

Contents

Preface xi

▶ PART I FUNDAMENTALS 1

- 1 Getting Started 3**
 - Why Use a Database? 4
 - What Is a Database System? 12
 - Summary 19
 - Review Questions 20
 - Exercises 21
 - Garden Glory Project Questions 21
 - James River Jewelry Project Questions 22
 - Mid-Western University Chemistry Department Project Questions 23

- 2 The Relational Model 25**
 - Relations 26
 - Types of Keys 28
 - Functional Dependencies and Normalization 35
 - The Problem of Null Values 44
 - Summary 44
 - Review Questions 45
 - Exercises 46
 - Garden Glory Project Questions 47
 - James River Jewelry Project Questions 48
 - Mid-Western University Chemistry Department Project Questions 49

- 3 Structured Query Language 51**
 - A Sample Database 52
 - SQL for Data Definition 54
 - SQL for Relational Query 59
 - SQL for Relational Data Modification 73
 - Summary 75
 - Review Questions 76

Exercises	78
Garden Glory Project Questions	78
James River Jewelry Project Questions	79
Mid-Western University Chemistry Department Project Questions	80

▶ PART II DATABASE DESIGN AND MANAGEMENT 83

4	Data Modeling and the Entity-Relationship Model	85
	The Requirements Stage	86
	The Entity-Relationship Data Model	87
	The E-R Model in UML	93
	Developing an Example E-R Diagram	94
	Summary	104
	Review Questions	105
	Exercises	106
	Garden Glory Project Questions	107
	James River Jewelry Project Questions	107
	Mid-Western University Chemistry Department Project Questions	107
5	Database Design	109
	Representing Entities with the Relational Model	110
	Representing Relationships	114
	Database Design at Heather Sweeney Designs	124
	Summary	128
	Review Questions	128
	Exercises	130
	Garden Glory Project Questions	130
	James River Jewelry Project Questions	131
	Mid-Western University Chemistry Department Project Questions	131
6	Database Administration	133
	The Database Processing Environment	134
	Concurrency Control	138
	Database Security	148
	Database Backup and Recovery	151
	Summary	156
	Review Questions	157
	Exercises	158
	Garden Glory Project Questions	159
	James River Jewelry Project Questions	160
	Mid-Western University Chemistry Department Project Questions	161

7	Advanced Topics	163
	Web Database Processing	164
	Database Processing and XML	171
	Distributed Database Processing	175
	OLAP and Data Mining	177
	Object-Relational Databases	180
	Summary	181
	Review Questions	182
	Exercises	184
	Appendix: Database Processing with Microsoft Access	185
	Glossary	219
	Index	227

Colin is a production supervisor for a small manufacturer in Seattle. Several years ago, Colin wanted to build a database to keep track of components in product packages. At the time, he was using a spreadsheet to perform this task, but he could not find a way to automate it. He had heard about Microsoft Access, and he tried to use it to solve his problem. After several days of frustration, he bought several popular Access books and attempted to learn from them. Ultimately, he gave up and hired a consultant who built an application that more or less met Colin's needs. Over time, Colin wanted to change his application, but he did not dare try.

Colin was a successful businessperson who was highly motivated to achieve his goals. A seasoned Windows user, he had been able to teach himself how to use Excel, PowerPoint, and a number of production-oriented application packages. He was flummoxed at his inability to use Access to solve his problem. "I'm sure I could do it, but I just don't have any more time to invest," he thought. This story is even more remarkable because it has occurred tens of thousands of times over the last decade.

Microsoft Corporation, Oracle Corporation, and other database management system (DBMS) vendors are aware of such scenarios and have invested millions of dollars in creating better graphical interfaces, hundreds of multi-panel wizards, and many sample applications. Unfortunately, such efforts treat the symptom and not the cause. In fact, most users have no clear idea of what wizards are doing on their behalf. As soon as these users require changes to database structure or to components such as forms and queries, they drown in a sea of complexity for which they are unprepared. With little understanding of the underlying fundamentals, these users grab at any straw that appears to lead in the direction they want. The consequence is poorly designed databases and applications that fail to meet the users' requirements.

Why can people like Colin learn to use a word processor or a spreadsheet product, yet fail when trying to learn to use a DBMS product? First, the underlying database concepts are unnatural to most people. Whereas everyone knows what paragraphs and margins are, no one knows what a relation is. Second, it seems like using a DBMS product ought to be easier than it is. "All I want to do is keep track of something. Why is it so hard?" people ask. Without knowledge of the relational model, breaking a sales invoice into five separate tables before storing the data is mystifying to business users.

► THE NEED FOR ESSENTIAL CONCEPTS

With today's technology, it is impossible to utilize a DBMS successfully without first learning fundamental concepts. After years of developing databases with business users, I believe that the following database concepts are essential.

- Fundamentals of the relational model
- Structured Query Language (SQL)
- Data modeling