DEPARTMENT OF INFORMATION TECHNOLOGY
U23ITT33-DATABASE MANAGEMENT SYSTEMS
QUESTION BANK
(2 Marks)

UNIT – I

1. Define database management system.

Database management system (DBMS) is a collection of interrelated data and a set of programs to access those data.

A database management system (DBMS) is a software package with computer programs that control the creation, maintenance, and the use of a database.

2.What is a data model? List the types of data model used. APRIL/MAY-2011

A **database model** is the theoretical foundation of a database and fundamentally determines in which manner data can be stored, organized, and manipulated in a database system. It thereby defines the infrastructure offered by a particular database system. The most popular example of a database model is the relational model.

Types of data model used

- Hierarchical model
- Network model
- Relational model
- Entity-relationship
- Object-relational model
- Object model

3.What are the disadvantages of file processing system? APRIL/MAY-2016

The disadvantages of file processing systems are

a) Data redundancy and inconsistency
b) Difficulty in accessing data
c) Data isolation
d) Integrity problems
e) Atomicity problems
f) Concurrent access anomalies

4. What are the advantages of using a DBMS? NOV/DEC-2014

The advantages of using a DBMS are

a) Controlling redundancy
b) Restricting unauthorized access
c) Providing multiple user interfaces
d) Enforcing integrity constraints.
e) Providing back up and recovery.

5. What are the levels of data abstraction? APRIL/MAY-2014,NOV/DEC-2017

a) Physical level
b) Logical level
c) View level

6.What are the components of storage manager?

The storage manager components include

a) Authorization and integrity manager
b) Transaction manager
c) File manager
d) Buffer manager

7.What is a data dictionary?

A data dictionary is a data structure which stores meta data about the structure of the database ie. The schema of the database.

8. What is embedded SQL? What are its advantages? April/May-2011

Embedded SQL is a method of combining the computing power of a programming language and the database manipulation capabilities of SQL. Embedded SQL statements are SQL statements written in line with the program source code of the host language. The embedded SQLstatements are parsed by an embedded SQL preprocessor and replaced by host-language calls to a code library. Programs written in any number of languages such as: C/C++, COBOL and Fortran.

9. Write short notes on relational model.

The relational model uses a collection of tables to represent both data and the relationships among those data. The relational model is an example of a record based model.

10.Define Relational Algebra.

- Relational Algebra is a procedural query language, Which takes instances of relations as input and yields instances of relations as output.
- It uses operators to perform queries. An operator can be either unary or binary.
- Relational algebra is performed recursively on a relation and intermediate results are also considered as relations
- Operations are Select, Project, Rename, Union, Set Difference, Cartesian Product, Intersection.

11.What is foreign key?

A relation schema r1 derived from an ER schema may include among its attributes the primary key of another relation schema r2.this attribute is called a ***foreign key*** from r1 referencing r2.

12.What are the categories of SQL command? APRIL/MAY-2016

SQL commands are divided in to the following categories:

1. data - definition language
2. data manipulation language
3. Data Query language
4. data control language
5. data administration statements
6. transaction control statements

13. What are aggregate functions? And list the aggregate functions supported by SQL?

Aggregate functions are functions that take a collection of values as input and return a single value. Aggregate functions supported by SQL are

Min, max, Total, Count, Average.

14.Define Super key and candidate key. APRIL/MAY-2014

A super key is a set of one or more attributes that, taken collectively, allow us to identify uniquely an entity in the entity set. For example, the social-security attribute of the entity set customer is sufficient to distinguish one customer entity from another. Similarly, the combination of customer-name and social security is a super key for the entity set customer.

Candidate key: If a relation schema has more than one key, then each is termed as a candidate key. Ex:Dno, Dname.

15.What are the different data types supported by SQL?

The basic datatypes available for attributes include numeric, Character string, Bit string Boolean, date and time. Char(n), varchar(n), int(n), float(n),Boolean

16.What is dynamic SQL?

It is convenient to write a program that can execute different SQL queries or updates dynamically at runtime.It is flexible.

17.Define physical schema and Logical schema.

Physical Schema: It Describes the database design at the physical level, Which is lowest level of abstraction describing how the data are actually stored.

Logical Schema: It describes the database design at the logical level, which describes what data are stored in the database and what relationship among the data.

18. Distinguish Embedded SQL and Dynamic SQL.

Embedded SQL.	Dynamic SQL.
How database will be accessed is predetermined in the embedded SQL statement.	How database will be accessed is determined at run time.
SQL statements are compiled at Compile time.	SQL statements are compiled at runtime.
It is less flexible.	It is more flexible.
Data is distributed uniformly.	Data is distributed non-uniformly.

19. What is the general form of a SQL query?

Select <Column List>

From <Table name>Where <condition>; Ex:Selectename from employee where eno=103;

20. List the different types of database-system users.

There are four different types pf database-system users, differentiated by the way that they expect to interact with the system.

- Application programmers, SophisticatedUsers, Specializedusers,Naïve users**.**

21. Define DDL.APRIL/MAY-2017

A database schema is specified by a set of definitions expressed by a special language called a data-definition language. The result of compilation of DDL statements is a set of tables that is stored in a special file called data dictionary. A data dictionary is a file that contains metadata-that is, data about data. The storage structure and access methods used by the database system are specified by a set of definitions in a special type of DDL called a data storage and definition language.

UNIT-II

1.What is an entity relationship model? APRIL/MAY-2014,APRIL/MAY-2016

The entity relationship model is a collection of basic objects called entities and relationship among those objects. An entity is a thing or object in the real world that is distinguishable from other objects.

2.What are attributes? Give examples.

An entity is represented by a set of attributes. Attributes are descriptive properties possessed by each member of an entity set.

Example: possible attributes of customer entity are customer name, customer id, Customer Street, customer city.

3.Define single valued and multi valued attributes.

Single valued attributes: attributes with a single value for a particular entity are called single valued attributes.

Multi valued attributes : Attributes with a set of value for a particular entity are called multivalued attributes.

4.what are stored and derived attributes?

Stored attributes: The attributes stored in a data base are called stored attributes.

Derived attributes: The attributes that are derived from the stored attributes are called derived attributes

5.What does the cardinality ratio specify?

Mapping cardinalities or cardinality ratios express the number of entities to which another entity can be associated. Mapping cardinalities must be one of the following:

- One to one
- One to many
- Many to one
- Many to many

6.Define weak and strong entity sets.

Weak entity set: entity set that do not have key attribute of their own are called weak entity sets. Strong entity set: Entity set that has a primary key is termed a strong entity set.

7.Define E-R model.NOV/DEC-2013

The entity-relationship data modal is based on perception of a real world that consists of a collection of basic objects, called entities, and of relationships among these objects. The overall logical structure of a database can be expressed graphically by an E-R diagram, which is built up from the following components:

- Rectangles, which represent entity sets.
- Ellipses, which represent attributes
- Diamonds, which represent relationships among entity sets

8. Define Functional Dependency. APRIL/MAY-2015

Functional dependencies are constraints on the set of legal relations. They allow us to express facts about the enterprise that we are modelling with our database.

Syntax: A -> B e.g.) account no -> balance for account table.

9.Define normalization. APRIL/MAY-2014

By decomposition technique we can avoid the Pitfalls in Relational Database Design. This process is termed as normalization

Pitfalls in Relational Database Design

1. Repetition of information
2. Inability to represent certain information

10. List the properties of decomposition.

1. Lossless join
2. Dependency Preservation
3. No repetition of information

11. Define First Normal Form.NOV/DEC-2015

If the Relation R contains only the atomic fields then that Relation R is in first normal form.

E.g.) R = (account no, balance) – first normal form.

12. Define Second Normal Form.

A relation schema R is in 2 NF with respect to a set F of FD's if for all FD's of the form A -> B, where A is contained in R and B is contained in R, and A is a superkey for schema R.

13. Define BCNF.APRIL / MAY-2013

A relation schema R is in BCNF with respect to a set F of FD's if for all FD's of the form A -> B, where A is contained in R and B is contained in R, at least one of the following holds:

1. A -> B is a trivial FD
2. A is a superkey for schema R.

14. Define 3 Normal Form.NOV/DEC -2014

A relation schema R is in 3 NF with respect to a set F of FD's if for all FD's of the form A -> B, where A is contained in R and B is contained in R, at least one of the following holds:

1. A -> B is a trivial FD
2. A is a superkey for schema R.

Each attribute in B – A is contained in a candidate key for R

15. Define Fourth Normal Form.

A relation schema R is in 4NF with respect to a set F of FD's if for all FD's of the form A ->> B (Multi valued Dependency), where A is contained in R and B is contained in R, at least one of the following holds:

1. A ->> B is a trivial MD
2. A is a superkey for schema R.

16. Define 5NF or Join Dependencies.

Let R be a relation schema and R1, R2, …, Rn be a decomposition of R. The join dependency *(R1, R2, …Rn) is used to restrict the set of legal relations to those for which R1,R2,…Rn is a lossless-join decomposition of R. Formally, if R= R1 U R2U …U Rn, we say that a relation r® satisfies the join dependency *(R1, R2, …Rn) if

R = A join dependency is trivial if one of the Ri is R itself.

17. What is meant by lossless-join decomposition? APRIL/MAY-2011

It claim the above decomposition is lossless. How can we decide whether decomposition islossless

1. Let *R* be a relation schema.
2. Let *F* be a set of functional dependencies on *R*.
3. Let R1and R2 form a decomposition of *R*.

4. The decomposition is a lossless-join decomposition of *R* if at least one of the following functional dependencies are in : a. R1∩ R2→ R1 b. R1∩ R2→ R2

18. What is meant by normalization of data?

It is a process of analyzing the given relation schemas based on their Functional Dependencies
(FDs) and primary key to achieve the properties

- Minimizing redundancy
- Minimizing insertion, deletion and updating anomalies

19.What is denormalization?

Denormalization is the process of storing the join of higher normal form relations as a base relation, which is in a lower normal form. It is used to optimize the read performance of a database by adding redundant data or by grouping data.

20.What is domainkey normal form? NOV/DEC-2014

Domain-key normal form(DKNF)is to specify the ultimate normal form that takes into account all possible types of dependencies and constraints.

A relation schema is said to be in DKNF if all constraints and dependencies that should hold on the valid relation states can be enforced simply by enforcing the domain constraints and key constraints on the relation.

UNIT –III

1.What are the ACID properties? APRIL/MAY-2011,MAY/JUNE-2013,NOV/DEC-2017

(*atomicity*, *consistency*, *isolation*, *durability*) is a set of properties that guarantee database transactions are processed reliably. In the context of databases, a single logical operation on the data is called a transaction. For example, a transfer of funds from one bank account to another, even though that might involve multiple changes (such as debiting one account and crediting another), is a single transaction.

2.What are two pitfalls (problem) of lock-based protocols? APRIL/MAY-2011

- Deadlock
- Starvation

3.What is transaction? NOV/DEC-2013

Collections of operations that form a single logical unit of work are called transactions. A Transaction includes one or more database access operations these are include insertion, deletion, modification or retrieval operations .This can placed between boundaries :Begin and End Transaction.

4.What are the states of transaction? NOV/DEC-2015

The states of transaction are
Active
Partially committed
Failed
Aborted
Committed
Terminated

5.What are the two types of serializability?

The two types of serializability is
Conflict serializability
View serializability

6.Define lock.

Lock is the most common used to implement the requirement is to allow a transaction to access a data item only if it is currently holding a lock on that item. Two types: Shared and Exclusive lock.

7.Define the phases of two phase locking protocol. APRIL/MAY-2013

Growing phase: a transaction may obtain locks but not release any lock.
Shrinking phase: a transaction may release locks but may not obtain any new locks.

8.**What is a recovery scheme?**

An integral part of a database system is a recovery scheme that can restore the database to the consistent state that existed before the failure.

9.What is the need for concurrency control?

Lost update problem
Temporary update problem
Incorrect summary problem
Unrepeatable read / Dirty read

10.What isschedule?APRIL/MAY-2017

A schedule S of n transactions T1,T2,...Tn is an ordering of the operations of transaction . Operations from different transaction can be interleaved in the schedule s. For each transaction Ti that participate in the schedule S, the operations of Ti in S must appear in the same order in which they occur in Ti.

11.What is Serializability?NOV/DEC-2014,APRIL/MAY-2014

Serializability of schedule is used to identify which schedules are correct when transaction executions have interleaving of their operations in the schedules. A schedule can be tested for correctness usinga)Conflict serializability b)View serializability

12.What is starvation?

Starvation occurs when a transaction cannot proceed for an indefinite period of time while other transactions in the system continue normally. This may occur if the waiting schemefor locked items are unfair, giving priority to some transaction over others.

13. Define atomicity and Durability?

Either all operations of the transaction are reflected properly in the database or none are. A transaction completes successfully then changes I t has made to the database persist even if there are system failure.

14. Define concurrency control? NOV/DEC-2016

Serializability of schedules generated by concurrently executing transactions can be ensure through one of a variety of mechanisms called concurrency control.

15. What are the tow operation access data in transaction?

Read(x)- transfer data item x from database.
Write(x)- transfer data item x from the local buffer.

16. When the schedule is called serial?

Each serial schedule consists of a sequence of instruction from various transaction where the instruction belonging to one single transaction appear together in that schedule.

17. Define Transaction Isolation.

Transaction-processing systems usually allow multiple transactions to run concurrently. Allowing multiple transactions to update data concurrently causes several complications with consistency of the data. Ensuring consistency in spite of concurrent execution of transactions requires extra work; it is far easier to insist that transactions run serially—that is, one at a time, each starting only after the previous one has completed.

18.What are the transaction isolation levels?

- Serializable
- Repeatable read
- Read committed
- Read uncommitted

19.Define Savepoint.

Transactions can record savepoints, and can be rolled back partially, up to a savepoint. This can be quite useful for deadlock handling, since transactions can be rolled back up to a point that permits release of required locks, and then restarted from that point. Programmers can also use savepoints to undo a transaction partially, and then continue execution; this approach can be useful to handle certain kinds of errors detected during the transaction execution.

20. What are the Recovery Algorithm used in transaction?

- Analysis pass: This pass determines which transactions to undo, which pages were dirty at the time of the crash, and the LSN from which the redo pass should start.
- Redo pass: This pass starts from a position determined during analysis, and performs a redo, repeating history, to bring the database to a state it was in before the crash.
- Undo pass: This pass rolls back all transactions that were incomplete at the time of crash.

UNIT – IV

1. Define RAID.APRIL/MAY-2013

It is collectively called redundant arrays of inexpensive disk, have been proposed to address the performance and reliability issues. RAID are used for their higher reliability and higher data transfer rate. RAID stands of independent, instead of inexpensive.

2. Define file organization.

A file is organized logically as a sequence of records. These records are mapped on to disk blocks. File are prided as a basic construct in operating system.

3. Define Query processing?

Query processing refers to the range of activities involved in extracting data form a database. These activities include translation of queries expressed in high-level database language into expression that can be implemented at the physical level of the file system.

4. What are the steps involved in query processing? NOV/DEC-2017

1. Parsing and translation
2. Optimization
3. Evaluation

5.What is indexing and What are the different kinds of indexing? APRIL/MAY-2016

Indexing is a technique for determining how quickly specific data can be found.

Types:

Binary search style indexing
B-Tree indexing
Inverted list indexing
Memory resident table
Table indexing

6.What is hashing? APRIL/MAY-2015

The conversion of a column's primary key value to a database page number on which the row will be stored. Retrieval operations that specify the key column value use the same hashing algorithm and can locate the row directly. Hashing provide fat retrieval for data that contains a unique key value.

7.Define B+ tree index structure. NOV/DEC-2016

The b+ tree index structure is the most widely used of several index Structures that maintain their efficiency despite insertion and deletion of data. A B+ tree index takes the form of a balanced tree in which ever path from the root of the tree to a leaf of the tree is the same length.

8.Define sparse index?

An index record is created for only some of the record. 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 that or equal to the search-key value.

9. Define Dense Index.

An index record appears for every search-key value in the file. The index record contains the search-key value and pointer to the first data record with that search key value.

10. What are the properties of B trees?

Balance tree. Inner nodes having pointer to its children leaf node having pointer to actual order of the tree is n then each node we can have n-1 elements.

11. What are the organization of records in files?

A file is organized logically as a sequence of records. There are several ways of organizing records in files.

Heap file organization.
Sequential file organization.
Hashing File organization.
Clustreing File organization.

12. Define ordered Indices. APRIL/MAY-2014, NOV/DEC-2017

In an ordered index, index entries are stored and sorted on the search key value. A file may have several indices on different search keys.

- The search key of primary index is usually the primary key.

13. Define File organization.

The database is stored as a collection files. Each file is a sequence of records. A record is a sequence of fields. Records are mapped onto disk block.

There are two approaches:

- Fixed length record.
- Variable length record.

14. What are the advantages of B+ tree index files?

- Automatically reorganize itself with small, local, changes during insertion and deletion.
- Reorganization of entire file is not required to maintain performance.
- If a root is not a leaf it must have at least 2 children.
- B+ tree index is a balanced tree in which every path from the root to a leaf is of same length.

15. What is called Mirroring?

- Data is written redundantly to two identical physical disks that are treated as one logical disk.
- When the data is read, it can be retrieved from the disk with shorter queuing,seek and rotational delays.
- If a disk fails, the other disk is used until the first is repaired.

16.What is bit and block level stripping?

Bit Level: It involves breaking up a byte of data into bits and spreading the bits of different disks.

Block level: It involves distributing blocks of file across disks. Blocks of a file can be striped across disks, giving rise to block level striping.

17. What are the factors to be taken into account when choosing a RAID level?

- Monetary Cost
- Performance: Number of I/O operations per second, and bandwidth during normal operation.
- Performance during Failure
- Performance during rebuild of failed disk including time taken to rebuild failed disk.

18.What is meant by query optimization? APRIL/MAY-2015,APRIL/MAY-2016

The phase that identifies an efficient execution plan for evaluating a query that has the least estimated cost is referred to as query optimization.

19. Draw and Evaluation plan for query optimization.

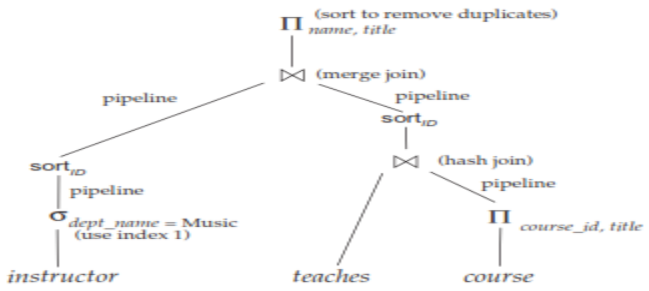

20.How to optimize the query processing using heuristics?

A drawback of cost-based optimization is the cost of optimization itself. Although the cost of query optimization can be reduced by clever algorithms, the number of different evaluation plans for a query can be very large, and finding the optimal plan from this set requires a lot of computational effort. Hence, optimizers use heuristicsto reduce the cost of optimization.

UNIT – V

1. Define distributed database.

- A Distributed database is a collection of multiple, logically interrelated databases distributed over a computer network.
- As the database is distributed, different users can access it without interfering with one another.
- DBMS must periodically synchronize the scattered database to make sure that they all have consistent data.
- Transaction may access data one or more sites.

1. **Define homogeneous and heterogeneous database**.

- In homogeneous, all sites have identical database management system software are aware of one another, and agree to cooperate in processing user's request.
- In heterogeneous database, different site may use different schemas, and different database management system software.

2. Define Replication. What are the advantages?

The system maintains several copies of the relation and stores each replicas at different site .

Advantages

- **Availability**
- **Parallelism**
- **Reduced data transfer**

3. Define Fragmentation.NOV/DEC-2017

The system partitions the relation into several fragments and stores each fragment at a different site. If relation r is fragmented r is divided into a number of fragments r1,r2,...rn.

4. What are the two types of fragmentation?

- Horizontal Fragmentation: It splits the relation by assigning each tuple of r to one or more fragments.
- Vertical Fragmentation: It Splits the relation by decomposing the schema R of relation **r.**

5. Define transparency and give its types. APRIL/MAY-2017

The user of a distributed database system required to know where the data are physically located or how the data can be accessed at the specific local site. This characteristic called transparency.

Types

- **Fragmentation Transparency.**
- **Replication Transparency.**
- **Location Transparency.**

6. Draw the system structure for distributed Transaction.

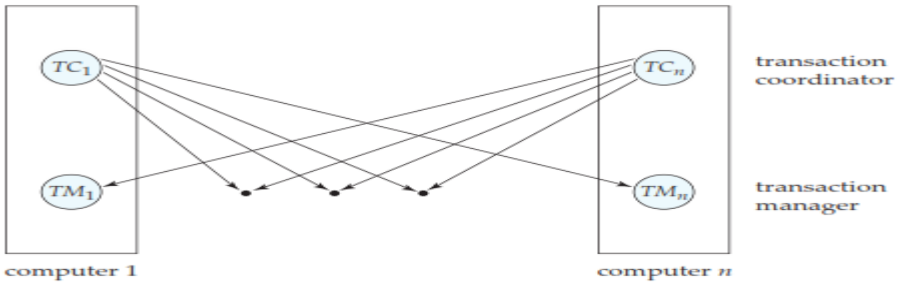

7. Define object oriented database system.

1.Build an object-oriented database system, that is, a database system that natively supports an object-oriented type system, and allows direct access to data from an object-oriented programming language using the native type system of the language.

2.Automatically convert data fromthe native type system of the programming language to a relational representation, and vice versa. Data conversion is specified using an object-relational mapping.

8. What is complex datatype? GiveEx.

Complex datatype is kind of datatype that represent whole data item that hide the subtitles. For example, addresses. While an entire address could be viewed as an atomic data item of type string, this view would hide details such as the street address, city, state, and postal code, which could be of interest to queries. On the other hand, if an addresswere represented by breaking it into the components (street address, city, state, and postal code), writing queries would be more complicated since they would have to mention each field.

9. What are the features of Object-relational database systems?

Each table stores the primary key (which may be inherited from a parent table) and the attributes that are defined locally. Inherited attributes (other than the primary key) do not need to be stored, and can be derived by means of a join with the supertable, based on the primary key.

10 What is object ODMG Model?

The **ODMG object model** is the data model upon which the object definition language (ODL) and object query language (OQL) are based.It is meant to provide a standard data model for object databases, just as SQL describes a standard data model for relational databases.

11. What is Objects and Literals?

Objects and literals are the basic building blocks of the object model. The main difference between the two is that an object has both an object identifier and a **state** (or current value), whereas a literal has a value (state) but no object identifier.

The object state can change over time by modifying the object value. A literal is basically a constant value, possibly having a complex structure, but it does not change

12. What are the five aspects of object?

- An **object** has five aspects: identifier, name, lifetime, structure, and creation.
- The **object identifier** is a unique system-wide identifier all individual objects will have unique names.
- **Lifetime** of an object specifies whether it is a *persistent object*
- The **structure** of an object specifies how the object is constructed by using
- Object **creation** refers to the manner in which an object can be created. This

13. What are the types of Literals?

There are three types of literals: atomic, structured, and collection.

Atomic literals :correspond to the values of basic data types and arepredefined.The basic data types of the object model include long, short.

Structured literals correspond roughly to values that are using the tuple constructor literals include Date, Interval, Time, and Timestamp.

Collection literals specify a literal value that is a collection of objects or values like array, Bag.

14. What is ODL?

The ODL is designed to support the semantic constructs of the ODMG object model and is independent of any particular programming language. Its main use is to create object specifications—that is, classes and interfaces. Hence, ODL is not a full programming language. A user can specify a database schema in ODL independentlyOf any programming language, and then use the specific language bindings to specify how ODL constructs can be mapped to constructs in specific programming languages, such as C++, Smalltalk, and Java.

15. Define OQL.

The object query language OQL is the query language proposed for the ODMG object model. It is designed to work closely with the programming languages for which an ODMG binding is defined, such as C++, Smalltalk, and Java. Hence, an OQL query embedded into one of these programming languages can return objects that match the type system of that language.

16. What is Structured, Semi structured and Unstructured Data in XML database?

- The information stored in databases is known as **structured data** because it is represented in a strict format.
- Additional attributes can be introduced in some of the newer data items at any time.
- **Unstructured data** because there is very limited indication of the type of data.

18.What are the three types of XML documents?NOV/DEC-2017

- Data-centric XML documents.
- Document-centric XML documents.
- XML documents.
-

19.Define information retrieval.

Textual data is unstructured, unlike the rigidly structured data in relational databases. The term **information retrieval** generally refers to the querying of unstructured textual data. Information-retrieval systems have much in common with database systems, in particular, the storage and retrieval of data on secondary storage.

20. Define Web crawler. How to index the web? NOV/DEC-2014

Web crawlers are programs that locate and gather information on the web. They recursively follow hyperlinks present in known documents to find other documents. Crawler's start from an initial set of URLs, which may be created manually. Each of the pages identified by these URLs is fetched from the web.

PART B

UNIT - I

1. Explain the Characteristics and Purpose of DBMS.
2. Explain the various views of Data with examples.
3. Discuss all the Data Models with examples.
4. Explain DBMS Architecture with an example.
5. What is Procedural Query Language? Explain the various operations performed using relational operators.
6. Explain Relational Databases with examples.
7. Draw a Schema and mention the various elements of the schema with examples
8. Explain the different types of keys used in DBMS.
9. Explain Embedded SQL with its applications.
10. Discuss the use of Dynamic SQL in DBMS in detail.
11. Explain the various types of SQL commands. Discuss the features of SQL.

UNIT – II

1. Draw an ER model by taking Hospital management/Banking System/University Database and explain all the relationship sets with cardinality.
2. Explain EER Model with a neat sketch for School Database.
3. Explain various functional dependencies of an ER model with an example
4. What is Normalization? Explain 1NF, 2NF, 3NF, BCNF, 4NF, 5NF with examples.
5. Explain the concept of Non-Loss Decomposition with an example
6. Discuss in detail the steps involved in the ER – to Relational mapping in the process of relational database design
7. Explain with suitable example, the constraints of specialization and generalization in ER data modeling.

UNIT – III

1. Discuss view Serializability and conflict Serializability.
2. Briefly describe two phase locking in concurrency control techniques.
3. Explain the concepts of concurrent execution in Transaction processing system.
4. Brief the ACID properties. Explain Transactions with SQL support for a banking application
5. What is concurrency control? How is it implemented in DBMS? Illustrate with a suitable example.
6. Briefly explain about Two phase commit and three phase commit protocols.
7. What is deadlock? How does it occur? How transactions be written to (i) Avoid deadlock (ii) Guarantee correct execution. Illustrate with suitable example.

8. (i) Narrate the actions that are considered for deadlock detection and the recovery from deadlock (ii) Discuss the properties of a transaction that ensure integrity of data in the database system.

9. Write Short notes on (i)Validation and Snapshot Isolation. (ii) Multiple Granularity locking. 10.What is Recovery? Explain various recovery techniques during transactions in detail

11.Write Short notes on (i) Shadow Paging. (ii) ARIES Algorithm.

UNIT – IV

1. Explain how the RAID systems improve performance and reliability.
2. What is RAID? List the different level in RAID technology and explain its Features with neat sketches.
3. Describe the structure of B+ tree and list the characteristics of a B+tree with indexing operation.
4. Explain the steps involved in Query Processing with a neat sketch.
5. Discuss static hashing and dynamic hashing with examples.
6. Discuss in detail about how records are represented in a file and how to organize them in a file.
7. Explain the cost estimation for query processing.
8. Discuss the selection, sorting and join operations using appropriate algorithms.
9. Explain the various heuristics involved in query optimization.
10. Explain B tree indexing with an example.

UNIT – V

1. Explain in detail about Distributed Databases with a neat sketch.
2. Discuss the query processing and optimization for a transaction in Distributed Databases.
3. Explain NOSQL databases with applications.
4. Explain about the various threats and risks in Database Management Systems.
5. Discuss about various Access Control Mechanisms and Efficient Methods to Secure the Databases.
6. Write Short notes on (i) CAP Theorem. (ii) Key Value Stores.
7. Write Short notes on (i) Column based systems (ii) Graph Databases
8. What is Statistical Database? Explain the security measures involved for having uniqueness of access.
9. Explain the encryption techniques used for securing a database system.
10. Explain Discretionary Access control based on Granting and Revoking Privileges.
11. Explain about the Federation of Distributed Database Systems

12.Explain SQL Injection Methods with an example.