

1. You are given a relation with 25 million integer values stored in a text file. Every line contains exactly one value. Generate this file with the provided Java-Program ueb02 gen (Attention: The generated file has 180 MB). Write a Java program, that answers the following SQL query in an efficient way:

```
select val, count(*) from r group by val;
```

The result should be a list of a respective value and its frequency in the 180 MB file.

Restriction:

To get closer to the real situation of a data warehouse, limit the main memory usage of the Java jvm to 30MB. For that you have to call Java with additional arguments as shown next:

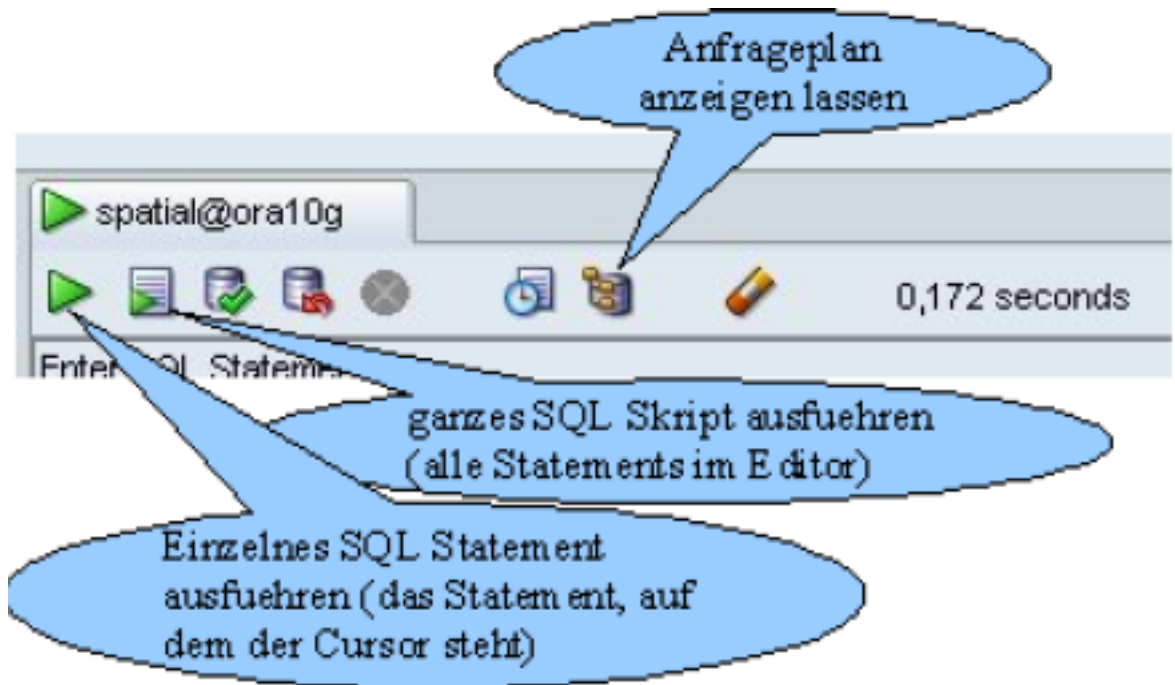
```
java -Xms30m -Xmx30m meineKlasse
```

If you exceed the maximum heap size, an `OutOfMemoryError` exception is thrown.

2. Explain the conceptual differences between transactional data bases (OLTP) and data warehouses (OLAP). For that, explain the term *transaction*.
3. Define the term data warehouse, distributed and federated data bases. Describe differences of these systems. How are data marts related to the former systems.
4. What are dimensions, facts and ratios in the context of data warehouses? Provide a self-chosen example.
5. Investigate on benchmarking of data bases. What benchmarks exist? Concentrate on TPC-H benchmarks.<sup>1</sup>
6. Preparation for exercise 3: Model a non-trivial beverage chain company. Data from different departments will be stored in the database system. Provide ER schema and SQL. Possible departments like:
  - (a) Human resource,
  - (b) Sales,
  - (c) Technical support,
  - (d) Vehicles management,
  - (e) Quality management,

---

<sup>1</sup>For further information visit <http://tpc.org/>



- (f) Distribution,
- (g) Product development,
- (h) Purchasing
- (i) ...

Do not take model from lecture slides!

you can logon to the database with sqldeveloper at home.<sup>2</sup> If you have an X server on your machine, just execute the following command in an X-term:

```
ssh -fX <login>@hippo.cs.uni-magdeburg.de /local/apps/sqldeveloper/sqldeveloper.sh
sqldeveloper is installed in the SUN pool:
SQLDeveloper: /local/apps/sqldeveloper/sqldeveloper.sh
SID: tix
Host: deanston.cs.uni-magdeburg.de
Port: 1521
```

<sup>2</sup>Download: [www.oracle.com/technology/products/database/sqldeveloper/index.html](http://www.oracle.com/technology/products/database/sqldeveloper/index.html)