

Reg. No. :

--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--

Question Paper Code : 71674

B.E./B.Tech. DEGREE EXAMINATION, APRIL/MAY 2017.

Third/Fifth/Eighth Semester

Computer Science and Engineering

CS 6302 — DATABASE MANAGEMENT SYSTEMS

(Common to Mechanical and Automation Engineering, Mechatronics Engineering,
Information Technology)

(Regulations 2013)

Time : Three hours

Maximum : 100 marks

Answer ALL questions.

PART A — (10 × 2 = 20 marks)

1. What are the desirable properties of decomposition?
2. Distinguish between key and super key.
3. What is a query execution plan?
4. Which cost component are used most often as the basis for cost function?
5. What is serializable schedule?
6. What type of locking needed for insert and delete operations?
7. Define replication transparency.
8. State the function of data marts.
9. Define support and confidence.
10. Distinguish between threats and risks.

PART B — (5 × 13 = 65 marks)

11. (a) Discuss the correspondence between the ER model construct and the relational model constructs. Show how each ER model construct can be mapped to the relational model. Discuss the option for mapping EER model construct.

Or

- (b) (i) Explain the overall architecture of the data base system in detail.(8)
(ii) List the operations of relational algebra and the purpose of each with example. (5)
12. (a) What is meant by semantic query optimization? How does it differ from other query optimization technique? Give example.

Or

- (b) Justify the need of embedded SQL. Consider the relation student (Reg No, name, mark, and grade). Write embedded dynamic SQL Program in C language to retrieve all the students' records whose mark is more than 90. (2+11)
13. (a) Discuss the violations caused by each of the following: dirty read, non repeatable read and phantoms with suitable example.

Or

- (b) Explain why timestamp-based concurrency control allows schedules that are not recoverable. Describe how it can be modified through buffering to disallow such schedules.
14. (a) (i) Compare and contrast the distributed databases and the centralized database systems. (8)
(ii) Describe the mobile database recovery schemes. (5)

Or

- (b) Explain what a RAID system is. How does it improve performance and reliability. Discuss the level 3 and level 4 of RAID. (3 + 4 + 6)
15. (a) (i) What are the basic crawling operations. Explain the processing steps in crawling procedure with example. (8)
(ii) Explain the process of querying XML data with an example. (5)

Or

- (b) Describe the various components of data warehouse and explain the different data models used to store data with example.

PART C — (1 × 15 = 15 marks)

16. (a) Consider the relation schema given in Figure 1. Design and draw an ER diagram that capture the information of this schema. (5)




Figure 1.

Write the following queries in relational algebra and SQL.

- (i) Find the names of employees who have borrowed a book published by McGraw-Hill. (5)
- (ii) Find the names of employees who have borrowed all books published by McGraw-Hill. (5)

Or

- (b) Trace the results of using the Apriori algorithm on the grocery store example with support threshold $s = 33.34\%$ and confidence threshold $c = 60\%$. Show the candidate and frequent itemsets for each database scan. Enumerate all the final frequent itemsets. Also indicate the association rules that are generated and highlight the strong ones, sort them by confidence.

Transaction ID	Items
T1	HotDogs, Buns, Ketchup
T2	HotDogs, Buns
T3	HotDogs, Coke, Chips
T4	Chips, Coke
T5	Chips, Ketchup
T6	HotDogs, Coke, Chips