:Course	**: B.Sc. (Hons) Computer Science IV sem (CBCS)**
Paper Name	**: Database Management Systems**
Unique Paper code	**: 32341403_OC**
Max Marks	**: 75**
Time	**: 3 hours**

Instructions:

1. Attempt any FOUR questions
2. All questions carry equal marks
3. The complete answer to a question MUST be uploaded in the form of a single PDF file.

Q1

The XYZ Co. deals with several pharmaceutical companies. Its data requirements include the following specifications:

- Patients are identified by SSN, and their names, addresses, and also age.
- Doctors are identified by SSN. For each doctor, the name, specialty and years of experience must be recorded.
- Each pharmaceutical company is identified by name and has a phone number.
- For each medicine, the trade name and formula must be recorded. Each medicine is sold by a given pharmaceutical company, and the trade name identifies a medicine uniquely from among the products of that company. If a pharmaceutical company is deleted, you need not keep track of its products any longer.
- Each pharmacy has a name, address, and phone number.
- Each pharmacy sells several drugs and has a price for each. A medicine could be sold at several pharmacies, and the price could vary from one pharmacy to another.
- Doctors prescribe medicine for patients. A doctor would prescribe one or more medicine for several patients, and a patient could obtain prescriptions from several doctors. Each prescription has a date and a quantity associated with it.

Draw an ER diagram that captures the above information. Clearly state any constraints that you assume. Also, extend the above ER diagram to EER using one specialization.

Q2

Consider the relations given below:

Doctor(SSN,Firstname,Lastname,Speciality,YearsOfExp,
PhoneNum)

Patient(SSN,Firstname,Lastname,Address,DOB,PrimaryDoc_SSN)

Medicine(TradeName,UnitPrice,GenericFlag)

Prescription(ID,Date,Doctor_SSN,Patient_SSN)

Prescription_Medicine(PrescriptionID,TradeName,Quantity)

where

`Medicine.GenericFlag` represents whether or not the medicine is generic (True or False).
`Patient.PrimaryDoc_SSN` is a foreign key to **`Doctor.SSN`**
`PrescriptionID` of **`Prescription_Medicine`** relation is a foreign key to **`ID`** attribute of **`Prescription`** relation.
`Prescription_Medicine.TradeName` refers to **`Medicine.TradeName`**
`Precription.Doctor_SSN` and **`Prescription.Patient_SSN`** are foreign keys to **`Doctor.SSN`** and **`Patient.SSN`** respectively.

Write the following queries in SQL and Relational Algebra.

- List the **`TradeName`** of generic medicine with unit price less than Rs 150.
- List the names of patients whose primary doctor is **`'Rakesh Sharma'`**
- List the names of doctors who are not primary doctors to any patient
- For medicines written in more than 20 prescriptions, report the trade name and the total quantity prescribed.
- List the **`SSN`** of patients who have **`'Paracetamol'`** and **`'Vitamin'`** trade names in one prescription
- List the **`SSN`** of distinct patients who have **`'Paracetamol'`** prescribed to them by a doctor named **`'Rakesh Sharma'`**.
- List the first and last name of patients who have no prescriptions written by doctors other than their primary doctors.

Q 3

Map the following E-R Diagram to relational schema :

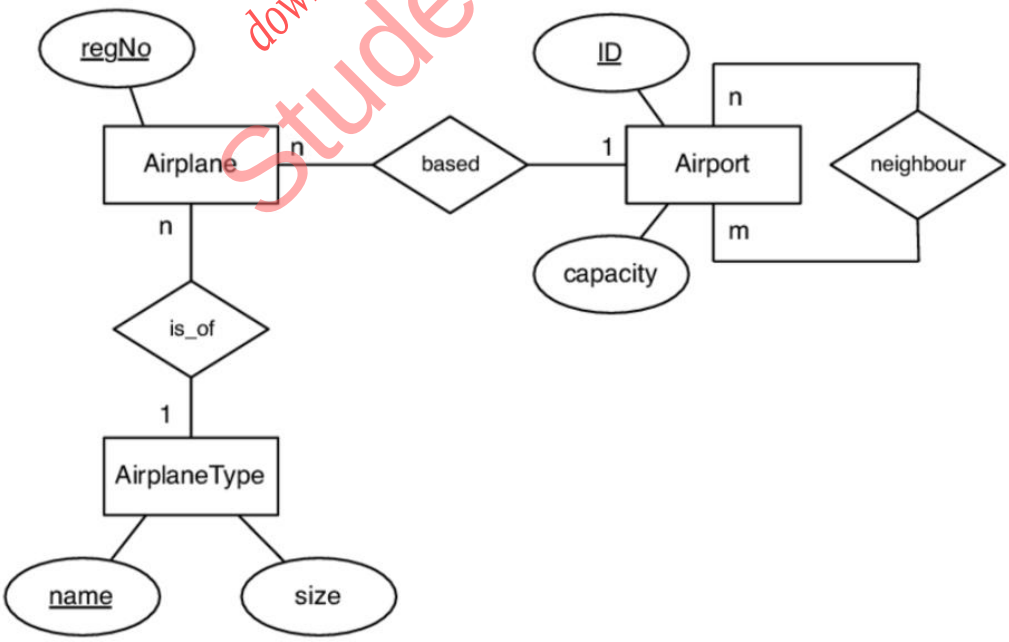

Now, write the **CREATE TABLE** command for all of the above tables in SQL ensuring that the following concepts are used at least once: integer, string, and date data, **NOT NULL** constraint, **CHECK** constraint, **PRIMARY KEY** constraint, **FOREIGN KEY** constraints (with **ON DELETE SET NULL** and **ON UPDATE CASCADE** constraints, if applicable).

Also, describe the role of recursive relationships in the context of the above diagram.

Q4

Given below are two sets of FDs for a relation **R(A,B,C,D,E)** :

F : A→ B, AB →C, D → AC, D→ E
G : A→ BC, D→ AE

Are F and G equivalent? Find the minimal cover for the set of dependencies F.

Now, Consider another relation **R = {A,B,C,D,E,F,G,H,I,J}** with the following set of FDs

F = { AB→C, A→DE, B→F, F→GH, D→IJ }

Find the key. Which highest normal form is this relation in? Why? Decompose it into 2NF and 3NF relations. Is this decomposition dependency preserving? Justify your answer.

Q5

A file comprising employee records has **emp_ID** as the primary key. It is searched with the help of B+ tree index with **order p = 3** for internal nodes and **order p_{leaf} = 2** for leaf nodes.

- If the values of **emp_ID** are inserted in the following order, show how the tree will expand and how many times the leaf node will split up.
  216, 182, 333, 115, 160, 235, 218, 332
- Show the tree after deleting the employees with **emp_ID** numbers 216, 182 and 333. Show each step.
- What will be the number of block accesses in the above indexing scheme?

Q6

How does DDL support the implementation of the three-schema architecture?

Consider the following two relations **R1** and **R2**:

R1		
P	**Q**	**S**
10	a	5
15	b	8
5	a	6

R2		
A	**B**	**C**
10	b	6
25	c	1
10	b	5

Show the result of the following relational queries:

- $R1 \bowtie_{R1.P=R2.A} R2$
- $R1 \ _{R1.Q=R2.B} \ R2$
- $R1 * R2$
- $R1 \bowtie_{(R1.P=R2.A \text{ AND } R1.R=R2.C)} R2$

Consider a relational schema comprising three relations as follows:

S represents Sailors(sid, name, rating, age)

B represents Boats (bid, name, color)

R represents Reserves (sid, bid, day, rname)

Draw a query tree to show a possible order of execution for the following relational expression:

$$\pi_{sid,sname,age}(\sigma_{age<30}(\sigma_{color='red'}(\sigma_{bid=bid}(B \times \sigma_{sid=sid}(S \times R)))))$$