

CS403

Database Management System

Important mcqs

Lec 23 - Physical Record and De-normalization

1. **What is a physical record?**

- a) The conceptual model of a database
- b) The actual data stored in a database on disk
- c) The logical model of a database
- d) The metadata associated with a database

Answer: b) The actual data stored in a database on disk

What is denormalization?

- a) The process of breaking normal form rules in a database to improve performance
- b) The process of improving the normalization of a database
- c) The process of removing redundant data from a database
- d) The process of optimizing database queries

Answer: a) The process of breaking normal form rules in a database to improve performance

What is the purpose of denormalization?

- a) To reduce the complexity of database queries
- b) To improve database performance
- c) To simplify the database design
- d) To increase data consistency in a database

Answer: b) To improve database performance

What are the risks of denormalization?

- a) Data redundancy and inconsistency
- b) Database performance degradation
- c) Increased query complexity
- d) All of the above

Answer: d) All of the above

What is data redundancy?

- a) The process of breaking normal form rules in a database to improve performance
- b) The duplication of data in a database
- c) The process of normalizing a database
- d) The metadata associated with a database

Answer: b) The duplication of data in a database

Which normal form does denormalization violate?

- a) First Normal Form
- b) Second Normal Form
- c) Third Normal Form

d) Fourth Normal Form

Answer: c) Third Normal Form

What is the primary goal of normalization?

- a) To eliminate data redundancy
- b) To improve database performance
- c) To simplify the database design
- d) To increase data inconsistency in a database

Answer: a) To eliminate data redundancy

Which of the following is a disadvantage of denormalization?

- a) Reduced query complexity
- b) Increased data redundancy
- c) Improved database performance
- d) Simplified database design

Answer: b) Increased data redundancy

What is the consequence of data inconsistency?

- a) Increased query performance
- b) Improved data quality
- c) Increased risk of errors and inaccuracies
- d) Simplified database design

Answer: c) Increased risk of errors and inaccuracies

Which of the following is an example of denormalization?

- a) Combining two tables into one
- b) Normalizing a database
- c) Adding a new column to a table
- d) Creating a new index on a table

Answer: a) Combining two tables into one

Lec 24 - Vertical Partitioning

1. What is vertical partitioning in a database?

- A) Dividing tables horizontally
- B) Dividing tables vertically
- C) Dividing data into multiple databases
- D) None of the above

Answer: B) Dividing tables vertically

Which of the following is a benefit of vertical partitioning?

- A) Improved data redundancy
- B) Improved data consistency
- C) Improved query performance
- D) Reduced storage space

Answer: C) Improved query performance

What is another name for vertical partitioning?

- A) Row partitioning
- B) Column partitioning
- C) Table partitioning
- D) None of the above

Answer: B) Column partitioning

When is vertical partitioning particularly useful?

- A) When a database contains a large number of rows
- B) When a database contains a small number of columns
- C) When certain columns are accessed more frequently than others
- D) When a database is used for infrequent data access

Answer: C) When certain columns are accessed more frequently than others

Which of the following is a potential drawback of vertical partitioning?

- A) Increased storage space
- B) Reduced query performance
- C) Increased data redundancy
- D) Reduced data consistency

Answer: A) Increased storage space

What is the main goal of vertical partitioning?

- A) To improve data redundancy
- B) To improve query performance
- C) To improve data consistency
- D) To reduce storage space

Answer: B) To improve query performance

Which of the following is an example of vertical partitioning?

- A) Splitting a table into multiple tables based on a date range
- B) Splitting a table into multiple tables based on location
- C) Splitting a table into multiple tables based on column values
- D) None of the above

Answer: C) Splitting a table into multiple tables based on column values

What is a potential challenge of managing a vertically partitioned database?

- A) Ensuring data consistency across partitions

- B) Managing large amounts of data within a single table
- C) Maintaining sufficient storage capacity
- D) None of the above

Answer: A) Ensuring data consistency across partitions

Which of the following is a benefit of vertical partitioning in a distributed database?

- A) Improved query performance
- B) Improved data redundancy
- C) Improved data consistency
- D) None of the above

Answer: A) Improved query performance

What is a key consideration when deciding whether to use vertical partitioning?

- A) The number of rows in the database
- B) The number of columns in the database
- C) The access patterns for the data
- D) The amount of available storage space

Answer: C) The access patterns for the data

Lec 25 - Rules of SQL Format

1. **What is the recommended way to format SQL keywords in a query?**

- a) All lowercase
- b) All uppercase
- c) Capitalize only the first letter
- d) It doesn't matter

Answer: b) All uppercase

What is the recommended way to format table and column aliases in a query?

- a) Use short, cryptic aliases to save space
- b) Use long, descriptive aliases to make the code more readable
- c) Don't use aliases at all
- d) It doesn't matter

Answer: b) Use long, descriptive aliases to make the code more readable

What is the recommended way to format subqueries in a query?

- a) Include them within the main query on a single line
- b) Break them out into separate, indented lines
- c) Don't use subqueries
- d) It doesn't matter

Answer: b) Break them out into separate, indented lines

What is the recommended way to format the SELECT clause in a query?

- a) Include all columns on a single line
- b) Break the columns out into separate lines
- c) Use a combination of both, depending on the length of the query
- d) It doesn't matter

Answer: b) Break the columns out into separate lines

What is the recommended way to format the FROM clause in a query?

- a) Include all tables on a single line
- b) Break the tables out into separate lines
- c) Use a combination of both, depending on the length of the query
- d) It doesn't matter

Answer: b) Break the tables out into separate lines

What is the recommended way to format the WHERE clause in a query?

- a) Include all conditions on a single line
- b) Break the conditions out into separate, indented lines
- c) Don't use the WHERE clause
- d) It doesn't matter

Answer: b) Break the conditions out into separate, indented lines

What is the recommended way to format comments in SQL code?

- a) Use single-line comments starting with #
- b) Use multi-line comments starting with /*
- c) Use both single-line and multi-line comments, depending on the context
- d) It doesn't matter

Answer: c) Use both single-line and multi-line comments, depending on the context

What is the recommended way to format keywords in a query that are not SQL

keywords?

- a) Use all lowercase
- b) Use all uppercase
- c) Capitalize only the first letter
- d) It doesn't matter

Answer: a) Use all lowercase

What is the recommended way to format a JOIN clause in a query?

- a) Include all tables on a single line
- b) Break the tables out into separate lines
- c) Use a combination of both, depending on the length of the query
- d) It doesn't matter

Answer: b) Break the tables out into separate lines

What is the recommended way to format a GROUP BY clause in a query?

- a) Include all columns on a single line
- b) Break the columns out into separate lines
- c) Use a combination of both, depending on the length of the query
- d) It doesn't matter

Answer: b) Break the columns out into separate lines

Lec 26 - Categories of SQL Commands

1. Which category of SQL commands is used to define and modify database objects such as tables and indexes?

- A. Data Manipulation Language (DML)
- B. Data Definition Language (DDL)
- C. Data Control Language (DCL)
- D. Transaction Control Language (TCL)

Answer: B

Which category of SQL commands is used to manipulate data in database objects?

- A. Data Manipulation Language (DML)
- B. Data Definition Language (DDL)
- C. Data Control Language (DCL)
- D. Transaction Control Language (TCL)

Answer: A

Which category of SQL commands is used to control access to the database?

- A. Data Manipulation Language (DML)
- B. Data Definition Language (DDL)
- C. Data Control Language (DCL)
- D. Transaction Control Language (TCL)

Answer: C

Which category of SQL commands is used to manage transactions and ensure data consistency?

- A. Data Manipulation Language (DML)
- B. Data Definition Language (DDL)
- C. Data Control Language (DCL)
- D. Transaction Control Language (TCL)

Answer: D

Which SQL command falls under the DDL category?

- A. SELECT
- B. UPDATE
- C. INSERT
- D. CREATE

Answer: D

Which SQL command falls under the DML category?

- A. ALTER
- B. DELETE
- C. DROP
- D. GRANT

Answer: B

Which SQL command falls under the DCL category?

- A. UPDATE
- B. REVOKE
- C. INSERT

D. DELETE

Answer: B

Which SQL command falls under the TCL category?

A. COMMIT

B. SELECT

C. WHERE

D. FROM

Answer: A

Which SQL command is used to create a new table in a database?

A. INSERT

B. UPDATE

C. DELETE

D. CREATE

Answer: D

Which SQL command is used to add a new row of data to an existing table?

A. DELETE

B. INSERT

C. UPDATE

D. ALTER

Answer: B

Lec 27 - Alter Table Statement

1. Which SQL command is used to modify the structure of an existing table?

- a. SELECT
- b. INSERT
- c. UPDATE
- d. ALTER TABLE

Answer: d. ALTER TABLE

What is the purpose of the ALTER TABLE statement in SQL?

- a. To create a new table
- b. To delete a table
- c. To modify the structure of an existing table
- d. To insert data into a table

Answer: c. To modify the structure of an existing table

Which keyword is used to add a new column to an existing table using the ALTER TABLE statement?

- a. ADD
- b. MODIFY
- c. DROP
- d. RENAME

Answer: a. ADD

Which keyword is used to modify the data type of a column using the ALTER TABLE statement?

- a. ADD
- b. MODIFY
- c. DROP
- d. RENAME

Answer: b. MODIFY

Which keyword is used to delete a column from an existing table using the ALTER TABLE statement?

- a. ADD
- b. MODIFY
- c. DROP
- d. RENAME

Answer: c. DROP

Which keyword is used to rename an existing table using the ALTER TABLE statement?

- a. ADD
- b. MODIFY
- c. DROP
- d. RENAME

Answer: d. RENAME

Which keyword is used to set a primary key constraint on a column using the ALTER TABLE statement?

- a. PRIMARY KEY

- b. FOREIGN KEY
- c. UNIQUE
- d. CHECK

Answer: a. PRIMARY KEY

Which keyword is used to add a foreign key constraint to an existing table using the ALTER TABLE statement?

- a. PRIMARY KEY
- b. FOREIGN KEY
- c. UNIQUE
- d. CHECK

Answer: b. FOREIGN KEY

Which keyword is used to set a unique constraint on a column using the ALTER TABLE statement?

- a. PRIMARY KEY
- b. FOREIGN KEY
- c. UNIQUE
- d. CHECK

Answer: c. UNIQUE

Which keyword is used to add a check constraint to an existing table using the ALTER TABLE statement?

- a. PRIMARY KEY
- b. FOREIGN KEY
- c. UNIQUE
- d. CHECK

Answer: d. CHECK

Lec 28 - Attribute Alias

1. What is attribute aliasing?

- a) A technique for renaming tables in a database
- b) The practice of giving multiple names to the same attribute in a database
- c) The process of aggregating data from multiple sources
- d) The use of machine learning algorithms to identify data patterns

Answer: b) The practice of giving multiple names to the same attribute in a database

What is the potential risk of attribute aliasing?

- a) Inaccurate data analysis
- b) Increased data security
- c) Faster query processing
- d) Improved data quality

Answer: a) Inaccurate data analysis

How can attribute aliasing be mitigated?

- a) By renaming all attributes in a database
- b) By documenting attribute aliases and ensuring consistency
- c) By removing all attribute aliases from a database
- d) By using a different database management system

Answer: b) By documenting attribute aliases and ensuring consistency

What is the impact of inconsistent attribute aliases?

- a) Faster query processing
- b) Improved data quality
- c) Confusion and errors when querying data
- d) More accurate data analysis

Answer: c) Confusion and errors when querying data

How can attribute aliases be managed effectively?

- a) By allowing users to create their own aliases
- b) By using different attribute names for different users
- c) By defining standard attribute names and enforcing them
- d) By allowing attribute aliases to be randomly generated

Answer: c) By defining standard attribute names and enforcing them

What is the role of data governance in attribute aliasing?

- a) To create more attribute aliases
- b) To remove all attribute aliases from a database
- c) To document attribute aliases and enforce consistency
- d) To randomly generate attribute aliases

Answer: c) To document attribute aliases and enforce consistency

What is the difference between an alias and a synonym?

- a) They are interchangeable terms
- b) An alias is a temporary name, while a synonym is a permanent name
- c) An alias is a name for a table, while a synonym is a name for a column
- d) An alias is a name for an attribute, while a synonym is a name for a table or view

Answer: d) An alias is a name for an attribute, while a synonym is a name for a table or view

Which of the following is an example of attribute aliasing?

- a) Using different attribute names in different databases

- b) Renaming a column in a table
- c) Combining data from multiple tables
- d) Creating a new database from scratch

Answer: b) Renaming a column in a table

How can attribute aliases affect data integration?

- a) By improving data integration
- b) By causing data integration to fail
- c) By speeding up data integration
- d) By reducing the need for data integration

Answer: b) By causing data integration to fail

What is the primary goal of managing attribute aliases?

- a) To increase data security
- b) To improve query performance
- c) To ensure consistency and accuracy in data analysis
- d) To reduce the amount of data stored in a database

Answer: c) To ensure consistency and accuracy in data analysis

Lec 29 - Data Manipulation Language

1. Which of the following DML command is used to insert new data into a table?

- a) SELECT
- b) INSERT
- c) UPDATE
- d) DELETE

Solution: b) INSERT

Which DML command is used to retrieve data from a table?

- a) INSERT
- b) UPDATE
- c) DELETE
- d) SELECT

Solution: d) SELECT

Which DML command is used to modify existing data in a table?

- a) INSERT
- b) UPDATE
- c) DELETE
- d) SELECT

Solution: b) UPDATE

Which DML command is used to delete data from a table?

- a) INSERT
- b) UPDATE
- c) DELETE
- d) SELECT

Solution: c) DELETE

Which of the following is not a type of SQL join used in DML?

- a) INNER JOIN
- b) OUTER JOIN
- c) UNION JOIN
- d) FULL OUTER JOIN

Solution: c) UNION JOIN

Which DML command is used to add a new column to an existing table?

- a) ADD COLUMN
- b) INSERT COLUMN
- c) UPDATE COLUMN
- d) ALTER TABLE

Solution: d) ALTER TABLE

Which DML command is used to rename an existing table?

- a) RENAME TABLE

- b) UPDATE TABLE
- c) ALTER TABLE
- d) MODIFY TABLE

Solution: a) RENAME TABLE

Which DML command is used to change the data type of an existing column in a table?

- a) MODIFY COLUMN
- b) CHANGE COLUMN
- c) ALTER COLUMN
- d) UPDATE COLUMN

Solution: c) ALTER COLUMN

Which of the following is not a valid data type in SQL used in DML?

- a) FLOAT
- b) BOOLEAN
- c) DECIMAL
- d) ARRAY

Solution: d) ARRAY

Which DML command is used to update multiple rows in a table at once?

- a) UPDATE
- b) INSERT
- c) DELETE
- d) REPLACE

Solution: a) UPDATE

Lec 30 - ORDER BY Clause

1. **What is the purpose of the ORDER BY clause in SQL?**

- A) To filter data based on specific criteria
- B) To group data based on specific columns
- C) To sort data in ascending or descending order
- D) To combine data from two or more tables

Answer: C) To sort data in ascending or descending order

Which keyword is used in the ORDER BY clause to sort data in descending order?

- A) DESC
- B) ASC
- C) ORDER
- D) BY

Answer: A) DESC

Can the ORDER BY clause be used with the SELECT statement?

- A) Yes
- B) No

Answer: A) Yes

Which clause comes after the ORDER BY clause in a SELECT statement?

- A) FROM
- B) WHERE
- C) GROUP BY
- D) HAVING

Answer: LIMIT or OFFSET clauses, but if they're not used then the query ends after ORDER BY clause.

What is the default sorting order used by ORDER BY clause in SQL?

- A) Descending order
- B) Ascending order
- C) No default sorting order
- D) Random order

Answer: B) Ascending order

Can multiple columns be used in the ORDER BY clause?

- A) Yes
- B) No

Answer: A) Yes

Which operator is used to separate multiple columns in the ORDER BY clause?

- A) AND
- B) OR
- C) COMMA
- D) COLON

Answer: C) COMMA

What is the use of the NULLS FIRST keyword in the ORDER BY clause?

- A) It sorts NULL values first in the result set.

- B) It sorts NULL values last in the result set.
- C) It is not a valid keyword in the ORDER BY clause.
- D) It has no effect on the sorting of NULL values.

Answer: A) It sorts NULL values first in the result set.

Can an alias name be used in the ORDER BY clause?

- A) Yes
- B) No

Answer: A) Yes

Can a subquery be used in the ORDER BY clause?

- A) Yes
- B) No

Answer: B) No

Lec 31 - Inner Join

1. What is the purpose of an Inner Join in SQL?

- a. To combine all rows from two tables
- b. To combine matching rows from two tables
- c. To combine non-matching rows from two tables
- d. None of the above

Answer: b. To combine matching rows from two tables

What is the syntax for an Inner Join in SQL?

- a. `SELECT * FROM table1, table2 WHERE table1.column = table2.column`
- b. `SELECT * FROM table1 JOIN table2 ON table1.column = table2.column`
- c. `SELECT * FROM table1 INNER JOIN table2 ON table1.column = table2.column`
- d. Both b and c

Answer: d. Both b and c

Which of the following types of join can result in NULL values?

- a. Inner Join
- b. Left Join
- c. Right Join
- d. Full Outer Join

Answer: d. Full Outer Join

In an Inner Join, what happens if there are duplicate values in the matching columns?

- a. Only one row is returned for each duplicate value
- b. All rows with duplicate values are returned
- c. An error is thrown
- d. None of the above

Answer: b. All rows with duplicate values are returned

Which of the following is an example of an Inner Join condition?

- a. `table1.column1 = table2.column2`
- b. `table1.column1 <> table2.column2`
- c. `table1.column1 > table2.column2`
- d. Both b and c

Answer: a. `table1.column1 = table2.column2`

In an Inner Join, what is the result if there are no matching values in either table?

- a. All rows from both tables are returned
- b. No rows are returned
- c. Only the rows from the first table are returned
- d. Only the rows from the second table are returned

Answer: b. No rows are returned

What is the difference between an Inner Join and a Left Join?

- a. Inner Join returns only matching rows, while Left Join returns all rows from the left table and matching rows from the right table
- b. Inner Join returns all rows from both tables, while Left Join returns only matching rows
- c. Inner Join and Left Join are the same thing
- d. None of the above

Answer: a. Inner Join returns only matching rows, while Left Join returns all rows from the left

table and matching rows from the right table

Which of the following keywords is used in an Inner Join to specify the columns to join on?

- a. ON
- b. WHERE
- c. JOIN
- d. FROM

Answer: a. ON

Which of the following operators is used in an Inner Join to combine multiple conditions?

- a. AND
- b. OR
- c. NOT
- d. XOR

Answer: a. AND

Which of the following statements is true about the order of tables in an Inner Join?

- a. The order does not matter
- b. The first table listed is always the left table
- c. The second table listed is always the right table
- d. Both b and c

Answer: a. The order does not matter.

Lec 32 - Application Programs

1. Which of the following is an example of an application program?

- a) Microsoft Windows operating system
- b) Google Chrome web browser
- c) Apache web server
- d) Python programming language

Answer: b) Google Chrome web browser

Which type of application program is used for creating and editing documents?

- a) Database management system
- b) Word processor
- c) Spreadsheet program
- d) Graphics software

Answer: b) Word processor

Which of the following is not an example of an application program?

- a) Microsoft Office Suite
- b) Adobe Photoshop
- c) Linux operating system
- d) QuickBooks accounting software

Answer: c) Linux operating system

Which type of application program is used for managing and organizing data?

- a) Word processor
- b) Database management system
- c) Web browser
- d) Media player

Answer: b) Database management system

Which type of application program is used for creating and delivering presentations?

- a) Graphics software
- b) Web browser
- c) Presentation software
- d) Video editing software

Answer: c) Presentation software

Which of the following is an example of open-source application software?

- a) Microsoft Office Suite
- b) Adobe Photoshop
- c) GIMP image editor
- d) Oracle Database

Answer: c) GIMP image editor

Which type of application program is used for performing complex mathematical calculations and analysis?

- a) Word processor
- b) Database management system
- c) Spreadsheet program
- d) Graphics software

Answer: c) Spreadsheet program

Which of the following is an example of proprietary application software?

- a) Apache web server

- b) Linux operating system
- c) QuickBooks accounting software
- d) GIMP image editor

Answer: c) QuickBooks accounting software

Which type of application program is used for managing and organizing email?

- a) Web browser
- b) Media player
- c) Email client
- d) Presentation software

Answer: c) Email client

Which of the following is an example of an educational application program?

- a) Adobe Photoshop
- b) Microsoft Office Suite
- c) Moodle learning management system
- d) QuickBooks accounting software

Answer: c) Moodle learning management system

Lec 33 - Designing Input Form

1. Which of the following is not a consideration when designing an input form?

- a) User needs
- b) Interface design principles
- c) Number of fields
- d) Security measures

Answer: c

What is the purpose of clear labels in an input form?

- a) To make the form look more attractive
- b) To confuse the user
- c) To provide guidance to the user
- d) None of the above

Answer: c

Which of the following should be provided for form fields where necessary?

- a) Validation
- b) Error messages
- c) Both a and b
- d) None of the above

Answer: c

What is the benefit of optimizing an input form for mobile devices?

- a) Increased user engagement
- b) Improved accessibility
- c) Higher conversion rates
- d) All of the above

Answer: d

What is the purpose of security measures in an input form?

- a) To protect user data
- b) To make the form look more professional
- c) To slow down form submission
- d) None of the above

Answer: a

Which of the following is not a principle of interface design?

- a) Consistency
- b) Simplicity
- c) Complexity
- d) Clarity

Answer: c

What is the purpose of appropriate sizing and spacing of form fields?

- a) To make the form look more attractive
- b) To improve usability
- c) To confuse the user
- d) None of the above

Answer: b

Which of the following is not a benefit of a well-designed input form?

- a) Improved user experience

- b) Increased accuracy of data collection
- c) Decreased loading time
- d) All of the above

Answer: c

Which of the following should be optimized in an input form for accessibility?

- a) Color contrast
- b) Font size
- c) Keyboard accessibility
- d) All of the above

Answer: d

What is the benefit of providing real-time feedback on form submission?

- a) Increased user engagement
- b) Improved user experience
- c) Higher conversion rates
- d) All of the above

Answer: d

Lec 34 - Data Storage Concepts

1. Which of the following is an example of structured data storage?

- A) Relational database
- B) Cloud storage
- C) Object storage
- D) NoSQL database

Answer: A

Which type of database allows for flexible and scalable data storage?

- A) Relational database
- B) Object-oriented database
- C) NoSQL database
- D) Hierarchical database

Answer: C

Which of the following is a disadvantage of using tape storage for backup?

- A) It is expensive
- B) It is not durable
- C) It has slow access times
- D) It requires regular maintenance

Answer: C

Which type of storage is ideal for storing large amounts of unstructured data?

- A) Block storage
- B) File storage
- C) Object storage
- D) Cache storage

Answer: C

Which of the following is a cloud storage provider?

- A) Oracle
- B) Amazon Web Services
- C) SAP
- D) IBM

Answer: B

Which of the following is an example of primary storage?

- A) Hard disk drive
- B) USB flash drive
- C) Magnetic tape
- D) RAM

Answer: D

Which of the following is a benefit of using RAID (redundant array of independent disks) for data storage?

- A) Faster data access times
- B) Increased data security
- C) Lower cost
- D) More efficient data backup

Answer: B

Which of the following is a common data storage protocol?

- B) FTP
- C) SMTP
- D) TCP/IP

Answer: B

Which of the following is a benefit of using a content delivery network (CDN) for data storage?

- A) Faster data access times
- B) Lower cost
- C) More efficient backup and recovery
- D) Increased data security

Answer: A

Which of the following is a common backup and recovery method?

- A) RAID
- B) Tape backup
- C) Cloud storage
- D) Object storage

Answer: B

Lec 35 - File Organizations

1. Which of the following file organizations is suitable for storing large amounts of data that need to be accessed in a sequential manner?

- a) Sequential organization
- b) Direct organization
- c) Indexed organization
- d) Hashed organization

Answer: a) Sequential organization

Which of the following file organizations is best suited for storing data that is frequently accessed randomly?

- a) Sequential organization
- b) Direct organization
- c) Indexed organization
- d) Hashed organization

Answer: b) Direct organization

In which of the following file organizations is data accessed using a key value?

- a) Sequential organization
- b) Direct organization
- c) Indexed organization
- d) Hashed organization

Answer: c) Indexed organization

Which of the following file organizations allows for faster access to data by using a hash function?

- a) Sequential organization
- b) Direct organization
- c) Indexed organization
- d) Hashed organization

Answer: d) Hashed organization

Which of the following file organizations requires the least amount of storage space?

- a) Sequential organization
- b) Direct organization
- c) Indexed organization
- d) Hashed organization

Answer: a) Sequential organization

Which of the following file organizations provides the fastest access to data?

- a) Sequential organization
- b) Direct organization
- c) Indexed organization
- d) Hashed organization

Answer: d) Hashed organization

Which of the following file organizations can be used for both sequential and random access to data?

- a) Sequential organization

- b) Direct organization
- c) Indexed organization
- d) Hashed organization

Answer: c) Indexed organization

Which of the following file organizations is not suitable for storing data that needs to be accessed randomly?

- a) Sequential organization
- b) Direct organization
- c) Indexed organization
- d) Hashed organization

Answer: a) Sequential organization

Which of the following file organizations provides faster access to data than sequential organization but slower than direct organization?

- a) Sequential organization
- b) Direct organization
- c) Indexed organization
- d) Hashed organization

Answer: c) Indexed organization

Which of the following file organizations provides a balance between storage space and access time?

- a) Sequential organization
- b) Direct organization
- c) Indexed organization
- d) Hashed organization

Answer: b) Direct organization

Lec 36 - Hashing

1. What is hashing?

- a) The process of encrypting data
- b) The process of converting data into a fixed-length value or key
- c) The process of compressing data
- d) The process of obfuscating data

Answer: b

What is a hash function?

- a) A function that converts data into a fixed-length value or key
- b) A function that compresses data
- c) A function that encrypts data
- d) A function that obfuscates data

Answer: a

What is the purpose of a hash function?

- a) To convert data into a fixed-length value or key
- b) To compress data
- c) To encrypt data
- d) To obfuscate data

Answer: a

Which of the following is an example of a hash algorithm?

- a) MD5
- b) RSA
- c) AES
- d) DES

Answer: a

What is a hash collision?

- a) When two different inputs produce the same hash output
- b) When a hash function fails to produce a fixed-length value or key
- c) When a hash function is too slow
- d) When a hash function is too complex

Answer: a

Which of the following is an advantage of hashing?

- a) Hashing allows for easy reverse engineering of data
- b) Hashing provides secure encryption of data
- c) Hashing provides a fixed-length representation of data
- d) Hashing compresses data to save storage space

Answer: c

Which of the following is a common use case for hashing?

- a) Digital signatures
- b) Image compression
- c) Audio encoding
- d) Video transcoding

Answer: a

What is a rainbow table?

- a) A precomputed table of hash values and corresponding input data

- b) A method of decrypting hashed data
- c) A type of hash function
- d) A method of obfuscating data

Answer: a

Which of the following is a potential issue with using hash functions for password storage?

- a) Hash collisions
- b) Slow computation time
- c) Hash cracking through brute force attacks
- d) None of the above

Answer: c

Which of the following is a way to mitigate the issue of hash cracking through brute force attacks?

- a) Using a stronger hash algorithm
- b) Salting the password before hashing
- c) Increasing the size of the hash value
- d) All of the above

Answer: b

Lec 37 - Index

1. What is an index in a database?

- a) A list of all data in the database
- b) A data structure used to improve data retrieval
- c) A backup of the entire database
- d) A list of all queries run on the database

Answer: b) A data structure used to improve data retrieval

Which of the following is a common type of database index?

- a) Binary search tree
- b) Linked list
- c) Queue
- d) Stack

Answer: a) Binary search tree

What is the purpose of an index in a search engine?

- a) To improve the relevance of search results
- b) To speed up the retrieval of search results
- c) To store all search queries
- d) To store all web pages on the internet

Answer: b) To speed up the retrieval of search results

In a database, what is a primary key index?

- a) An index that contains all data in the database
- b) An index that includes only unique values of a particular column
- c) An index that contains a copy of the entire database
- d) An index that is used to store backup data

Answer: b) An index that includes only unique values of a particular column

Which of the following is a disadvantage of using indexes in a database?

- a) Improved data retrieval performance
- b) Increased storage requirements
- c) Increased security risks
- d) Reduced data consistency

Answer: b) Increased storage requirements

What is an inverted index used for?

- a) Indexing databases
- b) Indexing search engines
- c) Indexing file systems
- d) Indexing social media platforms

Answer: b) Indexing search engines

What is a clustered index in a database?

- a) An index that stores data in a specific order based on a particular column
- b) An index that stores all data in the database
- c) An index that includes only unique values of a particular column
- d) An index that is used to store backup data

Answer: a) An index that stores data in a specific order based on a particular column

What is a non-clustered index in a database?

- a) An index that stores data in a specific order based on a particular column

- b) An index that stores all data in the database
- c) An index that includes only unique values of a particular column
- d) An index that is used to store backup data

Answer: c) An index that includes only unique values of a particular column

Which of the following is a common type of file system index?

- a) Binary search tree
- b) Linked list
- c) Queue
- d) Stack

Answer: b) Linked list

What is a hash index used for?

- a) Indexing databases
- b) Indexing search engines
- c) Indexing file systems
- d) Indexing social media platforms

Answer: a) Indexing databases

Lec 38 - Ordered Indices

1. What are ordered indices?

- A. A type of index used to store data in random order.
- B. A type of index used to store data in a specific order based on one or more columns.
- C. A type of index used only for range queries.
- D. A type of index that is not commonly used in databases.

Answer: B

Which data structure is commonly used to implement ordered indices?

- A. Hash table
- B. Linked list
- C. B-tree
- D. Queue

Answer: C

What is the primary benefit of using an ordered index?

- A. Improved data consistency
- B. Reduced storage requirements
- C. Faster data retrieval based on specific column values
- D. Improved data security

Answer: C

Can multiple ordered indices be created on the same table in a database?

- A. Yes
- B. No

Answer: A

What is a clustered index in a database?

- A. An index used to store data in random order
- B. An index used to store data in a specific order based on one or more columns
- C. An index used only for range queries
- D. An index that is not commonly used in databases

Answer: B

What is the difference between a primary key index and a clustered index?

- A. There is no difference.
- B. A primary key index includes only unique values, while a clustered index can include duplicates.
- C. A primary key index is used to ensure data consistency and integrity, while a clustered index is used for fast retrieval of data in a specific order.
- D. A primary key index is used for range queries, while a clustered index is used for specific value queries.

Answer: C

What is the downside of using too many ordered indices in a database?

- A. Increased storage requirements
- B. Slower data retrieval performance
- C. Increased data inconsistency
- D. Decreased data security

Answer: A

What is a composite index in a database?

- A. An index used to store data in random order

- B. An index used to store data in a specific order based on one column only
- C. An index used to store data in a specific order based on multiple columns
- D. An index used only for range queries

Answer: C

Can an ordered index be used for sorting data?

- A. Yes
- B. No

Answer: A

What is the most common type of data structure used to implement ordered indices?

- A. Hash table
- B. Linked list
- C. B-tree
- D. Queue

Answer: C

Lec 39 - Introduction to Views

1. What is a view in a database?

- a) A physical table
- b) A virtual table
- c) A temporary table
- d) A backup table

Answer: b) A virtual table

How is a view created in a database?

- a) By defining a CREATE TABLE statement
- b) By defining a SELECT statement as a table
- c) By defining a DELETE statement as a table
- d) By defining an ALTER statement as a table

Answer: b) By defining a SELECT statement as a table

What is the purpose of a view in a database?

- a) To display data from multiple tables
- b) To provide an additional layer of security
- c) To simplify complex queries
- d) All of the above

Answer: d) All of the above

Can a view be modified in a database?

- a) Yes, using the ALTER VIEW statement
- b) No, once it has been created it cannot be modified
- c) Yes, using the DROP VIEW statement
- d) No, it can only be dropped and recreated

Answer: a) Yes, using the ALTER VIEW statement

How does a view differ from a table in a database?

- a) A view is a physical table, while a table is a virtual table
- b) A view is a virtual table, while a table is a physical table
- c) A view can be modified, while a table cannot
- d) A view is a backup of a table

Answer: b) A view is a virtual table, while a table is a physical table

Can a view be used to enforce data constraints in a database?

- a) Yes, using the CHECK constraint
- b) No, views cannot enforce data constraints
- c) Yes, using the FOREIGN KEY constraint
- d) No, only tables can enforce data constraints

Answer: b) No, views cannot enforce data constraints

How does a view improve security in a database?

- a) By allowing users to access only the data they are authorized to see
- b) By encrypting all data in the database
- c) By preventing users from accessing any data in the database
- d) By providing a backup of all data in the database

Answer: a) By allowing users to access only the data they are authorized to see

Can a view be used in a JOIN operation in a database?

- a) Yes, views can be used in JOIN operations
- b) No, views cannot be used in JOIN operations
- c) Only if the view is created from a single table
- d) Only if the view is created from multiple tables

Answer: a) Yes, views can be used in JOIN operations

How does a view affect performance in a database?

- a) Views can improve performance by simplifying complex queries
- b) Views can degrade performance by adding an additional layer of complexity
- c) Views have no impact on performance
- d) Views can improve performance by speeding up data retrieval

Answer: b) Views can degrade performance by adding an additional layer of complexity

How is a view dropped in a database?

- a) By using the DROP TABLE statement
- b) By using the DELETE VIEW statement
- c) By using the ALTER VIEW statement
- d) By using the DROP VIEW statement

Answer: d) By using the DROP VIEW statement

Lec 40 - Introduction to Views part-2

1. What is a view in a database?

- a) A physical table
- b) A virtual table
- c) A temporary table
- d) A backup table

Answer: b) A virtual table

How is a view created in a database?

- a) By defining a CREATE TABLE statement
- b) By defining a SELECT statement as a table
- c) By defining a DELETE statement as a table
- d) By defining an ALTER statement as a table

Answer: b) By defining a SELECT statement as a table

What is the purpose of a view in a database?

- a) To display data from multiple tables
- b) To provide an additional layer of security
- c) To simplify complex queries
- d) All of the above

Answer: d) All of the above

Can a view be modified in a database?

- a) Yes, using the ALTER VIEW statement
- b) No, once it has been created it cannot be modified
- c) Yes, using the DROP VIEW statement
- d) No, it can only be dropped and recreated

Answer: a) Yes, using the ALTER VIEW statement

How does a view differ from a table in a database?

- a) A view is a physical table, while a table is a virtual table
- b) A view is a virtual table, while a table is a physical table
- c) A view can be modified, while a table cannot
- d) A view is a backup of a table

Answer: b) A view is a virtual table, while a table is a physical table

Can a view be used to enforce data constraints in a database?

- a) Yes, using the CHECK constraint
- b) No, views cannot enforce data constraints
- c) Yes, using the FOREIGN KEY constraint
- d) No, only tables can enforce data constraints

Answer: b) No, views cannot enforce data constraints

How does a view improve security in a database?

- a) By allowing users to access only the data they are authorized to see
- b) By encrypting all data in the database
- c) By preventing users from accessing any data in the database
- d) By providing a backup of all data in the database

Answer: a) By allowing users to access only the data they are authorized to see

Can a view be used in a JOIN operation in a database?

- a) Yes, views can be used in JOIN operations

- b) No, views cannot be used in JOIN operations
- c) Only if the view is created from a single table
- d) Only if the view is created from multiple tables

Answer: a) Yes, views can be used in JOIN operations

How does a view affect performance in a database?

- a) Views can improve performance by simplifying complex queries
- b) Views can degrade performance by adding an additional layer of complexity
- c) Views have no impact on performance
- d) Views can improve performance by speeding up data retrieval

Answer: b) Views can degrade performance by adding an additional layer of complexity

How is a view dropped in a database?

- a) By using the DROP TABLE statement
- b) By using the DELETE VIEW statement
- c) By using the ALTER VIEW statement
- d) By using the DROP VIEW statement

Answer: d) By using the DROP VIEW statement

Lec 41 - Updating Multiple Tables

1. Which SQL statement is used to update data in multiple tables at once?

- A) UPDATE
- B) ALTER
- C) JOIN
- D) SELECT

Answer: C) JOIN

Which type of join is used to update data in multiple tables based on a matching column?

- A) INNER JOIN
- B) OUTER JOIN
- C) LEFT JOIN
- D) RIGHT JOIN

Answer: A) INNER JOIN

What happens when you attempt to update data in multiple tables without using a join statement?

- A) The update fails
- B) Data in one table is updated but not the other
- C) Data in both tables is updated but the result may be inconsistent
- D) Data in both tables remains unchanged

Answer: C) Data in both tables is updated but the result may be inconsistent

Which keyword is used to specify the table to be updated in an SQL query?

- A) SET
- B) FROM
- C) WHERE
- D) UPDATE

Answer: D) UPDATE

When updating data in multiple tables, which of the following is NOT a recommended practice?

- A) Back up the database before making any changes
- B) Test the update on a small sample of data first
- C) Use a transaction to ensure all changes are made together
- D) Update tables in a random order

Answer: D) Update tables in a random order

Which type of subquery can be used to update data in a table based on values in another table?

- A) Correlated subquery
- B) Non-correlated subquery
- C) Nested subquery
- D) None of the above

Answer: A) Correlated subquery

In which scenario would you use a LEFT JOIN to update data in multiple tables?

- A) When you want to update data only in the first table

- B) When you want to update data only in the second table
- C) When you want to update data in both tables
- D) When you want to update data in the first table and keep all the rows in the second table

Answer: D) When you want to update data in the first table and keep all the rows in the second table

What is the purpose of a foreign key constraint in a database?

- A) To ensure that data in one table matches data in another table
- B) To prevent data duplication in a table
- C) To improve query performance
- D) To encrypt sensitive data

Answer: A) To ensure that data in one table matches data in another table

Which type of join returns all rows from both tables, even if there is no match?

- A) INNER JOIN
- B) OUTER JOIN
- C) LEFT JOIN
- D) RIGHT JOIN

Answer: B) OUTER JOIN

Which SQL statement can be used to roll back changes made during an update operation?

- A) ROLLBACK
- B) COMMIT
- C) SAVEPOINT
- D) UPDATE

Answer: A) ROLLBACK

Lec 42 - The Concept of a Transaction

1. **What does the concept of a transaction refer to in databases?**

- A) A single database operation
- B) A sequence of operations executed as a single unit of work
- C) A table modification
- D) A data query

Answer: B

Why is the concept of a transaction important in databases?

- A) It allows for faster data access
- B) It ensures data consistency and integrity
- C) It reduces database storage requirements
- D) It eliminates the need for data backups

Answer: B

Which of the following is an ACID property that a transaction must follow?

- A) Atomicity
- B) Security
- C) Availability
- D) Performance

Answer: A

What does atomicity mean in the context of transactions?

- A) The transaction must complete successfully or be rolled back completely
- B) The transaction must be isolated from other transactions
- C) The transaction must leave the database in a consistent state
- D) The transaction must be durable

Answer: A

Which ACID property ensures that a transaction leaves the database in a consistent state?

- A) Atomicity
- B) Consistency
- C) Isolation
- D) Durability

Answer: B

What does isolation mean in the context of transactions?

- A) The transaction must complete successfully or be rolled back completely
- B) The transaction must be isolated from other transactions
- C) The transaction must leave the database in a consistent state
- D) The transaction must be durable

Answer: B

Which ACID property ensures that a transaction is isolated from other transactions?

- A) Atomicity
- B) Consistency
- C) Isolation
- D) Durability

Answer: C

What does durability mean in the context of transactions?

- A) The transaction must complete successfully or be rolled back completely

- B) The transaction must be isolated from other transactions
- C) The transaction must leave the database in a consistent state
- D) The transaction must be durable

Answer: D

Which ACID property ensures that a transaction is durable even in the event of system failures?

- A) Atomicity
- B) Consistency
- C) Isolation
- D) Durability

Answer: D

Which of the following is a benefit of using transactions in databases?

- A) Faster data access
- B) Reduced storage requirements
- C) Improved data consistency and integrity
- D) Elimination of the need for data backups

Answer: C

Lec 43 - Incremental Log with Deferred Updates

1. **What is the primary goal of Incremental Log with Deferred Updates in a database management system?**

- a) To improve data consistency
- b) To improve data security
- c) To improve efficiency and performance
- d) To improve user experience

Answer: c) To improve efficiency and performance

What is the purpose of a log file in Incremental Log with Deferred Updates?

- a) To store all database records
- b) To record changes to the database as they occur
- c) To store metadata about the database
- d) To store backup copies of the database

Answer: b) To record changes to the database as they occur

In Incremental Log with Deferred Updates, when are changes applied to the database?

- a) As soon as they are made
- b) At the end of each transaction
- c) At a later time, after being recorded in the log file
- d) None of the above

Answer: c) At a later time, after being recorded in the log file

Which of the following benefits does Incremental Log with Deferred Updates offer?

- a) Improved data security
- b) Improved data consistency
- c) Improved efficiency and performance
- d) Improved ease of use

Answer: c) Improved efficiency and performance

Which of the following statements is true about Incremental Log with Deferred Updates?

- a) It is a technique used only in small databases
- b) It is a technique that must be used with all databases
- c) It is a technique that can improve the performance of large databases
- d) It is a technique that can only be used with certain types of databases

Answer: c) It is a technique that can improve the performance of large databases

What is the primary disadvantage of Incremental Log with Deferred Updates?

- a) It can lead to inconsistencies in the database
- b) It can increase the risk of data loss
- c) It can make it more difficult to manage the database
- d) It can be slower than other techniques

Answer: a) It can lead to inconsistencies in the database

Which of the following is an example of a situation where Incremental Log with Deferred Updates may not be appropriate?

- a) A financial institution that processes thousands of transactions per second
- b) A small business that processes a few transactions per day
- c) A healthcare organization that stores patient records

d) None of the above

Answer: b) A small business that processes a few transactions per day

Which of the following is a benefit of using a log file in Incremental Log with Deferred Updates?

- a) It allows for faster data entry
- b) It provides a backup of the database
- c) It can help to recover from system failures
- d) It reduces the need for database indexing

Answer: c) It can help to recover from system failures

What is the role of the transaction manager in Incremental Log with Deferred Updates?

- a) To manage the log file
- b) To ensure data consistency
- c) To apply changes to the database
- d) To ensure data security

Answer: b) To ensure data consistency

Which of the following best describes how Incremental Log with Deferred Updates improves performance?

- a) By reducing the amount of data that needs to be stored in memory
- b) By reducing the time it takes to apply changes to the database
- c) By reducing the amount of data that needs to be indexed
- d) By reducing the number of transactions that can be processed simultaneously

Answer: b) By reducing the time it takes to apply changes to the database

Lec 44 - Uncommitted Update Problem

1. What is the Uncommitted Update Problem?

- a) When a transaction updates a record, but the update is not yet committed
- b) When a transaction deletes a record, but the deletion is not yet committed
- c) When a transaction reads a record, but the record is not yet committed

Answer: a

What are the consequences of the Uncommitted Update Problem?

- a) Dirty reads
- b) Non-repeatable reads
- c) Phantom reads
- d) All of the above

Answer: d

Which concurrency control mechanisms are used to prevent the Uncommitted Update Problem?

- a) Locking
- b) Timestamps
- c) Both locking and timestamps
- d) None of the above

Answer: c

Which type of read anomaly can occur due to the Uncommitted Update Problem?

- a) Dirty read
- b) Non-repeatable read
- c) Phantom read
- d) All of the above

Answer: a

Which of the following is not a consequence of the Uncommitted Update Problem?

- a) Deadlock
- b) Dirty read
- c) Non-repeatable read
- d) Phantom read

Answer: a

Which of the following is a way to prevent the Uncommitted Update Problem?

- a) Increasing the transaction isolation level
- b) Decreasing the transaction isolation level
- c) Not using any concurrency control mechanism
- d) None of the above

Answer: a

Which of the following is an example of the Uncommitted Update Problem?

- a) A transaction updates a record, but the update is not yet committed
- b) A transaction reads a record, but the record is not yet committed
- c) A transaction deletes a record, but the deletion is not yet committed

Answer: a

Which of the following describes a dirty read?

- a) When a transaction reads a record that has been updated but not yet committed

- b) When a transaction reads a record that has been deleted but not yet committed
- c) When a transaction reads a record that has been inserted but not yet committed

Answer: a

Which of the following describes a non-repeatable read?

- a) When a transaction reads a record that has been updated but not yet committed
- b) When a transaction reads a record that has been deleted but not yet committed
- c) When a transaction reads the same record multiple times and gets different results

Answer: c

Which of the following is a potential consequence of using a low transaction isolation level?

- a) Reduced risk of the Uncommitted Update Problem
- b) Increased risk of the Uncommitted Update Problem
- c) No effect on the risk of the Uncommitted Update Problem

Answer: b

Lec 45 - Locking Idea

1. What is locking in a database management system?

- a. A mechanism to prevent transactions from accessing shared resources
- b. A mechanism to prevent transactions from accessing only exclusive resources
- c. A mechanism to allow transactions to access shared resources simultaneously
- d. A mechanism to allow transactions to access exclusive resources simultaneously

Answer: a

Which of the following is not a type of lock?

- a. Shared lock
- b. Exclusive lock
- c. Read lock
- d. Write lock

Answer: c

What is the purpose of a shared lock?

- a. To allow multiple transactions to read the same resource simultaneously
- b. To prevent multiple transactions from reading the same resource simultaneously
- c. To allow multiple transactions to write to the same resource simultaneously
- d. To prevent multiple transactions from writing to the same resource simultaneously

Answer: a

What is the purpose of an exclusive lock?

- a. To allow multiple transactions to read the same resource simultaneously
- b. To prevent multiple transactions from reading the same resource simultaneously
- c. To allow multiple transactions to write to the same resource simultaneously
- d. To prevent multiple transactions from writing to the same resource simultaneously

Answer: d

What is a deadlock?

- a. A situation where two or more transactions are waiting for each other to release locks
- b. A situation where a transaction is waiting for a lock that is held by another transaction
- c. A situation where a transaction is waiting for a lock that has already been released
- d. A situation where a transaction is waiting for a resource that is not available

Answer: a

What is a timeout in locking?

- a. A mechanism to release a lock after a specified time period
- b. A mechanism to acquire a lock after a specified time period
- c. A mechanism to prevent a transaction from acquiring a lock
- d. A mechanism to prevent a transaction from releasing a lock

Answer: a

What is the purpose of lock escalation?

- a. To reduce the number of locks held by a transaction
- b. To increase the number of locks held by a transaction
- c. To release all locks held by a transaction
- d. To prevent a transaction from acquiring any locks

Answer: a

What is the difference between a shared lock and an exclusive lock?

- a. A shared lock allows multiple transactions to read the same resource, while an exclusive lock

prevents multiple transactions from writing to the same resource simultaneously

b. A shared lock allows multiple transactions to write to the same resource simultaneously, while an exclusive lock prevents multiple transactions from reading the same resource simultaneously

c. A shared lock allows multiple transactions to access the same resource simultaneously, while an exclusive lock prevents any other transactions from accessing the resource until the lock is released

d. A shared lock allows a transaction to access a resource exclusively, while an exclusive lock allows multiple transactions to access the same resource simultaneously

Answer: c

What is the purpose of lock granularity?

a. To determine the size of the resource that will be locked

b. To determine the type of lock that will be used

c. To determine the duration of the lock

d. To determine the transaction that will hold the lock

Answer: a

What is the purpose of a lock manager in a database management system?

a. To manage the acquisition and release of locks

b. To manage the execution of transactions

c. To manage the storage of data

d. To manage the retrieval of data

Answer: a

