M.Sc. Examination 2018
Semester-IV
Computer Science
Course : MCSC-42
(Advanced DBMS)

Time : 3 Hours **Full Marks : 40**

Questions are of value as indicated in the margin

Answer Question No. 1 and **any four** from the rest

1. a) Differentiate between dense and space indices.
   b) What are the differences between relational-algebra expression and query execution plan?
   c) Show that there are $(2(n-1))!/(n-1)!$ different join orders with n relations.
   d) Although every conflict serializable schedule is view serializable, then why do we emphasize conflict serializability rather than view serializability? 2×4=8
2. a) State along with brief justifications whether the following schedule of three transactions (T1, T2, and T3) is conflict serializable, view serializable, recoverable, and cascadeless, Further answer whether it may be generated with 2PL and Timestamp without the Thomas' Write Rule. Finally draw the precedence graph for the following schedule.
   T1 : R(A), T1: R(B), T1: W(B), T2 : R(A), T2 : W(A), T2 : R(B), T3 : R(A), T1: R(C), T2 : R(C), T2:W(C), T3: R(B), T1: W(C), T1 : Commit, T3: R(C), T2: Commit, T3 : commit.
   b) What is strict 2PL? 7+1=8
3. a) "View serializability is less stringent than conflict serializability" – Justify.
   b) Briefly explain the redo and undo phases of log based recovery algorithm after a system crash.
   c) Does tree protocol ensure freedom from deadlock? Justify. 2+4+2=8
4. a) Compute the cost of $r \bowtie s$ using i) nested-loop join, ii) block nested-loop join, iii) Merge join, and iv) hash join, where r has 500 tuples with 20 tuples per block, s has 1000 tuples with 10 tuples per block and the buffer for the operation is 8 blocks long.
   b) How to evaluate disjunctive selection (complex selection).
   c) Distinguish between the materialization approach and the pipelining approach of query evaluation. 4+2+2=8
5. a) Explain how to estimate the size of a selection the form $\sigma_{A \le V}(r)$. If a histogram is available on attribute A, then how to use the histogram to get a more accurate estimate?
   b) Write down the equivalence rule of theta joins for associativity. Explain how to estimate the cost of $r(R) \bowtie s(S)$ when $R \cap S$ is a key for neither R *nor* S.
   c) Briefly explain heuristic query optimization. (2+1)+(1+2)+2=8

P.T.O.

6. a) Briefly explain different partitioning techniques in Parallel database.
   b) Briefly explain the parallel fragment-and-replicate join technique.
   c) How can a relation be reconstructed after horizontal and vertical fragmentations? 3+3+2=8
7. Write short notes on **any two** : 4×2=8
   a) Measures of query cost,
   b) Fuzy heckpointing,
   c) Multidimensional data,
   d) Two-phase commit protocol.

———