Q.1

a) Give examples of systems that use traditional file processing instead of a database approach.

Answer:

1. Small internal utility to locate files
2. Small single user application that does not require security (such as a customized calculator or a personal address and phone book)
3. Real-time navigation system (with heavy computation and very little data)
4. The students may think of others.

b) Explain heuristic based query optimization.

Answer: Page Number 547-548 of A. Silberchatz, H.F. Korth & S Sudershan, Database System Concept, McGraw Hill, 4th Edition, 2005

c) Given below are two sets of FDs for a relation R (A, B, C, D, E). Are they equivalent?

1. A → B, AB → C, D → AC, D → E

2. A → BC, D → AE

Answer:

Let us number the first set of FDs as:

1. A → B,
2. AB → C,
3. D → AC,
4. D → E

3 can be replaced by D → A and D → C.

From 1 and 2 together imply that 2 can be replaced by

A → B; AB → C = A → C.

Now we have D → A, A → C and so D → C is implied (by transitively)

The first set of FDs is thus equivalent to the following irreducible set.

The second given set of FD's A → BC and D → AE is clearly also equivalent to this irreducible set.

Thus, the two given sets are equivalent.

d) Explain various components of distributed databases.

Answer: Page Number 709-712 of A. Silberchatz, H.F. Korth & S Sudershan, Database System Concept, McGraw Hill, 4th Edition, 2005

e) Mention naming conventions to identify the foreign keys efficiently.

Answer:

The naming convention can be informally specified as FK followed by the referenced relationship or relation name, and optionally followed by the primary key of the referenced table. This foreign key naming convention is frequently used in large databases. Sometimes this produces very long column names but it pays off by facilitating readability for database professionals by using more descriptive names.

f) Explain techniques other than chaining to handle bucket overflow in external hashing?

Answer:

One can use techniques for handling collisions similar to those used for internal hashing. For example, if a bucket is full, the record which should be inserted in that bucket may be placed in the next bucket if there is space (open addressing). Another scheme is to use a whole overflow block for each bucket that is full. However, chaining seems to be the most appropriate technique for static external hashing.

g) Prove that cautions waiting avoids deadlock.

Answer:

In cautious waiting, a transaction Ti can wait on a transaction Tj (and hence Ti becomes blocked) only if Tj is not blocked at that time, say time b(Ti), when Ti waits. Later, at some time b(Tj) > b(Ti), Tj can be blocked and wait on another transaction Tk only if Tk is not blocked at that time. However, Tj cannot be blocked by waiting on an already blocked transaction since this is not allowed by the protocol. Hence, the wait-for graph among the blocked transactions in this system will follow the blocking times and will never have a cycle, and so deadlock cannot occur.

Q.2

a) A bank has many branches, and a large number of customers. A customer can open many different kinds of accounts with the bank. The bank keeps track of the customer with his SSN, name, address, and phone number. Age is a factor to check whether he is a major. There are different types of loans, each identified by a loan number. Customer can take out more than one type of loan, and all branches can give loans. Loans have a duration and interest rate. The account holder can enquire about the balance in his account.

Draw an ER Diagram for the bank. Make suitable assumptions and use them in showing maximum and minimum cardinality ratios.

Answer:

Solution:

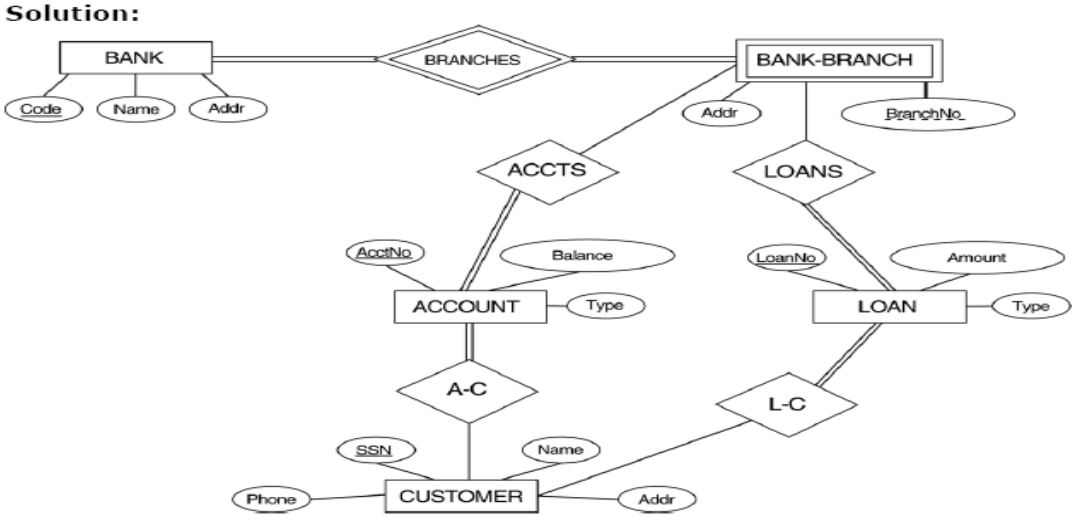

b) Explain domain constraints and key constraints in relational data model.
Answer: Page Number 153-158 of Text Book I

c) Give an example to illustrate the mapping of ER diagrams to tables.
Answer: Page Number 78 of Text Book I

Q.3
a) Define Fourth Normal Form. Consider a Relational Schema R=(A,B,C,D,E). Let M be the following set of Multi-Valued Dependencies:-
M= (A→→BC, B→→CD, E→→AD)
Give a lossless join decomposition of R into Fourth Normal Form. Justify your answer.
Answer:
Since A→→BC holds by complementation Rule A→→(R-A-BC) So by Fagin's Theorem R1(A,B,c) and R2(A,D,E) will be lossless decomposition of R.
Considering B→→CD we can prove that R1(B,C,D) and R2(B,A,E) will also be lossless decomposition of R.
Similarly E→→AD we can prove R1(E,A,D) and R2(E,B,C) will also be lossless decomposition of R.

b) How can the key and foreign key constraints be enforced by the DBMS?
Answer:
(i) Since V→WX, therefore V is the candidate key of R1
now R1∩R2={V} which is the candidate key of R1
therefore it is a loss less join decomposition of R.

(ii) R1∩R2={X} which is not the candidate key of R1 or R2. Therefore it is NOT a loss join decomposition of R

c) What is the difference between DROP and DELETE command?
Answer:
For each operation, we give the tuple calculus expression followed by the domain calculus expression.

i. {t.A, t.B, t.C, q.D, q.E | R(t) AND S(q) AND t.C=q.C }, { xyzvw | R(xyz) AND (EXISTS u) (S(uvw) AND z=u) }

ii. { t | R(t) OR S(t) },{ xyz | R(xyz) OR S(xyz) }
iii. { t.A, t.B, t.C, q.D, q.E, q.F | R(t) AND S(q) },(xyzuvw | R(xyz) AND S(uvw) }

Q.4
a) Consider the following tables
Works (Pname, Cname, Salary)
Lives (Pname, Street, City)
Locatedin (Cname, City)
Manager (Pname, Mgrname)

Where Pname = Person name, Cname = Company name, Mgrname = Manager name

Write the SQL for the following:

(i) List the names of the people who work for company Wipro along with the cities they live in.

(ii) Find the people who work for the company "Infosys" with salary more than Rs. 50000/-. List the names of the people, along with the streets and city addresses.

(iii) Find the names of the persons who live and work in the same city.

(iv) Find the names of the person who do not work for "Infosys".

(v) Find the persons whose salaries are more than that of all of the 'Oracle' employee.

(vi) Find the names of the companies that are located in every city where the company 'Infosys' is located.

Answer:

1. SELECT L.PName, City
   FROM Works W, Lives L
   Where CName = "Wipro" and W. PName = L.Pname;
2. SELECT L.PName, City, Street
   FROM WORKS W, Lives L
   Where W. CName = 'Infosys'
   AND W.PName = L. CName AND Salary > 50000;
3. SELECT W.PName
   FROM Works W, Lives L, LocatedIn I
   Where W.CName = I.CName AND W.PName = L.PName
   AND I.City = L.City
4. SELECT PName
   FROM Works
   Where CName ! = 'Infoysy';
5. SELECT PName
   FROM Works
   Where Salary > ALL
   (SELECT Salary
   FROM Works
   WHERE CNAME = "Oracle");
6. SELECT CName
   FROM LocatedIn I1
   Where NOT EXISTS
   (SELECT *
   FROM LocatedIn I2
   WHERE I2.CName = "Infosys" AND I2.City NOT IN
   SELECT City
   FROM LocatedIn I3
   WHERE i3.CName – I1.CName));

b. How can the key and foreign key constraints be enforced by the DBMS?

Answer:
One possible technique that is often used to check efficiently for the key constraint is to create an index on the combination of attributes that form each key (primary or secondary). Before inserting a new record (tuple), each index is searched to check that no value currently exists in the index that matches the key value in the new record. If this is the case, the record is inserted successfully.
For checking the foreign key constraint, an index on the primary key of each referenced relation will make this check relatively efficient. Whenever a new record is inserted in a referencing relation , its foreign key value is used to search the index for the primary key of the referenced relation, and if the referenced record exists, then the new record can be successfully inserted in the referencing relation.
For deletion of a referenced record, it is useful to have an index on the foreign key of each referencing relation so as to be able to determine efficiently whether any records reference the record being deleted.
If the indexes described above do not exist, and no alternative access structure (for example, hashing) is used in their place, then it is necessary to do linear searches to check for any of the above constraints, making the checks quite inefficient.

c. What is the difference between DROP and DELETE command?
Answer:
The drop command is used to drop named schema element such as tables, domains or constraints. One can also drop a schema. The delete command removes tuples from a relation. It includes a where clause to select the tuples to be deleted.

Q.5
a) Explain any two techniques used in query optimization for translations the SQL queries into relational algebra.
Answer: Page Number 553-566 of Text Book I

b) Explain the features of web databases.
Answer: Page Number 853-855, 45-49 of Text Book I

c) Explain various types of recovery techniques used in database systems.
Answer: Page Number 672-680 of Text Book I

Q.6
a) What do you mean by concurrency control Manager? What is the goal of Concurrency control schemes?
Answer:
Concurrency control is one of the major functions of DBMS, which is entrusted to its component called concurrency control manager (CCM). Concurrency control implies controlling the execution of concurrent transactions.
The goal of Concurrency control schemes is that the generated schedules are serializable and cascadeless, while at the same time ensuring as high degree of concurrency as possible.

b) Differentiate between the Shared and Exclusive locking mode with example.
Answer:
A shared mode lock is requested by a transaction T_i on the data item Q, when T_i needs to access the data item Q only in "Read" mode (not write mode).
An exclusive mode lock is requested by a transaction T_i on a data item Q when T_i needs to access the data item Q not only in "read" mode but also in "write".

c) Explain Two-phase locking protocol with example.
Answer:
In databases and transaction processing, two-phase locking (2PL) is a concurrency control method that guarantees serializability. It is also the name of the resulting set of database transaction schedules (histories). The protocol utilizes locks, applied by a transaction to data, which may block (interpreted as signals to stop) other transactions from accessing the same data during the transaction's life.
By the 2PL protocol locks are applied and removed in two phases:
1. Expanding phase: locks are acquired and no locks are released.
2. Shrinking phase: locks are released and no locks are acquired.

Q.7 Write short notes for any THREE of the following:
(i) Join Operation
(ii) BCNF
(iii) Tuple versus Domain relational calculus
(iv) Weak Entity

Answer:
(i) Page Number 183-188 of Text Book I
(ii) Page Number 370-371 of Text Book I
(iii) Page Number 201-212 of Text Book I
(ii) Page Number 76-77 of Text Book I

Text Books

1. R.Elmasri and S. Navathe, "Fundamental of Database Systems", Addison Wesley, 5th Edition, 2007

2. R Ramakrishnan & J Gehrke, Database Management Systems, McGraw Hill, Third Edition, 2002