

Contents

Foreword	xiii
About the Author	xv
About the Technical Reviewers	xvii
Acknowledgments	xix
Introduction	xxi
CHAPTER 1 What Do You Mean by Cost?	1
Optimizer Options	1
So What Is the Cost?	2
Transformation and Costing	5
WYSIWYG?	7
Summary	8
Test Cases	8
CHAPTER 2 Tables cans	9
Getting Started	9
Onwards and Upwards	14
Effects of Block Sizes.....	14
CPU Costing.....	16
The Power of CPU Costing.....	22
The BCHR Is Dead! Long Live the BCHR!	25
Parallel Execution	28
Index Fast Full Scan	31
Partitioning	34
Summary	39
Test Cases	40
CHAPTER 3 Single Table Selectivity	41
Getting Started	41
Null Values	44

Using Lists	45
10g Update	49
Range Predicates	50
10g Update	54
Two Predicates	55
Problems with Multiple Predicates	58
Summary	59
Test Cases	60
CHAPTER 4 Simple B-tree Access	61
Basics of Index Costing	61
Getting Started	63
Effective Index Selectivity	66
Effective Table Selectivity	67
clustering_factor.....	67
Putting It Together	69
Extending the Algorithm.....	71
The Three Selectivities.....	78
CPU Costing	81
Loose Ends	84
Summary	85
Test Cases	85
CHAPTER 5 The Clustering Factor	87
Baseline Example	87
Reducing Table Contention (Multiple Freelists).....	90
Reducing Leaf Block Contention (Reverse Key Indexes)	94
Reducing Table Contention (ASSM)	96
Reducing Contention in RAC (Freelist Groups).....	99
Column Order	101
Extra Columns	104
Correcting the Statistics	106
The sys_op_countchg() Technique.....	106
Informal Strategies	111
Loose Ends	112
Summary	112
Test Cases	113

CHAPTER 6	Selectivity Issues	115
Different Data Types	115	
Date Values	116	
Character Values.....	116	
Daft Data Types	118	
Leading Zeros	122	
Deadly Defaults	124	
Discrete Dangers	126	
10g Update	130	
Surprising sysdate	130	
Function Figures	132	
Correlated Columns	134	
Dynamic Sampling	137	
Optimizer Profiles	139	
Transitive Closure	141	
Constraint-Generated Predicates	144	
Summary	148	
Test Cases	148	
CHAPTER 7	Histograms	151
Getting Started	151	
Generic Histograms	157	
Histograms and Bind Variables	157	
When Oracle Ignores Histograms	160	
Frequency Histograms	162	
Faking Frequency Histograms	166	
Warning to Fakers	167	
“Height Balanced” Histograms	169	
The Arithmetic	172	
Data Problems Revisited	175	
Daft Datatypes	175	
Dangerous Defaults	178	
Summary	179	
Test Cases	180	

CHAPTER 8	Bitmap Indexes	181
Getting Started	181	
The Index Component	186	
The Table Component	188	
Bitmap Combinations	190	
Low Cardinality	192	
Null Columns	195	
CPU Costing	198	
Interesting Cases	200	
Multicolumn Indexes	200	
Bitmap Join Indexes	201	
Bitmap Transformations	202	
Summary	205	
Test Cases	205	
CHAPTER 9	Query Transformation	207
Getting Started	208	
Evolution	211	
Filtering	211	
Filter Optimization	215	
Scalar Subqueries	217	
Subquery Factoring	224	
Complex View Merging	230	
Pushing Predicates	232	
General Subqueries	234	
Subquery Parameters	236	
Categorization	237	
Semi-Joins	243	
Anti-Joins	246	
Anti-join Anomaly	248	
Nulls and Not In	249	
The ordered Hint	250	
Star Transformation Joins	252	
Star Joins	258	
The Future	260	
Summary	262	
Test Cases	263	

CHAPTER 10 Join Cardinality	265
Basic Join Cardinality	265
Biased Joins	269
Join Cardinality for Real SQL	271
Extensions and Anomalies	274
Joins by Range	275
Not Equal	276
Overlaps	278
Histograms	280
Transitive Closure	283
Three Tables	288
Nulls	291
Implementation Issues	295
Difficult Bits!	299
Features	301
An Alternative Viewpoint	303
Summary	304
Test Cases	305
CHAPTER 11 Nested Loops	307
Basic Mechanism	307
Worked Example	312
Sanity Checks	314
Summary	318
Test Cases	318
CHAPTER 12 Hash Joins	319
Getting Started	319
The Optimal Hash Join	323
The Onepass Hash Join	325
The Multipass Hash Join	331
Trace Files	335
Event 10104	336
Event 10053	338
Headaches	339
Traditional Costing	339
Modern Costing	340

Comparisons	341
Multitable Joins	347
Summary	350
Test Cases	350
CHAPTER 13 Sorting and Merge Joins	353
Getting Started	353
Memory Usage	359
CPU Usage	360
sort_area_retained_size	364
pga_aggregate_target	365
Real I/O	368
Cost of Sorts	370
10053 trace	371
Comparisons	375
Merge Joins	379
The Merge Mechanism	379
A Merge Join Without the First Sort	384
The Cartesian Merge Join	385
Aggregates and Others	387
Indexes	392
Set Operations	393
Final Warning	398
Summary	399
Test Cases	400
CHAPTER 14 The 10053 Trace File	403
The Query	404
The Execution Plan	405
The Environment	406
The Trace File	407
Parameter Settings	407
Query Blocks	411
Stored Statistics	412
Single Tables	414
Sanity Checks	416
General Plans	416
Join order[1]	417
Join order[2]	423

Join order[3]	424
Join order[4]	424
Join order[5]	425
Join order[6]	429
Join order[7]	429
Join order[8]	433
Join order[9]	435
Join order[10]	435
Join order[11]	436
Join order[12]	439
Join order[13]	441
Join order[14]	443
Join order[15]	444
Join order[16]	444
Join order[17]	445
Join order[18]	447
Join Evaluation Summary	449
Test Cases	451
APPENDIX A Upgrade Headaches	453
dbms_stats	453
Frequency Histograms	454
CPU Costing	455
Rounding Errors	455
Bind Variable Peeking	455
Nulls Across Joins	456
B-tree to Bitmap Conversions	456
Index Skip-Scans	456
AND-Equal	456
Index Hash Join	457
In-List Fixed	457
Transitive Closure	458
sysdate Arithmetic Fixed	458
Indexing Nulls	459
pga_aggregate_target	459
Sorting	460
Grouping	460
Sanity Checks	460
Going Outside the Limits	460
Type Hacking	460

optimizer_mode	461
Descending Indexes	461
Complex View Merging	461
Unnest Subquery	461
Scalar and Filter Subqueries	461
Parallel Query Changes x 2	462
Dynamic Sampling	462
Temporary Tables	462
Dictionary Stats	462
APPENDIX B Optimizer Parameters	465
optimizer_features_enable	465
The 10053 Trace File	467
v\$sql_optimizer_env	473
INDEX	475