

Database and Information Systems

Homework 6

November 22, 2005; Due November 29 at 1:30 PM

Problem 1 [50 points]:






Figure 1: B+ Tree for Problem 1

Consider the B+ tree index shown in Figure 1, which uses alternative(1) for data entries. Each intermediate node can hold up to five pointers and four key values. Each leaf node can hold up to four records (indicated as a key value plus a “*”), and leaf nodes are doubly linked to their predecessor and successor nodes in the index (although this is not shown in the figure). The minimum number of keys allowed in an intermediate node is two, and the minimum number of records allowed in a leaf node is also two. In other words, the *order* of the tree is two. Underscores (“___”) indicate unallocated entries.

1. Insert a record with search key 94 into the tree.



2. Delete the record with search key 38 from the tree.



Problem 2 [50 points]: For this part of the assignment, you will need to run Oracle on `eniac`, using the `sql` command, as in previous assignments. The goal is to see the benefits of indexing in action.

The Oracle setup on `eniac` has a series of sample tables created for running database benchmarks — the so-called TPC-H benchmark. Additionally, each of these tables may have several alternative index structures. We will be comparing performance depending on which indices are used. To do this assignment, you will need to repeat each query 5 times, time it with a stopwatch or the PC clock, and average the running times. (Averaging is necessary in case others are using `eniac` at the same time and affecting your running times.)

The `lineitem` table has nearly 300,000 rows and takes approximately 37MB of space. We have two indices in `lineitem`, each with a name: `sys_c001480` is an index over the attribute pair (`l_orderkey`, `l_linenumber`) and `lineno` is on the singleton attribute (`l_linenumber`).

1. Start by running the query:

```
select avg(l_linenumber)
  from zives.lineitem
 where l_linenumber >= 3;
```

which uses whatever indices Oracle prefers. Time the run and repeat it 5 times, averaging the total. What is that value?

Answer: about 0.1 seconds.

2. Next, repeat the query with a special *optimizer hint* that says you want to scan the full table and avoid indices:

```
select /*+ full(l) */ avg(l_linenumber)
  from zives.lineitem l
 where l_linenumber >= 3;
```

Average the running time for 5 runs and record the result. What can you say about the difference between the two runs? Do these results suggest that Oracle uses an index by default?

Answer: about 0.3 seconds. When the query avoids indices, it takes about 3 times longer. It suggests Oracle uses an index by default.

3. Try another hint, which makes use of the first index:

```
select /*+ index(l sys_c001480) */ avg(l_linenumber)
  from zives.lineitem l
 where l_linenumber >= 3;
```

Run 5 times and record the result. How did this compare to the first two cases? Did that particular index offer any benefits?

Answer: About 0.18 seconds. It is slower than the default case and faster than the no-index case. The particular index offers some benefits.

4. Which index was used in the default, un-hinted Oracle query? Generalizing from the Oracle SQL hints shown above, what is the query hint that is equivalent to that default query?

Answer: when use index `/*+ index(l lineno) */`, the averaged running time is about 0.1 seconds. So the default, un-hinted Oracle query might use `lineno` index.

5. For each of the two indices *and* for the non-indexed case, explain how what data and/or index structures will be examined in answering the query.

Answer:

The `lineno` index is sorted on `l_linenumber`, the system can locate the first index entry of `l_linenumber >=3`, skip the part of the index where `l_linenumber <3`. Partial indices need to be accessed.

The `sys_c001480` index is sorted first by `l_orderkey`, then by `l_linenumber`. Two data entries with same value of `l_linenumber` may be located in two different places in the index tree(or index table) due to `l_orderkey`. Thus, the entire index has to be accessed.

For the non-indexed case, all data entries in the `lineitem` relation will be scanned sequentially.

6. Run the two queries:

```
select /*+ index(l sys_c001480) */ count(distinct l_suppkey)
  from zives.lineitem l
 where l_linenumber < 400;
```

```
select /*+ full(l) */ count(distinct l_suppkey)
  from zives.lineitem l
 where l_linenumber < 400;
```

Report the averaged running times over 5 runs. Explain the relative differences (or lack thereof) in performance, based on your knowledge of index and data layout characteristics, as well as your knowledge of “covering” indices.

Answer:

The first query takes about 1 minutes 20 seconds; the second query takes about 45 seconds.

In the first query, since `sys_c001480` doesn't contains `l_suppkey`, when the system finishes an index access, it has to retrieve the actual `l_suppkey` value by another data entry access. The physical page of the data entry may be on the different page of the index entry. Thus, to obtain a single value of `l_suppkey`, two disk I/Os are required and take longer time.

The second query sequentially read every data entry once.