



(Pages : 3)

8801

Reg. No. :

Name :

Fifth Semester B.Tech. Degree Examination, December 2015

(2008 Scheme)

08.503 : DATABASE DESIGN (R)

Time : 3 Hours

Max. Marks : 100

PART - A

Answer all questions :

1. Explain and illustrate how primary index is different from clustered index.
2. Distinguish between physical data independence and logical data independence.
3. Can we always replace weak entities with multi-valued, composite attributes ? Why ?
4. Distinguish between declarative and procedural query languages.
5. Explain and illustrate the following concepts : super key, minimal super key, candidate key.
6. Explain and illustrate the difference between theta join, natural join and equi-join with real examples.
7. When do we say that a relation is in 1NF ? Give a typical 1NF relation.
8. Write a brief note on Fourth Normal Form.
9. Why concurrency control is needed in DBMS ?
10. What is meant by transaction roll-back ? List out the situations in which transaction roll-back is executed ?

(10×4=40 Marks)



P.T.O.



PART - B

Answer **one full** question from **each** Module :

Module - I

11. a) Explain and illustrate participation and cardinality constraints with typical real examples.

b) Design an E-R diagram for the following scenario :

A car company maintains details of cars, dealers and customers. Each car is identified by a vehicle identification number. Each car has a particular model and a version (for example, 'Maruti' offers the model 'Swift' in different versions : VXi, Lxi and so on). Each car is offered with zero or more number of options (like Air-bag, A/c and so on). The database needs to store information about models, versions and options, as well as information about dealers and customers.

Justify your design decisions. Document any assumptions made about entities, attributes and relationships.

OR

12. a) List out various types of users of a database system and give their roles/activities.

b) Consider a data file with 6200 records each of size 35 bytes. A multi-level secondary index is constructed on a field of the record and the size of this field is 12 bytes. The pointer field of the index record is 8 bytes. Assuming that the block size for the storage device is 512 bytes, compute the following :

blocking factor, number of data blocks, number of levels of index and the number of blocks at each level.

Module - II

13. a) Give a brief outline of various types of anomalies in a DB system.

b) Consider the R(A, B, C, D) with functional dependencies $A \rightarrow B$, $B \rightarrow C$, $C \rightarrow D$, $D \rightarrow A$

i) Compute closure of B with respect to the given dependencies.

ii) Identify key(s) of R

iii) Normalize R.

OR



14. a) Briefly discuss and illustrate GROUP BY and ORDER BY clauses of SQL with a real example.
- b) Consider the following relational schema. Primary keys are underlined and foreign keys bear the same name as primary keys. In addition, *advisor* is a foreign key referring to *empld* of *Faculty* and *preReqId* is a primary key referring to *courseId* of *Course*. A semester can be odd or even.

Student (rollNo, name, degree, year, sex, *advisor*)

Department(deptId, name, hod, phone)

Faculty(empld, name, sex, *deptId*, phone)

Course(courseId, cName, *deptId*)

Enrollment(rollNo, courseId, sem, year)

Teaching(empld, courseId, classRoom)

Pre-requisite(courseId, preReqId)



Express the following queries in relational algebra and SQL :

- i) Find names of faculty in 'Computer Science' department.
- ii) Find names of courses *not* offered by "Mathematics" department.
- iii) Find names of students who have enrolled for *all* the courses in the odd semester of the year "2012".
- iv) Find names of students who have *not* taken any course offered by his/her advisor (SQL only).

Module – III

15. a) What are the various types of failures in a database system ? Briefly discuss.
- b) What is system log used for ? What are the typical kinds of records in the system log ?
- c) For the following schedule, state, with suitable reasoning, whether it is serializable and/or conflict serializable, assuming that b_1 and b_2 are variables and T_1 and T_2 are transaction ids. Also show the precedence graph for the schedule.
- $read(T_1, b_2), read(T_2, b_1), write(T_1, b_1), write(T_2, b_1), commit(T_1), commit(T_2).$

OR

16. a) Briefly outline time-stamp based concurrency control.
- b) Discuss how two-phase locking ensures serializability.

(3x20=60 Marks)