B2.2-R3: INTRODUCTION TO DATABASE MANAGEMENT SYSTEMS

NOTE:

1. There are **TWO PARTS** in this Module/Paper. **PART ONE** contains **FOUR** questions and **PART TWO** contains **FIVE** questions.

2. **PART ONE** is to be answered in the **TEAR-OFF ANSWER SHEET** only, attached to the question paper, as per the instructions contained therein. **PART ONE** is **NOT** to be answered in the answer book.

3. Maximum time allotted for **PART ONE** is **ONE HOUR**. Answer book for **PART TWO** will be supplied at the table when the answer sheet for **PART ONE** is returned. However, candidates, who complete **PART ONE** earlier than one hour, can collect the answer book for **PART TWO** immediately after handing over the answer sheet for **PART ONE**.

TOTAL TIME: 3 HOURS **TOTAL MARKS: 100**

(PART ONE – 40; PART TWO – 60)

PART ONE
(Answer all the questions)

1. Each question below gives a multiple choice of answers. Choose the most appropriate one and enter in the "tear-off" answer sheet attached to the question paper, following instructions therein. (1 x 10)

1.1 Pick the odd one out
A) Primary key
B) Super key
C) Candidate key
D) Foreign key

1.2 Relational Algebra is
A) Data Definition Language
B) Meta Language
C) Procedural query language
D) Non procedural language

1.3 One of the following is a valid record -based data models
A) Object-oriented model
B) Relational model
C) Entity-relationship model
D) None of the above

1.4 One of the following steps is not involved in processing a query
A) Parsing and translation
B) Optimization
C) Evaluation
D) Distribution

1.5 Which one of the following describes the timestamp-based protocols correctly?
A) This protocol requires that each transaction issue lock and unlock requests in two phases.
B) This protocol employs only exclusive locks.
C) This protocol selects an ordering among transaction in advance.
D) None of the above

1.6 Which one of the following is not a valid relational database?
A) SYBASE
B) ORACLE
C) IMS
D) UNIFY

1.7 4NF is designed to cope with
A) transitive dependency
B) join dependency
C) multi valued dependency
D) none of these

1.8 Which one of the following is a valid join type?
A) natural
B) full outer join
C) on
D) using

1.9 Which one of the following is not a valid aggregation function in SQL?
A) avg
B) min
C) where
D) sum

1.10 Which of the following is not a valid unary operation in the relational algebra?
A) select
B) min
C) project
D) rename

2. Each statement below is either TRUE or FALSE. Choose the most appropriate one and ENTER in the "tear-off" sheet attached to the question paper, following instructions therein. (1 x 10)

2.1 A relationship is an association among several entities.
2.2 Physical data models are used to describe data at the highest level.
2.3 QBE is based on the tuple relational calculus.
2.4 The database schema and the database instance are the same thing.
2.5 Functional dependencies are constraints on the set of legal relations.
2.6 Integrity constraint guard against accidental damage to the database.
2.7 One-way to ensure serializability is to require that access to data items be done in a mutually I exclusive manner.
2.8 The cost of processing a query is not dependent on disk access.
2.9 The recovery scheme does not depend on the concurrency control scheme.
2.10 Deadlocks can be described precisely in terms of a directed graph.

3. Match words and phrases in column X with the closest related meaning/ word(s)/phrase(s) in column Y. Enter your selection in the "tear-off" answer sheet attached to the question paper, following instructions therein. (1 x 10)

X		Y	
3.1	Dense index	**A.**	data are represented by collection of records and relationship among data are represented by links
3.2	Transaction	**B.**	Query language based on both the relational algebra and the tuple relational calculus
3.3	Shadow Paging	**C.**	The index structure is the most widely used to several index structures that maintain their efficiency despite insertion and deletion of data
3.4	Referential integrity constraint	**D.**	A record appears for every search key value in the file
3.5	Committed	**E.**	A recovery technique
3.6	B+ tree index	**F.**	This ensures that a value that appears in one relation for a given set of attributes also appears for a certain set of attributes in another relation
3.7	Network Model	**G.**	The successful completion of a transaction
3.8	Entity	**H.**	A unit of program execution that accesses and possibly updates various data items
3.9	DMI	**I.**	A powerful declarative query language
3.10	Embedded SQL	**J.**	An object in the real world that is distinguishable from all other objects
		K.	The number of entities to which another entity can be associated via a relationship set
		L.	BCNF

4. Each statement below has a blank space to fit one of the word(s) or phrase(s) in the list below. Enter your choice in the "tear-off" answer sheet attached to the question paper, following instructions therein. (1 x 10)

A.	merge-join	**B.**	natural join	**C.**	starvation
D.	rollback	**E.**	from	**F.**	replicate
G.	cartesian product	**H.**	relational algebra	**I.**	fragmented
J.	ordered	**K.**	transaction	**L.**	division
M.	hash	**N.**	trigger	**O.**	super key

4.1 The _____ operation allows to combine information from any two relations.

4.2 A(n) _____ is a statement that is executed automatically by the system as a side effect of a modification to the database.

4.3 The _____ algorithm can be used to compute natural joins and equi-joins.

4.4 If a relation is _____ a copy of that is stored in two or more sites.

4.5 A(n) _____is a set of one or more attributes that taken collectively allows us to identify uniquely an entity in the entity set.

4.6 A(n) _____ is a collection of operations that performs a single logical function in a database application.

4.7 The _____ clause by itself defines a Cartesian product of the relations in the clause.

4.8 _____ indices are based on the values being distributed informally across a range of buckets.

4.9 The_____ is a situation where a transaction never completes its designated task.

4.10The_____ operation is suited to queries that include the phase "for all".

PART TWO

(Answer any **FOUR** questions**)**

5.

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.

b) Explain the advantages and disadvantages of Database Processing?

(10+5)

6.

a) List and explain with suitable example five primary relational algebra operators.

b) What is meant by Heuristic Optimization? Discuss the main heuristics that are applied to query optimization?

(10+5)

7.

a) Consider the insurance database given below:

person **(*driver-id*, *name*, *address*)**
car **(*license*, *model*, *year*)**
accident **(*report-number*, *date*, *location*)**
owns **(*driver-id*, *license*)**
participated **(*driver-id*, *car*, *report-number*, *damage-amount*)**

Construct the following SQL queries for this relational database.

i) Find the total number of people who owned cars that were involved in accidents in 2004.

ii) Find the number of accidents in which the cars belonging to "Thakre" were involved.

iii) Delete the Mazda belonging to "S Khan".

b) How does SQL allow implementation of entity and integrity constraints?

(9+6)

8.

a) List and explain Armstrong's Axioms.

b) Explain the purpose and utility of different normal forms. Specifically define and differentiate between third normal form and BCNF.

c) What is referential integrity? Explain with suitable examples.

(5+5+5)

9.

a) Explain ACID property of transactions.

b) What do you understand by lock granularity? Explain

c) Explain in brief working of two-phase locking protocol.

(5+5+5)