

Contents

1	Introduction	1
1.1	Goals of the Lecture	1
1.2	The Idea	1
1.3	Learning Map	2
1.4	Self Test Questions	3
	References	3
 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	8
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	Self Test Questions	13
	References	14
3	Enterprise Application Characteristics	15
3.1	Diverse Applications	15
3.2	OLTP Versus OLAP	15
3.3	Drawbacks of the Separation of OLAP from OLTP	16
3.4	The OLTP Versus OLAP Access Pattern Myth	16
3.5	Combining OLTP and OLAP Data	17
3.6	Enterprise Data Characteristics	17

3.7	Self Test Questions	18
	References	18
4	Changes in Hardware	19
4.1	Memory Cells	19
4.2	Memory Hierarchy	20
4.3	Cache Internals	21
4.4	Address Translation	22
4.5	Prefetching	23
4.6	Memory Hierarchy and Latency Numbers	23
4.7	Non-Uniform Memory Architecture	25
4.8	Scaling Main Memory Systems	26
4.9	Remote Direct Memory Access	27
4.10	Self Test Questions	27
	References	28
5	A Blueprint of SanssouciDB	29
5.1	Data Storage in Main Memory	29
5.2	Column-Orientation	29
5.3	Implications of Column-Orientation	30
5.4	Active and Passive Data	31
5.5	Architecture Overview	31
5.6	Self Test Questions	32
	Reference	33

Part II Foundations of Database Storage Techniques

6	Dictionary Encoding	37
6.1	Compression Example	38
6.1.1	Dictionary Encoding Example: First Names	39
6.1.2	Dictionary Encoding Example: Gender	39
6.2	Sorted Dictionaries	40
6.3	Operations on Encoded Values	40
6.4	Self Test Questions	41
7	Compression	43
7.1	Prefix Encoding	43
7.2	Run-Length Encoding	45
7.3	Cluster Encoding	46
7.4	Indirect Encoding	48
7.5	Delta Encoding	51
7.6	Limitations	52
7.7	Self Test Questions	52
	Reference	54

8	Data Layout in Main Memory	55
8.1	Cache Effects on Application Performance	55
8.1.1	The Stride Experiment	55
8.1.2	The Size Experiment.	57
8.2	Row and Columnar Layouts.	58
8.3	Benefits of a Columnar Layout	61
8.4	Hybrid Table Layouts	61
8.5	Self Test Questions.	62
	References	62
9	Partitioning	63
9.1	Definition and Classification	63
9.2	Vertical Partitioning	63
9.3	Horizontal Partitioning	64
9.4	Choosing a Suitable Partitioning Strategy	66
9.5	Self Test Questions.	66
	Reference	67
 Part III In-Memory Database Operators		
10	Delete	71
10.1	Example of Physical Delete.	71
10.2	Self Test Questions.	73
	Reference	73
11	Insert.	75
11.1	Example	75
11.1.1	INSERT without New Dictionary Entry	76
11.1.2	INSERT with New Dictionary Entry.	76
11.2	Performance Considerations.	79
11.3	Self Test Questions.	80
12	Update.	83
12.1	Update Types.	83
12.1.1	Aggregate Updates	83
12.1.2	Status Updates	84
12.1.3	Value Updates	84
12.2	Update Example.	84
12.3	Self Test Questions.	86
	References	87

13	Tuple Reconstruction	89
13.1	Introduction	89
13.2	Tuple Reconstruction in Row-Oriented Databases	89
13.3	Tuple Reconstruction in Column-Oriented Databases	90
13.4	Further Examples and Discussion	91
13.5	Self Test Questions	92
14	Scan Performance	95
14.1	Introduction	95
14.2	Row Layout: Full Table Scan	96
14.3	Row Layout: Stride Access	96
14.4	Columnar Layout: Full Column Scan	97
14.5	Additional Examples and Discussion	98
14.6	Self Test Questions	98
15	Select	99
15.1	Relational Algebra	99
15.1.1	Cartesian Product	99
15.1.2	Projection	99
15.1.3	Selection	100
15.2	Data Retrieval	100
15.3	Self Test Questions	102
16	Materialization Strategies	105
16.1	Aspects of Materialization	105
16.2	Example	106
16.3	Early Materialization	107
16.4	Late Materialization	108
16.5	Self Test Questions	111
	References	112
17	Parallel Data Processing	113
17.1	Hardware Layer	113
17.1.1	Multi-Core CPUs	114
17.1.2	Single Instruction Multiple Data	115
17.2	Software Layer	117
17.2.1	Amdahl's Law	117
17.2.2	Shared Memory	118
17.2.3	Message Passing	118
17.2.4	MapReduce	119
17.3	Self Test Questions	120
	References	120

18	Indices	121
18.1	Indices: A Query Optimization Approach	121
18.2	Technical Considerations	121
18.3	Inverted Index	123
18.4	Discussion	126
18.4.1	Memory Consumption	126
18.4.2	Lookup Performance	127
18.5	Self Test Questions	129
	Reference	129
19	Join	131
19.1	Join Execution in Main Memory	132
19.2	Hash-Join	133
19.2.1	Example Hash-Join	133
19.3	Sort-Merge Join	135
19.3.1	Example Sort-Merge Join	135
19.4	Choosing a Join Algorithm	136
19.5	Self Test Questions	137
20	Aggregate Functions	141
20.1	Aggregation Example Using the COUNT Function	141
20.2	Self Test Questions	143
21	Parallel Select	145
21.1	Parallelization	145
21.2	Self Test Questions	148
22	Workload Management and Scheduling	149
22.1	The Power of Speed	149
22.2	Scheduling	150
22.3	Mixed Workload Management	150
22.4	Self Test Questions	151
	Reference	152
23	Parallel Join	153
23.1	Partially Parallelized Hash-Join	153
23.2	Parallel Hash-Join	154
23.3	Self Test Questions	155
	References	155
24	Parallel Aggregation	157
24.1	Aggregate Functions Revisited	157
24.2	Parallel Aggregation Using Hashing	158

24.3 Self Test Questions	160
Reference	160

Part IV Advanced Database Storage Techniques

25 Differential Buffer	163
25.1 The Concept	163
25.2 The Implementation	163
25.3 Tuple Lifetime	165
25.4 Self Test Questions	166
References	166
26 Insert-Only	167
26.1 Definition of the Insert-Only Approach	167
26.2 Point Representation	168
26.3 Interval Representation	169
26.4 Concurrency Control: Snapshot Isolation	171
26.5 Insert-Only: Advantages and Challenges	172
26.6 Self Test Questions	173
Reference	174
27 The Merge Process	175
27.1 The Asynchronous Online Merge	176
27.1.1 Prepare Merge Phase	177
27.1.2 Attribute Merge Phase	177
27.1.3 Commit Merge Phase	178
27.2 Exemplary Attribute Merge of a Column	178
27.3 Merge Optimizations	180
27.3.1 Using the Main Store's Dictionary	180
27.3.2 Single Column Merge	181
27.3.3 Unified Table Concept	182
27.4 Self Test Questions	182
References	183
28 Logging	185
28.1 Logging Infrastructure	185
28.2 Logical Versus Dictionary-Encoded Logging	187
28.3 Example	189
28.4 Self Test Questions	191
References	192

29 Recovery	193
29.1 Reading Meta Data	193
29.2 Recovering the Database	194
29.3 Self Test Questions	194
Reference	195
30 On-the-Fly Database Reorganization	197
30.1 Reorganization in a Row Store	197
30.2 On-the-Fly Reorganization in a Column Store	198
30.3 Excursion: Multi-Tenancy Requires Online Reorganization	199
30.4 Hot and Cold Data	200
30.5 Self Test Questions	201
References	203
Part V Foundations for a New Enterprise Application Development Era	
31 Implications on Application Development	207
31.1 Optimizing Application Development for In-Memory Databases	207
31.1.1 Moving Business Logic into the Database	209
31.1.2 Stored Procedures	210
31.1.3 Example Application	211
31.2 Best Practices	212
31.3 Self Test Questions	213
32 Database Views	215
32.1 Advantages of Views	215
32.2 Layered Views Concept	216
32.3 Development Tools for Views	216
32.4 Self Test Questions	217
References	218
33 Handling Business Objects	219
33.1 Persisting Business Objects	219
33.2 Object-Relational Mapping	220
33.3 Self Test Questions	221
34 Bypass Solution	223
34.1 Transition Steps in Detail	224
34.2 Bypass Solution: Conclusion	227
34.3 Self Test Questions	228

Self Test Solutions	229
Glossary	283
Index	295