

**B.TECH/CSE/5<sup>TH</sup> SEM/CSEN 3102/2016**  
**DATABASE MANAGEMENT SYSTEMS**  
**(CSEN 3102)**

**Time Allotted : 3 hrs**

**Full Marks : 70**

*Figures out of the right margin indicate full marks.*

*Candidates are required to answer Group A and any 5 (five) from Group B to E, taking at least one from each group.*

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

**Group - A**  
**(Multiple Choice Type Questions)**

1. Choose the correct alternative for the following: **10 × 1 = 10**

- (i) The entity relationship modelling technique is  
 (a) top-down approach (b) bottom-up approach  
 (c) not related with database design (d) none of this.

(ii) Given the following relation instance

X	Y	Z
1	4	2
1	5	3
1	6	3
3	2	2

Which of the following functional dependencies are satisfied by the instance?

- (a)  $XY \rightarrow Z$  and  $Z \rightarrow Y$  (b)  $YZ \rightarrow X$  and  $Y \rightarrow Z$   
 (c)  $YZ \rightarrow X$  and  $X \rightarrow Z$  (d)  $XZ \rightarrow Y$  and  $Y \rightarrow X$ .
- (iii) Relational calculus is a  
 (a) procedural language (b) non- Procedural language  
 (c) data definition language (d) high level language.
- (iv) The view of total database content is  
 (a) conceptual view (b) internal view  
 (c) external view (d) physical View.
- (v) Which of the following statements are TRUE about an SQL query?

P: An SQL query can contain a HAVING clause even if it does not have a GROUP BY clause.

**B.TECH/CSE/5<sup>TH</sup> SEM/CSEN 3102/2016**

Q: An SQL query can contain a HAVING clause only if it has GROUP by clause.

R: All attributes used in the GROUP BY clause must appear in the SELECT clause.

S: Not all attributes used in the GROUP BY clause need to appear in the SELECT clause.

- (a) P and R (b) P and S (c) Q and R (d) Q and S.

(vi) An entity set that does not have sufficient attributes to form a primary key is a

- (a) strong entity set (b) weak entity set  
 (c) simple entity set (d) primary entity set.

(vii) Why does the following statement fail?

CREATE TABLE FRUITS&VEGETABLES (NAME VARCHAR2 (40));

- (a) The table should have more than one column defined  
 (b) NAME is a reserved word, which cannot be used as a column name  
 (c) The table name is invalid  
 (d) Column name cannot exceed 32 characters.

(viii) A transaction that completes its execution is said to be

- (a) committed (b) aborted  
 (c) rolled back (d) failed.

(ix) A table on the many side of a one to many or many to many relationship must:

- (a) be in Second Normal Form (2NF)  
 (b) be in Third Normal Form (3NF)  
 (c) have a single attribute key  
 (d) have a composite key

(x) Which of the following is not a property of transactions?

- (a) atomicity (b) concurrency  
 (c) isolation (d) durability.

**Group - B**

2. Consider a supplier data store which stores information about suppliers for a car manufacturer. The data store is supposed to contain the following types of information:

Suppliers:

- Every supplier has a unique supplier# (assigned by the car manufacturer for identification purposes);
- Every supplier supplies 1 or more parts;

- Every supplier has a unique name, city and postal code attribute;

Parts:

- Every part has a unique name, part# and price;
- Every part is supplied by one or more suppliers;
- Part# identifies uniquely a part.

Supplies:

- Every supply involves a supplier supplying a part;
  - Every supply has a unique quantity and a date;
  - Every supply is identified uniquely by its supplier, part and date
- Draw an entity-relationship diagram for this data store, showing clearly entities, relationships, attributes, cardinalities and keys.

**12**

3. (a) Which of the following plays an important role in representing information about the real world in a database? Explain briefly.
- The data definition language.
  - The data manipulation language.
  - The buffer manager.
  - The data model.
- (b) Explain the difference between logical and physical data independence.
- (c) 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?

**4 + 4 + 4 = 12****Group - C**

4. (a) Define the term functional dependency.
- (b) Consider the attribute set  $R = ABCDEF$  and the functional dependency set  $F = \{AD \rightarrow B, A \rightarrow E, C \rightarrow E, DEF \rightarrow A, F \rightarrow D\}$ . Find a candidate key of  $R$ . Given the attribute set  $R = ABCDEFGH$  and the functional dependency set  $F = \{BC \rightarrow GH, AD \rightarrow E, A \rightarrow H, E \rightarrow BCF, G \rightarrow H\}$ , decompose  $R$  into BCNF by decomposing in the order of the given functional dependencies.

**2 + (5 + 5) = 12**

5. Consider the following relational database:  
 Suppliers(supplier\_id, supplier\_name, address)  
 Parts(part\_id, part\_name, color)  
 Catalog(supplier\_id, part\_id, cost)  
 For each of the following queries,  
 (i) Find the names of suppliers who supply some red part  
 (ii) Find the supplier\_ids of suppliers who supply every part  
 (iii) Find the part\_ids of the most expensive parts supplied by suppliers named "Pipe Supplier".  
 Give an expression in  
 (a) the relational algebra  
 (b) the tuple relational calculus  
 (c) the domain relational calculus.

**(4 + 4 + 4) = 12****Group - D**

6. Based on the following schema answer the queries in SQL given below.  
 Student(st\_num, st\_name, major, grade, age)  
 Class(cname, meet\_at\_time, room, fac\_id)  
 Enrolled(st\_num, cname)  
 Faculty(fac\_id, fname, deptid)  
 (i) Find the names of all grade 10's who are enrolled in a class taught by "John Smith".  
 (ii) Find the names of all classes that either meet in room 5NE or have five or more students enrolled  
 (iii) Find the names of all students who are enrolled in two classes that meet at the same time.  
 (iv) Define a view part\_time\_students(st\_num, st\_name, major, grade, age) where a student is considered part-time if he or she takes less than 5 courses.

**(3 + 3 + 3 + 3) = 12**

7. (a) Define entity integrity and referential integrity. How does SQL allow specification of these?  
 (b) Discuss the strengths and weaknesses of the trigger mechanism. Contrast triggers with other integrity constraints supported by SQL.  
 (c) Explain the term stored procedure, and give examples why stored procedures are useful.

**4 + 4 + 4 = 12**

**Group - E**

8. In the schedule given below, the label  $R_i(X)$  indicates a read of element  $X$  by transaction  $T_i$ , and  $W_i(X)$  indicates a write of element  $X$  by transaction  $T_i$ .

(i) Draw the precedence graph for following schedule.

*$R_2(A) R_1(C) R_2(B) W_2(B) R_3(B) R_1(A) R_3(C) W_3(C) W_1(A)$*

(ii) Is the above schedule conflict-serializable? If so, what order of the three transactions defines a conflict-equivalent serial schedule?

(iii) What is the two-phase locking protocol? How does it guarantee serializability?

**(3 + 3 + 6) = 12**

9. (a) What is indexed sequential file organization? What are the applications of this organization?

(b) Describe B+ tree index files.

**8 + 4 = 12**