

**DBMS**  
**(INFO 3104)**

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) A relation R(P,Q,R,S) is decomposed into R1(P,Q) and R2(Q,R,S), where Q is the candidate key of R2. So, the decomposition is  
 (a) lossy (b) is lossless  
 (c) both (a) and (b) (d) none of these.
- (ii) The overall design of a database is called  
 (a) schema of the database (b) structure of the database  
 (c) the screen of the database (d) view of the database.
- (iii) If a schedule S can be transformed into a schedule S' by a series of swaps of non-conflicting instructions, we say that S and S' are  
 (a) conflict serializable (b) view equivalent  
 (c) conflict equivalent (d) view serializable.
- (iv) Relation R = (A,B,C,D) with AB as primary key. Choose one FD such that R should be in 1NF but not in 2NF.  
 (a) AB -> C (b) AB -> D (c) A -> D (d) AB -> CD.
- (v) It is an abstraction through which relationships are treated as higher level entities. What is it?  
 (a) Generalization (b) Specialization  
 (c) Aggregation (d) Inheritance.
- (vi) In multilevel indexing of a file containing n records, the last level contains how many blocks?  
 (a) n (b) n/2 (c) 1 (d) None.
- (vii) Precedence graphs help to find a  
 (a) serializable schedule (b) recoverable schedule  
 (c) deadlock free schedule (d) cascadeless schedule.

- (viii) Checkpoints are a part of  
 (a) recovery measures (b) security measures  
 (c) concurrency measures (d) authorization measures.
- (ix) An index on a non ordering key field is a  
 (a) clustering index (b) secondary index  
 (c) hash index (d) primary index.
- (x) Wait-die scheme for preventing deadlock is a  
 (a) pre-emptive scheme based on timestamp  
 (b) non pre-emptive scheme  
 (c) pre-emptive scheme  
 (d) non pre-emptive scheme based on timestamp.

**Group - B**

2. (a) Explain with diagram the functions of dbBuffer Cache, log buffer and library cache in the instance of Oracle server architecture. What happens in the background when user generates an update operation?  
 (b) Using three-schema architecture, explain the concept of logical data independence and physical data independence. **8 + 4 = 12**
3. (a) Produce an entity relationship diagram for a "Library Information System" with the following description.  
 The system is used to keep track of the books available for loan to card holders. Each book has an ISBN number (and no two books have the same ISBN number). In order to answer queries about the books in the library it is necessary for the system to keep track of the title and author of each book as well. Titles and authors are not generally unique. (Detailed information about authors won't be maintained, so you can assume that the "author" for a book will be represented by a short character string.)  
 There may be more than one copy of a book in the library. Each copy of the same book in the library has a different "copy number." In general, a "copy number" of a book is a small integer; if there are n copies of some book in the library then these copies have copy numbers from 1 to n inclusive. Each copy of a book might have a different storage location (building and floor number) and the system will be used to maintain this information as well. The system also stores information about cardholders (who are entitled to borrow books from the library). Each cardholder has a unique card number, as well as a name (first name, middle initial, and last name), telephone number, and

address (first line, second line, city, province, and postal code). Each cardholder can have zero or more copies of books out on loan; clearly, a copy of a book can be out on loan to at most one cardholder at a time. It is not necessary for this system to maintain information about a loan after the book has been returned to the library. Since the system will be used to produce reminder notices of overdue books it is necessary to remember the due date for each outstanding loan.

- (b) What is a weak entity? When can we convert a weak entity into an attribute?  
**8 + (2 + 2) = 12**

### Group – C

4. (a) Consider the relational database below, where the primary keys are underlined. Give an expression in the relational algebra to express each of the following queries:  
*employee* (person-name, street, city)  
*works* (person-name, company-name, salary)  
*company* (company-name, city)  
*manages* (person-name, manager-name)
- (i) Find the names, street address, and cities of residence of all employees who work for First Bank Corporation and earn more than \$10,000 per annum.  
(ii) Find the names of all employees in this database who live in the same city as the company for which they work.  
(iii) Find the names of all employees in this database who do not work for First Bank Corporation.  
(iv) Assume the companies may be located in several cities. Find all companies located in every city in which Small Bank Corporation is located.
- (b) Explain referential integrity constraint.  
**(2.5 × 4) + 2 = 12**

5. Consider the following relations:  
(i) **Employee** (Emp\_Code, Emp\_Name, Desig, Manager, Date\_of\_Joining, Basic, Dept\_Code)  
With Constraints: *Primary Key* is Emp\_Code  
*Foreign Key*: Manager references Employee(Emp\_Code),  
Dept\_Code references Department(Dept\_Code)  
(ii) **Department** (Dept\_Code, Dept\_Name, Location)  
With Constraints: *Primary Key* is Dept\_Code

Write the following queries in SQL:

- (i) List the names of the employees who are earning the lowest salary in each department;  
(ii) List the names of the employees who are earning more than the lowest salary of an employee in department 30;  
(iii) List average salaries of those departments which are employing at least 4 people;  
(iv) List the names of the employees who joined the company before their respective managers did.

$$(2 + 3 + 3 + 4) = 12$$

### Group – D

6. (a) A relation R with attributes A,B,C have the following FDs holding on it: Normalize R into BCNF.  
1)  $AB \rightarrow C, C \rightarrow B$   
2)  $A \rightarrow B, B \rightarrow C, C \rightarrow A$
- (b) (i) What is transitive dependency? Explain with an example.  
(ii) What problems occur if there is transitive dependency? Explain with proper examples.  
(iii) How is it corrected?  
**6 + (2 + 3 + 1) = 12**
7. (a) Consider a relation R with attributes subjects, instructor\_id and reference-books. A subject can be taught by many instructors and a subject can have many reference books. Is the relation in 4NF? Justify for or against.  
(b) Find the irreducible set of the following FD's given.  
(i)  $X \rightarrow YZ, XY \rightarrow P, XZ \rightarrow P$   
(ii)  $A \rightarrow B, B \rightarrow C, AB \rightarrow C, A \rightarrow C$   
**6 + (3 + 3) = 12**

### Group – E

8. (a) Consider the three transactions : T1, T2 and T3 and the schedules S<sub>1</sub> and S<sub>2</sub> given below. Draw the precedence graphs for S<sub>1</sub> and S<sub>2</sub>, and determine whether each schedule is serializable or not. If serializable, give its serial order(s).  
T<sub>1</sub>: r<sub>1</sub>(X); r<sub>1</sub>(Z); w<sub>1</sub>(X);  
T<sub>2</sub>: r<sub>2</sub>(Z); r<sub>2</sub>(Y); w<sub>2</sub>(Z); w<sub>2</sub>(Y);  
T<sub>3</sub>: r<sub>3</sub>(X); r<sub>3</sub>(Y); w<sub>3</sub>(Y);  
S<sub>1</sub>: r<sub>1</sub>(X); r<sub>2</sub>(Z); r<sub>1</sub>(Z); r<sub>3</sub>(X); r<sub>3</sub>(Y); w<sub>1</sub>(X); w<sub>3</sub>(Y); r<sub>2</sub>(Y); w<sub>2</sub>(Z); w<sub>2</sub>(Y);

$S_2: r_1(X); r_2(Z); r_3(X); r_1(Z); r_2(Y); r_3(Y); w_1(X); w_2(Z); w_3(Y); w_2(Y);$

- (b) Justify with example, for or against the Statement.
- (i) "Two phase locking does guarantee freedom from deadlock."
  - (ii) "Every cascadeless schedule is also recoverable".
- (4 + 4) + (2 + 2) = 12**

9. (a) Suppose you have a disk file with 40000 records. Each record length is 100 bytes. The size of each block is 1024 bytes. A primary index is prepared on an ordered key attribute. Each record in the index is of size 18 bytes.
- (i) Find out the blocking factor of the index.
  - (ii) Find out the number of blocks in the index.
  - (iii) Find how many block access is required to search a record.
- (b) What is the need of a log in a DBMS? Briefly describe the various types of records that are normally present in a log.
- (2 + 2 + 2) + (2 + 4) = 12**