

# Contents

|          |  |           |
|----------|--|-----------|
| <b>1</b> | <b>Database Systems</b>                      | <b>1</b>  |
| 1.1      | Why a Database System?                       | 1         |
| 1.2      | The Derby Database System                    | 6         |
| 1.3      | Database Engines                             | 8         |
| 1.4      | The SimpleDB Database System                 | 10        |
| 1.5      | The SimpleDB Version of SQL                  | 11        |
| 1.6      | Chapter Summary                              | 12        |
| 1.7      | Suggested Reading                            | 13        |
| 1.8      | Exercises                                    | 13        |
| <b>2</b> | <b>JDBC</b>                                  | <b>15</b> |
| 2.1      | Basic JDBC                                   | 15        |
| 2.2      | Advanced JDBC                                | 27        |
| 2.3      | Computing in Java vs. SQL                    | 41        |
| 2.4      | Chapter Summary                              | 44        |
| 2.5      | Suggested Reading                            | 46        |
| 2.6      | Exercises                                    | 46        |
| <b>3</b> | <b>Disk and File Management</b>              | <b>49</b> |
| 3.1      | Persistent Data Storage                      | 49        |
| 3.2      | The Block-Level Interface to the Disk        | 60        |
| 3.3      | The File-Level Interface to the Disk         | 61        |
| 3.4      | The Database System and the OS               | 65        |
| 3.5      | The SimpleDB File Manager                    | 66        |
| 3.6      | Chapter Summary                              | 71        |
| 3.7      | Suggested Reading                            | 75        |
| 3.8      | Exercises                                    | 75        |
| <b>4</b> | <b>Memory Management</b>                     | <b>79</b> |
| 4.1      | Two Principles of Database Memory Management | 79        |
| 4.2      | Managing Log Information                     | 81        |

|          |  |            |
|----------|--|------------|
| 4.3      | The SimpleDB Log Manager . . . . .           | 83         |
| 4.4      | Managing User Data . . . . .                 | 88         |
| 4.5      | The SimpleDB Buffer Manager . . . . .        | 93         |
| 4.6      | Chapter Summary . . . . .                    | 98         |
| 4.7      | Suggested Reading . . . . .                  | 101        |
| 4.8      | Exercises . . . . .                          | 102        |
| <b>5</b> | <b>Transaction Management . . . . .</b>      | <b>105</b> |
| 5.1      | Transactions . . . . .                       | 105        |
| 5.2      | Using Transactions in SimpleDB . . . . .     | 108        |
| 5.3      | Recovery Management . . . . .                | 110        |
| 5.4      | Concurrency Management . . . . .             | 123        |
| 5.5      | Implementing SimpleDB Transactions . . . . . | 145        |
| 5.6      | Chapter Summary . . . . .                    | 145        |
| 5.7      | Suggested Reading . . . . .                  | 150        |
| 5.8      | Exercises . . . . .                          | 151        |
| <b>6</b> | <b>Record Management . . . . .</b>           | <b>159</b> |
| 6.1      | Designing a Record Manager . . . . .         | 159        |
| 6.2      | Implementing a File of Records . . . . .     | 165        |
| 6.3      | SimpleDB Record Pages . . . . .              | 170        |
| 6.4      | SimpleDB Table Scans . . . . .               | 178        |
| 6.5      | Chapter Summary . . . . .                    | 184        |
| 6.6      | Suggested Reading . . . . .                  | 185        |
| 6.7      | Exercises . . . . .                          | 186        |
| <b>7</b> | <b>Metadata Management . . . . .</b>         | <b>189</b> |
| 7.1      | The Metadata Manager . . . . .               | 189        |
| 7.2      | Table Metadata . . . . .                     | 190        |
| 7.3      | View Metadata . . . . .                      | 193        |
| 7.4      | Statistical Metadata . . . . .               | 195        |
| 7.5      | Index Metadata . . . . .                     | 199        |
| 7.6      | Implementing the Metadata Manager . . . . .  | 205        |
| 7.7      | Chapter Summary . . . . .                    | 207        |
| 7.8      | Suggested Reading . . . . .                  | 210        |
| 7.9      | Exercises . . . . .                          | 211        |
| <b>8</b> | <b>Query Processing . . . . .</b>            | <b>213</b> |
| 8.1      | Relational Algebra . . . . .                 | 213        |
| 8.2      | Scans . . . . .                              | 217        |
| 8.3      | Update Scans . . . . .                       | 220        |
| 8.4      | Implementing Scans . . . . .                 | 221        |
| 8.5      | Pipelined Query Processing . . . . .         | 226        |
| 8.6      | Predicates . . . . .                         | 228        |
| 8.7      | Chapter Summary . . . . .                    | 229        |
| 8.8      | Suggested Reading . . . . .                  | 233        |
| 8.9      | Exercises . . . . .                          | 236        |

|           |                                      |     |
|-----------|--------------------------------------|-----|
| <b>9</b>  | <b>Parsing</b>                       | 239 |
| 9.1       | Syntax Versus Semantics              | 239 |
| 9.2       | Lexical Analysis                     | 240 |
| 9.3       | The SimpleDB Lexical Analyzer        | 241 |
| 9.4       | Grammars                             | 246 |
| 9.5       | Recursive-Descent Parsers            | 249 |
| 9.6       | Adding Actions to the Parser         | 250 |
| 9.7       | Chapter Summary                      | 260 |
| 9.8       | Suggested Reading                    | 262 |
| 9.9       | Exercises                            | 262 |
| <b>10</b> | <b>Planning</b>                      | 267 |
| 10.1      | Verification                         | 267 |
| 10.2      | The Cost of Evaluating a Query Tree  | 268 |
| 10.3      | Plans                                | 274 |
| 10.4      | Query Planning                       | 277 |
| 10.5      | Update Planning                      | 281 |
| 10.6      | The SimpleDB Planner                 | 284 |
| 10.7      | Chapter Summary                      | 288 |
| 10.8      | Suggested Reading                    | 289 |
| 10.9      | Exercises                            | 289 |
| <b>11</b> | <b>JDBC Interfaces</b>               | 295 |
| 11.1      | The SimpleDB API                     | 295 |
| 11.2      | Embedded JDBC                        | 297 |
| 11.3      | Remote Method Invocation             | 300 |
| 11.4      | Implementing the Remote Interfaces   | 305 |
| 11.5      | Implementing the JDBC Interfaces     | 306 |
| 11.6      | Chapter Summary                      | 309 |
| 11.7      | Suggested Reading                    | 309 |
| 11.8      | Exercises                            | 310 |
| <b>12</b> | <b>Indexing</b>                      | 313 |
| 12.1      | The Value of Indexing                | 313 |
| 12.2      | SimpleDB Indexes                     | 316 |
| 12.3      | Static Hash Indexes                  | 319 |
| 12.4      | Extendable Hash Indexes              | 322 |
| 12.5      | B-Tree Indexes                       | 327 |
| 12.6      | Index-Aware Operator Implementations | 345 |
| 12.7      | Index Update Planning                | 353 |
| 12.8      | Chapter Summary                      | 356 |
| 12.9      | Suggested Reading                    | 357 |
| 12.10     | Exercises                            | 358 |
| <b>13</b> | <b>Materialization and Sorting</b>   | 363 |
| 13.1      | The Value of Materialization         | 363 |
| 13.2      | Temporary Tables                     | 364 |

|           |  |            |
|-----------|--|------------|
| 13.3      | Materialization . . . . .                          | 364        |
| 13.4      | Sorting . . . . .                                  | 369        |
| 13.5      | Grouping and Aggregation . . . . .                 | 379        |
| 13.6      | Merge Joins . . . . .                              | 387        |
| 13.7      | Chapter Summary . . . . .                          | 389        |
| 13.8      | Suggested Reading . . . . .                        | 393        |
| 13.9      | Exercises . . . . .                                | 394        |
| <b>14</b> | <b>Effective Buffer Utilization . . . . .</b>      | <b>397</b> |
| 14.1      | Buffer Usage in Query Plans . . . . .              | 397        |
| 14.2      | Multibuffer Sorting . . . . .                      | 398        |
| 14.3      | Multibuffer Product . . . . .                      | 400        |
| 14.4      | Determining Buffer Allocation . . . . .            | 402        |
| 14.5      | Implementing Multibuffer Sorting . . . . .         | 403        |
| 14.6      | Implementing Multibuffer Product . . . . .         | 404        |
| 14.7      | Hash Joins . . . . .                               | 406        |
| 14.8      | Comparing the Join Algorithms . . . . .            | 412        |
| 14.9      | Chapter Summary . . . . .                          | 414        |
| 14.10     | Suggested Reading . . . . .                        | 415        |
| 14.11     | Exercises . . . . .                                | 416        |
| <b>15</b> | <b>Query Optimization . . . . .</b>                | <b>419</b> |
| 15.1      | Equivalent Query Trees . . . . .                   | 419        |
| 15.2      | The Need for Query Optimization . . . . .          | 426        |
| 15.3      | The Structure of a Query Optimizer . . . . .       | 430        |
| 15.4      | Finding the Most Promising Query Tree . . . . .    | 430        |
| 15.5      | Finding the Most Efficient Plan . . . . .          | 440        |
| 15.6      | Combining the Two Stages of Optimization . . . . . | 441        |
| 15.7      | Merging Query Blocks . . . . .                     | 449        |
| 15.8      | Chapter Summary . . . . .                          | 450        |
| 15.9      | Suggested Reading . . . . .                        | 452        |
| 15.10     | Exercises . . . . .                                | 452        |
|           | <b>Index . . . . .</b>                             | <b>455</b> |