

IV B.Tech II Semester Regular Examinations, Apr/May 2006

ADVANCED DATABASES

(Computer Science & Systems Engineering)

Time: 3 hours

Max Marks: 80

Answer any FIVE Questions
All Questions carry equal marks

1. Describe the main motivations for the development of distributed databases. [16]
2. Consider the global relations: PATIENT (NUMBER, NAME, SSN, AMOUNT-DUE, DEPT, DOCTOR, MED-TREATMENT) DEPARTMENT (DEPT, LOCATION, DIRECTOR) STAFF(STAFFNUM, DIRECTOR, TASK) Define their fragmentation as follows
 - (a) DEPARTMENT has a horizontal fragmentation by LOCATION, with two locations: Each department is conducted by one DIRECTOR.
 - (b) There are several staff members for each department, led by the departments director, STAFF has a horizontal fragmentation derived from that of DEPARTMENT and a semi-join on the DIRECTOR attribute. Which assumption is required in order to assure completeness and disjointness? Give also the reconstruction of global relations from fragments. [8+8]
3. Describe vertical fragmentation with respect to the design of database fragmentation. [16]
4. What is an operator tree of query? Draw operator trees for the following queries
 - (a) $PJ_{SNUM}SL_{AREA="North"}(SUPPLYJN_{DEPTNUM=DEPTNUM}DEPT)$
 - (b) $PJ_{EMP.NAME}(EMPJN_{DEPTNUM} = SL_{MGRNUM=373}DEPT)$ [4+5+6]
5.
 - (a) Explain the use of semi-join programs for join queries.
 - (b) Explain the rationale of semi-join reduction in distributed data bases. [8+8]
6. Explain the recovery of distributed transactions and compare with recovery in Centralized systems. [16]
7.
 - (a) Give an example of a distributed DBMS with three sites such that no two local waits-for graphs reveal a deadlock, yet there is a global deadlock.
 - (b) Suppose that a site does not get any response from another site for a long time. Can the first site tell whether the connecting link has failed or the Other site has failed? How is such a failure handled?
 - (c) Explain the need for a commit protocol in a distributed DBMS. [6+6+4]
8.
 - (a) Explain the operation of three-phase commitment protocol with a neat sketch.
 - (b) Explain the operation of quorum based three-phase commitment protocol with a neat sketch. [8+8]

IV B.Tech II Semester Regular Examinations, Apr/May 2006

ADVANCED DATABASES

(Computer Science & Systems Engineering)

Time: 3 hours

Max Marks: 80

Answer any FIVE Questions
All Questions carry equal marks

1. What is the basic function of a communication network? Explain the parameters that characterize this function. [16]
2. Discuss the various distributed database access primitives. [16]
3. (a) What are the various objectives of Data Distribution.
(b) Compare Top-Down and Bottom-Up approaches in the design of Data Distribution. [8+8]
4. Draw the operator tree for the following query and obtain the Canonical form of the same. $PJ_{SNUM}SL_{AREA="Noth"}(SUPPLYJN_{DEPTNUM=DETPNUM}DEPT)$ [16]
5. (a) Explain the methods for Optimization of general queries.
(b) Consider the join $RJNa=b S$; assume that R and S are at different sites, and disregard the cost of collecting the result of the join. Let $Co=0$ and $C1=1$. The following profiles are given:

 $Size(R) = 50; card(R) = 100; val(A[R]) = 50; size(A) = 3$
 $Size(S) = 5; card(S) = 50; val(B[S]) = 50; size(B) = 3$

 $RSJa=b S$ has selectivity $p=0.2$
 $SSJ b=a R$ has selectivity $p=0.8$

 Give the transmission cost of
 - i. Performing the join at the site of R using semi-join reduction
 - ii. Performing the join at the site of S using semi-join reduction
 - iii. Performing the join at the site of R without semi-join reduction
 - iv. Performing the join at the site of S without semi-join reduction Which is the best solution? Explain. [8+6]
6. Classify the failures in a centralized database. Explain the recovery procedures. [6+10]
7. (a) What is meant by serializability.
(b) Write a note on serializability in centralized database systems. [6+10]
8. (a) Explain the operation of two-phase commitment protocol at the time when all sites are inactive.

(b) Explain the features of primary copy locking approach. [10+6]

IV B.Tech II Semester Regular Examinations, Apr/May 2006

ADVANCED DATABASES

(Computer Science & Systems Engineering)

Time: 3 hours

Max Marks: 80

Answer any FIVE Questions
All Questions carry equal marks

1. List out the components of a Commercial DDBMS and mention the various services supported by the same. [16]
2. Explain derived horizontal fragmentation and mixed fragmentation. [16]
3. Explain how costs and benefits are evaluated while allocating the fragments of global relation. [16]
4. What is qualified relation? State the seven rules that define the result of applying the operations of relational algebra to qualified relations. [16]
5. Describe in detail how query processing carried out in SDD-1. [16]
6. What is meant by unilateral abort capability in 2-phase commit protocol? Identify the situations when a process is blocked in two phase commit protocol. [16]
7. Explain the following:
 - (a) Serializability in Centralized Database System.
 - (b) Serializability in Distributed Database System. [8+8]
8. (a) Explain the operation of non-blocking commitment protocol with a neat sketch.
(b) Explain the operation of non-blocking commitment protocol with site failures. [16]

IV B.Tech II Semester Regular Examinations, Apr/May 2006

ADVANCED DATABASES

(Computer Science & Systems Engineering)

Time: 3 hours

Max Marks: 80

**Answer any FIVE Questions
All Questions carry equal marks**

1. Give an overview of distributed databases. [16]
2. Discuss a read-only application at the following levels of distribution transparency.
 - (a) Fragmentation transparency.
 - (b) Location transparency.
 - (c) Local mapping transparency. [5+5+6]
3. Explain about primary fragmentation and derived horizontal fragmentation. [16]
4. What is operator graph Illustrate with an example how common sub expressions are determined. [16]
5. Explain the simpler representation of queries and a model for query Optimization. [16]
6. Define the terms: atomicity, consistency, isolation, durability, schedule, blind write, dirty read, unrepeatable read, recoverable schedule, avoids-cascading aborts schedule. [16]
7.
 - (a) Explain the conservative time stamp method.
 - (b) Explain the primitive and non-primitive methods for deadlock prevention. [16]
8.
 - (a) Describe with examples, the types of problems that can occur in distributed transactions environment when concurrent access to the database is allowed.
 - (b) Give the details of a mechanism for concurrency control that can be used to ensure that the types of problems discussed above cannot occur. [8+8]
