

CONTENTS

About the Author	xix
Introduction to the Fourth Edition.....	xxi
Chapter 1 Databases versus File Systems	1
1.1 Tables as Entities	4
1.2 Tables as Relationships	4
1.3 Rows versus Records.....	5
1.4 Columns versus Fields.....	6
1.5 Schema Objects	7
1.6 CREATE SCHEMA Statement	7
Chapter 2 Transactions and Concurrency Control	11
2.1 Sessions	11
2.2 Transactions and ACID.....	12
2.3 Concurrency Control	14
2.4 Pessimistic Concurrency Control	18
2.5 SNAPSHOT Isolation and Optimistic Concurrency.....	19
2.6 Logical Concurrency Control	21
2.7 Deadlock and Livelocks	21
Chapter 3 Schema Level Objects	23
3.1 CREATE SCHEMA Statement	23
3.2 CREATE PROCEDURE, CREATE FUNCTION, and CREATE TRIGGER.....	24
3.3 CREATE DOMAIN Statement	24
3.4 CREATE SEQUENCE	25
3.5 CREATE ASSERTION	26
3.6 Character Set Related Constructs.....	31
Chapter 4 Locating Data and Special Numbers	35
4.1 Exposed Physical Locators	35
4.2 Generated Identifiers.....	40
4.3 Sequence Generator Functions	42
4.4 Preallocated Values	43
4.5 Special Series.....	44

Chapter 5 Base Tables and Related Elements	51
5.1 CREATE TABLE Statement	53
5.2 Nested UNIQUE Constraints	62
5.3 CREATE ASSERTION Constraints	76
5.4 TEMPORARY Tables	77
5.5 Manipulating Tables	78
5.6 Avoiding Attribute Splitting	81
5.7 Modeling Class Hierarchies in DDL	83
5.8 Exposed Physical Locators	85
5.9 Auto-Incrementing Columns	85
5.10 Generated Identifiers	90
5.11 A Remark on Duplicate Rows	94
5.12 Other Schema Objects	96
5.13 Temporary Tables	97
5.14 CREATE DOMAIN Statement	98
5.15 CREATE TRIGGER Statement	99
5.16 CREATE PROCEDURE Statement	99
5.17 DECLARE CURSOR Statement	100
Chapter 6 Procedural, Semiprocedural, and Declarative Programming	105
6.1 Basics of Software Engineering	105
6.2 Cohesion	105
6.3 Coupling	106
6.4 The Big Leap	107
6.5 Rewriting Tricks	114
6.6 Functions for Predicates	118
6.7 Procedural versus Logical Decomposition	119
Chapter 7 Procedural Constructs	123
7.1 CREATE PROCEDURE	123
7.2 CREATE TRIGGER	124
7.3 CURSORS	127
7.4 SEQUENCES	141
7.5 Generated Columns	142
7.6 Table Functions	143

Chapter 8 Auxiliary Tables	145
8.1 The Series Table	145
8.2 Lookup Auxiliary Tables	151
8.3 Auxiliary Function Tables	159
8.4 Global Constants Tables	169
8.5 A Note on Converting Procedural Code to Tables	175
Chapter 9 Normalization	181
9.1 Functional and Multivalued Dependencies	183
9.2 First Normal Form (1NF)	184
9.3 Second Normal Form (2NF)	188
9.4 Third Normal Form (3NF)	189
9.5 Elementary Key Normal Form (EKNF)	191
9.6 Boyce-Codd Normal Form (BCNF)	192
9.7 Fourth Normal Form (4NF)	194
9.8 Fifth Normal Form (5NF)	194
9.9 Domain-Key Normal Form (DKNF)	196
9.10 Practical Hints for Normalization	204
9.11 Key Types	205
9.12 Practical Hints for Denormalization	208
Chapter 10 Numeric Data Types	215
10.1 Numeric Types	215
10.2 Numeric Type Conversion	220
10.3 Four Function Arithmetic	222
10.4 Arithmetic and NULLs	224
10.5 Converting Values to and from NULL	225
10.6 Mathematical Functions	228
10.7 Unique Value Generators	232
10.8 IP Addresses	235
Chapter 11 Temporal Data Types	237
11.1 Notes on Calendar Standards	237
11.2 SQL Temporal Data Types	240
11.3 INTERVAL Data Types	246

11.4	Temporal Arithmetic	249
11.5	The Nature of Temporal Data Models	250
Chapter 12	Character Data Types	255
12.1	Problems with SQL Strings	255
12.2	Standard String Functions	258
12.3	Common Vendor Extensions	259
12.4	Cutter Tables	268
12.5	Nested Replacement	269
Chapter 13	NULLs: Missing Data in SQL	271
13.1	Empty and Missing Tables	272
13.2	Missing Values in Columns	273
13.3	Context and Missing Values	275
13.4	Comparing NULLs	276
13.5	NULLs and Logic	277
13.6	Math and NULLs	281
13.7	Functions and NULLs	281
13.8	NULLs and Host Languages	281
13.9	Design Advice for NULLs	282
13.10	A Note on Multiple NULL Values	285
Chapter 14	Multiple Column Data Elements	289
14.1	Distance Functions	289
14.2	Storing an IPv4 Address in SQL	291
14.3	Storing an IPv6 Address in SQL	293
14.4	Currency and Other Unit Conversions	294
14.5	Social Security Numbers	295
14.6	Rational Numbers	298
Chapter 15	Table Operations	299
15.1	DELETE FROM Statement	299
15.2	INSERT INTO Statement	307
15.3	The UPDATE Statement	310
15.4	A Note on Flaws in a Common Vendor Extension	317
15.5	MERGE Statement	319

Chapter 16 Comparison or Theta Operators	323
16.1 Converting Data Types	323
16.2 Row Comparisons in SQL	326
16.3 IS [NOT] DISTINCT FROM Operator	328
Chapter 17 Valued Predicates	329
17.1 IS NULL	329
17.2 IS [NOT]{TRUE FALSE UNKNOWN} Predicate	330
17.3 IS [NOT] NORMALIZED Predicate	332
Chapter 18 CASE Expressions	333
18.1 The CASE Expression	333
18.2 Subquery Expressions and Constants	341
18.3 Rozenshtein Characteristic Functions	342
Chapter 19 LIKE and SIMILAR TO Predicates	345
19.1 Tricks with Patterns	346
19.2 Results with NULL Values and Empty Strings	347
19.3 LIKE Is Not Equality	348
19.4 Avoiding the LIKE Predicate with a Join	348
19.5 CASE Expressions and LIKE Search Conditions	349
19.6 SIMILAR TO Predicates	350
19.7 Tricks with Strings	352
Chapter 20 BETWEEN and OVERLAPS Predicates	355
20.1 The BETWEEN Predicate	355
20.2 OVERLAPS Predicate	358
Chapter 21 The [NOT] IN() Predicate	369
21.1 Optimizing the IN() Predicate	370
21.2 Replacing ORs with the IN() Predicate	373
21.3 NULLs and the IN() Predicate	374
21.4 IN() Predicate and Referential Constraints	376
21.5 IN() Predicate and Scalar Queries	377

Chapter 22 EXISTS() Predicate	381
22.1 EXISTS and NULLs	382
22.2 EXISTS and INNER JOINS	384
22.3 NOT EXISTS and OUTER JOINS	385
22.4 EXISTS() and Quantifiers	385
22.5 EXISTS() and Referential Constraints	386
22.6 EXISTS and Three-Valued Logic	387
Chapter 23 Quantified Subquery Predicates	389
23.1 Scalar Subquery Comparisons	389
23.2 Quantifiers and Missing Data	391
23.3 The ALL Predicate and Extrema Functions	393
23.4 The UNIQUE Predicate	394
23.5 The DISTINCT Predicate	395
Chapter 24 The Simple SELECT Statement	397
24.1 SELECT Statement Execution Order	397
24.2 One-Level SELECT Statement	397
Chapter 25 Advanced SELECT Statements	407
25.1 Correlated Subqueries	407
25.2 Infix INNER JOINS	411
25.3 OUTER JOINS	413
25.4 UNION JOIN Operators	425
25.5 Scalar SELECT Expressions	426
25.6 Old versus New JOIN Syntax	427
25.7 Constrained JOINS	428
25.8 Dr. Codd's T-Join	437
Chapter 26 Virtual Tables: VIEWS, Derived Tables, CTEs, and MQTs	445
26.1 VIEWS in Queries	445
26.2 Updatable and Read-Only VIEWS	446
26.3 Types of VIEWS	448
26.4 How VIEWS Are Handled in the Database Engine	453
26.5 WITH CHECK OPTION Clause	457

26.6 Dropping VIEWS	462
26.7 Hints on Using VIEWS versus TEMPORARY TABLES.	463
26.8 Using Derived Tables	466
26.9 Common Table Expressions	468
26.10 Recursive Common Table Expressions.	469
26.11 Materialized Query Tables	472
Chapter 27 Partitioning Data in Queries	473
27.1 Coverings and Partitions	473
27.2 Relational Division	478
27.3 Romley's Division	485
27.4 Boolean Expressions in an RDBMS	489
27.5 FIFO and LIFO Subsets	490
Chapter 28 Grouping Operations	493
28.1 GROUP BY Clause	493
28.2 GROUP BY and HAVING	495
28.3 Multiple Aggregation Levels	498
28.4 Grouping on Computed Columns	501
28.5 Grouping into Pairs	502
28.6 Sorting and GROUP BY	504
Chapter 29 Simple Aggregate Functions.	507
29.1 COUNT() Functions	508
29.2 SUM() Function	511
29.3 AVG() Function	512
29.4 Extrema Functions	517
29.5 The LIST() Aggregate Function	530
29.6 The PRD() Aggregate Function	532
29.7 Bitwise Aggregate Functions	536
Chapter 30 Advanced Grouping, Windowed Aggregation, and OLAP in SQL	539
30.1 Star Schema	540
30.2 GROUPING Operators	540
30.3 The Window Clause	544
30.4 Windowed Aggregate Functions	547

30.5 Ordinal Functions	547
30.6 Vendor Extensions	550
30.7 A Bit of History	553
Chapter 31 Descriptive Statistics in SQL	555
31.1 The Mode	555
31.2 The AVG() Function	556
31.3 The Median	557
31.4 Variance and Standard Deviation	572
31.5 Average Deviation	573
31.6 Cumulative Statistics	573
31.7 Cross Tabulations	582
31.8 Harmonic Mean and Geometric Mean	588
31.9 Multivariable Descriptive Statistics in SQL	589
31.10 Statistical Functions in SQL:2006	591
Chapter 32 Subsequences, Regions, Runs, Gaps, and Islands	595
32.1 Finding Subregions of Size (n)	596
32.2 Numbering Regions	597
32.3 Finding Regions of Maximum Size	598
32.4 Bound Queries	602
32.5 Run and Sequence Queries	603
32.6 Summation of a Series	607
32.7 Swapping and Sliding Values in a List	610
32.8 Condensing a List of Numbers	612
32.9 Folding a List of Numbers	612
32.10 Coverings	613
Chapter 33 Matrices in SQL	617
33.1 Arrays via Named Columns	617
33.2 Arrays via Subscript Columns	621
33.3 Matrix Operations in SQL	622
33.4 Flattening a Table into an Array	627
33.5 Comparing Arrays in Table Format	628

Chapter 34 Set Operations	631
34.1 UNION and UNION ALL	632
34.2 INTERSECT and EXCEPT	635
34.3 A Note on ALL and SELECT DISTINCT	640
34.4 Equality and Proper Subsets	640
Chapter 35 Subsets	643
35.1 Every N-th Item in a Table	643
35.2 Random Rows from a Table	644
35.3 The CONTAINS Operators	649
35.4 Gaps in a Series	654
35.5 Covering for Overlapping Intervals	656
35.6 Picking a Representative Subset	659
Chapter 36 Trees and Hierarchies in SQL	665
36.1 Adjacency List Model	666
36.2 The Path Enumeration Model	670
36.3 Nested Set Model of Hierarchies	673
36.4 Other Models for Trees and Hierarchies	680
Chapter 37 Graphs in SQL	681
37.1 Adjacency List Model Graphs	682
37.2 Split Node Nested Set Models for Graphs	693
37.3 Points inside Polygons	698
37.4 Graph Theory References	700
Chapter 38 Temporal Queries	701
38.1 Temporal Math	701
38.2 Personal Calendars	703
38.3 Time Series	704
38.4 Julian Dates	719
38.5 Other Temporal Functions	722
38.6 Weeks	723
38.7 Modeling Time in Tables	726

38.8 Calendar Auxiliary Table	729
38.9 Problems with the Year 2000.	731
Chapter 39 Optimizing SQL	737
39.1 Access Methods.	738
39.2 How to Index	740
39.3 Give Extra Information	744
39.4 Index Multiple Columns Carefully	745
39.5 Watch the IN Predicate	746
39.6 Avoid UNIONS	748
39.7 Prefer Joins over Nested Queries.	748
39.8 Use Fewer Statements	749
39.9 Avoid Sorting	750
39.10 Avoid CROSS JOINS.	754
39.11 Know Your Optimizer	754
39.12 Recompile Static SQL after Schema Changes.	756
39.13 Temporary Tables Are Sometimes Handy	757
39.14 Update Statistics	759
39.15 Do Not Trust Newer Features	760
References	763
General References	763
Logic	763
Mathematical Techniques.	763
Random Numbers	763
Scales and Measurements.	764
Missing Values	765
Regular Expressions.	765
Graph Theory.	766
Introductory SQL Books.	766
Optimizing Queries.	767
Temporal Data and the Year 2000 Problem.	767
SQL Programming Techniques	768
Classics	768
Updatable Views	769
Theory, Normalization, and Advanced Database Topics	769

Books on SQL-92 and SQL-99	770
Standards and Related Groups	770
Web Sites Related to SQL	771
Statistics	771
Temporal Databases	771
Miscellaneous Citations	772
Index	775