(v) A relation R(X, Y, Z), holds F= { XY$\rightarrow$Z, Z$\rightarrow$Y}. The candidate keys will be
(a) {XY} only (b) {XY} and {XZ}
(c) {XY}, {XZ} and {YZ} (d) X only.

(vi) At allocation schema the distribution transparency that exists is
(a) local mapping transparency
(b) fragmentation transparency
(c) location transparency
(d) none of these.

(vii) If every non-key attribute is functionally dependent on the primary key, the relation will be in
(a) First Normal Form (b) Second Normal Form
(c) Third Normal Form (d) Fourth Normal Form.

(viii) A lock that allows concurrent transactions to access different rows of the same table is known as a
(a) field-level lock (b) row-level lock
(c) table-level lock (d) database-level-lock.

(ix) TCL statements are
(a) grant and revoke
(b) commit and rollback
(c) commit, rollback and savepoint
(d) none of these.

(x) In distributed database, a minterm fragment is designed from
(a) simple predicates (b) minterm predicates
(c) minterm projection (d) none of these.

Group – B

2. (a) A relation R (X, P , C , T , G) with attributes is given.
The given set of FD's are
X$\rightarrow$ P
C$\rightarrow$ T
(X, P) $\rightarrow$ G
(i) Find out the candidate keys from the set of given FD's.

A single update to read cost ratio is 2.
If we take locality of reference as the objective of allocation, which of the above two design of allocation schema is better. Justify your answer

(b) Draw and explain the Architecture of distributed DBMS along with the distribution transparency at each level.

6 + 6= 12

7. (a) Consider relation R with attributes (id, dept_id, salary, post, proj). This is maintained in a distributed Database. Only 2 Departments are there T1 and T2.
Consider 3 applications. Application 1 manages the project information of employees of department T1. Application 2 deals with the salary and calculation of incentives of department T2, and application 3 maintains the average salary of the engineers of any department. Design the simple predicates , minterm predicates and the minterm fragments .
Consider the following distributed wait-for-graph (DWFG):

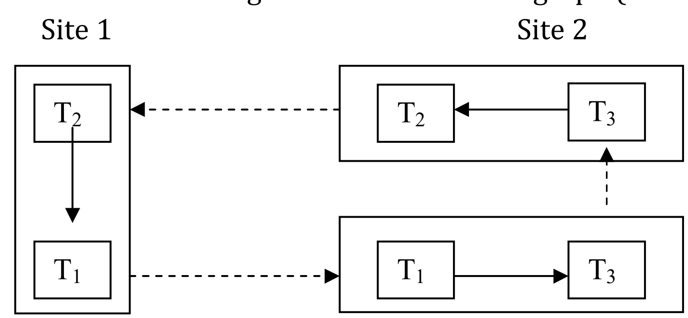

(b) Illustrate each step of detecting deadlock using distributed Deadlock Detection algorithm for the above graph. What is false deadlock?

6 + (5+1) = 12

Group – E

8. (a) Briefly analyze all cases where 2 phase commitment protocol deals with different kind of site failures, ensuring the consistency of database?

(b) Given the relational schema: **CLASSMATES** (RollNo, Name, RegistrationNo, Address, PhoneNo), **COURSES** (RollNo, CourseCode, CourseName), **TEACHERS** (CourseCode, TeacherCode, TeacherName)

Following are the horizontal fragments (F_i) created from the above relations:-

F_1: $\sigma_{\text{CourseCode = 'CSEM202'}}$(Courses)

F_2: $\sigma_{\text{CourseCode <> 'CSEM202'}}$(Courses)

F_3: $\sigma_{\text{CourseCode = 'CSEM202'} \wedge \text{CourseCode = 'CSEM201'}}$(Courses)

F_4: $\sigma_{\text{CourseCode IN ('CSEM401', 'CSEM410')}}$(Courses)

F_5: $\sigma_{\text{Name = 'Vikas'}}$(Classmates)

F_6: $\sigma_{\text{TeacherName = 'Rajeev'}}$(Teachers)

F_7: $\sigma_{\text{TeacherName <> 'Rajeev'}}$(Teachers)

Answer the following questions using reduction techniques involving Classmates, Courses and Teachers (take whole relations) and the fragments (F_1 to F_7). Queries should be written in both Relational Algebra and Structured Query Language. In addition, draw reduced tree for each query. Justify that the disjointness and completeness of the fragmented queries are preserved for the result that you have deduced.

i) What are the names of the classmates enrolled in the course CSEM202?

ii) Which course classmate Vikas is taking?

iii) What are the names of the classmates enrolled in both CSEM202 and CSEM201?

3 + (3x3) = 12

9. (a) Discuss the concept of checkpoints and cold restart with respect to distributed database.

(b) We have two very large tables R(A;B) and S(A;C) where all attributes are integers. The data in each table is randomly distributed across three servers N1, N2, and N3. Explain how a parallel DBMS can compute R 1R:A=S:A S in parallel using all three servers.

6 + 6 = 12

M.TECH/IT/2ND SEM /INFO 5201/2016
2016

ADVANCED DBMS
(INFO 5201)

Time Allotted : 3 hrs **Full Marks : 70**

Figures out of the right margin indicate full marks.

Candidates are required to answer Group A and <u>any 5 (five)</u> from Group B to E, taking <u>at least one</u> from each group.

Candidates are required to give answer in their own words as far as practicable.

Group – A
(Multiple Choice Type Questions)

1. Choose the correct alternatives for the following: **10 x 1=10**

(i) A relation R(A,B,C,D) is decomposed into R1(A,B) and R2(A,C,D), where A is the candidate key of R1. So, the decomposition is
(a) lossy (b) is lossless (c) Both a and b (d) None of these.

(ii) Assume transaction A holds a shared lock R. If transaction B also requests for a shared lock on R, then
(a) it'll result in deadlock situation
(b) it'll be rejected
(c) it'll be rejected
(d) it'll be granted when released by A.

(iii) In a relation R(A,B,C), if an MVD, A->> B|C holds, then
(a)B and C are dependent attributes
(b) B and C are independent of each other
(c) B determines C
(d) A, B and C are independent of each other.

(iv) Vertical fragmentation of global relation is subdivision of
(a) its attributes (b) its rows
(c) its records (d) its relation.

(ii) Find the canonical cover of F.

(iii) Find out in which normal form the relation is? Explain.

(iv) Convert the relation into its higher normal form such that dependency is preserved and lossless decomposition occurs. Explain.

(b) Explain why 4NF is a normal form more desirable than BCNF. Use Armstrong's axioms to prove the soundness of the decomposition rule.

(1½ x 4=6) + (3+3) = 12

3. (a) Distinguish between object oriented DBMS and RDBMS. Discuss the advantage of distributed database over centralized.

Consider a relational database as given below:

Train (<u>train-no</u>, train_name, start_station)

Coach(<u>coach-no, train-no</u>, type, price), where the underlined attributes are the primary keys. Write down the expressions in relational algebra and SQL for the following queries:

(i) List all the train names starting from station XYZ.

(ii) List the train number and price of all "3 A/C" (type) coaches with price below Rs.3025/-

(iii) List the price and type of all coaches of "Rajdhani" train.

(3+3) + (3×2) = 12

Group – C

4. (a) Consider the following two transactions :

T1 : read(A);
Read(B);
If A = 0 then B:= B+1
Write(B).

T2 : read(B);
Read(A);
If B=0 then A:=A+1;
Write(A).

Add lock and unlock instructions to transactions T1 and T2, so that they observe the two-phase locking protocol. Can the execution of these transactions result in a deadlock?

M.TECH/IT/2ND SEM /INFO 5201/2016

(ii) Find the canonical cover of F.

(iii) Find out in which normal form the relation is? Explain.

(iv) Convert the relation into its higher normal form such that dependency is preserved and lossless decomposition occurs. Explain.

(b) Explain why 4NF is a normal form more desirable than BCNF. Use Armstrong's axioms to prove the soundness of the decomposition rule.

(1½ x 4=6) + (3+3) = 12

3. (a) Distinguish between object oriented DBMS and RDBMS. Discuss the advantage of distributed database over centralized.

Consider a relational database as given below:

Train (train-no, train_name, start_station)

Coach(coach-no, train-no, type, price), where the underlined attributes are the primary keys. Write down the expressions in relational algebra and SQL for the following queries:

(i) List all the train names starting from station XYZ.

(ii) List the train number and price of all "3 A/C" (type) coaches with price below Rs.3025/-

(iii) List the price and type of all coaches of "Rajdhani" train.

(3+3) + (3×2) = 12

Group – C

4. (a) Consider the following two transactions :

T1 : read(A);
Read(B);
If A = 0 then B:= B+1
Write(B).

T2 : read(B);
Read(A);
If B=0 then A:=A+1;
Write(A).

Add lock and unlock instructions to transactions T1 and T2, so that they observe the two-phase locking protocol. Can the execution of these transactions result in a deadlock?

INFO 5201 3

(b) Explain the algorithm of 2PC protocol in distributed environment for both
Coordinator and Participants.

6 + 6 = 12

5. (a) Consider a distributed DBMS with two sites running transactions T1 and T2 with following data objects: A,B,C,D. Show whether the following execution sequences running at site 1 and site2 are conflict serializable or not. If the answer is affirmative, determine all possible total orders of transactions. If answer is negative, prove that there is no total order possible.

Execution 1:
Site 1: R1(A) R2(A) W1(A) W2(B)
Site 2: R1(C) R2(C) W2(C) W1(D)
Execution 2:
Site1: R1(B) R2(A) W2(A) W1(C)
Site2: W1(C) R1(D) R2(D) W1(D)

(b) Distinguish between the two classical examples of concurrency – the Uncommitted Dependency Problem and the Inconsistent Analysis Problem with example. How can Wait-for-Graph (WFG) be used in the scenario of deadlock detection and recovery.

6 + 6 = 12

Group – D

6. (a) Consider the following two allocation level design of fragments R1, R2 and R3

i) Allocation Design 1: R1 at site 1; R2 at site 2; R3 at site 3;

ii) Allocation Design 2 : R1 and R2 at site 1 ; R2 and R3 at site 3.
With the following applications (all with same frequency of execution).
A1, issued at site 1, reads 6 records of R1 and update 7 records of R3.
A2, issued at site 3, update 5 records of R3 and update 5 record of R2
A3, issued at site 2, reads 10 records of R1, update 4 records of R2.

M.TECH/IT/2ND SEM /INFO 5201/2016

(b) Explain the algorithm of 2PC protocol in distributed environment for both
Coordinator and Participants.

6 + 6 = 12

5. (a) Consider a distributed DBMS with two sites running transactions T1 and T2 with following data objects: A,B,C,D. Show whether the following execution sequences running at site 1 and site2 are conflict serializable or not. If the answer is affirmative, determine all possible total orders of transactions. If answer is negative, prove that there is no total order possible.

Execution 1:
Site 1: R1(A) R2(A) W1(A) W2(B)
Site 2: R1(C) R2(C) W2(C) W1(D)
Execution 2:
Site1: R1(B) R2(A) W2(A) W1(C)
Site2: W1(C) R1(D) R2(D) W1(D)

(b) Distinguish between the two classical examples of concurrency – the Uncommitted Dependency Problem and the Inconsistent Analysis Problem with example. How can Wait-for-Graph (WFG) be used in the scenario of deadlock detection and recovery.

6 + 6 = 12

Group – D

6. (a) Consider the following two allocation level design of fragments R1, R2 and R3

i) Allocation Design 1: R1 at site 1; R2 at site 2; R3 at site 3;

ii) Allocation Design 2 : R1 and R2 at site 1 ; R2 and R3 at site 3.
With the following applications (all with same frequency of execution).
A1, issued at site 1, reads 6 records of R1 and update 7 records of R3.
A2, issued at site 3, update 5 records of R3 and update 5 record of R2
A3, issued at site 2, reads 10 records of R1, update 4 records of R2.

INFO 5201 4