

CS3492-Database-Management-Systems question bank

REGULATION 2021

OBJECTIVES

- To learn the fundamentals of data models, relational algebra and SQL
- To represent a database system using ER diagrams and to learn normalization techniques
- To understand the fundamental concepts of transaction, concurrency and recovery processing
- To understand the internal storage structures using different file and indexing techniques which will help in physical DB design
- To have an introductory knowledge about the Distributed databases, NOSQL and database security

UNIT 1 RELATIONAL DATABASE 10

Purpose of Database System – Views of data – Data Models – Database System Architecture – Introduction to relational databases – Relational Model – Keys – Relational Algebra – SQL fundamentals Advanced SQL features – Embedded SQL– Dynamic SQL

UNIT 2 DATABASE DESIGN 8

Entity-Relationship model – E-R Diagrams – Enhanced-ER Model – ER-to-Relational Mapping – Functional Dependencies – Non-loss Decomposition – First, Second, Third Normal Forms, Dependency Preservation – Boyce/Codd Normal Form – Multi-valued Dependencies and Fourth Normal Form – Join Dependencies and Fifth Normal Form

UNIT 3 TRANSACTIONS

Transaction Concepts – ACID Properties – Schedules – Serializability – Transaction support in SQL – Need for Concurrency – Concurrency control – Two Phase Locking- Timestamp – Multiversion – Validation and Snapshot isolation– Multiple Granularity locking – Deadlock Handling – Recovery Concepts – Recovery based on deferred and immediate update – Shadow paging – ARIES Algorithm

UNIT 4 IMPLEMENTATION TECHNIQUES 9

RAID – File Organization – Organization of Records in Files – Data dictionary Storage – Column Oriented Storage– Indexing and Hashing –Ordered Indices – B+ tree Index Files – B tree Index Files – Static Hashing – Dynamic Hashing – Query Processing Overview – Algorithms for Selection, Sorting and join operations – Query optimization using Heuristics - Cost Estimation.

UNIT 5 ADVANCED TOPICS 9

Distributed Databases: Architecture, Data Storage, Transaction Processing, Query processing and optimization – NOSQL Databases: Introduction – CAP Theorem – Document Based systems – Key value Stores – Column Based Systems – Graph Databases. Database Security: Security issues – Access control based on privileges – Role Based access control – SQL Injection – Statistical Databasesecurity – Flow control – Encryption and Public Key infrastructures – Challenges

TOTAL: 45 PERIODS

OUTCOMES:**Upon completion of the course, the students will be able to:**

- Construct SQL Queries using relational algebra
- Design database using ER model and normalize the database
- Construct queries to handle transaction processing and maintain consistency of the database
- Compare and contrast various indexing strategies and apply the knowledge to tune the performance of the database
- Appraise how advanced databases differ from Relational Databases and find a suitable database for the given requirement.

TEXT BOOKS:

1. Abraham Silberschatz, Henry F. Korth, S. Sudharshan, —Database System Concepts, Seventh Edition, McGraw Hill, 2020.
2. Ramez Elmasri, Shamkant B. Navathe, —Fundamentals of Database Systems, Seventh Edition, Pearson Education, 2017

REFERENCES:

1. C.J.Date, A.Kannan, S.Swamynathan, —An Introduction to Database Systems, Eighth Edition, Pearson Education, 2006.

COURSE OUTCOME

At the end of course, students will be able to

CO1	Construct SQL Queries using relational algebra
CO2	Design database using ER model and normalize the database
CO3	Construct queries to handle transaction processing and maintain consistency of the database
CO4	Compare and contrast various indexing strategies and apply the knowledge to tune the performance of the database
CO5	Appraise how advanced databases differ from Relational Databases and find a suitable database for the given requirement.
CO6	To understand Query optimization using Heuristics and Cost Estimation

UNIT I

RELATIONAL DATABASE :

Purpose of Database System – Views of data – Data Models – Database System Architecture – Introduction to relational databases – Relational Model – Keys – Relational Algebra – SQL fundamentals Advanced SQL features – Embedded SQL– Dynamic SQL

PART A

1. What is database?(R)

A database is a basic electronic storage with collection of interrelated data, organized to provide efficient retrieval. Databases are organized by fields, records and files or tables.

2. What is DBMS? (R) (NOV 2008)

A database management system (DBMS) is a software package designed to define, manipulate, retrieve and manage data in a database.

3. List out some representative application of databases (U) (DEC 2007)

Databases are widely used. Here are some representative applications.

- Banking
- Airlines
- Universities
- Credit card transactions
- Telecommunication
- Finance
- Sales
- On-line retailers
- Manufacturing
- Human resources

4. What is the purpose of data base system? (U) (NOV 2014)

The primary goal of a DBMS is to provide an environment that is both convenient and efficient for people to storing and retrieving information.

5. What are problems caused by redundancy?(Nov/Dec-2017)

- The first is that storing values multiple times wastes space.
- The second problem is that when a field value changes, multiple occurrences need to be updated.
- The third problem occurs if we forget to change the values in any of the records. The database would then have inconsistent data.

6. Describe the three levels of data abstraction? (R)(Nov/Dec-2017)

There are three levels of abstraction:

- Physical level
- Logical level
- View level

7. What is a physical level? (R)

The lowest level of abstraction describes how data are stored.

8. What is a logical level? (R)

The next higher level of abstraction, describes what data are stored in database and what relationship among those data.

9. What is a view level? (R)

The highest level of abstraction describes only part of entire database.

10. What is a schema? (R)(MAY 2007)

The structure of a database system, described in a formal language supported by the database management system (DBMS), In a relational database, the schema defines the tables, the fields in each table, and the relationships between fields and tables. Schemas are generally stored in a data dictionary.

11. What are the different types of schemas? (R)

- Physical Schema
- Logical Schema
- External Schema

12. What is an instance? (R)

The collection of information stored in the database at a particular moment is called an instance of the database.

13. What do u mean by data model?(U) (MAY 2011)

Data Model can be defined as an integrated collection of concepts for describing and manipulating data, relationships between data, and constraints on the data in an organization.

14. What are the different data models? (R) (MAY 2012)

- Relational Data Model
- The Entity-Relationship Data Model
- Object-Based Data Model
- Semi structured Data Model

15. What is a relational model? (R) (MAY 2010)

The Relational model uses a collection of tables to represent both data and the relationships among those data. Each table has multiple columns, and each column has a unique Value.

16. What is an entity-relationship model? (R) (MAY 2008)

This data model is based on real world that consists of basic objects called entities and of relationship among these objects. Entities are described in a database by a set of attributes.

17. What is object based data model? (R)

This model is based on collection of objects. An object contains values stored in instance variables with in the object. An object also contains bodies of code that operate on the object. These bodies of code are called methods. Objects that contain same types of values and the same methods are grouped together into classes.

18. What is a semi structured data model? (R)

The semi structured data model permits the specification of data where individual data items of the same type may have different sets of attributes.

19. What is Meta data? (R) (MAY 2008)

Data dictionary contains Metadata-that is, data about data.

20. What is the use of data dictionary? (R)

The DDL gets input as some instructions (statements) and generates some output. The output of the DDL is placed in the data dictionary.

21. Define Relational Database. (R)

A relational database is based on the relational model and uses a collection of tables to represent both data and the relationships among those data. It also includes a DML and DDL.

22. What is data base design ? (R)

Database design mainly involves the design of the database schema. The design of a complete database application environment that meets the needs of the enterprise being modeled requires attention to a broder set of issues

23. What is physical data independence? (R)

Physical data independence deals with hiding the details of the storage structure from user applications. The application should not be involved with these issues, since there is no difference in the operation carried out against the data.

24. What are different access types in DML?(U)

The types of access are

- Retrieval of information stored in the database
- Insertion of new information into the database
- Deletion of information from the database
- Modification of information stored in the database

25. What is a procedural DML? (R)

Procedural DMLs require a user to specify what data are needed and how to get those data.

26. What is declarative DML? (R)

Declarative DMLs also referred to as nonprocedural DMLs require a user to specify what data are needed without specifying how to get those data.

27. Define Assertions(R)

An assertion is any condition that the database must always satisfy.

28. What is storage manager? (R)

A storage manager is a program module that provides the interface between the low level data stored in the database and the application programs and queries submitted to the system.

29. What are the components of storage manager? (U)

- Authorization and integrity manager
- Transaction manager
- File manager
- Buffer manager

30. List out the data structure used to implement the storage manager. (R)

- Data files
- Data dictionary
- Indices

31. What is the use of embedded SQL? (U) (MAY 2012) (NOV 2014)

A fundamental principle underlying embedded SQL, which we call **the dual-mode principle**, is that any SQL statement that can be used interactively can also be embedded in an application program.

32. Write short note on OPEN, FETCH and CLOSE statements. (R)

OPEN:

```
EXEC SQL OPEN <Cursor name>;
```

Opens the specified cursor. A set of rows is thus identified and becomes the current active set for the cursor.

FETCH:

```
EXEC SQL FETCH <Cursor name>  
        INTO <host variable reference  
        commalist>; Advances the specified cursor to the next  
        row in the active set.
```

CLOSE:

```
EXEC SQL CLOSE <Cursor name>;
```

Closes the specified cursor. The cursor now has no current active set.

33. Discuss about dynamic SQL. (R) (NOV 2014)

Dynamic SQL is a part of embedded SQL. It consists of a set of —dynamic statements— which themselves are compiled ahead of time- whose purpose is precisely to support the compilation and execution of regular SQL statements that are constructed at run time.

34. What is a candidate key? (R)

Let K be a set of attributes of relvar R. Then K is a candidate key for R if and only if it has both of the following properties: Uniqueness and Irreducibility

35. What is a primary key? (R)

It is possible for a given relation to have two or more candidate keys. In such a case, the relational model has historically required that exactly one of those keys be chosen as the primary key

36. What is an alternate key? (R)

It is possible for a given relation to have two or more candidate keys. In such a case, the relational model has historically required that exactly one of those keys be chosen as the primary key and the others are then called alternate keys.

37. What is a foreign key? (R)

A foreign key is a set of attributes of some relvar R2 whose values are required to match values of some candidate key of some relvar R1.

38. What is a trigger? (R)

Triggered procedures are precompiled procedures that are stored along with the database and invoked automatically whenever some specified event occurs.

39. What are the relational operators? (R)

Union, Intersect, Difference, Product, Select, Project, Join, Divide are the relational operators.

40. Define Cartesian product. (R)

Cartesian product of two relations a and b, a TIMES b, where a and b have no common

attribute names, to be a relation with a heading that is the union of the headings of a and b and with a body appearing in a and a tuple appearing in b.

41. What is a range variable? (R)

A range variable is a variable that —ranges over some specified relation. If range variable V ranges over relation r, then, at any given time, the expression —V denotes some tuple of r.

42. Discuss about query-by-example? (U)

Query-By-Example syntax, which is attractive and intuitively very simple, is based on the idea of making entries in blank tables.

43. What do you mean by type constraints? (R)

A type constraint is, precisely, a definition of the set of values that constitute a given type.

44. What do you mean by database constraints? (R)

A database constraint is a constraint on the values a given database is permitted to assume.

45. What do you mean by relation constraints? (R)

A relation constraint is a constraint on the values a given relvar is permitted to assume.

46. What do you mean by attribute constraints? (R)

An attribute constraint is a constraint on the values a given attribute is permitted to assume.

47. What is referential integrity? (R)

Referential integrity database must not contain any unmatched foreign key values.

48. What is a transition constraint? (R)

A transition constraint is a constraint on the legal transitions that a given variable—in particular, a given relvar or a given database—can make from one value to another.

49. Give the usage of the rename operation with an example.(MAY 2010) (U)

RA _S1.Name;S2.Name(_S1.Address = S2.Address(_S1(Students) _ _S2(Students))).

50. What are the desirable properties of decomposition?(Apr/May-2017)

- 1.LosslessDecomposition
- 2.DependencyPreservation
3. Lack of Data Redundancy

51. Distinguish between Key and Super Key? (Apr/May-2017)

Key A key is a single or combination of multiple fields. Its purpose is to access or retrieve data rows from table according to the requirement.

A **superkey** is a combination of attributes that can be uniquely used to identify a database record. A table might have many superkeys.

PART B

1. Explain the system structure of a database system with neat block diagram. **(R)** (DEC 2007), (MAY 2010) (MAY 2012)
2. What is a data model? Explain various data model for describing the design of a database at the logical level. **(U)** (**APRIL 2008**), (**MAY 2010**)
3. Explain the different between physical and logical data independence with an example. **(U)**(**APRIL 2008**)
4. Explain the characteristics of database approach. **(U)**(**APRIL 2008**)
5. Describe the three-schema architecture. Why do we need mapping between schema Levels? How do different schema definition languages support this architecture? **(R)** (**DEC 2008**)
6. What is the notation used in E-R diagram? Explain the E-R model structure with Example. **(U)**(**NOV 2014**)
7. Explain the role and functions of the database administrator. **(R)**
8. Explain the following **(R)** **(APRIL 2008)** (**NOV 2014**)
 - i. Data Definition Language.
 - ii. Data Manipulation Language.
 - iii. Data control Language.
 - iv. Views
9. Develop an Entity Relationship model for a library management system. Clearly State the problem Definition, Description, Business Rule and any assumption you make. **(AP)**(**MAY 2009**) (**NOV 2014**)
10. What is Embedded SQL? Give example **(R)**(**NOV 2016**)
11. State and explain the Architecture of DBMS. Draw the ER Diagram for banking Systems. (Home loan applications) **(Nov/Dec-2017)**
12. (i) State and explain the command DDL, DML, DCL with suitable example (7).
(ii) Justify the need of embedded SQL. Consider the relation Student (student no, name, mark and grade). Write the embedded SQL statement in C language to retrieve all the students records whose mark than 90. (6) **(Nov/Dec-2017)**
13. Discuss about Tuple Relational Calculus and Domain Relational Calculus. **(R)** (**DEC 2008**)
(MAY 2012)
14. What are aggregate functions? Explain five built-in aggregate functions. **(U)** (**MAY 2008**)
15. Consider the following relations for a company database Application. **(AP)**(**MAY 2009**)
Employee (**Eno**, Name, Sex, DOB, Doj, Designation, Basic_Pay, Dept_No)

Department (DeptNo,Name)

Project (ProjNo, Name, Dept_No)

Works for (Eno.ProjNo.Date, Hours)

The attributes specified for each relation is self-explanatory. However the business rules are stated as follows. A department can control a project. An employee can work on any number of projects on a day. However an employee cannot work more than once on a project he/she worked on that day. The primary key are underlined.

- (i). Identify the forging keys, Develop DDL to implement the above schema.
- (ii) Develop an SQL query to list the department number and the number of Employee in each department.
- (iii) Develop a View that will keep track of the department number, the number of employees in the department, and the total basic pay expenditure for each department.
- (iv) Develop an SQL query to list the details of employees who have worked in more than three projects on a day.

16. Explain how dangling tuple may arise and also explain problems that they Cause. (MAY 2008)

17. Briefly present a survey on Integrity and Security. (U) (MAY 2012)

18. Create an EMPLOYEE table and write the steps for various functions in an SQL Like add, update, delete, save and join the various attributes of an EMPLOYEE Table.

(C)(MAY 2007)

19. Explain the use of trigger with your own example. (U) (MAY 2010)

20. What is a view? How can it be created? Explain with an example. (R) (MAY 2010)

21. Briefly present a survey on Integrity and Security. (U) (MAY 2012)

22. Consider a student registration database comprising of the given below table schema.(AP)(Nov/Dec2015)

Student File

Student number	Student name	Address	Telephone
----------------	--------------	---------	-----------

Course File

Course Number	Description	Hours	Professor Number
---------------	-------------	-------	------------------

Professor File

Professor Number	Name	Office
------------------	------	--------

Registration File

Student Number	Course Number	Date
----------------	---------------	------

Consider a suitable example of tuples/records for the above mentioned tables and write DML statements (SQL) to answer for the queries listed below:

- (i) Which courses does a specific professor teach? (ii) What courses are taught by two specific professors?
- (iii) What teaches a specific course and where is his/her office?
- (iv) For a specific student number, in which courses is the student registered and what is his/her name?
- (v) Who are the professors for a specific student?
- (vi) Who are the students registered in a specific course?

23. Let relations $r_1(A,B,C)$ and $r_2(C,D,E)$ have the following properties r_1 has 20000 tuples r_2 has 45000 tuples, 25 tuples of r_1 , fit on one block and 30 tuples of r_2 fit on one block. Estimate the number of block transfers and seeks required, using each of the following join strategies for $r_1 \bowtie r_2$:

- (i) Nested loop join
- (ii) Block Nested loop join
- (iii) Merge join
- (iv) Hash join **(AP)(NOV 2016)**

24. Consider the relation schema given in fig 1. Design and draw an ER diagram that capture the information of this schema. **(5)(Apr/May-2017)**

Employee(empno,name,office,age)

Books(isbn,title,authors,publisher)

Loan (empno, isbn,date)

Write the following queries in relational algebra and SQL.

- (i) Find the names of employees who have borrowed a book published by McGraw-Hill(5)
- (ii) Find the names of employees who have borrowed all books published by McGraw-Hill(5)

UNIT II

DATABASE DESIGN

Entity-Relationship model – E-R Diagrams – Enhanced-ER Model – ER-to-Relational Mapping – Functional Dependencies – Non-loss Decomposition – First, Second, Third Normal Forms, Dependency Preservation – Boyce/Codd Normal Form – Multi-valued Dependencies and Fourth Normal Form – JoinDependencies and Fifth Normal Form
Fifth Normal Form

PART A

1. What is entity? (R)

It is a 'thing' in the real world with an independent existence.

2. What is relationship set? (R)

The collection (or set) of similar relationships.

3. What is entity set? (R)

It is a collection of all entities of particular entity type in the database.

4. What is weak entity? (R)(NOV 2008)(NOV 2016)

An entity set may not have sufficient attributes to form a primary key, and its primary key comprises of its partial key and primary key of its parent entity, then it is said to be Weak Entity set.

5. What is mapping cardinality? (R)

Mapping cardinality express the number of entities to which another entity can be associated via a relationship set.

6. What is functional dependence? (U) (Apr/May2015)

A Functional dependency is denoted by $X \rightarrow Y$ between two sets of attributes X and Y that are subsets of R specifies a constraint on the possible tuple that can form a relation state r of R. The constraint is for any two tuples t1 and t2 in r if $t1[X] = t2[X]$ then they have $t1[Y] = t2[Y]$. This means the value of X component of a tuple uniquely determines the value of component Y.

7. What is object oriented data model? (R)

The object –oriented data model is another data model that has seen increasing attention. The object-oriented model can be seen as extending the E-R model with notions of encapsulation, methods and object identity.

8. What do you mean by functional dependencies? (U) (MAY 2010)(nov2013)(Dec2015)

Let R be a relation variable and let X and Y be arbitrary subsets of the set of attributes of R. then we say that Y is functionally dependent on-in symbols,

$$X \rightarrow Y$$

(read —X functionally determines Y,| or simply — X arrow Y |)- if and only if in every possible legal value of R, each X Value has associated with it precisely one Y value.

9. What is a trivial dependency? (R) (MAY 2012)(nov 2013)

A dependency is trivial if it cannot possibly fail to be satisfied.

10. What is a nontrivial dependency? (R) (MAY 2012)

Nontrivial dependencies are the ones that correspond to —genuinell integrity constraints.

11. Define Heath's theorem. (R)

Let R{A,B,C} be a relvar, where A,B, and C are sets of attributes. If R satisfies the FD $A \rightarrow B$, then R is equal to the join of its projections on {A,B} and {A,C}.

12. Define first normal form.(R)

A relvar is in 1NF if and only if, in every legal value of that relvar, every tuple contains exactly one value for each attribute.

13. Define second normal form.(R)

A relvar is in 2NF if and only if it is in 1NF and every nonkey attribute is irreducibly dependent on the primary key.

14. Define third normal form. (R)

A relvar is in 3NF if and only if it is in 2NF and every nonkey attribute is nontransitively dependent on the primary key.

15. What is dependency preservation? (R)

Normalization procedure should decompose relvars into projections that are independent in Rissanen's sense has come to be known as dependency preservation.

16. Define Boyce/Codd normal form. (R) (MAY 2012)

A relvar is in BCNF if and only if every nontrivial, left-irreducible FD has a candidate key as its determinant (OR) A relvar is in BCNF if and only if every determinant is a candidate key.

17. What is multi-valued dependence? (R)

Let R be a relvar, and let A, B and C be subsets of the attributes of R. Then we say that B is multi-dependent on A-in symbols, $A \twoheadrightarrow B$

18. Define Fagin theorem. (R)

Let R{A,B,C} be a relvar, where A, B, and C are sets of attributes. Then R is equal to the join of its projections on {A, B} and {A, C} if and only if R satisfies the MVDs $A \twoheadrightarrow B|C$.

19. Define fourth normal form. (R)

Relvar R is in 4NF if and only if, whenever there exist subsets A and B of the attributes of R such that the nontrivial MVD $A \twoheadrightarrow B$ is satisfied, then all attributes of R are also functionally dependent on A.

20. What is a join dependency? (R)

Let R be a relvar, and let A, B,.... Z be subsets of the attributes of R. Then we say that R satisfies the JD

$*\{A, B \dots Z\}$

if and only if every legal value of R is equal to the join of its projections on A, B,.....Z.

21. Define fifth normal form. (R)

A relvar R is in 5NF-also called projection-join normal form(PJ/NF)-if and only if every nontrivial join dependency that is satisfied by R is implied by the candidate key(s) of R, where

a. The join dependency $*\{A, B, \dots Z\}$ on R is trivial if and only if atleast one of A, B,...Z is the set of all attributes of R.

b. The join dependency $*\{A, B, \dots Z\}$ on R is implied by the candidate key(s) of R if and

22. What is denormalization? (R)

Let R1, R2,....., Rn be a set of relvars. Then denormalizing those relvars means replacing them by their join R, such that for all i (i=1, 2,....., n) projecting R over the attributes of Ri is guaranteed to yield Ri again.

23. What is normalization? (R) (MAY 2010)

Normalization is the process of structuring relational database schema such that most ambiguity is removed.

24. What are the aspects of relational model?(R)

The Three aspects of relational model are:

- Structural aspect: The data in the database is perceived by the user as tables, and nothing but tables.
- Integrity aspect: Those tables satisfy certain integrity constraints.
- Manipulative aspect: The operators available to the user for manipulating those tables- for example, for purposes of data retrieval – are operators that derive tables from tables. Of those operators, three particularly important ones are select, project and join.

25. Mention the important points about relational databases. (R)

A set of important points about relational databases:

- a. Relational databases store data in the form of tables (logically).
- b. The rows of a table are called as tuples.
- c. The columns of a table are known as attributes.
- d. Every attribute has a data type associated with it.
- e. Every attribute has a domain which provides the set of all possible values that can be stored as values for that attribute.
- f. Tables are called as relations.
- g. The table names are called as relational variable

26. What is a catalog? (R)

The catalog contains detailed information, sometimes called descriptor information or Metadata, regarding the various objects that are of interest to the system itself.

27. What do you mean by weak entity set? (MAY 2010) (R)

A **weak entity** is an entity that cannot be uniquely identified by its attributes alone; therefore, it must use a foreign key in conjunction with its attributes to create a primary key. The foreign key is typically a primary key of an entity it is related to.

28. What is the difference between Data Integrity and Data Security? (NOV 2013) (U)

Data integrity and data security are two different aspects that make sure the usability of data is preserved all the time. Main difference between integrity and security is that integrity deals with the validity of data, while security deals with protection of data. Backing up, designing suitable user interfaces and error detection/correction in data are some of the means to preserve integrity, while authentication/authorization, encryptions and masking are some of the popular means of data security. Suitable control mechanisms can be used for both security and integrity.

29. Which operators are called as unary operators and why are they called so (NOV 2013) (U)

An operator that takes a single operand in an expression or a statement. The unary operators in C# are +, -,!, ~, ++, -- and the cast operator.

30. What is an Entity? (R)

It is a 'thing' in the real world with an independent existence.

31. What is a query? (NOV2013) (R)

A query is a statement requesting the retrieval of information.

32. Define query language? (R)

A query language is a language in which user requests information from the database.

PART B

1. What is the notation used in E-R diagram? Explain the E-R model structure with Example.

(U)(NOV 2014)

2. Develop an Entity Relationship model for a library management system. Clearly State the problem

Definition, Description, Business Rule and any assumption you make. **(AP)**

(MAY 2009) (NOV 2014)

3. Develop an Entity Relationship model preparation staff (chef) and finalize the customer_s bill. The Food preparation staffs (Chefs), with their touch-display interfaces to the system, are able to view orders sent to the kitchen by waiters. During preparation, they are able to let the waiter know the status of each item, and can send notifications when items are completed. The system should have full accountability and logging facilities, and should support Supervisor actions to account for exceptional circumstances, such as a meal being refunded or walked out on. **(AP) (Dec2015)**

4. Distinguish between lossless-join decomposition and dependency preserving decomposition(7)

5. State and explain the architecture of DBMS. Draw the ER diagram for banking systems.(Home loan Application)(13)(Nov/Dec-2017)

UNIT III

TRANSACTIONS

Transaction Concepts – ACID Properties – Schedules – Serializability – Transaction support in SQL – Need for Concurrency – Concurrency control – Two Phase Locking- Timestamp – Multiversion – Validation and Snapshot isolation– Multiple Granularity locking – Deadlock Handling – Recovery Concepts – Recovery based on deferred and immediate update – Shadow paging – ARIES Algorithm

PART – A

1. Define recovery in a database system.(R)

Recovery in a database system means, primarily, recovering the database itself; that is, restoring the database to a correct state after some failure has rendered the current state incorrect, or at least suspect.

2. What is a transaction? (MAY 2010)

A transaction is a logical unit of work; it begins with the execution of a BEGIN TRANSACTION operation, and ends with the execution of a COMMIT or ROLLBACK operation.

3. What is a commit point? (R)

A commit point corresponds to the end of a logical unit of work, and hence to a point at which the database is supposed to be in a correct state.

4. Discuss the ACID properties. (R) (MAY 2010) (NOV 2014)(Dec2015)

ACID properties are

- a. **Atomicity:** Transactions are atomic
- b. **Correctness:** Transactions transform a correct state of the database into another correct state, without necessarily preserving correctness at all intermediate points
- c. **Isolation:** Transactions are isolated from one another. That is, even though in general there will be many transactions running concurrently, any given transaction's updates all connected from all the rest, until that transaction commits. Another way of saying the same thing is that, for any two distinct transactions A and B, A might see B's updates or B might see A's updates, but not both.

d. Durability: Once a transaction commits, its updates persist in the database, even if there is a subsequent system crash.

5. What is a system failure? (R)

System failures (e.g., power outage), which affect all transactions currently in progress but do not physically damage the database. A system failure is sometimes called a soft crash.

6. What is a media failure? (R)

Media failures (e.g. head crash on the disk), which do cause damage to the database or some portion thereof, and affect at least those transactions currently using that portion. A media failure is sometimes called a hard crash.

7. What is a check point record? (R)

The checkpoint record contains a list of all transactions that were in progress at the time the checkpoint was taken.

8. Expand ARIES. (R)

The name ARIES stands for —Algorithms for Recovery and Isolation Exploiting Semantics.‡

9. What are the three broad phases of ARIES? (U)

ARIES operates in three broad phases:

- **Analysis:** Build the REDO and UNDO lists
- **Redo:** Start from a position in the log determined in the analysis phase and restore the database to the state it was in at the time of the crash.
- **Undo:** Undo the effects of transactions that failed to commit.

10. What is a two-phase commit? (U)

Two-phase commit is important whenever a given transaction can interact with several independent —resource managers,‡ each managing its own set of recoverable resources and maintaining its own recovery log.‡

11. What is the need for save points? (U)

It might be possible for a transaction to establish intermediate save points while it is executing, and subsequently to roll back to a previously established save point, if required, instead of having to roll back all the way to the beginning.

12. Define concurrency (MAY 2012) (R)

Concurrency refers to the fact that DBMSs typically allow many transactions to access the same database at the same time.

13. What are the three problems that any concurrency control mechanism must address? (R)

The three problems are:

- The lost update problem
- The uncommitted dependency problem
- The inconsistent analysis problem

14. What is the last update problem? (U)

Transaction A retrieves some tuple t at time t_1 ; transaction B retrieves that same tuple t at time t_2 ; transaction A updates the tuple at time t_3 ; and transaction B updates the same tuple at time t_4 ; Transaction A's update is lost at time t_4 , because transaction B overwrites it without even looking at it.

15. What is the uncommitted dependency problem? (U)

The uncommitted dependency problem arises if one transaction is allowed to retrieve-or, worse, update-a tuple that has been updated by another transaction but not yet committed by that other transaction.

16. What is deadlock? (NOV 2014) (R)

Deadlock is a situation in which two or more transactions are in a simultaneous wait state, each of them waiting for one of the others to release a lock before it can proceed.

17. What is serializability? (NOV2016) (NOV 2014) (R)

Serializability is the generally accepted —criterion for correctness— for the interleaved execution of a set of transactions; that is, such an execution is considered to be correct if and only if it is serializable.

18. State the two-phase locking theorem (R)

Two-phase locking theorem is —If all transactions obey the two-phase locking protocol, then all possible interleaved schedules are serializable.

19. State the two-phase locking protocol. (MAY 2012)(NOV 2013) (R)

The two-phase locking protocol is as follows

- i. Before operating on any object a transaction must acquire a lock on that object.
- ii. After releasing a lock, a transaction must never go on to

acquire any more locks.

20. What is an isolation level? (R)

The isolation level that applies to a given transaction might be defined as the degree of interference the transaction in question is prepared to tolerate on the part of concurrent transactions.

21. What do you mean by phantom problem? (R)

If transactions operate at less than the maximum isolation level is the so-called phantom problem

22. What is an intent locking protocol? (U)

The intent locking protocol, according to which no transaction is allowed to acquire a lock on a tuple before first acquiring a lock-probably an intent lock on the relvar that contains it.

23. List the four conditions for Deadlock(U) (NOV 2016)

A deadlock situation can arise if and only if the following four conditions hold simultaneously in a system-

Mutual Exclusion: At least one resource is held in a non-sharable mode that is only one process at a time can use the resource. If another process requests that resource, the requesting process must be delayed until the resource has been released.

Hold and Wait: There must exist a process that is holding at least one resource and is waiting to acquire additional resources that are currently being held by other processes.

No Preemption: Resources cannot be preempted; that is, a resource can only be released voluntarily by the process holding it, after the process has completed its task.

Circular Wait: There must exist a set $\{p_0, p_1, \dots, p_n\}$ of waiting processes such that p_0 is waiting for a resource which is held by p_1 , p_1 is waiting for a resource which is held by p_2, \dots
 p_{n-1} is waiting for a resource which is held by p_n and p_n is waiting for a resource which is held by p_0 .

24. What is a shadow copy scheme?

It is simple, but efficient, scheme called the shadow copy schemes. It is based on making copies of the database called shadow copies that one transaction is active at a time. The scheme also assumes that the database is simply a file on disk.

25. What type of locking needed for insert and delete operations (April/May-2017)

When you execute an INSERT, UPDATE, or DELETE statement, the database server uses exclusive locks. An exclusive lock means that no other users can update or delete the item until the database removes the lock

PART – B

1. Why is Recovery needed? Discuss any two Recovery Techniques.(U)
(MAY2009) (MAY 2012)
2. Write a relevant example discuss two phase Locking. (U)(MAY 2009)
3. What is Deadlock? List and discuss the four conditions for Deadlock. (R)
(MAY 2009)(MAY 2012) (NOV 2014) (NOV2016)
4. Give an example of a Scenario where two Phase Locking leads to Deadlock.(AN)(MAY 2009)
5. Explain testing for serializability with respect to concurrency control schemes. How will you determine, whether a schedule is serializable or not.
(U)(MAY 2008)
6. Explain the following protocols for concurrency control :(R)(MAY 2008)
 - a. Lock based protocols.
 - b. Time stamp based protocols.
7. Describe the three most common concurrent transaction execution problems.(R)(MAY 2007)
8. Explain how concurrency control can be used to avoid such problems?(U)
(MAY 2007)(NOV 2014)
9. Write a note on the following :

- a. Two-phase locking protocol (2PL).(U)(NOV 2016)
 - b. Lock based protocol.(U)
10. Explain the deferred and immediate-modification version of the log based Recovery scheme.
(U)(MAY 2007)
11. Write short notes on shadow paging.(R)(DEC 2007)
12. Describe about the testing of serializability. (U)(DEC 2007) (MAY 2012)
13. What is concurrency control? How is it implemented in DBMS? Explain. (R)(DEC 2007)
14. Explain various recovery techniques during transaction in detail. (R)(MAY 2007)
15. Explain the two-phase commit and three-phase commit protocols.(R)(Dec2015)
16. Consider the following schedules. The actions are listed in the order they are scheduled, and prefixed with the transaction name.(AN)
(Dec2015)
S1:T1:R(X),T2:R(X),T1:W(Y),T2:W(Y),T1:R(Y),T2:R(Y)
S2:T3:W(X),T1:R(X),T1:W(Y),T2:R(Z),T2:W(Z),T3:R(Z)
For each of the schedule, answer the following questions:
(i) what is the precedence graph for the schedule?(2)
(ii) is the schedule conflict –serializable?if so,what are all the conflict equivalent serial schedules?(7) (iii)is the schedule view-serializable?if so, what are all the view equivalent serial schedules?(7)
17. 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 instruction to transactions T1 and T2,so that they observe the two phase locking protocol. Can the execution of these transactions result in a deadlock?(AP) (NOV 2016)

UNIT IV

IMPLEMENTATION TECHNIQUES

RAID – File Organization – Organization of Records in Files – Data dictionary Storage – Column Oriented Storage– Indexing and Hashing –Ordered Indices – B+ tree Index Files – B tree Index Files – Static Hashing – Dynamic Hashing – Query Processing Overview – Algorithms for Selection, Sorting and join operations – Query optimization using Heuristics - Cost Estimation.

PART – A

1. What is cache memory?(U)

The cache is the fastest and most costly form of storage. Cache memory is small; its use is managed by the computer system hardware.

2. What is main memory?(U)

The storage medium used for data that are available to be operated on is main memory. The general-purpose machine instructions operate on main memory. Although main memory may contain many megabytes of data, or even hundreds of gigabytes of data in large server systems, it is generally too small for storing the entire database. The contents of main memory are usually lost if a power failure or system crash occurs.

3. What is flash memory? (MAY 2010)(U)

EEPROM (Electrically Erasable Programmable Read Only Memory)

- a. Data in flash memory survive from power failure
- b. Reading data from flash memory takes about 10 nano-secs (roughly as fast as from main memory), and writing data into flash memory is more complicated: write-once takes about 4-10 microsecs.
- c. To overwrite what has been written, one has to first erase the entire bank of the memory. It may support only a limited number of erase cycles (10^4 to 10^6).
- d. It has found its popularity as a replacement for disks for storing small volumes of data (5- 10 megabytes).

4. What is magnetic disk storage?(U)

Primary medium for long-term storage.

- Typically the entire database is stored on disk.
- Data must be moved from disk to main memory in order for the data to be operated on.
- After operations are performed, data must be copied back to disk if any changes were made.
- Disk storage is called direct access storage as it is possible to read data on the disk in any order (unlike sequential access).
- Disk storage usually survives power failures and system crashes

5. What is optical storage?(U)

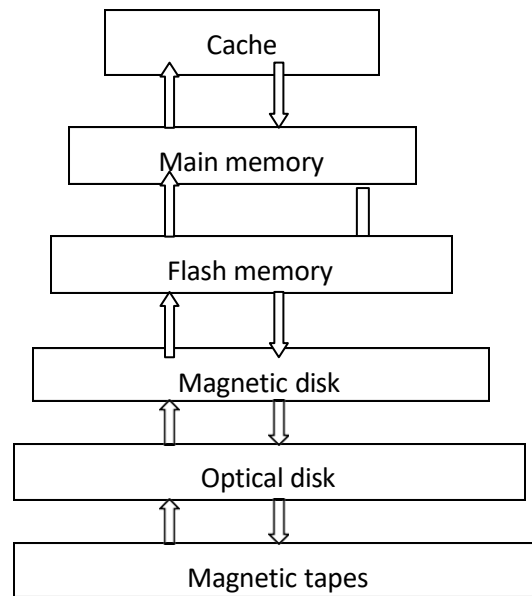
The most popular form of optical storage are CD-ROM (compact – disk read – only memory), WORM (write-once read-many)disk (for archival storage of data), and Juke box(containing a few drives and numerous disks loaded on demand).

6. What is tape storage?(U)

Tape storage used for primarily for backup and archival data.

- Cheaper, but much slower access, since tape must be read sequentially from the beginning.
- Used as protection from disk failures.

7. Draw the storage device hierarchy. (MAY 2010)(R)



8. What is SAN architecture?(R)

In the storage area network (SAN) architecture, large numbers of disks are connected by a high-speed network to a number of server computers.

9. What is NAS?(R)

Network attached storage (NAS) is an alternative to SAN. NAS is much like SAN, except that instead of the networked storage appearing to be a large disk, it provides a file system interface using networked file system protocols such as NFS or CIFS.

10. Define access time(R)

Access time is the time from when a read or write request is issued to when data transfer begins.

11. Define seek time.(R)

To access data on a given sector of a disk, the arm first must move so that it is positioned over the correct track, and then must wait for the sector to appear under it as the disk rotates. The time for repositioning the arm is called the seek time.

12. Define average seek time.(R)

The average seek time is the average of the seek times, measured over a sequence of random requests.

13. Define rotational latency time.(U)

Once the head has reached the desired track, the time spent waiting for the sector to be accessed to appear under the head is called the rotational latency time

14. Define average latency time.(U)

Average latency time of the disk is one-half the time for a full rotation of the disk.

15. Define data-transfer rate.(U)

The data-transfer rate is the rate at which data can be retrieved from or stored to the disk.

16. What is RAID? (NOV 2014)(R)

A variety of disk organization techniques, collectively called redundant arrays of independent disks (RAID), have been proposed to achieve improved performance and reliability.

17. What are the factors to be taken into account in choosing a RAID level?(R)

The factors to be taken into account in choosing a RAID level are

- Monetary cost of extra disk-storage requirements
- Performance requirements in terms of number of I/O operations
- Performance when a disk has failed.
- Performance during rebuild(that is, while the data in a failed disk are being rebuilt on a new disk)

18. What is heap file organization?(R)

Any record can be placed anywhere in the file where there is space for the record. There is no ordering of records. Typically, there is a single file for each relation.

19. What is sequential file organization?(R)

Records are stored in sequential order, according to the value of a —search key| of each record.

20. What is hashing file organization?(R)

A hash function is computed on some attribute of each record. The result of the hash function specifies in which block of the file the record should be placed,

21. What is a multitable clustering file organization?(U)

A multitable clustering file organization is a file organization that stores related records of two or more relations in each block. Such a file organization allows us to read records that would satisfy the join condition by using one block read.

22. What are the two basic kinds of indices?(U)

The Two basic kinds of indices are

- Ordered indices
- Hash indices

23. What are access types?(U)

The types of access that are supported efficiently. Access types can include finding records with a specified attribute value and finding records whose attribute values fall in a specified range.

24. What is access time?(U)

The time it takes to find a particular data item, or set of items, using the technique in question.

25. What is insertion time?(U)

The time it takes to insert a new data item. This value includes the time it takes to find the correct place to insert the new data item, as well as the time it takes to update the index structure.

26. What is deletion time?(U)

The time it takes to delete a data item. This value includes the time it takes to find the item to be deleted, as well as the time it takes to update the index structure.

27. What is space overhead?(U)

The additional space occupied by an index structure. Provided that the amount of additional space is moderate, it is usually worth-while to sacrifice the space to achieve improved performance.

28. What is a search key? (R)

An attribute or set of attributes used to look up records in a file is called a search key.

29. What is an clustering index (or) primary index? (NOV 2014)(U) Clustering index is an index whose search key also defines the sequential order of the file.

30. What is an nonclustering index (or) secondary index?(U)

Indices whose search key specifies an order different from the sequential order of the file are called nonclustering indices.

31. What is an index-sequential file?(U)

All files are ordered sequentially on some search key. Such files, with a clustering index on the search key, are called index-sequential files.

32. What is a dense index? (U) (MAY 2012)

An index record appears for every search-key value in the file. In a dense clustering index, the index record contains the search-key value and a pointer to the first data record with that search-key value. The rest of the records with the same search-key value would be stored sequentially after the first record, since, because the index is a clustering one, records are stored on the same search key. Dense index implementations may store a list of pointers to all records with the same search-key value; doing so is not essential for clustering indices.

33. What is a sparse index? (U) (MAY 2012)

An index record appears for only some of the search-key values. as is true in dense indices, each index record contains a search-key value and a pointer to the first data record with that search-key value. To locate a record, we find the index entry with the largest search-key value that is less than or equal to the search-key value for which we are looking. We start at the record pointed to by that index entry, and follow the pointers in the file until we find the desired record.

34. What is a multilevel indices?(U)

Indices with two or more levels are called multilevel indices.

35. What is a balanced tree?(U)

Balanced tree in which every path from the root of the tree to the leaf of the tree is of the same length.

36. Write the querying a B+ tree.(AP)

Procedure find (value V)

Set C = root node

While C is not a leaf node begin

Let K_i = smallest search-key
value, if any, greater than V If

there is no such value then begin

Let m = the number
of pointers in the
node Set C = node
pointed to by P_m

End End

Else set C = the node pointed to by P_i

If there is a key value $K_i = V$
Then pointer P_i directs us to the
desired record or bucket Else no
record with key value k exists.

37. Define hashing.(R)

Hashing allow us to avoid accessing an index structure. It provides a way of constructing indices.

38. Define a bucket.(R)

The term bucket to denote a unit of storage that can store one or more records. A bucket is typically a disk block, but could be chosen to be a smaller or larger than a disk block.

39. Define hash function.(U)

A hash function h is a function from K to B . Let h denote a hash function

40. What are the two different purposes of hashing?(R)

Hashing can be used for two different purposes.

- Hash file organization
- Hash index organization

41. What is a hash index?(R)

A hash index organizes the search keys, with their associated pointers, into a hash file structure.

42. What is dynamic hashing? (R) (NOV 2014)

Dynamic hashing techniques allow the hash function to be modified dynamically to accommodate the growth or shrinkage of the database.

43. What are the steps involved in query processing?(R)

The steps involved in processing a query are

- Parsing and translation
- Optimization
- Evaluation

44. Write the nested-loop join algorithm.(U)

For each tuple t_r in r do begin

 For each tuple t_s in s do begin

 Test pair (t_r, t_s) to see if they
 satisfy the join condition If

they do, add $t_r.t_s$ to the
result end

45. Write the block nested-loop join and its algorithm.(U)

Block nested-loop join, which is a variant of the nested – loop join where every block of the inner relation is paired with every block of the outer relation.

```
for each tuple  $t_r$  in r do begin
for each tuple  $t_s$  in s do begin
test pair  $(t_r, t_s)$  to test if they satisfy the given join condition ?
if test satisfied
add  $t_r . t_s$  to the result;
end inner loop
end outer loop
```

46. What is an indexed nested-loop join?(U)

Indexed nested loop join can be used with existing indices, as well as with temporary indices created for the sole purpose of evaluating the join.

47. What is a merge join?(U)

The merge-join algorithm can be used to compute natural joins and equi-joins.

48. What is a hash-join?(U)

The hash-join algorithm can be used to implement natural joins and equi-joins.

49. What are the contents of DBMS catalog?(R)


The DBMS catalog stores the following statistical information about database relations:

- n_r , the number of tuples in the relation r
- b_r , the number of blocks containing tuples of relation r
- l_r , the size of a tuple of relation r in bytes
- f_r , the blocking factor of relation r -that is, the number of tuples of relation r that fit into one block
- $V(A,r)$, the number of distinct values that appear in the relation r for attribute A . This value is the same as the size of (r) . If A is a key for relation r , $V(A,r)$ is n_r

50. What is a slotted page sheet. draw the diagram(R) (NOV 2013)

Record id = <page id, slot #>

Can move records on page without changing rid; so, attractive for fixed-length records too. Page is full when data space and slot array meet.



51. What is a histogram?(U)

In histogram the values for the attribute are divided into a number of ranges, and with each range the histogram associates the number of tuples whose attribute value lies in that range.

52. What is an index record?(U)

An index record, or index entry, consists of a search – key vale and pointers to one or more records with that value as their search-key value.

53. Define mean time to failure(MTTF)(R)

Mean time to failure is a measure of the reliability of the disk. The mean time to failure of a disk is the amount of time that, on average, we can expect the system to run continuously without any failure.

54. What is mirroring?(R)

The simplest approach to introducing redundancy is to duplicate every disk. This technique is called mirroring.

55. How dynamic hashing differ from static hashing? (U) (DEC2015) Static Hashing

- A bucket is a unit of storage containing one or more records (a bucket is typically a disk block).
- The file blocks are divided into M equal-sized buckets, numbered bucket0, bucket1... bucketM- 1. Typically, a bucket corresponds to one (or a fixed number of) disk block.
- In a hash file organization we obtain the bucket of a record directly from its search-key value using a hash

Dynamic Hashing

- Good for database that grows and shrinks in size
- Allows the hash function to be modified dynamically
- **Extendable hashing** – one form of dynamic hashing
- Hash function generates values over a large range —typically b -bit integers, with $b = 32$.
- At any time use only a prefix of the hash function to index into a table of bucket addresses.

- Let the length of the prefix be i bits, $0 \leq i \leq 32$.

56. Disadvantages of B TREE over B+ TREE(R)(NOV 2016)

In a B tree search keys and data stored in internal or leaf nodes. But in B+ tree data store only leaf nodes.

Searching any data in a B+ tree is very easy because all data are found in leaf nodes. In a B tree, data cannot be found in leaf nodes.

In a B tree, data may be found in leaf nodes or internal nodes. Deletion of internal nodes is very complicated. In a B+ tree, data is only found in leaf nodes. Deletion of leaf nodes is easy.

Insertion in B tree is more complicated than B+ tree.

B+ trees store redundant search key but B tree has no redundant value.

In a B+ tree, leaf nodes data are ordered as a sequential linked list but in B tree the leaf node cannot be stored using a linked list. Many database systems' implementations prefer the structural simplicity of a B+ tree.

57. What is a query execution plan?(Apr/May-2017)

Execution plan will be generated by Query optimizer with the help of statistics and Algebraizer\processor tree.

It is the result of Query optimizer and tells how to do\perform your work\requirement.

There are two different execution plans - Estimated and Actual.

Estimated execution plan indicates optimizer view.

Actual execution plan indicates what executed the query and how was it done.

58. Which cost component are used most often as the basis for cost function? (Apr/May-2017)

(i) **Access cost to secondary storage**

(ii) **Memory usage of cost**

(iii) **Storage cost**

59. What is replication transparency?(Apr/May-2017)(U)

Distribution Transparency. Distribution transparency is the property of **distributed databases** by the

virtue of which the internal details of the **distribution** are hidden from the users. The DDBMS designer may choose to fragment tables, **replicate** the fragments and store them at different sites.

60. State the Need for Query Optimization.

- Process of selecting an efficient execution plan for evaluating the query.
- refers to the process of finding lowest cost method of evaluating a given query.
- Need for being logically consistent because the least cost plan will always be consistently low.

PART B

1. Explain the security features provided in commercial query languages.(U) (MAY 2007)

2. What are the steps involved in query processing? How would you estimate the cost of the query? (U) (MAY 2007)

3. Explain the different properties of indexes in detail.(U)

4. Explain various hashing techniques.(R)

5. Explain the four important properties of transaction that a DBMS must ensure to maintain database.(U)(MAY 2007)

6. What is RAID? List the different level in RAID technology and explain its Features.(R)(MAY 2007), (MAY 2010) (NOV 2014)(NOV 2016)

7. What is concurrency control? How is it implemented in DBMS? Explain.(U)(MAY 2007)

8. Explain various recovery techniques during transaction in detail.(U) (MAY 2007)

9. Explain how the RAID systems improve performance and reliability. (U)(DEC 2007)

10. Describe the structure of B⁺ tree and list the characteristics of a B⁺tree.(U)

**(DEC 2008), (MAY 2012) (MAY 2010)
(NOV 2014)**

11. Explain the steps involved in processing a query.(R)(DEC 2008)

- 12.** Describe static hashing and dynamic hashing.(U)(DEC 2008). (NOV 2014)
- 13.** Describe in detail about how records are represented in a file and how to organize them in a file. (AP)(MAY 2012).
- 14.** Explain about spatial and mobile database (U)(NOV 2014)(NOV 2016)
- 15.** With suitable diagrams, discuss about the Raid Levels(Level0,Level1,Level 0+1,level3,Level and Level5).(U)(Dec2015)
- 16.** Explain the Architectural components of a data ware house and write about data marts. (R)(Dec2015)
- 17.** Explain the architecture distributed database system (U)(Nov 2016)
- 18.** Describe the benefits and drawbacks of a source driven architecture for gathering of data datawarehouse , as compared to a destination driven architecture (U)(NOV 2016)

UNIT-V

ADVANCED TOPICS

Distributed Databases: Architecture, Data Storage, Transaction Processing, Query processing and optimization – NOSQL Databases: Introduction – CAP Theorem – Document Based systems – Key value Stores – Column Based Systems – Graph Databases. Database Security: Security issues – Access control based on privileges – Role Based access control – SQL Injection – Statistical Database security – Flow control – Encryption and Public Key infrastructures – Challenges

PART-A

1. What are the different aspects of security problem?(R)

There are many aspects to the security problem. Here are some of them:

- i. Legal, social, and ethical aspects
- ii. Physical controls
- iii. Policy questions
- iv. Operational problems
- v. Hardware controls
- vi. Operational system support and finally

Issues that are the specific concern of the database system itself.

2. What is discretionary access control?(U)

In the case of discretionary control, a given user will typically have different access rights on different objects; further, there are few inherent limitations regarding which users can have which rights on which objects. Discretionary schemes are thus flexible.

3. What is mandatory access control?(U)

In the case of mandatory control, by contrast, each data object is labeled with a certain classification level, and each user is given a certain clearance level.

4. What are the contents of audit trail record?(R)

A typical audit trail record might contain the following information:

- Request(source text)
- Terminal from which the operation was invoked
- User who invoked the operation
- Date and time of the operation
- Relvar(s), tuples(s),attribute(s) affected
- Before images(old values)
- After images(new values)

5. What is entity integrity?(U)

No component of the primary key of any base relvar is allowed to accept nulls are called entity integrity.

6. What are views for?(U)

There are many reasons why view support is desirable. Here are some of them:

- Views provide a shorthand or —macro capability.
- Views allow the same data to be seen by different users in different ways at the same time.
- Views provide automatic security for hidden data.
- Views can provide logical data independence.

7. Define the golden rule.(R)

Golden Rule is

No update operation must ever assign to any relvar a value that causes its relvar predicate to evaluate to FALSE.

Or (a little loosely):

No relvar must ever be allowed to violate its own predicate.

8. Define a distributed database system.(U)

A distributed database system consists of a collection of sites, connected together via some kind of communications network, in which:

- a. Each site is a full database system site in its own right but
- b. The sites have agreed to work together so that a user at any site can access data anywhere in the network exactly as if the data were all stored at the user's own site.

9. Define a distributed database management system.(R)(NOV 2016)

A new software component at each site logically an extension of the local DBMS provides the necessary partnership functionality, and it is the combination of these new components together with the existing DBMSs that constitutes what is usually called the distributed database management system

10. What are the advantages of distributed databases?(U)

It enables the structure of the database to mirror the structure of the enterprise-local data can be kept locally, where it most logically belongs- while at the same time remote data can be accessed when necessary.

11. What is the fundamental principle of distributed database?(U)

The fundamental principle of distributed database is to the user, a distributed system should look exactly like a non-distributed system.

12. What are the objectives of distributed databases?(R)

The objectives of distributed databases are

- 1) Local autonomy
- 2) No reliance on a central site
- 3) Continuous operation
- 4) Location independence
- 5) Fragmentation independence
- 6) Replication independence
- 7) Distributed query processing
- 8) Distributed transaction management
- 9) Hardware independence
- 10) Operating system independence
- 11) Network independence
- 12) DBMS independence

13. What are the problems of distributed databases?(R)

The problems are

- Query processing
- Catalog management
- Update propagation
- Recovery
- Concurrency

14. Define client/server systems.(R)

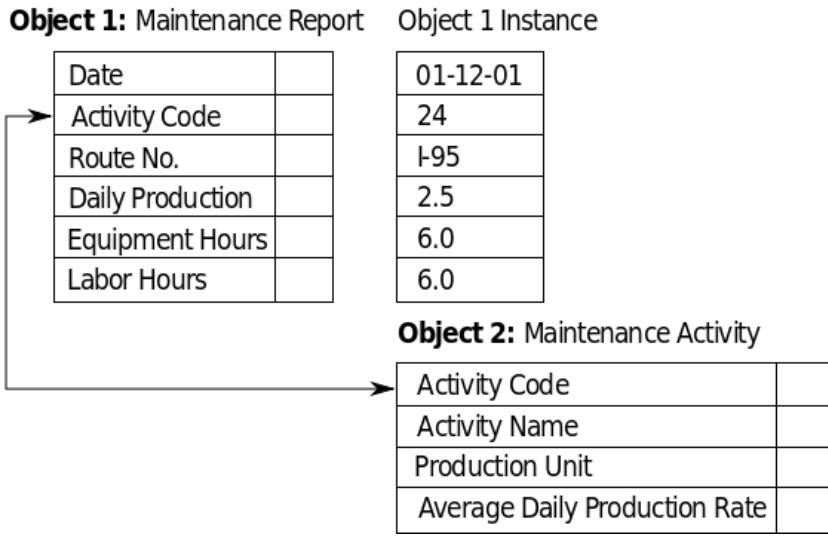
A client/server system is a distributed system in which

- a) Some sites are client sites and some are server sites
- b) All data resides at the server sites
- c) All applications execute at the client sites, and
- d) —the seams show

15. What is OODBMS ?(U)

Object-oriented database management systems (OODBMSs) combine database capabilities with object-oriented programming language capabilities. OODBMSs allow object-oriented programmers to develop the product, store them as objects, and replicate or modify existing objects to make new objects within the OODBMS. Because the database is integrated with the programming language, the programmer can maintain consistency within one environment, in that both the OODBMS and the programming language will use the same model of representation.

Object-Oriented Model



16. What is meant by Data warehousing? (NOV 2014)(U)

A data warehouse is a relational database that is designed for query and analysis rather than for transaction processing. It usually contains historical data derived from transaction data, but it can include data from other sources. It separates analysis workload from transaction workload and enables an organization to consolidate data from several sources.

17. What is meant by data mining? (NOV 2014)(R)

The overall goal of the data mining process is to extract information from a data set and transform it into an understandable structure for further use.

18. What is meant by information retrieval?(R)

An information retrieval process begins when a user enters a query into the system. Queries are formal statements of information needs, for example search strings in web search engines. In information retrieval a query does not uniquely identify a single object in the collection. Instead, several objects may match the query, perhaps with different degrees of relevancy.

19. What is meant by Relevancy ranking? (NOV 2014)(U)

Relevancy ranking is the process of sorting the document results so that those documents which are most likely to be relevant to your query are shown at the top.

20. Define Crawling. (NOV 2014)(R)

Web crawling is the process of search engines combing through web pages in order to properly index them. These —web crawlers| systematically crawl pages and look at the keywords contained on the page, the kind of content, all the links on the page, and then returns that information to the search engine__s server for indexing. Then they follow all the hyperlinks on the website to get to other websites. When a search engine user enters a query, the search engine will go to its index and return the most relevant search results based on the keywords in the search term. Web crawling is an automated process and provides quick, up to date data.

21. What is meant by XML Database?(R)

An XML database is a data persistence software system that allows data to be specified, and sometimes stored, in XML format. These data can then be queried, transformed, exported and returned to a calling system.

22. Define Data classification.(U)

Data classification is the process of organizing data into categories for its most effective and efficient use.

23. What is meant by Access Control? (NOV 2014)(R)

In the fields of physical security and information security, access control is the selective restriction of access to a place or other resource. The act of accessing may mean consuming, entering, or using. Permission to access a resource is called authorization. Locks and login credentials are two analogous mechanisms of access control.

24. What is a threat in dbms?(U)

When the data is been accessed by many people, the chances of data theft increases. In the past, database attacks were prevalent, but were less in number as hackers hacked the network more to show it was possible to hack and not to sell proprietary information. Another reason for database attacks is to gain money selling sensitive information, which includes credit card numbers, Social Security Numbers, etc.

25. What is Cryptography?(U)

Cryptography is a method of storing and transmitting data in a particular form so that only those for whom it is intended can read and process it. The term is most often associated with scrambling plaintext (ordinary text, sometimes referred to as cleartext) into ciphertext (a process called encryption), then back again (known as decryption).

26. Define Statistical database.(R)

A statistical database is a database used for statistical analysis purposes. It is an OLAP (online analytical processing), instead of OLTP (online transaction processing) system. Modern decision, and classical statistical databases are often closer to the relational model than the multidimensional model commonly used in OLAP systems.

27. Define Clustering in DBMS. (NOV 2014)(R)

Clustering, in the context of databases, refers to the ability of several servers or instances to connect to a single database. An instance is the collection of memory and processes that interacts with a database, which is the set of physical files that actually store data.

28. Write about the four types (Star, Snowflake, Galaxy and Fast constellation) of Data warehouse schemas.(DEC2015)(R)

1. STAR SCHEMA:Centralized Fact table connect the one or more denormalized data
2. SNOW FLAKE SCHEMA:Centralized Fact table connect the one or more normalized data
3. STAR FLAKE SCHEMA:One or more centralized fact table connect the single denormalized data
4. Fact Constellation Schema : For each star schema it is possible to construct fact constellation schema(for example by splitting the original star schema into more star schemes each of them describes facts on another level of dimension hierarchies). The fact constellation architecture contains multiple fact tables that share many dimension tables.

29. What is Association rule mining?(R)

Association rules are if/then statements that help uncover relationships between seemingly unrelated data in a relational database or other information repository. An example of an association rule would be "If a customer buys a dozen eggs, he is 80% likely to also purchase milk."

Applications:

- Cross-Marketing
- Basket Data Analysis
- Catalog design

PART B

1. Explain in detail Data Classification in DBMS. **(R)**
2. Discuss in detail Access Control mechanism. **(NOV 2014) (U)**
3. Explain about Distributed Databases. **(NOV 2014) (U)**
4. Discuss about Data warehousing and Data Mining. **(U)**
5. Explain in detail Information retrieval and Relevance Ranking. **(U)**
6. Describe about OODBMS and XML Database. **(R)**
7. Explain in detail Threats and risks in Database Management System. **(R)**
8. Neatly write the K-means algorithm and show the intermediate results in clustering the below given points into two clusters using K-means algorithm. P1(0,0),P2(1,10),P3:(2,20),P4:(1,15),P5:(1000,2000),P6:(1500,1500),P7:(1000,1250). **(U) (Dec2015)**
9. Discuss About The Access Control Mechanisms And Cryptography Methods To Secure the Databases. **(Dec2015) (U)**
10. What is Association rule mining? **(R)(NOV 2016)**
11. Suppose that you have been hired as a consultant to choose a database system for your client's application. For each of the following applications, state what type of database system (relational, persistent programming language based OODB, object relational; do not specify a commercial product) you would recommend. Justify your recommendation
 - i) A computer aided design system for manufacturer of airplanes
 - ii) A system to track contributions made to candidates for public office.
 - iii) an information system to support the making of movies. **(AN)(Nov 2016)**
12. Trace the results of using the Apriori algorithm on grocery store example to support threshold $s=33.34\%$ and confidence threshold $c=60\%$. Show the candidate and frequent itemsets for each database scan. Enumerate all the final frequent itemsets. Also indicate the association rules that are generated and highlight the strong ones, sort them by confidence.

Transaction Id	Items
T1	HotDogs,Buns,Ketchup

T2	HotDogs,Buns
T3	HotDogs,Coke,Chips
T4	Chips,Coke
T5	Chips,Ketchup
T6	HotDogs,Coke,Chips