

Max Planck Institut für Informatik  
Department D5  
Thomas Neumann

Campus E1 4, Room 418  
66123 Saarbrücken  
Phone: (0681) 9325 518  
Email: neumann@mpi-inf.mpg.de



# Query Optimization

## 3. Exercise, Summer 2009

Due 2009-05-13.

1. Give an example where GreedyJoinOrdering-1 does not find the optimal solution but IKKBZ does. Explain the IKKBZ solution.
2. How can we estimate the selectivity of a join predicate  $R_1.x = R_2.x$ ? Assume that we can pre-compute a few (suitable) values per relation if necessary.
3. Modify the program from the last exercise to construct join trees (with the tinydb HashJoin class) using Greedy Operator Ordering. Measure the runtime difference relative to the first program `select * from lineitem l, order o where l.l_orderkey=o.o_orderkey` on the *tpcd* dataset. Note: this exercise is due 2009-05-20.