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 A functional dependency of the form $X \rightarrow Y$ is trivial if

A) $Y \subseteq X$
B) $Y \subset X$
C) $X \subseteq Y$
D) $X \subset Y$ and $Y \subset X$

1.2 If every non-key attribute is functionality dependent on the primary key, then the relation will be in

A) first normal form
B) second normal form
C) third normal form
D) fourth normal form

The Q1.3 and Q1.4 are based on the following details. Consider the given schemes.

Branch_scheme = (Branch_name, assets, Branch_city)
Customer_scheme = (Customer_name, street, Customer_city)
Deposit_scheme = (Branch_name, account_number, Customer_name, balance)
Borrow_scheme = (Branch_name, loan_number, Customer_name, amount)
Client_scheme = (Customer_name, banker_name)

1.3 Which of the following queries finds the clients of banker Agassi and the city they live in?

A) $\pi_{Client.Customer\_name.Customer\_City}$ $(\sigma_{Client.Customer\_name=Customer.Customer\_name}$ $(\sigma_{Banker\_name="Agassi"}$ $(Client \times Customer)))$
B) $\pi_{Customer\_name.Customer\_City}$ $(\sigma_{Banker\_name="Agassi"}$ $(Client \times Customer))$
C) $\pi_{Client.Customer\_name.Customer\_City}$ $(\sigma_{Banker\_name="Agassi"}$ $(\sigma_{Client.Customer\_name=Customer.Customer\_name}$ $(Client \times Customer)))$
D) $\pi_{Customer\_name.Customer\_City}$ $(\sigma_{Banker\_name="Agassi"}$ $(Client \times Customer))$

1.4 Which of the following tuple relational calculus finds all the customers who have a loan amount of more than 1200?

A) { t (Customer_name) | t ε borrow Λ t[amount] >1200 }
B) { t | t (Customer_name) ε borrow Λ t[amount] >1200 }
C) { t | ∃ s ε borrow (t [Customer_name] = s [Customer_name] Λ s[amount] >1200) }
D) { t | ∃ s ε borrow (t [Customer_name] = s [Customer_name] Λ s[amount] >1200) }

1.5 Choose the most appropriate choice with respect to conceptual design.

A) Conceptual design is a documentation technique. Once the relation schemes are defined one can draw E-R diagrams from the relation schemes for documentation.
B) Conceptual design needs data volume and processing frequencies to determine the size of the data base.
C) Output of any conceptual design is an E-R diagram.
D) Conceptual design involves modeling the data requirements independent of the DBMS, operating system and the hardware.

1.6 Third normal form is inadequate in situations where the relation

A) has multiple candidate keys
B) has candidate keys that are composite
C) has overlapped candidate keys
D) none of the above

1.7 Manager's salary details are hidden from the employee. This is

A) Conceptual level data hiding
B) Physical level data hiding
C) External level data hiding
D) None of the above

1.8 If P and Q are predicates and P is the relational algebra expression, then which of the following equivalence are valid?

A) $\sigma_P\ (\sigma_Q(e)) = \sigma_Q(\sigma_P(e))$
B) $\sigma_P\ (\sigma_Q(e)) = \sigma_{P \wedge Q}(e)$
C) $\sigma_Q\ (\sigma_P(e)) = \sigma_{P \wedge Q}(e)$
D) None of the above

1.9 Consider the following set of functional dependencies on the scheme (A, B, C).

$A \rightarrow BC$
$B \rightarrow C$
$A \rightarrow B$
$AB \rightarrow C$

The canonical cover for this set is

A) $A \rightarrow BC$ and $B \rightarrow C$
B) $A \rightarrow BC$ and $AB \rightarrow C$
C) $A \rightarrow BC$ and $A \rightarrow B$
D) $A \rightarrow B$ and $\rightarrow C$

1.10 Assume transaction A holds a shared lock R. If transaction B also requests for a shared lock on R.

A) it will result in a deadlock situation
B) it will immediately be granted
C) it will immediately be rejected
D) it will be granted as soon as it is release by A

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 The recovery manager is responsible for implementing a particular strategy for concurring control.
2.2 We say that a non serial schedule is correct if it produces the same as result as some serial execution.
2.3 With concurrency control there will be many problems like multiple update problem, uncommitted dependency problem.
2.4 To ensuring serializability are locking, timestamp ordering and validation based protocols.
2.5 In tuple calculus, the variable represent the values drawn from specified relations.
2.6 Relational algebra is not a procedural language.
2.7 The result of relational algebraic expression is not a relation.
2.8 Referential integrity requires fully normalized tables.
2.9 Normalization need not always improve database performance.
2.10 Multi valued dependencies are the result of the 1NF which prohibited an attribute from having a set values.

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	Transaction that is cancelled before changes are applied.	**A.**	Abort transaction
3.2	Those controls that provide for restoring the database in case of system.	**B.**	Commit transaction
3.3	The maximum number of instances of one object set related to a single instance of other object set.	**C.**	Roll back
3.4	Identification of data elements, relationship and constraints for database.	**D.**	Backup and recovery control
3.5	Two transactions are mutually excluded from accessing the next record required to complete their transactions.	**E.**	Cardinality
3.6	Encoding data to make them unintelligible to unauthorized persons.	**F.**	Cartesian product
3.7	A relationship having a maximum cardinality of one in at least one direction.	**G.**	Conceptual database design
3.8	The process of converting a relation to a standard form.	**H.**	Conceptual level
3.9	Record of all activity that has affected a database during a period of time.	**I.**	Deadlock
3.10	Major limitation of file Management System.	**J.**	Encryption
		K.	Functional relationship
		L.	Functional dependency
		M.	Normalization
		N.	Transaction log
		O.	Data redundancy

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.	Degree of database	**B.**	Denormalization	**C.**	Primary key
D.	Candidate key	**E.**	Cardinality of database	**F.**	Recursive query
G.	Transitive dependency	**H.**	Domain key normal form (DKNF)	**I.**	Query resolution
J.	Secondary key	**K.**	Project	**L.**	Partial participation
M.	Projection	**N.**	Fourth normal form	**O.**	Join dependency
P.	Transaction Management	**Q.**	Transaction log	**R.**	Interleaved main memory

4.1 A function dependency between two or more non key attribute in a relation is ________.

4.2 ________ is that which ensure the transactions are either completed or cancelled so that the database is never left in an inconsistent state.

4.3 A query in which the output of the query is then used as input for the same query is ________.

4.4 ________ is the process of collecting the data needed to answer a query.

4.5 A query in which only some of the attributes in the source relation appears in the output is defined as ________.

4.6 A relation is in ________ if it is BCNF and it has at the most only one independent multi valued dependency.

4.7 For ________, the constraint is removed so that every occurrence of all the super class need not appear in the category.

4.8 Transforming data in blocks from the main memory to the cache memory enables an ________ unit to operate at its maximum speed.

4.9 ________ is the process of increasing redundancy in the database either for convenience or to improve performance.

4.10 In relational model, the columns of the table are known as ________.

PART TWO

(Answer any **FOUR** questions)

5.

a) Construct an ER diagram (including important attributes) for a car insurance database that includes data about Customer (Car owner), Cars, accident, drivers involved in accidents and injured drivers and/or passengers. Note that any customer can insure many cars, each car may have different drivers at different times and accidents typically involved one or more cars.

b) Consider relation R with 5 attributes A, B, C, D, E. Your are given following dependencies:

$A \rightarrow B$

$BC \rightarrow E$

$ED \rightarrow A$

i) List all key of R.

ii) Is R in 3NF or BCNF.

iii) Design the normalized table.

(9+6)

6.

a) What is the difference between a procedural and non-procedural language? How would you classify the relational algebra and relational calculus?

b) Define the five basic relational algebra operations. Define the Join, Intersection and Division operations in terms of these five basic operations.

c) In respect of project, we are given the following record types for the relational model of the database:

EMPLOYEE : with attributes ("Emp#" and "Name")
ASSIGNED_To: with attributes ('Project#', 'Emp#')
PROJECT : with attributes ('Project#', 'Project_Name', 'Chief_architect')

Express the following queries in relational algebra:

i) Get Emp# of employees working on project Comp 353

ii) Give details of employees (both number and name) working on project Comp 353

iii) Gather details of employees working on both Comp 353 and Comp 354

(5+5+5)

7.

a) What restrictions apply to the use of aggregate functions within the select?

b) Describe how the process of view resolution works and what restrictions are necessary to ensure that a view is updatable?

c) Consider the following Relation Schema. An employee can work in more than one department:

Emp (E-id, E_name, Salary)
Dept (d_id, d_name, manager_id, floor_number)

Write the following queries in SQL:

i) Print the name of all employees, who work on the 10^{th} floor and earn salary less than Rs. 50,000.

ii) Print the names of the departments that employee Santa work in.

iii) Print the names of all managers who manage three or more departments on the same floor.

iv) Print the names of all employees who work on floors where Jane Donald works.

v) Give every employee who works in the toys dept. at 10% raise in the salary.

(5+5+5)

8.

a) Explain the purpose and scope of Database security and explain the following in terms of providing security for a database: authorization, views, backup and recovery, Integrity, encryption and RAID technology.

b) Explain what is meant by transaction. Why are transactions important limit of operations in a DBMS?

c) Consistency and reliability aspects of transaction are due to the ACIDity properties of transaction. Discuss each of these properties and how they relate to the concurrency control and recovery mechanisms? Give examples to illustrate your answer.

(5+5+5)

9.

a) What is a time stamp? How do time stamp based protocols for concurrency control differ from locking based protocols?

b) Discuss how the log file is a fundamental feature in any recovery mechanism. Explain what is meant by forward and backward recovery and describe how the log file is used in forward and backward recovery. What is the significance of the write-ahead log protocol? How do check points affect the recovery protocol?

(7+8)