AB
21/6/13 AN

Reg. No. : | | | | | | | | | | | | |

Question Paper Code : 65030

5 Year M.Sc. DEGREE EXAMINATION, MAY/JUNE 2013.

Third Semester

Computer Technology

XCS 235/10677 SW 305 — DATABASE MANAGEMENT SYSTEMS

(Common to 5 Year M.Sc. – Information Technology/Software Engineering)

(Regulation 2003/2010)

Time : Three hours Maximum : 100 marks

Answer ALL questions.

PART A — (10 × 2 = 20 marks)

1. List four significant differences between file processing system and a DBMS.

2. List two reasons why null values might be introduced into the database.

3. How does the remapping of bad sectors by disk controllers affect data retrieval rates?

4. Define average latency time.

5. Write a query "to delete all instructors with a salary between 10,000 to 20,000".

6. Give the difference between first normal form and second normal form.

7. What is meant by transaction rollback?

8. Suppose that there is database system that never fails. Is a recovery manager required for this system?

9. What benefit does rigorous two phase locking provide?

10. What is shadow paging?

PART B — (5 × 16 = 80 marks)

11. (a) Draw an E— R Diagram for Banking Enterprises and University.

Or

(b) Draw the DBMS System Architecture and explain in detail.

12. (a) (i) Explain the various physical storage media with example.

(ii) In the sequential file organization, why is an overflow block used even if there is, at the moment, only one overflow record?

Or

(b) (i) Is it possible in general to have two clustering indices on the same relation for different search key? Explain your answer.

(ii) Discuss the various operations of B+ Tree.

13. (a) Explain detail in decomposition using Functional Dependencies.

Or

(b) Explain why 4NF is a normal form more desirable than BCNF.

14. (a) Describe how to incrementally maintain the results of the following operations, on both insertion and deletion.

(i) Union and set difference

(ii) Left outer join.

Or

(b) (i) List the ACID properties . Explain the usefulness of each.

(ii) Explain the distinction between the term serial schedule and serializable schedule.

15. (a) (i) Explain Recovery isolation levels with example.

(ii) When is a transaction rolled back under timestamp ordering, it is assigned a new timestamp? Why can it not simply keep its old timestamp?

Or

(b) (i) Show by example that there are schedules possible under the tree protocol that are not possible under the two phase locking protocol and vice versa.

(ii) Explain why log records for transactions on the undo-list must be processed in reverse order, whereas redo is performed in a forward direction.
