

SECOND EDITION

DATABASE SYSTEMS



DESIGN,
IMPLEMENTATION,
AND
MANAGEMENT

PETER ROB / CARLOS CORONEL

DATABASE SYSTEMS

Design, Implementation, and Management Second Edition

Peter Rob
Carlos Coronel

Middle Tennessee State University



A DIVISION OF COURSE TECHNOLOGY
ONE MAIN STREET, CAMBRIDGE, MA 02142

an International Thomson Publishing company ITP®

Cambridge • Albany • Bonn • Boston • Cincinnati • London • Madrid • Melbourne • Mexico City
New York • Paris • San Francisco • Singapore • Tokyo • Toronto • Washington

Executive Editor	James H. Edwards
Senior Project Manager	Christopher T. Doran
Production Editor	Barbara J. Worth
Manufacturing Coordinator	Tracy Megison
Marketing Coordinator	Daphne J. Meyers
Production Services and Composition	Jonathan Peck Typographers, Ltd.
Interior Design	Wendy Calmenson
Cover Design	Hannus Design

© 1995 by CTI.

A Division of Course Technology – ITP®

For more information contact:

Course Technology

One Main Street
Cambridge, MA 02142

International Thomson Publishing Europe

Berkshire House 168-173
High Holborn
London WC1V 7AA
England

Thomas Nelson Australia

102 Dodds Street
South Melbourne, 3205
Victoria, Australia

Nelson Canada

1120 Birchmount Road
Scarborough, Ontario
Canada M1K 5G4

International Thomson Editores

Campos Eliseos 385, Piso 7
Col. Polanco
11560 Mexico D.F. Mexico

International Thomson Publishing GmbH

Königswinterer Strasse 418
53227 Bonn
Germany

International Thomson Publishing Asia

211 Henderson Road
#05-10 Henderson Building
Singapore 0315

International Thomson Publishing Japan

Hirakawacho Kyowa Building, 3F
2-2-1 Hirakawacho
Chiyoda-ku, Tokyo 102
Japan

All rights reserved. This publication is protected by federal copyright law. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, or be used to make a derivative work (such as translation or adaptation), without prior permission in writing from Course Technology.

Trademarks

Course Technology and the open book logo are registered trademarks and CourseKits is a trademark of Course Technology. Custom Editions is a registered trademark of International Thomson Publishing, Inc.

ITP® The ITP logo is a registered trademark of International Thomson Publishing.

Some of the product names and company names used in this book have been used for identification purposes only and may be trademarks or registered trademarks of their respective manufacturers and sellers.

Disclaimer

CTI reserves the right to revise this publication and make changes from time to time in its content without notice.

ISBN 0-7895-0052-3

Printed in the United States of America

10 9 8 7 6 5 4

To Anne,
who remains my best friend after thirty-four years of marriage. And to our son, Peter William, who turned out to be the man we hoped he would be and who proved his wisdom by making Sheena our daughter in law. May their son, Adam Lee, be as fine a human being as his parents are. I dedicate this book to you, with love.

Peter Rob

To my parents for my upbringing.
To my wife, Victoria, who always brings out the best in me and who keeps me on the right track.
And to my son, Carlos Anthony, whose smiles give me one more reason to be happy.

Carlos Coronel

P R E F A C E

We are grateful that our first edition was well received and that many of its adopters took the time and effort to write comment cards to suggest improvements and, in many cases, to simply let us know that the book really performed its intended task well. That task was and remains to provide a solid and practical foundation for the design, implementation, and management of databases. The point we made in our first edition is just as valid as ever: It is as difficult to build a good building from a bad blueprint as it is to implement and manage a good database from a bad design. Because the database is the information system's core, this realization has spawned much interest in the topics we address in this book.

Good databases are very *practical* things, but their creation depends on understanding the important *theoretical* ideas and concepts that determine just how practical a database will be.

It's not easy to come up with the proper mix of theory and practice for students of database design, but we believe that this book, *Database Systems: Design, Implementation, and Management*, does it.

Database design is challenging enough without adding unnecessary complexity, but neither can it be treated simplistically. This book will help students achieve a high level of proficiency at design and implementation without getting them lost in a theoretical maze.

Begin with Basic Concepts

Part 1 of *Database Systems* offers six detailed chapters on database design concepts. It goes beyond the insular view of the database as an end in itself to look at the database as part of larger information systems, presenting the Database Development Life Cycle as part of the Systems Development Life Cycle.

Designing a good database requires modeling skills. Therefore, we make sure that Entity Relationship (E-R) modeling is covered well in Chapter 4. For example, we examine how important E-R topics such as recursive entities, composite entities, and entity supertypes and subtypes have a bearing on real-world database

design. We then build two complete E-R diagrams, one section at a time, to illustrate the modeling process. Next, we show how the E-R diagram becomes the basis for the database structure. Finally, we develop and illustrate twelve update and delete rules to ensure that the database structure helps to foster data integrity.

Because we are so interested in *practical* database design, Part 1 very carefully develops normalization procedures, tracing a complete procedure to show how, when necessary, good database tables can be created from bad ones.

Part I also includes a detailed chapter on SQL (Standard Query Language), from the creation of the database to the development of complex queries. Although some adopters may choose to limit their coverage to the basics of SQL table creation, we encourage full and detailed SQL coverage for several reasons. First, SQL merits considerable attention because it is the near-standard 4GL in the relational environment. Second, we have discovered that our students tend to be more marketable when they are able to “speak” SQL.

Tackle Detailed, Real-World Cases and Projects

Part 2 presents a detailed, real-world design case, in which the concepts of Part I are put to work. For example, relying on the comprehensive discussion of Entity Relationship modeling in Chapter 4 and the other important design concepts presented in Part 1, we proceed to build, one segment at a time, a complete E-R diagram for a real-world database design.

After the first detailed E-R diagram has been completed in Chapter 7, we *verify* the design in Chapter 8 to ensure that it meets all of the end users’ input and output requirements. The verification process also ensures that the proposed database meets the appropriate data integrity requirements. The verification process, too, illustrates our interest in the very important *practical* real-world aspects of database design.

The design developed in Chapters 7 and 8 traces the evolution of a university computer lab database. We had several good reasons for choosing a university computer lab operation as the source for the case:

1. Most MIS/CIS students have worked in a computer lab and are, therefore, familiar with its basic operations. Database design concepts and principles make more sense to novice designers when they find themselves in relatively familiar territory.
2. Effective lab management is based on principles and procedures that are just as applicable in any business environment. For example, the heart of the lab management system is its inventory module. Lab management must also address maintenance requirements and personnel assignments. And university regulations require a detailed reporting capability, thus allowing us to illustrate how database design is affected by output requirements.

3. Lab management is sufficiently complex to require the use of most of the design principles and procedures presented in Chapters 2 through 6. For example, the lab's conceptual design and its implementation must deal with multivalued attributes, composite entities, supertype/subtype relationships, and so on.

This detailed, nontrivial case study is a guide to full-featured database design and implementation. The wealth of problems presented in Part 2 provide plenty of design and implementation practice. Many of these problems are sufficiently challenging to serve as student design team projects.

Instructors may wish to appoint team leaders and project coordinators to tie each problem's design segments together in a single, full-featured design, so that students can be introduced to a realistic database-design work environment.

Be Up-to-Date in Outlook and Coverage

A look at older real-world databases often reveals that their "design" was driven by narrowly focused information requirements: as the need for information grew, the database was modified to include the additional data, rather like a house being built one room at a time. This process can hardly be described as "design."

Modern information requirements are simply too critical to be entrusted to this approach. *Database Systems* reflects the notion that a database must fit within an information system that serves a company's overall objectives and meets its operational needs. In short, we have learned that the database must serve strategic as well as tactical objectives.

To achieve this dual role, database professionals must broaden their emphasis from applications programming to the database design on which the applications are based. This is why we devote so much attention to good design techniques, the blending of basic systems analysis with careful E-R modeling, and comprehensive examples of real database design.

Database Systems is also up to date in specific topic coverage. In addition to database administration, Part 3 covers four other advanced, current database topics: transaction management, distributed databases, object orientation and the extended relational database, and client/server computing.

Combine Depth and Readability

This book gives careful, detailed coverage of real-world design problems. At the same time, unlike its many competitors, it is *readable*. We believe that even complex ideas can be more easily digested when explained in a conversational style, seasoned with many examples and illustrations. We are confident that our book offers readable, yet detailed and useable coverage.

It also provides carefully thought-out, flexible preview and review summaries highlighting the main points of each chapter. Some chapters require and therefore

receive more summary detail than others. For example, the applied design chapters (7 and 8) include a complete map of the procedures being used in the design, ensuring that readers understand the design details within the bigger picture.

To aid students in review, the book also includes a detailed glossary that frequently illustrates definitions with simple examples.

The following table summarizes the contents of *Database Systems*, from concepts and design to current practice, to a final look back at older database models (Part 4):

Part 1 Concepts and Design

the relational database model: basic concepts and components	Chapter 1	File Systems and Databases
	Chapter 2	The Relational Database model
	Chapter 3	An Introduction to Structured Query Language (SQL)
database design: concepts and procedures	Chapter 4	Entity Relationship (E-R) Modeling
	Chapter 5	Normalization of Database Tables
	Chapter 6	Database Design

Part 2 Practical Guide to Database Design and Implementation

implementing the database design concepts: a case	Chapter 7	The University Lab: Conceptual Design
	Chapter 8	The University Lab: Conceptual Design Verification, Logical Design, and Implementation

Part 3 Advanced Database Topics

complex database environments	Chapter 9	Transaction Management and Concurrency Control
	Chapter 10	Distributed Database Management Systems
new developments, and managing the database	Chapter 11	Object-Oriented Databases
	Chapter 12	Client/Server Systems
	Chapter 13	Database Administration

Part 4 Older Database Models

Chapter 14	The Hierarchical Database Model
Chapter 15	The Network Database Model

Acknowledgments

A book’s second edition allows authors the luxury of polishing the original writing. As we read and reread the first edition and reader comment cards, and as we used it in our classrooms, we reminded ourselves frequently of the need

to add new material, to reorganize and rephrase some existing material to make it more understandable, and to correct the occasional errors that, despite our best efforts, managed to creep into the first edition.

Authors receive the credit for writing books, yet they cannot get the job done without depending on technical, editorial, and moral support from a host of sources. For example, we are convinced that this book's first edition was well received because Frank T. Ruggirello's fingerprints are found all over our work. Although Frank's title was "publisher" in the Wadsworth empire, we know him to be one of the finest editors in the business. Truth compels us to admit that Frank's editorial whip sometimes produced two very tired and occasionally rebellious authors, but the first edition's reviews bear testimony to his sound judgment.

Producing the material for a second edition is a scary thing. After all, without proper editorial and technical support, it is easy to diminish the original work. Fortunately, we discovered that boyd & fraser's Vice President and Executive Editor, Jim Edwards, was as serious about producing an excellent book as we were. Jim's decision to accelerate the second edition's production to a 1995 release turned out to be manageable because he committed the necessary resources, including his valuable leadership. Thanks, Jim, for managing what can only be described as a class act.

The difficult work of bringing the revision together on an accelerated schedule was assigned to Chris Doran, Senior Project Manager, and Barb Worth, Production Editor. It is difficult to envision a finer team. The true mark of professionals is the ability to make the difficult look easy. By that standard and any others we might think of, Chris and Barb wear the professional badge naturally. We can't begin to express our thanks sufficiently for their patience, their constant good cheer, their can do approach, and their obviously high level of expertise. Our appreciation also goes out to Liz Powers, of Jonathan Peck Typographers, who managed the project through typesetting. It was a joy working with her.

Every technical book is reviewed by a group of people who are competent to judge the book's technical merits. Our first edition's reviewers not only provided a much-appreciated green light, but they also provided numerous excellent suggestions concerning technical content and topical coverage. Many of the suggestions caused us to reconsider existing material and to consider the inclusion of suggested topics and techniques. Many of our reviewers will discover that their comments were heeded: By the time the preproduction labors were finished, we had written well over ten thousand pages and explored three quite different sets of design approaches. Some chapters were rewritten five and six times. In spite of the considerable rewriting agony, we consider the final product to be worth the Herculean effort. So, while holding the reviewers blameless for any remaining shortcomings, we owe these reviewers many thanks: Carol Chrisman, Illinois State University; Kevin J. Gorman, University of North Carolina; Timothy Heintz, Marquette University; Herman Hoplin, Syracuse University; Robert Horton, University of Wisconsin, Whitewater; Dean James, Embry-Riddle University; Constance Knapp, Pace University; Bill Korn, University of Wisconsin, Eau Claire; Mary Ann Robbert, Bentley College; Stephen C. Solosky, Nassau Community

College; Francis J. Van Wetering, University of Nebraska; and Joseph Walls, University of Southern California.

The second edition reviewers, too, played an important role in the development of this book. They evaluated the coverage and accuracy of the new topics and they offered useful advice to ensure that the new material would fit within the text's framework. They also gave us the necessary feedback to help us rewrite or rearrange existing material. We thank these reviewers for their comments, suggestions, and encouragement: Roger Chiang, Syracuse University; Crist Costa, Rhode Island College; Sudesh M. Duggal, Northern Kentucky University; and Chang Koh, University of North Carolina at Greensboro.

Any academic author is indebted to students whose classroom struggles, comments, and questions lead to the initial attempt to write a better book in the first place. Their comments and questions caused us to revise, revise, and revise in order to clarify material and even to try completely different approaches. Augustus Lewis Scott, Jr.'s, and Wayne E. Gensemer's independent study project yielded important feedback that caused us to add more explanatory material to the E-R modeling chapter. We also owe thanks to Marino Gallego, the truly outstanding project coordinator for the 1992 database design class. Mr. Gallego *verified* many of the database design and other concepts covered in this book. In addition, Mr. Gallego's excellent 1992 independent study on object-oriented databases uncovered numerous library sources and raised many questions that helped get us started on Chapter 11. Sutee Sujitparapitaya also labored on an object-oriented database independent study and contributed sources and useful questions. The students who used our first edition in 1993 and in the Spring of 1994 were asked to respond to our question "How are we doing?" And they responded magnificently, many of them anonymously. However, Micky Stinnett, who wrote her honors thesis on database design after taking our database course, and Laura Witte, our 1993 outstanding Computer Information Systems Senior, gave particularly detailed comments that helped us tighten up the text's material. Many thanks.

A very special note of appreciation goes to IBM, affectionately known as Big Blue. IBM is a truly classy outfit that employs many classy people, two of whom in particular provided outstanding support. Michael D. Pennington helped us nail down some of DB2's details and proved that he is an exceptionally fine teacher. John Thompson was always available to answer questions, to help chase sources, to set up special training, to lend equipment, and to provide software. Bless you, and bless IBM.

B R I E F C O N T E N T S

Chapter 1	File Systems and Databases	1
Chapter 2	The Relational Database Model	43
Chapter 3	An Introduction to Structured Query Language (SQL)	85
Chapter 4	Entity Relationship (E-R) Modeling	145
Chapter 5	Normalization of Database Tables	209
Chapter 6	Database Design	237
Chapter 7	The University Lab: Conceptual Design	279
Chapter 8	The University Lab: Conceptual Design Verification, Logical Design, and Implementation	313
Chapter 9	Transaction Management and Currency Control	353
Chapter 10	Distributed Database Management Systems	379
Chapter 11	Object-Oriented Databases	423
Chapter 12	Client/Server Systems	481
Chapter 13	Database Administration	541
Chapter 14	The Hierarchical Database Model	589
Chapter 15	The Network Database Model	621
Appendix A	Creating the Database Structure	655
Appendix B	A Relational Schema for dBASE IV Users	661
Appendix C	University Computer Lab Report Formats	663
Appendix D	University Computer Lab Data Dictionary	673
Appendix E	Client/Server Network Infrastructure	683
	Glossary	689

C O N T E N T S

Preface xxi



Chapter 1 File Systems and Databases

1

Preview 1

- 1.1 Introducing the Database 2**
 - 1.1.1 Why Database Design Is Important 3
 - 1.1.2 A Practical Approach to Database Design 3
- 1.2 The Historical Roots of the Database: Files and File Systems 4**
- 1.3 A File System Critique 9**
 - 1.3.1 File System Data Management 10
 - 1.3.2 Structural and Data Dependence 11
 - 1.3.3 Field Definitions 12
 - 1.3.4 Data Duplication 13
- 1.4 Database Systems 14**
 - 1.4.1 Managing the Database System: A Shift in Focus 16
 - 1.4.2 Database Design and Modeling 17
- 1.5 Database Models 18**
 - 1.5.1 The Hierarchical Database Model 19
 - 1.5.2 The Network Database Model 25
 - 1.5.3 The Relational Database Model 30

Summary 34

Review Questions 36

Problems 37

References 40



Chapter 2 The Relational Database Model

43

Preview 43

- 2.1 A Logical View of Data** 44
 - 2.1.1 Entities and Attributes 44
 - 2.1.2 Tables and Their Characteristics 45
- 2.2 Keys: Links Between Tables** 49
- 2.3 Integrity Rules Revisited** 54
- 2.4 Relational Database Operators** 56
- 2.5 Relational Database Software Classification** 62
- 2.6 The Data Dictionary and the System Catalog** 63
- 2.7 Relationships Within the Relational Database** 64
- 2.8 Indexes** 71
 - Summary* 71
 - Review Questions* 73
 - Problems* 75
 - References* 83



Chapter 3 An Introduction to Structured Query Language (SQL)

85

Preview 85

- 3.1 Introduction to SQL** 86
- 3.2 Data Definition Commands** 87
 - 3.2.1 Creating the Database Structure 87
 - 3.2.2 Creating Table Structures 88
 - 3.2.3 SQL Integrity Constraints 94
- 3.3 Basic Data Management** 95
 - 3.3.1 Data Entry 95
 - 3.3.2 Checking the Table Contents 96
 - 3.3.3 Saving the Table Contents 98
 - 3.3.4 Adding Data to the Table 98
 - 3.3.5 Making a Correction 99
 - 3.3.6 Deleting Table Rows 100
 - 3.3.7 Restoring the Table Contents 100

3.4	Queries	101
3.4.1	Partial Listing of Table Contents	101
3.4.2	Logical Operators: AND, OR, and NOT	105
3.4.3	Special Operators	107
3.5	Advanced Data Management Commands	111
3.5.1	Changing a Column's Data Type	112
3.5.2	Changing Attribute Characteristics	112
3.5.3	Adding a Column to the Table	112
3.5.4	Entering Data into the New Column	113
3.5.5	Copying Parts of Tables	116
3.5.6	Deleting a Table from the Database	117
3.6	More Complex Queries and SQL Functions	117
3.6.1	Ordering a Listing	117
3.6.2	Eliminating Duplications	120
3.6.3	SQL's Numeric Functions	121
3.6.4	Grouping Data	125
3.6.5	Virtual Tables: Creating a View	126
3.6.6	SQL Indexes	127
3.6.7	Joining Database Tables	128
	<i>Summary</i>	130
	<i>Review Questions</i>	136
	<i>Problems</i>	139
	<i>References</i>	144



Chapter 4 Entity Relationship (E-R) Modeling

145

Preview 145

4.1	Basic Modeling Concepts	146
4.2	Data Models: Degrees of Data Abstraction	147
4.2.1	The Conceptual Model	147
4.2.2	The Internal Model	150
4.2.3	The External Model	151
4.2.4	The Physical Model	153
4.3	The Entity Relationship (E-R) Model	154
4.3.1	E-R Model Components	154
4.3.2	Entities	155
4.3.3	Attributes	155
4.3.4	Relationships	158

4.3.5	Connectivity	161
4.3.6	Cardinality	162
4.3.7	Existence Dependency	164
4.3.8	Relationship Participation	164
4.3.9	Weak Entities	167
4.3.10	Recursive Entities	167
4.3.11	Composite Entities	169
4.3.12	Entity Supertypes and Subtypes	171
4.3.13	E-R Component Summary	174
4.4	Developing an E-R Diagram	174
4.4.1	Converting an E-R Model into a Database Structure	181
4.4.2	General Rules Governing Relationships Among Tables	186
4.4.3	Converting Tiny College's E-R Model to a SQL Database	194
4.4.4	Building Another E-R Model	199
4.4.5	Considering Additional Requirements	200
	<i>Summary</i>	201
	<i>Review Questions</i>	202
	<i>Problems</i>	205
	<i>References</i>	208



Chapter 5 Normalization of Database Tables

209

Preview 209

5.1	Database Tables and Normalization	210
5.1.1	The Need for Normalization	210
5.1.2	Conversion to First Normal Form	213
5.1.3	Conversion to Second Normal Form	215
5.1.4	Conversion to Third Normal Form	217
5.1.5	The Boyce-Codd Normal Form (BCNF)	219
5.2	Normalization and Database Design	221
5.3	Higher-Level Normal Forms	224
5.4	Denormalization	226
	<i>Summary</i>	227
	<i>Review Questions</i>	229
	<i>Problems</i>	230
	<i>References</i>	235



Chapter 6 Database Design**237***Preview* 237

- 6.1 Changing Data into Information** 238
 - 6.2 The Information System** 239
 - 6.3 The Systems Development Life Cycle (SDLC)** 240
 - 6.3.1 Planning 240
 - 6.3.2 Analysis 242
 - 6.3.3 Detailed Systems Design 242
 - 6.3.4 Implementation 243
 - 6.3.5 Maintenance 243
 - 6.4 The Database Life Cycle (DBLC)** 244
 - 6.4.1 The Database Initial Study 244
 - 6.4.2 The Database Design Process 249
 - 6.4.3 Implementation and Loading 263
 - 6.4.4 Testing and Evaluation 268
 - 6.4.5 Operation 268
 - 6.4.6 Maintenance and Evolution 269
 - 6.5 A Special Note About Database Design Strategies** 269
 - 6.6 Centralized Versus Decentralized Design** 270
- Summary* 274
Review Questions 275
Problems 276
References 277



Chapter 7 The University Lab: Conceptual Design**279***Preview* 279

- 7.1 The Database Initial Study** 281
 - 7.1.1 UCL Objectives 281
 - 7.1.2 Organizational Structure 281
 - 7.1.3 Description of Operations 283
 - 7.1.4 Problems and Constraints 287
 - 7.1.5 System Objectives 289
 - 7.1.6 Scope and Boundaries 291

- 7.2 Database Design Phase: Conceptual Design 293**
 - 7.2.1 Information Sources and Users 293
 - 7.2.2 User Requirements 295
 - 7.2.3 Developing the Initial Entity Relationship Model 297
- Summary 307*
- Review Questions 309*
- Problems 309*



Chapter 8 The University Lab: Conceptual Design Verification, Logical Design, and Implementation

313

- Preview 313*
- 8.1 Completing the Conceptual Design:
Defining Attributes and Domains 314**
- 8.2 Evaluating the E-R Conceptual Model 332**
 - 8.2.1 Normalization 332
 - 8.2.2 E-R Model Verification 333
- 8.3 Logical Design 339**
 - 8.3.1 Tables 339
 - 8.3.2 Indexes and Views 340
- 8.4 Physical Design 340**
- 8.5 Implementation 342**
 - 8.5.1 Database Creation 342
 - 8.5.2 Database Loading and Conversion 343
 - 8.5.3 System Procedures 343
- 8.6 Testing and Evaluation 344**
 - 8.6.1 Performance Measures 344
 - 8.6.2 Security Measures 345
 - 8.6.3 Backup and Recovery Procedures 346
- 8.7 Operation 346**
 - 8.7.1 Database Is Operational 346
 - 8.7.2 Operational Procedures 346
- Summary 347*
- Review Questions 348*
- Problems 350*