

Contents

1	Transactions on the Logical Database	1
1.1	Transaction Server	1
1.2	Logical Database	2
1.3	Integrity of the Logical Database	4
1.4	Transactions	5
1.5	Transaction States	7
1.6	ACID Properties of Transactions	10
1.7	The Read-Write Model	13
1.8	The Key-Range Model	15
1.9	Savepoints and Partial Rollbacks	18
1.10	Multiple Granularity	20
	Problems	20
	Bibliographical Notes	23
2	Operations on the Physical Database	25
2.1	Data Structures and Processes on the Server	25
2.2	Database Pages and Files	27
2.3	Buffering and Fixing of Database Pages	31
2.4	Database States	33
2.5	Database Recovery and Checkpoints	34
2.6	Integrity of the Physical Database	36
2.7	Latching of Database Pages	37
2.8	Modifications on the Physical Structure	39
	Problems	42
	Bibliographical Notes	43
3	Logging and Buffering	45
3.1	Transaction Log	46
3.2	Physiological Logging	48
3.3	Active-Transaction Table and Modified-Page Table	50
3.4	Logging for the Read-Write and Key-Range Models	52
3.5	Logging Structure Modifications	55

3.6	Updating a Data Page	57
3.7	Write-Ahead Logging	58
3.8	The Commit Protocol	59
3.9	The Steal-and-No-Force Buffering Policy	60
3.10	Reducing Disk Access	61
	Problems	62
	Bibliographical Notes.....	63
4	Transaction Rollback and Restart Recovery	65
4.1	Transaction Aborts, Process Failures, and System Crashes	66
4.2	Partial and Total Rollbacks	67
4.3	Performing the Undo Actions	69
4.4	Checkpoints	72
4.5	Problems with Do-Undo-Redo Recovery	73
4.6	The ARIES Do-Redo-Undo Recovery Algorithm.....	75
4.7	Analysis Pass for Do-Redo-Undo Recovery	77
4.8	Redo Pass for Do-Redo-Undo Recovery	82
4.9	Undo Pass for Do-Redo-Undo Recovery	87
4.10	Recovery from a Process Failure	92
4.11	Recovering Structure Modifications	93
	Problems	96
	Bibliographical Notes.....	98
5	Transactional Isolation.....	101
5.1	Transaction Histories	102
5.2	Uncommitted Updates	105
5.3	Isolation Anomalies.....	106
5.4	Isolation Anomalies and Database Integrity	111
5.5	SQL Isolation Levels	115
5.6	Isolation and Transaction Rollback	116
5.7	Conflict-Serializability.....	118
5.8	Enforcing Isolation.....	120
	Problems	121
	Bibliographical Notes.....	123
6	Lock-Based Concurrency Control	125
6.1	Locks and the Lock Table	126
6.2	Acquiring and Releasing Locks	128
6.3	Locking Protocol for the Read-Write Model	131
6.4	Key-Range Locking.....	135
6.5	Deadlocks	144
6.6	Conditional Lock Requests	147
6.7	Algorithms for Read, Insert, and Delete Actions	148
6.8	Predicate Locking.....	153
	Problems	155
	Bibliographical Notes.....	157

7	B-Tree Traversals	159
7.1	Sparse B-Tree Indexes	160
7.2	Saved Paths	162
7.3	Traversals for Read Actions	165
7.4	Traversals for Insert Actions	170
7.5	Traversals for Delete Actions	176
7.6	Traversals for Undo Actions	180
	Problems	181
	Bibliographical Notes	182
8	B-Tree Structure Modifications	185
8.1	Top-Down Redo-Only Modifications	186
8.2	Tree-Height Increase and Page Split	188
8.3	Tree-Height Decrease, Page Merge, and Records Redistribute	191
8.4	Sequences of Redo-Only Modifications	195
8.5	Redoing B-Tree Structure Modifications	204
8.6	Bottom-Up Structure Modifications	207
8.7	B-Link Trees	212
8.8	Loading of a B-Tree	213
	Problems	215
	Bibliographical Notes	217
9	Advanced Locking Protocols	221
9.1	Key-Range Locking for Multiple Granules	222
9.2	Intention Locks and Multi-Granular Locking	223
9.3	Logical Locking vs. Physical Locking	227
9.4	Update-Mode Locks	228
9.5	Cursor Stability	230
9.6	Reads Without Locks	231
9.7	Recovery and Concurrent Transactions	232
	Problems	234
	Bibliographical Notes	236
10	Bulk Operations on B-Trees	237
10.1	Transactions with Bulk Actions	238
10.2	Locking Range Partitions	241
10.3	Bulk Reading	247
10.4	Bulk Insertion	248
10.5	Bulk Deletion	250
10.6	Logging and Undoing Bulk Actions	252
10.7	Complexity of Scanning a Key Range	253
	Problems	256
	Bibliographical Notes	257
11	Online Index Construction and Maintenance	259
11.1	Secondary-Key-Based Actions	260
11.2	Secondary B-Tree Indexes	262

11.3	Index Construction Without Quiescing Updates	264
11.4	Extraction of the Index Records	266
11.5	Concurrent Updates by Transactions	267
11.6	Populating the Index	272
11.7	Restoration of Clustering	274
	Problems	276
	Bibliographical Notes	277
12	Concurrency Control by Versioning	279
12.1	Transaction-Time Databases	279
12.2	Read-Only and Update Transactions	281
12.3	Transaction-Level Read Consistency	284
12.4	Statement-Level Read Consistency	286
12.5	Snapshot Isolation	287
12.6	Enforcing the Disjoint-Write Property	289
12.7	Version Management with B-Trees	290
12.8	Update Transactions on Versioned B-Trees	292
	Problems	294
	Bibliographical Notes	296
13	Distributed Transactions	299
13.1	Distributed Databases	300
13.2	Transactions on a Distributed Database	301
13.3	Atomic Commitment	306
13.4	Two-Phase Commit Protocol	307
13.5	Recovering a Server from a Failure	311
13.6	Isolation and Concurrency Control	312
13.7	Replicated Databases	314
13.8	Remote Backup Databases	316
13.9	Transactions on Key-Value Stores	319
	Problems	321
	Bibliographical Notes	324
14	Transactions in Page-Server Systems	327
14.1	Data Server	328
14.2	Page-Server Architecture	330
14.3	Client Caching and Page-Server State	331
14.4	Update Propagation	334
14.5	Cache Consistency	335
14.6	Log Management	338
14.7	Recovery from a Client Failure	341
14.8	Recovery from a Server Failure	343
14.9	Concurrency Control	345
14.10	Shared-Disks Systems	346
	Problems	347
	Bibliographical Notes	349

15	Processing of Write-Intensive Transactions	351
15.1	Group Commit	352
15.2	Online Page Relocation	354
15.3	Write-Optimized B-Trees.....	355
15.4	Merge Trees	360
15.5	Log-Structured Databases	363
15.6	Databases on Flash Memory	364
	Problems	367
	Bibliographical Notes.....	368
	References.....	371
	Index.....	381

Transaction Processing
Management of the Logical Database and its
Underlying Physical Structure

Sippu, S.; Soisalon-Soininen, E.

2014, XV, 392 p. 64 illus., Hardcover

ISBN: 978-3-319-12291-5