B.E./B.Tech. DEGREE EXAMINATION, APRIL/MAY 2011
Fourth Semester
Computer Science and Engineering
CS 2255 — DATABASE MANAGEMENT SYSTEMS
(Common to Information Technology)
(Regulation 2008)
Time : Three hours Maximum : 100 marks
Answer ALL questions

PART A — (10 × 2 = 20 marks)

1. Who is a DBA? What are the responsibilities of a DBA?
2. What is a data model? List the types of data models used.
3. What is embedded SQL? What are its advantages?
4. What is the difference between tuple relational calculus and domain relational
calculus?
5. What is meant by lossless-join decomposition?
6. A relation { } D C B A R , , , = has FD's { } A D D C C AB F → → → = , , . Is R is
in 3NF?
7. What are the ACID properties?
8. What are two pitfalls (problems) of lock-based protocols?
9. What are the advantages and disadvantages of indexed sequential file?
10. What is database tuning?

PART B — (5 × 16 = 80 marks)

11. (a) (i) With a neat diagram, explain the structure of a DBMS. (9)
(ii) Draw an E-R diagram for a small marketing company database, assuming your own data requirements. (7)
Or
(b) (i) Compare the features of file system with database system. (8)
(ii) Explain the differences between physical level, conceptual level and

view level of data abstraction. (4)
(iii) Mention any four major responsibilities of DBA. (4)

12. (a) (i) Consider the following relational database
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 to express each of the following queries :
Find the names and cities of residence of all employees who work for XYZ Bank. Find the names, street address, and cities of residence of all employees who work for XYZ Bank and earn more than Rs. 10,000 per annum.
Find the names of all employees in this database who live in the same city as the company for which they work.
Find the names of all employees who live in the same city and on the same street as do their managers. (4 × 3 = 12)
(ii) Define the term distributed database management system and mention the issues to be considered in the design of the same. (4)
Or
(b) (i) What are the relational algebra operations supported in SQL? Write the SQL statement for each operation. (12)
(ii) What is data integrity? Explain the types of integrity constraints. (4)

13. (a) (i) Explain 1NF, 2NF, 3NF and BCNF with suitable example. (8)
(ii) Consider the universal relation { } I H G F E D C B A R , , , , , , , , , = and the set of functional dependencies
( ) { } { } { } { } { } F B E D A C B A F → → → = , , , , { , { } { } { } { } } , , , J I D H G F → → .
What is the key for R? Decompose R into 2NF, then 3NF relations. (8)
Or
(b) What are the pitfalls in relational database design? With a suitable

example, explain the role of functional dependency in the process of normalization. (16)

14. (a) (i) Explain about immediate update and deferred update recovery techniques. (8)
(ii) Explain the concepts of serializability. (8)
Or
(b) (i) Explain Two-phase locking protocol. (8)
(ii) Describe about the deadlock prevention schemes. (8)

15. (a) (i) List the different levels in RAID technology and explain its features. (12)
(ii) Describe the different methods of implementing variable length records. (4)
Or
(b) (i) Explain the various indexing schemes used in database environment. (12)
(ii) Let relations ( ) C B A r , , 1 and ( ) E D C r , , 2 have the following properties : 1 r has 20,000 tuples, r2 has 45,000 tuples, 25 tuples of r1 fit on one block, and 30 tuples of r2 fit on one block. Estimate the number of block accesses required, using each of the following join strategies for r1 r2 : (4)
(1) Nested-loop join with r1 as outer relation
(2) Block nested-loop join with r1 as outer relation
(3) Merge join if r1 and r2 are initially sorted
(4) Hash join (assuming that no overflow occurs).