

Exercises with POSTGRESQL

The goal of this lab is to go over the last two lessons on schemas, views, and on query evaluation.

1 Connecting to the server

To connect to the server you will need to have your identifiers provided at the beginning of the lab. Once you have your identifier you will connect using the `psql` command in the terminal. The general syntax of the command is `psql -h hostname -p port -U user database`

In this lab we will use the following parameters: *hostname* is `teach.jachiet.com` *port* is `4444` *user* is the one given (of the form `tpXXX` where the `X` are numbers) *database* will be `unicode`.

Once you hit enter you should be asked for a password, type the one given at the beginning of the lab. Once you are successfully connected, I invite you to read <https://tomcam.github.io/postgres/> to learn how to manipulate the `psql` shell. Note that your user is configured in a way where it can only connect once, you cannot open two connection with the same user!

2 Query evaluation

2.1 Exploring the unicode dataset

In POSTGRESQL adding `EXPLAIN` at the beginning of a command tells POSTGRESQL to show the query plan instead of actually running the query. For this section we will use the `unicode` table in the `unicode` database of the server.

Exercise 1. What is the schema of the table?

Exercise 2. What are the indexes present?

Exercise 3. Design queries such that each of the operators below appears in the query plan of one of your queries. For that you will need to look at the data (number of records, which columns actually contain null, etc.).

- SeqScan
- Index Scan
- Index Only Scan
- Bitmap Index Scan
- Bitmap Heap Scan
- BitmapOr
- BitmapAnd
- Filter
- Nested Loop
- Hash Join
- Merge Join

Note that some of these operators might be complex to trigger. You need to think about how postgres optimizes to trigger these operations (e.g. `BitmapAnd` or the `Merge Join`).