

C

0200CST204122302

Pages: 4

Reg No.: _____

Name: _____

APJ ABDUL KALAM TECHNOLOGICAL UNIVERSITY

Fourth Semester B.Tech Degree (S, FE) Examination January 2024 (2019 Scheme)



Course Code: CST 204

Course Name: Database Management Systems

Max. Marks: 100

Duration: 3 Hours

PART A

(Answer all questions; each question carries 3 marks)

Marks

- | | | |
|----|--|---|
| 1 | Differentiate between Structured, Semi-structured, and Unstructured data. Give an example each. | 3 |
| 2 | Explain the three categories of Data Models. | 3 |
| 3 | Explain the "Cross-Reference approach" to mapping a Binary 1:1 Relationship Type that you employ when you map an ER Model into a Relational Schema, with the help of an example. | 3 |
| 4 | With the help of an example database, explain the usage of the set operations Union, Intersection, and set difference in Relational Algebra. | 3 |
| 5 | "Views simplify the specification of certain queries". Justify this statement | 3 |
| 6 | Differentiate between Assertions and Triggers. | 3 |
| 7 | Define the term "Functional Dependency". Give an example. | 3 |
| 8 | Give an algorithm to find the Minimal Cover for a set of Functional Dependencies. | 3 |
| 9 | Give six types of failures in a transaction processing system that necessitate Recovery procedures. | 3 |
| 10 | Explain the importance of transaction logging and checkpointing for efficient transaction processing. | 3 |

PART B

(Answer one full question from each module, each question carries 14 marks)

Module -1

- | | | |
|----|---|---|
| 11 | a) Explain the main four characteristics of the Database Approach that distinguish it from the traditional file-processing approach | 7 |
| | b) Categorize the different types of people who work in a database system environment. | 7 |

- 12 Design an ER diagram for a typical college library database and then map it into a relational database schema. List your assumptions and indicate the cardinality mappings. 14

Module -2

- 13 a) Consider the following schema and frame Relational Algebra queries for the following problems: 11

Suppliers (SID: integer, SName: string, Address: string) Parts (PID: integer, PName: string, Color: string) Catalog (SID: integer, PID: integer, Cost: real)

The key fields are underlined and the domain of each field is given after the field name.

- (i) Find the names of suppliers who supply **red** parts
 - (ii) Find the SIDs of suppliers who supply some **red** part **or** are at the address '221 Packer Ave'
 - (iii) Find the SIDs of suppliers who supply some **red** part **and** some **green** part
 - (iv) Find the SIDs of suppliers who supply **every red** part
- b) Differentiate between DELETE and DROP commands in SQL. Illustrate their usage. 3

- 14 a) Consider a company database having the following schema and frame Relational Algebra queries for the following problems. Primary keys are underlined. 10

EMPLOYEE (SSN, Name, SupervisorSSN, Dnum)

DEPARTMENT (Dnumber, Dname, MgrSSN)

PROJECT (Pnumber, Plocation, ControlDeptNum)

EMPLOYEE (Dnum) References DEPARTMENT (Dnumber) and

PROJECT (ControlDeptNum) References DEPARTMENT (Dnumber)

- (i) Find the names of all employees who are supervised by the supervisor of the employee named 'Smith'
 - (ii) List the numbers of the projects (Pnumbers) controlled by Smith's department
- b) Illustrate the usage of the SQL commands – ALTER, INSERT, DELETE and UPDATE 4

Module -3

- 15 Consider the following Database with two tables: 14

Table: Employees

EmployeeID INT PRIMARY KEY

FirstName VARCHAR(50)

LastName VARCHAR(50)

JobTitle VARCHAR(100)

Salary DECIMAL(10,2)

HireDate DATE

DepartmentID INT

Table: Departments

DepartmentID INT PRIMARY KEY

DepartmentName VARCHAR(100)

ManagerID INT

Foreign Key: Employees.DepartmentID references Departments.DepartmentID

Frame SQL queries for the following problems:

- (i) Calculate the average salary per department.
 - (ii) List the employees with the highest salary in each department:
 - (iii) Find departments with more than 25 employees.
 - (iv) Get the employee names starting with 'S' in alphabetical order
- 16 a) With the help of an example explain Single-level indexing and multi-level indexing. Also, compare and contrast single-level indexing with multi-level indexing. 8
- b) Explain a situation where a multi-level index would be significantly less effective than a single-level index, and vice versa 6

Module -4

- 17 a) Given the following FDs for the relation Employees (EmployeeID, DepartmentID, ManagerID, Salary): 6
- EmployeeID → DepartmentID
DepartmentID → ManagerID

ManagerID \rightarrow Salary

Identify any redundant FDs in the set and explain why they are redundant.

- b) Consider the following FDs for the relation Books (BookID, Title, Author, Publisher): 8

BookID \rightarrow Title, Author

Author \rightarrow Publisher

{Title, Publisher} \rightarrow BookID

Find a minimal cover for this set of FDs. Explain how you arrived at your answer.

- 18 a) Give an algorithm each for checking Lossless Join and Dependency Preserving Properties 6

- b) Suppose that we decompose the schema $R = (A, B, C, D, E)$ into (A, B, C) (A, D, E). 8

Show that this decomposition is a lossless decomposition if the following set F of functional dependencies holds:

$F = \{A \rightarrow BC, CD \rightarrow E, B \rightarrow D, E \rightarrow A\}$

Module -5

- 19 a) What is a serial schedule? Why are serial schedules unacceptable in practice? 4
- b) What is a conflict serializable schedule? Give an algorithm to check whether a schedule is conflict serializable or not. Check whether the following schedules are conflict serializable or not and find an equivalent serial schedule if possible. 10

(i) $r_2(X); w_2(X); r_1(X); w_1(X); r_1(Y); w_1(Y)$

(ii) $r_1(X); r_2(X); w_1(X); r_1(Y); w_2(X); w_1(Y)$

- 20 a) Explain the working of Binary Locks and Shared/Exclusive Locks 5
- b) Explain Two-Phase locking protocol and any three variants of it 9
