

Contents

Chapter 1 Introduction

- | | | | |
|----------------------------------|----|---|----|
| 1.1 Database-System Applications | 1 | 1.7 Database and Application Architecture | 21 |
| 1.2 Purpose of Database Systems | 5 | 1.8 Database Users and Administrators | 24 |
| 1.3 View of Data | 8 | 1.9 History of Database Systems | 25 |
| 1.4 Database Languages | 13 | 1.10 Summary | 29 |
| 1.5 Database Design | 17 | Exercises | 31 |
| 1.6 Database Engine | 18 | Further Reading | 33 |

PART ONE ■ RELATIONAL LANGUAGES

Chapter 2 Introduction to the Relational Model

- | | | | |
|---------------------------------------|----|----------------------------|----|
| 2.1 Structure of Relational Databases | 37 | 2.6 The Relational Algebra | 48 |
| 2.2 Database Schema | 41 | 2.7 Summary | 58 |
| 2.3 Keys | 43 | Exercises | 60 |
| 2.4 Schema Diagrams | 46 | Further Reading | 63 |
| 2.5 Relational Query Languages | 47 | | |

Chapter 3 Introduction to SQL

- | | | | |
|--|----|----------------------------------|-----|
| 3.1 Overview of the SQL Query Language | 65 | 3.7 Aggregate Functions | 91 |
| 3.2 SQL Data Definition | 66 | 3.8 Nested Subqueries | 98 |
| 3.3 Basic Structure of SQL Queries | 71 | 3.9 Modification of the Database | 108 |
| 3.4 Additional Basic Operations | 79 | 3.10 Summary | 114 |
| 3.5 Set Operations | 85 | Exercises | 115 |
| 3.6 Null Values | 89 | Further Reading | 124 |

Chapter 4 Intermediate SQL

- 4.1 Join Expressions 125
- 4.2 Views 137
- 4.3 Transactions 143
- 4.4 Integrity Constraints 145
- 4.5 SQL Data Types and Schemas 153
- 4.6 Index Definition in SQL 164
- 4.7 Authorization 165
- 4.8 Summary 173
- Exercises 176
- Further Reading 180

Chapter 5 Advanced SQL

- 5.1 Accessing SQL from a Programming Language 183
- 5.2 Functions and Procedures 198
- 5.3 Triggers 206
- 5.4 Recursive Queries 213
- 5.5 Advanced Aggregation Features 219
- 5.6 Summary 231
- Exercises 232
- Further Reading 238

PART TWO ■ DATABASE DESIGN

Chapter 6 Database Design Using the E-R Model

- 6.1 Overview of the Design Process 241
- 6.2 The Entity-Relationship Model 244
- 6.3 Complex Attributes 249
- 6.4 Mapping Cardinalities 252
- 6.5 Primary Key 256
- 6.6 Removing Redundant Attributes in Entity Sets 261
- 6.7 Reducing E-R Diagrams to Relational Schemas 264
- 6.8 Extended E-R Features 271
- 6.9 Entity-Relationship Design Issues 279
- 6.10 Alternative Notations for Modeling Data 285
- 6.11 Other Aspects of Database Design 291
- 6.12 Summary 292
- Exercises 294
- Further Reading 300

Chapter 7 Relational Database Design

- 7.1 Features of Good Relational Designs 303
- 7.2 Decomposition Using Functional Dependencies 308
- 7.3 Normal Forms 313
- 7.4 Functional-Dependency Theory 320
- 7.5 Algorithms for Decomposition Using Functional Dependencies 330
- 7.6 Decomposition Using Multivalued Dependencies 336
- 7.7 More Normal Forms 341
- 7.8 Atomic Domains and First Normal Form 342
- 7.9 Database-Design Process 343
- 7.10 Modeling Temporal Data 347
- 7.11 Summary 351
- Exercises 353
- Further Reading 360

PART THREE ■ APPLICATION DESIGN AND DEVELOPMENT

Chapter 8 Complex Data Types

- | | | | |
|--------------------------|-----|-----------------|-----|
| 8.1 Semi-structured Data | 365 | 8.5 Summary | 394 |
| 8.2 Object Orientation | 376 | Exercises | 397 |
| 8.3 Textual Data | 382 | Further Reading | 401 |
| 8.4 Spatial Data | 387 | | |

Chapter 9 Application Development

- | | | | |
|--|-----|-------------------------------------|-----|
| 9.1 Application Programs and User Interfaces | 403 | 9.7 Application Performance | 434 |
| 9.2 Web Fundamentals | 405 | 9.8 Application Security | 437 |
| 9.3 Servlets | 411 | 9.9 Encryption and Its Applications | 447 |
| 9.4 Alternative Server-Side Frameworks | 416 | 9.10 Summary | 453 |
| 9.5 Client-Side Code and Web Services | 421 | Exercises | 455 |
| 9.6 Application Architectures | 429 | Further Reading | 462 |

PART FOUR ■ BIG DATA ANALYTICS

Chapter 10 Big Data

- | | | | |
|---|-----|----------------------|-----|
| 10.1 Motivation | 467 | 10.5 Streaming Data | 500 |
| 10.2 Big Data Storage Systems | 472 | 10.6 Graph Databases | 508 |
| 10.3 The MapReduce Paradigm | 483 | 10.7 Summary | 511 |
| 10.4 Beyond MapReduce: Algebraic Operations | 494 | Exercises | 513 |
| | | Further Reading | 516 |

Chapter 11 Data Analytics

- | | | | |
|-----------------------------------|-----|-----------------|-----|
| 11.1 Overview of Analytics | 519 | 11.5 Summary | 550 |
| 11.2 Data Warehousing | 521 | Exercises | 552 |
| 11.3 Online Analytical Processing | 527 | Further Reading | 555 |
| 11.4 Data Mining | 540 | | |

PART FIVE ■ STORAGE MANAGEMENT AND INDEXING

Chapter 12 Physical Storage Systems

- 12.1 Overview of Physical Storage Media 559
- 12.2 Storage Interfaces 562
- 12.3 Magnetic Disks 563
- 12.4 Flash Memory 567
- 12.5 RAID 570
- 12.6 Disk-Block Access 577
- 12.7 Summary 580
- Exercises 582
- Further Reading 584

Chapter 13 Data Storage Structures

- 13.1 Database Storage Architecture 587
- 13.2 File Organization 588
- 13.3 Organization of Records in Files 595
- 13.4 Data-Dictionary Storage 602
- 13.5 Database Buffer 604
- 13.6 Column-Oriented Storage 611
- 13.7 Storage Organization in Main-Memory Databases 615
- 13.8 Summary 617
- Exercises 619
- Further Reading 621

Chapter 14 Indexing

- 14.1 Basic Concepts 623
- 14.2 Ordered Indices 625
- 14.3 B⁺-Tree Index Files 634
- 14.4 B⁺-Tree Extensions 650
- 14.5 Hash Indices 658
- 14.6 Multiple-Key Access 661
- 14.7 Creation of Indices 664
- 14.8 Write-Optimized Index Structures 665
- 14.9 Bitmap Indices 670
- 14.10 Indexing of Spatial and Temporal Data 672
- 14.11 Summary 677
- Exercises 679
- Further Reading 683

PART SIX ■ QUERY PROCESSING AND OPTIMIZATION

Chapter 15 Query Processing

- 15.1 Overview 689
- 15.2 Measures of Query Cost 692
- 15.3 Selection Operation 695
- 15.4 Sorting 701
- 15.5 Join Operation 704
- 15.6 Other Operations 719
- 15.7 Evaluation of Expressions 724
- 15.8 Query Processing in Memory 731
- 15.9 Summary 734
- Exercises 736
- Further Reading 740

Chapter 16 Query Optimization

- 16.1 Overview 743
- 16.2 Transformation of Relational Expressions 747
- 16.3 Estimating Statistics of Expression Results 757
- 16.4 Choice of Evaluation Plans 766
- 16.5 Materialized Views 778
- 16.6 Advanced Topics in Query Optimization 783
- 16.7 Summary 787
- Exercises 789
- Further Reading 794

PART SEVEN ■ TRANSACTION MANAGEMENT

Chapter 17 Transactions

- 17.1 Transaction Concept 799
- 17.2 A Simple Transaction Model 801
- 17.3 Storage Structure 804
- 17.4 Transaction Atomicity and Durability 805
- 17.5 Transaction Isolation 807
- 17.6 Serializability 812
- 17.7 Transaction Isolation and Atomicity 819
- 17.8 Transaction Isolation Levels 821
- 17.9 Implementation of Isolation Levels 823
- 17.10 Transactions as SQL Statements 826
- 17.11 Summary 828
- Exercises 831
- Further Reading 834

Chapter 18 Concurrency Control

- 18.1 Lock-Based Protocols 835
- 18.2 Deadlock Handling 849
- 18.3 Multiple Granularity 853
- 18.4 Insert Operations, Delete Operations, and Predicate Reads 857
- 18.5 Timestamp-Based Protocols 861
- 18.6 Validation-Based Protocols 866
- 18.7 Multiversion Schemes 869
- 18.8 Snapshot Isolation 872
- 18.9 Weak Levels of Consistency in Practice 880
- 18.10 Advanced Topics in Concurrency Control 883
- 18.11 Summary 894
- Exercises 899
- Further Reading 904

Chapter 19 Recovery System

- 19.1 Failure Classification 907
- 19.2 Storage 908
- 19.3 Recovery and Atomicity 912
- 19.4 Recovery Algorithm 922
- 19.5 Buffer Management 926
- 19.6 Failure with Loss of Non-Volatile Storage 930
- 19.7 High Availability Using Remote Backup Systems 931
- 19.8 Early Lock Release and Logical Undo Operations 935
- 19.9 ARIES 941
- 19.10 Recovery in Main-Memory Databases 947
- 19.11 Summary 948
- Exercises 952
- Further Reading 956

PART EIGHT ■ PARALLEL AND DISTRIBUTED DATABASES

Chapter 20 Database-System Architectures

- 20.1 Overview 961
- 20.2 Centralized Database Systems 962
- 20.3 Server System Architectures 963
- 20.4 Parallel Systems 970
- 20.5 Distributed Systems 986
- 20.6 Transaction Processing in Parallel and Distributed Systems 989
- 20.7 Cloud-Based Services 990
- 20.8 Summary 995
- Exercises 998
- Further Reading 1001

Chapter 21 Parallel and Distributed Storage

- 21.1 Overview 1003
- 21.2 Data Partitioning 1004
- 21.3 Dealing with Skew in Partitioning 1007
- 21.4 Replication 1013
- 21.5 Parallel Indexing 1017
- 21.6 Distributed File Systems 1019
- 21.7 Parallel Key-Value Stores 1023
- 21.8 Summary 1032
- Exercises 1033
- Further Reading 1036

Chapter 22 Parallel and Distributed Query Processing

- 22.1 Overview 1039
- 22.2 Parallel Sort 1041
- 22.3 Parallel Join 1043
- 22.4 Other Operations 1048
- 22.5 Parallel Evaluation of Query Plans 1052
- 22.6 Query Processing on Shared-Memory Architectures 1061
- 22.7 Query Optimization for Parallel Execution 1064
- 22.8 Parallel Processing of Streaming Data 1070
- 22.9 Distributed Query Processing 1076
- 22.10 Summary 1086
- Exercises 1089
- Further Reading 1093

Chapter 23 Parallel and Distributed Transaction Processing

- 23.1 Distributed Transactions 1098
- 23.2 Commit Protocols 1100
- 23.3 Concurrency Control in Distributed Databases 1111
- 23.4 Replication 1121
- 23.5 Extended Concurrency Control Protocols 1129
- 23.6 Replication with Weak Degrees of Consistency 1133
- 23.7 Coordinator Selection 1146
- 23.8 Consensus in Distributed Systems 1150
- 23.9 Summary 1162
- Exercises 1165
- Further Reading 1168

PART NINE ■ ADVANCED TOPICS

Chapter 24 Advanced Indexing Techniques

- | | | | |
|---|------|-------------------|------|
| 24.1 Bloom Filter | 1175 | 24.5 Hash Indices | 1190 |
| 24.2 Log-Structured Merge Tree and Variants | 1176 | 24.6 Summary | 1203 |
| 24.3 Bitmap Indices | 1182 | Exercises | 1205 |
| 24.4 Indexing of Spatial Data | 1186 | Further Reading | 1206 |

Chapter 25 Advanced Application Development

- | | | | |
|--|------|------------------------------------|------|
| 25.1 Performance Tuning | 1210 | 25.5 Distributed Directory Systems | 1240 |
| 25.2 Performance Benchmarks | 1230 | 25.6 Summary | 1243 |
| 25.3 Other Issues in Application Development | 1234 | Exercises | 1245 |
| 25.4 Standardization | 1237 | Further Reading | 1248 |

Chapter 26 Blockchain Databases

- | | | | |
|---|------|------------------------------|------|
| 26.1 Overview | 1252 | 26.6 Smart Contracts | 1269 |
| 26.2 Blockchain Properties | 1254 | 26.7 Performance Enhancement | 1274 |
| 26.3 Achieving Blockchain Properties via Cryptographic Hash Functions | 1259 | 26.8 Emerging Applications | 1276 |
| 26.4 Consensus | 1263 | 26.9 Summary | 1279 |
| 26.5 Data Management in a Blockchain | 1267 | Exercises | 1280 |
| | | Further Reading | 1282 |

PART TEN ■ APPENDIX A

Appendix A Detailed University Schema 1287

Index 1299

PART ELEVEN ■ ONLINE CHAPTERS

Chapter 27 Formal Relational Query Languages

Chapter 28 Advanced Relational Database Design

Chapter 29 Object-Based Databases

Chapter 30 XML

Chapter 31 Information Retrieval

Chapter 32 PostgreSQL