

Contents

- 1 Introduction 1**
 - 1.1 Goals of the Lecture 1
 - 1.2 The Idea 2
 - 1.3 Learning Map 3
 - 1.4 Self Test Questions 3
 - References 4
- Part I The Future of Enterprise Computing**
- 2 New Requirements for Enterprise Computing 7**
 - 2.1 Processing of Event Data 7
 - 2.1.1 Sensor Data 8
 - 2.1.2 Analysis of Game Events 9
 - 2.2 Combination of Structured and Unstructured Data 9
 - 2.2.1 Patient Data 10
 - 2.2.2 Airplane Maintenance Reports 10
 - 2.3 Social Networks and the Web 11
 - 2.4 Operating Cloud Environments 11
 - 2.5 Mobile Applications 12
 - 2.6 Production and Distribution Planning 12
 - 2.6.1 Production Planning 13
 - 2.6.2 Available-to-Promise Check 13
 - 2.7 Mathematical and Scientific Applications 13
 - 2.7.1 Proteomics Research 14
 - 2.7.2 Graph Data Processing 14
 - 2.7.3 SanssouciDB Data Scientist 15
 - 2.8 Self Test Questions 16
 - References 16
- 3 Enterprise Application Characteristics 17**
 - 3.1 Enterprise Data Sources 17
 - 3.2 OLTP vs. OLAP 17

3.3	Drawbacks of the Separation of OLAP from OLTP	18
3.4	The OLTP vs. OLAP Access Pattern Myth	18
3.5	Combining OLTP and OLAP Data	19
3.6	Enterprise Data is Sparse Data.....	20
3.7	Self Test Questions	20
	References	21
4	Changes in Hardware	23
4.1	Memory Cells	23
4.2	Memory Hierarchy	24
4.3	Cache Internals	25
4.4	Address Translation	26
4.5	Prefetching	27
4.6	Memory Hierarchy and Latency Numbers	27
4.7	Non-Uniform Memory Access	29
4.8	Scaling Main Memory Systems	30
4.9	Remote Direct Memory Access	30
4.10	Self Test Questions	31
	References	32
5	A Blueprint of SanssouciDB	33
5.1	Data Storage in Main Memory	33
5.2	Column-Orientation	34
5.3	Implications of Column-Orientation	34
5.4	Architecture Overview	35
	References	36
 Part II Foundations of Database Storage Techniques		
6	Dictionary Encoding	39
6.1	Compression Example	40
6.1.1	Dictionary Encoding Example: First Names	40
6.1.2	Dictionary Encoding Example: Gender	41
6.2	Sorted Dictionaries	42
6.3	Operations on Encoded Values	42
6.4	Self Test Questions	43
7	Compression	45
7.1	Prefix Encoding	45
7.1.1	Example	46
7.2	Run-Length Encoding.....	47
7.2.1	Example	47
7.3	Cluster Encoding	48
7.3.1	Example	48
7.4	Indirect Encoding	50
7.4.1	Example	51

7.5	Delta Encoding	53
7.5.1	Example	53
7.6	Limitations	54
7.7	Self Test Questions	54
	References	56
8	Data Layout in Main Memory	57
8.1	Cache Effects on Application Performance	57
8.1.1	The Stride Experiment	58
8.1.2	The Size Experiment	60
8.2	Row and Columnar Layouts	60
8.3	Benefits of a Columnar Layout	62
8.4	Hybrid Table Layouts	62
8.5	Self Test Questions	63
	References	63
9	Partitioning	65
9.1	Definition and Classification	65
9.2	Vertical Partitioning	65
9.3	Horizontal Partitioning	66
9.4	Choosing a Suitable Partitioning Strategy	68
9.5	Self Test Questions	68
	Reference	69

Part III In-Memory Database Operators

10	Delete	73
10.1	Example of Physical Delete	73
10.2	Self Test Questions	75
	Reference	75
11	Insert	77
11.1	Example	77
11.1.1	Inserting without New Dictionary Entry	78
11.1.2	Inserting with New Dictionary Entry	78
11.2	Performance Considerations	81
11.3	Self Test Questions	82
12	Update	85
12.1	Update Types	85
12.1.1	Aggregate Updates	86
12.1.2	Status Updates	86
12.1.3	Value Updates	86
12.2	Update Example	87
12.3	Self Test Questions	88
	References	89

13 Tuple Reconstruction	91
13.1 Tuple Reconstruction in Row-Oriented Databases	91
13.2 Tuple Reconstruction in Column-Oriented Databases	92
13.3 Further Examples and Discussion	93
13.4 Self Test Questions	94
14 Scan Performance	97
14.1 Introduction	97
14.2 Row Layout: Full Table Scan	98
14.3 Row Layout: Stride Access	98
14.4 Columnar Layout: Full Column Scan	99
14.5 Additional Examples and Discussion	100
14.6 Self Test Questions	101
15 Select	103
15.1 Relational Algebra	103
15.1.1 Cartesian product	103
15.1.2 Projection	103
15.1.3 Selection and Selectivity	104
15.2 Data Retrieval	105
15.3 Self Test Questions	107
16 Materialization Strategies	109
16.1 Aspects of Materialization	109
16.2 Example	110
16.3 Early Materialization	111
16.4 Late Materialization	112
16.5 Self Test Questions	115
References	116
17 Parallel Data Processing	117
17.1 Hardware Layer	117
17.1.1 Multi-Core CPUs	118
17.1.2 Single Instruction Multiple Data	119
17.2 Software Layer	121
17.2.1 Amdahl's Law	121
17.2.2 Shared Memory	122
17.2.3 Message Passing	122
17.2.4 MapReduce	123
17.3 Self Test Questions	123
References	124
18 Indices	125
18.1 Indices: A Query Optimization Approach	125
18.2 Technical Considerations	126
18.3 Inverted Index	127

18.4	Discussion	129
18.4.1	Memory Consumption	130
18.4.2	Lookup Performance	131
18.5	Self Test Questions	133
	Reference	133
19	Join	135
19.1	Join Execution in Main Memory	136
19.2	Hash Join	137
19.3	Example Hash Join	138
19.4	Sort-Merge Join	140
19.5	Example Sort-Merge Join	140
19.6	Choosing a Join Algorithm	142
19.7	Self Test Questions	142
	References	144
20	Aggregate Functions	145
20.1	Aggregation Example Using the COUNT Function	145
20.2	Other Aggregate Functions	147
20.3	Self Test Questions	147
21	Parallel Select	149
21.1	Context	149
21.2	Parallelization Example	149
21.3	Self Test Questions	152
22	Workload Management and Scheduling	153
22.1	The Importance of Speed	153
22.2	Scheduling	154
22.3	Mixed Workload Management	155
22.4	Self Test Questions	155
	Reference	156
23	Parallel Join	157
23.1	Partially Parallelized Hash Join	157
23.2	Parallel Hash Join	158
23.3	Self Test Questions	159
	References	159
24	Parallel Aggregation	161
24.1	Aggregate Functions Revisited	161
24.2	Parallel Aggregation Using Hashing	162
24.3	Self Test Questions	164
	Reference	164

Part IV Advanced Database Storage Techniques

25	Differential Buffer	167
25.1	The Concept	167
25.2	The Implementation	167
25.3	Tuple Lifetime	169
25.4	Self Test Questions	170
	References	171
26	Insert-Only	173
26.1	Definition of the Insert-Only Approach	173
26.2	Point Representation	174
26.3	Interval Representation	175
26.4	Concurrency Control: Snapshot Isolation	177
26.5	Insert-Only: Advantages and Challenges	178
26.6	Self Test Questions	179
	Reference	180
27	The Merge Process	181
27.1	The Asynchronous Online Merge	182
27.1.1	Prepare Merge Phase	183
27.1.2	Attribute Merge Phase	183
27.1.3	Commit Merge Phase	184
27.2	Exemplary Attribute Merge of a Column	184
27.3	Merge Optimizations	187
27.3.1	Using the Main Store's Dictionary	187
27.3.2	Single Column Merge	188
27.3.3	Unified Table Concept	188
27.4	Self Test Questions	189
	References	190
28	Aggregate Cache	191
28.1	Architecture	192
28.2	Cache Management	194
28.3	Updates and Deletes	194
28.4	Joins	195
28.5	Self Test Questions	195
	References	196
29	Logging	197
29.1	Logging Infrastructure	197
29.2	Logical vs. Dictionary-Encoded Logging	199
29.3	Example	202
29.4	Self Test Questions	204
	References	205

30 Recovery	207
30.1 Reading Meta Data	207
30.2 Recovering the Database	208
30.3 Self Test Questions	208
Reference	209
31 Database Views	211
31.1 Advantages of Views	211
31.2 Layered Views Concept	212
31.3 Self Test Questions	213
References	214
32 On-the-Fly Database Reorganization	215
32.1 Reorganization in a Row Store	215
32.2 On-the-Fly Reorganization in a Column Store	216
32.3 Excursion: Multi-Tenancy Requires Online Reorganization	217
32.4 Hot and Cold Data	218
32.4.1 Data Classification and Access	219
32.4.2 Application-Driven Classification	219
32.4.3 Database-Driven Classification	220
32.4.4 Application Transparent Classification	220
32.5 Self Test Questions	222
References	223

Part V Principles for Enterprise Application Development

33 Implications on Application Development	227
33.1 Optimizing Application Development for In-Memory Databases	227
33.1.1 Moving Business Logic into the Database	229
33.1.2 Stored Procedures	230
33.1.3 Example Application	231
33.2 Best Practices	232
33.3 Self Test Questions	233
34 Handling Business Objects	235
34.1 Object-Relational Mapping	237
34.2 Self Test Questions	237
References	238
35 Examples for New Enterprise Applications	239
35.1 Interactive Analyses of Point-of-Sales Data	239
35.1.1 Motivation	239
35.1.2 POS Explorer	240
35.2 High Performance In-Memory Genome Data Analysis	244
35.2.1 Application Layer	246
35.2.2 Platform Layer	249

35.2.3	Data Layer	250
35.2.4	Summary	251
	References	251
36	Bypass Solution	253
36.1	Transition Steps in Detail	254
36.2	Bypass Solution: Conclusion	257
36.3	Self Test Questions	258
	Self Test Solutions	259
	Glossary	305
	Index	311

A Course in In-Memory Data Management
The Inner Mechanics of In-Memory Databases
Plattner, H.
2014, XXIII, 312 p. 125 illus., 62 illus. in color.,
Hardcover
ISBN: 978-3-642-55269-4