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

PAPER ID : 1068   Roll No.

B. Tech.

(SEM. IV) EXAMINATION, 2008-09

DATA BASE MANAGEMENT SYSTEM

Time : 3 Hours]   *[Total Marks : 100*

Note :   Attempt all questions as per directions given thereof.
              All questions carry equal marks. Be precise in your
              answer. No second answer booklet will be provided.

1   Attempt any **two** of the following :

(a)   What are the different types of joins?

(b)   Suppose we informally define the data in the database
      of a department store as follows :

(i)   Each employee is represented. The data about
       an employee are his employee number, name,
       address, and the department he works for.

(ii)   Each department is represented. The data about
       a department are its name, employees, manager,
       and items sold.

(iii) Each item sold is represented. The data about
       an item are its name, manufacturer, price, model
       number(assigned by manufacturer), and an
       internal item number (assigned by the store)

(iv)   Each manufacturer is represented. The data about
       a manufacturer are its name, address, items
       supplied to the store, and their prices.

Give an entity relationship diagram for this database. Note that some information may be represented by attributes; other information may be represented by relationship.

(c) What are Armstrong rules? How do we say that they are complete and/or sound ?

2   Attempt any **two** of the followings :

(a) Suppose we have a database consisting of the three relations.

FREQUENT(DRINKER, BAR)
SERVES(BAR, BEER)
LIKES(DRINKER, BEER)

The first indicates the bars each drinker visits, the second tells what beers each bar seves, and the last indicates which beers each drinker likes to drink.

(i)   Print the bars that serve a beer that drinker Charles likes.

(ii) Print the drinkers that frequent at least one bar that serves a beer they like.

(iii) Print the drinker that frequent only that serve some beer that they like. (assuming each drinker likes at least one beer and frequent at least one bar).

(iv) Print the drinkers that frequent no bar that serves a beer that they like.

Express the above queries in the Relational Algebra.

(b) Consider the database Schema.
Emp=(ename, setof (Children), setof(Skills)
Children=(Name, Birthday)
Birthday=(Day, Month, Year)
Skill=(Type, Setof (Exams))
Exams=(year, city)
Write the following queries in SQL.

(i)   Find the name of all employees who have a child who has a birthday in March

(ii)   Find those employees who took an examination for the skill type "typing" in the city "Dayton"

(iii)   List all skill types in the relation emp.

(c)   When do you get constraints violation? What are the types of constraints?

3   Attempt any **two** of the following :

(a)   Suppose we have a database for an investment firm, consisting of the following attributes: B (broker), O (Office of the broker), I (investor), S (stock), Q (quantity of stock owned by an investor), and D (dividend paid by a stock), with the functional dependencies $S \rightarrow D$, $I \rightarrow B$, $IS \rightarrow Q$ and $B \rightarrow O$.

(i)   Find a key for the relation schema R=BOSQID.

(ii)   How many keys does relation schema R have ? - Justify.

(iii)   Find a lossless join decomposition of R into Boyce-Codd Normal form.

Find a decompostion of R into third normal form having a lossless join and preserving dependencies.

(b)   "Redundancy of data is many times beneficial". Justify the statement; also describe the situation when redundancy will mess up the current database status, at that instance of time what actions you will prefer to take ?

(c)   Two transactions are not interleaved in a history if every operation of one transaction precedes every operation of the other. Give an example of a serializable history H that has all of the following properties.

(i)   Transactions T1 and T2 are not interleaved in H.

(ii)   T1 precedes T2 in H; and

(iii) In any serial history equivalent to H2, T2 precedes T1.

The history may include more than two transactions.

4 Attempt any **two** of the following :

(a) Compare the shadow paging recovery scheme with the log based recovery in terms of ease of implementation and overhead cost.

(b) What is difference between in-place-updating and shadowing? Discuss the immediate recovery technique in both single user and multi user environments. What are advantages and disadvantages of immediate update?

(c) Under what conditions is it less expensive to avoid deadlock than to allow deadlocks to occur and then to detect them?

5 Attempt any **two** of the followings.

(a) Discuss different variants of 2-phase commit protocol applicable to database management system.

(b) How do optimistic concurrency control techniques differ from other concurrency control techniques? Why are they also called validation or certification techniques? Discuss the backward and forward validation.

(c) Write programs that implement a 2PL scheduler. Compare and contrast S2PL and D2PL.