

Question Paper

Exam Date & Time: 30-Apr-2024 (02:30 PM - 05:30 PM)



MANIPAL ACADEMY OF HIGHER EDUCATION

FOURTH SEMESTER B.TECH. (INFORMATION TECHNOLOGY) DEGREE EXAMINATIONS - APRIL / MAY 2024
SUBJECT: ICT 2221/ICT_2221- DATABASE SYSTEMS

Marks: 50

Duration: 180 mins.

Answer all the questions.

- 1A) Consider two entities: Student and Course. Each student can enroll in multiple courses. (5)
Conversely, each course can have multiple students.
i) Describe how an ER schema can be converted to tables for above relationship. Also explain how any weak entity set is converted to a strong entity set with above scenario.
ii) Let $R=(A,B,C)$ be the relational schema with the functional dependency $F=\{A \rightarrow B, B \rightarrow C\}$.
Decomposition of R : $R_1=(A,B)$, $R_2=(B,C)$.
Does this decomposition preserve the given dependencies?
(3+2 = 5 marks)

- 1B) Construct B+ Tree which is order of 3 and values are 1,4,7,10,17,21,31 show the result in step-by-step process. (3)

- 1C) Consider a double hashing scheme in which the primary hash function is $h_1(k) = k \bmod 23$, and the secondary hash function is $h_2(k) = 1 + (k \bmod 19)$. Assume that the table size is 23. Then the address returned by probe 1 in the probe sequence (assume that the probe sequence begins at probe 0) for key value $k = 90$ is _____. (2)

- 2A) Consider the following employee database. (5)
EmployeeInfo Table:

EmpID	EmpFname	EmpLname	Department	Project	Address	DOB	Gender
-------	----------	----------	------------	---------	---------	-----	--------

EmployeePosition Table:

EmpID	EmpPosition	DateOfJoining	Salary
-------	-------------	---------------	--------

- i) Write a query to fetch all employees who also hold the managerial position.
ii) Write a SQL query to retrieve employee details from EmployeeInfo table who have a date of joining in the EmployeePosition table.
iii) Write a query to fetch details of employees with the address as "DELHI(DEL)".
iv) Write a query to retrieve duplicate records from a table.
v) Write a query to retrieve EmpPosition along with total salaries paid for each of them.
- 2B) Consider the following transaction T_1 : $R(X)$, $R(X)$. Give an example of another transaction T_2 that, if run concurrently to transaction T_1 without some form of concurrency control, could interfere with T_1 to produce unrepeatable reads. Show the sequence of operations which would cause the problem. (3)
- 2C) Consider a database with the following schema: (2)
- | | |
|---|----------------------------|
| Person (<u>name</u> , age, gender) | name is a key |
| Frequents (<u>name</u> , <u>pizzeria</u>) | (name, pizzeria) is a key |
| Eats (<u>name</u> , <u>pizza</u>) | (name, pizza) is a key |
| Serves (<u>pizzeria</u> , <u>pizza</u> , price) | (pizzeria, pizza) is a key |
- Write relational algebra expressions for the following queries.

- i) Find the names of all females who eat both mushroom and pepperoni pizza.
 ii) Find the names of all people who frequent only pizzerias serving at least one pizza they eat.
- 3A) i) Assume we have the following application that models soccer teams, the games they play, and the players in each team. In the design, we want to capture the following: (5)
- We have a set of teams, each team has an ID (unique identifier), name, main stadium, and to which city this team belongs.
 - Each team has many players, and each player belongs to one team. Each player has a number (unique identifier), name, DoB, start year, and shirt number that he uses.
 - Teams play matches, in each match there is a host team and a guest team. The match takes place in the stadium of the host team.
 - For each match we need to keep track of the following:
 - o The date on which the game is played
 - o The final result of the match
 - o The players participated in the match. For each player, how many goals he scored, whether or not he took yellow card, and whether or not he took red card.
 - o During the match, one player may substitute another player. We want to capture this substitution and the time at which it took place.
 - Each match has exactly three referees. For each referee we have an ID (unique identifier), name, DoB, years of experience. One referee is the main referee and the other two are assistant referee.
- Design an ER diagram to capture the above requirements. State any assumptions you have that affects your. Make sure cardinalities and primary keys are clear. Also reduce the ER diagram to the Schema Diagram.
- ii) Consider the following functional dependencies, $F = \{AB \rightarrow C, AB \rightarrow D, C \rightarrow A, D \rightarrow B\}$ on the Relation $R(A, B, C, D)$. Decompose the R in to the highest Normal form possible.
- 3B) Identify all the drawbacks of Lock based protocol and demonstrate with an example a solution for each of the drawbacks. (3)
- 3C) Consider the concurrent schedule S: $R1(X), W2(Y), R2(Y), W1(Z), R2(K), W1(Y), W2(X)$. Show the working of the Validation based protocol on the schedule S. (2)
- 4A) i) Consider the Relation $R = (A, B, C, D)$, and set of functional dependencies $F = \{C \rightarrow D, C \rightarrow A, B \rightarrow C\}$. Check in which Normal Form the R is in? Further, decompose R in to the highest Normal Form possible. (5)
- ii) Consider the Relation $R(A, B, C, D, E)$ with the functional dependencies $BC \rightarrow ADE$ and $D \rightarrow B$. Find out all the possible candidate keys.
- 4B) Consider you are designing a database system that employs extendable hashing for efficient data storage and retrieval. The extendable hashing mechanism uses a dynamic directory to manage buckets, where each bucket can store up to 2 records. The system uses a hash function that considers the least significant bits (LSBs) of a record's key, with the number of bits equal to the current global depth (GD) of the directory. Initially, the global depth is set to 1, meaning the directory starts with 2 entries. Given the following sequence of record keys to be inserted into the database: 2, 3, 10, 11, 16, 20. Show the step by step process for the above insertion along with the content of Buckets and Bucket address table. (3)
- 4C) With a suitable example demonstrate the different Deadlock handling strategies, and also mention their drawbacks. (2)
- 5A) Consider a database schema given below for a **weather monitoring and prediction system** used by agricultural enterprises. This system collects weather data, analyzes it, and provides insights and forecasts to help farmers make informed decisions about crop management. (5)
- WeatherStations:** This table stores information about weather stations deployed in various locations.
- StationID (Primary Key)
 Location
 Latitude (Numeric field)
 Longitude (Numeric field)
 Elevation (Numeric field, in meters)
 InstallationDate
- WeatherData:** This table stores hourly weather data collected by weather stations.
- DataID (Primary Key)

StationID (Foreign Key referencing WeatherStations table)
 Timestamp
 Temperature (Numeric field, in Celsius)
 Humidity (Numeric field, percentage)
 Precipitation (Numeric field, in millimeters)
 WindSpeed (Numeric field, in meters per second)
 WindDirection (Categorical field, e.g., N, NE, E, SE, S, SW, W, NW)
 Pressure (Numeric field, in hectopascals)

Crops: This table stores information about different crops grown by farmers.

CropID (Primary Key)
 Name
 PlantingSeasonStart (Numeric field, day of the year)
 PlantingSeasonEnd (Numeric field, day of the year)
 HarvestingSeasonStart (Numeric field, day of the year)
 HarvestingSeasonEnd (Numeric field, day of the year)
 IdealTemperatureStart (Numeric field, in Celsius)
 IdealTemperatureEnd (Numeric field, in Celsius)
 IdealHumidityStart (Numeric field, percentage)
 IdealHumidityEnd (Numeric field, percentage)
 CropPlantings: This table stores information about specific crop plantings by farmers.

PlantingID (Primary Key)
 CropID (Foreign Key referencing Crops table)
 StationID (Foreign Key referencing WeatherStations table)
 PlantingDate
 HarvestingDate
 ExpectedYield (Numeric field, in tons per hectare)

CropSales: This table stores information about sales of crops by farmers.

SaleID (Primary Key)
 CropID (Foreign Key referencing Crops table)
 SaleDate
 QuantitySold (Numeric field, in kilograms)
 SalePricePerUnit (Numeric field, in currency)

SoilMoisture: This table stores information about soil moisture levels measured at different depths.

MoistureID (Primary Key)
 StationID (Foreign Key referencing WeatherStations table)
 Timestamp
 Depth (Numeric field, in centimeters)
 MoistureLevel (Numeric field, percentage)

- i) Write a SQL query to find the average temperature recorded by each weather station in the past week, along with the highest wind speed recorded during the same period.
- ii) Write a SQL query to find the top most crops with the highest expected yield among all plantings in each weather station's location.

- | | | |
|-----|--|-----|
| 5B) | Write a relational algebra to calculates the total sales amount for each crop sold, along with the average precipitation recorded during the month of sale on the Database Schema of weather monitoring and prediction system. | (3) |
| 5C) | With the suitable example discuss about the Static Hashing technique and it's deficiencies. | (2) |

-----End-----