



Query Optimization

5. Exercise

Due November 27, 2017, 9 AM

submit via email (radke@in.tum.de)

Exercise 1

Create the DP table (manually) for the relations A , B , C with cardinalities $|A| = 10$, $|B| = 20$, $|C| = 100$ and selectivities $f_{AB} = 0.5$, $f_{BC} = 0.1$ (cost function C_{out}). Mark the final table entries. Enumerate subsets in the integer order. Consider cross products.

Exercise 2

Using the program from the last exercise as basis, implement Greedy Operator Ordering. Print the partial steps together with their costs (e.g., $P = R_1 \bowtie R_2 200$, $Q = P \bowtie R_3 400$), as well as the final join tree.

Exercise 3

Load the [TPC H](#) data set. (You can use our [snapshot](#) of the data set, the `loadtpch*` script loads the data). Then, execute the following SQL query using the program implemented above:

```
select *
  from lineitem l, orders o, customers c
 where l.l_orderkey=o.o_orderkey
        and o.o_custkey=c.c_custkey
        and c.c_name="Customer#000014993".
```