



Query Optimization

4. Exercise

Due 17.11.2014, 9 AM

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

Exercise 1

1. Give an example query graph with join selectivities for which the greedy operator ordering (GOO) algorithm does not give the optimal (with regards to C_{out}) join tree. Specify the optimal join tree. Do not consider cross products.
2. Construct the join tree for the example from 1. using the IKKBZ-based heuristics. Give the C_{out} cost for the result and compare it with the GOO tree's cost.

Exercise 2

Using the program from the the last exercise as basis, construct the query graph for each connected component. Annotate the query graph as follows:

- annotate each node with the underlying relation, the list of predicates that can be pushed down to that relation, and the estimated result cardinality
- annotate each edge with the underlying join predicate, and the estimated selectivity

The *Table* and *Attribute* classes in *tinydb* allow to retrieve the domain size (i.e., the number of unique values) of attributes and the cardinality of base tables. Keep in mind that we will construct join trees for these query graphs in later exercises.