion versions of a basic Internet shopping cart.

Laura Biasci, University of South Florida, updated all of the screen captures and explanations of four tools and added a new section on Microsoft VisioPro for the comparisons of five data modeling CASE tools notation conventions for Appendix A. We want to thank her for gaining access to the five software programs so that she could compare and capture the notation conventions. We also thank John Russo of Wentworth Institute of Technology who prepared new and revised existing end-of-chapter questions. His work has been extremely important to making the seventh edition fresh and useful for all adopters.

We are also very grateful to the staff and associates of Prentice Hall for their support and guidance throughout this project. In particular, we wish to thank Executive Editors Bob Horan and David Alexander, who coordinated the planning for the text, Project Manager Lori Cerreto, who kept us on track and made sure everything was complete, Production Editor Suzanne Grappi, Senior Marketing Manager Sharon Koch, and Marketing Assistant Danielle Torio. We extend special thanks to Jennifer Welsch at BookMasters, Inc., whose excellent supervision of editing was invaluable.

Jeff and Mary are also thrilled to have had a chance to continue to work with and learn from Fred McFadden. Fred is the “father” of this text and was instrumental in setting its style for clarity, accessibility, careful organization, and mainstream modern content. Fred’s student-centered philosophy and disciplined writing set the standard for us. Fred has “come out of retirement” and has been extremely helpful in the seventh edition.

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 wall. In particular, we marvel at the commitment of Patty Hoffer, who has lived the lonely life of a textbook author’s spouse through seven editions over twenty years. Now, Larry Prescott has shown his support through three editions, including leading several “contemplation excursions” onto Tampa Bay. Much of the value of this text is due to their patience, encouragement, and love, but we alone bear the responsibility for any errors or omissions between the covers.

Jeffrey A. Hoffer
Mary B. Prescott
Fred R. McFadden

Brief Contents

Part I The Context of Database Management

- 1 The Database Environment 3
- 2 The Database Development Process 38

Part II Database Analysis

- 3 Modeling Data in the Organization 85
- 4 The Enhanced E-R Model and Business Rules 141

Part III Database Design

- 5 Logical Database Design and the Relational Model 187
- 6 Physical Database Design and Performance 238

Part IV Implementation

- 7 Introduction to SQL 287
- 8 Advanced SQL 326
- 9 The Client/Server Database Environment 356
- 10 The Internet Database Environment 398
- 11 Data Warehousing 436

Part V Advanced Database Topics

- 12 Data and Database Administration 495
- 13 Distributed Databases 548
- 14 Object-Oriented Data Modeling 579
- 15 Object-Oriented Database Development 616

Appendices

- A E-R Modeling Tools and Notation 641
- B Advanced Normal Forms 652
- C Data Structures 658
- D Object-Relational Databases 671

Glossary of Acronyms 677

Glossary of Terms 679

Credits 688

Index 689

Contents

Preface

XXV

Part I

The Context of Database Management

Part I Overview 2

1 The Database Environment 3

Learning Objectives 3

Introduction 3

Basic Concepts and Definitions 4

Data 5

Data Versus Information 5

Metadata 6

Database Management Systems 7

Traditional File Processing Systems 7

FILE PROCESSING SYSTEMS AT PINE VALLEY FURNITURE COMPANY 8

Disadvantages of File Processing Systems 9

Program-Data Dependence 9

Duplication of Data 9

Limited Data Sharing 9

Lengthy Development Times 9

Excessive Program Maintenance 9

The Database Approach 10

THE DATABASE APPROACH AT PINE VALLEY FURNITURE COMPANY 10

Enterprise Data Model 10

Relational Databases 12

Implementing Relational Databases 14

Introducing Internet Technology 14

A Database Application 15

The Range of Database Applications 16

Personal Databases 16

Workgroup Databases 17

Department Databases 19

Enterprise Databases 19

Internet, Intranet, and Extranet Databases 21

Summary of Database Applications 22

Advantages of the Database Approach 23

Program-Data Independence 23

Minimal Data Redundancy 23

3