Chair of applied computer science III
Prof. Dr. Guido Moerkotte
Email: moerkotte@uni-mannheim.de
UNIVERSITY OF MANNHEIM

Query Optimization — Exercise sheet 5

Exercise 1

Exercise 1 a)

Classify the join ordering algorithmns discussed in the lecture. Fill in the following table:

Name	Query Graph	Join-Tree	Cross products	cost functions	complexity	optimal	remarks

Exercise 1 b)

Consider the relations

$$|R_0| = 10, |R_1| = 2, |R_2| = 100$$

and the join selectivities

$$f_{R_0R_1} = 0.5, f_{R_1R_2} = 0.1$$

Note down the corresponding DP-tables for `DPsub` and `DPsize`.
Keep track of all plans written to the DP-table and mark the final plan.
Use C_{out} as the cost function.

Exercise 1 c)

Now, note down the DP-tables corresponding to the above query graph for `DPccp`.
The subscripts denote the order in which BFS (breadth-first search) visits the nodes.

Exercise 2

Exercise 2 a)

Implement either `DP-Linear-1`, `DP-Linear-2` or `DP-Bushy`. You may use the helper classes provided in the solution code.

Exercise 2 b)

Among how many join trees do `DP-Linear-1` (with cross products) and `DP-Bushy` select the optimal one?

Exercise 2 c)

Bonus
Compare the number of plans generated by the enumerator from the first exercise with the number generated partial trees (denoted by variable name CurrTree in the pseudocode) in `DP-bushy`.

Exercise 3

In the algorithms discussed in the lecture, We have already seen several enumeration techniques in the algorithms discussed in the lecture. This exercise is to recall these techniques.

Exercise 3 a)

How to enumerate all sizes for the *two sides of a plan* (i.e., all unordered pairs) for all plan sizes? Write down the order in which values created.

Exercise 3 b)

How to enumerate all possible values of a bitvector of size n in ascending order? Write down the order in which values created.

Exercise 3 c)

How to enumerate all subsets of a set of n values? Assume that a set is implemented as a bitvector. Write down the order in which values created.