

Final Assessment Test – July 2023



VIT
Vellore Institute of Technology
(Approved by the University Grants Commission, New Delhi, India, 1963)

Course: BCSE302L - Database Systems
Class NBR(s): 0682/0684/0686/0697/0702/0706/0710/
0719/0723/0724/0726/0730/0735/0737/0741/0743/
0747/0751/0755/0758/0762/0767/0770/0774/0943/
1372

Slot: A1+TA1

Time: Three Hours

Max. Marks: 100

KEEPING MOBILE PHONE/SMART WATCH, EVEN IN 'OFF' POSITION IS TREATED AS EXAM MALPRACTICE

Answer ALL Questions

(10 X 10 = 100 Marks)

1. a) Consider a typical online course registration system such as Coursera or Udacity. How do you enforce the concept of data abstraction and data independence by elaborating each level of abstraction with your example? [5]
- b) Describe briefly the function of the catalog in a DBMS. Show clearly how the catalog allows the DBMS to provide data independence and why data modelling is important? [5]

2. Major airline companies that provide passenger services in Singapore are UniAir, TransAsia Airways, Far Eastern Transport, Great China Airlines, etc. Singapore Federal Flight Administration (SFFA) keeps a database with lots of information on all airlines. The information consists of: Each airplane has an identification number, name of the contact person, and telephone number. For each aircraft identification number, capacity and model are recorded.

Each employee has an employee identification number, name, address, birthday, sex, position with the company and qualification.

Each route has a route identification number, origin, destination, classification (into domestic or international route), distance of the route and price charged per passenger.

Each airline keeps the information about their buy/sell transactions (for example, selling an airplane ticket is a sell transaction, and paying for maintenance is a buy transaction). Each transaction has a transaction identification number, date, description, and amount of money paid/received.

Design a suitable ER diagram for the above scenario with all possible entities, attributes and the corresponding relationship with its participations.

3. a) A relation *IADDR* is defined with attributes *NAME* (unique), *STREET*, *CITY*, *STATE* and *POSTCODE* for storing the customer address details. The functional dependencies of the relation are given below: [5]

POSTCODE → CITY, STATE

STREET, CITY, STATE → POSTCODE

NAME → STREET, CITY, STATE

Check if the given relation is in BCNF? 3NF? 2NF?



- b) Consider the following relation scheme $R(A, B, C, D, E, F, G, H, I, J)$ together with the following functional dependencies: $\{F \rightarrow H, E\}$, $\{C, G, A\} \rightarrow \{B, D, J\}$, $\{B\} \rightarrow \{A, I\}$, $\{B, G\} \rightarrow \{C\}$. Use Armstrong's axioms to prove formally that this set of functional dependencies implies that $\{B, G\} \rightarrow \{D\}$. [5]

4. A relation R has attributes A, B, C, D, E , and satisfies the following FDs:

$A \rightarrow BC$

$B \rightarrow D$

$CD \rightarrow E$

$E \rightarrow A$

- a) Show that (A, B, C) and (A, D, E) are a lossless decomposition of R . [3]
 b) Show that (A, B, C) and (C, D, E) are not a lossless decomposition of R . [3]
 c) Show that (A, B, C) and (A, D, E) are not a dependency preserving decomposition of R . [4]

5. Consider the following tables. Write a SQL to find out the customer id of any customer who has ordered the product "Queso Cabrales" with quantity >6. Transform the SQL into a collection of relational algebra which denotes different execution plans of the SQL. Also, visualize the answer with the help of the Query tree.

Orders

Order ID	Customer ID
10263	Wilmk
10311	Dumon

Products

Product ID	Product Name
11	Queso Cabrales
42	Singaporena Fried
69	Gudbrandsdalsost
72	Mozzarella di Giovanni

Orders Details

OrderID	ProductID	UnitPrice	Quantity
10263	11	21.00	12
10263	42	14.00	10
10263	72	34.80	5
10311	42	14.00	6
10311	69	28.80	7

6. Consider a B+-tree of order three. Show the results of entering one by one the keys that are three letter strings: (era, ban, bat, kin, day, log, rye, max, won, ace, ado, bug, cop, gas, let, fax) (in that order) to an initially empty B+-tree. Assume that you use lexicographic ordering to compare the strings. Show the state of the tree after every 4 insertions.

- 7. Consider the following 3 schedules for 3 concurrent transactions T1, T2, T3

S1 = {r2(c), r2(b), w2(b), r3(b), r3(d), r3(c), r1(a), w1(a), w3(b), w3(c), r2(a), r2(d), w2(d), r1(b), w1(b), w2(a)}

S2 = {r3(b), r3(c), r3(d), r1(a), w1(a), w3(b), w3(c), r2(c), r1(b), w1(b), r2(b), w2(b), r2(a), w2(a), r2(d), w2(d)}

S3 = {r1(a), w1(a), r2(c), r2(b), w2(b), r2(d), r2(a), w2(a), w2(d), r1(b), w2(b), r3(b), r3(c), w3(b), w3(c), r3(d)}

For each of the three interleaved schedules, determine if the schedule is serializable. If so give an equivalent serial schedule.

- 8. Discuss the timestamp ordering protocol for concurrency control. How does strict timestamp ordering differ from basic timestamp ordering?

9. a) Describe the shadow paging recovery technique. [5]

- b) Prove that the wait-die and wound-wait protocols avoid deadlock and starvation. [5]

10. a) Consider an e-commerce website with a global presence. Some of the key areas for data management are maintaining user sessions, secured and reliable financial data management, shopping cart management, recommendation of products, maintaining product catalog, generating reports, performing analytics, and logging user activities. For the above scenario, choose the appropriate database and justify in terms of CAP theorem and BASE properties. [5]

b) Which of the four NoSQL database types is more suitable for storing session details? Justify by explaining the features of the database chosen. [5]

⇔⇔⇔