DATABASE MANAGEMENT SYSTEMS
(INFO 2204)

Time Allotted: 2½ hrs **Full Marks : 60**

Figures out of the right margin indicate full marks.

Candidates are required to answer Group A and any 4 (four) from Group B to E, taking one from each group.

Candidates are required to give answer in their own words as far as practicable.

Group – A

1. Answer any twelve: **12 × 1 = 12**

Choose the correct alternative for the following

(i) Which of the following relational algebra operations do not require the participating tables to be union-compatible?
(a) Union (b) Intersection (c) Difference (d) Join.

(ii) A decomposition of a relation R into R1 and R2 is lossless if
(a) Common attributes between R1 and R2 is candidate key of both the decomposed relations
(b) Common attributes between R1 and R2 is candidate key of atleast one of the decomposed relations
(c) Common attributes between R1 and R2 is a non-key attribute of the decomposed relations
(d) Common attributes between R1 and R2 is key attribute of both the decomposed relations.

(iii) A multivalued dependency $A \rightarrow\rightarrow B$ is trivial if
(a) $A \cap B = \Phi$ (b) BUA=R (c) Either (a) or (b) (d) None of them

(iv) If a schedule S can be transformed into a schedule S' by a series of swaps of non-conflicting database read/write instructions, then S and S' are always:
(a) Conflict Equivalent to each other (b) View equivalent to each other
(c) Both (a) and (b) (d) None of these.

(v) **Statement 1:** It is always possible to normalize a relation to BCNF decomposition, preserving all functional dependencies and lossless join property.
Statement 2: It is always possible to normalize a relation to 3NF decomposition preserving all functional dependencies and lossless join property.
Considering the above two statements choose the right option fromthe followings:
(a) Only Statement 1 is correct (b) Only Statement 2 is correct
(c) Both Statement 1 and Statement 2 are correct (d) Both Statement 1 and Statement 2 are incorrect.

(vi) All aggregate functions ignore null values in their input collection. The exception is for which function?
(a) Count(attribute) (b) Count(*) (c) Avg (d) Sum.

(vii) A relation R(A,B,C,D,E,F) holds the following FDs
F={A-->BC, D-->F, B-->E}. Then AC^+ is
(a) {A,C,B,E} (b) {A,C,D} (c) {A,C, F,E} (d) None

(viii) Consider the following DDL statement:
Create table student(rollNo number(5) primary key, mobile number(10), dname varchar(5) references Department(dname));
Which of the following integrity constraints have been imposed on the table Employee using the above DDL statement:
(a) Entity Integrity constraint and Domain Integrity Constraint
(b) Entity Integrity constraint and Referential Integrity Constraint
(c) Entity Integrity constraint
(d) Entity Integrity constraint and Self-referential Integrity Constraint.

(ix) Consider a B+tree in which the maximum number of keys in a node is 5. What is the minimum number of keys in any non-rootnode?
(a) 1 (b) 2 (c) 3 (d) 4

(x) Emp(empcode, name, street, city, state, pincode).
For any pincode, there is only one city and state. Also, for given street, city and state, there is just one pincode. In normalization terms, Emp is a relation in
(a) 1 NF only (b) 2 NF and hence also in 1 NF
(c) 3NF and hence also in 2NF and 1NF (d) BCNF and hence also in 3NF, 2NF and 1NF

Fill in the blanks with the correct word

(xi) The _____ attribute and primary key of strong entity set acts as the composite primary key of the weak entity set.

(xii) '_____ %' matches any string of _____ characters.

(xiii) _____ graph is used for checking conflict serializability.

(xiv) During an update statement the required block is fetched from database (data file) to the _____.

(xv) Two phase locking protocol ensures _____.

Group - B

2. (a) Explain with the help of a diagram the Oracle Server architecture. *[(CO1)(Remember/LOCQ)]*

(b) Consider the following relations with underlined primary keys.
Product(P_code, Description, Stocking_date, QtyOnHand, MinQty, Price, Discount, V_code)
Vendor(V_code, Name, Address, Phone)
Here a vendor can supply more than one product but a product is supplied by only one vendor. Write SQL queries for the following:
(i) List the product description and the name of the vendor for which the QtyOnHand is less than MinQty.
(ii) List the names of all the vendors who supply more than one prod
(iii) List the details of the products whose prices exceed the average product price.
(iv) List the Name, Address and Phone of the vendors who are currently not supplying any product.

[(CO3)(Apply/IOCQ)]

4 + 8 = 12

3. (a) R(Sid,Pid,City) and S(Status,Pid,City) be the relations:

Sid	Pid	City
S1	P1	C1
S2	P2	C2
S3	P3	C3
S4	P4	C4

R

Status	Pid	City
T1	P1	C1
T2	P2	C2
T3	P3	C6

S

Compute the followings (if possible) for the relations above:
(i) R U S (ii) R X S
(iii) R(natural join)S (iv) R(right outer join)S. *[(CO2)(Apply/IOCQ)]*

(b) The LMF Zoo has many types of animals. Every type has a unique name. Every animal of the same type has a unique animal ID. Animals in two types may have the same animal ID. Animals also have age and gender. Animals may have diseases. The beginning time and the duration of a disease need to be recorded. A disease has a unique name. A type keeper takes care of only one type of animals. Every type may have many type keepers. A type keeper may or may not be familiar with diseases. But every disease must be handled by at least one type keeper. Type keepers have name, employee ID, ssn, address and phone number. All animals are in cages. Some cage may be empty. Every cage has a cage ID, space and height. A cage keeper may take care of many cages. Every non-empty cage must have at least one cage keeper. Empty cages don't need any cage keepers. Cage keepers have name, employee ID, ssn, address and phone number.

Model the entities and relationships (including attributes and properties of relationships) described below in an ER-diagram. Write down any assumptions you make. *[(CO2)(Analyse/IOCQ)]*

(1 + 1 + 1 + 2) + 7 = 12

Group - C

4. (a) Find the canonical cover of the following FD's given.
(i) AB→D, AC→D, A-->BC
(ii) Q→R, PQ→R, P→R, P-->Q. *[(CO2)(Analyse/IOCQ)]*

(b) Write a PL/SQL code to create an audit system, that will keep track of the records that are being deleted or updated on the user_accounts table. user_accounts table stores data about all the users created in the database. It maintains the account_id, password and account_creation_date. Whenever a record is deleted or updated from the user_accounts table, the original (old) value is stored in the accounts_audit table. When such an event is issued on the user_accounts table, a trigger is fired due to which the records are inserted in the accounts_audit table. The accounts_audit table maintains account_id information, operation being performed, user who is doing the operation and also the date on which such operation is taking place. *[(CO3)(Apply/IOCQ)]*

(2 + 2) + 8 = 12

5. (a) In a banking system there is a relation Employee Customer(employeeId, ename, customerId) which refers to the employee who are also customers of the bank. The constraints that hold on this relation are as follows:
(i) The names of each employee are different from each other i.e. no two employee have same name. That is, both employeeId as well as ename are unique for each entity.
(ii) customerId is also unique for each employee.
List the non-trivial Functional dependencies of this relation.
State in which highest normal form this relation exists? Argue in favour or against, that whether this relation can be further normalized to any other normal form or not. *[(CO2,CO6)(Evaluate/HOCQ)]*

(b) Consider relation R {ABCDEF} on which the following functional dependencies F are applicable
F = {A → B; BC → D; C → A; B → D; BE → F; CF → D}.
Find candidate keys of R. Show each steps. *[(CO2)(Analyse/IOCQ)]*

8 + 4 = 12

Group - D

6. (a) Consider a database maintaining the courses attended by the students during the summer recess. The student details consists of roll and coursed in which the student got enrolled. A student can get enrolled in more than one course. The course details consists of the course id, name of the course and course fees. Course id is the primary key of course details and acts as the foreign key in the student details. Course fee gets updated at intervals.
Write a PL/SQL code to create trigger which will be fired whenever an update in feestakes place. The updation details will be stored in the table fees-updation-details, having attributes course id, old fees, modified fees, date-of-change.
[(CO3)(Create/HOCQ)]

(b) Critically Comment:
(i) Concurrent execution of transaction can lead to inconsistencies.
(ii) Two phase locking can lead to a deadlock situation. *[(CO4)(Understand/LOCQ)],[(CO2)(Apply/IOCQ)]*

6 + (3 + 3) = 12

7. (a) Explain with example, what is a cascadeless schedule.
"Every cascadeless schedule is also Recoverable". Justify for or against this. *[(CO4)(Understand/LOCQ)]*

(b) Consider the concurrent schedule S of three transactions T1,T2 and T3, whereRi(A), Ri(B) are read operations and Wi(A), Wi(B) are write operations of Transaction Tion data items A and B accordingly:
S = R1(A) R3(C) W1(A) R2(A) W3(C) W2(A) R1(B) R2(C) W1(B) W2(C)
Find out whether the above concurrent schedule S is serializable or not – Justify your answer (Do not use Precedance graph). *[(CO4)(Apply/IOCQ)]*

(c) What are the disadvantages of concurrent execution of transactions? *[(CO4)(Understand/LOCQ)]*

4 + 5 + 3 = 12

Group - E

8. (a) Suppose that an unordered file with r=800000 records stored on a disk with block size 4096 bytes.File records are of fixed size with record length 1000 bytes.
(i) Find out the blocking factor and the number of blocks needed for the file. Suppose a **secondary index** is created on the above file, with non-ordering key field 12 bytes and block pointer 18 bytes.
(ii) Find the blocking factor for the index.
(iii) Also find the total number of blocks in the index.
(iv) Find how many blockaccess is required to search a particular record using the index. *[(CO5,CO6)(Apply/IOCQ)]*

(b) Convert the above secondary index to a multilevel index structure (such that the secondary index created in the above problem acts as the first level index).
(i) Calculate the number of levels required.
(ii) Calculate the number of block access required to search a particular record using this multilevel index structure.
[(CO5,CO6)(Apply/IOCQ)]

8 + 4 = 12

9. (a) Critically comment:
(i) "Primary index is a sparse index"
(ii) "Secondary index makes searching efficient in terms of block access"
(iii) "Wound and wait is a primitive deadlock prevention measure". *[(CO4,CO5)(Understand/LOCQ)]*

(b) Construct a B+ tree on the key field roll numbers of the student relation where andtheir corresponding grade points of four years are present. The order of internal node is 3 and that of leaf node is 2. Show each and every step. The roll numbers are as follows.
21, 110, 130, 12, 13, 117, 134, 290, 77,22, 18, 211,1 89, 23, 44, 55, 66
[(CO5)(Apply/IOCQ)]

(2 + 2 + 2) + 6 = 12

Cognition Level	LOCQ	IOCQ	HOCQ
Percentage distribution	23.95	57.3	18.75

Course Outcome (CO):

After the completion of the course students will be able to
1. Understand the need of DBMS over traditional file system and acquire the knowledge on overall database description, at three levels, namely, internal, conceptual, and external levels.
2. Deduce the constraints , i.e., the candidate keys, super-keys, that exists in a given real world problem and design the entity relationship diagram to graphically represent entities and their relationships to each other, typically used in computing in regard to the organization of data within databases or information systems.

3. Formulate a mathematical tool using relational algebra that operates on one or more relational tables and outputs a relational table as result, and design a normalized Database based on real-world situations, maintaining all constraints and manipulate database relations using SQL and PL/SQL.
4. Prove whether the ordering of concurrent transactions result in inconsistency of the database system or not.
5. Compare the number of block access required for searching a particular record, in an un indexed data file, with respect to a data file having (primary , secondary , clustering or multilevel) index structure.
6. Create a complete Normalized Database system, maintaining all the requirement specifications for a real life problem, and creating indexed relations for efficient accessing.

**LOCQ: Lower Order Cognitive Question; IOCQ: Intermediate Order Cognitive Question; HOCQ: Higher Order Cognitive Question.*