Hall Ticket No | | | | | | | | | | | Question Paper Code: ACSB08

INSTITUTE OF AERONAUTICAL ENGINEERING

(Autonomous)

Dundigal, Hyderabad - 500 043

MODEL QUESTION PAPER - I

B.Tech IV Semester End Examinations, May - 2020

Regulations: R18

DATABASE MANAGEMENT SYSTEMS

(Common to CSE & IT)

Time: 3 hours **Max. Marks: 70**

Answer ONE Question from each Unit
All Questions Carry Equal Marks
All parts of the question must be answered in one place only

MODULE – I

1.	a)	Explain the difference between external, internal, and conceptual schemas. How are these different schema layers related to the concepts of logical and physical data independence?.	**[7M]**
	b)	What are the responsibilities of a DBA? If we assume that the DBA is never interested in running his or her own queries, does the DBA still need to understand query optimization? Why?	**[7M]**
2.	a)	Construct an E-R diagram for a hospital with a set of patients and a set of medical doctors. Associate with each patient a log of the various tests and examinations conducted.	**[7M]**
	b)	Design an E-R diagram for keeping track of the exploits of your favorite sports team. You should store the matches played, the scores in each match, the players in each match and individual player statistics for each match. Summary statistics should be modeled as derived attributes.	**[7M]**

MODULE – II

3.	a)	Explain the statement that relational algebra operators can be composed. Why is the ability to compose operators important?	**[7M]**
	b)	Given two relations R1 and R2, where R1 contains N1 tuples, R2 contains N2 tuples, and N2 > N1 > 0, give the minimum and maximum possible sizes (in tuples) for the resulting relation produced by each of the following relational algebra expressions. In each case, state any assumptions about the schemas for R1 and R2 needed to make the expression meaningful: (1) R1∪R2, (2) R1∩R2, (3) R1−R2, (4) R1×R2, (5) σa=5(R1), (6) πa(R1), and (7) R1/R2.	**[7M]**
4.	a)	What is relational completeness? If a query language is relationally complete, can you write any desired query in that language?	**[7M]**

b) Consider the following database. Employee (employee-name, street, city) Works (employeename, company-name, salary) Company (company-name, city) Manager (employee-name, manager- name) Give an expression in the relational algebra, the tuple relational calculus, and the domain relational calculus, for the following query. Find the names of all employees who work for estate bank **[7M]**

MODULE – III

5. a) What is a view? How do views support logical data independence? How are views used for security? How are queries on views evaluated? Why does SQL restrict the class of views that can be updated? **[7M]**

b) Consider the following relations: Student(snum: integer, sname: string, major: string, level: string, age: integer) Class(name: string, meets at: string, room: string, fid: integer) Enrolled(snum: integer, cname: string) Faculty(fid: integer, fname: string, deptid: integer) Write the following queries in SQL. No duplicates should be printed in any of the answers. 1. Find the names of all Juniors (level = JR) who are enrolled in a class taught by I. Teach. 2. Find the age of the oldest student who is either a History major or enrolled in a course taught by I. Teach. 3. Find the names of all classes that either meet in room R128 or have five or more students enrolled. 4. Find the names of all students who are enrolled in two classes that meet at the same time **[7M]**

6. a) What is Normalization? Discuss what are the types of normalizations? Discuss the 1NF, 2NF, 3NF with example. **[7M]**

b) Consider the employee database , where the primary keys are Underlined. **[7M]**

employee(eid, empname, street,city)
works(empname ,companyname ,salary)
company(companyname,city)
manages(empname,management)

Give an expression in the relational algebra for each request

1) Find the names of all employees who work for First Bank Corporation.

2) Find the names, street addresses and cities of residence of all employees who work for First Bank Corporation and earn more than 200000 per annum

3) Find the names of all employees in this database who live in the same city as the company for which they work.

4) Find the names of all employees who earn more than every employees of small Bank Corpor.

MODULE – IV

7. a) What is the phantom problem? Can it occur in a database where the set of database objects is fixed and only the values of objects can be changed? **[7M]**

b) Consider a database with objects X and Y and assume that there are two transactions T1 and T2. Transaction T 1 reads objects X and Y and then writes object X. Transaction T 2 reads objects X and Y and then writes objects X and Y. **[7M]**

1. Give an example schedule with actions of transactions T 1 and T 2 on objects X and Y that results in a write-read conflict.

2. Give an example schedule with actions of transactions T 1 and T 2 on objects X and Y that results in a read-write conflict.

3. Give an example schedule with actions of transactions T 1 and T 2 on objects X and Y that results in a write-write conflict.

4. For each of the three schedules, show that Strict 2PL disallows the schedule.

8. a) Describe how a typical lock manager is implemented. Why must lock and unlock be atomic operations? What is the difference between a lock and a latch? What are convoys and how should a lock manager handle them?. **[7M]**

b) Compare the shadow-paging recovery scheme with the log-based recovery schemes in terms of ease of implementation and overhead cost? **[7M]**

MODULE – V

9. a) When is it preferable to use a dense index rather than a sparse index? Explain your answer with suitable example. **[7M]**

b) Consider the following relation: **[7M]**

Emp(eid: integer, sal: integer, age: real, did: integer)

There is a clustered index on eid and an unclustered index on age.

1. How would you use the indexes to enforce the constraint that eid is a key?

2. Give an example of an update that is definitely speeded up because of the available indexes. (English description is sufficient.)

3. Give an example of an update that is definitely slowed down because of the indexes. (English description is sufficient.)

4. Can you give an example of an update that is neither speeded up nor slowed down by the indexes?

10. a) What is the difference between a clustered index and an unclustered index? If an index contains data records as ‘data entries,’ can it be unclustered? **[7M]**

b) Suppose that we are using extendable hashing on a file that contains records with the following search- key values: 2,3,5,7,11,17,19,23,29,31 Show the extendable hash structure for this file if the hash function is h(x) = x mod 8 and buckets can hold three records **[7M]**

INSTITUTE OF AERONAUTICAL ENGINEERING

(Autonomous)

Dundigal, Hyderabad - 500 043

COURSE OBJECTIVES:

The course should enable the students to:

I	Understand the role of database management system in an organization and learn the database concepts.
II	Understand the role of database management system in an organization and learn the database concepts.
III	Construct database queries using relational algebra and calculus.
IV	Understand the concept of a database transaction and related database facilities.
V	Learn how to evaluate set of queries in query processing.

COURSE OUTCOMES (COs):

CO 1	Describe Purpose of Database Systems, View of Data, Data Models, Database Languages, Database Users, Various Components of overall DBS architecture, Various Concepts of ER Model, Basics of Relational Model.
CO 2	Determine Relational algebra, selection and projection, set operations, renaming, joins, division, examples of algebra queries, relational calculus: Tuple relational calculus, Domain relational calculus, expressive power of algebra and calculus.
CO 3	Understand SQL – Data Definition commands, Queries with various options, Mata manipulation commands, Views, Joins, views, integrity and security; Relational database design: Pitfalls of RDBD, Lossless join decomposition, Functional dependencies , Armstrong Axioms, Normalization for relational databases 1st , 2nd and 3rd normal forms, Basic definitions of MVDs and JDs, 4th and 5th normal forms.
CO 4	Explore the concept of Transaction, Transaction State, Implementation of Atomicity and Durability, Concurrent Executions, Serializability, Recoverability. Concurrency Control: Lock-Based Protocols, Timestamp-Based Protocols, Validation-Based Protocols, Multiple Granularities, Multiversion Schemes, Deadlock Handling. Recovery: Failure Classification, Storage Structure, Recovery and Atomicity, Log-Based Recovery, Shadow Paging, Recovery with Concurrent Transactions Buffer Management
CO 5	Knowledge the Physical Storage Media, Magnetic Disks, Storage Access, File Organization, Organization of Records in Files. Indexing and Hashing: Basic Concepts: Ordered Indices, B+-Tree Index Files, B-Tree Index Files, Static Hashing, Dynamic Hashing, Comparison of Ordered Indexing and Hashing. Query Processing: Overview, Measures of Query Cost.

COURSE LEARNING OUTCOMES (CLOs):

ACSB08.01	Describe the Purpose of Database Systems, Data Models, and View of Data.
ACSB08.02	Summarize the concept of Database Languages, Database Users.
ACSB08.03	Identify the Various Components of overall DBS architecture.
ACSB08.04	Use the concept of ER Model.
ACSB08.05	Describe Basics of Relational Model.
ACSB08.06	Determine Relational algebra, The Self variable.
ACSB08.07	Understand selection and projection, set operations.
ACSB08.08	Determine renaming, joins, division.
ACSB08.09	Use examples of algebra queries.
ACSB08.10	Illustrate Tuple relational calculus, Domain relational calculus, and also expressive power of algebra and calculus.
ACSB08.11	Understand SQL – Data Definition commands, Queries with various options.
ACSB08.12	Analyze the concept of Mata manipulation commands, Views, Joins, views.
ACSB08.13	Illustrate Calling a function, Returning multiple values from a function.
ACSB08.14	Contrast the Usage of Relational database design, Functional dependencies, Armstrong Axioms
ACSB08.15	Define Normalization, 2nd and 3rd Normalization, Basic definitions of MVDs and JDs, 4th and 5th normal forms.
ACSB08.16	Discuss the concept of Transaction, Transaction State.
ACSB08.17	Understand Atomicity and Durability, Concurrent Executions.
ACSB08.18	Summarize the concept of Serializability, Recoverability.
ACSB08.19	Discuss the Concurrency Control and various Protocols.
ACSB08.20	Understand the concept of Multiversion Schemes, Deadlock Handling. Recovery: Failure Classification, Storage Structure, Recovery and Atomicity, Log-Based Recovery, Shadow Paging, Recovery with Concurrent Transactions Buffer Management.
ACSB08.21	Knowledge about the Physical Storage Media, Magnetic Disks, Storage Access
ACSB08.22	Apply Working with File Organization, Organization of Records in Files.
ACSB08.23	Understand Ordered Indices, B+-Tree Index Files, B-Tree Index Files, Static Hashing, Dynamic Hashing.
ACSB08.24	Comparison of Ordered Indexing and Hashing.
ACSB08.25	Illustrate Query Processing: Overview, Measures of Query Cost.

MAPPING OF SEMESTER END EXAMINATION - COURSE OUTCOMES

SEE Question No		Course Learning Outcomes		Course Outcomes	Blooms Taxonomy Level
1	a	ACSB08.03	Describe the Purpose of Database Systems, Data Models, and View of Data	CO 1	Understand
	b	ACSB08.01	Summarize the concept of Database Languages, Database Users.	CO 1	Understand
2	a	ACSB08.02	Use the concept of ER Model.	CO 1	Remember
	b	ACSB08.02	Use the concept of ER Model.	CO 1	Understand
3	a	ACSB08.06	Describe Basics of Relational Model.	CO 2	Understand
	b	ACSB08.06	Determine Creating a class, The Self variable.	CO 2	Remember
4	a	ACSB08.08	Determine Creating a class, The Self variable..	CO 2	Understand
	b	ACSB08.07	Understand selection and projection, set operations.	CO 2	Understand
5	a	ACSB08.11	Analyze the concept of Mata manipulation commands, Views, Joins, views.	CO 3	Understand
	b	ACSB08.12	Understand SQL – Data Definition commands, Queries with various options.	CO 3	Understand
6	a	ACSB08.13	Define Normalization, 2nd and 3rd Normalization, Basic definitions of MVDs and JDs, 4th and 5th normal forms.	CO 3	Understand
	b	ACSB08.14	Contrast the Usage of Relational database design, Functional dependencies, Armstrong Axioms,	CO 3	Understand
7	a	ACSB08.17	Understand Atomicity and Durability, Concurrent Executions.	CO 4	Understand
	b	ACSB08.18	Discuss the concept of Transaction, Transaction State.	CO 4	Understand
8	a	ACSB08.18	Discuss the Concurrency Control and various Protocols.	CO 4	Understand
	b	ACSB08.19	Understand the concept of Multiversion Schemes, Deadlock Handling. Recovery: Paging.	CO 4	Understand
9	a	ACSB08.22	Understand Ordered Indices, B+-Tree Index Files, B-Tree Index Files, Static Hashing, Dynamic Hashing.	CO 5	Understand
	b	ACSB08.24	Understand Ordered Indices, B+-Tree Index Files, B-Tree Index Files, Static Hashing, Dynamic Hashing.	CO 5	Understand
10	a	ACSB08.25	Comparison of Ordered Indexing and Hashing.	CO 5	Understand
	b	ACSB08.24	Comparison of Ordered Indexing and Hashing.	CO 5	Understand

Signature of Course Coordinator **HOD, IT**