

(12)

ISLAMIC UNIVERSITY OF TECHNOLOGY (IUT)
ORGANISATION OF ISLAMIC COOPERATION (OIC)
Department of Computer Science and Engineering (CSE)

SEMESTER FINAL EXAMINATION
 DURATION: 3 HOURS

WINTER SEMESTER, 2022-2023
 FULL MARKS: 150

CSE 4307: Database Management Systems

Programmable calculators are not allowed. Do not write anything on the question paper.

Answer all 6 (six) questions. Figures in the right margin indicate full marks of questions whereas corresponding CO and PO are written within parentheses.

1. a) Briefly explain the major characteristics of Big Data. What is a "blind zone" in this context? 7
(CO1)
(PO1)
- b) Define the following terms: 4 × 2
(CO1)
(PO1)
 - i. Database schema
 - ii. Relation instance
 - iii. Self-referencing
 - iv. Lossy decomposition
2. Consider the following relations (Here, underline means primary key, $fk<x>$ means foreign key referencing entity x):

Employee (EID, EName, Department, Position, JoinDate)

Project (PID, PName, StartDate, EndDate, Budget)

Assignment (EID{ $fk<Employee>$ }, PID{ $fk<Project>$ }, HoursWorked, AssignmentDate)

- a) Write SQL statements for the following queries: 5 × 2
(CO1)
(PO1)
 - i. Find the employees' names who have not been assigned to any projects in the year 2023.
 - ii. List the project IDs, project names, and the total number of assigned employees for each project in descending order of total number of assigned employees.
 - iii. List the names of employees whose average work hours are greater than the average work hours across all employees.
 - iv. List all employees' names, their departments, and the number of projects they are assigned for the year 2023. If they have not been assigned to any projects, display 0.
 - v. Find the name of the longest project in terms of duration.
- b) Create a view that will show all the project IDs, corresponding length of the projects, and budgets in the descending order of the budget. Can you issue any DML statement through this view? Justify. 5
(CO1)
(PO1)
- c) Assume that the current datatype of EID is integer. The database administrator intends to change the datatype to varchar2(12). Present the sequential operations required to execute this modification (You are not required to provide the exact queries). 10
(CO1)
(PO1)
3. Consider the following scenario:

Citizen Information: It stores information for all citizens of the country including ID, Name, Date of Birth (DOB), Blood Group (BG), and Profession. Professions are organized as 2-layer hierarchy to include both generic and specific name of the profession. For instance, a "Civil Engineer" is a specific name for the category of the profession "Engineer".

Driving License and Accidents: Each citizen may have at most one driving license having the basic information such as Name, DOB, Type of Vehicle (i.e. Light or Heavy), Issue Date, Expire Date. (Note: a certain level of redundancy is acceptable here.) Whenever any accident occurs, it is recorded with short description, date and time, ID of the driver, and place of the accident.

a) Draw the Entity Relationship Diagram (ER-D) using standard symbols.

15
(CO2)
(PO3)

b) Implement the ER-D using standard DDL statements.

15
(CO1)
(PO1)

c) Considering your design in Question 3.b), write a PL/SQL function using the following information:

10
(CO1)
(PO1)

Input	Driver ID (or similar)
Output	Status (i.e. Good, average, bad)
Algorithm	If the total number of accidents by this driver during the last 6 months exceeds 10 then his/her status is "Bad", if it is between 2 to 10 then it is "Average" otherwise it is "Good".

4. a) Given a relation $R(A, B, C, D, E, F)$ and the following functional dependencies (FDs):

$$BC \rightarrow ADEF, D \rightarrow B, A \rightarrow C, E \rightarrow F$$

15
(CO3)
(PO2)

Show the steps to determine all possible candidate keys for the above relation. Also, show if the given relation is in BCNF.

b) Given a relation $R(A, B, C, D, E, F)$ and the following functional dependencies (FDs):

$$AC \rightarrow BDEF, CD \rightarrow E$$

10
(CO3)
(PO2)

Show the steps to determine if the given relation is in 3NF. If the relation is not in 3NF, proceed to decompose it into 3NF.

5. Consider the following indexing for the database file in Figure 1 to answer the following questions:




Figure 1: Indexing for Question 5.

a) Create bitmap indices for the search-key "Designation" of the Database File.

10
(CO1)
(PO1)

10
(CO1)
(PO1)
5
(CO1)
(PO1)
20
(CO1)
(PO1)

- b) Describe the steps to execute the following index update operation for the specific example provided in Figure 1.
- i. Insert (11543, Mr.Moo, Physics, Prof.) into the Database file.
 - ii. Delete (11231, Ms.Lee , Botany, Asst.Prof.) from the Database file.
- c) Can sparse indexing be applied using the search key "Designation" while maintaining the same sequence of records in the Database file? Provide a brief explanation for your answer.
6. Consider the following B+ tree from Figure 2. Note that B+ tree diagram for this problem omits pointers for convenience.



Figure 2: B+ Tree of order $n = 4$ for Question 6.

Also, consider the following assumptions:

- A left pointer in an internal node guides towards keys less than ($<$) its corresponding key, while a right pointer guides towards keys greater than equal (\leq).
- A leaf node underflows when the number of keys goes below $\lceil \frac{d-1}{2} \rceil$.
- An internal node underflows when the number of pointers goes below $\lceil \frac{d}{2} \rceil$.

Perform the following operations sequentially and draw the tree after each operation.

- a) Insert 23.
- b) Insert 15.
- c) Update 40 to 16.
- d) Delete 5.