

# CS3743 Exercise Questions (3)

November 21, 2007

1. [15] Consider a disk drive with the following parameters:

block size	2048 bytes
inter-block gap	128 bytes
# of 2-side platters	5
tracks per surface	300
blocks per track	100
rotation speed	5400 rpm
avg. seek time	6 msec

- (a) Find both the gross and effective total capacity (in byte) of the disk drive.
- (b) Calculate the block transfer time and the average rotational delay.
- (c) Calculate the access time for random and for sequential access of 60 blocks.
2. [20] Consider a file that contains 1,000,000 employee records, where each record has following fields (with their respective sizes in byte in parentheses): EID (12), Name (20), Address (25), City (10), Phone (10), and Salary (32), Department (6), Title (14). The EID field contains unique values. Assume that the block size is 2048 bytes, a block address has 8 bytes and the file is packed into pages *unspanned*.

Answer the following questions and show your calculation.

- (a) What is the total size (in block) of the file?
- (b) Suppose no fast access path (i.e., hashing or indexing) is available, how many block I/Os does it take to find an employee for a given EID? What about for a given name? (Hint: You need to consider possible ways in which records are stored.)
- (c) Suppose a simple (one-level) clustering index on City field is available and there are 5,000 different cities. Find the size (in block) of the index and the number of block I/Os needed to retrieve employee records for a given city.
- (d) Assume a B+ tree primary index is available on EID. The index nodes are 68% full on average, with perhaps the exception of the root. Find the number of levels and the total number of blocks in the B+ tree, and the I/O cost of finding an employee with a given EID.

3. [30] Consider two relations R(A, B, C, D) and S(C, A, E, F). The primary key for R is R.A and that of S is S.C. R.C is a foreign key referencing S.C with no null value. Both relations contains fixed length records and are stored in blocks unspanned. A tuple of R has 120 bytes and a tuple of S has 99 bytes. Block size is 2KB (that is, 2048 bytes). Parameters of the two relations are given in following tables. All indexes are B+ trees.

relation	records	blocks	Index	levels	type	attribute	distinct values
R	1,000,000	58,824	R.A	3	primary	R.B	4,000
S	300,000	15,000	R.B	3	secondary	R.C	15,000
			R.D	2	secondary	R.D	3,000
			S.C	3	primary	S.A	25,000
						S.E	10,000
						S.F	10

- (a) List all possible methods for evaluating the following query, together with their costs in number of I/O blocks (ignore the cost of writing the final results).

```
select *
from R
where A <= 75,000 and B = 'abc' and D = d
```

Assume the values in R.A range from 1 to 1,000,000. Also assumes that the values in R.A, R.B, and R.D are independently distributed in R.

- (b) Consider the join  $R \bowtie_{R.C=S.C} S$ .
- Estimate the size of the result in blocks.
  - Assume the buffer has 202 blocks. Identify the best nested-loop join method for this join, and calculate its I/O cost in blocks (without the cost of writing the final result).

4. [15] Consider the following SQL query.

```
select r.A, s.C, u.C, s.E
from R r, S s, S u
where r.D < 100 and r.C = s.C and s.F = u.F
and u.E = 'a' and r.A = s.A
```

Give the initial and an optimal logical evaluation plans (in both relational algebraic expressions and query trees).

5. [10] Consider the following schedule.

$$S : r_1(X), r_2(Y), r_2(X), w_2(X), r_1(Y), w_1(X), w_1(Y), w_2(Y)$$

- Is  $S$  serializable? Explain your answer.
- Add lock/unlock requests to  $S$  to make it a Strict 2PL schedule and describe its outcome.

6. [10] Assume that a recovery system performs immediate updates and the log contains the following entries immediately before the system crashes.

<T1, Start>  
<T1, X, 5, 25>  
<T2 Start>  
<T2, Z, 10, 15>  
<T2, Commit>  
<Checkpoint>  
<T3, Start>  
<T3, Z, 15, 30>  
<T1, Y, 25, 14>  
<T1, Commit>  
<T3, X, 25, 40>  
<T3, Y, 14, 70>

What would the system do at the database restart and what values would the data have right after the system is restarted?