

Contents

1	Introduction	1
1.1	Problem Statement	6
1.2	Contribution and Scope	8
1.3	Thesis Organization	11
Part I Background of Transactional and Analytical Systems in Logical Database Design and Benchmarking		
2	Enterprise Data Management for Transaction and Analytical Processing	15
2.1	Data Models for Transaction and Analytical Processing	17
2.1.1	Transaction Processing Systems	18
2.1.2	Analytical Processing Systems	28
2.2	Relational Database Design	35
2.2.1	Relational Database Schemas in Transaction and Analytical Processing	36
2.2.2	Normalization	38
2.2.3	Physical Database Design	39
2.2.4	Database Storage	41
2.3	Summary	43
3	Benchmarks for Transaction and Analytical Processing Systems	45
3.1	Transaction Processing Versus Analytical Processing	46
3.1.1	OLTP and OLAP Workload Characteristics	46
3.1.2	Blurring the Border Between OLTP and OLAP	47
3.2	Benchmark Classification	49
3.2.1	Transaction Processing System Benchmarks	50
3.2.2	Analytical Processing System Benchmarks	52
3.2.3	Mixed Workload Benchmarking	54
3.2.4	Other Database Benchmarks	56

- 3.3 Key Criteria for the Value of Benchmarks 58
 - 3.3.1 Established Benchmarks and the Benchmark Properties 59
 - 3.3.2 Benchmark Measures 60
- 3.4 Summary 61

Part II Towards a Benchmark for Mixed Workloads and Its Application in Evaluating Database Schemas

- 4 Combined Transaction Processing and Reporting Benchmark 65**
 - 4.1 Creation of a Hybrid Benchmark 66
 - 4.2 The Benchmark Scenario 67
 - 4.2.1 The Order-to-Cash Process 69
 - 4.2.2 Conceptual Data Model and Database Schema 70
 - 4.3 The Benchmark Queries 74
 - 4.3.1 Transactions 74
 - 4.3.2 Analytical Queries 79
 - 4.3.3 CBTR Query Shares and Database Access 83
 - 4.4 CBTR Measures and Parameters 84
 - 4.4.1 The Throughput and Response Time Measures 85
 - 4.4.2 Scaling 85
 - 4.4.3 Workload Mix 87
 - 4.5 CBTR and the Benchmark Properties 89
 - 4.6 Summary 91
- 5 Database Schema Variants for Mixed OLTP and OLAP 93**
 - 5.1 Database Design Variation Levels 94
 - 5.2 Database Schema Variants 96
 - 5.2.1 Document-Oriented Schema 97
 - 5.2.2 Snowflake Schema 99
 - 5.2.3 Third Normal Form Schema Variant 102
 - 5.3 Summary 107

Part III Implementation, Evaluation, and Discussion

- 6 The CBTR Tool Chain 111**
 - 6.1 The Benchmark Run 111
 - 6.2 Visualization of Results 115
 - 6.3 Limitations and Opportunities 118
 - 6.4 Summary 119
- 7 Evaluation of Mixing the Workload and Variation of the Database Schema 121**
 - 7.1 General Test Setup 121
 - 7.2 Impact of Adding OLAP to OLTP 122
 - 7.3 Impact of Database Schema Variation 124

- 7.4 Database Schemas Under Varying Workload Mixes 129
- 7.5 Summary 131
- 8 Conclusion** 133
 - 8.1 Discussion 135
 - 8.2 Future Work 137

- A Related Activities and Publications** 139

- B Implementation** 141

- C Evaluation Results** 145

- Bibliography** 153



<http://www.springer.com/978-3-642-38069-3>

Benchmarking Transaction and Analytical Processing
Systems

The Creation of a Mixed Workload Benchmark and its
Application

Bog, A.

2014, XIII, 164 p., Hardcover

ISBN: 978-3-642-38069-3