

Assignment 3

Index Tuning

Database Tuning

Start date: November 13, 2014

Due date: November 25, 23:59

Grading: 1 point

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.

1. Which index data structures (e.g., B^+ -tree index) are supported?
2. Discuss how the system supports clustered indexes, in particular:
 - a) How do you create a clustered index on `ssnum`? Show the query.¹
 - b) Are clustered indexes on non-key attributes supported, e.g., on `name`? Show the query.
 - c) Is the clustered index dense or sparse?
 - d) How does the system deal with overflows in clustered indexes? How is the fill factor controlled?
 - e) Discuss any further characteristics of the system related to clustered indexes that are relevant to a database tuner?
3. Discuss how the system supports non-clustered indexes, in particular:
 - a) How do you create a non-clustered index on `(dept,salary)`? Show the query.¹

¹Give the queries for creating a hash index *and* a B^+ -tree index if both of them are supported.

- 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-clustered indexes that are relevant to a database tuner?
4. 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.