

Contents

Introduction	xvii
Part I: Approaching Relational Database Modeling	1
Chapter 1: Database Modeling Past and Present	3
Grasping the Concept of a Database	4
Understanding a Database Model	5
What Is an Application?	5
The Evolution of Database Modeling	6
File Systems	7
Hierarchical Database Model	8
Network Database Model	8
Relational Database Model	9
Relational Database Management System	11
The History of the Relational Database Model	11
Object Database Model	12
Object-Relational Database Model	14
Examining the Types of Databases	14
Transactional Databases	15
Decision Support Databases	15
Hybrid Databases	16
Understanding Database Model Design	16
Defining the Objectives	17
Looking at Methods of Database Design	20
Summary	21
Chapter 2: Database Modeling in the Workplace	23
Understanding Business Rules and Objectives	24
What Are Business Rules?	25
The Importance of Business Rules	26
Incorporating the Human Factor	27
People as a Resource	27
Talking to the Right People	29
Getting the Right Information	30

Contents

Dealing with Unfavorable Scenarios	32
Computerizing a Pile of Papers	32
Converting Legacy Databases	33
Homogenous Integration of Heterogeneous Databases	33
Converting from Spreadsheets	33
Sorting Out a Messed-up Database	34
Summary	34
Chapter 3: Database Modeling Building Blocks	35
Information, Data and Data Integrity	37
Understanding the Basics of Tables	37
Records, Rows, and Tuples	39
Fields, Columns and Attributes	40
Datatypes	42
Simple Datatypes	42
Complex Datatypes	46
Specialized Datatypes	47
Constraints and Validation	47
Understanding Relations for Normalization	48
Benefits of Normalization	49
Potential Normalization Hazards	49
Representing Relationships in an ERD	49
Crows Foot	50
One-to-One	51
One-to-Many	52
Many-to-Many	53
Zero, One, or Many	55
Identifying and Non-Identifying Relationships	57
Understanding Keys	58
Primary Keys	59
Unique Keys	59
Foreign Keys	60
Understanding Referential Integrity	63
Understanding Indexes	64
What Is an Index?	65
Alternate Indexing	65
Foreign Key Indexing	65
Types of Indexes	66
Different Ways to Build Indexes	68
Introducing Views and Other Specialized Objects	69
Summary	70
Exercises	70

Part II: Designing Relational Database Models	71
Chapter 4: Understanding Normalization	73
What Is Normalization?	74
The Concept of Anomalies	74
Dependency, Determinants, and Other Jargon	76
Defining Normal Forms	80
Defining Normal Forms the Academic Way	80
Defining Normal Forms the Easy Way	81
1st Normal Form (1NF)	82
1NF the Academic Way	82
1NF the Easy Way	83
2nd Normal Form (2NF)	89
2NF the Academic Way	89
2NF the Easy Way	89
3rd Normal Form (3NF)	96
3NF the Academic Way	96
3NF the Easy Way	97
Beyond 3rd Normal Form (3NF)	103
Why Go Beyond 3NF?	104
Beyond 3NF the Easy Way	104
One-to-One NULL Tables	104
Beyond 3NF the Academic Way	107
Boyce-Codd Normal Form (BCNF)	108
4th Normal Form (4NF)	111
5th Normal Form (5NF)	116
Domain Key Normal Form (DKNF)	121
Summary	122
Exercises	122
Chapter 5: Reading and Writing Data with SQL	123
Defining SQL	124
The Origins of SQL	125
SQL for Different Databases	125
The Basics of SQL	126
Querying a Database Using SELECT	127
Basic Queries	127
Filtering with the WHERE Clause	130
Precedence	132
Sorting with the ORDER BY Clause	134

Contents

Aggregating with the GROUP BY Clause	135
Join Queries	137
Nested Queries	141
Composite Queries	143
Changing Data in a Database	144
Understanding Transactions	144
Changing Database Metadata	145
Summary	148
Exercises	149
Chapter 6: Advanced Relational Database Modeling	151
Understanding Denormalization	152
Reversing Normal Forms	152
Denormalizing Beyond 3NF	153
Denormalizing 3NF	157
Denormalizing 2NF	160
Denormalizing 1NF	161
Denormalization Using Specialized Database Objects	162
Denormalization Tricks	163
Understanding the Object Model	165
Introducing the Data Warehouse Database Model	167
Summary	169
Exercises	170
Chapter 7: Understanding Data Warehouse Database Modeling	171
The Origin of Data Warehouses	172
The Relational Database Model and Data Warehouses	173
Surrogate Keys in a Data Warehouse	174
Referential Integrity in a Data Warehouse	174
The Dimensional Database Model	175
What Is a Star Schema?	176
What Is a Snowflake Schema?	178
How to Build a Data Warehouse Database Model	182
Data Warehouse Modeling Step by Step	183
How Long to Keep Data in a Data Warehouse?	183
Types of Dimension Tables	184
Understanding Fact Tables	190
Summary	191
Exercises	192

Chapter 8: Building Fast-Performing Database Models	193
The Needs of Different Database Models	194
Factors Affecting OLTP Database Model Tuning	194
Factors Affecting Client-Server Database Model Tuning	195
Factors Affecting Data Warehouse Database Model Tuning	196
Understanding Database Model Tuning	197
Writing Efficient Queries	198
The SELECT Command	200
Filtering with the WHERE Clause	202
The HAVING and WHERE Clauses	204
Joins	205
Auto Counters	206
Efficient Indexing for Performance	206
Types of Indexes	207
How to Apply Indexes in the Real World	207
When Not to Use Indexes	209
Using Views	210
Application Caching	211
Summary	212
Exercises	213
Part III: A Case Study in Relational Database Modeling	215
Chapter 9: Planning and Preparation Through Analysis	217
Steps to Creating a Database Model	219
Step 1: Analysis	219
Step 2: Design	220
Step 3: Construction	220
Step 4: Implementation	220
Understanding Analysis	221
Analysis Considerations	222
Potential Problem Areas and Misconceptions	224
Normalization and Data Integrity	224
More Normalization Leads to Better Queries	224
Performance	224
Generic and Standardized Database Models	225
Putting Theory into Practice	225
Putting Analysis into Practice	225
Company Objectives	226

Contents

Case Study: The OLTP Database Model	229
Establishing Company Operations	229
Discovering Business Rules	232
Case Study: The Data Warehouse Model	243
Establishing Company Operations	244
Discovering Business Rules	248
Project Management	253
Project Planning and Timelines	253
Budgeting	255
Summary	256
Exercises	257
Chapter 10: Creating and Refining Tables During the Design Phase	259
A Little More About Design	260
Case Study: Creating Tables	262
The OLTP Database Model	262
The Data Warehouse Database Model	265
Case Study: Enforcing Table Relationships	269
Referential Integrity	269
Primary and Foreign Keys	270
Using Surrogate Keys	271
Identifying versus Non-Identifying Relationships	272
Parent Records without Children	272
Child Records with Optional Parents	273
The OLTP Database Model with Referential Integrity	274
The Data Warehouse Database Model with Referential Integrity	279
Normalization and Denormalization	282
Case Study: Normalizing an OLTP Database Model	283
Denormalizing 2NF	284
Denormalizing 3NF	285
Denormalizing 1NF	286
Denormalizing 3NF Again	287
Deeper Normalization Layers	289
Case Study: Backtracking and Refining an OLTP Database Model	295
Example Application Queries	298
Case Study: Refining a Data Warehouse Database Model	308
Summary	316
Exercises	317

Chapter 11: Filling in the Details with a Detailed Design	319
Case Study: Refining Field Structure	320
The OLTP Database Model	320
The Data Warehouse Database Model	323
Understanding Datatypes	329
Simple Datatypes	329
ANSI (American National Standards Institute) Datatypes	330
Microsoft Access Datatypes	331
Specialized Datatypes	331
Case Study: Defining Datatypes	332
The OLTP Database Model	332
The Data Warehouse Database Model	336
Understanding Keys and Indexes	338
Types of Indexes	339
What, When, and How to Index	342
When Not to Create Indexes	342
Case Study: Alternate Indexing	343
The OLTP Database Model	343
The Data Warehouse Database Model	345
Summary	352
Exercises	352
Chapter 12: Business Rules and Field Settings	353
What Are Business Rules Again?	354
Classifying Business Rules in a Database Model	355
Normalization, Normal Forms, and Relations	355
Classifying Relationship Types	356
Explicitly Declared Field Settings	357
Storing Code in the Database	358
Stored Procedure	360
Stored Function	362
Event Trigger	363
External Procedure	364
Macro	364
Case Study: Implementing Field Level Business Rules in a Database Model	364
Table and Relation Level Business Rules	364
Individual Field Business Rules	364
Field Level Business Rules for the OLTP Database Model	364
Field Level Business Rules for the Data warehouse Database Model	370

Contents

Encoding Business Rules	373
Encoding Business Rules for the OLTP Database Model	373
Encoding Business Rules for the Data Warehouse Database Model	374
Summary	379
Part IV: Advanced Topics	381
Chapter 13: Advanced Database Structures and Hardware Resources	383
Advanced Database Structures	384
What and Where?	384
Views	384
Materialized Views	384
Indexes	385
Clusters	385
Auto Counters	385
Partitioning and Parallel Processing	385
Understanding Views	386
Understanding Materialized Views	387
Understanding Types of Indexes	390
BTree Index	391
Bitmap Index	392
Hash Keys and ISAM Keys	393
Clusters, Index Organized Tables, and Clustered Indexes	393
Understanding Auto Counters	393
Understanding Partitioning and Parallel Processing	393
Understanding Hardware Resources	396
How Much Hardware Can You Afford?	396
How Much Memory Do You Need?	396
Understanding Specialized Hardware Architectures	396
RAID Arrays	397
Standby Databases	397
Replication	399
Grids and Computer Clustering	400
Summary	401
Glossary	403
Appendix A: Exercise Answers	421
Appendix B: Sample Databases	435
Index	443