



THE UNIVERSITY OF THE WEST INDIES

EXAMINATION OF APRIL/MAY 2002

Code and Name of Course: CS35A - Database Management Systems

Paper:

Date and Time: Friday May 3, 2002 1:00 p.m. - 3:00 p.m.

Duration 2 HRS

INSTRUCTIONS TO CANDIDATES: This paper has 3 page(s) and 4

questions

Instructions: Answer Question 1 and any two others.

Question 1

[40 Marks]

- a) You have been asked to design a database for a Car Insurance Company

A car insurance company needs a database to streamline its operations. The insurance company need to keep track of all its customers, their name, address phone number. A customer can insure a number of cars at the company. Each car must be uniquely identified and the information kept must include its licence number, make, model, colour, value. The type of insurance e.g. third party or comprehensive and the date that it was taken out must also be recorded.

A log must be kept of the accidents that each car has had. For each accident a record of the estimate damage, the date of the accident, the amount paid, and a brief description of the accident is kept.

The company employs a number of agents that handle the files. Each insurance file is assigned to one of these agents. The company needs to keep track of the employees name, address, contact number, position, and salary. All employees of a particular position are paid the same salary. They also need to know at any point in time which agent is responsible for which file.

- i Draw an Entity-Relationship Diagram for the information above. State any assumptions that you made and any constraints that you were unable to represent.
[9 Marks]
- ii Convert the diagram into a set of relations normalised up to at least third normal form, and where possible BCNF. You must give reasons for converting the relations in the form that you did.
[6 Marks]

Question 1 (con't)

- b) Consider the following set of relations for a hospital; information is kept on the doctors, patients, salary scale and, for each patient, all the visits they have made. The patient will only see one doctor per visit. **[10 Marks]**

Doctor(Did, Dname, Specialization)
SalaryScale(Specialization, Salary)
Patient(Pid, Pname, Pphone)
Vistis(Pid, Did, Vdate, Problem, VCost)

Formulate the following queries in SQL:

- i Find all the problems that 'Mary Smith' has had.
 - ii Find the name and specialization of all the doctors that patient 'John Brown' has seen.
 - iii For each patient find the total amount of money they have paid to the hospital.
 - iv Find the name of all doctors who earn more than the average salary of all doctors.
 - v Increase the salary of 'Dr. John Brown' by 10%.
- c) Assume that the log of a DBMS has the following entries:

<T1, BEGIN>
<T1, X, 500, 600>
<T1, Y, 1000, 1500>
<T1, COMMIT>
<T2, BEGIN>
<T2, P, 300, 200>
<T3, BEGIN>
<T3, X, 600, 400>
<CHECKPOINT>
<T4, BEGIN>
<T4, Y, 1000, 500>
<T2, Q, 600, 800>
<T5, BEGIN>
<T5, Z, 900, 700>
<T2, COMITT>
<T4, COMITT>

-----SYSTEM FAILURE-----

N.B. For an entry of the form: <T1, X, 500, 600>: the first entry is the name of the transaction, the second is the name of the data item, the third is the old value of the data item and the fourth is the new value of the data item.

At restart time, what actions would be performed by the transaction manager to ensure atomicity i.e. explain, for each transaction, what would need to be done to ensure that the database remains consistent. **[10 Marks]**

- d) It is well understood that redundancy is a bad thing in databases, why then in distributed systems is replication used? **[5 Marks]**

Question 2

- a) Describe the undesirable properties that a database that has not been normalised up to BCNF would exhibit. [4 Marks]
- b) What four properties should a decomposition ideally have? Discuss. [6 Marks]

Question 3

- a) Using the same set of tables in Question 1 b) translate the following into **both** relational algebra and relational calculus queries: [7 Marks]
- i Find all the problems that 'Mary Smith' has had.
 - ii Find the name and specialization of all the doctors that patient 'John Brown' has seen.
- b) Describe, using a relational algebra expression, how doing selections as early as possible can improve the efficiency of query processing. [3 Marks]

Question 4

- a) Describe how shared and exclusive locks are used to enforce concurrency control. Explain, using an example, the problem(s) introduced by these locks. [6 Marks]
- b) *Distribution makes concurrency control even more difficult.* Explain. [4 Marks]