Multiple Choice Questions & Answers

1. What is the primary goal of schema refinement in database design?

A) Eliminating all redundancy

B) Maximizing data storage

C) Optimizing query performance

D) Reducing data integrity

Answer: A) Eliminating all redundancy

2. Which problem is caused by redundancy in a database schema?

A) Improved data consistency

B) Enhanced data retrieval

C) Increased storage efficiency

D) Data update anomalies

Answer: D) Data update anomalies

3. In the context of database design, what is decomposition?

A) The process of combining tables

B) The process of simplifying queries

C) The process of breaking a relation into smaller relations

D) The process of optimizing indexes

Answer: C) The process of breaking a relation into smaller relations

4. What are some common problems related to decomposition in database design?

A) Loss of data integrity

B) Improved query performance

C) Simplified schema

D) Increased storage space

Answer: A) Loss of data integrity

5. What is the primary purpose of reasoning about functional dependencies in schema refinement?

A) To identify unique keys

B) To maximize storage efficiency

C) To minimize data consistency

D) To simplify queries

Answer: A) To identify unique keys

6. Which normal form requires that a relation be in 1NF and no partial dependency exists?

A) FIRST normal form

B) SECOND normal form

C) THIRD normal form

D) BCNF (Boyce-Codd Normal Form)

Answer: C) THIRD normal form

7. What does BCNF (Boyce-Codd Normal Form) address in a relation schema?

A) Partial dependency

B) Multivalued dependency

C) Functional dependency

D) Lossless join

Answer: C) Functional dependency

8. What does it mean for a decomposition of a relation to be in lossless join form?

A) No data loss occurs when joining the decomposed relations

B) Data redundancy is eliminated

C) It is in 3NF

D) It is fully normalized

Answer: A) No data loss occurs when joining the decomposed relations

9. What are multi-valued dependencies in the context of database design?

A) Dependencies between attributes in a relation

B) Dependencies between tables

C) Dependencies involving multiple values

D) Dependencies between keys

Answer: C) Dependencies involving multiple values

10. Which normal form addresses multi-valued dependencies in a relation schema?A) SECOND normal form

B) THIRD normal form

C) BCNF (Boyce-Codd Normal Form)

D) FOURTH normal form

Answer: D) FOURTH normal form

11. What is the primary focus of the FOURTH normal form (4NF)?

A) Eliminating all redundancy

B) Addressing multi-valued dependencies

C) Ensuring lossless join decomposition

D) Simplifying queries

Answer: B) Addressing multi-valued dependencies

12. Which normal form is sometimes referred to as Project-Join Normal Form (PJNF)?

A) SECOND normal form

B) THIRD normal form

C) BCNF (Boyce-Codd Normal Form)

D) FIFTH normal form

Answer: D) FIFTH normal form

13. What is the primary goal of the FIFTH normal form (5NF)?

A) Eliminating all redundancy

B) Ensuring lossless join decomposition

C) Addressing join dependencies

D) Simplifying queries

Answer: C) Addressing join dependencies

14. Which normal form ensures that all join dependencies are satisfied?

A) FIRST normal form

B) SECOND normal form

C) BCNF (Boyce-Codd Normal Form)

D) FIFTH normal form

Answer: D) FIFTH normal form

15. In schema refinement, what is the purpose of identifying functional dependencies?

A) 1To improve data consistency

B) To maximize storage efficiency

C) To minimize query performance

D) To eliminate data redundancy

Answer: D) To eliminate data redundancy

16. Which problem does redundancy in a database schema often lead to?

A) Improved data consistency

B) Enhanced data retrieval

C) Data update anomalies

D) Simplified schema

Answer: C) Data update anomalies

17. What is the primary goal of normalization in schema refinement?

A) To increase data redundancy

B) To maximize data inconsistency

C) To eliminate data anomalies

D) To minimize data storage

Answer: C) To eliminate data anomalies

18. Which normal form allows a relation to have no partial dependencies?

A) FIRST normal form

B) SECOND normal form

C) THIRD normal form

D) BCNF (Boyce-Codd Normal Form)

Answer: C) THIRD normal form

19. What does BCNF (Boyce-Codd Normal Form) aim to eliminate in a relation sch1ema?

A) Partial dependency

B) Multivalued dependency

C) Functional dependency

D) Data anomalies

Answer: A) Partial dependency

20. In schema refinement, what does "lossless join decomposition" mean?

A) Data loss occurs when joining decomposed relations

B) No data loss occurs when joining decomposed relations

C) Data redundancy is introduced

D) Data is lost during queries

Answer: B) No data loss occurs when joining decomposed relations

21. What are multi-valued dependencies?

A) Dependencies between attributes in a relation

B) Dependencies between tables

C) 2Dependencies involving multiple values

D) Dependencies between keys

Answer: C) Dependencies involving multiple values

22. Which normal form addresses multi-valued dependencies in a relation schema?

A) SECOND normal form

B) THIRD normal form

C) BCNF (Boyce-Codd Normal Form)

D) FOURTH normal form

Answer: D) FOURTH normal form

23. What is the primary focus of the FOURTH normal form (4NF)?

A) Eliminating all redundancy

B) Addressing multi-valued dependencies

C) Ensuring lossless join decomposition

D) Simplifying queries

Answer: B) Addressing multi-valued dependencies

24. Which normal form is sometimes referred to as Project-Join Normal Form (PJNF)?

A) SECOND normal form

B) THIRD normal form

C) BCNF (Boyce-Codd Normal Form)

D) FIFTH normal form

Answer: D) FIFTH normal form

25. What is the primary goal of the FIFTH normal form (5NF)?

A) Eliminating all redundancy

B) 2Ensuring lossless join decomposition

C) Addressing join dependencies

D) Simplifying queries

Answer: C) Addressing join dependencies

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

A) A database backup

B) A database query

C) A sequence of database operations

D) A database schema

Answer: C) A sequence of database operations

27. In the transaction state model, what is the "committed" state of a transaction?

A) The transaction has been aborted

B) 2The transaction is in progress

C) The transaction has been successfully completed

D) The transaction is waiting

Answer: C) The transaction has been successfully completed

28. What is the primary goal of ensuring atomicity and durability in database transactions?

A) To improve query performance

B) To enhance data retrieval

C) To ensure that transactions are executed in parallel

D) To maintain data consistency and integrity

Answer: D) To maintain data consistency and integrity

29. What is a concurrent execution of transactions in a database system?

A) A single transaction executing sequentially

B) Multiple transactions executing independently

C) A transaction in the "committed" state

D) A transaction that has failed

Answer: B) Multiple transactions executing independently

30. In the context of transactions, what does "serializability" refer to?

A) The ability to execute transactions in isolation

B) The ability to execute transactions sequentially

C) The ability to execute transactions concurrently

D) The ability to recover from failures

Answer: B) The ability to execute transactions sequentially

31. What does the term "recoverability" mean in the context of transactions?

A) The ability to execute transactions in isolation

B) The ability to execute transactions concurrently

C) The ability to recover from failures

D) The ability to execute transactions sequentially

Answer: C) The ability to recover from failures

32. How is isolation typically implemented in a database system?

A) Using timestamps

B) Using locks

C) Using validation-based protocols

D) Using log-based recovery

Answer: B) Using locks

33. What is the purpose of testing for serializability in database transactions?

A) To improve data retrieval

B) To ensure data consistency

C) To validate data integrity

D) To check if concurrent transactions are serializable

Answer: D) To check if concurrent transactions are serializable

34. What is a lock-based protocol in database concurrency control?

A) A protocol that uses timestamps

B) A protocol that uses locks to control access

C) A protocol that uses validation checks

D) A protocol that uses log-based recovery

Answer: B) A protocol that uses locks to control access

35. What is a timestamp-based protocol in database concurrency control?

A) A protocol that uses locks to control access

B) A protocol that uses timestamps to order transactions

C) A protocol that uses validation checks

D) A protocol that uses log-based recovery

Answer: B) A protocol that uses timestamps to order transactions

36. What is the primary goal of validation-based protocols in database concurrency control?

A) To improve query performance

B) To enhance data retrieval

C) To ensure data consistency

D) To maintain data redundancy

Answer: C) To ensure data consistency

37. What is meant by "multiple granularity" in database concurrency control?

A) The ability to use multiple database systems

B) The ability to control transactions at different levels

C) The ability to perform multiple queries simultaneously

D) The ability to recover from multiple failures

Answer: B) The ability to control transactions at different levels

38. What is the relationship between recovery and atomicity in database transactions?

A) They are unrelated concepts

B) Atomicity is a subset of recovery

C) Recovery is a subset of atomicity

D) They are synonymous terms

Answer: B) Atomicity is a subset of recovery

39. What is the purpose of log-based recovery in database systems?

A) To improve data retrieval

B) To recover lost data

C) To enhance query performance

D) To maximize data redundancy

Answer: B) To recover lost data

40. In log-based recovery, what is a "redo log"?

A) A log that records only failed transactions

B) A log that records successful transactions

C) A log that records changes made by transactions

D) A log that records queries

Answer: C) A log that records changes made by transactions

41. What does "recovery with concurrent transactions" refer to in database management?

A) The ability to recover from multiple failures

B) The ability to recover data while other transactions are in progress

C) The ability to execute transactions concurrently

D) The ability to recover lost data

Answer: B) The ability to recover data while other transactions are in progress

42. What is the primary purpose of a transaction in a database system?

A) To recover lost data

B) To perform data validation

C) To execute queries

D) To group database operations into a single unit

Answer: D) To group database operations into a single unit

43. In the transaction state model, what is the "active" state of a transaction?

A) The transaction has been committed

B) The transaction is in progress

C) The transaction has been aborted

D) The transaction is waiting

Answer: B) The transaction is in progress

44. What is the significance of ensuring atomicity in a database transaction?

A) To maximize data inconsistency

B) To allow partial execution of transactions

C) To improve data retrieval

D) To maintain data consistency

Answer: D) To maintain data consistency

45. What is a "concurrent execution" of transactions in a database system?

A) A single transaction executing sequentially

B) Multiple transactions executing independently

C) A transaction in the "committed" state

D) A transaction that has failed

Answer: B) Multiple transactions executing independently

46. What is the primary goal of serializability in database transactions?

A) To maximize data inconsistency

B) To improve query performance

C) To maintain data consistency

D) To allow concurrent execution of transactions

Answer: C) To maintain data consistency

47. What does "recoverability" ensure in the context of transactions?

A) The ability to execute transactions sequentially

B) The ability to execute transactions in isolation

C) The ability to recover from failures

D) The ability to execute transactions concurrently

Answer: C) The ability to recover from failures

48. How is isolation typically implemented in a database system?

A) Using timestamps

B) Using locks

C) Using validation-based protocols

D) Using log-based recovery

Answer: B) Using locks

49. What is the primary purpose of testing for serializability in database transactions?

A) To maximize data inconsistency

B) To ensure data consistency

C) To validate data integrity

D) To check if concurrent transactions are serializable

Answer: D) To check if concurrent transactions are serializable

50. What is a lock-based protocol in database concurrency control?

A) A protocol that uses timestamps

B) A protocol that uses locks to control access

C) A protocol that uses validation checks

D) A protocol that uses log-based recovery

Answer: B) A protocol that uses locks to control access

52. What is a timestamp-based protocol in database concurrency control?

A) A protocol that uses locks to control access

B) A protocol that uses timestamps to order transactions

C) A protocol that uses validation checks

D) A protocol that uses log-based recovery

Answer: B) A protocol that uses timestamps to order transactions

53. What is the primary goal of validation-based protocols in database concurrency control?

A) To maximize data inconsistency

B) To enhance data retrieval

C) To ensure data consistency

D) To maintain data redundancy

Answer: C) To ensure data consistency

54. What is meant by "multiple granularity" in database concurrency control?

A) The ability to use multiple database systems

B) The ability to control transactions at different levels

C) The ability to perform multiple queries simultaneously

D) The ability to recover from multiple failures

Answer: B) The ability to control transactions at different levels

55. What is the relationship between recovery and atomicity in database transactions?

A) They are unrelated concepts

B) Atomicity is a subset of recovery

C) Recovery is a subset of atomicity

D) They are synonymous terms

Answer: B) Atomicity is a subset of recovery

56. What is the purpose of log-based recovery in database systems?

A) To maximize data redundancy

B) To recover lost data

C) To enhance query performance

D) To improve data retrieval

Answer: B) To recover lost data

57. In log-based recovery, what is a "redo log"?

A) A log that records only failed transactions

B) A log that records successful transactions

C) A log that records changes made by transactions

D) A log that records queries

Answer: C) A log that records changes made by transactions

58. What does "recovery with concurrent transactions" refer to in database management?

A) The ability to recover from multiple failures

B) The ability to recover data while other transactions are in progress

C) The ability to execute transactions concurrently

D) The ability to recover lost data

Answer: B) The ability to recover data while other transactions are in progress

59. What is the primary purpose of a transaction in a database system?

A) To recover lost data

B) To perform data validation

C) To execute queries

D) To group database operations into a single unit

Answer: D) To group database operations into a single unit

60. In the transaction state model, what is the "active" state of a transaction?

A) The transaction has been committed

B) The transaction is in progress

C) The transaction has been aborted

D) The transaction is waiting

Answer: B) The transaction is in progress

61. What is the significance of ensuring atomicity in a database transaction?

A) To maximize data inconsistency

B) To allow partial execution of transactions

C) To improve data retrieval

D) To maintain data consistency

Answer: D) To maintain data consistency

62. What is a "concurrent execution" of transactions in a database system?

A) A single transaction executing sequentially

B) Multiple transactions executing independently

C) A transaction in the "committed" state

D) A transaction that has failed

Answer: B) Multiple transactions executing independently

63. What is the primary goal of serializability in database transactions?

A) To maximize data inconsistency

B) To improve query performance

C) To maintain data consistency

D) To allow concurrent execution of transactions

Answer: C) To maintain data consistency

64. What does "recoverability" ensure in the context of transactions?

A) The ability to execute transactions sequentially

B) The ability to execute transactions in isolation

C) The ability to recover from failures

D) The ability to execute transactions concurrently

Answer: C) The ability to recover from failures

65. How is isolation typically implemented in a database system?

A) Using timestamps

B) Using locks

C) Using validation-based protocols

D) Using log-based recovery

Answer: B) Using locks

66. What is the primary purpose of testing for serializability in database transactions?

A) To maximize data inconsistency

B) To ensure data consistency

C) To validate data integrity

D) To check if concurrent transactions are serializable

Answer: D) To check if concurrent transactions are serializable

67. What is a lock-based protocol in database concurrency control?

A) A protocol that uses timestamps

B) A protocol that uses locks to control access

C) A protocol that uses validation checks

D) A protocol that uses log-based recovery

Answer: B) A protocol that uses locks to control access

68. What is a timestamp-based protocol in database concurrency control?

A) A protocol that uses locks to control access

B) A protocol that uses timestamps to order transactions

C) A protocol that uses validation checks

D) A protocol that uses log-based recovery

Answer: B) A protocol that uses timestamps to order transactions

69. What is the primary goal of validation-based protocols in database concurrency control?

A) To maximize data inconsistency

B) To enhance data retrieval

C) To ensure data consistency

D) To maintain data redundancy

Answer: C) To ensure data consistency

70. What is meant by "multiple granularity" in database concurrency control?

A) The ability to use multiple database systems

B) The ability to control transactions at different levels

C) The ability to perform multiple queries simultaneously

D) The ability to recover from multiple failures

Answer: B) The ability to control transactions at different levels

71. What is the relationship between recovery and atomicity in database transactions?

A) They are unrelated concepts

B) Atomicity is a subset of recovery

C) Recovery is a subset of atomicity

D) They are synonymous terms

Answer: B) Atomicity is a subset of recovery

72. What is the purpose of log-based recovery in database systems?

A) To maximize data redundancy

B) To recover lost data

C) To enhance query performance

D) To improve data retrieval

Answer: B) To recover lost data

73. In log-based recovery, what is a "redo log"?

A) A log that records only failed transactions

B) A log that records successful transactions

C) A log that records changes made by transactions

D) A log that records queries

Answer: C) A log that records changes made by transactions

74. What does "recovery with concurrent transactions" refer to in database management?

A) The ability to recover from multiple failures

B) The ability to recover data while other transactions are in progress

C) The ability to execute transactions concurrently

D) The ability to recover lost data

Answer: B) The ability to recover data while other transactions are in progress

75. What is the primary purpose of a transaction in a database system?

A) To recover lost data

B) To perform data validation

C) To execute queries

D) To group database operations into a single unit

Answer: D) To group database operations into a single unit

76. What is external storage in the context of databases?

A) Temporary storage for query results

B) Storage on external hard drives

C) Storage on RAM

D) Storage on SSDs

Answer: B) Storage on external hard drives

77. Which of the following is a common file organization method in databases?

A) Hash tables

B) Linked lists

C) Stacks

D) Arrays

Answer: A) Hash tables

78. What is the purpose of indexing in database systems?

A) To improve query performance

B) To increase data redundancy

C) To enhance data retrieval

D) To optimize data storage

Answer: A) To improve query performance

79. In the context of indexing, what is a "cluster index"?

A) An index based on a hashing algorithm

B) An index that organizes data in a cluster

C) An index that groups similar data together

D) An index that orders data by a clustering key

Answer: D) An index that orders data by a clustering key

80. What are primary indexes in a database?

A) Indexes that are created first

B) Indexes that are used for primary keys

C) Indexes that contain primary data

D) Indexes that are optional

Answer: B) Indexes that are used for primary keys

81. What is the purpose of secondary indexes in a database?

A) To provide an alternative access path to data

B) To store data in secondary storage

C) To reduce data redundancy

D) To eliminate primary indexes

Answer: A) To provide an alternative access path to data

82. Which data structure is commonly used for implementing indexes?

A) Linked lists

B) Arrays

C) Hash tables

D) Trees

Answer: D) Trees

83. What is hash-based indexing in database systems?

A) Indexing based on alphabetical order

B) Indexing using a cryptographic hash function

C) Indexing based on the data's hash code

D) Indexing using a linear hash function

Answer: C) Indexing based on the data's hash code

84. What is tree-based indexing in databases?

A) Indexing using a binary tree structure

B) Indexing based on tree rings

C) Indexing using a tree hierarchy

D) Indexing based on a forest structure

Answer: A) Indexing using a binary tree structure

85. When comparing file organizations, which aspect is typically evaluated?

A) Data redundancy

B) Query performance

C) Data storage cost

D) Data retrieval speed

Answer: B) Query performance

86. What is the purpose of indexes in a database's performance tuning?

A) To slow down query execution

B) To reduce data redundancy

C) To improve query performance

D) To increase data storage cost

Answer: C) To improve query performance

87. What is a common intuition for understanding tree indexes?

A) Data is organized in a linear fashion

B) Data is arranged in a circular manner

C) Data is structured like a forest

D) Data is organized hierarchically

Answer: D) Data is organized hierarchically

88. What does ISAM stand for in the context of indexed access methods?

A) Indexed Sequential Access Mechanism

B) In-memory Storage and Access Method

C) Indexed Sequential Access Model

D) Indexed Storage and Allocation Method

Answer: A) Indexed Sequential Access Mechanism

89. What is a B+ tree in the context of database indexing?

A) A binary tree with balanced height

B) A binary tree with unbalanced height

C) A tree structure with branching factors

D) A tree structure with only leaf nodes

Answer: A) A binary tree with balanced height

90. What is the primary goal of external storage in a database system?

A) To store temporary data

B) To store data permanently

C) To store query results

D) To store data in RAM

Answer: B) To store data permanently

91. Which file organization method is commonly used for indexing?

A) Linked lists

B) Stacks

C) Hash tables

D) Trees

Answer: D) Trees

92. How does indexing impact data retrieval in a database?

A) It slows down data retrieval

B) It has no impact on data retrieval

C) It enhances data retrieval

D) It increases data redundancy

Answer: C) It enhances data retrieval

93. What does a cluster index do in a database?

A) It organizes data into clusters

B) It groups similar data together

C) It orders data by a clustering key

D) It eliminates data redundancy

Answer: C) It orders data by a clustering key

94. What distinguishes primary indexes from other types of indexes in a database?

A) They are created first

B) They contain primary data

C) They are optional

D) They are used for primary keys

Answer: D) They are used for primary keys

95. What is the primary function of secondary indexes in a database?

A) To provide an alternative access path to data

B) To store data in secondary storage

C) To reduce data redundancy

D) To eliminate primary indexes

Answer: A) To provide an alternative access path to data

96. Which data structure is commonly used for implementing indexes in a database?

A) Linked lists

B) Arrays

C) Hash tables

D) Trees

Answer: D) Trees

97. What characterizes hash-based indexing in a database?

A) Indexing based on alphabetical order

B) Indexing using a cryptographic hash function

C) Indexing based on the data's hash code

D) Indexing using a linear hash function

Answer: C) Indexing based on the data's hash code

98. What is tree-based indexing in databases?

A) Indexing using a binary tree structure

B) Indexing based on tree rings

C) Indexing using a tree hierarchy

D) Indexing based on a forest structure

Answer: A) Indexing using a binary tree structure

99. When comparing file organizations, what is typically evaluated?

A) Data redundancy

B) Query performance

C) Data storage cost

D) Data retrieval speed

Answer: B) Query performance

100. How do indexes contribute to performance tuning in a database?

A) By slowing down query execution

B) By reducing data redundancy

C) By improving query performance

D) By increasing data storage cost

Answer: C) By improving query performance

101. What is a common intuition for understanding tree indexes?

A) Data is organized in a linear fashion

B) Data is arranged in a circular manner

C) Data is structured like a forest

D) Data is organized hierarchically

Answer: D) Data is organized hierarchically

102. What does ISAM stand for in the context of indexed access methods?

A) Indexed Sequential Access Mechanism

B) In-memory Storage and Access Method

C) Indexed Sequential Access Model

D) Indexed Storage and Allocation Method

Answer: A) Indexed Sequential Access Mechanism

103. What is a B+ tree in the context of database indexing?

A) A binary tree with balanced height

B) A binary tree with unbalanced height

C) A tree structure with branching factors

D) A tree structure with only leaf nodes

Answer: A) A binary tree with balanced height

104. In the context of indexing, what is the purpose of a hash function?

A) To organize data alphabetically

B) To calculate a checksum for data

C) To generate a unique key for data

D) To compress data for storage

Answer: C) To generate a unique key for data

105. What type of data structure is often used to implement B+ trees in databases?

A) Linked lists

B) Arrays

C) Stacks

D) Node-based structures

Answer: D) Node-based structures

106. What does "ISAM" stand for in the context of database indexing?

A) Indexed Sequential Access Mechanism

B) In-memory Storage and Access Model

C) Indexed Sequential Allocation Method

D) Indexed Storage and Allocation Mechanism

Answer: A) Indexed Sequential Access Mechanism

107. In a B+ tree, where are the data values typically stored?

A) In the internal nodes

B) In the leaf nodes

C) In separate data files

D) In the root node

Answer: B) In the leaf nodes

108. How do B+ trees contribute to efficient data retrieval in databases?

A) They have unbalanced heights

B) They store data redundantly

C) They use a linear search approach

D) They maintain a balanced structure

Answer: D) They maintain a balanced structure

109. What is the primary goal of using external storage in a database system?

A) To store query results

B) To store temporary data

C) To improve data retrieval speed

D) To store data permanently

Answer: D) To store data permanently

110. Which file organization method is suitable for high-speed data retrieval?

A) Linked lists

B) Stacks

C) Hash tables

D) Trees

Answer: D) Trees

111. In the context of indexing, what is the primary role of a clustering key?

A) To organize data into clusters

B) To eliminate data redundancy

C) To group similar data together

D) To calculate hash codes

Answer: A) To organize data into clusters

112. What is external storage in the context of databases?

A) Temporary storage for query results

B) Storage on external hard drives

C) Storage on RAM

D) Storage on SSDs

Answer: B) Storage on external hard drives

113. Which of the following is a common file organization method in databases?

A) Hash tables

B) Linked lists

C) Stacks

D) Arrays

Answer: A) Hash tables

114. What is the purpose of indexing in database systems?

A) To improve query performance

B) To increase data redundancy

C) To enhance data retrieval

D) To optimize data storage

Answer: A) To improve query performance

115. In the context of indexing, what is a "cluster index"?

A) An index based on a hashing algorithm

B) An index that organizes data in a cluster

C) An index that groups similar data together

D) An index that orders data by a clustering key

Answer: D) An index that orders data by a clustering key

116. What are primary indexes in a database?

A) Indexes that are created first

B) Indexes that are used for primary keys

C) Indexes that contain primary data

D) Indexes that are optional

Answer: B) Indexes that are used for primary keys

117. What is the purpose of secondary indexes in a database?

A) To provide an alternative access path to data

B) To store data in secondary storage

C) To reduce data redundancy

D) To eliminate primary indexes

Answer: A) To provide an alternative access path to data

118. Which data structure is commonly used for implementing indexes?

A) Linked lists

B) Arrays

C) Hash tables

D) Trees

Answer: D) Trees

119. What is hash-based indexing in database systems?

A) Indexing based on alphabetical order

B) Indexing using a cryptographic hash function

C) Indexing based on the data's hash code

D) Indexing using a linear hash function

Answer: C) Indexing based on the data's hash code

120. What is tree-based indexing in databases?

A) Indexing using a binary tree structure

B) Indexing based on tree rings

C) Indexing using a tree hierarchy

D) Indexing based on a forest structure

Answer: A) Indexing using a binary tree structure

121. When comparing file organizations, which aspect is typically evaluated?

A) Data redundancy

B) Query performance

C) Data storage cost

D) Data retrieval speed

Answer: B) Query performance

122. What is the purpose of indexes in a database's performance tuning?

A) To slow down query execution

B) To reduce data redundancy

C) To improve query performance

D) To increase data storage cost

Answer: C) To improve query performance

123. What is a common intuition for understanding tree indexes?

A) Data is organized in a linear fashion

B) Data is arranged in a circular manner

C) Data is structured like a forest

D) Data is organized hierarchically

Answer: D) Data is organized hierarchically

124. What does ISAM stand for in the context of indexed access methods?

A) Indexed Sequential Access Mechanism

B) In-memory Storage and Access Method

C) Indexed Sequential Access Model

D) Indexed Storage and Allocation Method

Answer: A) Indexed Sequential Access Mechanism

125. What is a B+ tree in the context of database indexing?

A) A binary tree with balanced height

B) A binary tree with unbalanced height

C) A tree structure with branching factors

D) A tree structure with only leaf nodes

Answer: A) A binary tree with balanced height