
employee (*person_name*, *street*, *city*)
works (*person_name*, *company_name*, *salary*)
company (*company_name*, *city*)

Figure 2.17 Employee database.

Practice Exercises

- 2.1 Consider the employee database of Figure 2.17. What are the appropriate primary keys?
- 2.2 Consider the foreign-key constraint from the *dept_name* attribute of *instructor* to the *department* relation. Give examples of inserts and deletes to these relations that can cause a violation of the foreign-key constraint.
- 2.3 Consider the *time_slot* relation. Given that a particular time slot can meet more than once in a week, explain why *day* and *start_time* are part of the primary key of this relation, while *end_time* is not.
- 2.4 In the instance of *instructor* shown in Figure 2.1, no two instructors have the same name. From this, can we conclude that *name* can be used as a superkey (or primary key) of *instructor*?
- 2.5 What is the result of first performing the Cartesian product of *student* and *advisor*, and then performing a selection operation on the result with the predicate $s_id = ID$? (Using the symbolic notation of relational algebra, this query can be written as $\sigma_{s_id=ID}(student \times advisor)$.)
- 2.6 Consider the employee database of Figure 2.17. Give an expression in the relational algebra to express each of the following queries:
 - a. Find the name of each employee who lives in city “Miami”.
 - b. Find the name of each employee whose salary is greater than \$100000.
 - c. Find the name of each employee who lives in “Miami” and whose salary is greater than \$100000.
- 2.7 Consider the bank database of Figure 2.18. Give an expression in the relational algebra for each of the following queries:
 - a. Find the name of each branch located in “Chicago”.
 - b. Find the ID of each borrower who has a loan in branch “Downtown”.

```

branch(branch_name, branch_city, assets)
customer (ID, customer_name, customer_street, customer_city)
loan (loan_number, branch_name, amount)
borrower (ID, loan_number)
account (account_number, branch_name, balance)
depositor (ID, account_number)

```

Figure 2.18 Bank database.

2.8 Consider the employee database of Figure 2.17. Give an expression in the relational algebra to express each of the following queries:

- a. Find the ID and name of each employee who does not work for “BigBank”.
- b. Find the ID and name of each employee who earns at least as much as every employee in the database.

2.9 The **division operator** of relational algebra, “ \div ”, is defined as follows. Let $r(R)$ and $s(S)$ be relations, and let $S \subseteq R$; that is, every attribute of schema S is also in schema R . Given a tuple t , let $t[S]$ denote the projection of tuple t on the attributes in S . Then $r \div s$ is a relation on schema $R - S$ (that is, on the schema containing all attributes of schema R that are not in schema S). A tuple t is in $r \div s$ if and only if both of two conditions hold:

- t is in $\Pi_{R-S}(r)$
- For every tuple t_s in s , there is a tuple t_r in r satisfying both of the following:
 - a. $t_r[S] = t_s[S]$
 - b. $t_r[R - S] = t$

Given the above definition:

- a. Write a relational algebra expression using the division operator to find the IDs of all students who have taken all Comp. Sci. courses. (Hint: project *takes* to just ID and *course_id*, and generate the set of all Comp. Sci. *course_ids* using a select expression, before doing the division.)
- b. Show how to write the above query in relational algebra, without using division. (By doing so, you would have shown how to define the division operation using the other relational algebra operations.)

Exercises

- 2.10** Describe the differences in meaning between the terms *relation* and *relation schema*.
- 2.11** Consider the *advisor* relation shown in the schema diagram in Figure 2.9, with *s_id* as the primary key of *advisor*. Suppose a student can have more than one advisor. Then, would *s_id* still be a primary key of the *advisor* relation? If not, what should the primary key of *advisor* be?
- 2.12** Consider the bank database of Figure 2.18. Assume that branch names and customer names uniquely identify branches and customers, but loans and accounts can be associated with more than one customer.
- What are the appropriate primary keys?
 - Given your choice of primary keys, identify appropriate foreign keys.
- 2.13** Construct a schema diagram for the bank database of Figure 2.18.
- 2.14** Consider the employee database of Figure 2.17. Give an expression in the relational algebra to express each of the following queries:
- Find the ID and name of each employee who works for “BigBank”.
 - Find the ID, name, and city of residence of each employee who works for “BigBank”.
 - Find the ID, name, street address, and city of residence of each employee who works for “BigBank” and earns more than \$10000.
 - Find the ID and name of each employee in this database who lives in the same city as the company for which she or he works.
- 2.15** Consider the bank database of Figure 2.18. Give an expression in the relational algebra for each of the following queries:
- Find each loan number with a loan amount greater than \$10000.
 - Find the ID of each depositor who has an account with a balance greater than \$6000.
 - Find the ID of each depositor who has an account with a balance greater than \$6000 at the “Uptown” branch.
- 2.16** List two reasons why null values might be introduced into a database.
- 2.17** Discuss the relative merits of imperative, functional, and declarative languages.
- 2.18** Write the following queries in relational algebra, using the university schema.
- Find the ID and name of each instructor in the Physics department.

- b. Find the ID and name of each instructor in a department located in the building “Watson”.
- c. Find the ID and name of each student who has taken at least one course in the “Comp. Sci.” department.
- d. Find the ID and name of each student who has taken at least one course section in the year 2018.
- e. Find the ID and name of each student who has not taken any course section in the year 2018.

Further Reading

E. F. Codd of the IBM San Jose Research Laboratory proposed the relational model in the late 1960s ([Codd (1970)]). In that paper, Codd also introduced the original definition of relational algebra. This work led to the prestigious ACM Turing Award to Codd in 1981 ([Codd (1982)]).

After E. F. Codd introduced the relational model, an expansive theory developed around the relational model pertaining to schema design and the expressive power of various relational languages. Several classic texts cover relational database theory, including [Maier (1983)] (which is available free, online), and [Abiteboul et al. (1995)].

Codd’s original paper inspired several research projects that were formed in the mid to late 1970s with the goal of constructing practical relational database systems, including System R at the IBM San Jose Research Laboratory, Ingres at the University of California at Berkeley, and Query-by-Example at the IBM T. J. Watson Research Center. The Oracle database was developed commercially at the same time.

Many relational database products are now commercially available. These include IBM’s DB2 and Informix, Oracle, Microsoft SQL Server, and Sybase and HANA from SAP. Popular open-source relational database systems include MySQL and PostgreSQL. Hive and Spark are widely used systems that support parallel execution of queries across large numbers of computers.

Bibliography

[Abiteboul et al. (1995)] S. Abiteboul, R. Hull, and V. Vianu, *Foundations of Databases*, Addison Wesley (1995).

[Codd (1970)] E. F. Codd, “A Relational Model for Large Shared Data Banks”, *Communications of the ACM*, Volume 13, Number 6 (1970), pages 377–387.

[Codd (1982)] E. F. Codd, “The 1981 ACM Turing Award Lecture: Relational Database: A Practical Foundation for Productivity”, *Communications of the ACM*, Volume 25, Number 2 (1982), pages 109–117.