

CHAPTER 6



Formal Relational Query Languages

Practice Exercises

- 6.1 Write the following queries in relational algebra, using the university schema.
- Find the titles of courses in the Comp. Sci. department that have 3 credits.
 - Find the IDs of all students who were taught by an instructor named Einstein; make sure there are no duplicates in the result.
 - Find the highest salary of any instructor.
 - Find all instructors earning the highest salary (there may be more than one with the same salary).
 - Find the enrollment of each section that was offered in Autumn 2009.
 - Find the maximum enrollment, across all sections, in Autumn 2009.
 - Find the sections that had the maximum enrollment in Autumn 2009

Answer:

- $\Pi_{title}(\sigma_{dept_name = 'Comp. Sci' \wedge credits=3}(course))$
- $\Pi_{ID}(\sigma_{IID = 'Einstein'}(takes \bowtie \rho_{FI}(IID, course_id, section_id, semester, year)teaches))$
Assuming the set version of the relational algebra is used, there is no need to explicitly remove duplicates. If the multiset version is used, the grouping operator can be used without any aggregation to remove duplicates. For example given relation $r(A, B)$ possibly containing duplicates, $_{A,B}\mathcal{G}(r)$ would return a duplicate free version of the relation.
- $\mathcal{G}_{\max(salary)}(instructor)$

- d. $instructor \bowtie (\mathcal{G}_{\max(salary)} \text{ as } salary(instructor))$
 Note that the above query renames the maximum salary as salary, so the subsequent natural join outputs only instructors with that salary.
- e. $course_id, section_id \mathcal{G}_{\text{count}(*)} \text{ as } enrollment(\sigma_{year=2009 \wedge semester=Autumn}(takes))$
- f. $t1 \leftarrow course_id, section_id \mathcal{G}_{\text{count}(*)} \text{ as } enrollment(\sigma_{year=2009 \wedge semester=Autumn}(takes))$
 $result = \mathcal{G}_{\max(enrollment)}(t1)$
- g. $t2 \leftarrow \mathcal{G}_{\max(enrollment)} \text{ as } enrollment(t1)$
 where $t1$ is as defined in the previous part of the question.
 $result = t1 \bowtie t2$

6.2 Consider the relational database of Figure 6.22, where the primary keys are underlined. Give an expression in the relational algebra to express each of the following queries:

- Find the names of all employees who live in the same city and on the same street as do their managers.
- Find the names of all employees in this database who do not work for “First Bank Corporation”.
- Find the names of all employees who earn more than every employee of “Small Bank Corporation”.

Answer:

- $\Pi_{person_name} ((employee \bowtie manages) \bowtie (manager_name = employee2.person_name \wedge employee.street = employee2.street \wedge employee.city = employee2.city)(\rho_{employee2}(employee)))$
- The following solutions assume that all people work for exactly one company. If one allows people to appear in the database (e.g. in *employee*) but not appear in *works*, the problem is more complicated. We give solutions for this more realistic case later.
 $\Pi_{person_name} (\sigma_{company_name \neq \text{“First Bank Corporation”}}(works))$
 If people may not work for any company:
 $\Pi_{person_name}(employee) - \Pi_{person_name}(\sigma_{(company_name = \text{“First Bank Corporation”})}(works))$
- $\Pi_{person_name}(works) - (\Pi_{works.person_name}(works) \bowtie (\sigma_{(works.salary \leq works2.salary \wedge works2.company_name = \text{“Small Bank Corporation”})}(\rho_{works2}(works))))$

6.3 The natural outer-join operations extend the natural-join operation so that tuples from the participating relations are not lost in the result of the join.

Describe how the theta-join operation can be extended so that tuples from the left, right, or both relations are not lost from the result of a theta join.

Answer:

- a. The left outer theta join of $r(R)$ and $s(S)$ ($r \bowtie_{\theta}^{\leftarrow} s$) can be defined as $(r \bowtie_{\theta} s) \cup ((r - \Pi_R(r \bowtie_{\theta} s)) \times (null, null, \dots, null))$
The tuple of nulls is of size equal to the number of attributes in S .
 - b. The right outer theta join of $r(R)$ and $s(S)$ ($r \bowtie_{\theta}^{\rightarrow} s$) can be defined as $(r \bowtie_{\theta} s) \cup ((null, null, \dots, null) \times (s - \Pi_S(r \bowtie_{\theta} s)))$
The tuple of nulls is of size equal to the number of attributes in R .
 - c. The full outer theta join of $r(R)$ and $s(S)$ ($r \bowtie_{\theta}^{\leftarrow\rightarrow} s$) can be defined as $(r \bowtie_{\theta} s) \cup ((null, null, \dots, null) \times (s - \Pi_S(r \bowtie_{\theta} s))) \cup ((r - \Pi_R(r \bowtie_{\theta} s)) \times (null, null, \dots, null))$
The first tuple of nulls is of size equal to the number of attributes in R , and the second one is of size equal to the number of attributes in S .
- 6.4 (Division operation):** 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 . 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.)

Answer:

- a. $\Pi_{ID}(\Pi_{ID, course_id}(takes) \div \Pi_{course_id}(\sigma_{dept_name='Comp. Sci'}(course)))$
- b. The required expression is as follows:
 $r \leftarrow \Pi_{ID, course_id}(takes)$

$$s \leftarrow \Pi_{course_id}(\sigma_{dept_name='Comp. Sci'}(course))$$

$$\Pi_{ID}(takes) - \Pi_{ID}((\Pi_{ID}(takes) \times s) - r)$$

In general, let $r(R)$ and $s(S)$ be given, with $S \subseteq R$. Then we can express the division operation using basic relational algebra operations as follows:

$$r \div s = \Pi_{R-S}(r) - \Pi_{R-S}((\Pi_{R-S}(r) \times s) - \Pi_{R-S,S}(r))$$

To see that this expression is true, we observe that $\Pi_{R-S}(r)$ gives us all tuples t that satisfy the first condition of the definition of division. The expression on the right side of the set difference operator

$$\Pi_{R-S}((\Pi_{R-S}(r) \times s) - \Pi_{R-S,S}(r))$$

serves to eliminate those tuples that fail to satisfy the second condition of the definition of division. Let us see how it does so. Consider $\Pi_{R-S}(r) \times s$. This relation is on schema R , and pairs every tuple in $\Pi_{R-S}(r)$ with every tuple in s . The expression $\Pi_{R-S,S}(r)$ merely reorders the attributes of r .

Thus, $(\Pi_{R-S}(r) \times s) - \Pi_{R-S,S}(r)$ gives us those pairs of tuples from $\Pi_{R-S}(r)$ and s that do not appear in r . If a tuple t_j is in

$$\Pi_{R-S}((\Pi_{R-S}(r) \times s) - \Pi_{R-S,S}(r))$$

then there is some tuple t_s in s that does not combine with tuple t_j to form a tuple in r . Thus, t_j holds a value for attributes $R - S$ that does not appear in $r \div s$. It is these values that we eliminate from $\Pi_{R-S}(r)$.

6.5 Let the following relation schemas be given:

$$R = (A, B, C)$$

$$S = (D, E, F)$$

Let relations $r(R)$ and $s(S)$ be given. Give an expression in the tuple relational calculus that is equivalent to each of the following:

- $\Pi_A(r)$
- $\sigma_{B=17}(r)$
- $r \times s$
- $\Pi_{A,F}(\sigma_{C=D}(r \times s))$

Answer:

- $\{t \mid \exists q \in r (q[A] = t[A])\}$
- $\{t \mid t \in r \wedge t[B] = 17\}$
- $\{t \mid \exists p \in r \exists q \in s (t[A] = p[A] \wedge t[B] = p[B] \wedge t[C] = p[C] \wedge t[D] = q[D] \wedge t[E] = q[E] \wedge t[F] = q[F])\}$

- d. $\{t \mid \exists p \in r \exists q \in s (t[A] = p[A] \wedge t[F] = q[F] \wedge p[C] = q[D])\}$
- 6.6 Let $R = (A, B, C)$, and let r_1 and r_2 both be relations on schema R . Give an expression in the domain relational calculus that is equivalent to each of the following:
- $\Pi_A(r_1)$
 - $\sigma_{B=17}(r_1)$
 - $r_1 \cup r_2$
 - $r_1 \cap r_2$
 - $r_1 - r_2$
 - $\Pi_{A,B}(r_1) \bowtie \Pi_{B,C}(r_2)$

Answer:

- $\{ \langle t \rangle \mid \exists p, q (\langle t, p, q \rangle \in r_1) \}$
 - $\{ \langle a, b, c \rangle \mid \langle a, b, c \rangle \in r_1 \wedge b = 17 \}$
 - $\{ \langle a, b, c \rangle \mid \langle a, b, c \rangle \in r_1 \vee \langle a, b, c \rangle \in r_2 \}$
 - $\{ \langle a, b, c \rangle \mid \langle a, b, c \rangle \in r_1 \wedge \langle a, b, c \rangle \in r_2 \}$
 - $\{ \langle a, b, c \rangle \mid \langle a, b, c \rangle \in r_1 \wedge \langle a, b, c \rangle \notin r_2 \}$
 - $\{ \langle a, b, c \rangle \mid \exists p, q (\langle a, b, p \rangle \in r_1 \wedge \langle q, b, c \rangle \in r_2) \}$
- 6.7 Let $R = (A, B)$ and $S = (A, C)$, and let $r(R)$ and $s(S)$ be relations. Write expressions in relational algebra for each of the following queries:
- $\{ \langle a \rangle \mid \exists b (\langle a, b \rangle \in r \wedge b = 7) \}$
 - $\{ \langle a, b, c \rangle \mid \langle a, b \rangle \in r \wedge \langle a, c \rangle \in s \}$
 - $\{ \langle a \rangle \mid \exists c (\langle a, c \rangle \in s \wedge \exists b_1, b_2 (\langle a, b_1 \rangle \in r \wedge \langle c, b_2 \rangle \in r \wedge b_1 > b_2)) \}$

Answer:

- $\Pi_A(\sigma_{B=17}(r))$
 - $r \bowtie s$
 - $\Pi_A(s \bowtie (\Pi_{r.A}(\sigma_{r.b=d.b}(r \times \rho_d(r))))))$
- 6.8 Consider the relational database of Figure 6.22 where the primary keys are underlined. Give an expression in tuple relational calculus for each of the following queries:
- Find all employees who work directly for “Jones.”
 - Find all cities of residence of all employees who work directly for “Jones.”

6 Chapter 6 Formal Relational Query Languages

- c. Find the name of the manager of the manager of “Jones.”
- d. Find those employees who earn more than all employees living in the city “Mumbai.”

Answer:

a.

$$\{t \mid \exists m \in \text{manages} (t[\text{person_name}] = m[\text{person_name}] \wedge m[\text{manager_name}] = \text{'Jones'})\}$$

b.

$$\{t \mid \exists m \in \text{manages} \exists e \in \text{employee} (e[\text{person_name}] = m[\text{person_name}] \wedge m[\text{manager_name}] = \text{'Jones'} \wedge t[\text{city}] = e[\text{city}])\}$$

c.

$$\{t \mid \exists m1 \in \text{manages} \exists m2 \in \text{manages} (m1[\text{manager_name}] = m2[\text{person_name}] \wedge m1[\text{person_name}] = \text{'Jones'} \wedge t[\text{manager_name}] = m2[\text{manager_name}])\}$$

d.

$$\{t \mid \exists w1 \in \text{works} \neg \exists w2 \in \text{works} (w1[\text{salary}] < w2[\text{salary}] \wedge \exists e2 \in \text{employee} (w2[\text{person_name}] = e2[\text{person_name}] \wedge e2[\text{city}] = \text{'Mumbai'}))\}$$

6.9 Describe how to translate join expressions in SQL to relational algebra.

Answer: A query of the form

```
select A1, A2, ..., An
from R1, R2, ..., Rm
where P
```

can be translated into relational algebra as follows:

$$\Pi_{A1, A2, \dots, An}(\sigma_P(R1 \times R2 \times \dots \times Rm))$$

An SQL join expression of the form

R1 natural join R2

can be written as $R1 \bowtie R2$.

An SQL join expression of the form

R1 join R2 on (P)

can be written as $R1 \bowtie_P R2$.

An SQL join expression of the form

$$R1 \text{ join } R2 \text{ using } (A1, A2, \dots, An)$$

can be written as $\Pi_S(R1 \bowtie_{R1.A1=R2.A1 \wedge R1.A2=R2.A2 \wedge \dots \wedge R1.An=R2.An} R2)$ where S is $A1, A2, \dots, An$ followed by all attributes of $R1$ other than $R1.A1, R1.A2, \dots, R1.An$, followed by all attributes of $R2$ other than $R2.A1, R2.A2, \dots, R2.An$,

The outer join versions of the SQL join expressions can be similarly written by using $\bowtie\!\!\!\diagup$, $\bowtie\!\!\!\diagdown$ and $\bowtie\!\!\!\diagup\!\!\!\diagdown$ in place of \bowtie .¹

The most direct way to handle subqueries is to extend the relational algebra. To handle where clause subqueries, we need to allow selection predicates to contain nested relational algebra expressions, which can reference correlation attributes from outer level relations. Scalar subqueries can be similarly translated by allowing nested relational algebra expressions to appear in scalar expressions. An alternative approach to handling such subqueries used in some database systems, such as Microsoft SQL Server, introduces a new relational algebra operator called the Apply operator; see Chapter 30, page 1230-1231 for details. Without such extensions, translating subqueries into standard relational algebra can be rather complicated.

¹The case of outer joins with the **using** clause is a little more complicated; with a right outer join it is possible that $R1.A1$ is null, but $R2.A1$ is not, and the output should contain the non-null value. The SQL **coalesce** function can be used, replacing S by **coalesce**($R1.A1, R2.A1$), **coalesce**($R1.A2, R2.A2$), ... **coalesce**($R1.An, R2.An$), followed by the other attributes of $R1$ and $R2$.

