

Code: AC14 / AT11 Subject: DATABASE MANAGEMENT SYSTEMS

AMIETE – CS/IT (OLD SCHEME)

Time: 3 Hours

DECEMBER 2011

Max. Marks: 100

NOTE: There are 9 Questions in all.

- Please write your Roll No. at the space provided on each page immediately after receiving the Question Paper.
- 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. A database schema is specified by a set of definitions expressed by a special language called as _____.
- (A) Data Definition language (B) Data Manipulation language
(C) Data Control language (D) Data Query language
- b. A _____ is set of one or more attributes, taken collectively that allow us to identify uniquely an entity in the entity set
- (A) Candidate key (B) Primary key
(C) Super key (D) Weak key
- c. The clause that corresponds to the Cartesian-product operation of the relational algebra is
- (A) select (B) from
(C) where (D) none of these
- d. If α is a set of attributes and $\beta \subseteq \alpha$, then $\alpha \rightarrow \beta$ holds. This rule is known as
- (A) Union rule (B) Transitivity rule
(C) Decomposition rule (D) Reflexivity rule
- e. Which RAID level is also known as block-interleaved parity organization?
- (A) RAID level 0 (B) RAID level 3
(C) RAID level 4. (D) RAID level 5
- f. The organization in which the records of several different relations can be stored in the same file refers to.
- (A) Hashing file organization (B) Clustering file organization
(C) Heap file organization (D) B tree file organization

Code: AC14 / AT11 Subject: DATABASE MANAGEMENT SYSTEMS

- g. In query processing, name the lowest level operator to access data?
- (A) File scan (B) Index scan
(C) Searching (D) Sorting
- h. Ensuring durability is the responsibility of a component of the database system called the
- (A) Transaction-management component
(B) Concurrency-control component
(C) Recovery-management component
(D) Query-optimization component
- i. The phenomenon, in which a single transaction failure leads to a series of transaction rollbacks, is called _____.
- (A) Cascading rollback (B) Uncommitted rollback
(C) Conflict rollback (D) Topological rollback
- j. The protocol that ensures that any conflicting **read** and **write** operations are executed in timestamp order is
- (A) Graph-Based locking protocol (B) Multiple-Granularity protocol
(C) Two-Phase Locking protocol (D) Timestamp-Ordering protocol

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

- Q.2** a. What do you mean by Data Abstraction? Explain the difference between Physical level, Logical level and View level of data abstraction? (6)
- b. Construct an E-R diagram for a car-insurance company whose customers own one or more cars each. Each car has associated with it zero to any number of recorded accidents. (10)
- Q.3.** a. Explain the following relational algebra operations with their notations by suitable example
- (i) SELECT
(ii) SET DIFFERENCE
(iii) PROJECT
(iv) ASSIGNMENT (2×4)
- b. Let the following relation schemas be given:
 $R = (A, B, C)$
 $S = (D, E, F)$
 Let relations $r(R)$ and $s(S)$ be given. Give an expression in the tuple relational calculus that is equivalent to each of the following:
- (i) $\Pi_A(r)$
 (ii) $\sigma_{B=19}(r)$
 (iii) $r \times s$
 (iv) $\Pi_{A,F}(\sigma_{C=D}(r \times s))$ (8)

Code: AC14 / AT11 Subject: DATABASE MANAGEMENT SYSTEMS

- Q.4** a. 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 of all employees who work for ABCL Corporation.
 (ii) Find all employees in the database who live in the same cities as the companies for which they work.
 (iii) Find those companies whose employees earn a higher salary, on average, than the average salary at ABCL Corporation.
 (iv) Give all managers of ABCL Corporation a 10 percent raise.
 (v) Delete all tuples in the *works* relation for employees of BOB Corporation. **(10)**
- b. Consider the following account schema.
account(branch_name, account_number, balance)
 Write an SQL query, without using a **with** clause, to find all branches where the total account deposit is less than the average total account deposit at all branches,
 (i) Using a nested query in the **from** clause.
 (ii) Using a nested query in a **having** clause. **(6)**
- Q.5** a. Explain the normal form called as Boyce-Codd (BCNF)? Prove that any relation schema with two attributes is in BCNF? **(8)**
- b. What is referential integrity? What is the need for triggers and how are they implemented in SQL? Give Example. **(8)**
- Q.6** a. Discuss Byte-String Representation technique for implementing variable-length records? **(8)**
- b. Since indices speed-up query processing, why might they not be kept on several search keys? List as many reasons as possible. **(4)**
- c. Discuss the various cost components for Query execution? **(4)**
- Q.7** a. What role Sorting play in database systems? How sorting accomplished in database system? Discuss the different phases of external sorting? **(8)**
- b. Consider the following two transactions:
 T_1 : read(A);
 read(B);
 if A = 0 **then** B := B + 1;
 write(B).
 T_2 : read(B);
 read(A);
 if B = 0 **then** A := A + 1;
 write(A).

Code: AC14 / AT11 Subject: DATABASE MANAGEMENT SYSTEMS

Let the consistency requirement be $A = 0 \vee B = 0$, with $A = B = 0$ the initial values.

- (i) Show that every serial execution involving these two transactions preserves the consistency of the database.
- (ii) Show a concurrent execution of T_1 and T_2 that produces a nonserializable schedule.
- (iii) Is there a concurrent execution of T_1 and T_2 that produces a serializable schedule? (3+3+2)

Q.8 a. What is “Two-Phase locking protocol”? Explain with the help of example of a schedule, how the protocol ensures a schedule to be conflict-serializable, but not Cascade-less? (8)

b. Define Deadlock? What are the necessary four conditions for a deadlock to occur? Discuss the different methods that can be used for deadlock prevention? (8)

Q.9 Explain the following:

- (i) Functions of Database Administrator (DBA)
- (ii) Embedded SQL
- (iii) Loss-Less Join decomposition
- (iv) ACID properties (4×4)