

COMPUTER SCIENCE

CORE-6 (DATABASE SYSTEM)

FILL IN THE BLANKS

1. A database is a _____ of related data. Answer: Collection
2. The three main types of database users are _____, _____, and _____. Answer: End users, Application programmers, Database administrators
3. Database system concepts and architecture are used to define the _____ and _____ of a database system. Answer: Structure, Functionality
4. The _____ data model is used to represent data at the physical level. Answer: Internal
5. A database _____ defines the structure of the data stored in a database. Answer: Schema
6. A database _____ represents a snapshot of the data in the database at a specific moment in time. Answer: Instance
7. The Entity-Relationship (ER) model is used for _____ modeling of data. Answer: Conceptual
8. In the ER model, an entity type is a _____ category of objects. Answer: Distinct
9. The attributes of an entity type are also called its _____. Answer: Properties
10. A _____ key is an attribute that uniquely identifies an entity within an entity set. Answer: Primary
11. In the ER model, a relationship type is a _____ between entity types. Answer: Association
12. The set of all entities of the same entity type is called an _____. Answer: Entity set
13. In a relationship type, _____ define the roles that entity types play in the relationship. Answer: Roles
14. _____ are constraints that limit the number of entities that can be involved in a relationship. Answer: Cardinality constraints
15. A weak entity type is an entity type that does not have a _____ key attribute. Answer: Primary
16. In the ER model, _____ are used to give a name to entity types, attributes, and relationship types. Answer: Naming conventions
17. The Enhanced Entity-Relationship (EER) model extends the ER model with concepts like _____ and _____. Answer: Specialization, Generalization
18. In EER modeling, a(n) _____ is a subset of entities of an entity type, grouped together for some specific purpose. Answer: Subtype
19. A(n) _____ is a higher-level entity type that encompasses one or more lower-level entity types. Answer: Supertype
20. _____ is a type of relationship in the EER model where an entity can participate in multiple roles within the same relationship. Answer: Recursive relationship
21. A(n) _____ is a condition that must hold for all entities in a relationship. Answer: Participation constraint
22. In an EER model, a(n) _____ attribute is an attribute that can have multiple values for a single entity. Answer: Multivalued
23. A(n) _____ key is a combination of attributes that uniquely identifies an entity within an entity set. Answer: Composite
24. In the EER model, _____ constraints specify that an entity of a subtype must also be an entity of its supertype. Answer: Inheritance

25. The process of converting an EER model into relational tables is called _____. Answer: Mapping
26. A _____ is a property that defines some aspect of the state of an entity. Answer: Attribute
27. In the ER model, a _____ key is an attribute that can be used to uniquely identify an entity. Answer: Candidate
28. In EER modeling, a(n) _____ entity type is one that inherits attributes from more than one supertype. Answer: Overlapping
29. _____ constraints specify the minimum and maximum number of entities that can participate in a relationship. Answer: Cardinality
30. In an EER model, a _____ specialization is one where an entity can belong to more than one subtype simultaneously. Answer: Multiple
31. Database design theory focuses on creating databases that are _____, efficient, and minimize data redundancy. Answer: Structured
32. In the context of normalization, a functional dependency is a relationship between two sets of attributes: a _____ and a _____. Answer: Determinant, Dependent
33. A superkey is a set of one or more attributes that can uniquely identify a(n) _____ in a relation. Answer: Tuple
34. The highest normal form that a relation can achieve is called _____. Answer: Fifth Normal Form (5NF)
35. A relation is in _____ if it is in 1NF and no non-prime attribute is partially dependent on a superkey. Answer: Second Normal Form (2NF)
36. A relation is in _____ if it is in 2NF and there are no partial dependencies between non-prime attributes and the superkey. Answer: Third Normal Form (3NF)
37. Boyce-Codd Normal Form (BCNF) is a stronger form of _____. Answer: 3NF
38. A _____ dependency occurs when one or more multivalued attributes are functionally dependent on another attribute in a relation. Answer: Multivalued
39. In Fourth Normal Form (4NF), a relation is free from _____ dependencies. Answer: Multivalued
40. A _____ dependency is a constraint where the presence of certain values in a relation requires the presence of certain values in another relation. Answer: Join
41. Fifth Normal Form (5NF) is also known as _____. Answer: Project-Join Normal Form (PJNF)
42. In 5NF, each relation is a _____ of a join dependency. Answer: Projection
43. A _____ key is a minimal superkey, which means no proper subset of it can uniquely identify a tuple in a relation. Answer: Candidate
44. A relation is in 1NF if it contains only _____ values in each attribute. Answer: Atomic
45. A _____ dependency occurs when an attribute is functionally dependent on a part of a composite key. Answer: Partial
46. Functional dependencies are typically represented using a notation like $A \rightarrow B$, where A is the _____ and B is the _____. Answer: Determinant, Dependent
47. To bring a relation to 2NF, you need to remove _____ dependencies. Answer: Partial
48. A relation is in BCNF if, for every non-trivial functional dependency $X \rightarrow Y$, X is a _____. Answer: Superkey
49. In 4NF, a relation is free from _____ dependencies. Answer: Multivalued
50. Fifth Normal Form (5NF) deals with _____ dependencies. Answer: Join
51. In a multivalued dependency $X \twoheadrightarrow Y$, X is called the _____ set, and Y is called the _____ set. Answer: Left-hand, Right-hand
52. A candidate key is a set of attributes that can uniquely identify a tuple, and it also satisfies the _____ property. Answer: Irreducibility

53. A _____ is a set of attributes that is a candidate key but is not part of any candidate key. Answer: Superfluous attribute
54. The process of bringing a relation to a higher normal form without introducing any new attributes or dependencies is called _____. Answer: Lossless decomposition
55. In BCNF, every non-trivial functional dependency is determined by a _____. Answer: Superkey
56. A _____ key is a candidate key that is chosen to be the primary key of a relation. Answer: Primary
57. Fifth Normal Form (5NF) ensures _____ of attributes in a relation. Answer: Orthogonality
58. To eliminate partial dependencies in a relation, you need to bring it to _____. Answer: 2NF
59. A relation is in BCNF if, for every non-trivial functional dependency $X \rightarrow Y$, X is a _____ key. Answer: Super
60. The goal of normalization is to reduce data _____ and improve data integrity in a database. Answer: Redundancy
61. The relational data model represents data as a collection of _____. Answer: Tables
62. In the relational model, a _____ is a named collection of related attributes. Answer: Relation
63. A _____ is a set of values of the same type. Answer: Domain
64. SQL stands for Structured Query _____ Language. Answer: Language
65. The SQL _____ statement is used to create a new table. Answer: CREATE TABLE
66. The _____ constraint ensures that each row in a table has a unique value for a specified column. Answer: UNIQUE
67. The _____ data type in SQL is used to store variable-length character strings. Answer: VARCHAR
68. The primary key is a combination of one or more _____ in a table. Answer: Attributes/Columns
69. The SQL _____ statement is used to retrieve data from a database. Answer: SELECT
70. The _____ operation in SQL is used to add new rows to a table. Answer: INSERT
71. The SQL _____ statement is used to remove rows from a table. Answer: DELETE
72. The SQL _____ statement is used to modify existing data in a table. Answer: UPDATE
73. In SQL, the _____ keyword is used to filter rows based on a condition in the WHERE clause. Answer: WHERE
74. The SQL _____ clause is used to sort the result set in ascending or descending order. Answer: ORDER BY
75. The SQL _____ function is used to count the number of rows in a result set. Answer: COUNT
76. The _____ operation in relational algebra selects rows from a table based on a condition. Answer: SELECT
77. The _____ operation in relational algebra eliminates duplicate rows from a table. Answer: PROJECT
78. In SQL, a _____ key is a key that is used to establish a relationship between two tables. Answer: Foreign
79. The SQL _____ statement is used to specify the structure of a table, including its columns and data types. Answer: CREATE TABLE
80. The SQL _____ constraint ensures that a column in a table cannot have a NULL value. Answer: NOT NULL

81. The _____ operation in relational algebra combines two relations to create a new relation. Answer: JOIN
82. The SQL _____ statement is used to change the structure of an existing table. Answer: ALTER TABLE
83. The SQL _____ clause is used to group rows in a result set based on the values in one or more columns. Answer: GROUP BY
84. The _____ operation in relational algebra returns all the rows that are in one relation but not in another. Answer: MINUS
85. In SQL, the _____ statement is used to remove a table from the database. Answer: DROP TABLE
86. The SQL _____ constraint ensures that the values in a column are unique across all rows in a table. Answer: UNIQUE
87. The _____ operation in relational algebra combines two relations and returns only the rows that are common to both. Answer: INTERSECT
88. The SQL _____ data type is used to store fixed-length character strings. Answer: CHAR
89. The SQL _____ clause is used to filter rows based on multiple conditions in the WHERE clause. Answer: AND
90. The _____ operation in relational algebra returns all the rows that are in both relations. Answer: UNION
91. Transaction processing involves the execution of a series of _____ on a database. Answer: Operations
92. A _____ is a logical unit of work that consists of one or more database operations. Answer: Transaction
93. ACID stands for Atomicity, Consistency, Isolation, and _____. Answer: Durability
94. The _____ property of transactions ensures that a transaction's changes to the database are permanent. Answer: Durability
95. A _____ is a schedule of transactions that preserves the consistency of the database. Answer: Serializable schedule
96. In a two-phase locking protocol, a transaction can acquire locks but cannot release any locks until it has reached the _____ phase. Answer: Commit
97. A _____ lock allows multiple transactions to read the same data concurrently but prevents write access until the lock is released. Answer: Shared
98. The _____ property of transactions ensures that a transaction appears to be the only one accessing the data. Answer: Isolation
99. The _____ control technique allows transactions to execute concurrently while ensuring the consistency of the database. Answer: Concurrency
100. The _____ property of transactions ensures that the database starts in a consistent state and ends in a consistent state after a transaction. Answer: Consistency
101. _____ is the process of undoing the effects of a failed transaction. Answer: Rollback
102. The _____ lock allows a transaction to have exclusive access to data and prevents other transactions from accessing it. Answer: Exclusive
103. In a _____ schedule, transactions are executed one after the other with no overlap. Answer: Serial
104. The _____ property of transactions ensures that a transaction is either fully completed or fully aborted. Answer: Atomicity
105. A _____ is a set of rules that govern the order in which transactions can access data. Answer: Synchronization point
106. The _____ timestamp ordering protocol assigns a unique timestamp to each transaction and uses these timestamps to determine the order of execution. Answer: Time-stamp
107. The _____ rule ensures that a transaction that reads a data item after another transaction has written it will not be rolled back. Answer: Strict

108. The _____ rule ensures that if a transaction T2 writes a data item that transaction T1 reads, then T1 must read T2's write. Answer: Cascading
109. The _____ property of transactions ensures that the database remains in a consistent state even in the presence of hardware or software failures. Answer: Recoverability
110. A _____ is a data structure used to record information about transactions and their status. Answer: Transaction log
111. The _____ control technique uses locks to ensure that only one transaction can access a data item at a time. Answer: Lock-based
112. The _____ rule ensures that a transaction cannot read a data item while it is being written by another transaction. Answer: No-read
113. A _____ schedule is one in which the final state of the database is the same as if the transactions had executed one after the other. Answer: Serializable
114. The _____ rule ensures that a transaction cannot write a data item while it is being read by another transaction. Answer: No-write
115. The _____ timestamp ordering protocol uses timestamps to determine the order in which transactions are allowed to access data items. Answer: Wait
116. The _____ control technique uses timestamps to order transactions and determine which transactions can proceed. Answer: Time-stamp ordering
117. The _____ rule ensures that a transaction cannot read a data item that is being written by another transaction. Answer: No-fault
118. The _____ rule ensures that a transaction cannot write a data item that is being written by another transaction. Answer: No-updates
119. The _____ property of transactions ensures that a transaction's changes to the database are reflected even after a system crash. Answer: Durability
120. The _____ control technique uses timestamps to order transactions and resolve conflicts. Answer: Time-stamp ordering

Short Type

Introduction to Database and Database Users:

1. What is a database?

- Answer: A database is a structured collection of data that is organized and stored for efficient retrieval and manipulation.

2. Who are the primary users of a database system?

- Answer: The primary users of a database system include end-users, application programmers, and database administrators.

3. What is the role of a database administrator (DBA)?

- Answer: A DBA is responsible for managing and maintaining the database system, including ensuring data integrity, security, and performance.

Database System Concepts and Architecture:

4. What is a DBMS, and what does it stand for?

- Answer: DBMS stands for Database Management System. It is software that facilitates the creation, maintenance, and use of a database.

5. Name two popular types of database models.

- Answer: Two popular database models are the relational model and the NoSQL model.

6. What is a schema in a database system?

- Answer: A schema defines the structure of the database, including tables, relationships, and constraints.

Conceptual Modeling and Database Design (ER Model):

7. What is an Entity in the Entity-Relationship (ER) model?

- Answer: An entity represents a real-world object or concept that has data attributes.

8. What is an attribute in the context of the ER model?

- Answer: An attribute is a property or characteristic of an entity.

9. **What is a key attribute in the ER model?**
 - Answer: A key attribute is an attribute that uniquely identifies an entity within its entity set.
10. **What is a relationship in the ER model?**
 - Answer: A relationship represents an association between two or more entities.
11. **What is a weak entity type in the ER model?**
 - Answer: A weak entity type is an entity type that doesn't have a primary key attribute on its own and relies on another entity (called the owner) for identification.
12. **What are structural constraints in the ER model?**
 - Answer: Structural constraints specify rules and restrictions on how entities and relationships can be combined and used in the database.

ER Naming Conventions:

13. **Why is it important to follow naming conventions in a database?**
 - Answer: Naming conventions help maintain consistency and clarity in a database schema, making it easier to understand and maintain.
14. **Give an example of a common naming convention for primary keys.**
 - Answer: Commonly, primary keys are named as "TableNameID" or "TableName_PK," where TableName is the name of the table.

Enhanced Entity-Relationship (EER) Model:

15. **What does the Enhanced Entity-Relationship (EER) model extend from the ER model?**
 - Answer: The EER model extends the ER model by adding more advanced modeling concepts, such as specialization, generalization, and aggregation.
16. **What is specialization in the EER model?**
 - Answer: Specialization is the process of defining subtypes or subclasses of an entity type based on certain characteristics.
17. **What is an attribute in the EER model?**
 - Answer: An attribute in the EER model is similar to an attribute in the ER model, representing a property or characteristic of an entity.
18. **Explain the concept of aggregation in the EER model.**
 - Answer: Aggregation in the EER model represents a whole-part relationship between entities, where one entity (the whole) is composed of other entities (the parts).

Functional Dependencies and Normal Forms:

1. **What is a functional dependency in a database?**
 - Answer: A functional dependency in a database occurs when one attribute's value uniquely determines another attribute's value.
2. **What is the primary objective of normalization in database design?**
 - Answer: The primary objective of normalization is to reduce data redundancy and improve data integrity in a database.
3. **What is the first normal form (1NF)?**
 - Answer: 1NF requires that all attributes in a relation have atomic (indivisible) values and there are no repeating groups or arrays.
4. **What is the second normal form (2NF)?**
 - Answer: 2NF requires that a relation is in 1NF, and all non-key attributes are fully functionally dependent on the primary key.
5. **What is Boyce-Codd Normal Form (BCNF)?**
 - Answer: BCNF is a stronger form of 2NF, where every non-trivial functional dependency in the relation is dependent on a superkey.
6. **What is the third normal form (3NF)?**
 - Answer: 3NF requires that a relation is in 2NF and that there are no transitive dependencies between non-key attributes.
7. **What is a superkey in the context of normalization?**

- Answer: A superkey is a set of attributes that can uniquely identify a tuple (row) in a relation.

8. What is a candidate key in a relation?

- Answer: A candidate key is a minimal superkey, meaning it is a minimal set of attributes that uniquely identifies tuples.

Multivalued Dependency and Fourth Normal Form:

9. What is a multivalued dependency in a relation?

- Answer: A multivalued dependency exists when an attribute set can have multiple independent sets of values associated with it.

10. What is Fourth Normal Form (4NF)?

- Answer: 4NF is a normal form that deals with multivalued dependencies, ensuring that they are appropriately represented in a relation.

Join Dependencies and Fifth Normal Form:

11. What is a join dependency in the context of normalization?

- Answer: A join dependency exists when a relation can be decomposed into multiple smaller relations and then recombined through a natural join to obtain the original relation.

12. What is Fifth Normal Form (5NF)?

- Answer: 5NF, also known as Project-Join Normal Form (PJNF), addresses join dependencies and ensures that relations can be reconstructed from their decomposed parts using natural joins.

13. What is a lossless join property in normalization?

- Answer: A lossless join property ensures that when decomposing a relation and then recombining it through a natural join, the resulting relation is the same as the original relation.

14. Why is achieving 5NF important in database design?

- Answer: Achieving 5NF is essential to ensure that a database schema is free from anomalies related to join operations and that it maintains data integrity.

15. What are some common methods to achieve 5NF?

- Answer: Common methods to achieve 5NF include decomposition into smaller relations based on join dependencies and ensuring the lossless join property.

Relational Model Concepts:

1. What is the primary key in a relational database?

- Answer: The primary key is a unique identifier for each record in a table, ensuring data integrity.

2. What is a foreign key in the relational model?

- Answer: A foreign key is a field in one table that links to the primary key in another table, establishing relationships between tables.

3. What is a relation in the relational model?

- Answer: A relation is a table that represents a specific entity or concept, consisting of rows and columns.

4. What is normalization in the context of the relational model?

- Answer: Normalization is the process of organizing data in a relational database to minimize data redundancy and improve data integrity.

Basic SQL Queries:

5. What does SQL stand for?

- Answer: SQL stands for Structured Query Language.

6. What is the purpose of the SELECT statement in SQL?

- Answer: The SELECT statement is used to retrieve data from one or more tables in a database.

7. What SQL keyword is used to filter records in a SELECT statement?

- Answer: The WHERE clause is used to filter records based on specified conditions.

8. **What is the result of a SQL INNER JOIN operation?**

- Answer: An INNER JOIN returns only the rows that have matching values in both joined tables.

SQL Data Definition and Data Types:

9. **What is the purpose of the CREATE TABLE statement in SQL?**

- Answer: The CREATE TABLE statement is used to define the structure of a new table in a database.

10. **What is the difference between CHAR and VARCHAR data types in SQL?**

- Answer: CHAR stores fixed-length character data, while VARCHAR stores variable-length character data.

Constraints in SQL:

11. **What is a primary key constraint in SQL?**

- Answer: A primary key constraint enforces the uniqueness and integrity of values in a specific column or set of columns.

12. **What is a foreign key constraint in SQL?**

- Answer: A foreign key constraint enforces referential integrity by ensuring that values in one column match values in another column.

Retrieval Queries in SQL:

13. **How do you retrieve all records from a table in SQL?**

- Answer: To retrieve all records, you can use the SQL statement: SELECT * FROM table_name;

14. **What SQL clause is used to sort the result set in ascending order?**

- Answer: The ORDER BY clause is used to sort the result set in ascending order.

15. **What is the purpose of the GROUP BY clause in SQL?**

- Answer: The GROUP BY clause is used to group rows that have the same values into summary rows.

INSERT, DELETE, and UPDATE Statements in SQL:

16. **What is the purpose of the INSERT statement in SQL?**

- Answer: The INSERT statement is used to add new rows to a table.

17. **What is the purpose of the DELETE statement in SQL?**

- Answer: The DELETE statement is used to remove rows from a table.

18. **What is the purpose of the UPDATE statement in SQL?**

- Answer: The UPDATE statement is used to modify existing records in a table.

Relational Algebra and Relational Calculus:

19. **What is the primary focus of relational algebra?**

- Answer: Relational algebra focuses on the operations used to manipulate and retrieve data from relations (tables).

20. **What is the result of the SELECT operation in relational algebra?**

- Answer: The SELECT operation retrieves rows from a relation that satisfy a given condition.

21. **What is the result of the PROJECT operation in relational algebra?**

- Answer: The PROJECT operation selects specific columns from a relation while eliminating duplicates.

Binary Relations in SQL:

22. **What is the purpose of the JOIN operation in SQL?**

- Answer: The JOIN operation combines rows from two or more tables based on a related column.

23. **What is the result of a SQL LEFT JOIN operation?**

- Answer: A LEFT JOIN returns all rows from the left table and matching rows from the right table; if there are no matches, NULL values are included for the right table.

24. **What is the purpose of the DIVISION operation in relational algebra?**

- Answer: The DIVISION operation returns all values from one column that have corresponding values in another column.

Introduction to Transaction Processing:

1. What is transaction processing in the context of databases?

- Answer: Transaction processing involves managing a series of related database operations as a single unit to ensure data consistency and integrity.

2. What is the primary goal of transaction processing systems (TPS)?

- Answer: The primary goal of TPS is to ensure the correctness, reliability, and integrity of data in a database while allowing concurrent access by multiple users.

3. What is a transaction in the context of database management?

- Answer: A transaction is a sequence of one or more database operations (such as inserts, updates, or deletes) that are treated as a single unit of work.

Transaction and System Concepts:

4. What is the ACID property in the context of transactions?

- Answer: ACID stands for Atomicity, Consistency, Isolation, and Durability. It describes the key properties that transactions should exhibit.

5. Explain the Atomicity property of transactions.

- Answer: Atomicity ensures that a transaction is treated as an indivisible unit, meaning that it either completes in its entirety or has no effect on the database.

Properties of Transactions:

6. What is the Consistency property of transactions?

- Answer: Consistency ensures that a transaction brings the database from one consistent state to another, adhering to integrity constraints.

7. What is the Isolation property of transactions?

- Answer: Isolation guarantees that concurrent transactions do not interfere with each other and that each transaction sees a consistent state of the database.

8. What is the Durability property of transactions?

- Answer: Durability ensures that once a transaction is committed, its effects are permanent and survive system failures.

Recoverability:

9. What is the significance of recoverability in transaction processing?

- Answer: Recoverability ensures that even in the event of a system crash or failure, a database can be restored to a consistent state.

Serializability:

10. What does serializability mean in the context of transactions?

- Answer: Serializability ensures that the execution of concurrent transactions produces the same result as if they were executed serially, one after another.

Concurrency Control Techniques:

11. What is the purpose of concurrency control in transaction processing?

- Answer: Concurrency control ensures that multiple transactions can execute concurrently without violating the ACID properties.

12. Name two common concurrency control techniques.

- Answer: Two common concurrency control techniques are locking and timestamp-based ordering.

Locking Techniques for Concurrency Control:

13. What is a lock in the context of concurrency control?

- Answer: A lock is a mechanism used to restrict access to a resource (e.g., a database record) by allowing only one transaction to modify it at a time.

14. What is a shared lock, and when is it used?

- Answer: A shared lock allows multiple transactions to read a resource simultaneously but prevents any of them from modifying it.

15. What is an exclusive lock, and when is it used?

- Answer: An exclusive lock prevents other transactions from accessing a resource, ensuring that only one transaction can modify it at a time.

Concurrency Control based on Time-Stamp Ordering:

16. What is time-stamp ordering in concurrency control?

- Answer: Time-stamp ordering assigns a unique timestamp to each transaction and uses these timestamps to determine the order in which transactions are allowed to access resources.

17. What is the purpose of a transaction timestamp in time-stamp ordering?

- Answer: A transaction timestamp is used to record when a transaction starts and to order transactions based on their timestamps.

LONG TYPE

Introduction to Database and Database Users:

1. What is a database, and why is it essential in modern computing?
2. Explain the role of a Database Management System (DBMS) in managing data.
3. What are the key components of a DBMS architecture?
4. Describe the different types of database users and their responsibilities.
5. How does data redundancy affect the efficiency and reliability of a database system?

Database System Concepts and Architecture: Data Models, Schema, and Instances:

6. Define the terms "data model," "schema," and "data instance" in the context of databases.
7. Compare and contrast the hierarchical, network, and relational data models.
8. Explain the advantages and disadvantages of a relational data model.
9. What is a database schema, and how does it differ from a database instance?
10. Discuss the importance of data independence in a database system.

Conceptual Modeling and Database Design: Entity-Relationship (ER) Model:

11. What is the Entity-Relationship (ER) model, and why is it used in database design?
12. Define the concepts of entity types and entity sets in the ER model.
13. Explain the significance of attributes in the ER model. Give examples.
14. What is a key attribute in the context of the ER model? Provide an example.
15. Describe the purpose of relationship types and relationship sets in the ER model.

Conceptual Modeling and Database Design: Weak Entity Types, ER Naming

Conventions:

16. What is a weak entity type, and how does it differ from a strong entity type?
17. Provide an example of a weak entity type and its associated identifying relationship.
18. Explain the importance of naming conventions in the ER model. Give some common naming conventions.
19. How do roles and structural constraints enhance the expressiveness of the ER model?
20. Discuss the role of cardinality and participation constraints in relationship sets.

Enhanced Entity-Relationship (EER) Model:

21. What is the Enhanced Entity-Relationship (EER) model, and why is it used for complex data modeling?
22. Describe the concept of specialization and generalization in the EER model.
23. Explain the purpose of aggregation in the EER model. Provide an example.
24. How does the EER model handle multivalued attributes and derived attributes?
25. What are the key features of the EER model that extend beyond the basic ER model?

Database Design Theory and Normalization: Functional Dependencies:

26. Define functional dependencies in the context of database design.
27. Explain the difference between a superkey, a candidate key, and a primary key.
28. What is a partial dependency, and how does it relate to normalization?
29. Describe the process of determining the closure of an attribute set with respect to a set of functional dependencies.
30. Why is it essential to minimize redundancy in a database through normalization?

Normal Forms based on Primary Keys: Second and Third Normal Forms:

31. Define the Second Normal Form (2NF) and explain how it eliminates partial dependencies.
32. Discuss the concept of transitive dependencies and their role in the Third Normal Form (3NF).
33. Provide an example of a table that is in 2NF but not in 3NF.
34. What are the advantages and disadvantages of achieving higher levels of normalization?
35. How does normalization impact data insertion, update, and retrieval operations?

Boyce-Codd Normal Form, Multivalued Dependency, and Fourth Normal Form:

36. Define the Boyce-Codd Normal Form (BCNF) and its relationship to candidate keys.
37. Explain multivalued dependencies and their significance in database design.
38. Provide an example of a table that is in BCNF but violates multivalued dependencies.
39. What is the Fourth Normal Form (4NF), and when is it applicable in database design?
40. How does BCNF differ from 4NF, and in what scenarios would you use each?

Join Dependencies and Fifth Normal Form:

41. Describe join dependencies and how they relate to database design.
42. Explain the concept of Fifth Normal Form (5NF) and its relationship to join dependencies.
43. Provide an example of a table that satisfies 5NF.
44. When is it beneficial to decompose a relation to achieve 5NF?
45. Discuss the trade-offs between achieving higher normalization levels and query performance.

Relational Model Concepts:

46. Define the key components of the relational model, including relations, tuples, and attributes.
47. Explain the concept of integrity constraints in the relational model.
48. What are domain constraints, and why are they important in defining attributes?
49. Discuss the role of primary keys and foreign keys in maintaining referential integrity.
50. How does the relational model support data manipulation through relational algebra and SQL?

Basic SQLs, SQL Data Definition and Data Types, Constraints in SQL:

51. What are the fundamental SQL operations, and how do they relate to the relational algebra?
52. Describe the SQL data definition language (DDL) and its role in creating and modifying database structures.
53. Provide examples of common SQL data types and their use cases.
54. Explain the purpose of constraints in SQL, including primary key, foreign key, and check constraints.
55. How can you enforce referential integrity in SQL databases using foreign keys?

Retrieval Queries in SQL, INSERT, DELETE, UPDATE Statements in SQL:

56. Write an SQL query to retrieve all employees who earn more than \$50,000 per year.
57. Explain the SQL SELECT statement and its various clauses, such as WHERE and GROUP BY.
58. How do you use the SQL INSERT statement to add new records to a table?
59. Discuss the SQL DELETE statement and its implications for data integrity.
60. Describe the SQL UPDATE statement and its role in modifying existing data.

Relational Algebra and Relational Calculus: Unary Relational Operations: SELECT and PROJECT:

61. What is relational algebra, and how does it differ from SQL?
62. Explain the unary relational operation SELECT (σ) and its use in filtering rows.
63. Describe the unary relational operation PROJECT (π) and its purpose in selecting specific columns.
64. Provide an example of using SELECT and PROJECT operations to retrieve data from a relation.
65. How does relational algebra help in formulating complex database queries?

Binary Relation: JOIN and DIVISION:

66. Define the binary relational operation JOIN (\bowtie) and its role in combining two relations.
67. Explain the different types of joins, including INNER JOIN, LEFT JOIN, and RIGHT JOIN.
68. What is the purpose of the DIVISION operation in relational algebra?
69. Provide an example of using JOIN and DIVISION operations in a database query.
70. Discuss the efficiency considerations when performing JOIN operations on large datasets.

Introduction to Transaction Processing Concepts and Theory:

71. What is transaction processing, and why is it critical in database systems?
72. Define the concept of a transaction and its properties in the context of database management.
73. Explain the ACID properties and their importance in ensuring data consistency.
74. How does a system handle the recovery of data after a transaction failure?
75. Discuss the concept of concurrency control and its challenges in multi-user database systems.

Transaction and System Concepts: Properties of Transactions:

76. Describe the properties of a transaction, including atomicity, consistency, isolation, and durability.
77. How does the atomicity property ensure that a transaction is all-or-nothing?
78. Explain the role of the consistency property in maintaining data integrity.
79. Discuss isolation levels in transactions and their impact on concurrent access to data.
80. Why is durability important in ensuring that committed data survives system failures?

Recoverability, Serializability, Concurrency Control Techniques:

81. What is the concept of recoverability in transaction processing, and how is it achieved?
82. Define the serializability property and its significance in controlling transaction execution.
83. Explain the difference between strict two-phase locking and deadlock prevention.
84. Discuss the challenges of distributed databases and their impact on concurrency control.
85. How do optimistic concurrency control techniques differ from pessimistic ones?

Locking Techniques for Concurrency Control:

86. Describe the purpose of locks in controlling concurrent access to database resources.
87. Explain the difference between shared locks and exclusive locks in a locking system.
88. What is deadlock, and how can it be detected and resolved in a locking environment?
89. Discuss the advantages and disadvantages of using lock-based concurrency control.
90. How can lock granularity impact the performance of a database system?

Concurrency Control based on Time-Stamp Ordering:

91. What is time-stamp ordering, and how does it ensure serializability in transactions?
92. Explain the concept of a global time-stamp and how it is used in time-stamp ordering.
93. Discuss the optimistic concurrency control approach using time-stamps.
94. How can time-stamp-based concurrency control handle conflicts and ensure data consistency?
95. Compare and contrast time-stamp ordering with traditional lock-based concurrency control.