5 Attempt any two parts of following : 2×10

(a) The MGL protocol states that a transaction T can unlock a node N, only if none of the children of node N are still locked by transaction T. Show that without this condition, the MGL protocol would be incorrect.

(b) Using the bank example, write relational-algebra queries to find the accounts held by more than two customers in the following ways:

(i) Using an aggregate function.

(ii) Without using any aggregate functions.

Let the following relation schemas be given:

R = (A, B, C)

S = (D, E, F)

Let relations r(R) and s(S) be given. Give an expression in the tuple relational calculus that is equivalent to each of the following :

(1) $\Pi_A(r)$

(2) $\sigma_{B=17}(r)$

(3) r × s

(4) $\Pi_{A,F}\left(\sigma_{c=D}(r \times s)\right)$

(c) Assuming a suitable relational schema for a banking system, write an SQL query, without using a with clause, to find all branches where the total account deposit is less than the average total account deposit at all branches.

(i) Using a nested query in the from clause

(ii) Using a nested query in a having clause.

PAPER ID : 1068

TCS-404/TIT-404 Printed Pages : 4

Paper ID and Roll No. to be filled in your Answer Book

Roll No.

B. Tech.

(SEM. IV) (EVEN SEM.) EXAMINATION, 2013

DATABASE MANAGEMENT SYSTEM

Time : 3 Hours] [Total Marks : 100

1 Attempt any four parts of the following : 4×5

(a) List four significant differences between a file-processing system and a DBMS.

(b) Define the concept of aggregation. Give two examples of where this concept is useful.

(c) Consider a two-dimensional integer array of size n × m that is to be used in your favourite programming language. Using the array as an example, illustrate the difference (i) between the three levels of data abstraction, and (i) between a schema and instances.

(d) Construct an E-R diagram for a hospital with a set of patients and a set of medical doctors. Associate with each patient a log of the various tests and examinations conducted.

(e) Draw the overall structure of DBMS and explain its various components.

(f) Define the following terms :

(i) Data redundancy and consistency

(ii) Referential Integrity

(iii) Data atomicity

(iv) Domain constraints

(v) Data models.

2 Attempt any four parts of the following : 4×5

(a) Describe the wait-die and wound-wait protocols for deadlock prevention.

(b) Discuss the selection, projection and join operator of relational algebra with a suitable example.

(c) Discuss two multi-version techniques for concurrency control.

(d) What is a view in SQL, and how is it defined? Discuss the problems that may arise when one attempts to update a view. How views are typically implemented?

(e) Describe conceptually how an SQL query will be executed by specifying the conceptual order of executing each of the six clauses.

(f) Describe the circumstances in which you would choose to use embedded SQL rather than SQL alone or only a general-purpose programming language.

3 Attempt any two parts of the following : 2×10

(a) Define first, second, and third normal forms when only primary keys are considered. How do the general definitions of 2NF and 3NF, which consider all keys of a relation, differ from those that consider only primary keys?

(b) Define Boyce-Codd normal form. How does it differ from 3NF? Why is it considered a stronger form of 3NF?

(c) Consider the universal relation R = {A, B, C, D, E, F, G, H, I, J} and the set of functional dependencies F = {{A, B}} $\rightarrow$ {C}, {A} $\rightarrow$ {D, E}, {B} $\rightarrow$ {F}, {F} $\rightarrow$ {G, H}, {D} $\rightarrow$ {I, J}}. What is the key for R? Decompose R into 2NF and then 3NF relations.

4 Attempt any two parts of the following : 2×10

(a) What is a serial schedule? What is a serializable schedule? Why a serial schedule is considered correct? Why a serializable schedule is considered correct?

(b) What is the system log used for? What are the typical kinds of records in a system log? What are transaction commit points, and why are they important?

(c) Why an explicit transaction end statement is needed in SQL2 but not an explicit begin statement ?