

Assignment 3

Index Tuning

Database Tuning

Due date: April 10, 2019, 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_ss2019`
- User/Password: you should have received them via email

The database server (`biber.cosy.sbg.ac.at`) is accessible only from within the university network. If you would like to work from home, please connect to `fanny.cosy.sbg.ac.at` via `ssh`. Java, the PostgreSQL client, and 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 `#dbt2019`¹ (preferred way of communication)
- Upon request via email (`dkocher@cs.sbg.ac.at`)

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

Choose one of the following database management systems:

- PostgreSQL (latest stable version: 11)
- Oracle (latest stable version: 19c)
- SQL Server (latest stable version: 2017)
- IBM DB2 UDB (latest stable version: 11.1)

You are allowed to choose older versions of the respective systems (for example, we run PostgreSQL 9.6 on `biber.cosy.sbg.ac.at`). Therefore, mention *precisely* the version of your system of choice in the report.

¹<https://dbteaching.slack.com>

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 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:

- How do you create a clustering index on `ssnum`? Show the query.²
- Are clustering indexes on non-key attributes supported, e.g., on `name`? Show the query.
- Is the clustering index dense or sparse?
- How does the system deal with overflows in clustering indexes? How is the fill factor controlled?
- 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:

- How do you create a combined, non-clustering index on `(dept,salary)`? Show the query.²
- 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)`?
- 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) is supported? How large is the default disk page? Can it be changed?

Important: Reference your information sources (footnotes, bibliography, ...).

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

Reminder: Additional questions about the involved topics/techniques will be asked during the meeting.

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

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