Table of Contents

About the Authors ... xi

About the Technical Reviewer ... xiii

Acknowledgments ... xv

Introduction ... xvii

Chapter 1: Why Optimize? ... 1

- What Do We Mean by Optimization? ... 1
- Why It Is Difficult: Imperative and Declarative ... 2
- Optimization Goals ... 5
- Optimizing Processes ... 7
  - Optimizing OLTP and OLAP ... 8
  - Database Design and Performance ... 8
  - Application Development and Performance ... 10
  - Other Stages of the Lifecycle ... 10
- PostgreSQL Specifics ... 11
- Summary ... 12

Chapter 2: Theory: Yes, We Need It! ... 13

- Query Processing Overview ... 13
  - Compilation ... 13
  - Optimization and Execution ... 14
- Relational, Logical, and Physical Operations ... 15
  - Relational Operations ... 15
  - Logical Operations ... 19

Queries as Expressions: Thinking in Sets 20
Operations and Algorithms 20
Summary 21
Chapter 3: Even More Theory: Algorithms 23
Algorithm Cost Models 23
Data Access Algorithms 24
Storage Structures 25
Full Scan 26
Index-Based Table Access 27
Index-Only Scan 28
Comparing Data Access Algorithms 28
Index Structures 31
What Is an Index? 31
B-Tree Indexes 33
Why Are B-Trees Used So Often? 35
Bitmaps 35
Other Kinds of Indexes 36
Combining Relations 37
Nested Loops 37
Hash-Based Algorithms 39
Sort-Merge Algorithm 41
Comparing Algorithms 42
Summary 42
Chapter 4: Understanding Execution Plans 43
Putting Everything Together: How an Optimizer Builds an Execution Plan 43
Reading Execution Plans 44
Understanding Execution Plans 49
What Is Going On During Optimization? 49
Why Are There So Many Execution Plans to Choose From? 50

How Are Execution Costs Calculated? 51
How Can the Optimizer Be Led Astray? 54
Summary 55
Chapter 5: Short Queries and Indexes 57
Which Queries Are Considered Short? 57
Choosing Selection Criteria 59
Index Selectivity 60
Unique Indexes and Constraints 61
Indexes and Non-equal Conditions 65
Indexes and Column Transformations 65
Indexes and the *like* Operator 70
Using Multiple Indexes 72
Compound Indexes 73
How Do Compound Indexes Work? 74
Lower Selectivity 76
Using Indexes for Data Retrieval 76
Covering Indexes 77
Excessive Selection Criteria 79
Partial Indexes 83
Indexes and Join Order 85
When Are Indexes Not Used 88
Avoiding Index Usage 89
Why Does PostgreSQL Ignore My Index? 89
Let PostgreSQL Do Its Job! 92
How to Build the Right Index(es)? 97
To Build or Not to Build 97
Which Indexes Are Needed? 98
Which Indexes Are Not Needed? 98
Indexes and Short Query Scalability 99
Summary 100

Chapter 6: Long Queries and Full Scans 101
Which Queries Are Considered Long? 101
Long Queries and Full Scans 103
Long Queries and Hash Joins 104
Long Queries and the Order of Joins 105
What Is a Semi-join? 105
Semi-joins and Join Order 107
More on Join Order 109
What Is an Anti-join? 112
Semi- and Anti-joins Using the JOIN Operator 113
When Is It Necessary to Specify Join Order? 116
Grouping: Filter First, Group Last 118
Grouping: Group First, Select Last 125
Using SET operations 128
Avoiding Multiple Scans 132
Conclusion 138
Chapter 7: Long Queries: Additional Techniques 139
Structuring Queries 139
Temporary Tables and CTEs 140
Temporary Tables 140
Common Table Expressions (CTEs) 142
Views: To Use or Not to Use 147
Why Use Views? 154
Materialized Views 154
Creating and Using Materialized Views 154
Refreshing Materialized Views 156
Create a Materialized View or Not? 156
Do Materialized Views Need to Be Optimized? 158
Dependencies 159
Partitioning 160

Parallelism 165
Summary 166
Chapter 8: Optimizing Data Modification 167
What Is DML? 167
Two Ways to Optimize Data Modification 167
How Does DML Work? 168
Low-Level Input/Output 168
The Impact of Concurrency Control 169
Data Modification and Indexes 172
Mass Updates and Frequent Updates 173
Referential Integrity and Triggers 174
Summary 175
Chapter 9: Design Matters 177
Design Matters 177
Why Use a Relational Model? 181
Types of Databases 182
Entity-Attribute-Value Model 182
Key-Value Model 183
Hierarchical Model 184
Combining the Best of Different Worlds 185
Flexibility vs. Efficiency and Correctness 185
Must We Normalize? 187
Use and Misuse of Surrogate Keys 189
Summary 195
Chapter 10: Application Development and Performance 197
Response Time Matters 197
World Wide Wait 198
Performance Metrics 199
Impedance Mismatch 200

The Road Paved with Good Intentions 200
Application Development Patterns 201
"Shopping List Problem" 203
Interfaces 205
Welcome to the World of ORM 205
In Search of a Better Solution 207
Summary 210
Chapter 11: Functions 211
Function Creation 211
Internal Functions 212
User-Defined Functions 212
Introducing Procedural Language 213
Dollar Quoting 214
Function Parameters and Function Output: Void Functions 215
Function Overloading 216
Function Execution 218
Function Execution Internals 220
Functions and Performance 223
How Using Functions Can Worsen Performance 224
Any Chance Functions Can Improve Performance? 226
Functions and User-Defined Types 226
User-Defined Data Types 226
Functions Returning Composite Types 227
Using Composite Types with Nested Structure 231
Functions and Type Dependencies 235
Data Manipulation with Functions 236
Functions and Security 238
What About Business Logic? 239
Functions in OLAP Systems 240

Parameterizing 240

No Explicit Dependency on Tables and Views 241

Ability to Execute Dynamic SQL 241

Stored Procedures 241

Functions with No Results 241

Functions and Stored Procedures 242

Transaction Management 242

Exception Processing 243

Summary 244

Chapter 12: Dynamic SQL 245

What Is Dynamic SQL 245

Why It Works Better in Postgres 245

What About SQL Injection? 246

How to Use Dynamic SQL in OLTP Systems 246

How to Use Dynamic SQL in OLAP Systems 252

Using Dynamic SQL for Flexibility 256

Using Dynamic SQL to Aid the Optimizer 263

FDWs and Dynamic SQL 266

Summary 267

Chapter 13: Avoiding the Pitfalls of Object-Relational Mapping 269

Why Application Developers Like NORM 269

ORM vs. NORM 270

NORM Explained 272

Implementation Details 278

Complex Searches 283

Updates 286

Insert 287

Update 288

Delete 290

Why Not Store JSON?! 291
Performance Gains 291
Working Together with Application Developers 292
Summary 292
Chapter 14: More Complex Filtering and Search 293
Full Text Search 293
Multidimensional and Spatial Search 295
Generalized Index Types in PostgreSQL 295
GIST Indexes 296
Indexes for Full Text Search 296
Indexing Very Large Tables 297
Indexing JSON and JSONB 298
Summary 302
Chapter 15: Ultimate Optimization Algorithm 303
Major Steps 303
Step-by-Step Guide 304
Step 1: Short or Long? 304
Step 2: Short 304
Step 3: Long 306
Step 4: Incremental Updates 306
Step 5: Non-incremental Long Query 306
But Wait, There Is More! 307
Summary 308
Chapter 16: Conclusion 309
Index 311