

Contents

Preface

xxiii

PART ONE Introduction

1

1 Overview of Databases and Transactions

3

1.1 What Are Databases and Transactions? 3

1.2 Features of Modern Database and Transaction Processing Systems 6

1.3 Major Players in the Implementation and Support of Database and Transaction Processing Systems 7

1.4 Decision Support Systems—OLAP and OLTP 9

2 The Big Picture

13

© 2.1 Case Study: A Student Registration System 13

2.2 Introduction to Relational Databases 14

2.3 What Makes a Program a Transaction—The ACID Properties 20

Bibliographic Notes 25

Exercises 25

PART TWO Database Management

29

3 The Relational Data Model

31

3.1 What Is a Data Model? 31

3.2 The Relational Model 35

3.2.1 Basic Concepts 35

3.2.2 Integrity Constraints 38

3.3 SQL—Data Definition Sublanguage 46

3.3.1 Specifying the Relation Type 46

3.3.2	The System Catalog	46
3.3.3	Key Constraints	47
3.3.4	Dealing with Missing Information	48
3.3.5	Semantic Constraints	49
3.3.6	User-Defined Domains	53
3.3.7	Foreign-Key Constraints	53
3.3.8	Reactive Constraints	56
3.3.9	Database Views	59
3.3.10	Modifying Existing Definitions	60
3.3.11	SQL-Schemas	62
3.3.12	Access Control	63
	Bibliographic Notes	65
	Exercises	66

4	Conceptual Modeling of Databases with Entity-Relationship Diagrams and the Unified Modeling Language	69
4.1	Conceptual Modeling with the E-R Approach	70
4.2	Entities and Entity Types	70
4.3	Relationships and Relationship Types	73
4.4	Advanced Features in Conceptual Data Modeling	78
4.4.1	Entity Type Hierarchies	78
4.4.2	Participation Constraints	81
4.4.3	The Part-of Relationship	83
4.5	From E-R Diagrams to Relational Database Schemas	86
4.5.1	Representation of Entities	86
4.5.2	Representation of Relationships	88
4.5.3	Representing IsA Hierarchies in the Relational Model	90
4.5.4	Representation of Participation Constraints	92
4.5.5	Representation of the Part-of Relationship	94
★ 4.6	UML: A New Kid on the Block	95
4.6.1	Representing Entities in UML	96
4.6.2	Representing Relationships in UML	97
4.6.3	Advanced Modeling Concepts in UML	101
4.6.4	Translation to SQL	105
4.7	A Brokerage Firm Example	106
4.7.1	An Entity-Relationship Design	106
★ 4.7.2	A UML Design	110
CS 4.8	Case Study: A Database Design for the Student Registration System	111
4.8.1	The Database Part of the Requirements Document	112
4.8.2	The Database Design	113

- 4.9 Limitations of Data Modeling Methodologies 119
- Bibliographic Notes 123
- Exercises 123

5 Relational Algebra and SQL 127

- 5.1 Relational Algebra: Under the Hood of SQL 128
 - 5.1.1 Basic Operators 128
 - 5.1.2 Derived Operators 137
- 5.2 The Query Sublanguage of SQL 147
 - 5.2.1 Simple SQL Queries 148
 - 5.2.2 Set Operations 154
 - 5.2.3 Nested Queries 157
 - 5.2.4 Quantified Predicates 163
 - 5.2.5 Aggregation over Data 164
 - 5.2.6 A Query Evaluation Algorithm for SQL with Aggregates 170
 - 5.2.7 Join Expressions in the FROM Clause 173
 - 5.2.8 More on Views in SQL 174
 - 5.2.9 Materialized Views 177
 - 5.2.10 The Null Value Quandary 181
- 5.3 Modifying Relation Instances in SQL 182
 - 5.3.1 Inserting Data 182
 - 5.3.2 Deleting Data 184
 - 5.3.3 Updating Existing Data 185
 - 5.3.4 Updates on Views 185
 - Bibliographic Notes 187
 - Exercises 188

6 Database Design with the Relational Normalization Theory 193

- 6.1 The Problem of Redundancy 193
- 6.2 Decompositions 195
- 6.3 Functional Dependencies 198
- 6.4 Properties of Functional Dependencies 200
- 6.5 Normal Forms 207
 - 6.5.1 The Boyce-Codd Normal Form 208
 - 6.5.2 The Third Normal Form 210
- 6.6 Properties of Decompositions 211
 - 6.6.1 Lossless and Lossy Decompositions 212
 - 6.6.2 Dependency-Preserving Decompositions 215
- 6.7 An Algorithm for BCNF Decomposition 219

6.8	Synthesis of 3NF Schemas	222	
6.8.1	Minimal Cover	222	
6.8.2	3NF Decomposition through Schema Synthesis	224	
6.8.3	BCNF Decomposition through 3NF Synthesis	227	
★ 6.9	The Fourth Normal Form	229	
★ 6.10	Advanced 4NF Design	235	
6.10.1	MVDs and Their Properties	235	
6.10.2	The Difficulty of Designing for 4NF	236	
6.10.3	A 4NF Decomposition How-To	240	
6.11	Summary of Normal Form Decomposition	242	
Ⓢ 6.12	Case Study: Schema Refinement for the Student Registration System	242	
6.13	Tuning Issues: To Decompose or Not to Decompose?	245	
	Bibliographic Notes	246	
	Exercises	247	
7	Triggers and Active Databases		251
7.1	What Is a Trigger?	251	
7.2	Semantic Issues in Trigger Handling	252	
7.3	Triggers in SQL	256	
7.4	Avoiding a Chain Reaction	264	
	Bibliographic Notes	265	
	Exercises	265	
8	Using SQL in an Application		267
8.1	What Are the Issues Involved?	267	
8.2	Embedded SQL	268	
8.2.1	Status Processing	271	
8.2.2	Sessions, Connections, and Transactions	273	
8.2.3	Executing Transactions	274	
8.2.4	Cursors	276	
8.2.5	Stored Procedures on the Server	282	
8.3	More on Integrity Constraints	285	
8.4	Dynamic SQL	286	
8.4.1	Statement Preparation in Dynamic SQL	287	
★ 8.4.2	Prepared Statements and the Descriptor Area	290	
8.4.3	Cursors	293	
8.4.4	Stored Procedures on the Server	293	
8.5	JDBC and SQLJ	294	
8.5.1	JDBC Basics	294	

	8.5.2 Prepared Statements	297
	8.5.3 Result Sets and Cursors	297
	8.5.4 Obtaining Information about a Result Set	300
	8.5.5 Status Processing	300
	8.5.6 Executing Transactions	301
	8.5.7 Stored Procedures on the Server	302
	8.5.8 An Example	303
	8.5.9 SQLJ: Statement-Level Interface to Java	303
★ 8.6	ODBC	307
	8.6.1 Prepared Statements	309
	8.6.2 Cursors	309
	8.6.3 Status Processing	312
	8.6.4 Executing Transactions	312
	8.6.5 Stored Procedures on the Server	313
	8.6.6 An Example	313
8.7	Comparison	315
	Bibliographic Notes	316
	Exercises	316

PART THREE Optimizing DBMS Performance 319

9 Physical Data Organization and Indexing 321

9.1	Disk Organization	322
	9.1.1 RAID Systems	326
9.2	Heap Files	329
9.3	Sorted Files	333
9.4	Indices	337
	9.4.1 Clustered versus Unclustered Indices	340
	9.4.2 Sparse versus Dense Indices	342
	9.4.3 Search Keys Containing Multiple Attributes	344
9.5	Multilevel Indexing	347
	9.5.1 Index-Sequential Access	350
	9.5.2 B ⁺ Trees	353
9.6	Hash Indexing	363
	9.6.1 Static Hashing	364
★ 9.6.2	Dynamic Hashing Algorithms	366
★ 9.7	Special-Purpose Indices	375
	9.7.1 Bitmap Indices	375
	9.7.2 Join Indices	376

9.8	Tuning Issues: Choosing Indices for an Application	377
	Bibliographic Notes	378
	Exercises	378

10 The Basics of Query Processing **383**

10.1	Overview of Query Processing	383
10.2	External Sorting	384
10.3	Computing Projection, Union, and Set Difference	388
10.4	Computing Selection	390
	10.4.1 Selections with Simple Conditions	391
	10.4.2 Access Paths	393
	10.4.3 Selections with Complex Conditions	395
10.5	Computing Joins	396
	10.5.1 Computing Joins Using Simple Nested Loops	397
	10.5.2 Sort-Merge Join	400
	10.5.3 Hash Join	402
★ 10.6	Multirelational Joins	403
10.7	Computing Aggregate Functions	405
	Bibliographic Notes	405
	Exercises	405

11 An Overview of Query Optimization **409**

11.1	Query Processing Architecture	409
11.2	Heuristic Optimization Based on Algebraic Equivalences	411
11.3	Estimating the Cost of a Query Execution Plan	414
11.4	Estimating the Size of the Output	422
11.5	Choosing a Plan	424
	Bibliographic Notes	429
	Exercises	429

12 Database Tuning **433**

12.1	Disk Caches	434
	12.1.1 Tuning the Cache	435
12.2	Tuning the Schema	437
	12.2.1 Indices	437
	12.2.2 Denormalization	444
	12.2.3 Repeating Groups	446
	12.2.4 Partitioning	446
12.3	Tuning the Data Manipulation Language	447
12.4	Tools	451

- 12.5 Managing Physical Resources 451
- 12.6 Influencing the Optimizer 453
 - Bibliographic Notes 455
 - Exercises 455

PART FOUR Advanced Topics in Databases 459

13 Relational Calculus, Visual Query Languages, and Deductive Databases 461

- 13.1 Tuple Relational Calculus 461
- 13.2 Understanding SQL through Tuple Relational Calculus 471
- 13.3 Domain Relational Calculus and Visual Query Languages 474
- 13.4 Visual Query Languages: QBE and PC Databases 479
- 13.5 The Relationship between Relational Algebra and the Calculi 486
- ★ 13.6 Deductive Databases 488
 - 13.6.1 Limitations of Relational Query Languages 488
 - 13.6.2 Recursive Queries in SQL 490
 - 13.6.3 Datalog 496
 - Bibliographic Notes 510
 - Exercises 510

14 Object Databases 515

- 14.1 Limitations of the Relational Data Model 515
- 14.2 Object Databases versus Relational Databases 521
- 14.3 The Conceptual Object Data Model 523
 - 14.3.1 Objects and Values 523
 - 14.3.2 Classes 525
 - 14.3.3 Types 526
 - 14.3.4 Object-Relational Databases 529
- 14.4 Objects in SQL:1999 and SQL:2003 530
 - 14.4.1 Row Types 531
 - 14.4.2 User-Defined Types 531
 - 14.4.3 Objects 533
 - 14.4.4 Querying User-Defined Types 534
 - 14.4.5 Updating User-Defined Types 535
 - 14.4.6 Reference Types 538
 - 14.4.7 Inheritance 540
 - 14.4.8 Collection Types 540

14.5	The ODMG Standard	543
14.5.1	ODL—The ODMG Object Definition Language	546
14.5.2	OQL—The ODMG Object Query Language	552
14.5.3	Transactions in ODMG	557
14.5.4	Object Manipulation in ODMG	557
14.5.5	Language Bindings	558
★ 14.6	Common Object Request Broker Architecture	562
14.6.1	CORBA Basics	563
14.6.2	CORBA and Databases	569
	Bibliographic Notes	573
	Exercises	575
15	XML and Web Data	579
15.1	Semistructured Data	579
15.2	Overview of XML	582
15.2.1	XML Elements and Database Objects	585
15.2.2	XML Attributes	587
15.2.3	Namespaces	589
15.2.4	Document Type Definitions	594
15.2.5	Inadequacy of DTDs as a Data Definition Language	596
15.3	XML Schema	599
15.3.1	XML Schema and Namespaces	599
15.3.2	Simple Types	603
15.3.3	Complex Types	608
15.3.4	Putting It Together	616
15.3.5	Shortcuts: Anonymous Types and Element References	616
15.3.6	Integrity Constraints	620
15.4	XML Query Languages	627
15.4.1	XPath: A Lightweight XML Query Language	628
★ 15.4.2	XSLT: A Transformation Language for XML	637
15.4.3	XQuery: A Full-Featured Query Language for XML	649
15.4.4	SQL/XML	668
	Bibliographic Notes	679
	Exercises	680
16	Distributed Databases	687
16.1	The Application Designer's View of the Database	688
16.2	Distributing Data among Different Databases	691
16.2.1	Partitioning	692
16.2.2	Updates and Partitioning	695

- 16.2.3 Replication 696
- 16.3 Query Planning Strategies 698
 - 16.3.1 Global Query Optimization 698
 - 16.3.2 Strategies for a Multidatabase System 705
 - 16.3.3 Tuning Issues: Database Design and Query Planning in a Distributed Environment 706
- Bibliographic Notes 707
- Exercises 707

17 OLAP and Data Mining 711

- 17.1 OLAP and Data Warehouses—Old and New 711
- 17.2 A Multidimensional Model for OLAP Applications 713
- 17.3 Aggregation 717
 - 17.3.1 Drilling, Slicing, Rolling, and Dicing 718
 - 17.3.2 The CUBE Operator 721
- 17.4 ROLAP and MOLAP 725
- 17.5 Implementation Issues 727
- 17.6 Populating a Data Warehouse 728
- 17.7 Data Mining Tasks 730
- 17.8 Mining Associations 731
- 17.9 Classification and Prediction Using Decision Trees 734
- ★ 17.10 Classification and Prediction Using Neural Nets 744
- 17.11 Clustering 752
 - Bibliographic Notes 757
 - Exercises 757

PART FIVE Transaction Processing 761

18 ACID Properties of Transactions 763

- 18.1 Consistency 764
 - 18.1.1 Checking Integrity Constraints 765
 - 18.1.2 A Transaction as a Unit of Work 766
- 18.2 Atomicity 767
- 18.3 Durability 768
- 18.4 Isolation 769
- 18.5 The ACID Properties 773
 - Bibliographic Notes 774
 - Exercises 774

19 Models of Transactions

777

- 19.1 Flat Transactions 777
- 19.2 Providing Structure within a Transaction 779
 - 19.2.1 Savepoints 779
 - 19.2.2 Distributed Transactions 781
 - 19.2.3 Nested Transactions 785
- 19.3 Structuring an Application as Multiple Transactions 787
 - 19.3.1 Chained Transactions 788
 - 19.3.2 Sagas and Compensation 791
 - 19.3.3 Declarative Transaction Demarcation 793
 - ★ 19.3.4 Multilevel Transactions 796
 - 19.3.5 Transaction Scheduling with Recoverable Queues 799
 - 19.3.6 Workflows and Workflow Management Systems 804
- Bibliographic Notes 809
- Exercises 810

20 Implementing Isolation

813

- 20.1 Schedules and Schedule Equivalence 815
 - 20.1.1 Serializability 819
 - 20.1.2 Conflict Equivalence and View Equivalence 821
 - 20.1.3 Serialization Graphs 822
- 20.2 Recoverability, Cascaded Aborts, and Strictness 824
- 20.3 Models for Concurrency Control 827
- 20.4 A Strategy for Immediate-Update Pessimistic Concurrency Controls 829
 - 20.4.1 Conflict Avoidance 830
 - 20.4.2 Deadlocks 832
- 20.5 Design of an Immediate-Update Pessimistic Concurrency Control 834
 - 20.5.1 An Implementation Using Lock Sets and Wait Sets 834
 - 20.5.2 Two-Phase Locking 836
 - 20.5.3 Lock Granularity 838
- 20.6 Objects and Semantic Commutativity 839
 - ★ 20.6.1 Partial Operations and Backward-Commutativity 840
- 20.7 Atomicity, Recoverability, and Compensating Operations 842
- 20.8 Isolation in Structured Transaction Models 848
 - 20.8.1 Savepoints 848
 - 20.8.2 Chained Transactions 849
 - 20.8.3 Recoverable Queues 849

- 20.8.4 Nested Transactions 850
- ★ 20.8.5 Multilevel Transactions 851
- 20.9 Other Concurrency Controls 856
 - 20.9.1 Timestamp-Ordered Concurrency Controls 856
 - 20.9.2 Optimistic Concurrency Controls 859
- Bibliographic Notes 863
- Exercises 863

21 Isolation in Relational Databases 869

- 21.1 Conflicts in a Relational Database 869
 - 21.1.1 Phantoms 870
 - 21.1.2 Predicate Locking 872
- 21.2 Locking and the SQL Isolation Levels 875
 - 21.2.1 Lost Updates, Cursor Stability, and Update Locks 880
 - © 21.2.2 Case Study: Correctness and NonSERIALIZABLE Schedules—
The Student Registration System 883
 - 21.2.3 Serializable, SERIALIZABLE, and Correct 887
- 21.3 Granular Locking: Intention Locks and Index Locks 887
 - 21.3.1 Index Locks: Granular Locking without Phantoms 890
 - ★ 21.3.2 Granular Locking in an Object Database 900
- 21.4 Tuning Transactions 901
- 21.5 Multiversion Concurrency Controls 903
 - 21.5.1 Read-Only Multiversion Concurrency Control 904
 - 21.5.2 Read-Consistency Multiversion Concurrency Controls 906
 - © 21.5.3 Case Study: SNAPSHOT Isolation 906
- Bibliographic Notes 912
- Exercises 912

22 Atomicity and Durability 919

- 22.1 Crash, Abort, and Media Failure 919
- 22.2 Immediate-Update Systems and Write-Ahead Logs 921
 - 22.2.1 Performance and Write-Ahead Logging 925
 - 22.2.2 Checkpoints and Recovery 928
 - ★ 22.2.3 Logical and Physiological Logging 934
- 22.3 Recovery in Deferred-Update Systems 936
- 22.4 Recovery from Media Failure 937
- Bibliographic Notes 941
- Exercises 941

PART SIX	Distributed Applications and the Web	945
23	Architecture of Transaction Processing Systems	947
23.1	Transaction Processing in a Centralized System	947
23.1.1	Organization of a Single-User System	947
23.1.2	Organization of a Centralized Multiuser System	949
23.2	Transaction Processing in a Distributed System	950
23.2.1	Organization of a Distributed System	951
23.2.2	Sessions and Context	957
23.2.3	Queued Transaction Processing	960
23.3	The TP Monitor: An Overview	961
23.3.1	The Services Provided by a TP Monitor	963
23.4	The TP Monitor: Global Atomicity and the Transaction Manager	965
23.5	The TP Monitor: Remote Procedure Call	967
23.5.1	Implementation of Remote Procedure Call	968
23.5.2	Directory Services	970
23.5.3	The Transaction Manager and Transactional RPC	971
23.6	The TP Monitor: Peer-to-Peer Communication	974
23.6.1	Establishing a Connection	975
23.6.2	Distributed Commitment	976
23.7	The TP Monitor: Event Communication	977
23.7.1	Event Broker	979
23.8	Storage Architectures	981
23.9	Transaction Processing on the Internet	982
23.9.1	Architectures for C2B Transaction Processing Systems on the Internet	983
23.10	Web Application Servers—J2EE	985
23.10.1	Enterprise Java Beans	986
23.10.2	The EJB Container	992
23.10.3	Using Java Beans	1001
	Bibliographic Notes	1002
	Exercises	1003
24	Implementing Distributed Transactions	1005
24.1	Implementing the ACID Properties	1005
24.2	Atomic Termination	1007
24.2.1	The Two-Phase Commit Protocol	1008
24.2.2	Dealing with Failures in the Two-Phase Commit Protocol	1013
24.2.3	The Peer-to-Peer Atomic Commit Protocol	1020
24.3	Transfer of Coordination	1021

24.3.1	The Linear Commit Protocol	1022
24.3.2	Two-Phase Commit without a Prepared State	1023
24.4	Distributed Deadlock	1023
24.5	Global Serialization	1024
24.6	When Global Atomicity Cannot Be Guaranteed	1026
24.6.1	Weaker Commit Protocols	1027
24.7	Replicated Databases	1028
24.7.1	Synchronous-Update Replication Systems	1031
24.7.2	Asynchronous-Update Replication Systems	1033
24.8	Distributed Transactions in the Real World	1038
	Bibliographic Notes	1038
	Exercises	1038

25 Web Services

1043

25.1	The Basic Idea	1043
25.2	Web Basics	1047
25.3	Hypertext Transfer Protocol	1048
25.4	SOAP: Message Passing	1051
25.4.1	SOAP and Remote Procedure Call	1053
25.4.2	SOAP Extensibility	1055
25.4.3	SOAP Faults	1061
25.4.4	SOAP Binding	1061
25.5	WSDL: Specifying Web Services	1063
25.5.1	The Abstract Level	1063
25.5.2	The Concrete Level	1067
25.5.3	Putting It All Together	1077
25.5.4	WSDL Version 2.0	1078
25.6	BPEL: Specifying Business Processes	1080
25.6.1	Communication	1082
25.6.2	Processes	1088
25.6.3	Structured Activities	1092
25.6.4	Links	1095
25.6.5	BPEL and WS-Addressing	1097
25.6.6	Handling Errors	1099
25.6.7	Handling Multiple Requests	1105
25.6.8	Front-End and Back-End Systems	1108
25.6.9	Interacting with a Web Service: Projection of a BPEL Process	1109
25.7	UDDI: Publishing and Discovering Information about Services	1110
25.7.1	Data Structures in the UDDI Registry	1111
25.7.2	The Inquiry Interface (Query Language)	1117
25.7.3	The Publisher Interface (Update Language)	1120

25.7.4	Some Final Observations about UDDI	1123
25.8	WS-Coordination: Transactional Web Services	1123
	Bibliographic Notes	1131
	Exercises	1131
26	Security and Electronic Commerce	1135
26.1	Authentication, Authorization, and Encryption	1135
26.2	Encryption	1136
26.3	Digital Signatures	1141
26.4	Key Distribution and Authentication	1143
26.4.1	The Kerberos Protocol: Tickets	1144
26.4.2	Nonces	1148
26.5	Authorization	1149
26.6	Authenticated Remote Procedure Call	1151
26.7	Electronic Commerce	1152
26.8	The Secure Sockets Layer Protocol: Certificates	1153
26.9	Passport: Single Sign-On	1155
26.10	Keeping Credit Card Numbers Private	1157
★ 26.11	The Secure Electronic Transaction Protocol: Dual Signatures	1158
★ 26.12	Goods Atomicity, Certified Delivery, and Escrow	1162
★ 26.13	Electronic Cash: Blind Signatures	1165
26.14	Security in XML-Based Web Services	1171
26.14.1	Encryption and Signatures—XML Encryption and XML Signature	1171
26.14.2	Encrypting and Signing SOAP Messages—WS-Security	1176
26.14.3	SAML: Authentication, Authorization, and Single Sign-On	1180
	Bibliographic Notes	1184
	Exercises	1185
	Bibliography	1187
	Index	1203
	Appendices, available on the Web¹	
A	An Overview of Transaction Processing	A-1
A.1	Isolation	A-1
A.1.1	Serializability	A-2

¹ Available on the Web at <http://www.aw-bc.com/kifer>