DATABASE MANAGEMENT SYSTEMS

UNIT 1

PART A

1. Define DBMS. **(May 2000, Nov 2009)**

   DBMS is a collection of interrelated data and a set of programs to access those data.

2. What are the four main characteristics that differentiate the database approach from the file processing approach? **(Nov 2019, Nov 2020, May 2021)**

   i. Self describing nature of a database system
   ii. Insulation between programs and data and data abstraction
   iii. Support of multiple views of data
   iv. Sharing of data and multiuser transaction processing

3. What is meant by an instance of the database and schema? **(Nov 2010)**

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

   The overall design of the database is called the database **schema**.

4. Specify the various levels of abstraction in a database. **(Nov 2010, May 2011, May 2014, Nov 2020, May 2021)**

   Physical Level: It is the lowest level of abstraction that describes how the data are actually stored.

   Logical Level: It is the next higher level of abstraction that describes what data are stored in the database and what relationships exist among those data..

   View Level: It is the highest level of abstraction that describes only part of the entire database.

5. What is data model? **(May 2011, Nov 2011, Nov 2012, May 2019, May 2022)**

   Data Model is a collection of conceptual tools for describing data, data relationships, data semantics and consistency constraints.

6. Differentiate a weak entity set from a strong entity set. **(Nov 2009, May 2011)**

   An entity set that may not have sufficient attributes to form a primary key is called as **weak entity set.**

   An entity set that has a primary key is called as **strong entity set.**

7. What is Data manipulation language? **(Nov 2021)**

Data-manipulation language (DML) The SQL DML provides the ability to query information from the database and to insert tuples into, delete tuples from, and modify tuples in the database.

8. What are the various types of keys in the database? **(Nov 2004, Nov 2010, Nov 2020, May 2021)**
   - ✓ Super Key
   - ✓ Candidate Key
   - ✓ Primary Key
   - ✓ Foreign Key
9. Distinguish between primary key and candidate key. **(Nov 2006, May 2008)**

   Primary Key is an attribute which is unique and not null, can identify an instance of the entity set. It is chosen by the database designer.

   Eg.: {cid, name, street, city, pincode}

   Here cid is the primary key.

   Candidate Key is a minimal super key which have more than one attribute that uniquely identify an instance of an entity set.

   Eg.: {cid} and {name, street}
10. What is embedded SQL? **(May 2003, May 2011)**

    When SQL is embedded with any programming language like C, C++, Java, it is called as **embedded SQL**.

PART B

1. What is a data model? Classify the different types of data models. **(Nov 2011, Nov 2012, Nov 2013)**
2. Consider the following schema

   Suppliers (<u>sid</u>: integer, sname: string, address: string)

   Parts (<u>pid</u>: integer, pname: string, color: string)

   Catalog (<u>sid</u>: integer, <u>pid</u>: integer, cost: real)

   The key fields are underlined, and the domain of each field is listed after the field name. Therefore sid is the key for Suppliers, pid is the key for Parts, and sid and pid together form the key for Catalog. The Catalog relation lists the prices charged for parts by Suppliers. Write the following queries in relational algebra:

   i. Find the sids of suppliers who supply some red part or are at 221 Packer Street.
   ii. Find the sids of suppliers who supply some red part red or green part.
   iii. Find the pids of parts supplied by at least two different suppliers. (Nov 2019)

3. Explain the system architecture of a database system with neat block diagram. **(May 2011, May 2012, May 2013, Nov 2019, Nov 2020, May 2021, May 2022)**
4. Consider the schema given in question no. 2 and write the following queries in SQL:
   i. Find the names of suppliers who supply some red part.
   ii. Find the sids of suppliers who supply some red part and some green part.
   iii. Find the pids of parts supplied by at least two different suppliers.
   iv. Find the sids of suppliers who supply every red part. (Nov 2019)
5. What are the several parts of SQL query language? What are the basic built in types used during SQL create statement? State and given example for the basic structure of SQL queries. **(Nov 2020, May 2021)**
6. Consider the following relations :

   EMPLOYEE (ENO, NAME, DATE_BORN, GENDER, DATE_JOINED, DESIGNATION, BASIC_PAY, DEPARTMENT_NUMBER) DEPARTMENT (DEPARTMENT NUMBER, NAME)

   Write SQL queries to perform the following:
   i. List the details of employees belonging to department number 'CSE'.
   ii. List the employee number, employee name, department number and department name of all employees.
   iii. List the department number and number of employees in each department.
   iv. List the details of employees who earn less than the average basic pay of all employees. (Nov 2021)
7. Outline equi-join, left outer join, right outer join and full outer join operations in relational algebra with an example. **(Nov 2021)**
8. Consider the following relations :

   Sailors (sid:integer, sname:string, rating:integer, age:real)

   Boats (*bid:*integer, *bname*:string, *color*:string)

   Reserves (*sid*:integer, *bid*:integer, *day*:date)

   Write the SQL statement for the following queries :
   i. Find all sailors with a rating above 7.
   ii. Find the sids of sailors who have reserved a red boat.
   iii. Find the colors of boats reserved by lubber.
   iv. Find the names of sailors who have reserved at least one boat. (Nov 2020, May 2021)

UNIT 2

PART A

1. For a binary relationship set *R* between entity sets *A* and *B*, list the mapping cardinalities. **(May 2011, Nov 2021)**
   i. One-to-One
   ii. One-to-Many
   iii. Many-to-One
   iv. Many-to-Many
2. State the types of attributes in ER model. **(May 2022)**
   i. Simple Attribute
   ii. Composite Attribute
   iii. Single-valued Attribute
   iv. Multi-valued Attribute
   v. Base Attribute
   vi. Derived Attribute
   vii. Null value Attribute
   viii. Key Attribute
3. What is E-R diagram? **(Nov 2020, May 2021)**

   ER diagram can express the overall logical structure of a database graphically. ER diagrams are simple and clear.
4. Draw the symbols used in an entity relationship diagram for representing an entity set, weak entity set, attribute and multivalued attribute. **(Nov 2021)**

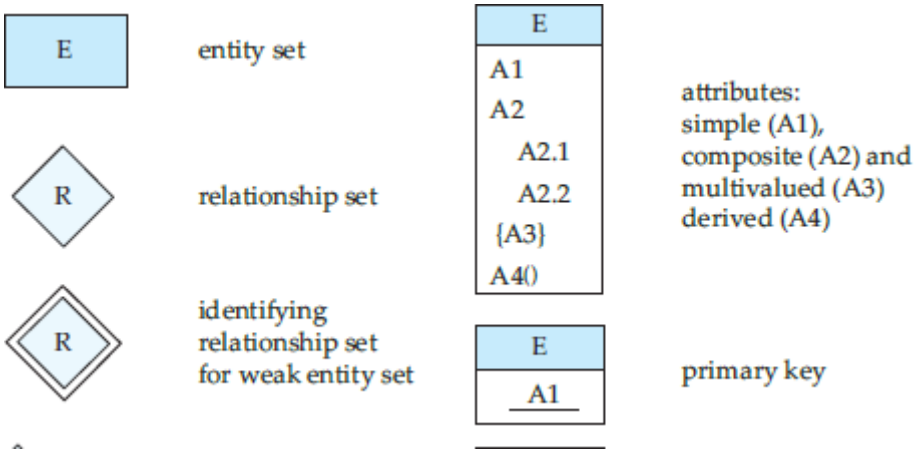

5. Define functional dependency. **(Nov 2009, Nov 2010, Nov 2011, Nov 2012, Nov 2013, May 2014, Nov 2020, May 2021)**

Functional dependency requires that the value for a certain set of attributes determines uniquely the value for another set of attributes. In a given relation R, X and Y are attributes. Attribute Y is functionally dependent on attribute X if each value of X determines exactly one value of Y which is represented as X -> Y. ie) X determines Y or Y is functionally dependent on X. X -> Y does not imply y -> X.

6. Give the properties of decomposition. **(May 2019)**
   i. Lossless join
   ii. Dependency Preservation
   iii. Repetition of information
7. State the various pitfalls in relational database design. **(Nov 2002, Nov 2009, May 2010)**

   A bad design may lead to
   i. Repetition of information- that leads to insertion, deletion, updation problems.
   ii. Inability to represent certain information.
8. What is normalization? Write the types of all normal forms. **(May 2013, May 2014)**

   Normalization is a process of analyzing the given relation schema based on the functional dependencies and primary keys to achieve the desirable properties of
   - ✓ Minimizing redundancy
   - ✓ Minimizing insert, delete and update anomalies.
   - ✓ Improving consistency

   Types
   - First Normal Form
   - Second Normal Form
   - Third Normal Form
   - Boyce Codd Normal Form
   - Fourth Normal Form
   - Fifth Normal Form
9. Prove that any relation schema with two attributes is in BCNF **(Nov 2011)**

   Consider a relation schema R = {A, B} with two attributes. The only possible non-trivial FDs are {A} -> {B} and {B} -> {A}. There are four possible cases:
   i. No FD holds in R. In this case, the key is {A, B} and the relation satisfies BCNF.
   ii. Only {A} -> {B} holds. In this case, the key is {A} and the relation satisfies BCNF.
   iii. Only {B} -> {A} holds In this case, the key is {B} and the relation satisfies BCNF.

iv. Both {A} -> {B} and {B} -> {A} hold. In this case, there are two keys {A} and {B} and the relation satisfies BCNF.

10. Consider a relation R = {A, B, C, D, E} with the following dependencies:

AB -> C CD -> E DE -> B

Is AB a candidate key of this relation? If not, is ABD? Explain your answer. **(May 2010)**

No, AB+ = {A, B, C}, a proper subset of {A, B, C, D, E}

Yes, ABD+ = {A, B, C, D, E}

PART B

1. Construct an ER diagram for a Car insurance company whose customers own one or more cars each. Each car has associated with it zero to any number of recorded accidents. Convert the designed ER into a relational design. **(Nov 2018, Nov 2020, Nov 2021, May 2022)**
2. A university registrar's office maintains data about the following entities:
   (1) courses, including number, title, credits, syllabus, and prerequisites;
   (2) course offerings, including course number, year, semester. section number, instructor, timings, and classroom;
   (3) students, including student-id, name, and program; and
   (4) instructors, including identification number, name, department, and title. Further, the enrollment of students in courses and grades awarded to students in each course they are enrolled for must be appropriately modeled. Model an entity relationship diagram for the registrar's office. **(May 2018, Nov 2021)**

3. Consider the following information about a university database:
   i. Professors have an SSN, a name, an age, a rank, and a research specialty.
   ii. Projects have a project number, a sponsor name (e.g., NSF), a starting date, an ending date, and a budget.
   iii. Graduate students have an SSN, a name, an age, and a degree program (e.g., M.S. or Ph.D.).
   iv. Each project is managed by one professor (known as the project's principal investigator).
   v. Each project is worked on by one or more professors (known as the project's co-investigators).
   vi. Professors can manage and/or work on multiple projects.
   vii. Each project is worked on by one or more graduate students (known as the project's research assistants).
   viii. When graduate students work on a project, a professor must supervise their work on the project. Graduate students can work on multiple projects, in which case they will have a (potentially different) supervisor for each one.

ix. Departments have a department number, a department name, and a main office.

x. Departments have a professor (known as the chairman) who runs the department.

xi. Professors work in one or more departments, and for each department that they work in, a time percentage is associated with their job.

xii. Graduate students have one major department in which they are working on their degree.

xiii. Each graduate student has another, more senior graduate student (known as a student advisor) who advises him or her on what courses to take.

Design and draw an ER diagram that captures the information about the university. Use only the basic ER model here, that is, entities, relationships, and attributes. Be sure to indicate any key and participation constraints. **(Nov 2019)**

4. i) State the Armstrong axioms.

   (ii) Define BCNF and justify a relation R with two attributes is in BCNF. **(Nov 2021)**

5. Explain the following terms briefly : attribute, domain, entity relationship, entity set, relationship set, one-to-many relationship, many-to-many relationship, participation constraint, overlap constraint, covering constraint, weak entity set, aggregation and role indicator. **(Nov 2020, May 2021)**

6. Discuss in detail the steps involved in ER-to-Relational Mapping in the process of relational database design. **(Nov 2019)**

7. i) Write an algorithm to find closure of functional dependents. (5)

   ii) Compute the closure of the following set F of functional dependencies for relation schema R = {A, B, C, D, E}.

   A -> BC

   CD -> E

   B -> D

   E -> A

   List the candidate keys for R. **(Nov 2020, May 2021)**

8. (i) Consider the following relation:

   R(U, V, W, X, Y, Z)

   All the attributes of relation R are atomic. The primary key of relation R is combination of U and V. The following functional dependencies hold:

   $UV \rightarrow W$

   $U \rightarrow X$

   $V \rightarrow Y$

$Y \rightarrow Z$

Is relation R normalized? If yes, justify the relation is normalized. If no, state reasons and normalize the same.

(ii) Consider the following relation:

STUDENT (ROLLNUMBER, NAME, DOB, GENDER, BRANCH_CODE, BRANCH_NAME)

The primary key of the relation is ROLLNUMBER. The following functional dependencies hold:

ROLLNUMBER → NAME, DOB, GENDER, BRANCH_CODE

BRANCH CODE → BRANCH_NAME

Is relation STUDENT normalized? If yes, justify the relation is normalized. If no, state reasons and normalize the same. **(Nov 2021)**

9. Consider the table User_Personal and answer to queries given below.

UserID	**U_email**	**Fname**	**Lname**	**City**	**State**	**Zip**
MA12	Mani@ymail.com	MANISH	JAIN	BILASPUR	CHATISGARH	458991
PO45	Pooja.g@gmail.co	POOJA	MAGG	KACCH	GUJRAT	832212
LA33	Lavle98@jj.com	LAVLEEN	DHALLA	RAIPUR	CHATISGARH	853578
CH99	Cheki9j@ih.com	CHIMAL	BEDI	TRICHY	TAMIL NADU	632011
DA74	Danu58@g.com	DANY	JAMES	TRICHY	TAMIL NADU	645018

i. Is this table in First Normal Form – 1NF? Justify and normalize to 1NF if needed.

ii. Is this table in Second Normal Form – 2NF? Justify and normalize to 2NF if needed.

iii. Is User_Personal in Third Normal Form - 3NF? Justify and normalize to 1NF if needed. (Nov **2019)**

10. What is normalization? Justify the need for normalization with example. **(Nov 2009, May 2011, Nov 2011, May 2012, Nov 2012, May 2019, Nov 2019)**

UNIT 3

PART A

1. List down the various states of a transaction. **(May 2011, May 2019)**

The states of a transaction are

i. Active

ii. Partially Committed

iii. Failed

iv. Aborted

v. Committed

2. List the ACID properties and its usefulness. **(Nov 2006, Nov 2007, Nov 2008, Nov 2009, May 2010, May 2011, May 2012, Nov 2012, May 2013, Nov 2020, May 2021)**

   A – Atomicity

   C – Consistency

   I – Isolation

   D - Durability

3. What are serializable schedules? **(Nov 2021)**

   Consistency of the database ensured under concurrent execution by making sure that any schedule that is executed has the same effect as a schedule that could have occurred without any concurrent execution. That is, the schedule should, in some sense, be equivalent to a serial schedule. Such schedules are called **serializable schedules**.

4. Write notes on starvation. **(May 2011)**

   If a transaction never gets a lock that is currently hold by some other transaction then the transaction is said to be starved or **starvation** or livelock.

5. What do you mean by concurrency control? **(Nov 2009)**

   The database system must control the interaction among the concurrent transactions to prevent them from destroying the consistency of the database. This mechanism is called **concurrency control.**

6. List the commonly used concurrency control techniques. **(Nov 2011)**

   i. Lock based protocol
   ii. Timestamp based protocol
   iii. Validation based protocol
   iv. Multiple granularity
   v. Serializability
   vi. Deadlock handling

7. What benefit does strict two-phase locking provide? What are the disadvantages of it? **(Nov 2020, May 2021)**

   Benefits

   Ensures conflict serializability

   Prevents any other transaction from reading the uncommitted data

   Recovery is very easy.

   Disadvantages

   Does not ensure freedom from deadlock

   Cascading rollback may occur.

8. Define deadlock. **(May 2008)**

   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.

9. Name the four conditions for deadlock. **(Nov 2021)**

   The four necessary **conditions** for a deadlock situation are mutual exclusion, no preemption, hold and wait and circular set.

10. List the SQL statements used for transaction control. **(May 2009, Nov 2011)**

    Commit, Rollback, Savepoint, Set Transaction

PART B

1. What is a transaction? Draw the state diagram corresponding to a transaction and present an outline of the same. **(Nov 2019, Nov 2021, May 2022)**
2. Discuss in detail about the ACID properties of a transaction. **(Nov 2019, Nov 2021, May 2022)**
3. Discuss in detail about the testing of serializability. **(May 2019)**
4. Discuss elaborately the two phase locking protocol that ensures serializability. **(Nov 2019, May 2022)**
5. Narrate the actions that are considered for deadlock detection and the recovery from deadlock. **(Nov 2019)**
6. What are the two approaches of deadlock prevention? Explain in detail with suitable example. **(Nov 2020, May 2021)**
7. Explain deferred and immediate modification versions of the log based recovery scheme. **(May 2019)**
8. (i) Outline the isolation levels specified by the SQL standard with an example.
   (ii) Outline the SQL statements used for transaction control. **(Nov 2021)**
9. State and explain the transaction isolation level. **(Nov 2020, May 2021)**

UNIT 4

PART A

1. List the different levels in RAID technology and specify its feature. **(Nov 2010, May 2013)**
   i. RAID level 0 – Non redundancy block striping
   ii. RAID level 1 – Disk mirroring with block striping
   iii. RAID level 2 – Memory style error correcting code with bit striping
   iv. RAID level 3 – Bit interleaved parity
   v. RAID level 4 – Block interleaved parity
   vi. RAID level 5 – Block interleaved distributed parity

vii. RAID level 6 – P + Q redundancy scheme

2. What are the advantages of file organization? **(Nov 1999, May 2000)**
   i. The ability to access any location in a storage medium without having to access prior location.
   ii. Provide faster access
3. How do you organize the records in files? **(Nov 2006, Nov 2009, Nov 2011, May 2014, Nov 2021)**
   i. Physical database organization
   ii. Heap file organization
   iii. Sequential file organization
   iv. Hashing file organization
   v. Clustering file organization
   vi. Indexing
      - Ordered Indices
         Dense Index
         Sparse Index
         Multilevel Index
      - Hash Indices
4. What is ordered index? Give example. **(May 2009, Nov 2011)**

   Files ordered sequentially based on some search key is known as **ordered index**.

   Eg.: Primary Index
5. What are the factors needed to evaluate the technique of ordered indexing and hashing? **(Nov 2020, May 2021)**
   i. Access types
   ii. Access time
   iii. Insertion time
   iv. Deletion time
   v. Space Overhead
6. Differentiate dense index and sparse index. **(May 2007, May 2008, May 2009, May 2010, Nov 2010, Nov 2011, May 2019)**

   Dense index – An index record appears for every search - key value in the file.

   Insertion and deletion is very difficult.

   Sparse index - An index record appears for only some of the search - key values.

   It requires less space and less maintenance overhead for insertion and deletion.
7. List out the mechanisms to avoid collision during hashing. **(Nov 2016)**

Open addressing – Linear probing

Chaining

Multiple hashing

8. What is Metadata? **(Nov 2021)**

   Metadata is "data that provides information about other data", but not the content of the data, such as the text of a message or the image itself.

9. What are the steps involved in query processing? **(May 2011)**
   i. Parsing and translation
   ii. Optimization
   iii. Evaluation

10. Which cost components contribute to query execution? **(Nov 2019)**
    i. Access cost to secondary storage.
    ii. Disk storage cost.
    iii. Computation cost.
    iv. Memory usage cost
    v. Communication cost.

PART B

1. What is RAID? List the different levels in RAID technology and explain its features. **(Nov 2010, May 2011, Nov 2011, May 2019, Nov 2019, May 2022)**
2. Describe the procedure for index update for single level indices with example. **(Nov 2020, May 2021)**
3. a) What is an index record? Outline dense index and sparse index with an example.
   b) Outline the factors used to evaluate indexing and hashing techniques. **(Nov 2021)**
4. Describe the structure of B^+ tree and give the algorithm for search in the B^+ tree with example. **(May 2003, Nov 2003, Nov 2006, Nov 2007, May 2008, Nov 2010, May 2011, May 2019)**
5. Construct a B^+ tree for the following set of key values (2, 3, 5, 7, 11, 17, 19, 23, 29, 31). Assume that the tree is initially empty and values are added in ascending order. Construct B^+ tree for the cases where the number of pointers that will fit in one node is as follows:
   i. Four
   ii. Six
   iii. Eight **(Nov 2020, May 2021)**
6. Explain about static and dynamic hashing with example. **(May 2002, May 2003, Nov 2006, May 2007, May 2008, Nov 2008, May 2009, Nov 2009, Nov 2010, Nov 2011, Nov 2020, May 2021)**

7. What is query processing? Outline the steps involved in processing a query with a diagram. **(May 2002, May 2003, Nov 2004, May 2007, Nov 2007, May 2008, Nov 2008, Nov 2011, Nov 2019, Nov 2021)**
8. With simple algorithms, explain the computing of Nested-loop join and Block nested-loop join. **(Nov 2019)**

UNIT 5

PART A

1. Define a distributed database management system. **(May 2018)**

   A **distributed database system** consists of loosely coupled sites (computer) that share no physical components and each site is associated a database system.
2. Specify the types of fragmentation in distributed databases. **(May 2022)**
   i. Horizontal Fragmentation
   ii. Vertical Fragmentation
3. Outline the motivation of Replication in a distributed database environment. **(Nov 2021)**

   The system maintains several identical replicas (copies) of the relation, and stores each replica at a different site. The alternative to **replication** is to store only one copy of relation *r*.
4. State the storage device hierarchy. **(May 2011, Nov 2020, May 2021)**
   i. Magnetic tapes
   ii. Optical disk
   iii. Magnetic disk
   iv. Flash memory
   v. Main memory
   vi. Cache memory
5. Mention two features of multimedia databases. **(May 2019)**

   i) The **multimedia database** systems are to be used when it is required to administrate huge amounts of multimedia data objects of different types of media (optical storage, video, tapes, audio records, etc.) so that they can be used (that is, efficiently accessed and searched) for as many applications as needed.

   ii) The objects of Multimedia Data are: text, images, graphics, sound recordings, video recordings, signals, etc. that are digitalized and stored.
6. Compare sequential access devices versus random access devices with an example. **(May 2019)**

Sequential access devices	Random access devices

Must be accessed from the beginning.	It is possible to read data from any location.
Eg:-Tape storage	Eg:-Disk storage
Data is faster	Access to data is much slower
Cheaper than disk	Expensive when compared with disk

7. List information types of documents necessary for relevance ranking of documents in IR. **(Nov 2019)**

   Unstructured data, images, audio recordings, video – strips, maps.

8. What one could understand from allocation schema? **(Nov 2019)**

   An **allocation schema** describes the allocation of fragments to sites of the DDBS, hence, it is a mapping that specifies for each fragment the sites at which it is stored. If a fragment is stored at more than one site, it is said to be replicated.

9. What are Ontologies? **(Nov 2021)**

   Ontologies are hierarchical structures that reflect relationships between concepts.

   The most common relationship is the is-a relationship; for example, a leopard *is-a* mammal, and a mammal *is-a* animal.

   Ontologies have been defined for specific areas to deal with terminology relevant to those areas. For example, ontologies have been created to standardize terms used in businesses; this is an important step in building a standard infrastructure for handling order processing and other inter-organization flow of data.

   It is also possible to build ontologies that link multiple languages. For example, WordNets have been built for different languages, and common concepts between languages can be linked to each other. Such a system can be used for translation of text. In the context of information retrieval, a multilingual ontology can be used to implement a concept-based search across documents in multiple languages.

10. What is the difference between a false positive and false drop? **(Nov 2020, May 2021)**

    Each keyword may be contained in a large number of documents; hence, a compact representation is critical to keep space usage of the index low. Thus, the sets of documents for a keyword are maintained in a compressed form. So that storage space is saved, the index is sometimes stored such that the retrieval is approximate; a few relevant documents may not be retrieved (called a **false drop** or **false negative**), or a few irrelevant documents may be retrieved (called a **false positive**).

    False drop or **false negative**

i. False negatives may occur when documents are ranked, as a result of relevant documents receiving a low ranking.
ii. False negative depends on how many documents are examined.
iii. Measure the recall as a function of the number of documents fetched.

False positive

i. False positives may occur because irrelevant documents get higher rankings than relevant documents.
ii. This too depends on how many documents are examined.
iii. One option is to measure precision as a function of number of documents fetched.

PART B

1. Explain the architecture of distributed databases. **(May 2016, May 2017, May 2022)**
2. a) Outline the two basic types of fragmentation and replication in a distributed database environment with an example.

   b) Compare the features of Object based and Object-relational databases. **(Nov 2021)**
3. State and explain the persistent programming languages. (Nov 2020, May 2021)
4. Present an outline of Document Type Declaration, XML schema, path expressions and XQuery language. **(Nov 2021, May 2022)**
5. Explain in detail about the deductive DB and spatial DB. **(May 2019)**
6. a) Illustrate the usage of OQL, the DMG's query language**.**

   b) Brief on the methods to store XML documents. **(Nov 2019)**
7. Explain in detail about the deductive DB and spatial DB. **(May 2019)**
8. How effectiveness of retrieval is measured? Discuss. **(Nov 2019)**
9. Give the DTD or XML Schema for an XML representation of the following nested-relational schema:

   Emp = (ename, ChildrenSet setof(Children), SkillsSet setof(Skills))

   Children = (name, Birthday)

   Birthday = (day, month, year)

   Skills = (type, ExamsSet setof(Exams))

   Exams = (year, city) **(Nov 2016)**
10. 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.

a. A computer-aided design system for a manufacturer of airplanes.
b. A system to track contributions made to candidates for public office.
c. An information system to support the making of movies. **(Nov 2016)**

11. A car-rental company maintains a vehicle database for all vehicles in its current fleet. For all vehicles, it includes the vehicle identification number, license number, manufacturer, model, date of purchase, and color. Special data are included for certain types of vehicles:
    - Trucks: cargo capacity
    - Sports cars: horsepower, renter age requirement
    - Vans: number of passengers
    - Off-road vehicles: ground clearance, drivetrain (four- or two-wheel drive)

    Construct an object-oriented database schema definition for this database. Use inheritance where appropriate. **(Nov 2015)**