

Appendixes

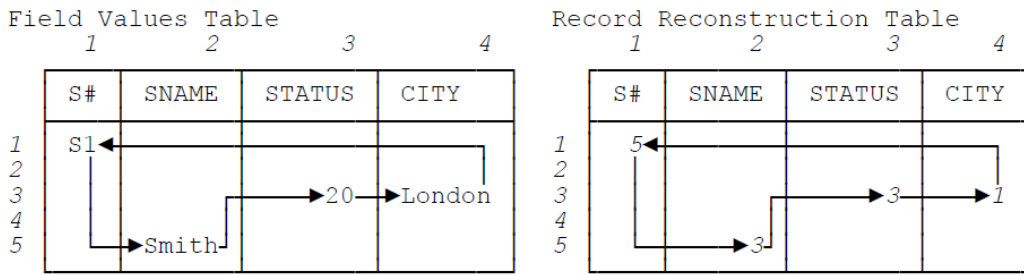
Appendix A Exercises

Exercise 1: Use the following Field Values Table and Record Reconstruction Table to reconstruct the suppliers file:

	1	2	3	4
	S#	SNAME	STATUS	CITY
1	S1	Adams	10	Athens
2	S2	Blake	20	London
3	S3	Clark	20	London
4	S4	Jones	30	Paris
5	S5	Smith	30	Paris

	1	2	3	4
	S#	SNAME	STATUS	CITY
1	5	4	4	5
2	4	5	2	4
3	2	2	3	1
4	3	1	1	2
5	1	3	5	3

The following diagram should serve to remind you how the reconstruction algorithm works (it shows the pointer rings for the record obtained by starting at cell [1,1] in each of the two tables):



The “first” reconstructed record is thus as shown below. You should be able to fill in the rest (begin with cell [2,1] in the Field Values Table, then cell [3,1], then cell [4,1], and finally cell [5,1]—in other words, proceed down column 1).

	1	2	3	4
	S#	SNAME	STATUS	CITY
1	S1	Smith	20	London
2				
3				
4				
5				

Your answer should look like Fig. 3.2.

Exercise 2: Use the following suppliers file and corresponding Permutation Table to build a Record Reconstruction Table:

File		1	2	3	4																										
	<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="width: 5%;"></th> <th style="width: 15%;">S#</th> <th style="width: 20%;">SNAME</th> <th style="width: 15%;">STATUS</th> <th style="width: 45%;">CITY</th> </tr> </thead> <tbody> <tr> <td style="text-align: right;">1</td> <td>S4</td> <td>Clark</td> <td>20</td> <td>London</td> </tr> <tr> <td style="text-align: right;">2</td> <td>S5</td> <td>Adams</td> <td>30</td> <td>Athens</td> </tr> <tr> <td style="text-align: right;">3</td> <td>S2</td> <td>Jones</td> <td>10</td> <td>Paris</td> </tr> <tr> <td style="text-align: right;">4</td> <td>S1</td> <td>Smith</td> <td>20</td> <td>London</td> </tr> <tr> <td style="text-align: right;">5</td> <td>S3</td> <td>Blake</td> <td>30</td> <td>Paris</td> </tr> </tbody> </table>		S#	SNAME	STATUS	CITY	1	S4	Clark	20	London	2	S5	Adams	30	Athens	3	S2	Jones	10	Paris	4	S1	Smith	20	London	5	S3	Blake	30	Paris
	S#	SNAME	STATUS	CITY																											
1	S4	Clark	20	London																											
2	S5	Adams	30	Athens																											
3	S2	Jones	10	Paris																											
4	S1	Smith	20	London																											
5	S3	Blake	30	Paris																											

	1	2	3	4																											
Permutation Table																															
	<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="width: 5%;"></th> <th style="width: 15%;">S#</th> <th style="width: 20%;">SNAME</th> <th style="width: 15%;">STATUS</th> <th style="width: 45%;">CITY</th> </tr> </thead> <tbody> <tr> <td style="text-align: right;">1</td> <td>4</td> <td>2</td> <td>3</td> <td>2</td> </tr> <tr> <td style="text-align: right;">2</td> <td>3</td> <td>5</td> <td>1</td> <td>1</td> </tr> <tr> <td style="text-align: right;">3</td> <td>5</td> <td>1</td> <td>4</td> <td>4</td> </tr> <tr> <td style="text-align: right;">4</td> <td>1</td> <td>3</td> <td>2</td> <td>3</td> </tr> <tr> <td style="text-align: right;">5</td> <td>2</td> <td>4</td> <td>5</td> <td>5</td> </tr> </tbody> </table>		S#	SNAME	STATUS	CITY	1	4	2	3	2	2	3	5	1	1	3	5	1	4	4	4	1	3	2	3	5	2	4	5	5
	S#	SNAME	STATUS	CITY																											
1	4	2	3	2																											
2	3	5	1	1																											
3	5	1	4	4																											
4	1	3	2	3																											
5	2	4	5	5																											

Here's the algorithm:

Step 1: Let PT be the Permutation Table. Build a table RRT with the same number of rows and columns as PT and with all cells empty.

Step 2: For all records in the user file, do *Step 3*.

Step 3: For all columns of PT , do *Step 4*.

Step 4: Let the current record of the user file be the r th record, and let the current column of PT be the j th column. Let cell $[i,j]$ of PT be the cell of column j that contains the record number r . At cell $[i,j]$ of RRT , place the value i' , where cell $[i',j+1]$ of PT is the cell of column $j+1$ that contains the record number r . If column j is the last column, take column $j+1$ as the first column.

After this algorithm has been executed, table RRT is the desired Record Reconstruction Table:

	1	2	3	4
	S#	SNAME	STATUS	CITY
1				
2				
3				
4				
5				

Your answer should look like the Record Reconstruction Table shown in Exercise 1.


Exercise 3: Use the following suppliers file to build a corresponding Field Values Table:

File					Field Values Table				
	1	2	3	4		1	2	3	4
	S#	SNAME	STATUS	CITY		S#	SNAME	STATUS	CITY
1	S1	Smith	20	London	1				
2	S2	Jones	10	Paris	2				
3	S3	Blake	30	Paris	3				
4	S4	Clark	20	London	4				
5	S5	Adams	30	Athens	5				


Your answer should look like the Field Values Table shown in **Exercise 1**. Now construct a corresponding Permutation Table and (using that Permutation Table) a corresponding Record Reconstruction Table:

Permutation Table					Record Reconstruction Table				
	1	2	3	4		1	2	3	4
	S#	SNAME	STATUS	CITY		S#	SNAME	STATUS	CITY
1					1				
2					2				
3					3				
4					4				
5					5				

Does your Record Reconstruction Table look like Fig. 3.5? If not, why not?

SIMPLY CLEVER


We will turn your CV into an opportunity of a lifetime



Do you like cars? Would you like to be a part of a successful brand?
We will appreciate and reward both your enthusiasm and talent.
Send us your CV. You will be surprised where it can take you.

Send us your CV on
www.employerforlife.com



Exercise 4: Use the following Permutation Table to build a corresponding Inverse Permutation Table:

Permutation Table

	1	2	3	4
	S#	SNAME	STATUS	CITY
1	4	2	3	2
2	3	5	1	1
3	5	1	4	4
4	1	3	2	3
5	2	4	5	5

Inverse Permutation Table

	1	2	3	4
	S#	SNAME	STATUS	CITY
1				
2				
3				
4				
5				

Recall that if you think of any given permutation as a vector V , then the inverse permutation V' can be obtained in accordance with the rule that if $V[i] = i'$, then $V'[i'] = i$. Your answer should look like the Inverse Permutation Table shown in **Exercise 5**.

Exercise 5: Use the following Inverse Permutation Table to build a Record Reconstruction Table:

	1	2	3	4
S#	4	3	2	2
SNAME	5	1	4	1
STATUS	2	4	1	4
CITY	1	5	3	3
	3	2	5	5

	1	2	3	4
S#				
SNAME				
STATUS				
CITY				

Here's the algorithm:

Go to cell $[i, 1]$ of the Inverse Permutation Table. Let that cell contain the value r ; also, let the next cell to the right, cell $[i, 2]$, contain the value r' . Go to the r th row of the Record Reconstruction Table and place the value r' in cell $[r, 1]$.

Executing this algorithm for $i = 1, 2, \dots, 5$ yields the entire S# column of the Record Reconstruction Table. The other columns are built analogously. Your answer should look like the Record Reconstruction Table shown in **Exercise 1**.

I joined MITAS because
I wanted **real responsibility**

The Graduate Programme
for Engineers and Geoscientists
www.discovermitas.com




Real work
International opportunities
Three work placements



Month 16
I was a construction
supervisor in
the North Sea
advising and
helping foremen
solve problems





Exercise 6: Given the following suppliers file—

File

	1	2	3	4
	S#	SNAME	STATUS	CITY
1	S4	Clark	20	London
2	S5	Adams	30	Athens
3	S2	Jones	10	Paris
4	S1	Smith	20	London

—check that the following Field Values Table and Record Reconstruction Table are correct:

Field Values Table					Record Reconstruction Table				
	1	2	3	4		1	2	3	4
	S#	SNAME	STATUS	CITY		S#	SNAME	STATUS	CITY
1	S1	Adams	10	Athens	1	4	4	4	4
2	S2	Clark	20	London	2	3	2	2	3
3	S4	Jones	20	London	3	2	1	3	1
4	S5	Smith	30	Paris	4	1	3	1	2

Exercise 7: Given the following suppliers file—

File

	1	2	3	4
	S#	SNAME	STATUS	CITY
1	S4	Clark	20	London
2	S5	Adams	30	Athens
3	S2	Jones	10	Paris
4	S1	Smith	20	London
5	S3	Blake	30	Paris
6	S6	Brady	30	Athens
7	S7	Patel	40	Haifa

—check that the following Field Values Table and Record Reconstruction Table are correct:

Field Values Table

	1	2	3	4
	S#	SNAME	STATUS	CITY
1	S1	Adams	10	Athens
2	S2	Blake	20	Athens
3	S3	Brady	20	Haifa
4	S4	Clark	30	London
5	S5	Jones	30	London
5	S6	Patel	30	Paris
5	S7	Smith	40	Paris

Record Reconstruction Table

	1	2	3	4
	S#	SNAME	STATUS	CITY
1	7	4	6	5
2	5	6	4	6
3	2	5	5	7
4	4	3	1	1
5	1	1	2	4
5	3	7	7	2
5	6	2	3	3

Exercise 8: Use the following Field Values and Record Reconstruction Tables to reconstruct the shipments file, starting at cell [1,1] of each of the two tables for the first record in that reconstruction and continuing down column 1. Then do the same thing again, but this time going down column 2; and then again, going down column 3; and then again, going down column 4.

	1	2	3	4
	S#	P#	J#	QTY
1	S1	P1	J1	100
2	S1	P1	J1	100
3	S2	P1	J1	200
4	S2	P1	J1	200
5	S2	P2	J2	200
6	S3	P2	J2	200
7	S3	P3	J2	500
8	S3	P3	J2	500
9	S3	P3	J2	500

	1	2	3	4
	S#	P#	J#	QTY
1	2	1	2	2
2	8	2	3	6
3	3	3	4	1
4	4	7	5	3
5	5	8	1	8
6	1	9	6	9
7	6	4	7	4
8	7	5	8	5
9	9	6	9	7

Your answers should be as shown in Fig. 6.5.

ie business school

#1 EUROPEAN BUSINESS SCHOOL
FINANCIAL TIMES 2013

#gobeyond

MASTER IN MANAGEMENT

Because achieving your dreams is your greatest challenge. IE Business School's Master in Management taught in English, Spanish or bilingually, trains young high performance professionals at the beginning of their career through an innovative and stimulating program that will help them reach their full potential.

- Choose your area of specialization.
- Customize your master through the different options offered.
- Global Immersion Weeks in locations such as London, Silicon Valley or Shanghai.

Because you change, we change with you.

www.ie.edu/master-management | mim.admissions@ie.edu |



Exercise 9: Use the following Inverse Permutation Table to build a “preferred” Record Reconstruction Table for shipments:

Inverse Permutation Table
1 2 3 4

	S#	P#	J#	QTY
1	1	2	2	3
2	2	8	5	1
3	3	3	3	4
4	4	4	7	7
5	5	5	8	8
6	6	1	1	2
7	7	6	9	9
8	8	7	4	5
9	9	9	6	6

Record Reconstruction Table

	S#	P#	J#	QTY
1				
2				
3				
4				
5				
6				
7				
8				
9				

Your answer should look like Fig. 6.4.

Exercise 10: Given the following shipments file, show a Permutation Table corresponding to the following sort orders:

- For column S# : S# - P# - J#
- For column P# : P# - J# - S#
- For column J# : J# - S# - P#
- For column QTY : QTY - S# - P# - J#

File	1	2	3	4
	S#	P#	J#	QTY
1	S1	P1	J1	200
2	S1	P3	J2	100
3	S2	P1	J1	200
4	S2	P1	J2	500
5	S2	P2	J2	500
6	S3	P1	J1	100
7	S3	P2	J2	500
8	S3	P3	J1	200
9	S3	P3	J2	200

Permutation Table	1	2	4	5
	S#	P#	J#	QTY
1				
2				
3				
4				
5				
6				
7				
8				
9				

Use this Permutation Table to build a corresponding Inverse Permutation Table and a corresponding Record Reconstruction Table:

Inverse Permutation Table	1	2	3	4
	S#	P#	J#	QTY
1				
2				
3				
4				
5				
6				
7				
8				
9				

Record Reconstruction Table	1	2	3	4
	S#	P#	J#	QTY
1				
2				
3				
4				
5				
6				
7				
8				
9				

Your Record Reconstruction Table should look like Fig. 6.6. Check that this Record Reconstruction Table does exhibit the desired behavior regarding major-to-minor orderings.

Exercise 11: Use the following Field Values Table and Record Reconstruction Table to reconstruct the parts file. Start with column 5 in order to obtain the result in ascending city name sequence.

Field Values Table

	1	2	3	4	5
	P#	PNAME	COLOR	WEIGHT	CITY
1	P1	Bolt [1:1]	Blue [1:2]	12.0 [1:2]	London [1:3]
2	P2	Cam [2:2]	Green [3:3]	14.0 [3:3]	Oslo [4:4]
3	P3	Cog [3:3]	Red [4:6]	17.0 [4:5]	Paris [5:6]
4	P4	Nut [4:4]		19.0 [6:6]	
5	P5	Screw [5:6]			
6	P6				

Record Reconstruction Table

	1	2	3	4	5
	P#	PNAME	COLOR	WEIGHT	CITY
1	4	3	2	1	1
2	1	1	4	6	4
3	5	6	5	2	6
4	6	4	1	4	3
5	2	2	3	5	2
6	3	5	6	3	5

To remind you, here's the crucial revision to the reconstruction algorithm:

Consider cell $[i,j]$ of the Record Reconstruction Table. Instead of going to cell $[i,j]$ of the Field Values Table, go to cell $[i',j]$ of that table, where cell $[i',j]$ is that unique cell within column j of that table that contains a row range that includes row i .

Your answer should look like Fig. 7.2, except that the records should appear in ascending city name sequence.

Exercise 12: Use the following Field Values Table and Record Reconstruction Table to reconstruct the parts file. Start with column 5 in order to obtain the result in ascending city name sequence.

Field Values Table

	1	2	3	4	5
	P#	PNAME	COLOR	WEIGHT	CITY
1	P1	Bolt [1:1]	Blue [1:2]	12.0 [1:2]	London [1:3]
2	P2	Cam [2:2]	Green [3:3]	14.0 [3:3]	Oslo [4:4]
3	P3	Cog [3:3]	Red [4:6]	17.0 [4:5]	Paris [5:6]
4	P4	Nut [4:4]		19.0 [6:6]	
5	P5	Screw [5:6]			
6	P6				

Record Reconstruction Table

	1	2	3	4	5
	P#	PNAME	COLOR	WEIGHT	CITY
1	4	1•3	1•2	1•1	1•1
2	1	2•1	1•4	1•6	1•4
3	5	3•6	2•5	2•2	1•6
4	6	4•4	3•1	3•4	2•3
5	2	5•2	3•3	3•5	3•2
6	3	5•5	3•6	4•3	3•5

To remind you, in those columns of the Record Reconstruction Table that include two row numbers instead of one, the first is the number of the desired row within the Field Values Table, and the second is the number of the next row to be inspected within the Record Reconstruction Table. As in **Exercise 11**, your answer should look like Fig. 7.2, except that the records should appear in ascending city name sequence.

“I studied English for 16 years but...
...I finally learned to speak it in just six lessons”
Jane, Chinese architect

ENGLISH OUT THERE

Click to hear me talking before and after my unique course download



Exercise 13: Given the following bill-of-materials file—

File

	1	2	3
	MAJOR_P#	MINOR_P#	QTY
1	P3	P4	3
2	P1	P3	4
3	P2	P4	8
4	P1	P4	1
5	P2	P5	6
6	P3	P6	4
7	P1	P2	2
8	P5	P6	3
9	P2	P3	3

—check that the following Field Values Table and Record Reconstruction Table are correct:

Field Values Table
1 2 3

	MAJOR_P#	MINOR_P#	QTY
1	P1	P2	1
2	P1	P3	2
3	P1	P3	3
4	P2	P4	3
5	P2	P4	3
6	P2	P4	4
7	P3	P5	4
8	P3	P6	6
9	P5	P6	8

Record Reconstruction Table
1 2 3

	MAJOR_P#	MINOR_P#	QTY
1	1	2	3
2	2	6	1
3	4	3	4
4	3	1	7
5	5	9	9
6	7	4	2
7	6	8	8
8	8	7	6
9	9	5	5

Note: The Record Reconstruction Table is intended to reflect the following sort orders:

- For column MAJOR_P#: MAJOR_P# - MINOR_P# - QTY
- For column MINOR_P#: MINOR_P# - MAJOR_P# - QTY
- For column QTY: QTY - MAJOR_P# - MINOR_P#

Exercise 14: Given the Field Values Table and Record Reconstruction Table from **Exercise 13**, check that the following condensed and expanded versions (respectively) are correct.

Field Values Table

	1	2	3
	MAJOR_P#	MINOR_P#	QTY
1	P1 [1:3]	P2 [1:1]	1 [1:1]
2	P2 [4:6]	P3 [2:3]	2 [2:2]
3	P3 [7:8]	P4 [4:6]	3 [3:5]
4	P5 [9:9]	P5 [7:7]	4 [6:7]
5		P6 [8:9]	6 [8:8]
6			8 [9:9]
	-----	-----	-----

Record Reconstruction Table

	1	2	3
	MAJOR_P#	MINOR_P#	QTY
1	1•1	1•2	1•3
2	1•2	2•6	2•1
3	1•4	2•3	3•4
4	2•3	3•1	3•7
5	2•5	3•9	3•9
6	2•7	3•4	4•2
7	3•6	4•8	4•8
8	3•8	5•7	5•6
9	4•9	5•5	6•5

Exercise 15: Use the following Field Values Table and Record Reconstruction Table to reconstruct the bill-of-materials file. Start with column 1 in order to obtain the result in sequence by minor part number within major part number.

Field Values Table

	1	2	
	MAJOR_P#	+ MINOR_P#	QTY
1	P1	[1:3] [:]	1 [1:1]
2	P2	[4:6] [1:1]	2 [2:2]
3	P3	[7:8] [2:3]	3 [3:5]
4	P4	[:] [4:6]	4 [6:7]
3	P5	[9:9] [7:7]	6 [8:8]
4	P6	[:] [8:9]	8 [9:9]

Record Reconstruction Table

	1	2	3
	MAJOR_P#	MINOR_P#	QTY
1	1▪1	2▪2	1▪3
2	1▪2	3▪6	2▪1
3	1▪4	3▪3	3▪4
4	2▪3	4▪1	3▪7
5	2▪5	4▪9	3▪9
6	2▪7	4▪4	4▪2
7	3▪6	5▪8	4▪8
8	3▪8	6▪7	5▪6
9	5▪9	6▪5	6▪5

Your answer should be a file that's a direct image of relation MMQ as shown in Fig. 8.1.

Excellent Economics and Business programmes at:



**university of
 groningen**





“The perfect start
 of a successful,
 international career.”

CLICK HERE
 to discover why both socially
 and academically the University
 of Groningen is one of the best
 places for a student to be

www.rug.nl/feb/education



Exercise 16: Given the following suppliers file—

File

	1	2	3	4
	S#	SNAME	STATUS	CITY
1	S1	Smith	20	London
2	S2	Jones	10	Paris
3	S3	Blake	30	Paris
4	S4	Clark	20	London
5	S5	Adams	30	Athens

—check that the following Field Values Table and Record Reconstruction Table are correct:

Field Values Table

	1	2	3	4
	S#	SNAME	STATUS	CITY
1	S1	Adams [1:1]	10 [1:1]	Athens [1:1]
2	S2	Blake [2:2]	20 [2:3]	London [2:3]
3	S3	Clark [3:3]	30 [4:5]	Paris [4:5]
4	S4	Jones [4:4]		
5	S5	Smith [5:5]		

Record Reconstruction Table

	1	2	3	4
	S#	SNAME	STATUS	CITY
1	5	1•5	1•4	1•5
2	4	2•4	2•2	2•1
3	2	3•3	2•3	2•4
4	3	4•1	3•5	3•2
5	1	5•2	3•1	3•3

Likewise, given the following shipments file—

File

	1	2	3	4
	S#	P#	J#	QTY
1	S1	P1	J1	200
2	S1	P3	J2	100
3	S2	P1	J1	200
4	S2	P1	J2	500
5	S2	P2	J2	500
6	S3	P1	J1	100
7	S3	P2	J2	500
8	S3	P3	J1	200
9	S3	P3	J2	200

—check that the following Field Values Table and Record Reconstruction Table are correct:

Field Values Table

	1	2	3	4
	S#	P#	J#	QTY
1	S1 [1:2]	P1 [1:4]	J1 [1:4]	100 [1:2]
2	S2 [3:5]	P2 [5:6]	J2 [5:9]	200 [3:6]
3	S3 [6:9]	P3 [7:9]		500 [7:9]

Record Reconstruction Table

	1	2	3	4
	S#	P#	J#	QTY
1	1•2	1•1	1•2	1•2
2	1•8	1•2	1•3	1•6
3	2•3	1•3	1•4	2•1
4	2•4	1•7	1•5	2•3
5	2•5	2•8	2•1	2•8
6	3•1	2•9	2•6	2•9
7	3•6	3•4	2•7	3•4
8	3•7	3•5	2•8	3•5
9	3•9	3•6	2•9	3•7

Finally, check that the following merged Field Values Table is correct:

Field Values Table

	1	2	3	4	5	6	7
	S#	SNAME	STATUS	CITY	P#	J#	QTY
1	S1 [1:2]	Adams [1:1]	10 [1:1]	Athens [1:1]	P1 [1:4]	J1 [1:4]	100 [1:2]
2	S2 [3:5]	Blake [2:2]	20 [2:3]	London [2:3]	P2 [5:6]	J2 [5:9]	200 [3:6]
3	S3 [6:9]	Clark [3:3]	30 [4:5]	Paris [4:5]	P3 [7:9]		500 [7:9]
4	S4 [:]	Jones [4:4]					
5	S5 [:]	Smith [5:5]					

Appendix B References and Bibliography

What follows is a consolidated list of references for the entire book. Let me immediately apologize for the embarrassingly large number of references to publications for which I'm the author or a coauthor; such a state of affairs is more or less unavoidable, however, given the nature of this book and the history of my own involvement in this field.

1. M. M. Astrahan et al.: "System R: Relational Approach to Database Management," *ACM Transactions on Database Systems* 1, No. 2 (June 1976).
2. E. Babb: "Implementing a Relational Database by Means of Specialized Hardware," *ACM Transactions on Database Systems* 4, No. 1 (March 1979).
3. R. G. G. Cattell: *Object Data Management* (revised edition). Reading, Mass.: Addison-Wesley (1994).
4. R. G. G. Cattell and Douglas K. Barry (eds.): *The Object Database Standard: ODMG 3.0*. San Francisco, Calif.: Morgan Kaufmann (2000).
5. E. F. Codd: "Derivability, Redundancy, and Consistency of Relations Stored in Large Data Banks," IBM Research Report RJ599 (August 19th, 1969).
6. E. F. Codd: "A Relational Model of Data for Large Shared Data Banks," *Communications of the ACM* 13, No. 6 (June 1970). Republished in "Milestones of Research," *Communications of the ACM* 26, No. 1 (January 1982).
7. E. F. Codd: "Relational Completeness of Data Base Sublanguages," in Randall J. Rustin (ed.), *Data Base Systems: Courant Computer Science Symposia Series 6*. Englewood Cliffs, N.J.: Prentice-Hall (1972).
8. E. F. Codd: "Recent Investigations into Relational Data Base Systems," Proc. 1974 Congress, Stockholm, Sweden (1974).
9. E. F. Codd: *The Relational Model for Database Management Version 2*. Reading, Mass.: Addison-Wesley (1990).
10. E. F. Codd and C. J. Date: "Interactive Support for Nonprogrammers: The Relational and Network Approaches," Proc. ACM SIGMOD Workshop on Data Description, Access, and Control (Vol. II), Ann Arbor, Mich. (May 1974). Republished in C. J. Date, *Relational Database: Selected Writings*. Reading, Mass.: Addison-Wesley (1986).
11. Hugh Darwen (writing as Andrew Warden): "The Keys of the Kingdom," in C. J. Date, *Relational Database Writings 1985-1989*. Reading, Mass.: Addison-Wesley (1990).
12. Hugh Darwen: "The Nullologist in Relationland; or, Nothing Really Matters," in C. J. Date (with Hugh Darwen), *Relational Database Writings 1989-1991*. Reading, Mass.: Addison-Wesley (1992).
13. Hugh Darwen: "The Duplicity of Duplicate Rows," in C. J. Date (with Hugh Darwen), *Relational Database Writings 1989-1991*. Reading, Mass.: Addison-Wesley (1992).
14. Data Base Task Group of CODASYL Programming Language Committee: *Report* (April 1971).
15. C. J. Date: "A Critique of the SQL Database Language," in *Relational Database: Selected Writings*. Reading, Mass.: Addison-Wesley (1986).
16. C. J. Date: "What's Wrong with SQL?," in *Relational Database Writings 1985-1989*. Reading, Mass.: Addison-Wesley (1990).
17. C. J. Date: "SQL Dos and Don'ts," in *Relational Database Writings 1985-1989*. Reading, Mass.: Addison-Wesley (1990).
18. C. J. Date: "NOT Is Not 'Not!' (Notes on Three-Valued Logic and Related Matters)," in *Relational Database Writings 1985-1989*. Reading, Mass.: Addison-Wesley (1990).

Download free eBooks at bookboon.com

19. C. J. Date: “EXISTS Is Not ‘Exists!’ (Some Logical Flaws in SQL),” in *Relational Database Writings 1985-1989*. Reading, Mass.: Addison-Wesley (1990).
20. C. J. Date: “Why Duplicate Rows Are Prohibited,” in *Relational Database Writings 1985-1989*. Reading, Mass.: Addison-Wesley (1990). A considerably updated version of this article is “Double Trouble, Double Trouble” (in two parts), published on the website *www.dbdebunk.com* (April 2002).
21. C. J. Date: “Support for the Conceptual Schema: The Relational and Network Approaches,” in *Relational Database Writings 1985-1989*. Reading, Mass.: Addison-Wesley (1990).
22. C. J. Date: “Three-Valued Logic and the Real World,” in C. J. Date (with Hugh Darwen), *Relational Database Writings 1989-1991*. Reading, Mass.: Addison-Wesley (1992).
23. C. J. Date: “Oh No Not Nulls Again,” in C. J. Date (with Hugh Darwen), *Relational Database Writings 1989-1991*. Reading, Mass.: Addison-Wesley (1992).
24. C. J. Date: “Tables with No Columns” and “More on DEE and DUM,” both in *Relational Database Writings 1991-1994*. Reading, Mass.: Addison-Wesley (1995).
25. C. J. Date: “An Inverted List System: DATACOM/DB,” “A Hierarchic System: IMS,” and “A Network System: IDMS,” all in *Relational Database Writings 1991-1994*. Reading, Mass.: Addison-Wesley (1995).
26. C. J. Date: “Storage Structures and Access Methods,” Appendix A of *An Introduction to Database Systems* (6th edition). Reading, Mass.: Addison-Wesley (1995).
27. C. J. Date: “The Normal Is So ... Interesting” (in two parts), *Database Programming & Design* 10, No. 11 (November 1997) and No. 12 (December 1997).
28. C. J. Date: “Why the ‘Object Model’ Is Not a Data Model,” in C. J. Date (with Hugh Darwen and David McGoveran), *Relational Database Writings 1994-1997*. Reading, Mass.: Addison-Wesley (1998).
29. C. J. Date: “Object Identifiers vs. Relational Keys,” in C. J. Date (with Hugh Darwen and David McGoveran), *Relational Database Writings 1994-1997*. Reading, Mass.: Addison-Wesley (1998).
30. C. J. Date: “Don’t Mix Pointers and Relations!” and “Don’t Mix Pointers and Relations—Please!”, both in C. J. Date (with Hugh Darwen and David McGoveran), *Relational Database Writings 1994-1997*. Reading, Mass.: Addison-Wesley (1998).
31. C. J. Date: “Grievous Bodily Harm” (in two parts), *Database Programming & Design* 11, No. 5 (May 1998) and No. 6 (June 1998).
32. C. J. Date: *An Introduction to Database Systems* (7th edition). Reading, Mass.: Addison-Wesley (2000).
33. C. J. Date: “The Relational Model,” Part II of reference [32].
34. C. J. Date: *WHAT Not HOW: The Business Rules Approach to Application Development*. Reading, Mass.: Addison-Wesley (2000).
35. C. J. Date: *The Database Relational Model: A Retrospective Review and Analysis*. Reading, Mass.: Addison-Wesley (2001).
36. C. J. Date: “Constraints and Predicates: A Brief Tutorial” (in three parts), published on the websites *www.dbdebunk.com* (May 2001) and *www.BRCCommunity.com* (May/September/November 2001).
37. C. J. Date and E. F. Codd: “The Relational and Network Approaches: Comparison of the Application Programming Interfaces,” Proc. ACM SIGMOD Workshop on Data Description, Access, and Control (Vol. II), Ann Arbor, Mich. (May 1974). Republished in C. J. Date, *Relational Database: Selected Writings*. Reading, Mass.: Addison-Wesley (1986).
38. C. J. Date and E. F. Codd: “Much Ado About Nothing,” in C. J. Date, *Relational Database Writings 1991-1994*. Reading, Mass.: Addison-Wesley (1995).
39. C. J. Date (with Hugh Darwen): *A Guide to the SQL Standard* (4th edition). Reading, Mass.: Addison-Wesley (1997).

40. C. J. Date and Hugh Darwen: *Foundation for Future Database Systems: The Third Manifesto* (2nd edition). Reading, Mass.: Addison-Wesley (2000).
41. C. J. Date and Hugh Darwen: "Subtyping and Inheritance," Part IV of reference [39].
42. C. J. Date, Hugh Darwen, and Nikos A. Lorentzos: *Temporal Data and the Relational Model*. San Francisco, Calif.: Morgan Kaufmann (2002, to appear).
43. C. J. Date, Hugh Darwen, and David McGoveran: "Nothing to Do with the Case," in *Relational Database Writings 1994-1997*. Reading, Mass.: Addison-Wesley (1998).
44. C. J. Date, Hugh Darwen, and David McGoveran: "Up to a Point, Lord Copper," in *Relational Database Writings 1994-1997*. Reading, Mass.: Addison-Wesley (1998).
45. C. J. Date (with Colin J. White): *A Guide to DB2* (4th edition). Reading, Mass.: Addison-Wesley (1992).
46. Keith Devlin: *The Math Gene*. New York, N.Y.: Basic Books (2000).
47. Andrew Eisenberg and Jim Melton: "SQL:1999, Formerly Known as SQL3," *ACM SIGMOD Record* 28, No. 1 (March 1999).
48. Ronald Fagin, Jurg Nievergelt, Nicholas Pippenger, and H. Raymond Strong: "Extendible Hashing—A Fast Access Method for Dynamic Files," *ACM Transactions on Database Systems* 4, No. 3 (September 1979).
49. Richard Finkelstein: "Sybase IQ: Expressly for the Warehouse," *Database Programming & Design* 9, No. 12 (December 1996).
50. Hector Garcia-Molina, Jeffrey D. Ullman, and Jennifer Widom: *Database System Implementation*. Upper Saddle River, N.J.: Prentice Hall (2000).
51. Anil K. Garg and C. C. Gotlieb: "Order-Preserving Key Transformations," *ACM Transactions on Database Systems* 11, No. 2 (June 1986).
52. Robert C. Goldstein and Alois J. Strnad: "The MacAIMS Data Management System," Proc. 1970 ACM SICFIDET Workshop on Data Description and Access, Houston, Texas (November 1970).
53. International Organization for Standardization (ISO): *Database Language SQL*, Document ISO/IEC 9075:1999. Also available as American National Standards Institute (ANSI) Document ANSI NCITS.135-1999.
54. Won Kim: "On Marrying Relations and Objects: Relation-Centric and Object-Centric Perspectives," *Data Base Newsletter* 22, No. 6 (November/December 1994).
55. Donald E. Knuth: *The Art of Computer Programming. Volume III: Sorting and Searching*. Reading, Mass.: Addison-Wesley (1973).
56. Vincent Y. Lum: "Multi-Attribute Retrieval with Combined Indexes," *Communications of the ACM* 13, No. 11 (November 1970).
57. William C. McGee: "The IMS/VS System," *IBM Sys. J.* 16, No. 2 (June 1977).
58. David McGoveran: "Nothing from Nothing" (in four parts), in C. J. Date (with Hugh Darwen and David McGoveran), *Relational Database Writings 1994-1997*. Reading, Mass.: Addison-Wesley (1998).
59. T. H. Merrett: "Why Sort/Merge Gives the Best Implementation of the Natural Join," *ACM SIGMOD Record* 13, No. 2 (January 1983).
60. R. Morris: "Scatter Storage Techniques," *Communications of the ACM* 11, No. 1 (January 1968).
61. James K. Mullin: "Retrieval-Update Speed Tradeoffs Using Combined Indices," *Communications of the ACM* 14, No. 12 (December 1971).
62. George Polya: *How To Solve It* (2nd edition). Princeton, N.J.: Princeton University Press (1971).
63. U.S. Patent and Trademark Office: *Value-Instance-Connectivity Computer-Implemented Database*. US Patent No. 6,009,432 (December 28th, 1999).

64. Ben Shneiderman: "Reduced Combined Indexes for Efficient Multiple Attribute Retrieval," *Information Systems* 2, No. 4 (1976).
65. Alois J. Strnad: "The Relational Approach to the Management of Data Bases," Proc. 1971 IFIP Congress, Ljubljana, Yugoslavia (August 1971).
66. Peter J. Titman: "An Experimental Data Base System Using Binary Relations," Proc. IFIP TC-2 Working Conference on Data Base Management Systems (eds., Klimbie & Koffeman), Cargèse, Corsica (April 1974). Amsterdam, Netherlands: North-Holland (1974).
67. Transaction Processing Council (TPC): "TPC-H Result Highlights - HP 9000 Superdome Enterprise Server - Executive Summary," www.tpc.org/results/individual_results/HP/hp_tpch_sd_750_es.prof.
68. Patrick Valduriez: "Join Indices," *ACM Transactions on Database Systems* 12, No. 2 (June 1987).
69. Gio Wiederhold: *Database Design* (2nd edition). New York, N.Y.: McGraw-Hill (1983).



LIGS University
based in Hawaii, USA

is currently enrolling in the
Interactive Online **BBA, MBA, MSc,**
DBA and PhD programs:

- ▶ enroll **by October 31st, 2014** and
- ▶ **save up to 11%** on the tuition!
- ▶ pay in 10 installments / 2 years
- ▶ Interactive **Online education**
- ▶ visit www.ligsuniversity.com to find out more!

Note: LIGS University is not accredited by any nationally recognized accrediting agency listed by the US Secretary of Education. More info [here](#).

