U.S.N.										

B.M.S. College of Engineering, Bengaluru-560019

Autonomous Institute Affiliated to VTU

August 2024 Supplementary Examinations

Programme: B E **Semester: IV**
Branch: Artificial Intelligence and Machine Learning **Duration: 3 hrs.**
Course Code: 20AM4PCDBM **Max Marks: 100**
Course: Data Base Management Systems

Instructions: 1. Answer any FIVE full questions, choosing one full question from each unit.
2. Missing data, if any, may be suitably assumed.

UNIT - I

1 a) Consider the database worksOn to answer the following queries: **10**

WorksOn Database:

emp (eno, ename, bdate, title, salary, dno)

proj (pno, pname, budget, dno)

dept (dno, dname, mgreno)

workson (eno, pno, resp, hours)

i. Write an SQL query that returns the employee name, project name, employee title, and hours for all works on records.
ii. Write an SQL query that returns the employee numbers and salaries of all employees in the 'Consulting' department ordered by descending salary.
iii. Write an SQL query that returns the project name, department name, and budget for all projects with a budget < $50,000.
iv. Write an SQL query that returns the project name, hours worked, and project number for all works on records where hours > 10.
v. Write an SQL query that returns the employees (name only) in department 'D1' ordered by decreasing salary.

b) Discuss the main categories of data models. What are the basic differences between the relational model, the object model, and the XML model? What is the difference between a database schema and a database state? **10**

OR

2 a) What is the difference between logical data independence and physical data independence? Which one is harder to achieve? Why? **4**

b) Discuss the different types of user-friendly database languages and interfaces and the types of users who typically use each. **6**

c) Consider the salesman database to answer the following questions: **10**

salesman

salesman_id	name	city	commission
5001	James Hoog	New York	0.15
5002	Nail Knite	Paris	0.13
5005	Pit Alex	London	0.11
5006	Mc Lyon	Paris	0.14
5003	Lauson Hen		0.12
5007	Paul Adam	Rome	0.13

customer

customer_id	customer_name	city	grade	salesman_id
3002	Nick Rimando	New York	100	5001
3005	Graham Zusi	California	200	5002
3001	Brad Guzan	London		
3004	Fabian Johns	Paris	300	5006
3007	Brad Davis	New York	200	5001
3009	Geoff Camero	Berlin	100	
3008	Julian Green	London	300	5002
3003	Jozy Altidor	Moncow	200	5007

order

order_no	purch_amt	order_date	customer_id	salesman_id
70001	150.5	2016-10-05	3005	5002
70009	270.65	2016-09-10	3001	
70002	65.26	2016-10-05	3002	5001
70004	110.5	2016-08-17	3009	
70007	948.5	2016-09-10	3005	5002
70005	2400.6	2016-07-27	3007	5001
70008	5760	2016-09-10	3002	5001
70010	1983.43	2016-10-10	3004	5006
70003	2480.4	2016-10-10	3009	
70012	250.45	2016-06-27	3008	5002
70011	75.29	2016-08-17	3003	5007

i. Find those salesmen with all information who gets the commission within a range of 0.12 and 0.14.

ii. Find the name and city of those customers and salesmen who lives in the same city.

iii. Find the names of all customers along with the salesmen who works for them.

iv. Display all those orders by the customers not located in the same cities where their salesmen live.

v. Display all the orders issued by the salesman 'Paul Adam' from the orders table.

UNIT - II

3 a) You have just moved to Washington, D.C., to work for the U.S. House of Representatives (the "House") as a database specialist. For your first job, they want you to design a database to keep track of votes taken in the House. (A counterpart to you working for the U.S. Senate is designing a similar database for tracking votes in the Senate.) **10**

The database will track votes taken in the House during the current two-year congressional session. It should record each U.S. state (for instance, Texas) with its name, number of representatives, and region in which the state is located (northeast, mid-atlantic, midwest, and so forth). Each congress-creature, er, I mean representative, in the House is to be described by his or her name, the district (by district number) that he or she represents, the year when he or she was first elected, and the political party to which he or she belongs (for instance, Republican, Democrat, Independent, Green, Reform, Other). The database should track each bill (legislation) with its name, the date on which its vote was taken, whether the bill passed or failed (so the domain is yes and no), and its sponsors (the representatives who proposed the bill). The database should track how each representative voted on each bill (yes, no, abstain, absent).

i. Design an entity-relationship (E-R) schema diagram for the above enterprise. Be careful to ensure that each of the attributes would be restricted to legal values (no pun. . .). State clearly any assumptions that you make. Also state any business rules—logic to which the database should adhere—that are not captured in your E-R diagram.

ii. Do you have any derived attributes in your diagram? Would these be problematic in the implemented database?

b) Define the following terms: entity, attribute, attribute value, relationship instance, composite attribute, multivalued attribute, derived attribute, complex attribute, key attribute, and value set (domain). **10**

UNIT - III

4 a) Given the below relations, Employee and Department.

Employee	Department
Ram	Production
Raju	Sales
Ramu	Production

Department	Head
Production	Brown
Purchasing	Blak

Find the following and justify your answer.

i. Cartesian Product
ii. Natural Join
iii. Full Outer Join
iv. Left Outer Join
v. Right Outer Join

b) Consider the following Relations: **5**

Film (Name, Producer, Fyear, Rating)
Artists (Artist, Ayear)
Plays (Artist, Name)
Producers (Producer, Pyear)
Write the Relational Algebra for the following queries:

1) Find films made before 1920 and rating greater than 5.0
2) Find films acted by artist 'ABC' and produced by 'XYZ'
3) Find all films and ratings
4) Find all artists and producers
5) Display the name and year of the film produced by 'XYZ' after 1990

c) Consider the following relational database schema consisting of the four relation schemas: **10**

passenger (**pid,** pname, pgender, pcity)
agency (**aid,** aname, acity)
flight (**fid,** fdate, time, src, dest)
booking (**pid, aid, fid,** fdate)
Answer the following questions using relational algebra queries;

i. Get the complete details of all flights to New Delhi.
ii. Get the details about all flights from Chennai to New Delhi.
iii. Find only the flight numbers for passenger with pid 123 for flights to Chennai before 06/11/2020.

[Hint: Given conditions are pid, dest, and fdate. To get the flight id for a passenger given a pid, we have two tables flight and booking to be joined with necessary conditions. From the result, the flight id can be projected]

iv. Find the passenger names for passengers who have bookings on at least one flight.
v. Find the passenger names for those who do not have any bookings in any flights.

[Hint: here applied a set difference operation. The set difference operation returns only pids that have no booking. The result is joined with passenger table to get the passenger names.]

vi. Get the details of flights that are scheduled on both dates 01/12/2020 and 02/12/2020 at 16:00 hours.

[Hint: the requirement is for flight details for both dates in common. Hence, set union is used between the temporary relations generated from application of various conditions.]

UNIT - IV

5 a) Examine the table shown below. **10**

i. Why is this table not in 2NF?

ii. Describe and illustrate the process of normalizing the data shown in this table to third n (3NF).

iii. Identify the primary, (alternate) and foreign keys relation

staffNo	*branchNo*	*branchAddress*	*name*	*position*	*hoursPerWeek*
S4555	B002	City Center Plaza, Seattle, WA 98122	Ellen Layman	Assistant	16
S4555	B004	16 – 14th Avenue, Seattle, WA 98128	Ellen Layman	Assistant	9
S4612	B002	City Center Plaza, Seattle, WA 98122	Dave Sinclair	Assistant	14
S4612	B004	16 – 14th Avenue, Seattle, WA 98128	Dave Sinclair	Assistant	10

b) Discuss four informal guidelines with examples that may be used as measures to determine the quality of relation schema design **10**

UNIT - V

6 a) Illustrate two-phase locking with a schedule containing three transactions. Argue that 2PL ensures serializability. Also, argue that 2PL can lead to deadlock. **12**

b) Consider the following two transactions: **8**

T1: read(A);

read(B);

if A = 0 then B := B + 1;

write(B).

T2: read(B);

read(A);

if B = 0 then A := A + 1;

write(A).

Let the consistency requirement be A = 0 ∨ B = 0, with A = B = 0 the initial values.

i. Show that every serial execution involving these two transactions preserves the consistency of the database.

ii. Show a concurrent execution of T1 and T2 that produces a non serializable schedule.

iii. Is there a concurrent execution of T1 and T2 that produces a serializable schedule?

OR

7 a) During execution, a transaction passes through several states, until it finally commits or aborts. List all possible sequences of states through which a transaction may pass. Explain why each state transition may occur. Also, explain with an example the properties that must be satisfied by a transaction. **10**

b) Consider the following two transactions: **5**

T31: read(A);
read(B);
if A = 0 then B := B + 1;
write(B).
T32: read(B);
read(A);
if B = 0 then A := A + 1;
write(A).

Add lock and unlock instructions to transactions T31 and T32, so that they observe the two-phase locking protocol. Can the execution of these transactions result in a deadlock?

c) Are the following two schedules conflict equivalent? **5**

T1	T2
R (A)	
W (A)	
	R (A)
	W (A)
R (B)	
W (B)	

Schedule S1

T1	T2
R (A)	
W (A)	
R (B)	
W (B)	
	R (A)
	W (A)

Schedule S2

******