

## Assignment #3

Due: 2:00pm CST on October 6, 2006

Out: September 20, 2006

**NOTE:** Please submit a hard copy of your homework. Bring it to the lecture table at the beginning of the lecture on 6th October, 2006.

The hard copy should be as clearly readable as possible. You may be subtracted points for unreadability and ugly presentation.

**I2CS Students:** You should email your solutions to the TA(ylee11@uiuc.edu) in the **pdf** or the **word** format. Please send the file as attachment with your email by 2PM UIUC time (CST). I2CS students in other time zones should note that the deadline is according to CST.

## Problem 1 (30 pts)

Suppose you have 2 relations,  $R(A, B, C)$  and  $S(B, C, D, E)$ . You have a clustered unique (no duplicate keys) B+-tree index on attribute A for relation R. Assume this index is kept entirely in memory (i.e., you do not need to read it from disk).

For relation S, you have two indexes: (i) a non-clustered non-unique B+-tree index for attribute B, and (ii) a clustered non-unique B+-tree index for attribute C. Assume that these two indexes are also kept in memory. Also, assume that all of the tuples of S that agree on attribute C are stored in sequentially adjacent blocks on disk (that is, if more than one block is needed to store all of the tuples with some value of C, then these blocks will be sequentially located on the disk).

Other relevant data:

- 2000 tuples of R are stored per block on disk.
- $N(R) = 2,000,000$  (number of tuples of R)
- 200 tuples of S are stored per block on disk.
- $N(S) = 200,000$  (number of tuples of S)
- $V(B,S)=40,000$  (image size of attribute B in S)
- $V(C,S)=800$  (image size of attribute C in S)

You want to execute the following query:

2

```
SELECT *
FROM R, S
WHERE (R.B=S.B) AND (R.C=S.C)
```

We present you with two query plans:

**Plan 1:**

```
For every block BL of R, retrieved using the clustered index on A for R
  For every tuple r of BL
    Use the index on B for S to retrieve all
      of the tuples s of S such that s.B=r.B
    For each of these tuples s, if s.C=r.C, output r.A, r.B, r.C,
                                          s.B, s.C, s.D, s.E
```

**Plan 2:**

```
For every block BL of R, retrieved using the clustered index on A for R
  For every tuple r of BL
    Use the index on C for S to retrieve all
      of the tuples s of S such that s.C=r.C
    For each of these tuples s, if s.B=r.B, output r.A, r.B, r.C,
                                          s.B, s.C, s.D, s.E
```

- a) Analyze the above two plans carefully in terms of their behavior regarding accesses to disk, and explain which of the plans is therefore better. Be sure to include in your analysis which accesses to disk are sequential accesses and which ones are random accesses.
- b) Can you think of an alternative plan that is better than the above two in terms of disk access? You do not need to submit detail computations for this question.

## Problem 2 (10 pts)

Consider the join  $R \bowtie_{R.a=S.b} S$ , given the following information about the relations to be joined. The cost metric is the number of page I/Os, and the cost of writing out the result should be uniformly ignored.

- Relation  $R$  contains 10,000 tuples and has 10 tuples per page.
- Relation  $S$  contains 2,000 tuples and also has 10 tuples per page.
- Attribute  $b$  of relation  $S$  is the primary key for  $S$ .
- Neither relation has any indexes built on it.
- 52 buffer pages are available.

Can we use **hybrid hash join**? If yes, what is the cost? If no, why?

### Problem 3 (20 pts)

Query optimization has been considered as a key technique for the realization of the relational model.

- Why does the relational DBMS (in particular) need query optimization? Why this was not an issue for earlier DBMS (e.g., of the network model).
- System R has established the *Selinger-style* query optimization. What are the main techniques in this framework?
- Suppose we wish to compute the following:

$$\tau_b(R(a, b) \bowtie S(b, c) \bowtie T(c, d))$$

That is, we join the three relations and produce the result sorted on attribute  $b$ . Assume that we do **NOT** join  $R$  and  $T$  first, as that is a Cartesian product.

What are all the subexpressions and *interesting orders* that System R optimizer would consider?

### Problem 4 (15 pts)

Consider a database  $DB$ .  $DB$  has two relations  $R1$  and  $R2$ . The relation  $R1$  contains tuples  $t1$  and  $t2$ , while  $R2$  contains tuples  $t3$ ,  $t4$ , and  $t5$ . Assume that the database  $DB$ , relations, and tuples form a hierarchy of lockable database elements.

Tell the sequence of lock requests and the response of the locking scheduler, which is described in the paper “Granularity of Locks and Degrees of Consistency in a Shared Data Base”, to the following sequence of request. You may assume all requests occur just before they are needed, and all unlocks occur at the end of the transaction.

r1(t1); w2(t2); r2(t3); w1(t4)

[w2(t2) represents the creation of t2 by transaction T2.]

### Problem 5 (25 pts)

This problem studies serializability and degree of consistency. Examine the schedule below. There are three transactions T1, T2, and T3.

	T1	T2	T3
0	start		
1	read X		
2		start	
3		read Y	
4	write X		
5			start
6			read X
7			write X
8			commit
9		read X	
10		write Y	
11		write X	
12		commit	
13	read Y		
14	write Y		
15	commit		

- a) Draw the precedence graph for this schedule.
- b) What is the equivalent serialization order for this schedule? If no order is possible, state NONE.
- c) Assume that transaction T1 did not run at all. Would the precedence graph be different in this case? Explain.
- d) What is the equivalent serialization order for this schedule? If no order is possible, state NONE.
- e) When all transactions run in the above schedule, identify the transactions with degree 3 consistency (using Definition 1 in the paper, “Granularity of Locks and Degrees of Consistency in a Shared Data Base”). Answer the same question when transaction T1 did not run at all. How is the degree of consistency relevant to serializability?