

CS3492 Database Management Systems Study Materials

UNIT I (RELATIONAL DATABASES)

2 Marks:

1. Question: What is a database system?

Answer: A database system is a software application that allows users to store, organize, and manage data.

2. Question: What are the different views of data in a database system?

Answer: The different views of data in a database system are:

Physical view: This view shows how the data is actually stored on the disk.

Logical view: This view shows how the data is organized into tables and relationships.

User view: This view shows how the data is presented to the user.

3. Question: What are the different data models in a database system?

Answer: The different data models in a database system are:

Hierarchical model: This model organizes data into a tree-like structure.

Network model: This model organizes data into a graph-like structure.

Relational model: This model organizes data into tables.

4. Question: What are the different components of a database system architecture?

Answer: The different components of a database system architecture are:

Database engine: This component manages the data in the database.

Query processor: This component processes user queries and returns the results.

Transaction manager: This component ensures that transactions are completed successfully.

Question: What is a relational database?

Answer: A relational database is a database that uses the relational model to organize data.

5. Question: What is the relational model?

Answer: The relational model is a data model that organizes data into tables. Each table has a number of columns, where each column represents a different attribute of the data.

6. Question: What are keys in a relational database?

Answer: Keys in a relational database are unique identifiers for rows in a table. The most common types of keys are primary keys, foreign keys, and candidate keys.

7. Question: What is relational algebra?

Answer: Relational algebra is a set of operations that can be used to manipulate relational data.

8. Question: What are some of the advanced features of SQL?

Answer: Some of the advanced features of SQL include:

Subqueries: Subqueries are queries that are embedded within other queries.

Views: Views are virtual tables that are based on the results of a query.

Stored procedures: Stored procedures are pre-compiled SQL statements that can be executed multiple times.

Functions: Functions are user-defined SQL statements that can be used to perform complex calculations.

Question: What is embedded SQL?

Answer: Embedded SQL is a programming language that allows users to embed SQL statements within their code.

9. Question: What is dynamic SQL?

Answer: Dynamic SQL is a type of SQL that allows users to generate and execute SQL statements at runtime.

Detail Question:

1. What are the advantages and disadvantages of using relational databases?
2. What are the different types of database systems?
3. What are the different components of a database system architecture?
4. What are the different types of SQL statements?
5. What are the different types of database indexes?
6. What are the different types of database transactions?
7. What are the different types of database locking mechanisms?
8. What are the different types of database recovery mechanisms?
9. What are the different types of database performance tuning techniques?

UNIT II (DATABASE DESIGN)

2Marks:

1. Question: What is the Entity-Relationship (ER) model?

Answer: The ER model is a data modeling technique that is used to design relational databases. It uses a graphical notation to represent entities, relationships, and attributes.

2. Question: What are the different types of ER diagrams?

Answer: The different types of ER diagrams are:

Entity-relationship diagrams (ERDs): ERDs are used to model the entities and relationships in a database.

Enhanced entity-relationship diagrams (EERDs): EERDs are ERDs that include additional features, such as specialization and generalization.

3. Question: What is ER-to-relational mapping?

Answer: ER-to-relational mapping is the process of converting an ER diagram into a relational database schema.

4. Question: What are functional dependencies?

Answer: Functional dependencies are constraints between attributes in a table. They specify that a change in the value of one attribute will always result in a change in the value of another attribute.

5. Question: What is non-loss decomposition?

Answer: Non-loss decomposition is the process of breaking a table down into smaller tables without losing any information.

6. Question: What are the first, second, and third normal forms?

Answer: The first, second, and third normal forms are normalization criteria that are used to ensure that a relational database schema is well-designed.

7. Question: What is dependency preservation?

Answer: Dependency preservation is the principle that a normalization process should not remove any valid dependencies from a relational database schema.

8. Question: What is Boyce/Codd normal form?

Answer: Boyce/Codd normal form (BCNF) is a normalization criterion that is stricter than the first, second, and third normal forms. It ensures that all dependencies in a relational database schema are fully-functional dependencies.

9. Question: What are multivalued dependencies and fourth normal form?

Answer: Multivalued dependencies are a type of dependency that can occur in a relational database schema. Fourth normal form (4NF) is a normalization criterion that is stricter than BCNF and takes multivalued dependencies into account.

10. Question: What are join dependencies and fifth normal form?

Answer: Join dependencies are a type of dependency that can occur in a relational database schema. Fifth normal form (5NF) is a normalization criterion that is stricter than 4NF and takes join dependencies into account.

Detail Questions:

1. What are the advantages and disadvantages of using the ER model?
2. What are the different types of ER relationships?
3. What are the different steps involved in ER-to-relational mapping?
4. What are the different types of functional dependencies?
5. What are the different types of non-loss decomposition?
6. How do you normalize a relational database schema to the first, second, and third normal forms?

7. How do you ensure dependency preservation?
8. How do you normalize a relational database schema to BCNF?
9. How do you normalize a relational database schema to 4NF?
10. How do you normalize a relational database schema to 5NF?

UNIT III (TRANSACTIONS)

2Marks:

1. Question: What is a transaction?

Answer: A transaction is a logical unit of work that consists of one or more database operations.

2. Question: What are the ACID properties of transactions?

Answer: The ACID properties of transactions are:

Atomicity: All operations in a transaction must either succeed or fail as a whole.

Consistency: A transaction must leave the database in a consistent state.

Isolation: Transactions must be isolated from each other, so that one transaction cannot see the uncommitted changes of another transaction.

Durability: Once a transaction is committed, its changes must be permanent.

3. Question: What are schedules?

Answer: Schedules are the sequences in which the operations of concurrent transactions are executed.

4. Question: What is serializability?

Answer: Serializability is a property of schedules that ensures that the results of concurrent transactions are the same as the results of executing the transactions serially in some order.

5. Question: What is transaction support in SQL?

Answer: SQL provides a number of features for supporting transactions, such as the BEGIN TRANSACTION, COMMIT, and ROLLBACK statements.

Detail Questions:

1. What are the different types of concurrency control mechanisms?
2. What are the advantages and disadvantages of each type of concurrency control mechanism?
3. How does two-phase locking work?
4. How does timestamp ordering work?
5. How does multiversion concurrency control work?
6. How does validation and snapshot isolation work?
7. How does multiple granularity locking work?
8. How can deadlocks be handled in a database system?
9. What are the different types of database recovery mechanisms?
10. How does shadow paging work?
11. How does the ARIES algorithm work?

UNIT IV (IMPLEMENTATION TECHNIQUES)

2Marks:

1. Question: What is RAID?

Answer: RAID (Redundant Array of Independent Disks) is a technology that combines multiple disk drives into a single logical unit. RAID can be used to improve the performance and reliability of a storage system.

2. Question: What are the different types of RAID levels?

Answer: There are many different types of RAID levels, each with its own advantages and disadvantages. Some common RAID levels include:

RAID 0: RAID 0 stripes data across multiple disks, which can improve performance but reduces reliability.

RAID 1: RAID 1 mirrors data across multiple disks, which provides high reliability but reduces performance.

RAID 5: RAID 5 stripes data across multiple disks and stores parity information on a separate disk, which provides good performance and reliability.

RAID 6: RAID 6 stripes data across multiple disks and stores two copies of parity information on separate disks, which provides even better reliability than RAID 5.

3. Question: What are the different types of file organization?

Answer: The different types of file organization include:

Sequential file organization: Sequential file organization stores data records in a sequential order.

Indexed sequential file organization: Indexed sequential file organization is a variation of sequential file organization that includes an index to improve the performance of search operations.

Hash file organization: Hash file organization stores data records in buckets based on their hash values.

B-tree file organization: B-tree file organization stores data records in a balanced tree structure.

Question: What are the different types of data dictionary storage?

Answer: The different types of data dictionary storage include:

In-memory data dictionary storage: In-memory data dictionary storage stores the data dictionary in memory, which provides the best performance but uses more memory resources.

On-disk data dictionary storage: On-disk data dictionary storage stores the data dictionary on disk, which uses less memory resources but reduces performance.

Hybrid data dictionary storage: Hybrid data dictionary storage stores the most frequently accessed parts of the data dictionary in memory and the rest of the data dictionary on disk.

4. Question: What is column oriented storage?

Answer: Column oriented storage stores data in columns instead of rows. This can improve the performance of analytical queries that aggregate data across multiple columns.

5. Question: What are the different types of indexing?

Answer: The different types of indexing include:

B-tree index: B-tree index is a balanced tree structure that is used to index data records.

Hash index: Hash index stores data records in buckets based on their hash values.

Bitmap index: Bitmap index stores data records in a bit array, where each bit represents a different value.

5. Question: What is query processing?

Answer: Query processing is the process of executing a SQL query. It involves parsing the query, optimizing the query plan, and executing the query plan.

6. Question: What is query optimization?

Answer: Query optimization is the process of choosing the best query plan for a given query. The goal of query optimization is to minimize the execution time of the query.

Detail Questions:

1. What are the different factors to consider when choosing a RAID level?
2. What are the advantages and disadvantages of different types of file organization?
3. How do you design a data dictionary for a database system?
4. What are the benefits of using column oriented storage?
5. How do you choose the right type of index for a given column?
6. What are the different steps involved in query processing?
7. What are the different query optimization techniques?

UNIT V (ADVANCED TOPICS)

2MARKS:

1. Question: What is a distributed database?

Answer: A distributed database is a database that is spread across multiple computers.

2. Question: What are the different types of distributed database architectures?

Answer: The different types of distributed database architectures are:

Client-server architecture: Client-server architecture is a distributed database architecture where the database server is located on a separate computer from the database client.

Peer-to-peer architecture: Peer-to-peer architecture is a distributed database architecture where each node in the network is both a client and a server.

3. Question: What are the different ways to store data in a distributed database?

Answer: The different ways to store data in a distributed database are:

Replication: Replication is a data storage technique where the same data is stored on multiple nodes in the network.

Fragmentation: Fragmentation is a data storage technique where the data is divided into smaller pieces and stored on different nodes in the network.

4. Question: What are the challenges of transaction processing in a distributed database?

Answer: The challenges of transaction processing in a distributed database include:

Ensuring the atomicity of transactions: Atomicity ensures that all operations in a transaction are either committed or aborted as a whole.

Ensuring the consistency of transactions: Consistency ensures that a transaction leaves the database in a consistent state.

Ensuring the isolation of transactions: Isolation ensures that transactions do not interfere with each other.

5. Question: What are the different NoSQL database models?

Answer: The different NoSQL database models are:

Document-based databases: Document-based databases store data in JSON or XML documents.

Key-value stores: Key-value stores store data in key-value pairs.

Column-based databases: Column-based databases store data in columns instead of rows.

Graph databases: Graph databases store data in graphs, where the nodes represent entities and the edges represent relationships between entities.

6. Question: What is the CAP theorem?

Answer: The CAP theorem states that a distributed system cannot simultaneously guarantee consistency, availability, and partition tolerance.

7. Question: What are some of the security issues associated with databases?

Answer: Some of the security issues associated with databases include:

Unauthorized access to data

Data breaches

SQL injection attacks

8. Question: What are some of the ways to secure a database?

Answer: Some of the ways to secure a database include:

Access control based on privileges

Role-based access control

SQL injection prevention

Statistical database security

Flow control

Encryption and public key infrastructures

Detail Questions:

1. What are the advantages and disadvantages of using distributed databases?
2. What are the different types of distributed database queries?
3. How do you optimize distributed database queries?
4. What are the different types of NoSQL databases?
5. What are the advantages and disadvantages of using different types of NoSQL databases?
6. How do you choose the right type of NoSQL database for a given application?
7. What are the different database security threats?
8. How do you implement access control based on privileges?
9. How do you implement role-based access control?
10. How do you prevent SQL injection attacks?
11. What are the different statistical database security techniques?
12. What are the different flow control mechanisms for databases?
13. How do you implement encryption and public key infrastructures for database security?

