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

Query Optimization — Exercise sheet 10

Exercise 1

Exercise 1 a)

Describe the Cheung formula.
Compute the Cheung formula for the values $N = 1000, m = 100, k = 15$

Exercise 2

Exercise 2 a)

Read the chapter on *Genetic Algorithms* in the *Probabilistic Algorithms* section in the script.

Exercise 2 b)

Implement the genetic algorithm for the join ordering problem. It is sufficient to consider only left deep trees.
Note, on a high level, it is easy to implement the algorithm. However you will need a significant number of helper functions.
Feel free to take the solution code and only implement the high level level algorithm (`Run` function).

Exercise 3

Queries with multiple predicates may require the evaluation of multiple indices. In what order should you process the predicates?

Exercise 4

Instead of using B-Tree one can also use hashing based data structures. For instance, one can build on the idea of extendible hashing
`https://en.wikipedia.org/wiki/Extendible_hashing`
For queries with what type of predicates can you apply such data structures?

Exercise 5

You are given a relation with 6 tuples. These tuples are equally distributed over 3 pages. Compute the average number of page accesses for reading 2 tuples. Assume that all tuples have the same probability to be read.

Exercise 6

BONUS

Write a program that allows you to measure some physicial properties of your disk.