

1 Motivation

In the previous notes we talked about how SQL is a declarative programming language. This means that you specify what you want, but you don't have to specify how to do it. This is great from a user's perspective as it makes the queries much easier to write. As database engineers, however, we often want a language that is more expressive. When we study query optimization in a few weeks we're going to want a way to express the many different valid plans a database can use to execute a query. For this we will use **Relational Algebra**, a procedural programming language (meaning that the query specifies exactly what operators to use and in what order).

2 Relational Algebra Introduction

All of the operators in relational algebra take in a relation and output a different relation. A basic query looks like this:

$$\pi_{name}(dogs)$$

The π operator picks only the columns that it wants to advance to the next operator (just like SQL SELECT). In this case, the operator takes the `dogs` relation in as a parameter and returns a relation that only has the `name` column of the `dogs` relation. An important fact about relational algebra is that the relations are sets of tuples, meaning that they cannot have duplicates in them. If the `dogs` relation is initially:

name	age
Scooby	10
Buster	15
Buster	20

The query above would return:

name
Scooby
Buster

Initially the two Busters are different because they have different ages, but once you get rid of the age column, they become duplicates, so only one remains in the output relation.

Let's formally introduce the relational algebra operators.

3 Projection (π)

We have already been introduced to the **projection operator** which selects only the columns specified. The columns are specified in the subscript of the operator like almost all parameters to

operators. The projection operator is relational algebra's version of the SQL SELECT clause.

We now can express SQL queries involving just the SELECT and FROM clauses with relational algebra. For example the SQL query:

```
SELECT name FROM dogs ;
```

Can be represented with the expression we introduced in section 2:

$$\pi_{name}(dogs)$$

Note that there is no operator equivalent to the FROM operator in relational algebra because the parameters of these operators specify which tables we pull from.

4 Selection (σ)

The **selection operator** is used to filter rows based on a certain condition. Don't let the name confuse you - this operator is equivalent to SQL's WHERE clause, **not** its SELECT clause. Let's try to express the following query in terms of relational algebra:

```
SELECT name, age FROM dogs WHERE age = 12;
```

The equivalent relational algebra expression is:

$$\sigma_{age=12}(\pi_{name,age}(dogs))$$

Another correct expression for that query is:

$$\pi_{name,age}(\sigma_{age=12}(dogs))$$

This illustrates the beauty of relational algebra. There is only one (reasonable) way to write SQL for what the query is trying to accomplish, but we can come up with multiple different expressions in relational algebra that get the same result. In the first expression we select only the columns we want first, and then we filter out the rows we don't want. In the second we filter the rows first and then select the columns. We will soon learn ways to evaluate which of these plans is better!

The selection operator also supports compound predicates. The \wedge symbol corresponds to the AND keyword in SQL and the \vee symbol corresponds to the OR keyword. For example,

```
SELECT name, age FROM dogs WHERE age = 12 AND name = 'Timmy' ;
```

is equivalent to

$$\pi_{name,age}(\sigma_{age=12 \wedge name='Timmy'}(dogs))$$

5 Union (\cup)

The first way we will learn how to combine data from different relations is with the **union operator**. Just like the UNION clause in SQL, we take all the rows from each tuple and combine them removing duplicates along the way. As an example, say we have a dogs table:

name	age
Scooby	10
Buster	15
Garfield	20

and a cats table that looks like this:

name	age
Tom	8
Garfield	10

The expression:

$$\pi_{name}(dogs) \cup \pi_{name}(cats)$$

would return:

name
Scooby
Buster
Tom
Garfield

Note that Garfield only shows up once because relations are sets of tuples and remove all duplicates as a result. Additionally, note that one rule for all of these set operators is that they must operate on relations that have the same number of attributes (columns), and the attributes must have the same type. It would not be legal to union a relation with two columns with a relation that only has one column nor would it be legal to union a relation with a column of strings with another relation with one column of ints.

6 Set Difference ($-$)

Another set operator is the **set difference** operator. Set difference is equivalent to the SQL clause EXCEPT. It returns every row in the first table except the rows that also show up in the second table. If you run:

$$\pi_{name}(dogs) - \pi_{name}(cats)$$

expression on the dogs and cats table introduced in the previous section you would get:

name
Scooby
Buster

Garfield does not show up because he is in the cats table, and none of the cats' names will show up because it is only possible for rows from the first relation to show up in the output.

7 Intersection (\cap)

Intersection is just like the INTERSECT SQL operator in that it only keeps rows that occur in both tables in the intersection. If you run:

$$\pi_{name}(dogs) \cap \pi_{name}(cats)$$

on the tables introduced in section 5 you would get:

name
Garfield

because Garfield is the only name to occur in both tables.

8 Joins (\bowtie)

We haven't yet discussed how to represent joins in relational algebra - let's fix that! To join two tables together, write the left table on the left of the \bowtie operator, put the join condition in the subscript, and put the right operator on the right side. To join together the cats table with the dogs table on the name column, you would write:

$$cats \bowtie_{cats.name=dogs.name} dogs$$

If you don't specify the join condition, it becomes a natural join. Recall from the SQL notes that a natural join joins together all columns from each table with the same name. Therefore, you could also write the same query as above like:

$$cats \bowtie dogs$$

Just like the selection operator σ , the join operator \bowtie also supports the compound predicate operators \wedge (AND) and \vee (OR).

9 Rename (ρ)

The last operator we will discuss is the **rename operator** which essentially accomplishes the same thing as aliasing in SQL. If you wanted to avoid having to include the table name for the rest of the expression like you would for the expressions in the join section, you could instead write:

$$cats \bowtie_{name=dname} \rho_{name \rightarrow dname}(dogs)$$

This expression renames the dogs relation's name column to dname first, so there is no conflict in column names. You can no longer use a natural join anymore because the columns do not have the same name, but you no longer need to specify which relation the column is coming from if you want to include other operators.

10 Practice Questions

Given the following two relations:

`teams(teamid, name)`

`players(playerid, name, teamid, position)`

Answer the following questions:

1. Write an expression that finds the name and playerid of every player that plays the “center” position.
2. Write an expression that finds the name of every player that plays on the “Warriors”. How would this expression change if we renamed players' teamid column to pteamid?
3. Write an expression that finds the teamid of all teams that do not have any players.

11 Solutions

1. $\pi_{name, playerid}(\sigma_{position='center'}(players))$. We first filter out the rows for players who aren't centers, then we project only the columns that we need.
2. $\pi_{players.name}(\sigma_{teams.name='Warriors'}(teams \bowtie_{teams.teamid=players.teamid} players))$. We first join together the teams and players table to get all the information that we need, then we filter out the rows that aren't for players who play for the Warriors, then we finally project the only column that we're looking for.
3. $\pi_{teamid}(teams) - \pi_{teamid}(players)$. All teams must be in the teams table so we first get all their teamids. Then we subtract any teamid that appears in the players table, because if that teamid appears in the players table it implies that the team has a player on it. We are then left with only teamids of teams that don't have any players.