

**AMIETE – CS/IT**

Time: 3 Hours

**JUNE 2013**

Max. Marks: 100

*PLEASE WRITE YOUR ROLL NO. AT THE SPACE PROVIDED ON EACH PAGE IMMEDIATELY AFTER RECEIVING THE QUESTION PAPER.*

**NOTE: There are 9 Questions in all.**

- Question 1 is compulsory and carries 20 marks. Answer to Q.1 must be written in the space provided for it in the answer book supplied and nowhere else.
- The answer sheet for the Q.1 will be collected by the invigilator after 45 Minutes of the commencement of the examination.
- Out of the remaining EIGHT Questions answer any FIVE Questions. Each question carries 16 marks.
- Any required data not explicitly given, may be suitably assumed and stated.

**Q.1 Choose the correct or the best alternative in the following: (2×10)**

- a. Which one of the following is not necessarily a characteristic of a database system?
- (A) Self describing nature                      (B) Data independence  
(C) Data accuracy                                (D) Support for multiple views of the data
- b. Which one of the following are important component of the DBMS?
- (A) Query optimizer                              (B) Database manager  
(C) File Manager                                 (D) All of these
- c. Which one of the following is not included in the classical E-R Model?
- (A) Entities                                        (B) Relationships  
(C) Integrity constraints                        (D) Attribute
- d. Which one of the following is not correct?
- (A) Each entity must have attributes.  
(B) Each Relationship must have attributes.  
(C) Some attributes must be able to uniquely identify each entity instance.  
(D) Identifiers of the entities participating in the relationship are sufficient to uniquely identify each relationship instance.
- e. Which one of the following is correct?
- (A) A relation may have duplicate rows.  
(B) A relation may have duplicate values in a column.  
(C) A relation may have ordering associated with the rows.  
(D) A relation may have non-atomic attributes values.
- f. Which one of the following is correct?
- (A) Each candidate key is a primary key.  
(B) Each primary key is a foreign key.  
(C) Each foreign key is a primary key in some relation.  
(D) Each foreign key is a secondary key.

- g. Which one of the following is not correct?
- (A) A selection selects one or more rows of a relation.  
 (B) A projection selects one or more columns of a relation.  
 (C) A join glues each any one row of one relation with all the rows of the other.  
 (D) A difference gives all the rows in the first relation that are not in the second.
- h. Which one of the following operators require elimination of duplicates?
- (A) Projection (B) Intersection  
 (C) Difference (D) Join
- i. Which one of the following is correct?
- (A) All functional dependencies are many to many relationship.  
 (B) All functional dependencies are many to one relationship.  
 (C) All functional dependencies are one to one relationship.  
 (D) None of these.
- j. Which one of the following are true? A set of functional dependencies is irreducible if and only if
- (A) The left hand side of every FD is just one attribute.  
 (B) The right hand side of every FD is just one attribute.  
 (C) Both the left and right hand side of every FD is just one attribute.  
 (D) None of these.

---

**Answer any FIVE Questions out of EIGHT Questions.**  
**Each question carries 16 marks.**

---

- Q.2** a. With the help of examples, explain the following terms briefly:  
 entity set, one-to-many relationship, participation constraint, weak entity set. (8)
- b. Which of the following plays an important role in *representing* information about the real world in a database? Explain briefly.
- (i) The data definition language. (4)  
 (ii) The data manipulation language.
- c. List four significant differences between a file-processing system and a DBMS. (4)
- Q.3** a. Using suitable example, explain the following relational operations:  
 SELECT, JOIN, DIVISION (6)
- b. Explain the entity integrity and referential integrity constraints. (6)
- c. Define foreign key. What is this concept used for? (4)
- Q.4** a. What is a view? How is it different from a table? (4)
- b. Consider the following employee database, where the primary keys are underlined.  
*employee* (employee-name, street, city)  
*works* (employee-name, company-name, salary)  
*company* (company-name, city)

*manages (employee-name, manager-name)*

Give an expression in SQL for each of the following queries.

- (i) Find the names, street addresses, and cities of residence of all employees who work for first Bank Corporation and earn more than \$10,000.
- (ii) Find all employees in the database who live in the same cities as the companies for which they work.
- (iii) Find all employees in the database who live in the same cities and on the same streets as do their managers. **(12)**

**Q.5** a. Explain with the help of examples, the concept of insertion anomalies and deletion anomalies. **(8)**

b. Given the relational schema  $R(A,B,C,D,E)$  and given the following set of FDs defined on  $R$

$F = \{A \rightarrow BC, CD \rightarrow E, AC \rightarrow E, B \rightarrow D, E \rightarrow AB\}$

- (i) Determine a lossless join decomposition of  $R$ .
- (ii) Determine decomposition of  $R$  which is not lossless join. **(8)**

**Q.6** a. What is collision in hashing? What are the various ways of collision resolution? **(8)**

b. What is a B-tree? Describe the structure of B-tree nodes. **(8)**

**Q.7** a. Let relations  $r1(A,B,C)$  and  $r2(C,D,E)$  have the following properties:  $r1$  has 20,000 tuples,  $r2$  has 45,000 tuples, 25 tuples of  $r1$  fit on one block, and 30 tuples of  $r2$  fit on one block. Estimate the number of block accesses required, using each of the following join strategies for  $r1$  &  $r2$ . **(8)**

- (i) Nested-loop join
- (ii) Hash join

b. With the help of a diagram, explain the steps required in processing of a high-level query. **(8)**

**Q.8** a. List the ACID properties. Explain the usefulness of each. **(6)**

b. During its execution, a transaction passes through several states, until it finally commits or aborts. List all possible sequences of states through which a transaction may pass. Explain why each state transition may occur. **(6)**

c. What is a cascadeless schedule? Why is cascadelessness of schedules desirable? Are there any circumstances under which it would be desirable to allow non cascadeless schedules? Explain your answer. **(4)**

**Q.9** a. Briefly explain recovery techniques based on immediate update. **(6)**

b. What is the difference between stable storage and disk? **(4)**

c. What are the roles of the Analysis, Redo, and Undo phases in ARIES? **(6)**

- (i) What is done during Analysis?
- (ii) What is done during Redo?
- (iii) What is done during Undo?