UNIVERSITY COLLEGE LONDON

University of London

EXAMINATION FOR INTERNAL STUDENTS

For the following qualifications :-

M. Sc.

D6: Database Systems

COURSE CODE	:	COMPOOD6
DATE	:	18-MAY-01
TIME	:	14.30
TIME ALLOWED	:	2 hours 30 minutes

Answer any three out of the following five questions

1. Answer the following three parts.

   a. What is a data model ? Illustrate your answer using the relational data model.

   [11 marks]

   b. By default SQL does not remove duplicate rows. Discuss the main reason for this approach and illustrate your answer with an SQL query.

   How can you remove duplicate rows in SQL ?

   [11 marks]

   c. Write a sentence or two explaining each of the following:

      1. physical database
      2. growth independence
      3. DDL and DML
      4. DBA
      5. Web information retrieval

   [11 marks]

[Total 33 marks]

2. Answer the following three parts.

a. Describe how an Entity Relationship. Diagram can be translated into a relational schema, illustrating your answer by a simple example.

[11 marks]

b. Show how ISA relationships can be used to model inheritance between entity types.

How can the Universal Relation Schema Assumption (URSA) help disambiguate naming conflicts due to multiple inheritance ?

[11 marks]

c. Assume that we allow in relations null values of type "value exists but is unknown". For example, we may not know the age or salary of an employee.

Using simple examples illustrate how the AVG and COUNT aggregate functions should work in the presence of such null values.

[11 marks]

[Total 33 marks]

3. Answer the following three parts.

   a. Briefly argue why the relational algebra (or equivalently SQL) cannot express the transitive closure operation.

      Why do you think the designers of SQL decided to exclude the transitive closure operation from the language ?

      [11 marks]

   b. Given a fixed relation, say r, over a relation schema, R, whose attributes are PARENT and CHILD give a simple algorithm which computes the transitive closure of the relation r.

      You are allowed to use the following relational algebra operators in your algorithm: projection, natural join, union and renaming.

      [11 marks]

   c. What are the fundamental differences between the relational algebra and SQL ?

      Explain briefly why SQL, as opposed to the relational algebra, has become the industry standard relational database query language.

      [11 marks]

[Total 33 marks]

4. Answer the following three parts.

a. Briefly explain the main difference between Third Normal Form (3NF) and Boyce-Codd Normal Form (BCNF) through a simple example.

Explain why in practice, a database designer might be content with a 3NF database schema that is not in BCNR

[11 marks]

b. A simple key is a key which consists of only a single attribute.

Give a proof (or a convincing argument) of the assertion that a relation schema R that is in Third Normal Form (3NF) with respect to a set of functional dependencies F over R and such that all the keys for R with respect to F are simple is also in Boyce-Codd Normal Form (BCNF) with respect to F.

[11 marks]

c. Two sets of FDs F and G both over a relation schema R are said to be logically equivalent if they are covers of each other, i.e. the closure of F with respect to Armstrong's axiom system is equal to the closure of G, also with respect to Armstrong's axiom system.

Give an efficient algorithm for testing whether a set of FDs F over R is logically equivalent to a set of FDs G over R.

[11 marks]

[Total 33 marks]

5. Answer the following three parts.

a. What are the two types of null value that are most common ? (Give an example of their use.)

To what extent does SQL support these types of null values ?

[11 marks]

b. Assume a relation schema in the database containing personal information, with primary key PersonID, and additional attributes such as PersonName, PersonAdrress and PersonEmail.

Suppose that, for each person who is married, you would also like to record who their spouse is and also their spouse's personal information. Suggest two different ways of modelling this situation.

[11 marks]

c. Prove the soundness of the decomposition inference rule using the reflexivity, augmentation and transitivity inference rules of Armstrong's axiom system.

Recall that the decomposition rule states that: if F I- X -> YZ then F ~- X -> Y and F |- X ---> Z, where F is a set of functional dependencies over R and X, Y, Z are subsets of schema(R).

[11 marks]

[Total 33 marks]

END OF PAPER