Roll No. | | | | | | | | | | |

B.E (FT) END SEMESTER EXAMINATIONS – MAY / JUNE 2024

Computer Science and Engineering

Fourth Semester

CS6106 & DATABASE MANAGEMENT SYSTEMS

(Regulation 2018 - RUSA)

Time: 3 Hours | Answer ALL Questions | Max. Marks 100

CO 1	Model an application's data requirements using conceptual modeling tools like ER diagrams and design database schemas based on the conceptual model
CO 2	Formulate solutions to a broad range of query problems using relational algebra/ SQL
CO 3	Demonstrate an understanding of normalization theory and apply such knowledge to the normalization of a database
CO 4	Run transactions and estimate the procedures for controlling the consequences of concurrent data access
CO 5	Discuss the basic database storage structures and access techniques: file and page organizations, indexing methods including B-tree, and hashing
CO 6	Point out the basics of query evaluation techniques and query optimization

BL – Bloom's Taxonomy Levels

(L1 - Remember, L2 - Understand, L3 - Apply, L4 - Analyze, L5 - Evaluate, L6 - Create)

PART- A (10 x 2 = 20 Marks)

(Answer all Questions)

Q. No	Questions	Marks	CO	BL
1	What is physical data independence? Why is it important in database systems?	2	1	L4
2	List out the constraints on Generalizations in ER model.	2	1	L2
3	What are the integrity constraints used in designing a relational database?	2	2	L1
4	Compare embedded SQL versus dynamic SQL.	2	2	L4
5	List out the ACID properties.	2	4	L1
6	Give the distinction between serial schedule and serializable schedule.	2	4	L4
7	What are the factors to be considered to determine which transaction to roll back to break the deadlock?	2	4	L1
8	What are the main heuristics that are applied during query optimization?	2	6	L3
9	When it is preferable to use a dense index rather than a sparse index?	2	4	L4
10	List the guidelines for database tuning.	2	5	L1

PART – B (8 x 8 = 64 marks)

(Answer any 8 questions)

Q. No	Questions	Marks	CO	BL
11	Illustrate the system structure of database and explain its components.	8	1	L6

12	Construct an E-R diagram for keeping track of the exploits of India men's national cricket team. You should store the matches played, the scores in each match, the players in each match, and individual player statistics for each match.	**8**	1	L6
13	Consider the following relations: Suppliers(sid:integer,sname:string,address:string,phoneNo:integer) Parts(pid:integer,pname:string,color:string) Catalog(sid:integer,pid:integer,cost:real) The Catalog relation lists the prices charged for parts by Suppliers. Write the following queries in relational algebra. (i) Find the names of suppliers who supply only some red part. (ii) Find the sids of suppliers who supply some yellow part or are at 221 Packer Street. (iii) Find the sids of suppliers who supply every part. (iv) Find the sids of suppliers who supply some red but not green part.	**2** **2** **2** **2**	2	L3
14	(i) Consider a relation schema R = (A, B, C, D, E) with the set of functional dependencies F = {A -> BC, CD -> E, B -> D, E->A}. If we decompose the relation R into R1 (A, B, C) and R2 (A, D, E) then show this decomposition is a lossless decomposition. (ii) Why 3NF is desirable than BCNF with a suitable example.	**4** **4**	3	L6 L4
15	Consider the schema given in Question No.13 and answer the following queries in SQL: (i) Find the pids of parts supplied by at least two different suppliers. (ii) Find the pids of the most expensive parts supplied by the supplier named Sham. (iii) Find the sids who supplies least number of parts. (iv) Find the average number of parts supplied by every supplier.	**2** **2** **2** **2**	2	L6
16	(i) Explain the states of a transaction it undergoes during execution. (ii) Consider the schedule of a set of transactions T_1, T_2, T_3, T_4. Is the given below schedule is conflict serializable? Explain your answer. T_1: read(Y), read(Z), read(U), read(U), write(U), write(Y), write(Z) T_2: read(X), read(Y), write(Y) T_3: read(Y), write(Y), read(Z), write(Z) T_4: read(V), read(W), read(W)	**4** **4**	4	L2 L6
17	Consider the following two transactions: T1: read(A);	**8**	4	L5

	T2: read(B);
	read(A);
	if B=10 then A:=A+40;
	write(A)
	Add lock and unlock instructions to the transactions T1 and T2, so that they observe the two-phase locking protocol. Can the execution of these transactions result in deadlock? Explain your answer.			

18	(i) Explain the deferred update recovery? Why does this method not require the UNDO - type log entries? (ii) Consider the following log corresponding to a particular schedule at the point of a system crash for four transactions T1, T2, T3, T4 and T5. Suppose that we use the immediate update protocol with checkpoints. Illustrate the recovery process from the system crash. Specify which transactions are ROLLEDBACK, which operations in the log are REDONE and which (if any) are UNDONE, and whether any cascading ROLLBACK takes place. <T1, Start> <T1, D, 20, 20> <T1, Commit > <Checkpoint {T1}> <T2, Start> <T2, B, 12, 30> <T2, Commit> <T4, Start> <T4, F, 25, 35> <T3, Start> <T3, C, 30, 50> <T4, A, 30, 20> <T3, Commit> <T5, Start> <T5, G,10,70> <T4, E, 15, 25> System Crash ←	4 4	4	L2 L5
19	Explain the cost estimation for select operations with suitable examples.	8	6	L2
20	(i) Illustrate the steps involved in query processing. (ii) Let relations r1(A, B, C) and r2(C, D, E) have the following properties: r1 has 10,000 tuples, r2 has 5,000 tuples, 25 tuples of r1 fit on one block, and 50 tuples of r2 fit on one block. The relation r1 is to be used as build input. Partition it into five partitions, each of size 20 blocks. This partitioning can be done in one pass. Similarly, partition of the relation r2 into five partitions, each of size 80 blocks. This is also done in one pass. Estimate the number of block accesses required using hash join strategy for r1 ⋈ r2.	4 4	6	L1 L5
21	Assume that the extendable hashing on a file is used for storing the following search-key values: 2,3,5,7,10,14,17,20,27,35. Show that the extendable hash structure for this file, if the hash function is h(x) = x mod 5 and buckets can hold two records.	8	5	L3
22	Describe the architecture of a typical data warehouse and the design issues in building the data ware house.	8	5	L2

PART – C (2 x 8 = 16 marks)

Q. No	Questions	Marks	CO	BL
23	Consider the schema given in Question No.13 and answer the following questions in SQL: (i) Create a TRIGGER that will check the phoneNo is valid or not during insertion or updation in the Suppliers table. If the phoneNo is less than 10 digits, then the display following message "Phone Number is not valid". (ii) Create a view to store the details of the suppliers only supplies green and red parts and the total number of parts that the supplier supplies.	4 4	2	L6
24	Normalize the relation DB(PatientNo, PatientName, appNo, time, doctor) which holds the following functional dependencies: PatientNo -> PatientName PatientNo,appNo -> time,doctor time -> appNo	8	3	L6