Name: _____ **Student ID:** _____

Exam Duration: 60 minutes

Points: 5 points per question

Note: Each question may have multiple correct answers. You can have one page of notes.

1. Why is an ERD primarily useful in database design? (5 points)

A. It guarantees all business rules are enforced in queries.

B. It helps visualize entities and their relationships.

C. It ensures automatic indexing of all tables.

D. It generates efficient SQL queries.

2. How many errors are in this SQL query? (5 points)

Query: SELECT name age FROM users WHERE username = "admin";

A. 0

B. 1

C. 2

D. 3

3. Which set of attributes best fits an ERD for a library system? (5 points)

A. Entities: Book, Author, Library | Attributes: Title, Biography, Location

B. Entities: Member, Staff, Transaction | Attributes: Age, Salary, Timestamp

C. Entities: Book, Member, Loan | Attributes: Title, Name, Due Date

D. Entities: Book, Genre, Publisher | Attributes: Title, Category, Address

4. Which SQL constraints best enforce data integrity? (5 points)

A. UserID: NOT NULL, Email: UNIQUE, Name: PRIMARY KEY

B. UserID: PRIMARY KEY, Email: UNIQUE, Name: NOT NULL

C. UserID: UNIQUE, Email: NOT NULL, Name: UNIQUE

D. UserID: PRIMARY KEY, Email: NOT NULL, Name: UNIQUE

5. Which approach is best for generating unique IDs in one database? (5 points)

A. Auto-increment for sequential numeric IDs

B. Custom ID format based on business rules

C. UUID for globally unique identifiers

D. Using a manually entered primary key

6. How does relational algebra operate on data? (5 points)

A. It manipulates sets of tuples.

B. It processes tuples separately.

C. It splits attributes into relations.

D. It performs only calculations.

7. Which transformation optimizes the following relational algebra query? (5 points)

Query: π_name(σ_age > 30 (Employee ⋈ Department))

A. π_name(Employee ⋈ σ_age > 30 (Department))

B. σ_age > 30 (π_name (Employee ⋈ Department))

C. π_name(σ_age > 30 Employee ⋈ Department)

D. Employee ⋈ π_name(σ_age > 30 Department)

8. Which query correctly retrieves departments with more than 10 employees? (5 points)

A. SELECT department, COUNT*() FROM employees WHERE COUNT()* > 10;

B. SELECT department FROM employees GROUP BY department HAVING COUNT*() > 10;*

C. SELECT department FROM employees WHERE COUNT() > 10 GROUP BY department;

D. SELECT department, COUNT*() FROM employees HAVING COUNT()* > 10 GROUP BY department;

9. Which JOIN retrieves only matching rows from both tables? (5 points)

Query: SELECT employees.name, departments.name FROM employees ___ JOIN departments ON employees.dept_id = departments.id;

A. LEFT

B. RIGHT

C. INNER

D. FULL

10. Which statement correctly describes normalization levels in databases? (5 points)

A. Unnormalized form (UNF) allows duplicate and multi-valued attributes.

B. First normal form (1NF) eliminates duplicate rows but allows repeating groups.

C. Second normal form (2NF) requires 1NF and ensures all non-key attributes depend on the full primary key.

D. Third normal form (3NF) requires 2NF and ensures all attributes depend only on the primary key.

11. Which of the following statements correctly describes a property of database transactions? (5 points)

A. Atomicity ensures a transaction can be divided and executed in parts.

B. Consistency allows the database to reflect temporary states during a transaction.

C. Isolation ensures that transactions do not interfere with each other's intermediate states.

D. Durability means a transaction can be rolled back after it is committed.

12. What does the acronym ACID stand for in database transactions? (5 points)

A. Atomicity, Concurrency, Isolation, Durability

B. Accuracy, Consistency, Integrity, Durability

C. Atomicity, Consistency, Isolation, Durability

D. Automation, Consistency, Integration, Durability

13. What is the primary purpose of Role-Based Access Control (RBAC) in database systems? (5 points)

A. To back up and restore data automatically

B. To encrypt all user passwords in the database

C. To restrict user access based on predefined roles and permissions

D. To normalize the database structure to reduce redundancy

14. What is the main purpose of encryption in database systems? (5 points)

A. To reduce query execution time

B. To compress large data for storage efficiency

C. To prevent unauthorized access by converting data into unreadable formats

D. To automatically generate backups of the database

15. According to the CAP theorem in distributed databases, which three properties cannot be fully achieved at the same time? (5 points)

A. Concurrency, Availability, Performance

B. Consistency, Availability, Partition Tolerance

C. Compatibility, Accuracy, Persistence

D. Connectivity, Access Control, Privacy

16. How does the concept of "Consistency" differ between ACID and CAP in database systems? (5 points)

A. In both ACID and CAP, consistency ensures data is encrypted during storage.

B. In ACID, consistency refers to valid state transitions; in CAP, it means all nodes see the same data at the same time.

C. In ACID, consistency ensures availability across nodes; in CAP, it ensures transaction rollback.

D. In both ACID and CAP, consistency guarantees that operations are isolated from each other.

17. Why are column-oriented databases often preferred over row-oriented databases for aggregation queries? (5 points)

A. Column databases store indexes separately, making joins faster.

B. Column databases compress data better, reducing storage for images.

C. Column databases allow fast access to specific columns needed for aggregation, improving performance.

D. Column databases are optimized for storing relationships between tables.

18. Based Access Control (RBAC) and User-Based Access Control (UBAC)? (5 points)

A. RBAC assigns permissions directly to users, while UBAC uses predefined roles.

B. RBAC is more dynamic because it changes access based on user behavior.

C. UBAC assigns permissions directly to individual users, while RBAC assigns permissions to roles that users are assigned to.

D. UBAC is only used in cloud databases, while RBAC is for local databases.

19. Design an Entity-Relationship Diagram (ERD) for a database that manages customer orders for a plant store. (5 points)
Your ERD should include:

- At least **three entities** (e.g., Customer, Order, Plant)
- Each entity must have **at least four attributes**
- At least **two relationships** between the entities (e.g., a customer places orders, an order includes plants)
  Label primary keys and indicate cardinality for each relationship.

20. If the plant watering system from the team project application fails, what are three potential database-related root causes? (5 points)