Code No: R22053   R10   SET - 1

II B. Tech II Semester Regular Examinations August - 2014
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 is data model? List and explain different data models.

    b) Explain the difference between external, internal, and conceptual schemas. How are these
       different schema layers related to the concepts of logical and physical data independence?

                                                                            (8M+7M)

2. a) Discuss in detail about the main steps in the database design and clearly focus in detail about
       the goal of each step? In which steps is the ER model mainly used?

    b) Draw and explain E-R diagram of an Airline reservation system?   (8M+7M)

3. a) Discuss in detail about integrity constraint over relations.

    b) What is join operation in relational algebra? Discuss in detail about variants of joins.

                                                                            (8M+7M)

4. a) Consider the following relational schema. An employee can work in more than one
       department; the pct time field of the Works relation shows the percentage of time that a given
       employee works in a given department.
       Emp(eid: integer, ename: string, age: integer, salary: real)
       Works(eid: integer, did: integer, pct time: integer)
       Dept(did: integer, budget: real, managerid: integer)
       Write the following queries in SQL:
       i) Print the names and ages of each employee who works in both the Hardware department
          and the Software department.
       ii) For each department with more than 20 full-time-equivalent employees (i.e., where the
           part-time and full-time employees add up to at least that many full-time employees), print
           the did together with the number of employees that work in that department.
       iii) Print the name of each employee whose salary exceeds the budget of all of the
             departments that he or she works in.
       iv) Find the manager ids of managers who manage only departments with budgets greater than
           $1,000,000.
       v) Find the enames of managers who manage the departments with the largest budget.

    b) What is a trigger? What are the three parts? Differentiate row-level and statement-level
       triggers.   (8M+7M)

Code No: R22053   R10   SET - 1

5. a) What is join dependency? How is it different to that of multi-valued dependency and functional dependency? Give an example for join dependencies and multi-valued dependencies.

   b) Contrast 3NF decomposition method with BCNF decomposition method illustratively.

   (8M+7M)

6. a) What is a serializable schedule? What is a recoverable schedule? What is a schedule that avoids cascading aborts? What is a strict schedule?

   b) Discuss in detail about the phases the recovery manager proceeds when the system is restarted after a crash.   (8M+7M)

7. a) Describe in detail about algorithms for updating single level indices.

   b) Give comparison of different file organizations.   (8M+7M)

8. Describe a B+ tree for the following set of key values:(2,3,5,7,11,17,19,23,29,31)
   Assume that the tree is initially empty and values are added in ascending order.
   i) Construct B+ tree for the case where the number of pointer that will fit in one node is four.
   ii) Show the step involved to find records with a search-key value of 11.   (8M+7M)

Code No: R22053   R10   SET - 2

II B. Tech II Semester Regular Examinations August - 2014
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?

    b) Which of the following plays an important role in representing information about the real world in a database? Explain briefly about:
    i) The data definition language.
    ii) The data manipulation language.
    iii) The buffer manager.
    iv) The data model.   (7M+8M)

2. a) Describe in detail about conceptual design with ER model.
    b) Construct E-R diagram for a banking enterprise.   (8M+7M)

3. a) Using the following schema represent the following queries using Tuple relational calculus :
    PROJECT (Projectnum, Project Name, Project Type, Project Manager)
    EMPLOYEE (Empnum, Empname)
    ASSIGNED_TO (Projectnum, Empnum)
    Find Employee details working on a project name starts with 'L'
    List all the employee details who are working under project manager "Clevee"
    List the employees who are still not assigned with any project.
    List the employees who are working in more than one project.

    b) What is view, updatable and non-updatable views? Explain the advantages of view in maintaining the security.   (8M+7M)

Code No: R22053   **R10**   SET - 2

4. a) What are nested queries? What is correlation in nested queries? How would you use the operators IN, EXISTS, UNIQUE, ANY and ALL in writing nested queries? Why are they useful? Illustrate your answer by showing how to write the division operator in SQL.

    b) Compare constraints and triggers and also give examples for each.   (8M+7M)

5. a) Consider a relation R with five attributes ABCDE. You are given the following dependencies: A $\rightarrow$ B, BC $\rightarrow$ E, and ED $\rightarrow$ A.

    List all keys for R.

    Is R in 3NF?

    Is R in BCNF?

    b) What is decomposition? Describe problems related to decomposition.   (9M+6M)

6. a) Is every conflict serializable schedule is serializable? Explain.

    b) Explain different types of failures that arise due to loss of non-volatile storage.   (8M+7M)

7. a) What is an index? Discuss important properties of an index that affect the efficiency of searches using the index.

    b) Describe in detail about different RAID levels.   (8M+7M)

8. a) What are the main differences between ISAM and B+ tree indexes?

    b) Describe a B+ tree for the following set of key values:(2,3,5,7,11,17,19,23,29,31) Assume that the tree is initially empty and values are added in ascending order. Construct B+ tree for the case where the number of pointer that will fit in one node is six.   (6M+9M)

Code No: R22053   R10   SET - 3

II B. Tech II Semester Regular Examinations August - 2014
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 application programs? Discuss in detail about database access for application programs.
   b) Explain the difference between logical and physical data independence. What is logical data independence and why is it important?   (8M+7M)

2. a) A university database contains information about professors (identified by social security number, or SSN) and courses (identified by courseid). Professors teach courses; each of the following situations concerns the Teaches relationship set. For each situation, draw an E-R diagram that describes it (assuming that no further constraints hold).
      Professors can teach the same course in several semesters, and each offering must be recorded.
      Professors can teach the same course in several semesters, and only the most recent such offering needs to be recorded.
      Every professor must teach some course.
      Every professor teaches exactly one course (no more, no less).
      Every professor teaches exactly one course (no more, no less), and every course must be taught by some professor.
   b) Explain the difference between weak entity and strong entity set? How to represent the strong entity and weak entity set through E-R diagram.   (10M+5M)

3. a) What is the difference between a candidate key and the primary key for a given relation? What is a super key?
   b) Discuss in detail about integrity constraints over relations.   (8M+7M)

4. a) Discuss the strengths and weaknesses of the trigger mechanism. Contrast triggers with other integrity constraints supported by SQL.
   b) Explain with examples about Logical connectivity's – AND, OR and NOT in detail. (8M+7M)

Code No: R22053 $\boxed{\text{R10}}$ $\boxed{\text{SET - 3}}$

5. a) Consider the following relation R(A,B,C,D,E)
       And FD's
       A→BC C→A   D→E   F→A   E→D
       Is the decomposition of R into R_1(A,C,D), R_2(B,C,D) and R_3(E,F,D) lossless? Explain the requirements of lossless decomposition.
    b) Explain in detail about Dependency-Preserving Decomposition. Explain why it is important.
                                                                        (8M+7M)

6. a) What overheads are associated with lock-based concurrency control? Discuss blocking and aborting overheads specifically?
    b) Define these terms: atomicity, consistency, isolation, durability, schedule, blind write, dirty read, unrepeatable read, serializable schedule, recoverable schedule, avoids-cascading- aborts schedule.   (8M+7M)

7. a) What are the causes of bucket overflow in a hash file organization? What can be done to reduce the occurrence of bucket overflows?
    b) Discuss about multilevel indices in detail.   (8M+7M)

8. a) Explain main characteristics of a B+ tree in detail. Discuss operations on B+ trees.
    b) Describe a B+ tree for the following set of key values:(2,3,5,7,11,17,19,23,29,31)
       Assume that the tree is initially empty and values are added in ascending order. Construct B+ tree for the case where the number of pointer that will fit in one node is six.   (6M+9M)

Code No: R22053   **R10**   SET - 4

II B. Tech II Semester Regular Examinations August - 2014
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) Why would choose a database system instead of simply storing data in operating system files? When would it make sense not to use a database system?
    b) Shows the structures of a typical DBMS based on the relational data model and explain in detail.   (8M+7M)

2. a) Construct an E-R diagram for a university registrar's office. The office maintains data about each class, including the instructor the enrollment, and the time and place of the class meetings. For each student-class pair, a grade is recorded. Document all assumptions that you make about the mapping constraints.
    b) Explain the difference between weak entity and strong entity sets? How to represent the strong and weak entity set through E-R diagrams.   (8M+7M)

3. a) What restrictions are necessary to ensure that view is updatable? State any three advantages and three disadvantages of views.
    b) Discuss in detail about the set operations of relational algebra and explain with examples.
                                                                      (8M+7M)

4. a) The following relations keep track of airline flight information:
       Flights(flno: integer, from: string, to: string, distance: integer, departs: time, arrives: time, price: integer) Aircraft(aid: integer, aname: string, cruisingrange: integer)
       Certified(eid: integer, aid: integer)
       Employees(eid: integer, ename: string, salary: integer)
       Note that the Employees relation describes pilots and other kinds of employees as well; every pilot is certified for some aircraft, and only pilots are certified to fly. Write each of the following queries in SQL.
       i) Find the names of aircraft such that all pilots certified to operate them earn more than 80,000.
       ii) For each pilot who is certified for more than three aircraft, find the eid and the maximum cruisingrange of the aircraft that he (or she) is certified for.
       iii) Find the names of pilots whose salary is less than the price of the cheapest route from Los Angeles to Honolulu.
       iv) For all aircraft with cruising range over 1,000 miles, find the name of the aircraft and the average salary of all pilots certified for this aircraft.
    b) Explain in detail about different aggregative operators in SQL with example queries.
                                                                      (8M+7M)

Code No: R22053   **R10**   SET - 4

5. a) What is schema refinement? Discuss the problems caused by redundancy.
    b) Give a set of FDs for the relation schema R(A,B,C,D) with primary key AB under which R is in 2NF but not in 3NF.   (8M+7M)

6. a) Explain various types of lock based concurrency control with a neat sketch and examples.
    b) Describe in detail about shadow paging recovery technique. Under what circumstances does it not require a log?   (8M+7M)

7. a) Explain the distinction between closed and open hashing. Discuss the relative merits of each technique in database applications.
    b) On what factors techniques for indexing and hashing must be evaluated? Explain.   (8M+7M)

8. Explain all the operations on B+ tree by taking a sample example.   (15M)