

Coursework Instructions Milestone 3: Querying a Database

The goal of the third part of your IDA project is to

- query your PostgreSQL database using the SQL query language;
- practice queries in relational algebra.

For this milestone, you will have an occasion to apply the material of the lectures on relational algebra and on the SQL query language.

Requirement Queries

As the first step in your requirement collection you have specified 10 queries from the domain of your application. Now, try to express them in SQL.

It may be the case that some of your queries were too ambitious so that they cannot be captured by SQL. In that case

1. explain why you think the query cannot be written as SQL query;
2. identify other queries that can retrieve useful data for your original query and describe how that data has to be processed further to obtain answers for the original query.

New Queries Over the Database

In addition to the queries stemming from the requirements, write new queries over your database schema that are interesting in your application domain. The queries may be simple `SELECT` statements or `INSERT`, `UPDATE`, `DELETE` statements that contain a `SELECT` statement.

For each query, first write in English what it is supposed to do, and then give the SQL code. Groups of three should write at least 24, groups of four at least 32 essentially different queries. (Note: This means, that all in all you write $10+24$ or $10+32$ queries.)

Among those queries at least two thirds should be *complex* in the sense that they contain at least one of the following features:

- joins involving the same table twice,
- aggregation with group by and having,
- nesting with aggregation,
- nested negation, involving NOT EXISTS or NOT IN,
- outer join,
- usage of an auxiliary view.

Queries in Relational Algebra

Choose a number of queries you have written in SQL and express them in relational algebra (linear notation is fine). Groups of three students should write 6 such queries and groups of four students should write 8. For the each algebra query, say which is the corresponding SQL query.

The point of this exercise is to give you one more occasion to practise relational algebra queries and to receive feedback.

Deliverable

The deliverable will be a document consisting of two parts:

1. a brief report on the translation of the requirement queries into SQL, showing for each query in English the SQL translation, or discussing the difficulties encountered;
2. a brief report containing the new queries and for each query an explanation in English of
 - what it is supposed to do, and
 - why the query is interesting in your application;
3. a brief report with the relational algebra queries.

Note: All queries must be tested and in a form that they can be executed by PostgreSQL. SQL queries containing a syntax error do not count as part of the submission.

Deadline and Submission

The work is to be submitted by publishing it on the web site of your IDA group. The deadline is

Tuesday, 30 November, 11:30pm.