

Code No: R22053

R10**SET - 1****II B. Tech II Semester Supplementary Examinations, April/May-2017****DATA BASE MANAGEMENT SYSTEMS**

(Com. to CSE, IT)

Time: 3 hours

Max. Marks: 75

Answer any **FIVE** Questions
All Questions carry **Equal** Marks
~~~~~

1. a) 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? (8M)  
b) What are the main benefits of using a DBMS to manage data in applications involving extensive data access? (7M)
2. a) How the choice between using aggregation or a ternary relationship is determined in ER diagram design? (8M)  
b) What is a weak entity? What are class hierarchies? What is aggregation? Give an example scenario motivating the use of each of these ER model design constructs. (7M)
3. Consider the following schema: (Assume the data) (15M)  
Sailors(sid: integer, sname: string, rating: integer, age: real)  
Boats(bid: integer, bname: string, color: string)  
Reserves (sid: integer, bid: integer, day: date)  
Write relational algebra queries to i) Find the names of sailors who have reserved a red boat. ii) Find the names of sailors who have reserved at least one boat  
iii) Find the names of sailors who have reserved at least two boats.
4. How can we translate an ER diagram into SQL statements to create tables? How are entity sets mapped into relations? How is relationship sets mapped? How are constraints in the ER model, weak entity sets, class hierarchies, and aggregation handled? (15M)
5. a) Define 'Multivalued dependencies and Join dependencies. Discuss the use of such dependencies for database design. (10M)  
b) Give a set. Of FDs for the relation schema R(A,B, C,D) with primary key AB under which R is in 1NF but not in 2NF. (5M)
6. a) Show that the two phase locking protocol ensures conflict serializability and that transactions can be serialized according to their lock points. (8M)  
b) Compare the deferred and immediate modification versions of the log based recovery schemes, in terms of ease of implementation and over head cost. (7M)
7. How does a hash-based index handle an equality query? Discuss the use of the hash function in identifying a bucket to search. Given a bucket number, explain how the record is located on disk. (15M)
8. Describe the B+ tree insertion algorithm, and explain how it eliminates. Over flow pages. Under what conditions can an insert increase the height of the tree? (15M)