

# Table of Contents

<b>1</b>	<b>The Worlds of Database Systems</b>	<b>1</b>
1.1	The Evolution of Database Systems . . . . .	1
1.1.1	Early Database Management Systems . . . . .	2
1.1.2	Relational Database Systems . . . . .	3
1.1.3	Smaller and Smaller Systems . . . . .	3
1.1.4	Bigger and Bigger Systems . . . . .	4
1.1.5	Information Integration . . . . .	4
1.2	Overview of a Database Management System . . . . .	5
1.2.1	Data-Definition Language Commands . . . . .	5
1.2.2	Overview of Query Processing . . . . .	5
1.2.3	Storage and Buffer Management . . . . .	7
1.2.4	Transaction Processing . . . . .	8
1.2.5	The Query Processor . . . . .	9
1.3	Outline of Database-System Studies . . . . .	10
1.4	References for Chapter 1 . . . . .	12
<b>I</b>	<b>Relational Database Modeling</b>	<b>15</b>
<b>2</b>	<b>The Relational Model of Data</b>	<b>17</b>
2.1	An Overview of Data Models . . . . .	17
2.1.1	What is a Data Model? . . . . .	17
2.1.2	Important Data Models . . . . .	18
2.1.3	The Relational Model in Brief . . . . .	18
2.1.4	The Semistructured Model in Brief . . . . .	19
2.1.5	Other Data Models . . . . .	20
2.1.6	Comparison of Modeling Approaches . . . . .	21
2.2	Basics of the Relational Model . . . . .	21
2.2.1	Attributes . . . . .	22
2.2.2	Schemas . . . . .	22
2.2.3	Tuples . . . . .	22
2.2.4	Domains . . . . .	23
2.2.5	Equivalent Representations of a Relation . . . . .	23

2.2.6	Relation Instances . . . . .	24
2.2.7	Keys of Relations . . . . .	25
2.2.8	An Example Database Schema . . . . .	26
2.2.9	Exercises for Section 2.2 . . . . .	28
2.3	Defining a Relation Schema in SQL . . . . .	29
2.3.1	Relations in SQL . . . . .	29
2.3.2	Data Types . . . . .	30
2.3.3	Simple Table Declarations . . . . .	31
2.3.4	Modifying Relation Schemas . . . . .	33
2.3.5	Default Values . . . . .	34
2.3.6	Declaring Keys . . . . .	34
2.3.7	Exercises for Section 2.3 . . . . .	36
2.4	An Algebraic Query Language . . . . .	38
2.4.1	Why Do We Need a Special Query Language? . . . . .	38
2.4.2	What is an Algebra? . . . . .	38
2.4.3	Overview of Relational Algebra . . . . .	39
2.4.4	Set Operations on Relations . . . . .	39
2.4.5	Projection . . . . .	41
2.4.6	Selection . . . . .	42
2.4.7	Cartesian Product . . . . .	43
2.4.8	Natural Joins . . . . .	43
2.4.9	Theta-Joins . . . . .	45
2.4.10	Combining Operations to Form Queries . . . . .	47
2.4.11	Naming and Renaming . . . . .	49
2.4.12	Relationships Among Operations . . . . .	50
2.4.13	A Linear Notation for Algebraic Expressions . . . . .	51
2.4.14	Exercises for Section 2.4 . . . . .	52
2.5	Constraints on Relations . . . . .	58
2.5.1	Relational Algebra as a Constraint Language . . . . .	59
2.5.2	Referential Integrity Constraints . . . . .	59
2.5.3	Key Constraints . . . . .	60
2.5.4	Additional Constraint Examples . . . . .	61
2.5.5	Exercises for Section 2.5 . . . . .	62
2.6	Summary of Chapter 2 . . . . .	63
2.7	References for Chapter 2 . . . . .	65
<b>3</b>	<b>Design Theory for Relational Databases</b> . . . . .	<b>67</b>
3.1	Functional Dependencies . . . . .	67
3.1.1	Definition of Functional Dependency . . . . .	68
3.1.2	Keys of Relations . . . . .	70
3.1.3	Superkeys . . . . .	71
3.1.4	Exercises for Section 3.1 . . . . .	71
3.2	Rules About Functional Dependencies . . . . .	72
3.2.1	Reasoning About Functional Dependencies . . . . .	72
3.2.2	The Splitting/Combining Rule . . . . .	73

3.2.3	Trivial Functional Dependencies . . . . .	74
3.2.4	Computing the Closure of Attributes . . . . .	75
3.2.5	Why the Closure Algorithm Works . . . . .	77
3.2.6	The Transitive Rule . . . . .	79
3.2.7	Closing Sets of Functional Dependencies . . . . .	80
3.2.8	Projecting Functional Dependencies . . . . .	81
3.2.9	Exercises for Section 3.2 . . . . .	83
3.3	Design of Relational Database Schemas . . . . .	85
3.3.1	Anomalies . . . . .	86
3.3.2	Decomposing Relations . . . . .	86
3.3.3	Boyce-Codd Normal Form . . . . .	88
3.3.4	Decomposition into BCNF . . . . .	89
3.3.5	Exercises for Section 3.3 . . . . .	92
3.4	Decomposition: The Good, Bad, and Ugly . . . . .	93
3.4.1	Recovering Information from a Decomposition . . . . .	94
3.4.2	The Chase Test for Lossless Join . . . . .	96
3.4.3	Why the Chase Works . . . . .	99
3.4.4	Dependency Preservation . . . . .	100
3.4.5	Exercises for Section 3.4 . . . . .	102
3.5	Third Normal Form . . . . .	102
3.5.1	Definition of Third Normal Form . . . . .	102
3.5.2	The Synthesis Algorithm for 3NF Schemas . . . . .	103
3.5.3	Why the 3NF Synthesis Algorithm Works . . . . .	104
3.5.4	Exercises for Section 3.5 . . . . .	105
3.6	Multivalued Dependencies . . . . .	105
3.6.1	Attribute Independence and Its Consequent Redundancy . . . . .	106
3.6.2	Definition of Multivalued Dependencies . . . . .	107
3.6.3	Reasoning About Multivalued Dependencies . . . . .	108
3.6.4	Fourth Normal Form . . . . .	110
3.6.5	Decomposition into Fourth Normal Form . . . . .	111
3.6.6	Relationships Among Normal Forms . . . . .	113
3.6.7	Exercises for Section 3.6 . . . . .	113
3.7	An Algorithm for Discovering MVD's . . . . .	115
3.7.1	The Closure and the Chase . . . . .	115
3.7.2	Extending the Chase to MVD's . . . . .	116
3.7.3	Why the Chase Works for MVD's . . . . .	118
3.7.4	Projecting MVD's . . . . .	119
3.7.5	Exercises for Section 3.7 . . . . .	120
3.8	Summary of Chapter 3 . . . . .	121
3.9	References for Chapter 3 . . . . .	122

<b>4</b>	<b>High-Level Database Models</b>	<b>125</b>
4.1	The Entity/Relationship Model . . . . .	126
4.1.1	Entity Sets . . . . .	126
4.1.2	Attributes . . . . .	126
4.1.3	Relationships . . . . .	127
4.1.4	Entity-Relationship Diagrams . . . . .	127
4.1.5	Instances of an E/R Diagram . . . . .	128
4.1.6	Multiplicity of Binary E/R Relationships . . . . .	129
4.1.7	Multiway Relationships . . . . .	130
4.1.8	Roles in Relationships . . . . .	131
4.1.9	Attributes on Relationships . . . . .	134
4.1.10	Converting Multiway Relationships to Binary . . . . .	134
4.1.11	Subclasses in the E/R Model . . . . .	135
4.1.12	Exercises for Section 4.1 . . . . .	138
4.2	Design Principles . . . . .	140
4.2.1	Faithfulness . . . . .	140
4.2.2	Avoiding Redundancy . . . . .	141
4.2.3	Simplicity Counts . . . . .	142
4.2.4	Choosing the Right Relationships . . . . .	142
4.2.5	Picking the Right Kind of Element . . . . .	144
4.2.6	Exercises for Section 4.2 . . . . .	145
4.3	Constraints in the E/R Model . . . . .	148
4.3.1	Keys in the E/R Model . . . . .	148
4.3.2	Representing Keys in the E/R Model . . . . .	149
4.3.3	Referential Integrity . . . . .	150
4.3.4	Degree Constraints . . . . .	151
4.3.5	Exercises for Section 4.3 . . . . .	151
4.4	Weak Entity Sets . . . . .	152
4.4.1	Causes of Weak Entity Sets . . . . .	152
4.4.2	Requirements for Weak Entity Sets . . . . .	153
4.4.3	Weak Entity Set Notation . . . . .	155
4.4.4	Exercises for Section 4.4 . . . . .	156
4.5	From E/R Diagrams to Relational Designs . . . . .	157
4.5.1	From Entity Sets to Relations . . . . .	157
4.5.2	From E/R Relationships to Relations . . . . .	158
4.5.3	Combining Relations . . . . .	160
4.5.4	Handling Weak Entity Sets . . . . .	161
4.5.5	Exercises for Section 4.5 . . . . .	163
4.6	Converting Subclass Structures to Relations . . . . .	165
4.6.1	E/R-Style Conversion . . . . .	166
4.6.2	An Object-Oriented Approach . . . . .	167
4.6.3	Using Null Values to Combine Relations . . . . .	168
4.6.4	Comparison of Approaches . . . . .	169
4.6.5	Exercises for Section 4.6 . . . . .	171
4.7	Unified Modeling Language . . . . .	171

4.7.1	UML Classes . . . . .	172
4.7.2	Keys for UML classes . . . . .	173
4.7.3	Associations . . . . .	173
4.7.4	Self-Associations . . . . .	175
4.7.5	Association Classes . . . . .	175
4.7.6	Subclasses in UML . . . . .	176
4.7.7	Aggregations and Compositions . . . . .	177
4.7.8	Exercises for Section 4.7 . . . . .	179
4.8	From UML Diagrams to Relations . . . . .	179
4.8.1	UML-to-Relations Basics . . . . .	179
4.8.2	From UML Subclasses to Relations . . . . .	180
4.8.3	From Aggregations and Compositions to Relations . . . . .	181
4.8.4	The UML Analog of Weak Entity Sets . . . . .	181
4.8.5	Exercises for Section 4.8 . . . . .	183
4.9	Object Definition Language . . . . .	183
4.9.1	Class Declarations . . . . .	184
4.9.2	Attributes in ODL . . . . .	184
4.9.3	Relationships in ODL . . . . .	185
4.9.4	Inverse Relationships . . . . .	186
4.9.5	Multiplicity of Relationships . . . . .	186
4.9.6	Types in ODL . . . . .	188
4.9.7	Subclasses in ODL . . . . .	190
4.9.8	Declaring Keys in ODL . . . . .	191
4.9.9	Exercises for Section 4.9 . . . . .	192
4.10	From ODL Designs to Relational Designs . . . . .	193
4.10.1	From ODL Classes to Relations . . . . .	193
4.10.2	Complex Attributes in Classes . . . . .	194
4.10.3	Representing Set-Valued Attributes . . . . .	195
4.10.4	Representing Other Type Constructors . . . . .	196
4.10.5	Representing ODL Relationships . . . . .	198
4.10.6	Exercises for Section 4.10 . . . . .	198
4.11	Summary of Chapter 4 . . . . .	200
4.12	References for Chapter 4 . . . . .	202

## II Relational Database Programming 203

<b>5</b>	<b>Algebraic and Logical Query Languages</b>	<b>205</b>
5.1	Relational Operations on Bags . . . . .	205
5.1.1	Why Bags? . . . . .	206
5.1.2	Union, Intersection, and Difference of Bags . . . . .	207
5.1.3	Projection of Bags . . . . .	208
5.1.4	Selection on Bags . . . . .	209
5.1.5	Product of Bags . . . . .	210
5.1.6	Joins of Bags . . . . .	210

5.1.7	Exercises for Section 5.1 . . . . .	212
5.2	Extended Operators of Relational Algebra . . . . .	213
5.2.1	Duplicate Elimination . . . . .	214
5.2.2	Aggregation Operators . . . . .	214
5.2.3	Grouping . . . . .	215
5.2.4	The Grouping Operator . . . . .	216
5.2.5	Extending the Projection Operator . . . . .	217
5.2.6	The Sorting Operator . . . . .	219
5.2.7	Outerjoins . . . . .	219
5.2.8	Exercises for Section 5.2 . . . . .	222
5.3	A Logic for Relations . . . . .	222
5.3.1	Predicates and Atoms . . . . .	223
5.3.2	Arithmetic Atoms . . . . .	223
5.3.3	Datalog Rules and Queries . . . . .	224
5.3.4	Meaning of Datalog Rules . . . . .	225
5.3.5	Extensional and Intensional Predicates . . . . .	228
5.3.6	Datalog Rules Applied to Bags . . . . .	228
5.3.7	Exercises for Section 5.3 . . . . .	230
5.4	Relational Algebra and Datalog . . . . .	230
5.4.1	Boolean Operations . . . . .	231
5.4.2	Projection . . . . .	232
5.4.3	Selection . . . . .	232
5.4.4	Product . . . . .	235
5.4.5	Joins . . . . .	235
5.4.6	Simulating Multiple Operations with Datalog . . . . .	236
5.4.7	Comparison Between Datalog and Relational Algebra . . . . .	238
5.4.8	Exercises for Section 5.4 . . . . .	238
5.5	Summary of Chapter 5 . . . . .	240
5.6	References for Chapter 5 . . . . .	241
<b>6</b>	<b>The Database Language SQL</b> . . . . .	<b>243</b>
6.1	Simple Queries in SQL . . . . .	244
6.1.1	Projection in SQL . . . . .	246
6.1.2	Selection in SQL . . . . .	248
6.1.3	Comparison of Strings . . . . .	250
6.1.4	Pattern Matching in SQL . . . . .	250
6.1.5	Dates and Times . . . . .	251
6.1.6	Null Values and Comparisons Involving NULL . . . . .	252
6.1.7	The Truth-Value UNKNOWN . . . . .	253
6.1.8	Ordering the Output . . . . .	255
6.1.9	Exercises for Section 6.1 . . . . .	256
6.2	Queries Involving More Than One Relation . . . . .	258
6.2.1	Products and Joins in SQL . . . . .	259
6.2.2	Disambiguating Attributes . . . . .	260
6.2.3	Tuple Variables . . . . .	261

6.2.4	Interpreting Multirelation Queries . . . . .	262
6.2.5	Union, Intersection, and Difference of Queries . . . . .	265
6.2.6	Exercises for Section 6.2 . . . . .	267
6.3	Subqueries . . . . .	268
6.3.1	Subqueries that Produce Scalar Values . . . . .	269
6.3.2	Conditions Involving Relations . . . . .	270
6.3.3	Conditions Involving Tuples . . . . .	271
6.3.4	Correlated Subqueries . . . . .	273
6.3.5	Subqueries in FROM Clauses . . . . .	274
6.3.6	SQL Join Expressions . . . . .	275
6.3.7	Natural Joins . . . . .	276
6.3.8	Outerjoins . . . . .	277
6.3.9	Exercises for Section 6.3 . . . . .	279
6.4	Full-Relation Operations . . . . .	281
6.4.1	Eliminating Duplicates . . . . .	281
6.4.2	Duplicates in Unions, Intersections, and Differences . . . . .	282
6.4.3	Grouping and Aggregation in SQL . . . . .	283
6.4.4	Aggregation Operators . . . . .	284
6.4.5	Grouping . . . . .	285
6.4.6	Grouping, Aggregation, and Nulls . . . . .	287
6.4.7	HAVING Clauses . . . . .	288
6.4.8	Exercises for Section 6.4 . . . . .	289
6.5	Database Modifications . . . . .	291
6.5.1	Insertion . . . . .	291
6.5.2	Deletion . . . . .	292
6.5.3	Updates . . . . .	294
6.5.4	Exercises for Section 6.5 . . . . .	295
6.6	Transactions in SQL . . . . .	296
6.6.1	Serializability . . . . .	296
6.6.2	Atomicity . . . . .	298
6.6.3	Transactions . . . . .	299
6.6.4	Read-Only Transactions . . . . .	300
6.6.5	Dirty Reads . . . . .	302
6.6.6	Other Isolation Levels . . . . .	304
6.6.7	Exercises for Section 6.6 . . . . .	306
6.7	Summary of Chapter 6 . . . . .	307
6.8	References for Chapter 6 . . . . .	308
<b>7</b>	<b>Constraints and Triggers</b> . . . . .	<b>311</b>
7.1	Keys and Foreign Keys . . . . .	311
7.1.1	Declaring Foreign-Key Constraints . . . . .	312
7.1.2	Maintaining Referential Integrity . . . . .	313
7.1.3	Deferred Checking of Constraints . . . . .	315
7.1.4	Exercises for Section 7.1 . . . . .	318
7.2	Constraints on Attributes and Tuples . . . . .	319

7.2.1	Not-Null Constraints . . . . .	319
7.2.2	Attribute-Based CHECK Constraints . . . . .	320
7.2.3	Tuple-Based CHECK Constraints . . . . .	321
7.2.4	Comparison of Tuple- and Attribute-Based Constraints . . . . .	323
7.2.5	Exercises for Section 7.2 . . . . .	323
7.3	Modification of Constraints . . . . .	325
7.3.1	Giving Names to Constraints . . . . .	325
7.3.2	Altering Constraints on Tables . . . . .	326
7.3.3	Exercises for Section 7.3 . . . . .	327
7.4	Assertions . . . . .	328
7.4.1	Creating Assertions . . . . .	328
7.4.2	Using Assertions . . . . .	329
7.4.3	Exercises for Section 7.4 . . . . .	330
7.5	Triggers . . . . .	332
7.5.1	Triggers in SQL . . . . .	332
7.5.2	The Options for Trigger Design . . . . .	334
7.5.3	Exercises for Section 7.5 . . . . .	337
7.6	Summary of Chapter 7 . . . . .	339
7.7	References for Chapter 7 . . . . .	339
<b>8</b>	<b>Views and Indexes</b>	<b>341</b>
8.1	Virtual Views . . . . .	341
8.1.1	Declaring Views . . . . .	341
8.1.2	Querying Views . . . . .	343
8.1.3	Renaming Attributes . . . . .	343
8.1.4	Exercises for Section 8.1 . . . . .	344
8.2	Modifying Views . . . . .	344
8.2.1	View Removal . . . . .	345
8.2.2	Updatable Views . . . . .	345
8.2.3	Instead-Of Triggers on Views . . . . .	347
8.2.4	Exercises for Section 8.2 . . . . .	349
8.3	Indexes in SQL . . . . .	350
8.3.1	Motivation for Indexes . . . . .	350
8.3.2	Declaring Indexes . . . . .	351
8.3.3	Exercises for Section 8.3 . . . . .	352
8.4	Selection of Indexes . . . . .	352
8.4.1	A Simple Cost Model . . . . .	352
8.4.2	Some Useful Indexes . . . . .	353
8.4.3	Calculating the Best Indexes to Create . . . . .	355
8.4.4	Automatic Selection of Indexes to Create . . . . .	357
8.4.5	Exercises for Section 8.4 . . . . .	359
8.5	Materialized Views . . . . .	359
8.5.1	Maintaining a Materialized View . . . . .	360
8.5.2	Periodic Maintenance of Materialized Views . . . . .	362
8.5.3	Rewriting Queries to Use Materialized Views . . . . .	362



8.5.4	Automatic Creation of Materialized Views . . . . .	364
8.5.5	Exercises for Section 8.5 . . . . .	365
8.6	Summary of Chapter 8 . . . . .	366
8.7	References for Chapter 8 . . . . .	367
<b>9</b>	<b>SQL in a Server Environment</b>	<b>369</b>
9.1	The Three-Tier Architecture . . . . .	369
9.1.1	The Web-Server Tier . . . . .	370
9.1.2	The Application Tier . . . . .	371
9.1.3	The Database Tier . . . . .	372
9.2	The SQL Environment . . . . .	372
9.2.1	Environments . . . . .	373
9.2.2	Schemas . . . . .	374
9.2.3	Catalogs . . . . .	375
9.2.4	Clients and Servers in the SQL Environment . . . . .	375
9.2.5	Connections . . . . .	376
9.2.6	Sessions . . . . .	377
9.2.7	Modules . . . . .	378
9.3	The SQL/Host-Language Interface . . . . .	378
9.3.1	The Impedance Mismatch Problem . . . . .	380
9.3.2	Connecting SQL to the Host Language . . . . .	380
9.3.3	The DECLARE Section . . . . .	381
9.3.4	Using Shared Variables . . . . .	382
9.3.5	Single-Row Select Statements . . . . .	383
9.3.6	Cursors . . . . .	383
9.3.7	Modifications by Cursor . . . . .	386
9.3.8	Protecting Against Concurrent Updates . . . . .	387
9.3.9	Dynamic SQL . . . . .	388
9.3.10	Exercises for Section 9.3 . . . . .	390
9.4	Stored Procedures . . . . .	391
9.4.1	Creating PSM Functions and Procedures . . . . .	391
9.4.2	Some Simple Statement Forms in PSM . . . . .	392
9.4.3	Branching Statements . . . . .	394
9.4.4	Queries in PSM . . . . .	395
9.4.5	Loops in PSM . . . . .	396
9.4.6	For-Loops . . . . .	398
9.4.7	Exceptions in PSM . . . . .	400
9.4.8	Using PSM Functions and Procedures . . . . .	402
9.4.9	Exercises for Section 9.4 . . . . .	402
9.5	Using a Call-Level Interface . . . . .	404
9.5.1	Introduction to SQL/CLI . . . . .	405
9.5.2	Processing Statements . . . . .	407
9.5.3	Fetching Data From a Query Result . . . . .	408
9.5.4	Passing Parameters to Queries . . . . .	410
9.5.5	Exercises for Section 9.5 . . . . .	412

9.6	JDBC . . . . .	412
9.6.1	Introduction to JDBC . . . . .	412
9.6.2	Creating Statements in JDBC . . . . .	413
9.6.3	Cursor Operations in JDBC . . . . .	415
9.6.4	Parameter Passing . . . . .	416
9.6.5	Exercises for Section 9.6 . . . . .	416
9.7	PHP . . . . .	416
9.7.1	PHP Basics . . . . .	417
9.7.2	Arrays . . . . .	418
9.7.3	The PEAR DB Library . . . . .	419
9.7.4	Creating a Database Connection Using DB . . . . .	419
9.7.5	Executing SQL Statements . . . . .	419
9.7.6	Cursor Operations in PHP . . . . .	420
9.7.7	Dynamic SQL in PHP . . . . .	421
9.7.8	Exercises for Section 9.7 . . . . .	422
9.8	Summary of Chapter 9 . . . . .	422
9.9	References for Chapter 9 . . . . .	423
<b>10</b>	<b>Advanced Topics in Relational Databases</b>	<b>425</b>
10.1	Security and User Authorization in SQL . . . . .	425
10.1.1	Privileges . . . . .	426
10.1.2	Creating Privileges . . . . .	427
10.1.3	The Privilege-Checking Process . . . . .	428
10.1.4	Granting Privileges . . . . .	430
10.1.5	Grant Diagrams . . . . .	431
10.1.6	Revoking Privileges . . . . .	433
10.1.7	Exercises for Section 10.1 . . . . .	436
10.2	Recursion in SQL . . . . .	437
10.2.1	Defining Recursive Relations in SQL . . . . .	437
10.2.2	Problematic Expressions in Recursive SQL . . . . .	440
10.2.3	Exercises for Section 10.2 . . . . .	443
10.3	The Object-Relational Model . . . . .	445
10.3.1	From Relations to Object-Relations . . . . .	445
10.3.2	Nested Relations . . . . .	446
10.3.3	References . . . . .	447
10.3.4	Object-Oriented Versus Object-Relational . . . . .	449
10.3.5	Exercises for Section 10.3 . . . . .	450
10.4	User-Defined Types in SQL . . . . .	451
10.4.1	Defining Types in SQL . . . . .	451
10.4.2	Method Declarations in UDT's . . . . .	452
10.4.3	Method Definitions . . . . .	453
10.4.4	Declaring Relations with a UDT . . . . .	454
10.4.5	References . . . . .	454
10.4.6	Creating Object ID's for Tables . . . . .	455
10.4.7	Exercises for Section 10.4 . . . . .	457

10.5	Operations on Object-Relational Data . . . . .	457
10.5.1	Following References . . . . .	457
10.5.2	Accessing Components of Tuples with a UDT . . . . .	458
10.5.3	Generator and Mutator Functions . . . . .	460
10.5.4	Ordering Relationships on UDT's . . . . .	461
10.5.5	Exercises for Section 10.5 . . . . .	463
10.6	On-Line Analytic Processing . . . . .	464
10.6.1	OLAP and Data Warehouses . . . . .	465
10.6.2	OLAP Applications . . . . .	465
10.6.3	A Multidimensional View of OLAP Data . . . . .	466
10.6.4	Star Schemas . . . . .	467
10.6.5	Slicing and Dicing . . . . .	469
10.6.6	Exercises for Section 10.6 . . . . .	472
10.7	Data Cubes . . . . .	473
10.7.1	The Cube Operator . . . . .	473
10.7.2	The Cube Operator in SQL . . . . .	475
10.7.3	Exercises for Section 10.7 . . . . .	477
10.8	Summary of Chapter 10 . . . . .	478
10.9	References for Chapter 10 . . . . .	480

### **III Modeling and Programming for Semistructured Data 481**

<b>11</b>	<b>The Semistructured-Data Model 483</b>
11.1	Semistructured Data . . . . . 483
11.1.1	Motivation for the Semistructured-Data Model . . . . . 483
11.1.2	Semistructured Data Representation . . . . . 484
11.1.3	Information Integration Via Semistructured Data . . . . . 486
11.1.4	Exercises for Section 11.1 . . . . . 487
11.2	XML . . . . . 488
11.2.1	Semantic Tags . . . . . 488
11.2.2	XML With and Without a Schema . . . . . 489
11.2.3	Well-Formed XML . . . . . 489
11.2.4	Attributes . . . . . 490
11.2.5	Attributes That Connect Elements . . . . . 491
11.2.6	Namespaces . . . . . 493
11.2.7	XML and Databases . . . . . 493
11.2.8	Exercises for Section 11.2 . . . . . 495
11.3	Document Type Definitions . . . . . 495
11.3.1	The Form of a DTD . . . . . 495
11.3.2	Using a DTD . . . . . 499
11.3.3	Attribute Lists . . . . . 499
11.3.4	Identifiers and References . . . . . 500
11.3.5	Exercises for Section 11.3 . . . . . 502

11.4	XML Schema . . . . .	502
11.4.1	The Form of an XML Schema . . . . .	502
11.4.2	Elements . . . . .	503
11.4.3	Complex Types . . . . .	504
11.4.4	Attributes . . . . .	506
11.4.5	Restricted Simple Types . . . . .	507
11.4.6	Keys in XML Schema . . . . .	509
11.4.7	Foreign Keys in XML Schema . . . . .	510
11.4.8	Exercises for Section 11.4 . . . . .	512
11.5	Summary of Chapter 11 . . . . .	514
11.6	References for Chapter 11 . . . . .	515
<b>12</b>	<b>Programming Languages for XML</b>	<b>517</b>
12.1	XPath . . . . .	517
12.1.1	The XPath Data Model . . . . .	518
12.1.2	Document Nodes . . . . .	519
12.1.3	Path Expressions . . . . .	519
12.1.4	Relative Path Expressions . . . . .	521
12.1.5	Attributes in Path Expressions . . . . .	521
12.1.6	Axes . . . . .	521
12.1.7	Context of Expressions . . . . .	522
12.1.8	Wildcards . . . . .	523
12.1.9	Conditions in Path Expressions . . . . .	523
12.1.10	Exercises for Section 12.1 . . . . .	526
12.2	XQuery . . . . .	528
12.2.1	XQuery Basics . . . . .	530
12.2.2	FLWR Expressions . . . . .	530
12.2.3	Replacement of Variables by Their Values . . . . .	534
12.2.4	Joins in XQuery . . . . .	536
12.2.5	XQuery Comparison Operators . . . . .	537
12.2.6	Elimination of Duplicates . . . . .	538
12.2.7	Quantification in XQuery . . . . .	539
12.2.8	Aggregations . . . . .	540
12.2.9	Branching in XQuery Expressions . . . . .	540
12.2.10	Ordering the Result of a Query . . . . .	541
12.2.11	Exercises for Section 12.2 . . . . .	543
12.3	Extensible Stylesheet Language . . . . .	544
12.3.1	XSLT Basics . . . . .	544
12.3.2	Templates . . . . .	544
12.3.3	Obtaining Values From XML Data . . . . .	545
12.3.4	Recursive Use of Templates . . . . .	546
12.3.5	Iteration in XSLT . . . . .	549
12.3.6	Conditionals in XSLT . . . . .	551
12.3.7	Exercises for Section 12.3 . . . . .	551
12.4	Summary of Chapter 12 . . . . .	553

12.5	References for Chapter 12 . . . . .	554
<b>IV</b>	<b>Database System Implementation</b>	<b>555</b>
<b>13</b>	<b>Secondary Storage Management</b>	<b>557</b>
13.1	The Memory Hierarchy . . . . .	557
13.1.1	The Memory Hierarchy . . . . .	557
13.1.2	Transfer of Data Between Levels . . . . .	560
13.1.3	Volatile and Nonvolatile Storage . . . . .	560
13.1.4	Virtual Memory . . . . .	560
13.1.5	Exercises for Section 13.1 . . . . .	561
13.2	Disks . . . . .	562
13.2.1	Mechanics of Disks . . . . .	562
13.2.2	The Disk Controller . . . . .	564
13.2.3	Disk Access Characteristics . . . . .	564
13.2.4	Exercises for Section 13.2 . . . . .	567
13.3	Accelerating Access to Secondary Storage . . . . .	568
13.3.1	The I/O Model of Computation . . . . .	568
13.3.2	Organizing Data by Cylinders . . . . .	569
13.3.3	Using Multiple Disks . . . . .	570
13.3.4	Mirroring Disks . . . . .	571
13.3.5	Disk Scheduling and the Elevator Algorithm . . . . .	571
13.3.6	Prefetching and Large-Scale Buffering . . . . .	573
13.3.7	Exercises for Section 13.3 . . . . .	573
13.4	Disk Failures . . . . .	575
13.4.1	Intermittent Failures . . . . .	576
13.4.2	Checksums . . . . .	576
13.4.3	Stable Storage . . . . .	577
13.4.4	Error-Handling Capabilities of Stable Storage . . . . .	578
13.4.5	Recovery from Disk Crashes . . . . .	578
13.4.6	Mirroring as a Redundancy Technique . . . . .	579
13.4.7	Parity Blocks . . . . .	580
13.4.8	An Improvement: RAID 5 . . . . .	583
13.4.9	Coping With Multiple Disk Crashes . . . . .	584
13.4.10	Exercises for Section 13.4 . . . . .	587
13.5	Arranging Data on Disk . . . . .	590
13.5.1	Fixed-Length Records . . . . .	590
13.5.2	Packing Fixed-Length Records into Blocks . . . . .	592
13.5.3	Exercises for Section 13.5 . . . . .	593
13.6	Representing Block and Record Addresses . . . . .	593
13.6.1	Addresses in Client-Server Systems . . . . .	593
13.6.2	Logical and Structured Addresses . . . . .	595
13.6.3	Pointer Swizzling . . . . .	596
13.6.4	Returning Blocks to Disk . . . . .	600

13.6.5	Pinned Records and Blocks . . . . .	600
13.6.6	Exercises for Section 13.6 . . . . .	602
13.7	Variable-Length Data and Records . . . . .	603
13.7.1	Records With Variable-Length Fields . . . . .	604
13.7.2	Records With Repeating Fields . . . . .	605
13.7.3	Variable-Format Records . . . . .	607
13.7.4	Records That Do Not Fit in a Block . . . . .	608
13.7.5	BLOBs . . . . .	608
13.7.6	Column Stores . . . . .	609
13.7.7	Exercises for Section 13.7 . . . . .	610
13.8	Record Modifications . . . . .	612
13.8.1	Insertion . . . . .	612
13.8.2	Deletion . . . . .	614
13.8.3	Update . . . . .	615
13.8.4	Exercises for Section 13.8 . . . . .	615
13.9	Summary of Chapter 13 . . . . .	615
13.10	References for Chapter 13 . . . . .	617
<b>14</b>	<b>Index Structures</b> . . . . .	<b>619</b>
14.1	Index-Structure Basics . . . . .	620
14.1.1	Sequential Files . . . . .	621
14.1.2	Dense Indexes . . . . .	621
14.1.3	Sparse Indexes . . . . .	622
14.1.4	Multiple Levels of Index . . . . .	623
14.1.5	Secondary Indexes . . . . .	624
14.1.6	Applications of Secondary Indexes . . . . .	625
14.1.7	Indirection in Secondary Indexes . . . . .	626
14.1.8	Document Retrieval and Inverted Indexes . . . . .	628
14.1.9	Exercises for Section 14.1 . . . . .	631
14.2	B-Trees . . . . .	633
14.2.1	The Structure of B-trees . . . . .	634
14.2.2	Applications of B-trees . . . . .	637
14.2.3	Lookup in B-Trees . . . . .	639
14.2.4	Range Queries . . . . .	639
14.2.5	Insertion Into B-Trees . . . . .	640
14.2.6	Deletion From B-Trees . . . . .	642
14.2.7	Efficiency of B-Trees . . . . .	645
14.2.8	Exercises for Section 14.2 . . . . .	646
14.3	Hash Tables . . . . .	648
14.3.1	Secondary-Storage Hash Tables . . . . .	649
14.3.2	Insertion Into a Hash Table . . . . .	649
14.3.3	Hash-Table Deletion . . . . .	650
14.3.4	Efficiency of Hash Table Indexes . . . . .	651
14.3.5	Extensible Hash Tables . . . . .	652
14.3.6	Insertion Into Extensible Hash Tables . . . . .	653

14.3.7	Linear Hash Tables . . . . .	655
14.3.8	Insertion Into Linear Hash Tables . . . . .	657
14.3.9	Exercises for Section 14.3 . . . . .	659
14.4	Multidimensional Indexes . . . . .	661
14.4.1	Applications of Multidimensional Indexes . . . . .	661
14.4.2	Executing Range Queries Using Conventional Indexes . . . . .	663
14.4.3	Executing Nearest-Neighbor Queries Using Conventional Indexes . . . . .	664
14.4.4	Overview of Multidimensional Index Structures . . . . .	664
14.5	Hash Structures for Multidimensional Data . . . . .	665
14.5.1	Grid Files . . . . .	665
14.5.2	Lookup in a Grid File . . . . .	666
14.5.3	Insertion Into Grid Files . . . . .	667
14.5.4	Performance of Grid Files . . . . .	669
14.5.5	Partitioned Hash Functions . . . . .	671
14.5.6	Comparison of Grid Files and Partitioned Hashing . . . . .	673
14.5.7	Exercises for Section 14.5 . . . . .	673
14.6	Tree Structures for Multidimensional Data . . . . .	675
14.6.1	Multiple-Key Indexes . . . . .	675
14.6.2	Performance of Multiple-Key Indexes . . . . .	676
14.6.3	<i>kd</i> -Trees . . . . .	677
14.6.4	Operations on <i>kd</i> -Trees . . . . .	679
14.6.5	Adapting <i>kd</i> -Trees to Secondary Storage . . . . .	681
14.6.6	Quad Trees . . . . .	681
14.6.7	R-Trees . . . . .	683
14.6.8	Operations on R-Trees . . . . .	684
14.6.9	Exercises for Section 14.6 . . . . .	686
14.7	Bitmap Indexes . . . . .	688
14.7.1	Motivation for Bitmap Indexes . . . . .	689
14.7.2	Compressed Bitmaps . . . . .	691
14.7.3	Operating on Run-Length-Encoded Bit-Vectors . . . . .	693
14.7.4	Managing Bitmap Indexes . . . . .	693
14.7.5	Exercises for Section 14.7 . . . . .	695
14.8	Summary of Chapter 14 . . . . .	695
14.9	References for Chapter 14 . . . . .	697
<b>15</b>	<b>Query Execution</b>	<b>701</b>
15.1	Introduction to Physical-Query-Plan Operators . . . . .	703
15.1.1	Scanning Tables . . . . .	703
15.1.2	Sorting While Scanning Tables . . . . .	704
15.1.3	The Computation Model for Physical Operators . . . . .	704
15.1.4	Parameters for Measuring Costs . . . . .	705
15.1.5	I/O Cost for Scan Operators . . . . .	706
15.1.6	Iterators for Implementation of Physical Operators . . . . .	707
15.2	One-Pass Algorithms . . . . .	709

15.2.1	One-Pass Algorithms for Tuple-at-a-Time Operations . . .	711
15.2.2	One-Pass Algorithms for Unary, Full-Relation Operations	712
15.2.3	One-Pass Algorithms for Binary Operations . . . . .	715
15.2.4	Exercises for Section 15.2 . . . . .	718
15.3	Nested-Loop Joins . . . . .	718
15.3.1	Tuple-Based Nested-Loop Join . . . . .	719
15.3.2	An Iterator for Tuple-Based Nested-Loop Join . . . . .	719
15.3.3	Block-Based Nested-Loop Join Algorithm . . . . .	719
15.3.4	Analysis of Nested-Loop Join . . . . .	721
15.3.5	Summary of Algorithms so Far . . . . .	722
15.3.6	Exercises for Section 15.3 . . . . .	722
15.4	Two-Pass Algorithms Based on Sorting . . . . .	723
15.4.1	Two-Phase, Multiway Merge-Sort . . . . .	723
15.4.2	Duplicate Elimination Using Sorting . . . . .	725
15.4.3	Grouping and Aggregation Using Sorting . . . . .	726
15.4.4	A Sort-Based Union Algorithm . . . . .	726
15.4.5	Sort-Based Intersection and Difference . . . . .	727
15.4.6	A Simple Sort-Based Join Algorithm . . . . .	728
15.4.7	Analysis of Simple Sort-Join . . . . .	729
15.4.8	A More Efficient Sort-Based Join . . . . .	729
15.4.9	Summary of Sort-Based Algorithms . . . . .	730
15.4.10	Exercises for Section 15.4 . . . . .	730
15.5	Two-Pass Algorithms Based on Hashing . . . . .	732
15.5.1	Partitioning Relations by Hashing . . . . .	732
15.5.2	A Hash-Based Algorithm for Duplicate Elimination . . .	732
15.5.3	Hash-Based Grouping and Aggregation . . . . .	733
15.5.4	Hash-Based Union, Intersection, and Difference . . . . .	734
15.5.5	The Hash-Join Algorithm . . . . .	734
15.5.6	Saving Some Disk I/O's . . . . .	735
15.5.7	Summary of Hash-Based Algorithms . . . . .	737
15.5.8	Exercises for Section 15.5 . . . . .	738
15.6	Index-Based Algorithms . . . . .	739
15.6.1	Clustering and Nonclustering Indexes . . . . .	739
15.6.2	Index-Based Selection . . . . .	740
15.6.3	Joining by Using an Index . . . . .	742
15.6.4	Joins Using a Sorted Index . . . . .	743
15.6.5	Exercises for Section 15.6 . . . . .	745
15.7	Buffer Management . . . . .	746
15.7.1	Buffer Management Architecture . . . . .	746
15.7.2	Buffer Management Strategies . . . . .	747
15.7.3	The Relationship Between Physical Operator Selection and Buffer Management . . . . .	750
15.7.4	Exercises for Section 15.7 . . . . .	751
15.8	Algorithms Using More Than Two Passes . . . . .	752
15.8.1	Multipass Sort-Based Algorithms . . . . .	752



15.8.2	Performance of Multipass, Sort-Based Algorithms . . . .	753
15.8.3	Multipass Hash-Based Algorithms . . . . .	754
15.8.4	Performance of Multipass Hash-Based Algorithms . . . .	754
15.8.5	Exercises for Section 15.8 . . . . .	755
15.9	Summary of Chapter 15 . . . . .	756
15.10	References for Chapter 15 . . . . .	757
<b>16</b>	<b>The Query Compiler</b>	<b>759</b>
16.1	Parsing and Preprocessing . . . . .	760
16.1.1	Syntax Analysis and Parse Trees . . . . .	760
16.1.2	A Grammar for a Simple Subset of SQL . . . . .	761
16.1.3	The Preprocessor . . . . .	764
16.1.4	Preprocessing Queries Involving Views . . . . .	765
16.1.5	Exercises for Section 16.1 . . . . .	767
16.2	Algebraic Laws for Improving Query Plans . . . . .	768
16.2.1	Commutative and Associative Laws . . . . .	768
16.2.2	Laws Involving Selection . . . . .	770
16.2.3	Pushing Selections . . . . .	772
16.2.4	Laws Involving Projection . . . . .	774
16.2.5	Laws About Joins and Products . . . . .	776
16.2.6	Laws Involving Duplicate Elimination . . . . .	777
16.2.7	Laws Involving Grouping and Aggregation . . . . .	777
16.2.8	Exercises for Section 16.2 . . . . .	780
16.3	From Parse Trees to Logical Query Plans . . . . .	781
16.3.1	Conversion to Relational Algebra . . . . .	782
16.3.2	Removing Subqueries From Conditions . . . . .	783
16.3.3	Improving the Logical Query Plan . . . . .	788
16.3.4	Grouping Associative/Commutative Operators . . . . .	790
16.3.5	Exercises for Section 16.3 . . . . .	791
16.4	Estimating the Cost of Operations . . . . .	792
16.4.1	Estimating Sizes of Intermediate Relations . . . . .	793
16.4.2	Estimating the Size of a Projection . . . . .	794
16.4.3	Estimating the Size of a Selection . . . . .	794
16.4.4	Estimating the Size of a Join . . . . .	797
16.4.5	Natural Joins With Multiple Join Attributes . . . . .	799
16.4.6	Joins of Many Relations . . . . .	800
16.4.7	Estimating Sizes for Other Operations . . . . .	801
16.4.8	Exercises for Section 16.4 . . . . .	802
16.5	Introduction to Cost-Based Plan Selection . . . . .	803
16.5.1	Obtaining Estimates for Size Parameters . . . . .	804
16.5.2	Computation of Statistics . . . . .	807
16.5.3	Heuristics for Reducing the Cost of Logical Query Plans .	808
16.5.4	Approaches to Enumerating Physical Plans . . . . .	810
16.5.5	Exercises for Section 16.5 . . . . .	813
16.6	Choosing an Order for Joins . . . . .	814

16.6.1	Significance of Left and Right Join Arguments . . . . .	815
16.6.2	Join Trees . . . . .	815
16.6.3	Left-Deep Join Trees . . . . .	816
16.6.4	Dynamic Programming to Select a Join Order and Grouping	819
16.6.5	Dynamic Programming With More Detailed Cost Functions	823
16.6.6	A Greedy Algorithm for Selecting a Join Order . . . . .	824
16.6.7	Exercises for Section 16.6 . . . . .	825
16.7	Completing the Physical-Query-Plan . . . . .	826
16.7.1	Choosing a Selection Method . . . . .	827
16.7.2	Choosing a Join Method . . . . .	829
16.7.3	Pipelining Versus Materialization . . . . .	830
16.7.4	Pipelining Unary Operations . . . . .	830
16.7.5	Pipelining Binary Operations . . . . .	830
16.7.6	Notation for Physical Query Plans . . . . .	834
16.7.7	Ordering of Physical Operations . . . . .	837
16.7.8	Exercises for Section 16.7 . . . . .	838
16.8	Summary of Chapter 16 . . . . .	839
16.9	References for Chapter 16 . . . . .	841
<b>17</b>	<b>Coping With System Failures</b>	<b>843</b>
17.1	Issues and Models for Resilient Operation . . . . .	843
17.1.1	Failure Modes . . . . .	844
17.1.2	More About Transactions . . . . .	845
17.1.3	Correct Execution of Transactions . . . . .	846
17.1.4	The Primitive Operations of Transactions . . . . .	848
17.1.5	Exercises for Section 17.1 . . . . .	851
17.2	Undo Logging . . . . .	851
17.2.1	Log Records . . . . .	851
17.2.2	The Undo-Logging Rules . . . . .	853
17.2.3	Recovery Using Undo Logging . . . . .	855
17.2.4	Checkpointing . . . . .	857
17.2.5	Nonquiescent Checkpointing . . . . .	858
17.2.6	Exercises for Section 17.2 . . . . .	862
17.3	Redo Logging . . . . .	863
17.3.1	The Redo-Logging Rule . . . . .	863
17.3.2	Recovery With Redo Logging . . . . .	864
17.3.3	Checkpointing a Redo Log . . . . .	866
17.3.4	Recovery With a Checkpointed Redo Log . . . . .	867
17.3.5	Exercises for Section 17.3 . . . . .	868
17.4	Undo/Redo Logging . . . . .	869
17.4.1	The Undo/Redo Rules . . . . .	870
17.4.2	Recovery With Undo/Redo Logging . . . . .	870
17.4.3	Checkpointing an Undo/Redo Log . . . . .	872
17.4.4	Exercises for Section 17.4 . . . . .	874
17.5	Protecting Against Media Failures . . . . .	875

17.5.1	The Archive . . . . .	875
17.5.2	Nonquiescent Archiving . . . . .	875
17.5.3	Recovery Using an Archive and Log . . . . .	878
17.5.4	Exercises for Section 17.5 . . . . .	879
17.6	Summary of Chapter 17 . . . . .	879
17.7	References for Chapter 17 . . . . .	881
<b>18</b>	<b>Concurrency Control</b>	<b>883</b>
18.1	Serial and Serializable Schedules . . . . .	884
18.1.1	Schedules . . . . .	884
18.1.2	Serial Schedules . . . . .	885
18.1.3	Serializable Schedules . . . . .	886
18.1.4	The Effect of Transaction Semantics . . . . .	887
18.1.5	A Notation for Transactions and Schedules . . . . .	889
18.1.6	Exercises for Section 18.1 . . . . .	889
18.2	Conflict-Serializability . . . . .	890
18.2.1	Conflicts . . . . .	890
18.2.2	Precedence Graphs and a Test for Conflict-Serializability . . . . .	892
18.2.3	Why the Precedence-Graph Test Works . . . . .	894
18.2.4	Exercises for Section 18.2 . . . . .	895
18.3	Enforcing Serializability by Locks . . . . .	897
18.3.1	Locks . . . . .	898
18.3.2	The Locking Scheduler . . . . .	900
18.3.3	Two-Phase Locking . . . . .	900
18.3.4	Why Two-Phase Locking Works . . . . .	901
18.3.5	Exercises for Section 18.3 . . . . .	903
18.4	Locking Systems With Several Lock Modes . . . . .	905
18.4.1	Shared and Exclusive Locks . . . . .	905
18.4.2	Compatibility Matrices . . . . .	907
18.4.3	Upgrading Locks . . . . .	908
18.4.4	Update Locks . . . . .	909
18.4.5	Increment Locks . . . . .	911
18.4.6	Exercises for Section 18.4 . . . . .	913
18.5	An Architecture for a Locking Scheduler . . . . .	915
18.5.1	A Scheduler That Inserts Lock Actions . . . . .	915
18.5.2	The Lock Table . . . . .	918
18.5.3	Exercises for Section 18.5 . . . . .	921
18.6	Hierarchies of Database Elements . . . . .	921
18.6.1	Locks With Multiple Granularity . . . . .	921
18.6.2	Warning Locks . . . . .	922
18.6.3	Phantoms and Handling Insertions Correctly . . . . .	926
18.6.4	Exercises for Section 18.6 . . . . .	927
18.7	The Tree Protocol . . . . .	927
18.7.1	Motivation for Tree-Based Locking . . . . .	927
18.7.2	Rules for Access to Tree-Structured Data . . . . .	928

18.7.3	Why the Tree Protocol Works . . . . .	929
18.7.4	Exercises for Section 18.7 . . . . .	932
18.8	Concurrency Control by Timestamps . . . . .	933
18.8.1	Timestamps . . . . .	934
18.8.2	Physically Unrealizable Behaviors . . . . .	934
18.8.3	Problems With Dirty Data . . . . .	935
18.8.4	The Rules for Timestamp-Based Scheduling . . . . .	937
18.8.5	Multiversion Timestamps . . . . .	939
18.8.6	Timestamps Versus Locking . . . . .	941
18.8.7	Exercises for Section 18.8 . . . . .	942
18.9	Concurrency Control by Validation . . . . .	942
18.9.1	Architecture of a Validation-Based Scheduler . . . . .	942
18.9.2	The Validation Rules . . . . .	943
18.9.3	Comparison of Three Concurrency-Control Mechanisms . . . . .	946
18.9.4	Exercises for Section 18.9 . . . . .	948
18.10	Summary of Chapter 18 . . . . .	948
18.11	References for Chapter 18 . . . . .	950
<b>19</b>	<b>More About Transaction Management</b>	<b>953</b>
19.1	Serializability and Recoverability . . . . .	953
19.1.1	The Dirty-Data Problem . . . . .	954
19.1.2	Cascading Rollback . . . . .	955
19.1.3	Recoverable Schedules . . . . .	956
19.1.4	Schedules That Avoid Cascading Rollback . . . . .	957
19.1.5	Managing Rollbacks Using Locking . . . . .	957
19.1.6	Group Commit . . . . .	959
19.1.7	Logical Logging . . . . .	960
19.1.8	Recovery From Logical Logs . . . . .	963
19.1.9	Exercises for Section 19.1 . . . . .	965
19.2	Deadlocks . . . . .	966
19.2.1	Deadlock Detection by Timeout . . . . .	967
19.2.2	The Waits-For Graph . . . . .	967
19.2.3	Deadlock Prevention by Ordering Elements . . . . .	970
19.2.4	Detecting Deadlocks by Timestamps . . . . .	970
19.2.5	Comparison of Deadlock-Management Methods . . . . .	972
19.2.6	Exercises for Section 19.2 . . . . .	974
19.3	Long-Duration Transactions . . . . .	975
19.3.1	Problems of Long Transactions . . . . .	976
19.3.2	Sagas . . . . .	978
19.3.3	Compensating Transactions . . . . .	979
19.3.4	Why Compensating Transactions Work . . . . .	980
19.3.5	Exercises for Section 19.3 . . . . .	981
19.4	Summary of Chapter 19 . . . . .	982
19.5	References for Chapter 19 . . . . .	983

<b>20 Parallel and Distributed Databases</b>	<b>985</b>
20.1 Parallel Algorithms on Relations . . . . .	985
20.1.1 Models of Parallelism . . . . .	986
20.1.2 Tuple-at-a-Time Operations in Parallel . . . . .	989
20.1.3 Parallel Algorithms for Full-Relation Operations . . . . .	989
20.1.4 Performance of Parallel Algorithms . . . . .	990
20.1.5 Exercises for Section 20.1 . . . . .	993
20.2 The Map-Reduce Parallelism Framework . . . . .	993
20.2.1 The Storage Model . . . . .	993
20.2.2 The Map Function . . . . .	994
20.2.3 The Reduce Function . . . . .	995
20.2.4 Exercises for Section 20.2 . . . . .	996
20.3 Distributed Databases . . . . .	997
20.3.1 Distribution of Data . . . . .	997
20.3.2 Distributed Transactions . . . . .	998
20.3.3 Data Replication . . . . .	999
20.3.4 Exercises for Section 20.3 . . . . .	1000
20.4 Distributed Query Processing . . . . .	1000
20.4.1 The Distributed Join Problem . . . . .	1000
20.4.2 Semijoin Reductions . . . . .	1001
20.4.3 Joins of Many Relations . . . . .	1002
20.4.4 Acyclic Hypergraphs . . . . .	1003
20.4.5 Full Reducers for Acyclic Hypergraphs . . . . .	1005
20.4.6 Why the Full-Reducer Algorithm Works . . . . .	1006
20.4.7 Exercises for Section 20.4 . . . . .	1007
20.5 Distributed Commit . . . . .	1008
20.5.1 Supporting Distributed Atomicity . . . . .	1008
20.5.2 Two-Phase Commit . . . . .	1009
20.5.3 Recovery of Distributed Transactions . . . . .	1011
20.5.4 Exercises for Section 20.5 . . . . .	1013
20.6 Distributed Locking . . . . .	1014
20.6.1 Centralized Lock Systems . . . . .	1015
20.6.2 A Cost Model for Distributed Locking Algorithms . . . . .	1015
20.6.3 Locking Replicated Elements . . . . .	1016
20.6.4 Primary-Copy Locking . . . . .	1017
20.6.5 Global Locks From Local Locks . . . . .	1017
20.6.6 Exercises for Section 20.6 . . . . .	1019
20.7 Peer-to-Peer Distributed Search . . . . .	1020
20.7.1 Peer-to-Peer Networks . . . . .	1020
20.7.2 The Distributed-Hashing Problem . . . . .	1021
20.7.3 Centralized Solutions for Distributed Hashing . . . . .	1022
20.7.4 Chord Circles . . . . .	1022
20.7.5 Links in Chord Circles . . . . .	1024
20.7.6 Search Using Finger Tables . . . . .	1024
20.7.7 Adding New Nodes . . . . .	1027

20.7.8	When a Peer Leaves the Network . . . . .	1030
20.7.9	When a Peer Fails . . . . .	1030
20.7.10	Exercises for Section 20.7 . . . . .	1031
20.8	Summary of Chapter 20 . . . . .	1031
20.9	References for Chapter 20 . . . . .	1033

## V Other Issues in Management of Massive Data 1035

<b>21</b>	<b>Information Integration</b>	<b>1037</b>
21.1	Introduction to Information Integration . . . . .	1037
21.1.1	Why Information Integration? . . . . .	1038
21.1.2	The Heterogeneity Problem . . . . .	1040
21.2	Modes of Information Integration . . . . .	1041
21.2.1	Federated Database Systems . . . . .	1042
21.2.2	Data Warehouses . . . . .	1043
21.2.3	Mediators . . . . .	1046
21.2.4	Exercises for Section 21.2 . . . . .	1048
21.3	Wrappers in Mediator-Based Systems . . . . .	1049
21.3.1	Templates for Query Patterns . . . . .	1050
21.3.2	Wrapper Generators . . . . .	1051
21.3.3	Filters . . . . .	1052
21.3.4	Other Operations at the Wrapper . . . . .	1053
21.3.5	Exercises for Section 21.3 . . . . .	1054
21.4	Capability-Based Optimization . . . . .	1056
21.4.1	The Problem of Limited Source Capabilities . . . . .	1056
21.4.2	A Notation for Describing Source Capabilities . . . . .	1057
21.4.3	Capability-Based Query-Plan Selection . . . . .	1058
21.4.4	Adding Cost-Based Optimization . . . . .	1060
21.4.5	Exercises for Section 21.4 . . . . .	1060
21.5	Optimizing Mediator Queries . . . . .	1061
21.5.1	Simplified Adornment Notation . . . . .	1061
21.5.2	Obtaining Answers for Subgoals . . . . .	1062
21.5.3	The Chain Algorithm . . . . .	1063
21.5.4	Incorporating Union Views at the Mediator . . . . .	1067
21.5.5	Exercises for Section 21.5 . . . . .	1068
21.6	Local-as-View Mediators . . . . .	1069
21.6.1	Motivation for LAV Mediators . . . . .	1069
21.6.2	Terminology for LAV Mediation . . . . .	1070
21.6.3	Expanding Solutions . . . . .	1071
21.6.4	Containment of Conjunctive Queries . . . . .	1073
21.6.5	Why the Containment-Mapping Test Works . . . . .	1075
21.6.6	Finding Solutions to a Mediator Query . . . . .	1076
21.6.7	Why the LMSS Theorem Holds . . . . .	1077
21.6.8	Exercises for Section 21.6 . . . . .	1078

21.7	Entity Resolution . . . . .	1078
21.7.1	Deciding Whether Records Represent a Common Entity . . . . .	1079
21.7.2	Merging Similar Records . . . . .	1081
21.7.3	Useful Properties of Similarity and Merge Functions . . . . .	1082
21.7.4	The R-Swoosh Algorithm for ICAR Records . . . . .	1083
21.7.5	Why R-Swoosh Works . . . . .	1086
21.7.6	Other Approaches to Entity Resolution . . . . .	1086
21.7.7	Exercises for Section 21.7 . . . . .	1087
21.8	Summary of Chapter 21 . . . . .	1089
21.9	References for Chapter 21 . . . . .	1091
<b>22</b>	<b>Data Mining</b>	<b>1093</b>
22.1	Frequent-Itemset Mining . . . . .	1093
22.1.1	The Market-Basket Model . . . . .	1094
22.1.2	Basic Definitions . . . . .	1095
22.1.3	Association Rules . . . . .	1097
22.1.4	The Computation Model for Frequent Itemsets . . . . .	1098
22.1.5	Exercises for Section 22.1 . . . . .	1099
22.2	Algorithms for Finding Frequent Itemsets . . . . .	1100
22.2.1	The Distribution of Frequent Itemsets . . . . .	1100
22.2.2	The Naive Algorithm for Finding Frequent Itemsets . . . . .	1101
22.2.3	The A-Priori Algorithm . . . . .	1102
22.2.4	Implementation of the A-Priori Algorithm . . . . .	1104
22.2.5	Making Better Use of Main Memory . . . . .	1105
22.2.6	When to Use the PCY Algorithm . . . . .	1106
22.2.7	The Multistage Algorithm . . . . .	1107
22.2.8	Exercises for Section 22.2 . . . . .	1109
22.3	Finding Similar Items . . . . .	1110
22.3.1	The Jaccard Measure of Similarity . . . . .	1110
22.3.2	Applications of Jaccard Similarity . . . . .	1110
22.3.3	Minhashing . . . . .	1112
22.3.4	Minhashing and Jaccard Distance . . . . .	1113
22.3.5	Why Minhashing Works . . . . .	1113
22.3.6	Implementing Minhashing . . . . .	1114
22.3.7	Exercises for Section 22.3 . . . . .	1115
22.4	Locality-Sensitive Hashing . . . . .	1116
22.4.1	Entity Resolution as an Example of LSH . . . . .	1117
22.4.2	Locality-Sensitive Hashing of Signatures . . . . .	1118
22.4.3	Combining Minhashing and Locality-Sensitive Hashing . . . . .	1121
22.4.4	Exercises for Section 22.4 . . . . .	1122
22.5	Clustering of Large-Scale Data . . . . .	1123
22.5.1	Applications of Clustering . . . . .	1123
22.5.2	Distance Measures . . . . .	1125
22.5.3	Agglomerative Clustering . . . . .	1128
22.5.4	$k$ -Means Algorithms . . . . .	1130

22.5.5	$k$ -Means for Large-Scale Data . . . . .	1132
22.5.6	Processing a Memory Load of Points . . . . .	1133
22.5.7	Exercises for Section 22.5 . . . . .	1136
22.6	Summary of Chapter 22 . . . . .	1137
22.7	References for Chapter 22 . . . . .	1139
<b>23</b>	<b>Database Systems and the Internet</b>	<b>1141</b>
23.1	The Architecture of a Search Engine . . . . .	1141
23.1.1	Components of a Search Engine . . . . .	1142
23.1.2	Web Crawlers . . . . .	1143
23.1.3	Query Processing in Search Engines . . . . .	1146
23.1.4	Ranking Pages . . . . .	1146
23.2	PageRank for Identifying Important Pages . . . . .	1147
23.2.1	The Intuition Behind PageRank . . . . .	1147
23.2.2	Recursive Formulation of PageRank — First Try . . . . .	1148
23.2.3	Spider Traps and Dead Ends . . . . .	1150
23.2.4	PageRank Accounting for Spider Traps and Dead Ends . . . . .	1153
23.2.5	Exercises for Section 23.2 . . . . .	1154
23.3	Topic-Specific PageRank . . . . .	1156
23.3.1	Teleport Sets . . . . .	1156
23.3.2	Calculating A Topic-Specific PageRank . . . . .	1158
23.3.3	Link Spam . . . . .	1159
23.3.4	Topic-Specific PageRank and Link Spam . . . . .	1160
23.3.5	Exercises for Section 23.3 . . . . .	1161
23.4	Data Streams . . . . .	1161
23.4.1	Data-Stream-Management Systems . . . . .	1162
23.4.2	Stream Applications . . . . .	1163
23.4.3	A Data-Stream Data Model . . . . .	1164
23.4.4	Converting Streams Into Relations . . . . .	1165
23.4.5	Converting Relations Into Streams . . . . .	1166
23.4.6	Exercises for Section 23.4 . . . . .	1168
23.5	Data Mining of Streams . . . . .	1169
23.5.1	Motivation . . . . .	1169
23.5.2	Counting Bits . . . . .	1171
23.5.3	Counting the Number of Distinct Elements . . . . .	1175
23.5.4	Exercises for Section 23.5 . . . . .	1176
23.6	Summary of Chapter 23 . . . . .	1177
23.7	References for Chapter 23 . . . . .	1179
	<b>Index</b>	<b>1183</b>