

Assignment 2

Query Tuning

Database Tuning

Start date: October 30, 2014

Due date: November 11, 23:59

Grading: 1 point

In this assignment you will gain hands-on experience in rewriting slow queries and in experimentally evaluating the rewritten queries.

Task 1: Create a database with the following database schema:

- Employee(ssnum, name, manager, dept, salary, numfriends)
 - unique index on `ssnum`
 - unique index on `name`
 - index on `dept`
- Student(ssnum, name, course, grade)
 - unique indexes on `ssnum`
 - unique indexes on `name`
- Techdept(dept, manager, location)
 - unique index on `dept`
 - a manager may manage multiple departments
 - a location may contain multiple departments

Task 2: Fill the database with 100k employees, 100k students, and 10 technical departments. Only about 10% of the employees are in a technical department. The types of the attributes should make sense (e.g., `ssnum` should be an integer), but the values need not be meaningful (e.g., names can be random strings).

Task 3: Choose two types of queries that might be hard for your database to optimize. Taking queries from the lecture notes is OK.

NOTE: For at least one of your queries rewriting should make a difference.

Task 4: Rewrite the queries and consult the execution plans of the original and the rewritten query.

Task 5: Run the original and the rewritten query and measure the runtime.

Report:

- Describe your instance (data types, how did you fill the tables?).
- Give the original and the rewritten queries.
- Show and explain the execution plans.
- Report and briefly discuss the runtime results from your experiment.

Please indicate the average time per group member that was spent solving this assignment. The time that you indicate will have *no* impact on your grade.