Relational Databases

Characteristics of a Relational Table

TABLE 3.1 CHARACTERISTICS OF A RELATIONAL TABLE

- 1 A table is perceived as a two-dimensional structure composed of rows and columns.
- 2 Each table row (tuple) represents a single entity occurrence within the entity set.
- 3 Each table column represents an attribute, and each column has a distinct name.
- 4 Each row/column intersection represents a single data value.

TABLE 3.1 CHARACTERISTICS OF A RELATIONAL TABLE (CONTINUED)

- 5 All values in a column must conform to the same data format. For example, if the attribute is assigned an integer data format, all values in the column representing that attribute must be integers.
- 6 Each column has a specific range of values known as the **attribute domain**.
- 7 The order of the rows and columns is immaterial to the DBMS.
- 8 Each table must have an attribute or a combination of attributes that uniquely identifies each row.

STUDENT Table Attribute Values

Database name: Ch03_TinyCollege Table name: STUDENT

	STU_NUM	STU_LNAME	STU_FNAME	STU_INIT	STU_DOB	STU_HRS	STU_CLASS
•	321452	Bowser	William	С	12-Feb-1972	42	So
	324257	Smithson	Anne	K	15-Nov-1977	81	Jr
1.733	324258	Brewer	Juliette		23-Aug-1966	36	So
8.00	324269	Oblonski	Walter	Н	16-Sep-1973	66	Jr
138	324273	Smith	John	D	30-Dec-1955	102	Sr
	324274	Katinga	Raphael	P	21-Oct-1976	114	Sr
	324291	Robertson	Gerald	Т	08-Apr-1970	120	Sr
100	324299	Smith	John	В	30-Nov-1983	15	Fr

	hla	1.1.1.1.1		STU_GPA	STU_TRANSFER	DEPT_CODE	STU_PHONE	PROF_NUM
STUDENT	Die,		•	2.84	No	BIOL	2134	205
continued				3.27	Yes	CIS	2256	222
Contraction Contraction				2.26	Yes	ACCT	2256	228
			2822	3.09	No	CIS	2114	222
			13.221	2.11	Yes	ENGL	2231	199
		1.000		3.15	No	ACCT	2267	228
				3.87	No	EDU	2267	311
				2.92	No	ACCT	2315	230
STU_HRS STU_CLASS STU_DOB	= Cree = Stuc = Stuc	dit hou dent cl dent da	irs ass ate	earned ification of birth	STU_GPA STU_PHONE PROF_NUM	= Grade po = 4-digit ca = Number who is th	oint average ampus phone of the profess e student's ad	extension or visor

Keys

- Consists of one or more attributes that determine other attributes
 Not the other way B is functionally depended on A
 - A determines B: STU_NUM determines STU_LNAME
- Primary key (PK)
 - an attribute (or a combination of attributes) that uniquely identifies any given entity (row)
- Composite key <u>stu_lname + stu_fname → stu_dob</u>
 - Composed of more than one attribute
- Key attribute
 - Any attribute that is part of a key

Keys

- Foreign key (FK)
 - Attribute that is linked to the other table
 - An attribute whose values match primary key values in the related table
- Referential integrity
 - FK contains a value that refers to an existing valid tuple (row) in another relation

Simple Relational Database: Example

Foreign key: VEND_CODE PROD_CODE PROD_CODE PROD_DESCRIPT												
		PROD_CODE	PRO	D_D	DES	CRIPT	PROD_PRICE	PROD	ON_HAND	VEND_C	ODE	
	+	001278-AB	Claw hamm	er	0.290.25 2000.0		\$12.95	0.08263	23		232	
123-21UUY Houselite cl GER-34256 Sledge harr		nain	sav	w, 16-in. bar	\$189.99		4		235			
		Sledge ham	mer	,16	-lb. head	\$18.63	122.2	6		231		
1	+	SRE-657UG	Rat-tail file				\$2.99		15		232	
	+	ZZX/3245Q	Steel tape,	12-ft. length		ngth	\$6.79		8		235	
						_		link				
						VEND_COD	E VEND_CO		VEND_AR	EACODE	VEND.	_PHO
bl	e	name: VENI	DOR	•	Ŧ	VEND_COD	E VEND_CO 30 Shelly K. Sr	link NTACT nithson	VEND_AR	EACODE	VEND. 555-12	_PH(
ıbl	ena	name: VENI ry key: VEN	DOR D_CODE	•	*	VEND_COD	E VEND_CO 30 Shelly K. Sr 31 James Johr	link NTACT nithson	VEND_AR 608 615	EACODE	VEND, 555-12 123-45	_PHC 234 336
ıbl in ore	ena	name: VENI ry key: VEN n key: none	DOR D_CODE	•	+++++++++++++++++++++++++++++++++++++++	VEND_COD	E VEND_CO 30 Shelly K. Sr 31 James John 32 Annelise Cr	link NTACT nithson Ison	VEND_AR 608 615 608	EACODE	VEND, 555-12 123-45 224-21	_PHC 34 36 34
ıbl 'in pre	e na ig	name: VENI ry key: VEN 3n key: none	DOR D_CODE	•	+++++++++++++++++++++++++++++++++++++++	VEND_COD 2 2 2	E VEND_CO 30 Shelly K. Sr 31 James John 32 Annelise Cr 33 Candice Wa	link NTACT nithson nson rystall allace	VEND_AR 608 615 608 904	EACODE	VEND, 555-12 123-45 224-21 342-65	_PHC 34 36 34
ıbl rim pre	e na ig	name: VENI ry key: VEN 3n key: none	DOR D_CODE	•	+ + + +	VEND_COD 2 2 2 2 2	E VEND_CO 30 Shelly K. Sr 31 James John 32 Annelise Cr 33 Candice Wa 34 Arthur Jone	link NTACT nithson nson rystall allace	VEND_AR 608 615 608 904 615	EACODE	VEND, 555-12 123-45 224-21 342-65 123-33	_PHC 234 336 34 34 34

Null Values

- No data entry
- Not permitted in primary key
- Should be avoided in other attributes
- Can represent
 - An unknown attribute value
 - A known, but missing, attribute value
 - A "not applicable" condition
- Can create problems in logic and using formulas

A Dummy Variable Value Used as a Flag

AGENT_CODE	AGENT_AREA CODE	AGENT_PHONE	AGENT_LNAME	AGENT_YTD_SALES
-99	000	000-000	None	\$0.00

Integrity Rules

ENTITY INTEGRITY	DESCRIPTION
Requirement	All primary key entries are unique, and no part of a primary key may be null.
Purpose	Guarantees that each entity will have a unique identity and ensures that foreign key values can properly reference primary key values.
Example	No invoice can have a duplicate number, nor can it be null. In short, all invoices are uniquely identified by their invoice number.
REFERENTIAL INTEGRITY	DESCRIPTION
Requirement	A foreign key may have either a null entry—as long as it is not a part of its table's primary key—or an entry that matches the primary key value in a table to which it is related. (Every non-null foreign key value <i>must</i> reference an <i>existing</i> primary key value.)
Purpose	Makes it possible for an attribute NOT to have a corresponding value, but it will be impossible to have an invalid entry. The enforcement of the referential integrity rule makes it impossible to delete a row in one table whose primary key has mandatory matching foreign key values in another table.
Example	A customer might not (yet) have an assigned sales representative (number), but it will be impossible to have an invalid sales representative (number).

Data Dictionary

- Used to provide detailed accounting of all tables found within the user/designer-created database
- Contains (at least) all the attribute names and characteristics for each table in the system
- Contains metadata—data about data
- Sometimes described as "the database designer's database" because it records the design decisions about tables and their structures

A Sample Data Dictionary

TABLE NAME	ATTRIBUTE NAME	CONTENTS	ТҮРЕ	FORMAT	RANGE	REQUIRED	PK Or FK	FK REFERENCED TABLE
CUSTOMER	CUS_CODE CU_LNAME CUS_FNAME CUS_INITIAL CUS_RENEW_DATE AGENT_CODE	Customer acct. code Customer last name Customer first name Customer initial Customer insurance renewal date Agent code	CHAR(5) VCCHAR(20) VCHAR(20 CHAR(1) DATE CHAR(3)	99999 Xxxxxxx Xxxxxxx X dd-mmm-yyyy 999	10000-99999 100-999	Y Y Y	PK FK	AGENT
AGENT	AGENT_CODE AGENT_AREACODE AGENT_PHONE AGENT_LNAME AGENT_YTD_SALES	Agent code Agent area code Agent telephone number Agent last name Agent year-to-date sales	CHAR(3) CHAR(3) CHAR(8) VCHAR(20) NUMBER(9,2)	999 999 999-9999 Xxxxxxx 9,999,999.99	0.00-9,999,999.99	Y Y Y Y Y	РК	

FK	= Foreign key
РК	= Primary key
CHAR	= Fixed character length data (1-255 characters)
VARCHAR	= Variable character length data (1-2,000 characters)
NUMBER	= Numeric data. NUMBER(9,2) is used to specify numbers with two decimal places and up to nine digits, including the decimal places. Some RDBMSs permit the use of a MONEY or CURRENCY data type.

Relationships within the Relational Database

- 1:M relationship
 - Relational modeling ideal
 - Should be the norm in any relational database design
- M:N relationships
 - Must be avoided because they lead to data redundancies
- 1:1 relationship
 - Rare in most relational database designs

The 1:1 Relationship Between PROFESSOR and DEPARTMENT



The Implemented 1:1 Relationship Between PROFESSOR and DEPARTMENT

Database name: Ch03_TinyCollege

Table name: PROFESSOR Primary key: EMP_NUM Foreign key: DEPT_CODE

	EMP_NUM	DEPT_CODE	PROF_OFFICE	PROF_EXTENSION	PROF_HIGH_DEGREE
	103	HIST	DRE 156	6783	Ph.D.
	104	ENG	DRE 102	5561	MA
	105	ACCT	KLR 229D	8665	Ph.D.
	106	MKT/MGT	KLR 126	3899	Ph.D.
	110	BIOL	AAK 160	3412	Ph.D.
	114	ACCT	KLR 211	4436	Ph.D.
	155	MATH	AAK 201	4440	Ph.D.
	160	ENG	DRE 102	2248	Ph.D.
	162	CIS	KLR 203E	2359	Ph.D.
	191	MKT/MGT	KLR 409B	4016	DBA
NA.	195	PSYCH	AAK 297	3550	Ph.D.
	209	CIS	KLR 333	3421	Ph.D.
	228	CIS	KLR 300	3000	Ph.D.
The second	297	MATH	AAK 194	1145	Ph.D.
	299	ECON/FIN	KLR 284	2851	Ph.D.
	301	ACCT	KLR 244	4683	Ph.D.
	335	ENG	DRE 208	2000	Ph.D.
	342	SOC	BBG 208	5514	Ph.D.
	387	BIOL	AAK 230	8665	Ph.D.
	401	HIST	DRE 156	6783	MA
	425	ECON/FIN	KLR 284	2851	MBA
	435	ART	BBG 185	2278	Ph.D.

Table name: DEPARTMENT Primary key: DEPT_CODE Foreign key: EMP_NUM

For	eig	gn key: EMP	_NUM				
		DEPT_CODE	DEPT_NAME	SCHOOL_CODE	EMP_NUM	DEPT_ADDRESS	DEPT_EXTENSION
	+	ACCT	Accounting	BUS	114	KLR 211, Box 52	3119
	+	ART	Fine Arts	A&SCI	435	BBG 185, Box 128	2278
	+	BIOL	Biology	A&SCI	387	AAK 230, Box 415	4117
	-	CIS	Computer Info. Systems	BUS	209	KLR 333, Box 56	3245
1.1.1	+	ECON/FIN	Economics/Finance	BUS	299	KLR 284, Box 63	3126
	+	ENG	English	A&SCI	160	DRE 102, Box 223	1004
	+	HIST	History	A&SCI	103	DRE 156, Box 284	1867
	+	MATH	Mathematics	A&SCI	297	AAK 194, Box 422	4234
	-	MKT/MGT	Marketing/Management	BUS	106	KLR 126, Box 55	3342
	+	PSYCH	Psychology	A&SCI	195	AAK 297, Box 438	4110
	+	SOC	Sociology	A&SCI	342	BBG 208, Box 132	2008

The 1:M Relationship Between PAINTER and PAINTING



The Implemented 1:M Relationship Between PAINTER and PAINTING

Table name: PAINTER Primary key: PAINTER_NUM Foreign key: none

Database name: Ch03_Museum

	PAINTER_NUM	PAINTER_LNAME	PAINTER_FNAME	PAINTER_INITIAL
	123	Ross	Georgette	Р
+	126	Itero	Julio	G

Table name: PAINTING Primary key: PAINTING_NUM Foreign key: PAINTER_NUM							
	PAINTING_NUM	PAINTING_TITLE	PAINTER_NUM				
	1338	Dawn Thunder	123				
	1339	Vanilla Roses To Nowhere	123				
	1340	Tired Flounders	126				
1341 Hasty Exit 123							
	1342	Plastic Paradise	126				

The 1:M Relationship Between COURSE and CLASS



The Implemented 1:M Relationship Between COURSE and CLASS

Table name: COURSE Primary key: CRS_CODE Foreign key: none

Database name: Ch03_TinyCollege

	CRS_CODE	DEPT_CODE	CRS_DESCRIPTION	CRS_CREDIT
•	ACCT-211	ACCT	Accounting I	3
+	ACCT-212	ACCT	Accounting II	3
+	CIS-220	CIS	Intro. to Microcomputing	3
+	CIS-420	CIS	Database Design and Implementation	4
+	QM-261	CIS	Intro. to Statistics	3
+	QM-362	CIS	Statistical Applications	4

Table name: CLASS Primary key: CLASS_CODE Foreign key: CRS_CODE

	CLASS_CODE	CRS_CODE	CLASS_SECTION	CLASS_TIME	CLASS_ROOM	PROF_NUM
+	10012	ACCT-211	1	MVVF 8:00-8:50 a.m.	BUS311	105
+	10013	ACCT-211	2	MVVF 9:00-9:50 a.m.	BUS200	105
+	10014	ACCT-211	3	TTh 2:30-3:45 p.m.	BUS252	342
+	10015	ACCT-212	1	MVVF 10:00-10:50 a.m.	BUS311	301
+	10016	ACCT-212	2	Th 6:00-8:40 p.m.	BUS252	301
+	10017	CIS-220	1	MVVF 9:00-9:50 a.m.	KLR209	228
+	10018	CIS-220	2	MVVF 9:00-9:50 a.m.	KLR211	114
+	10019	CIS-220	3	MVVF 10:00-10:50 a.m.	KLR209	228
+	10020	CIS-420	1	VV 6:00-8:40 p.m.	KLR209	162
+	10021	QM-261	1	MVVF 8:00-8:50 a.m.	KLR200	114
+	10022	QM-261	2	TTh 1:00-2:15 p.m.	KLR200	114
+	10023	QM-362	1	MVVF 11:00-11:50 a.m.	KLR200	162
+	10024	QM-362	2	TTh 2:30-3:45 p.m.	KLR200	162

The M:N Relationship

- Can be implemented by breaking it up to produce a set of 1:M relationships
- Can avoid problems inherent to M:N relationship by creating a composite entity or bridge entity

The ERD's M:N Relationship Between STUDENT and CLASS



Sample Student Enrollment Data

STUDENT'S LAST NAME	SELECTED CLASSES
Bowser	Accounting 1, ACCT-211, code 10014 Intro. to Microcomputing, CIS-220, code 10018 Intro. To Statistics, QM-261, code 10021
Smithson	Accounting 1, ACCT-211, code 10014 Intro. to Microcomputing, CIS-220, code 10018 Intro. To Statistics, QM-261, code 10021

The M:N Relationship Between STUDENT and CLASS

Table name: STUDENT Primary key: STU_NUM Foreign key: none

Database name: Ch03_CollegeTry

1223	STU_NUM	STU_LNAME	CLASS_CODE
	321452	Bowser	10014
	321452	Bowser	10018
	321452	Bowser	10021
	324257	Smithson	10014
	324257	Smithson	10018
	324257	Smithson	10021

Table name: CLASS Primary key: CLASS_CODE Foreign key: STU_NUM

	CLASS_CODE	STU_NUM	CRS_CODE	CLASS_SECTION	CLASS_TIME	CLASS_ROOM	PROF_NUM
•	10014	321452	ACCT-211	3	TTh 2:30-3:45 p.m.	BUS252	342
and.	10014	324257	ACCT-211	3	TTh 2:30-3:45 p.m.	BUS252	342
THE.	10018	321452	CIS-220	2	MVVF 9:00-9:50 a.m.	KLR211	114
	10018	324257	CIS-220	2	MVVF 9:00-9:50 a.m.	KLR211	114
	10021	321452	QM-261	1	MVVF 8:00-8:50 a.m.	KLR200	114
	10021	324257	QM-261	1	MVVF 8:00-8:50 a.m.	KLR200	114

Linking Table

- Implementation of a composite entity
- Yields required M:N to 1:M conversion
- Composite entity table must contain at least the primary keys of original tables
- Linking table contains multiple occurrences of the foreign key values
- Additional attributes may be assigned as needed

Converting the M:N Relationship into Two 1:M Relationships

Table name: STUDENT Primary key: STU_NUM Foreign key: none

		STU_NUM	STU_LNAME
•	+	321452	Bowser
	+	324257	Smithson

Database name: Ch03_CollegeTry2

Table name: ENROLL Primary key: CLASS_CODE + STU_NUM Foreign key: CLASS_CODE, STU_NUM

	CLASS_CODE	STU_NUM	ENROLL_GRADE
•	10014	321452	С
	10014	324257	B
	10018	321452	A
	10018	324257	в
	10021	321452	С
	10021	324257	С

Composite Entity Table for Linking

Table name: CLASS Primary key: CLASS_CODE Foreign key: CRS_CODE

	CLASS_CODE	CRS_CODE	CLASS_SECTION	CLASS_TIME	CLASS_ROOM	PROF_NUM
+	10014	ACCT-211	3	TTh 2:30-3:45 p.m.	BUS252	342
+	10018	CIS-220	2	MVVF 9:00-9:50 a.m.	KLR211	114
+	10021	QM-261	1	MVVF 8:00-8:50 a.m.	KLR200	114

Changing the M:N Relationship to Two 1:M Relationships



The Expanded Entity Relationship Model



The Relational Schema for the Ch03_TinyCollege Database



Data Redundancy Revisited

- Data redundancy leads to data anomalies
 - Such anomalies can destroy database effectiveness
- Foreign keys
 - Control data redundancies by using common attributes shared by tables
 - Crucial to exercising data redundancy control
- Sometimes, data redundancy is necessary

A Small Invoicing System

Table name: CUSTOMER Primary key: CUS_CODE Foreign key: none

Database name: Ch03_SaleCo

		CUS_CODE	CUS_LNAME	CUS_FNAME	CUS_INITIAL	CUS_AREACODE	CUS_PHONE
	+	10010	Ramas	Alfred	A	615	844-2573
	+	10011	Dunne	Leona	к	713	894-1238
	(+)	10012	Smith	Kathy	~	615	894-2285
	+	10013	Olovvski	Paul	F	615	894-2180
1	[+]	10014	Orlando	Myron	All the second se	615	222-1672
		10015	O'Brian	Amy	B	713	442-3381
	-	10016	Brown	James	G	615	297-1228
		10017	vvilliams	George	the second second second	615	290-2556
	-	10018	Farriss	Anne	G	713	382-7185
	[+]	10019	Smith	Olette	к	615	297-3809

Table name: INVOICE Primary key: INV_NUMBER Foreign key: CUS_CODE

	INV_NUMBER	CUS_CODE	INV_DATE
+	1001	10014	08-Mar-04
+	1002	10011	08-Mar-04
(+)	1003	10012	08-Mar-04
-	1004	10011	09-Mar-04

Table name: LINE Primary key: INV_NUMBER + LINE_NUMBER Foreign keys: INV_NUMBER, PROD_CODE

	INV_NUMBER	LINE_NUMBER	PROD_CODE	LINE_UNITS	LINE_PRICE
	1001	1	123-21UUY	1	\$189.99
	1001	2	SRE-657UG	3	\$2.99
	1002	1	QER-34256	2	\$18.63
	1003	1	ZZX/3245Q	1	\$6.79
14.37	1003	2	SRE-657UG	1	\$2.99
	1003	3	001278-AB	1	\$12.95
	1004	1	001278-AB	1	\$12.95
1 Sector	1004	2	SRE-657UG	2	\$2.99

Table name: PRODUCT Primary key: PROD_CODE Foreign key: none

		PROD_CODE	PROD_DESCRIPT	PROD_PRICE	PROD_ON_HAND	VEND_CODE
	(+)	001278-AB	Clavv hammer	\$12.95	23	232
	(+)	123-21UUY	Houselite chain saw, 16-in. bar	\$189.99	4	235
1.000		QER-34256	Sledge hammer, 16-lb. head	\$18.63	6	231
	(+)	SRE-657UG	Rat-tail file	\$2.99	15	232
10000	[+]	ZZX/3245Q	Steel tape, 12-ft. length	\$6.79	8	235

The Relational Schema for the Invoicing System



Relational Database Operators

Relational algebra

- Defines theoretical way of manipulating table contents using relational operators:
 - SELECT UNION
 - PROJECT DIFFERENCE
 - JOIN PRODUCT
 - INTERSECT DIVIDE
- Use of relational algebra operators on existing tables (relations) produces new relations

Union

Combines all rows from two tables, excluding duplicate rows

 Tables must have the same attribute characteristics (union-compatible)

	P_CODE	P_DESCRIPT	PRICE			P_CODE	P_DESCRIPT	PRICE
•	123456	Flashlight	\$5.26	UNION	•	345678	Microwave	\$160.00
	123457	Lamp	\$25.15			345679	Dishwasher	\$500.00
	123458	Box Fan	\$10.99					
	213345	9v battery	\$1.92				olde	
	254467	100W bulb	\$1.47			y,	eius	
	311452	Powerdrill	\$34.99					
						P_CODE	P_DESCRIPT	PRICE
					•	123456	Flashlight	\$5.26
						123457	Lamp	\$25.15
						123458	Box Fan	\$10.99
						213345	9v battery	\$1.92
						254467	100W bulb	\$1.47
						311452	Powerdrill	\$34.99
						345678	Microwave	\$160.00
						345679	Dishwasher	\$500.00

Intersect

Yields only the rows that appear in both tablesTables must be union-compatible



Difference

 Yields all rows in one table not found in the other table—that is, it subtracts one table from the other

Tables must be union-compatible



Product

Yields all possible pairs of rows from two tables

Also known as the Cartesian product

	P_CODE	P_DESCRIPT	PRICE					STORE	E A	ISLE S	SHELF	1		
•	123456	Flashlight	\$5.26	PRO	DUC	г		23	W	. 5	;			
10000	123457	Lamp	\$25.15					24	ĸ	9)			
10000	123458	Box Fan	\$10.99					25	Z	6	;			
	213345	9v battery	\$1.92											
	254467	100W bulb	\$1.47											
	311452	Powerdrill	\$34.99					yield	s					
				Г	P_	CODE	P_D	ESCRIP	T	PRICE	STO	DRE	AISLE	SHELF
					123	3456	Flas	hlight		\$5.26	3 23		w	5
					123	3456	Flas	hlight		\$5.26	3 24		к	9
					123	3456	Flas	hlight		\$5.26	25		Z	6
					123	3457	Lam	р		\$25.15	5 23		w.	5
					123	3457	Lam	р		\$25.15	5 24		к	9
					123	3457	Lam	р		\$25.15	5 25		Z	6
					123	3458	Box	Fan		\$10.99	23		w	5
					123	3458	Box	Fan		\$10.99	24		к	9
					123	3458	Box	Fan		\$10.99	25		Z	6
					213	3345	9v b	attery		\$1.92	23		w	5
					213	3345	9v b	attery		\$1.92	24		к	9
				L	213	3345	9v b	attery		\$1.92	25		Z	6
				L	311	452	Pow	erdrill		\$34.99	23		w	5
				L	311	452	Pow	erdrill		\$34.99	3 24		к	9
				L	311	452	Pow	erdrill		\$34.99	25		Z	6
					254	467	100	/V bulb		\$1.47	23		w	5
					254	467	100	/V bulb		\$1.47	24		к	9
					254	467	100	AV bulb		\$1.47	25		Z	6

Select

- Yields values for all rows found in a table (SELECT ALL)
- Can be used to list either all row values or it can yield only those row values that match a specified criterion
- Yields a horizontal subset of a table



Project

- Yields all values for selected attributes
- Yields a vertical subset of a table



Join

- Allows us to combine information from two or more tables
- Real power behind the relational database, allowing the use of independent tables linked by common attributes

Two Tables That Will Be Used in Join Examples

Ia	ble name: C	USTOMER		
	CUS_CODE	CUS_LNAME	CUS_ZIP	AGENT_CODE
•	1132445	Walker	32145	231
	1217782	Adares	32145	125
	1312243	Rakowski	34129	167
	1321242	Rodriguez	37134	125
	1542311	Smithson	37134	421
	1657399	Vanloo	32145	231

Table name: AGENT

	AGENT_CODE	AGENT_PHONE
•	125	6152439887
	167	6153426778
	231	6152431124
	333	9041234445

Natural Join

- Links tables by selecting only rows with common values in their common attribute(s)
- Result of a three-stage process:
 - 1. PRODUCT of the tables is created
 - 2. SELECT is performed on Step 1 output to yield only the rows for which the AGENT_CODE values are equal
 - Common column(s) are called join column(s)
 - 3. PROJECT is performed on Step 2 results to yield a single copy of each attribute, thereby eliminating duplicate columns

Natural Join, Step 1: PRODUCT

	CUS_CODE	CUS_LNAME	CUS_ZIP	CUSTOMER.AGENT_CODE	AGENT.AGENT_CODE	AGENT_PHONE
	132445	Walker	32145	231	125	6152439887
1	1132445	Walker	32145	231	167	6153426778
1	1132445	Walker	32145	231	231	6152431124
1	1132445	Walker	32145	231	333	9041234445
1	1217782	Adares	32145	125	125	6152439887
1	1217782	Adares	32145	125	167	6153426778
1	1217782	Adares	32145	125	231	6152431124
1	1217782	Adares	32145	125	333	9041234445
1	1312243	Rakowski	34129	167	125	6152439887
1	1312243	Rakowski	34129	167	167	6153426778
1	1312243	Rakowski	34129	167	231	6152431124
1	1312243	Rakowski	34129	167	333	9041234445
1	1321242	Rodriguez	37134	125	125	6152439887
1	1321242	Rodriguez	37134	125	167	6153426778
1	1321242	Rodriguez	37134	125	231	6152431124
1	1321242	Rodriguez	37134	125	333	9041234445
1	1542311	Smithson	37134	421	125	6152439887
1	1542311	Smithson	37134	421	167	6153426778
1	1542311	Smithson	37134	421	231	6152431124
1	1542311	Smithson	37134	421	333	9041234445
1	1657399	Vanloo	32145	231	125	6152439887
1	1657399	Vanloo	32145	231	167	6153426778
1	1657399	Vanloo	32145	231	231	6152431124
1	1657399	Vanloo	32145	231	333	9041234445

Natural Join, Step 2: SELECT

Γ	CUS_CODE	CUS_LNAME	CUS_ZIP	CUSTOMER.AGENT_CODE	AGENT.AGENT_CODE	AGENT_PHONE
•	1217782	Adares	32145	125	125	6152439887
	1321242	Rodriguez	37134	125	125	6152439887
	1312243	Rakowski	34129	167	167	6153426778
	1132445	Walker	32145	231	231	6152431124
	1657399	Vanloo	32145	231	231	6152431124

Natural Join, Step 3: PROJECT

and the second se	C03_CODE	CUS_ENAME	CUS_ZIP	AGENT_CODE	AGENT_PHONE
	1217782	Adares	32145	125	6152439887
	1321242	Rodriguez	37134	125	6152439887
	1312243	Rakowski	34129	167	6153426778
	1132445	Walker	32145	231	6152431124
4	1657399	Vanloo	32145	231	6152431124

Natural Join (continued)

- Final outcome yields table that
 - Does not include unmatched pairs
 - Provides only copies of matches
- If no match is made between the table rows,
 - the new table does not include the unmatched row
- The column on which we made the JOIN—that is, AGENT_CODE—occurs only once in the new table

Outer Join

- Matched pairs are retained and any unmatched values in other table are left null
- In outer join for tables CUSTOMER and AGENT, two scenarios are possible:
 - Left outer join
 - Yields all rows in CUSTOMER table, including those that do not have a matching value in the AGENT table
 - Right outer join
 - Yields all rows in AGENT table, including those that do not have matching values in the CUSTOMER table

Left Outer Join

▶ 1217782 Adares 32145 125 1321242 Rodriguez 37134 125 13121242 Relevandei 34129 167	6152439887 6152439887
1321242 Rodriguez 37134 125	6152439887
1212242 Bekeweki 24120 467	
1312243 Rakuwski 34129 167	6153426778
1132445 Walker 32145 231	6152431124
1657399 Vanloo 32145 231	6152431124
1542311 Smithson 37134 421	

Right Outer Join

		1		
 CUS_CODE	CUS_LNAME	CUS_ZIP	AGENT_CODE	AGENT_PHONE
1217782	Adares	32145	125	6152439887
1321242	Rodriguez	37134	125	6152439887
1312243	Rakowski	34129	167	6153426778
1132445	Walker	32145	231	6152431124
1657399	Vanloo	32145	231	6152431124
			333	9041234445

Outer Join

- Especially useful to detect referential integrity problems
 - Example: when converting large spreadsheets or other non-relational data into relational tables
- Left and Right refer to the order in which tables are listed in the SQL command

DIVIDE

 DIVIDE requires the use of one single-column table and one two-column table



Find common values of A and B