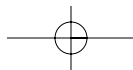
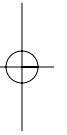
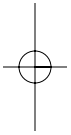
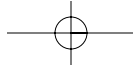


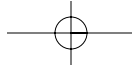
Contents at a Glance

Foreword	xv
About the Authors	xvii
About the Technical Reviewers	xix
Acknowledgments	xxi
Introduction	xxiii
CHAPTER 1 Introduction to Database Concepts.....	1
CHAPTER 2 Data Modeling	33
CHAPTER 3 Conceptual Data Modeling	71
CHAPTER 4 The Normalization Process	121
CHAPTER 5 Implementing the Base Table Structures	181
CHAPTER 6 Protecting the Integrity of Your Data	273
CHAPTER 7 Securing Access to the Data	335
CHAPTER 8 Table Structures and Indexing	395
CHAPTER 9 Coding for Concurrency	439
CHAPTER 10 Code-Level Architectural Decisions	489
CHAPTER 11 Database Interoperability	541
APPENDIX A Codd's 12 Rules for an RDBMS	573
APPENDIX B Datatype Reference	581
INDEX	613

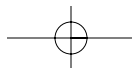


Contents

Foreword	xv
About the Authors	xvii
About the Technical Reviewers	xix
Acknowledgments	xxi
Introduction	xxiii
CHAPTER 1 Introduction to Database Concepts	1
Database Design Phases	2
Conceptual	3
Logical	4
Implementation	5
Physical	6
Relational Data Structures	6
Database and Schema	7
Tables, Rows, and Columns	7
The Information Principle	11
Domains	13
Metadata	14
Keys	14
Missing Values (NULLs)	21
Relationships	22
Foreign Keys	22
Types of Relationships	23
Data Access Language (SQL)	28
Understanding Dependencies	29
Functional Dependency	30
Determinant	31
Multivalued Dependency	31
Summary	32



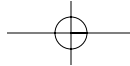
CHAPTER 2	Data Modeling	33
	Introduction to Data Modeling	33
	Entities	35
	Entity Naming	36
	Attributes	38
	Primary Key	39
	Alternate Keys	41
	Foreign Keys	42
	Domains	43
	Naming	45
	Relationships	46
	Identifying Relationship	47
	Nonidentifying Relationship	48
	Optional Identifying Relationship	49
	Cardinality	51
	Role Names	52
	Other Types of One-to-N Relationships	54
	Subtypes	56
	Many-to-Many Relationship	58
	Verb Phrases (Relationship Names)	60
	Descriptive Information	62
	Alternative Modeling Methodologies	64
	Information Engineering	64
	Chen ERD	66
	Management Studio Database Diagrams	67
	Best Practices	68
	Summary	69
CHAPTER 3	Conceptual Data Modeling	71
	Understanding the Requirements	72
	Documenting the Process	73
	Requirements Gathering	75
	Client Interviews	75
	Questions to Be Answered	76
	Existing Systems and Prototypes	80
	Other Types of Documentation	81
	Identifying Objects and Processes	82
	Identifying Entities	84
	Relationships Between Entities	92
	Identifying Attributes and Domains	99



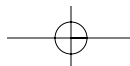
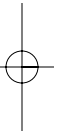
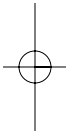
Identifying Business Rules and Processes	112
Identifying Business Rules	112
Identifying Fundamental Processes	114
Finishing the Conceptual Model	116
Identifying Obvious Additional Data Needs	116
Review with the Client	117
Repeat Until the Customer Agrees with Your List of Objects	118
Best Practices	118
Summary	119
CHAPTER 4 The Normalization Process	121
Why Normalize?	122
Eliminating Duplicated Data	122
Avoiding Unnecessary Coding	122
Keeping Tables Thin	122
Maximizing Clustered Indexes	123
Lowering the Number of Indexes Per Table	123
How Far to Normalize	123
The Process of Normalization	124
Entity and Attribute Shape: First Normal Form	125
All Attributes Must Be Atomic	125
All Instances in an Entity Must Contain the Same	
Number of Values	129
All Occurrences of an Entity Type in an Entity	
Must Be Different	131
Programming Anomalies Avoided by First Normal Form	132
Clues That Existing Data Is Not in First Normal Form	136
Relationships Between Attributes	137
Second Normal Form	138
Third Normal Form	144
Boyce-Codd Normal Form	151
Multivalued Dependencies in Entities	155
Fourth Normal Form	156
Fifth Normal Form	169
Denormalization	171
Best Practices	171
Summary	172
Bonus Example	173
The Story of the Book So Far	179

CHAPTER 5	Implementing the Base Table Structures	181
	The Design Process	181
	Reviewing the Logical Design	185
	Transforming the Design	186
	Naming Concerns	186
	Dealing with Subtypes	190
	Choosing Primary Keys	195
	Domain Specification	201
	Setting up Schemas	213
	Reviewing the “Final” Implementation Model	214
	Property Tables	215
	Implementing the Design	217
	Basic Table Creation	218
	Uniqueness Keys	228
	Default Constraints	233
	Relationships (Foreign Keys)	239
	Large-Value Datatype Columns	251
	Collation (Sort Order)	253
	Computed Columns	255
	Implementing Complex Datatypes	257
	Documentation	266
	Best Practices	270
	Summary	271
CHAPTER 6	Protecting the Integrity of Your Data	273
	Best Practices	274
	Constraints	276
	Example Schema	277
	Basic Syntax	278
	Constraints Based on Functions	283
	Handling Errors Caused by Constraints	286
	Programmatic Data Protection	289
	DML Triggers	289
	Stored Procedures	326
	Programmatic Data Protection Outside the RDBMS	329
	More Best Practices	332
	Summary	333
	The Continuing Story of the Book So Far	333

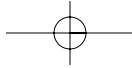
CHAPTER 7	Securing Access to the Data	335
	Controlling Data Access	337
	Principals and Securables	337
	Database Security Overview	339
	Controlling Object Access Via Coded Objects	357
	Views and Table-Valued Functions	370
	Obfuscating Data	377
	Keeping an Eye on Users	380
	Watching Table History Using Triggers	381
	DDL Triggers	385
	Logging with Profiler	388
	Best Practices	391
	Summary	392
CHAPTER 8	Table Structures and Indexing	395
	Physical Database Structure	396
	Files and Filegroups	396
	Extents and Pages	399
	Indexes Overview	402
	Basic Index Structure	402
	Index Types	404
	Basics of Index Creation	409
	Basic Index Usage	411
	Advanced Index Usage Scenarios	427
	Foreign Key Indexes	428
	Using Indexed Views to Optimize Denormalizations	432
	Best Practices	435
	Summary	436
CHAPTER 9	Coding for Concurrency	439
	What Is Concurrency?	439
	Query Optimization Basics	441
	OS and Hardware Issues	443
	Transactions	444
	Transaction Syntax	445
	Compiled SQL Server Code	453
	SQL Server Concurrency Controls	460
	Locks	460
	Isolation Levels	465



Coding for Integrity and Concurrency	475
Pessimistic Locking	476
Optimistic Locking	478
Logical Unit of Work	485
Best Practices	487
Summary	488
CHAPTER 10 Code-Level Architectural Decisions	489
Data-Access Strategies	489
Ad Hoc SQL	490
Stored Procedures	501
Opinions	512
Choosing Between T-SQL and CLR	514
Good Reasons to Use .NET	515
Hosting the CLR	516
Using the .NET CLR for SQL Server Objects	518
Guidelines and Opinions	536
Best Practices	537
Summary	538
CHAPTER 11 Database Interoperability	541
Step One: Datatypes	541
Step Two: Identifier Rules	543
Step Three: Basic SQL Statements	545
The DELETE Statement	546
The INSERT Statement	549
The SELECT Statement	552
The UPDATE Statement	557
Step Four: Creating Database Objects	560
Creating Tables	560
Creating Indexes	564
Creating Views	566
Creating Triggers	567
Creating Procedures and Functions	570
Best Practices	571
Summary	572



APPENDIX A	Codd's 12 Rules for an RDBMS	573
	Rule 1: The Information Rule	573
	Rule 2: Guaranteed Access Rule	574
	Rule 3: Systematic Treatment of NULL Values	574
	Rule 4: Dynamic On-Line Catalog Based on the Relational Model	575
	Rule 5: Comprehensive Data Sublanguage Rule	575
	Rule 6: View Updating Rule	576
	Rule 7: High-Level Insert, Update, and Delete	576
	Rule 8: Physical Data Independence	576
	Rule 9: Logical Data Independence	577
	Rule 10: Integrity Independence	578
	Rule 11: Distribution Independence	579
	Rule 12: Non-Subversion Rule	579
	Summary	580
APPENDIX B	Datatype Reference	581
	Precise Numeric Data	582
	Integer Numbers	583
	Decimal Values	585
	Approximate Numeric Data	589
	Date and Time Data	591
	smalldatetime	591
	datetime	591
	Using User-Defined Datatypes to Manipulate Dates and Times	592
	Character Strings	594
	char(length)	595
	varchar(length)	596
	varchar(max)	596
	text	598
	Unicode Character Strings: nchar, nvarchar, nvarchar(max), ntext	598
	Binary Data	598
	binary(length)	599
	varbinary(length)	600
	varbinary(max)	600
	image	601



xiv ■ **CONTENTS**

Other Datatypes	601
rowversion (a.k.a. timestamp)	601
uniqueidentifier	602
cursor	605
table	605
XML	607
sql_variant Data	607
Summary	611

■ INDEX	613
----------------------	-----

