



Inside Microsoft® SQL Server® 2008: T-SQL Querying

*Itzik Ben-Gan, Lubor Kollar,
Dejan Sarka, Steve Kass (Solid
Quality Mentors)*

To learn more about this book, visit Microsoft Learning at
<http://www.microsoft.com/MSPress/books/12804.aspx>

9780735626034

Microsoft®
Press

© 2009 Itzik Ben-Gan; Dejan Sarka; Lubor Kollar; Steve Kass.

Table of Contents

Foreword	xiii
Acknowledgments	xv
Introduction	xix
1 Logical Query Processing	1
Logical Query Processing Phases	2
Logical Query Processing Phases in Brief	3
Sample Query Based on Customers/Orders Scenario	5
Logical Query Processing Phase Details	7
Step 1: The FROM Phase	7
Step 2: The WHERE Phase	11
Step 3: The GROUP BY Phase	12
Step 4: The HAVING Phase	13
Step 5: The SELECT Phase	14
Step 6: The Presentation ORDER BY Phase	16
Further Aspects of Logical Query Processing	20
Table Operators	20
OVER Clause	29
Set Operators	31
Conclusion	33
2 Set Theory and Predicate Logic	35
An Example of English-to-Mathematics Translation	35
Well-Definedness	37
Equality, Identity, and Sameness	39
Mathematical Conventions	39
Numbers	41
Context	41
Functions, Parameters, and Variables	43
Instructions and Algorithms	43

 **What do you think of this book? We want to hear from you!**

Microsoft is interested in hearing your feedback so we can continually improve our books and learning resources for you. To participate in a brief online survey, please visit:

www.microsoft.com/learning/booksurvey/

Set Theory	44
Notation for Sets	45
Well-Definedness of Sets	46
Domains of Discourse	46
Faithfulness	49
Russell's Paradox	52
Ordered Pairs, Tuples, and Cartesian Products	53
The Empty Set(s)	54
The Characteristic Function of a Set	55
Cardinality	56
Order	57
Set Operators	61
Set Partitions	63
Generalizations of Set Theory	64
Predicate Logic	65
Logic-Like Features of Programming Languages	65
Propositions and Predicates	66
The Law of Excluded Middle	68
And, Or, and Not	68
Logical Equivalence	70
Logical Implication	70
Quantification	72
Alternatives and Generalizations	73
Relations	75
The Reflexive, Symmetric, and Transitive Properties	75
A Practical Application	77
Conclusion	81
3 The Relational Model	83
Introduction to the Relational Model	83
Relations, Tuples and Types	84
The Relational Model: A Quick Summary	89
Relational Algebra and Relational Calculus	90
Basic Operators	90
Relational Algebra	91
Relational Calculus	102
T-SQL Support	103
Data Integrity	104
Declarative Constraints	105
Other Means of Enforcing Integrity	109

Normalization and Other Design Topics	111
Normal Forms Dealing with Functional Dependencies	112
Higher Normal Forms	119
Denormalization	122
Generalization and Specialization	124
Conclusion	126
4 Query Tuning	127
Sample Data for This Chapter	127
Tuning Methodology	131
Analyze Waits at the Instance Level	134
Correlate Waits with Queues	143
Determine Course of Action	145
Drill Down to the Database/File Level	145
Drill Down to the Process Level	148
Tune Indexes and Queries	169
Tools for Query Tuning	171
Cached Query Execution Plans	171
Clearing the Cache	171
Dynamic Management Objects	172
STATISTICS IO	172
Measuring the Run Time of Queries	173
Analyzing Execution Plans	174
Hints	185
Traces/Profiler	186
Database Engine Tuning Advisor	187
Data Collection and Management Data Warehouse	187
Using SMO to Clone Statistics	187
Index Tuning	187
Table and Index Structures	188
Index Access Methods	197
Analysis of Indexing Strategies	244
Fragmentation	256
Partitioning	258
Preparing Sample Data	259
Data Preparation	259
TABLESAMPLE	265
An Examination of Set-Based vs. Iterative/Procedural Approaches and a Tuning Exercise	268
Conclusion	276

5 Algorithms and Complexity	277
Do You Have a Quarter?	278
How Algorithms Scale	279
An Example of Quadratic Scaling	280
An Algorithm with Linear Complexity	280
Exponential and Superexponential Complexity	281
Sublinear Complexity	282
Constant Complexity	283
Technical Definitions of Complexity	283
Comparing Complexities	285
Classic Algorithms and Algorithmic Strategies	286
Algorithms for Sorting	287
String Searching	289
A Practical Application	290
Identifying Trends in Measurement Data	291
The Algorithmic Complexity of LISLP	291
Solving the Longest Increasing Subsequence Length Problem in T-SQL	292
Conclusion	295
6 Subqueries, Table Expressions, and Ranking Functions	297
Subqueries	298
Self-Contained Subqueries	298
Correlated Subqueries	302
Misbehaving Subqueries	314
Uncommon Predicates	316
Table Expressions	318
Derived Tables	318
Common Table Expressions	321
Analytical Ranking Functions	330
Row Number	332
Rank and Dense Rank	352
Tile Number	354
Auxiliary Table of Numbers	359
Missing and Existing Ranges (Also Known as Gaps and Islands)	363
Missing Ranges (Gaps)	366
Existing Ranges (Islands)	375
Conclusion	387

7	Joins and Set Operations	389
	Joins	389
	Old Style vs. New Style	389
	Fundamental Join Types	390
	Further Examples of Joins	402
	Sliding Total of Previous Year	417
	Join Algorithms	421
	Separating Elements	429
	Set Operations	435
	UNION	436
	EXCEPT	437
	INTERSECT	439
	Precedence of Set Operations	440
	Using INTO with Set Operations	441
	Circumventing Unsupported Logical Phases	441
	Conclusion	443
8	Aggregating and Pivoting Data	445
	OVER Clause	445
	Tiebreakers	448
	Running Aggregations	451
	Cumulative Aggregations	453
	Sliding Aggregations	457
	Year-to-Date (YTD)	459
	Pivoting	460
	Pivoting Attributes	460
	Relational Division	465
	Aggregating Data	466
	Unpivoting	470
	Custom Aggregations	473
	Custom Aggregations Using Pivoting	474
	User Defined Aggregates (UDA)	476
	Specialized Solutions	487
	Histograms	499
	Grouping Factor	503
	Grouping Sets	506
	Sample Data	507
	The GROUPING SETS Subclause	508

- The CUBE Subclause 511
- The ROLLUP Subclause 512
- Grouping Sets Algebra 514
- The GROUPING_ID Function 518
- Materialize Grouping Sets 521
- Sorting 524
- Conclusion 525

- 9 TOP and APPLY 527**
 - SELECT TOP 527
 - TOP and Determinism 529
 - TOP and Input Expressions 530
 - TOP and Modifications 531
 - TOP on Steroids 534
 - APPLY 535
 - Solutions to Common Problems Using TOP and APPLY 537
 - TOP *n* for Each Group 537
 - Matching Current and Previous Occurrences 543
 - Paging 547
 - Random Rows 552
 - Median 554
 - Logical Transformations 556
 - Conclusion 559

- 10 Data Modification 561**
 - Inserting Data 561
 - Enhanced VALUES Clause 561
 - SELECT INTO 563
 - BULK Rowset Provider 565
 - Minimally Logged Operations 567
 - INSERT EXEC 590
 - Sequence Mechanisms 595
 - GUIDs 600
 - Deleting Data 601
 - TRUNCATE vs. DELETE 601
 - Removing Rows with Duplicate Data 601
 - DELETE Using Joins 603

Updating Data	606
UPDATE Using Joins.	606
Updating Large Value Types	610
SELECT and UPDATE Statement Assignments.	611
Merging Data	616
MERGE Fundamentals.	617
Adding a Predicate	621
Multiple WHEN Clauses	623
WHEN NOT MATCHED BY SOURCE	624
MERGE Values.	626
MERGE and Triggers	627
OUTPUT Clause.	628
INSERT with OUTPUT.	629
DELETE with OUTPUT	630
UPDATE with OUTPUT.	632
MERGE with OUTPUT	634
Composable DML	636
Conclusion.	638
11 Querying Partitioned Tables	639
Partitioning in SQL Server.	639
Partitioned Views.	639
Partitioned Tables	640
Conclusion.	657
12 Graphs, Trees, Hierarchies, and Recursive Queries.	659
Terminology	659
Graphs	659
Trees	660
Hierarchies.	661
Scenarios	661
Employee Organizational Chart	661
Bill of Materials (BOM)	663
Road System	666
Iteration/Recursion	670
Subordinates.	671
Ancestors.	681

Subgraph/Subtree with Path Enumeration	685
Sorting	688
Cycles	691
Materialized Path	694
Maintaining Data	695
Querying	701
Materialized Path with the HIERARCHYID Data Type	706
Maintaining Data	708
Querying	715
Further Aspects of Working with HIERARCHYID	719
Nested Sets	730
Assigning Left and Right Values	731
Querying	737
Transitive Closure	740
Directed Acyclic Graph	740
Conclusion	755
Appendix A: Logic Puzzles	757
Index	779



What do you think of this book? We want to hear from you!

Microsoft is interested in hearing your feedback so we can continually improve our books and learning resources for you. To participate in a brief online survey, please visit:

www.microsoft.com/learning/booksurvey/