Assignment 3

Index Tuning

Database Tuning

Due date: April 24, 2018, 23:55
Grading: 5 points

Notes

- It is suggested that you also have a look at the report template before you start working on the assignment.

Access Parameters for PostgreSQL

- Host: `biber.cosy.sbg.ac.at`
- Port: `5432`
- Database: `dbtuning_ss2018`
- User/Password: you should have received them via email

The database server (`biber.cosy.sbg.ac.at`) is accessible only from inside the university network. If you would like to work from home, please connect to `fanny.cosy.sbg.ac.at` via `ssh`. Java and the PostgreSQL client as well as Python are installed on this machine.

Support

If there are any ambiguities or problems of understanding regarding the assignment, you have the following possibilities to clarify them:

- Slack channel **#dbt**[1] (preferred way of communication)
- Office hours: Wednesday, 10am - 11am, Office 0.26 (ground floor)

In this assignment you will study the indexing capabilities of a database management systems of your choice.

Choose one of the following database management systems:

- PostgreSQL 9
- Oracle 12c
- SQL Server 2012
- IBM DB2 UDB V9

Consider the table `Employee(ssnum,name,dept,salary)`, where `ssnum` is a key. For the system of your choice answer the following questions in your report.

[1] `https://dbteaching.slack.com`

1 Index Data Structures

Which index data structures (e.g., B^+-tree index) are supported?

2 Clustering Indexes

Discuss how the system supports clustering indexes, in particular:

a) How do you create a clustering index on `ssnum`? Show the query.[2]

b) Are clustering indexes on non-key attributes supported, e.g., on `name`? Show the query.

c) Is the clustering index dense or sparse?

d) How does the system deal with overflows in clustering indexes? How is the fill factor controlled?

e) Discuss any further characteristics of the system related to clustering indexes that are relevant to a database tuner.

3 Non-Clustering Indexes

Discuss how the system supports non-clustering indexes, in particular:

a) How do you create a non-clustering index on `(dept,salary)`? Show the query.[1]

b) Can the system take advantage of covering indexes? What if the index covers the query, but the condition is not a prefix of the attribute sequence `(dept,salary)`?

c) Discuss any further characteristics of the system related to non-clustering indexes that are relevant to a database tuner?

4 Key Compression and Page Size

If your system supports B^+-trees, what kind of key compression (if any) does it support? How large is the default disk page? Can it be changed?

Important: Reference your information sources.

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.

Grading scheme:

For each item, i.e., 1, 2a 2b, ... 3c, and 4, 0.5 points

Important: If the grading scheme is unclear, ask the lecturer!

[2]Give the queries for creating a hash index *and* a B^+-tree index if both of them are supported.