Query Optimization

3. Exercise

Due 12.05.2014, 9 AM

submit via email (`Andrey.Gubichev@in.tum.de`)

Exercise 1

The selectivity estimations (from the previous homework) are far from perfect. Construct specific examples (database schema, concrete instances of relations and selections/joins), where our estimations are very "bad", i.e. for some queries (give examples of SQL queries) the logical plan will be suboptimal (w.r.t. C_{out}), if we use these estimations. Give two examples (one for selections, one for joins).

Exercise 2

Give an example query instance where the optimal join tree (using C_{out}) is bushy and includes a cross product. Note: the query graph should be connected!

Exercise 3

Using the program from the first exercise as a basis, implement a program that parses SQL queries, translates them into *tinydb* execution plans, and executes the query. Note: a canonical translation of the joins is fine, but push all predicates of the form $attr = const$ down to the base relations