

Homework 2 for Distributed Database Systems

(Issue: 2012 / 5 / 5 Due: 2012 / 5 / 16)

1. Consider a database consisting of the following two relations

Employees (eid:integer, did:integer, sal:real)

Department (did:integer, mgrid:integer, budget:integer)

The mgrid field of Department is the eid of the manager. Each of these relations contains 20-byte tuples, and the sal and budget fields both contain uniformly distributed values in the range 0 to 1,000,000. The Employees relation contains 100,000 pages, the Departments relation contains 5000 pages, and each processor has 100 buffer pages of 4000 bytes each. The cost of one page I/O is t_d and the cost of shipping one page is t_s . There are no indexes.

The database is stored in distributed DBMS with 10 sites. The Departments tuples are horizontally partitioned across the 10 sites by did, with the same number of tuples assigned to each site, and with no particular order to how tuples are assigned to sites. The Employees tuples are similarly partitioned, by sal ranges, with $sal \leq 100,000$ assigned to the first site, $100,000 < sal \leq 200,000$ assigned to the second site, etc. In addition, the partition $sal \leq 100,000$ is frequently accessed and infrequently updated, and it is therefore replicated at every site. No other Employees partition is replicated.

- 1) Give the cost for computing the natural join of Employees and Departments using the strategy of shipping all fragments of the smaller relation to every site containing tuples of the larger relation.
 - 2) Describe the best plan and its cost for each of the following queries:
 - a) Find the highest paid employee.
 - b) Find the highest paid employee, with salary greater than 450,000 and less than 550,000.
 - c) Find the highest paid manager for those departments stored at the query site.
2. Suppose that an Employees relation (the schema of Employees is described in Question 1) is stored in Hong Kong and the tuples with $sal \leq 100,000$ are replicated at Beijing.

Consider the following three options for lock management:

- All locks managed at a single site called Shenzheng;
- Primary copy: the copy of Employees at Hong Kong is chosen as the primary copy;
- Fully distributed

For each of the lock management options, explain locks are set at which site for the following queries, and state which site the page is read from:

- 1) A query submitted at Shanghai wants to read a page containing Employees tuples with $sal \leq 50,000$.
- 2) A query submitted at Hong Kong wants to read a page containing Employees tuples with $sal \leq 50,000$.
- 3) A query submitted at Beijing wants to read a page containing Employees tuples with $sal \leq 50,000$.
- 4) An update transaction, "Give all employees a 10% raise", is issued in the DDBS described in Question 1. Describe the sites visited and the locks obtained.
(Hint: ROWA policy is used for replications, and the conditions of the original partitioning of Employees must still be satisfied after update.)