Answer any three out of the following five questions

1. Answer the following three parts.

   a. What are the main reasons that the relational model has become so successful ?

   [11 marks]

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

   [11 marks]

   c. Briefly discuss the differences between finding information in a relational database using SQL and finding information in the World-Wide-Web using a search engine.

   [11 marks]

[Total 33 marks]

2. Answer the following three parts.

   a. What is the Universal Relation Schema Assumption (URSA) ? Discuss the importance of the URSA in the context of ISA relationships.

   [11 marks]

   b. Why do multi-valued attributes in an Entity Relationship Diagram (ERD) violate First Normal Form ? Show with the aid of an ERD fragment how you would express a multi-valued attribute such as a set of authors using only single-valued attributes.

   [11 marks]

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

      1. physical data independence
      2. growth independence
      3. optional and mandatory classification of relationships
      4. COMMIT and ROLLBACK
      5. precision and recall

   [11 marks]

[Total 33 marks]

3. Answer the following three parts.

   a. In the network and hierarchical data models records are linked together via pointers and querying the database is done via “pointer chasing”. How are entities related in the relational model ? Discuss the advantages of querying related entities in the relational model versus pointer chasing as a means of navigating through a database.

   [11 marks]

   b. Express the natural join of two relations, one over a schema having attributes A and B and the other over a schema having attributes B and C, using the renaming, Cartesian product, selection and projection operators. Express this join query in SQL.

   [11 marks]

   c. What are the two types of null value that are most common ? Give an example of how information can be lost when joining two relations having null values in the joined columns and suggest a solution.

   [11 marks]

[Total 33 marks]

4. Answer the following three parts.

   a. Recall that the manipulative part for a *nested relational database* includes the two operators $NEST$ and $UNNEST$. The operator $NEST$ transforms a nested relation into a "more deeply" nested relation while the operator $UNNEST$ transforms a nested relation into a "flatter" nested relation. Explain with the aid of an example how these operators are used, given that you have available a flat relational database.

   [11 marks]

   b. In the presence of NULLs (interpreted as in SQL) we can redefine a functional dependency (FD) X $\rightarrow$ Y to hold in a relation r over relation schema R if any two rows in r that do *not* have a NULL value in any attribute of X and, in addition, have the same X-values, also have the same Y-values. Is the transitivity inference rule sound for such FD with NULLs ?

   [11 marks]

   c. Define Boyce-Codd Normal Form (BCNF) using an example. Why is BCNF sometimes referred to as the "ultimate normal form" in the presence of functional dependencies.

   [11 marks]

   [Total 33 marks]

5. Answer the following three parts.

   a. Give an efficient (i.e. polynomial time) algorithm for finding one key for a relation schema R with respect to a set of functional dependencies F over R.

   [11 marks]

   b. Argue that if a relation schema R is in Third Normal Form but not in Boyce-Codd Normal Form with respect to a set of functional dependencies F, then it must have at least two distinct keys for R with respect to F which overlap, i.e. such that their intersection is nonempty.

   [11 marks]

   c. Why do we need object-relational databases ?

   [11 marks]

[Total 33 marks]

END OF PAPER