

Databases

Understanding, using and developing

Course and exercises

Jean-Luc Hainaut

Professor at the Computer Science Faculty
Research director at the PReCISE center
University of Namur

2nd edition

Dunod



Big RDBMS vendor observing with a greedy interest an emerging technology in data management¹.

1. Allegoric reference to section 1.7.8. *Cronos devours his children. Linogravure*, B. Hainaut, 2004

Table of contents

FOREWORD	27
CHAPTER 1 • MOTIVATION AND INTRODUCTION	33
1.1 Of data and men	33
1.2 A first approach to data	34
1.3 Using a database	38
1.4 First observations	39
1.5 Database management systems	40
1.6 Challenges of today's databases	41
1.7 A short history	42
1.7.1 The collectivist model: CODASYL DBTG	43
1.7.2 The single mindset model: IMS	44
1.7.3 Birth of the relational model	44
1.7.4 A new actor: the end-user	45
1.7.5 The great model war	45
1.7.6 The <i>Fall of Icarus</i>	46
1.7.7 The relational model: the sequel	46
1.7.8 Cronos devours his children	47

PART 1
CONCEPTS OF DATABASES

CHAPTER 2 • CONCEPTS OF DATABASES	49
2.1 Introduction	49
2.2 Tables, rows and columns	50
2.3 The <i>null</i> value	51
2.4 Unique keys	52
2.5 Foreign keys	54
2.6 Schema and contents of a database	55
2.7 Graphical representation of a schema	55
2.8 A sample database	57
2.9 Note on foreign keys	58
2.10 Data update and integrity constraints	61
2.10.1 Mandatory columns	61
2.10.2 Uniqueness constraints	61
2.10.3 Referential constraints	61
2.11 Redundancies in a table	63
2.11.1 The phenomenon	64
2.11.2 Normalization through decomposition	65
2.12 Physical structures	66
2.13 Summary of the chapter	68
2.14 Further reading	69
2.15 How to practice databases?	69
2.16 Exercises	70
CHAPTER 3 • RELATIONAL MODEL AND NORMALISATION	73
3.1 The phenomenon of internal redundancy	73
3.2 The database relational model	74
3.2.1 Basic concepts	74
3.2.2 Algebraic operators	75
3.3 Unique keys and their properties	78
3.4 Functional dependencies	79
3.4.1 The phenomenon	79
3.4.2 A complex example	80
3.4.3 The AFD (Attribute and FD) graph of a relation	81
3.4.4 Properties of functional dependencies	81

3.5	Inclusion constraints and foreign keys	84
3.6	Computing the unique keys of a relation	86
3.6.1	Unique keys and functional dependencies	86
3.6.2	Relation with a single unique key	87
3.6.3	Relation with multiple unique keys	88
3.7	Decomposition of a relation	89
3.8	Normalization of a relation	91
3.8.1	Revisiting the internal redundancy problem	91
3.8.2	Origin of redundancies	92
3.8.3	Normalized relation	93
3.8.4	Normalization of a relation	93
3.8.5	Dealing with an irreducible kernel	96
3.8.6	DBMS and functional dependencies	98
3.9	Summary of the chapter	99
3.10	Further reading	100
3.11	Exercises	100
CHAPTER 4 • IMPLEMENTATION OF DATA STRUCTURES		103
4.1	Introduction	103
4.2	Architecture of a computer system	104
4.3	External memory: the magnetic disk	105
4.3.1	Geometry of a disk	105
4.3.2	Reading and writing data on a disk	107
4.3.3	Optimization: exchange buffer and read ahead	109
4.4	Organization of a storage space	111
4.5	Sequential processing of a file	115
4.5.1	Sequential file reading	116
4.5.2	Sequential file sorting	118
4.6	Primary and secondary indexes	122
4.7	Indexed sequential file organization	124
4.7.1	Structure of an indexed sequential file	124
4.7.2	Reading and modifying records in an indexed sequential file	125
4.7.3	Characteristics and performances of an indexed sequential file	127
4.7.4	Case study	134
4.8	Hash file organization	135
4.8.1	Principles of hash files	135
4.8.2	Addressing functions	136
4.8.3	Overflow management	138
4.8.4	Characteristics and performances of a hash file	139

4.8.5	Comparison of primary index techniques	142
4.9	Secondary indexes	143
4.9.1	Structure of a secondary index with pointer arrays	143
4.9.2	Performances of a secondary index with pointer arrays	144
4.9.3	Structure and performances of a bitmap index	147
4.9.4	Augmented secondary indexes	147
4.9.5	Availability of primary and secondary indexing techniques	148
4.10	Index-based access or sequential scan?	148
4.11	Clustering techniques	149
4.11.1	Clustering index	149
4.11.2	Clusters	150
4.11.3	Performances of clustering techniques	150
4.12	Summary of the chapter	151
4.13	Further reading	152
4.14	Exercises	152
 CHAPTER 5 • DATABASE MANAGEMENT SYSTEMS		155
5.1	Introduction	155
5.2	Functions of database management systems	156
5.3	Architecture of a DBMS	158
5.4	Transactions	160
5.5	Concurrency control	162
5.5.1	The lost update scenario	163
5.5.2	Lock-based concurrency control	163
5.5.3	Serialisability and 2-phase locking	164
5.5.4	Deadlock	165
5.6	Protection against failures	165
5.6.1	Backup and journals	165
5.6.2	Cold recovery	166
5.6.3	Warm recovery	166
5.7	Access control	167
5.7.1	Access control models	167
5.7.2	Private data protection - Statistical inference	168
5.8	The DBMS market	170
5.9	Further reading	171

PART 2
USING DATABASES

CHAPTER 6 • THE SQL-DDL LANGUAGE	173
6.1 Introduction	173
6.2 SQL DDL	174
6.3 Creating a schema	174
6.4 Creating a table	175
6.4.1 Columns and column types	175
6.4.2 Unique keys	177
6.4.3 Foreign keys	178
6.4.4 Mandatory/optional columns	178
6.4.5 A complete example	178
6.4.6 Synthetic form of constraints	179
6.5 Deleting a table	179
6.6 Adding, removing and modifying a column	180
6.7 Adding and discarding a constraint	180
6.8 Physical structures	181
6.9 Summary of the chapter	182
6.10 Further reading	182
CHAPTER 7 • THE SQL-DML LANGUAGE (1)	183
7.1 Introduction	183
7.2 Elementary queries	185
7.2.1 Simple extraction	185
7.2.2 Selecting rows	186
7.2.3 Reducing duplicated rows	187
7.2.4 More complex conditions	188
7.3 Extracted data and derived data	190
7.3.1 SQL functions	191
7.3.2 System registers	193
7.4 Statistical functions	194
7.5 Sub-queries	196
7.5.1 Conditions based on sub-queries	196
7.5.2 Sub-queries and compound foreign keys	197
7.5.3 Beware of negative association conditions	197
7.5.4 Multiple references to a table	200
7.6 Set quantifiers	202

7.7	Exercises	204
CHAPTER 8 • THE SQL-DML LANGUAGE (2)		207
8.1	Introduction	207
8.2	Data extraction from several tables	208
8.2.1	Joining several tables	208
8.2.2	Join conditions and selection conditions	209
8.2.3	Join and non-matching rows	210
8.3	Set operators	211
8.4	The relational product	212
8.5	Queries on cyclic data structures	213
8.6	Complements on joins	217
8.6.1	Subquery or join?	218
8.6.2	Derived data in a join	220
8.6.3	Generalized joins	220
8.6.4	Interpretation of the result of a join	221
8.7	Grouped data extraction	224
8.7.1	Notion of grouped data	224
8.7.2	Group selection and row selection	226
8.7.3	Groups and joins	227
8.7.4	Grouping criterion	228
8.7.5	Beware of multi-level grouping	229
8.7.6	Can grouped data be avoided?	230
8.8	Row order in a result set	230
8.9	Interpretation of a query	232
8.10	Modifying data	232
8.10.1	Row insertion	233
8.10.2	Row deletion	234
8.10.3	Row update	234
8.10.4	Data update and referential constraint	235
8.10.5	Determinism of modification queries	237
8.11	Modifying data structures	240
8.12	Summary of the chapter	241
8.13	Further reading	241
8.14	Exercises	241

CHAPTER 9 • ADVANCED SQL	245
9.1 Introduction	245
9.2 Access control	246
9.3 SQL views	248
9.3.1 Principles and objectives of views	248
9.3.2 Defining and using a view	248
9.3.3 Views as interface for specific user requirements	249
9.3.4 Views as access control support	250
9.3.5 Views as database evolution support	250
9.3.6 Views as a help to complex queries writing	250
9.3.7 Data modification through views	251
9.4 Extension of SFW queries	252
9.4.1 Extension of the select clause	252
9.4.2 Extension of the from clause	252
9.5 Recursive queries	258
9.6 Object extension of SQL3	260
9.6.1 Complex data types (<i>row</i> and <i>array</i>)	261
9.6.2 User-defined type (UDT)	262
9.6.3 Typed table	262
9.6.4 Types hierarchy	262
9.6.5 Typed tables hierarchy	264
9.6.6 Referencing rows	265
9.6.7 Final remarks	266
9.7 Predicates (check)	267
9.8 SQL procedures (stored procedures)	268
9.9 Triggers	269
9.10 The catalog	274
9.11 SQL interfaces for application programs	277
9.11.1 Classical static SQL interface	278
9.11.2 SQLJ: a static interface for Java	282
9.11.3 Classical dynamic SQL interface	284
9.11.4 JDBC: a dynamic CLI interface for Java	287
9.11.5 Comparison of interaction models	289
9.11.6 A security problem: SQL code injection	290
9.12 SQL and incomplete information	292
9.12.1 Introduction	292
9.12.2 the SQL <code>null</code> value	292
9.12.3 The 3-value logic of SQL	292
9.12.4 Propagation of <code>null</code> in SQL	293

9.12.5	Propagation of <code>unknown</code> in SQL	293
9.12.6	Incomplete information problems in SQL	295
9.12.7	Two recommendations	299
9.13	Summary of the chapter	300
9.14	Further reading	301

PART 3

DEVELOPING DATABASES

CHAPTER 10 • BUILDING A DATABASE	303	
10.1	From craft to method	303
10.2	Roles in databases	304
10.3	Two design paths	305
10.4	A simple approach	305
10.5	Tools for database development	306
CHAPTER 11 • THE BASIC ENTITY-RELATIONSHIP MODEL	307	
11.1	Introduction	307
11.2	Entity types	308
11.3	Attributes	309
11.4	Relationship types	310
11.4.1	Roles of a relationship type	310
11.4.2	Properties of a relationship type	311
11.4.3	Functional class of a relationship type	311
11.4.4	Optional and mandatory roles	314
11.4.5	Cardinality constraints of a relationship type	314
11.4.6	Cyclic relationship types	314
11.4.7	Graphical representation of instances	316
11.5	A matter of vocabulary	316
11.6	Identifiers	317
11.6.1	Hybrid identifiers	317
11.6.2	Composition of identifiers	319
11.6.3	Minimal and implicit identifiers	320
11.6.4	Importance of the concept of identifier	321
11.7	Integrity constraints	322
11.7.1	Basic constraints	322
11.7.2	Additional constraints	322
11.8	Information contents of a schema	322

11.9	Examples	324
11.9.1	An administrative structure	324
11.9.2	Library management	325
11.9.3	Travel by train	326
11.10	Summary of the chapter	327
11.11	Further reading	328
11.12	Drawing a conceptual schema with a tool	328
11.13	Exercises	330
CHAPTER 12 • DESIGNING A CONCEPTUAL SCHEMA		333
12.1	Introduction	333
12.2	A first example	335
12.3	Decomposition of the problem statement	336
12.3.1	Notion of elementary proposition	336
12.3.2	Cardinalities	337
12.3.3	General and particular propositions	338
12.3.4	Beware of irreducible complex proposition!	339
12.3.5	N-ary reducible propositions	341
12.4	Relevance of a proposition	342
12.5	Representation of a proposition	343
12.5.1	New entity type and its first attribute	343
12.5.2	New relationship type between two existing entity types	343
12.5.3	Attribute of an existing entity type	344
12.5.4	New entity type and its first two attribute	344
12.5.5	Restructuring to make integration possible	344
12.5.6	Multivalued properties	348
12.5.7	Practical recommendation	351
12.6	Redundancy of propositions	352
12.6.1	Explicit redundancy	352
12.6.2	Expression variants	352
12.6.3	Indirect redundancy	352
12.6.4	Apparent redundancy	353
12.6.5	Synonyms et homonyms	354
12.7	Non-contradiction of propositions	354
12.8	Integrity constraints	356
12.9	Schema normalization	356
12.9.1	Schema simplification	356
12.9.2	Reducing internal redundancies	357
12.10	Schema documentation	359

12.11	Schema completeness	359
12.12	Summary of the chapter	360
12.13	Further reading	361
12.14	Exercises	361
CHAPTER 13 • PRODUCING A DATABASE SCHEMA		365
13.1	Introduction	365
13.2	Representation of entity types and attributes	366
13.3	Representation of relationship types	366
13.3.1	One-to-many relationship types	366
13.3.2	One-to-one relationship types	368
13.3.3	Many-to-many relationship types	370
13.3.4	Cyclic relationship types	371
13.4	Representation of identifiers	371
13.5	Translation of names	372
13.6	Summary of the representation rules	373
13.7	Physical structures	374
13.8	Coding data structures in SQL	376
13.9	Summary of the chapter	377
13.10	Further reading	378
13.11	Producing a database schema with a tool	378
13.12	Exercises	379
CHAPTER 14 • DATABASE METHODOLOGY		381
14.1	Introduction	381
14.2	Methods, processes et products	382
14.3	Database development methods	383
14.4	Development models	388
14.5	Transformation techniques	390
14.6	Tool-supported database methodologies	390
CHAPTER 15 • THE EXTENDED ENTITY-RELATIONSHIP MODEL		391
15.1	Introduction	391
15.2	The basic model - Reminder	392
15.3	Semantics	393
15.4	Model, schema and population	394

15.5	Stereotypes	396
15.6	Database schema	397
15.7	Attributes	397
15.7.1	Compound attributes	397
15.7.2	Multivalued attributes	398
15.7.3	Complex attributes	398
15.7.4	Attribute types and domains	399
15.7.5	Attribute cardinality	400
15.8	Relationship types	401
15.8.1	Functional relationship types	401
15.8.2	N-ary relationship types	401
15.8.3	Relationship types with multi-ET roles	402
15.8.4	Relationship type attributes	403
15.8.5	Relationship types of relationship types	403
15.8.6	Generalizing role cardinality constraints	404
15.8.7	Computing the size of relationship types	404
15.9	Entity types	405
15.9.1	Generalization-specialization relation (<i>is-a</i> relation)	405
15.9.2	Properties and set constraints of subtypes	406
15.9.3	Inheritance	407
15.9.4	Multiple supertypes and inheritance	409
15.9.5	Constraints on inherited components	411
15.9.6	Multiple distributions	412
15.9.7	Intrinsic and rule-based subtype distribution	413
15.9.8	Entity types and Java classes	414
15.9.9	On the bad usage of <i>is-a</i> relations	415
15.9.10	<i>Is-a</i> relations between other components of a schema	416
15.10	Identifiers	416
15.10.1	Entity type identifiers	416
15.10.2	Entity type cyclic identifiers	418
15.10.3	Entity type global identifiers	419
15.10.4	Relationship type identifiers	420
15.10.5	Complex attribute identifiers	422
15.11	Integrity constraints	423
15.11.1	Domain constraints	424
15.11.2	Inter-values constraints	424
15.11.3	Existence constraints	424
15.11.4	Dependency constraints	426
15.12	Additional integrity constraints	427
15.12.1	Expressing general constraints	427

15.12.2	Remarks on constraint languages	428
15.12.3	Cyclic constraints	429
15.12.4	Redundancy constraints	431
15.13	Names	432
15.13.1	Naming rules	432
15.13.2	Objects with identical names, with no name and stéréotypes	432
15.13.3	Attribute names	433
15.13.4	Relationship type names	434
15.14	Generic object classes	435
15.14.1	Generic entity-types	435
15.14.2	Generic relationship types	436
15.15	Operations and methods	439
15.16	Example of Entity-relationship schema	440
15.17	Metadata and metaschemas	440
15.17.1	A first approach: the database catalog	441
15.17.2	Metaschema of the basic Entity-relationship model	442
15.18	Semantics of the Entity-relationship model	443
15.18.1	Entity types and <i>is-a</i> relations	444
15.18.2	Relationship types	444
15.18.3	Attributes	445
15.18.4	Complex attributes	445
15.18.5	Mandatory attributes and roles	446
15.18.6	Identifiers	446
15.18.7	Other constraints	448
15.18.8	A complete example	449
15.18.9	Application of the relational theory to the Entity-relationship model	449
15.18.10	Relational interpretation vs relational translation	451
15.19	Summary of the chapter	451
15.20	Further reading	452
15.21	Tool-supported creation of a conceptual schema	453
15.22	Exercises	454
CHAPTER 16 • UML CLASS DIAGRAMS		459
16.1	Introduction	459
16.2	UML and databases	460
16.3	Extension mechanisms	461
16.4	Object classes	462
16.5	Binary associations	462

16.6	N-ary associations	463
16.6.1	Maximal multiplicity	464
16.6.2	Minimal multiplicity	464
16.7	Attributes	466
16.8	Operations	467
16.9	Qualified associations	467
16.10	Association classes	469
16.11	Agregation and composition associations	469
16.12	Constraints	470
16.13	The constraint language OCL	471
16.14	A matter of vocabulary	474
16.15	Example of database UML class diagram	474
16.16	UML class diagrams revisited (DB-UML)	474
16.17	Semantics of class diagrams in DB-UML	476
16.18	Summary of the chapter	477
16.19	Further reading	478
16.20	Tool-supported creation of a DB-UML conceptual schema	478
 CHAPTER 17 • CONCEPTUAL ANALYSIS OF THE APPLICATION DOMAIN		479
17.1	Introduction	479
17.2	Text analysis	482
17.2.1	Existential propositions	482
17.2.2	Binary relational propositions	483
17.2.3	N-ary relational propositions	486
17.2.4	Integrating a proposition in a schema	490
17.3	Design patterns	491
17.4	Structured document analysis	493
17.4.1	Analysis of organization documents	494
17.4.2	Analysis of electronic forms	495
17.4.3	Analysis of paper forms	496
17.5	Naming schema objects	497
17.6	Fixing schema errors	498
17.6.1	Syntactically incorrect constructs	499
17.6.2	Non-satisfiable constructs	500
17.7	Normalization of a conceptuel schema	506
17.7.1	Simplification of non minimal constructs	508
17.7.2	Explicitation of poorly expressive constructs	512
17.7.3	Reduction of over-specifications	517

17.7.4	Processing abnormal constructs	519
17.7.5	Unifying irregular constructs	521
17.7.6	Elimination of constraint redundancies	522
17.7.7	Elimination of structural redundancies	524
17.7.8	Elimination of internal redundancies	527
17.7.9	Improving the graphical presentation	530
17.7.10	Complying with current standards	533
17.8	Schema evaluation	536
17.8.1	Schema paraphrasing	537
17.8.2	Database prototyping	537
17.8.3	Structural inference	538
17.8.4	Error heuristics	538
17.8.5	Remark	539
17.9	Schema integration	540
17.9.1	Integration scenarios	540
17.9.2	Principles of binary schema integration	541
17.9.3	Correspondance identification	542
17.9.4	Schema unification	546
17.9.5	Schema merging	551
17.9.6	Resolving schema conflicts	551
17.9.7	Complements	552
17.9.8	Case study	553
17.10	Tool-supported conceptual analysis	558
17.10.1	Schema transformation	558
17.10.2	Schema analysis	560
17.10.3	Schema integration	560
17.11	Summary of the chapter	561
17.12	Further reading	562
17.13	Exercises	563
CHAPTER 18 • LOGICAL DESIGN OF A RELATIONAL DATABASE		569
18.1	Introduction	569
18.2	The relational logical model	571
18.2.1	Basic relational model	571
18.2.2	Enriched relational model	572
18.3	Representation of attributes	574
18.3.1	Representation of simple attributes	574
18.3.2	Representation of compound attributes	574
18.3.3	Representation of multivalued attributes	577
18.3.4	Representation of complex attributes	578

18.3.5	Representation of entity attributes	578
18.3.6	Translation of types and domains	579
18.3.7	Complements	580
18.4	Functional relationship types revisited	581
18.5	Representation of complex relationship types	582
18.5.1	N-arity, attributes and constraints	582
18.5.2	Representation of relationship types with multi-ET roles	582
18.5.3	Foreign keys and missing identifiers	583
18.6	Representation of <i>is-a</i> relations	584
18.6.1	The three basic techniques	585
18.6.2	Representation of <i>is-a</i> through materialization	586
18.6.3	Representation of <i>is-a</i> through downward inheritance	588
18.6.4	Representation of <i>is-a</i> through upward inheritance	590
18.6.5	Comparison of techniques	591
18.6.6	Multiple hierarchies and correlated foreign keys	592
18.7	Complements	593
18.7.1	Optional identifiers	593
18.7.2	Complex primary identifiers	594
18.7.3	Translation of conceptual constraints	594
18.7.4	Naming relational objects	595
18.7.5	<i>Delete</i> and <i>update modes</i> of foreign keys	595
18.7.6	About roles with cardinality [1-N]	596
18.8	Representation of a conceptual schema	597
18.8.1	Selection of best representations	597
18.8.2	Building the transformation plan	598
18.8.3	Limitations of a transformation plan	599
18.8.4	Validation of a transformation plan	600
18.9	View representation	601
18.10	Relational logical design: an example	601
18.11	Using SQL3 constructs	602
18.11.1	The object-relational model	605
18.11.2	Representation of complex attributes	605
18.11.3	Representation of <i>is-a</i> hierarchies	606
18.11.4	Example of object-relational design	607
18.11.5	Method extension	607
18.12	Tool-supported logical design	607
18.12.1	Schema transformation	607
18.12.2	Schema analysis	609
18.13	Summary of the chapter	609
18.14	Further reading	610

18.15 Exercises	611
CHAPTER 19 • PHYSICAL DESIGN OF A RELATIONAL DATABASE	617
19.1 Introduction	617
19.2 The concept of database efficiency	619
19.3 Sixteen rules for straightforward physical design	620
19.4 Index selection	620
19.4.1 Indexes derived from the logical structures	620
19.4.2 Indexes derived from data processing	621
19.4.3 Discarding indexes	622
19.4.4 Index maintenance	623
19.5 Index technology	624
19.6 Storage space selection	624
19.7 Advanced optimization	625
19.8 Evaluation of processing times and volumes	628
19.9 Summary of the chapter	629
19.10 Further reading	629
19.11 Exercises	630
CHAPTER 20 • CODING A RELATIONAL DATABASE	633
20.1 Introduction	633
20.2 Declarative expression of data structures	634
20.3 SQL constructs for expressing constraints	635
20.4 Coding non standard behaviours	638
20.5 Basic integrity constraints	639
20.6 Additional integrity constraints	640
20.7 Constraints derived from Entity-relationship constructs	643
20.8 Dynamic integrity constraints	645
20.9 Translation of <i>is-a</i> relations	645
20.10 View translation	649
20.11 Summary of the chapter	649
20.12 Code generation with a tool	650
20.13 Exercises	651
CHAPTER 21 • DATABASE REVERSE ENGINEERING	653
21.1 Introduction	653

21.2	General DBRE methodology	657
21.3	Physical extraction	658
21.4	Recovering the logical schema	659
21.4.1	Information sources and analysis techniques	659
21.4.2	Recovering implicit constructs	663
21.5	Logical schema conceptualization	665
21.6	Summary of the chapter	669
21.7	Further reading	670
21.8	Tool-supported DBRE	670
CHAPTER 22 • CASE STUDIES		671
22.1	Introduction	671
22.2	Zoo tycoon	671
22.2.1	Problem statement	672
22.2.2	Conceptual analysis	672
22.2.3	Production of the relational schema	674
22.2.4	The SQL code	674
22.2.5	Exploitation	675
22.3	Integrated circuits	676
22.3.1	Problem statement	676
22.3.2	Conceptual analysis	676
22.3.3	Relation schema and sample data	681
22.4	Reverse engineering a database	682
22.4.1	Extraction of the physical schema	682
22.4.2	Extraction of the logical schema	683
22.4.3	Conceptualization	684
22.5	Other case studies	686
BIBLIOGRAPHY		687
INDEX		695

APPENDICES

These appendices are available on the book website.

APPENDIX A - EXERCISES AND SOLUTIONS

A.1	Chapter 1	1
A.2	Chapter 2	1
A.3	Chapter 3	4
A.4	Chapter 4	10
A.5	Chapter 5	15
A.6	Chapter 6	16
A.7	Chapter 7	18
A.8	Chapter 8	29
A.9	Chapter 9	42
A.10	Chapter 10	44
A.11	Chapter 11	45
A.12	Chapter 12	49
A.13	Chapter 13	67
A.14	Chapter 14	71
A.15	Chapter 15	71
A.16	Chapter 16	83
A.17	Chapter 17	84
A.18	Chapter 18	102
A.19	Chapter 19	116
A.20	Chapter 20	119

APPENDIX B - SQL, SET THEORY AND LOGIC

B.1	SQL et mathematics	1
B.2	Sets and SQL	2
B.3	Multi-sets and SQL	7
B.4	Logic and SQL	10
B.5	Summary of the chapter	21
B.6	Further reading	21
B.7	Exercises	22

APPENDIX C - ADVANCED DATABASE APPLICATIONS

C.1	Introduction	1
C.2	Ordered structures	1
C.3	Active databases	4
C.4	Temporal databases	8
C.5	Code generation	18
C.6	Exercises	24

APPENDIX D - GRAPHICAL CONVENTIONS FOR SCHEMA REPRESENTATION

D.1	Introduction	1
D.2	The basic ER model	1
D.3	The extended ER model	2
D.4	The relational model	5
D.5	The object-relational model	6
D.6	The physical model	8

APPENDIX E - SCHEMA TRANSFORMATIONS

E.1	Introduction	1
E.2	General characteristics of transformations	2
E.3	Semantics preservation	4
E.4	Catalog of transformations	5
E.5	Mutation transformations	9
E.6	Transformation of entity type attributes	10
E.7	Transformation relationship type attributes	23
E.8	Transformations of entity types	26
E.9	Transformation of <i>is-a</i> relations	35
E.10	Transformation of relationship types	48
E.11	Transformations of constraints	58
E.12	Further reading	59
E.13	References and bibliography	60

APPENDIX F - CASE STUDIES - SEASON 2

F.1	Introduction	1
F.2	The case study portfolio	2
F.3	Air flights	3
F.4	Schema integration	8

F.5	Conceptual normalisation	17
F.6	Representation of graph structures	22
F.7	Key management	44
F.8	File migration	50
F.9	Data dictionary	79
F.10	Reverse engineering of a Hypercard stack (Formal Concept Analysis)	87
F.11	SQL and ontologies	98

APPENDIX G - LOGICAL DESIGN OF OBJECT-RELATIONAL DATABASES

G.1	Introduction	1
G.2	The object-relational logical model	2
G.3	Entity types and identifiers	4
G.4	Representation of complex attributes	5
G.5	Representation of binary relationship types	6
G.6	Representation of complex relationship types	10
G.7	Representation of is-a relations	12
G.8	Complements	17
G.9	Transformation of a conceptual schema	17
G.10	Logical design: example	19
G.11	Summary of the chapter	20
G.12	Further reading	21
G.13	Exercises	23

APPENDIX H - DATA STRUCTURE TECHNOLOGY - ADVANCED TOPICS

H.1	Introduction	1
H.2	Sequential processing of a file	1
H.3	Indexed sequential file organization	11
H.4	Hash file organization	53
H.5	Secondary indexes	64
H.6	Analysis of Microsoft SQL Server data structures	71
H.7	Exercises	73

SUGGESTED COURSES

This book and its complementary materials is intended to be used as the support of two complementary courses: *introductory* and *advanced*.

Introductory course

Understanding databases

- Chapter 1 - Motivation and introduction
- Chapter 2 - Concepts of databases

Using databases

- Chapter 6 - The SQL-DDL Language
- Chapter 7 - The SQL-DML Language (1)
- Chapter 8 - The SQL-DML Language (2)

Developing databases

- Chapter 10 - Building a database
- Chapter 11 - The basic Entity-relationship model
- Chapter 12 - Designing a conceptual model
- Chapter 13 - Producing a database schema

Case studies

- Chapter 22 - Case studies

Complete course

Understanding databases

- Chapter 1 - Motivation and introduction
- Chapter 2 - Concepts of databases
- Chapter 3 - Relational model and normalisation
- Chapter 4 - Implementation of data structures
- Chapter 5 - Database management systems

Using databases

- Chapter 6 - The SQL-DDL Language
- Chapter 7 - The SQL-DML Language (1)
- Chapter 8 - The SQL-DML Language (2)
- Chapter 9 - Advanced SQL

Developing databases

- Chapter 14 - Database methodology
- Chapter 15 - The extended Entity-relationship model
- Chapter 16 - UML class diagrams
- Chapter 17 - Conceptual analysis of the application domain
- Chapter 18 - Logical design of a relational database
- Chapter 19 - Physical design of a relational database
- Chapter 20 - Coding a relational database
- Chapter 21 - Database reverse engineering
- Appendix G - Logical design of object-relational databases

Case studies

- Chapter 22 - Case studies
- Appendix C - Advanced database applications
- Appendix F - Case studies - Season 2