

**VIT**

SCHOOL OF COMPUTER SCIENCE AND ENGINEERING
CONTINUOUS ASSESSMENT TEST - II
FALL SEMESTER 2025-2026

REG.NO.: 2UMIL0128

SLOT: B2

Programme Name & Branch : M.Tech (Integrated) Computer Science and Engineering
Course Code and Course Name : CSI1001 - Principles of Database Systems
Faculty Name(s) : Prof. Sudhakar. P, Prof. Harini. N, Prof. Divya. M
Class Number(s) : 1638, 1641, 1635
Date of Examination : 06-Oct-2025 09:30 AM - 11:00 AM
Exam Duration : 90 minutes
Maximum Marks: 50

General instruction(s):

- Answer All Questions
- M - Max mark; CO - Course Outcome; BL - Blooms Taxonomy Level (1 - Remember, 2 - Understand, 3 - Apply, 4 - Analyse, 5 - Evaluate, 6 - Create)
- Course Outcomes:
CO3 - Analyze and improve a database design by normalization.
CO5 - Analyze the basics of query evaluation and heuristic query optimization techniques.

Q. No	Question	Module	Marks	CO	BL																												
1.	Consider the following Football relation Football = (Match, Player, Ground, Schedule, Goal) and the set of functional dependencies: Match → Player, Ground Ground, Schedule → Goal Player → Schedule Goal → Match Question: Compute the candidate keys of relation "Football". Justify your answer step-by-step with appropriate explanation.	3	10	3	3																												
2.	Consider the following table <table><tr><th>VisitNo</th><th>PatientID</th><th>PatientName</th><th>DoctorID</th><th>DoctorName</th><th>Department</th><th>AppointmentDate</th></tr><tr><td>V001</td><td>P101</td><td>Alex</td><td>D01</td><td>Dr. Prasath</td><td>Cardiology</td><td>01-10-2024</td></tr><tr><td>V002</td><td>P101</td><td>Alex</td><td>D02</td><td>Dr. Kiran</td><td>Neurology</td><td>15-10-2024</td></tr><tr><td>V003</td><td>P102</td><td>Mike</td><td>D01</td><td>Dr. Prasath</td><td>Cardiology</td><td>02-10-2024</td></tr></table> Question: Normalise the given table till BCNF. Discuss your answer step-by-step at each normal form.	VisitNo	PatientID	PatientName	DoctorID	DoctorName	Department	AppointmentDate	V001	P101	Alex	D01	Dr. Prasath	Cardiology	01-10-2024	V002	P101	Alex	D02	Dr. Kiran	Neurology	15-10-2024	V003	P102	Mike	D01	Dr. Prasath	Cardiology	02-10-2024	3	10	3	3
VisitNo	PatientID	PatientName	DoctorID	DoctorName	Department	AppointmentDate																											
V001	P101	Alex	D01	Dr. Prasath	Cardiology	01-10-2024																											
V002	P101	Alex	D02	Dr. Kiran	Neurology	15-10-2024																											
V003	P102	Mike	D01	Dr. Prasath	Cardiology	02-10-2024																											
3.	Consider the following relations in the Healthcare scenario: Patients(PatientID, Name, Age, Gender) Visits(VisitID, PatientID, DoctorID, VisitDate) Doctors(DoctorID, Name, Specialization)	4	10	5	4																												

**VIT**

Vellore Institute of Technology

SCHOOL OF COMPUTER SCIENCE AND ENGINEERING
CONTINUOUS ASSESSMENT TEST - II
FALL SEMESTER 2025-2026

REG.NO.:

SLOT: B2

To retrieve the patient name along with doctor name who have consulted, the following query was executed. <pre>SELECT p.Name, d.Name FROM Patients p JOIN Visits v ON p.PatientID = v.PatientID JOIN Doctors d ON v.DoctorID = d.DoctorID WHERE d.Specialization = 'Cardiology';</pre> Patient, Visits and Doctors table has 10,000, 50,000 and 500 records, respectively. A block can hold 100, 100, 50 records from Patient, Visits and Doctors table, respectively. a) Compute the query cost involved to execute the given query (with out indexing). Explain your answer with the query tree b) How the query cost change if Doctor table is indexed based on specialization, Visit table is indexed based on DoctorID and Patients table is indexed based on PatientID. The join applied is 'Block Nested'. Compare the query cost with and without indexing.												
4. Consider the following transactions <table border="1" data-bbox="303 878 1125 1102"> <thead> <tr> <th>T1</th> <th>T2</th> </tr> </thead> <tbody> <tr> <td>Read (Customer_A)</td> <td>Read (Customer_A)</td> </tr> <tr> <td>Write (Customer_A)</td> <td>Write (Customer_A)</td> </tr> <tr> <td>Read (Customer_B)</td> <td>Write (Customer_B)</td> </tr> </tbody> </table> Note: Customer_A and Customer_B is a specific record in the relation. a) Check whether any conflicts exist in the given transactions ? If no, Why ? If yes, list the conflicts. b) Check the conflict serializability with the precedence graph for the given scenario and explain the solution to overcome, if conflict exists.	T1	T2	Read (Customer _A)	Read (Customer _A)	Write (Customer _A)	Write (Customer _A)	Read (Customer _B)	Write (Customer _B)	4	10	5	3
T1	T2											
Read (Customer _A)	Read (Customer _A)											
Write (Customer _A)	Write (Customer _A)											
Read (Customer _B)	Write (Customer _B)											
5. Consider a transaction is executing following SQL statements <pre>START TRANSACTION; UPDATE InsuranceAccount SET Balance = Balance - 500 WHERE PatientID = 101; UPDATE PatientBills SET Status = 'Paid' WHERE PatientID=101 and BillID = 'B202309'; COMMIT;</pre> How ACID properties ensure that the above transaction execution leads to consistent state? Explain each property with the given transaction.	4	10	5	3								