KATHMANDU UNIVERSITY
End Semester Examination [C]
June/July 2024

Marks Scored:

Level : B.E./B.Sc.
Year : II

Course : COMP 232
Semester : II

Exam Roll No. :

Time: 30 mins.

F. M. : 10

Registration No.:

Date : 05 JUL 2024

SECTION "A"
[20Q. × 0.5 = 10 marks]

Choose and encircle the most appropriate option from each set of choices

1. What does data independence mean?
   a. Data is defined separately and not included in programs.
   b. Programs are not dependent on the physical attributes of data
   c. Programs are dependent on the logical attributes of data
   d. Data can be stored non-contiguously in the memory.

2. Which database level is closest to the users?
   a. External b. Internal c. Physical d. Conceptual

3. Which of the following is not comparison operator?
   a. <> b. < c. >= d. =<

4. Which of the following in **TRUE**?
   a. In an ER diagram, a weak entity set is indicated by a double-lined rectangle connected to a rectangle.
   b. In an ER diagram, a weak entity set is indicated by a double-lined rectangle connected to a double-lined diamond.
   c. In an ER diagram, a weak entity set is indicated by a double-lined diamond connected to a rectangle.
   d. In an ER diagram, a weak entity set is indicated by a double-lined diamond connected to a double-lined rectangle.

5. _____ is the process of organizing data into related tables.
   a. Normalization b. Generalization
   c. Specialization d. Mapping cardinality

6. For the LIKE operator, which of the following is **NOT** True?
   a. % matches zero or more characters
   b. _ matches exactly one character
   c. \% matches the % symbol
   d. '_ _ _ %' matches any string of at most three characters

7. Centralizing the integrity checking directly under the DBMS _____ duplication and ensures the consistency and validity of the database.
   a. increases b. reduces c. ignores d. ensures

8. Consider the following statements issued by a DBA in that order. If USER1 and PG is a database user and the TABLE1 table is owned by PG, what would be the outcome after executing the statements?

```
CREATE ROLE role1;
GRANT SELECT, INSERT ON pg. table1 TO role1;
GRANT role1 TO user1;
GRANT SELECT ON pg. table1 TO user1;
REVOKE SELECT ON pg.table1 FROM user1;
```

a. The REVOKE statement would give an error because the SELECT privilege has already been granted to the role ROLE1
b. The REVOKE statement would remove the SELECT privilege from SCOTT as well as from the role ROLE1
c. USER1 would be able to query the PG.TABLE1 table
d. USER1 would not be able to query the PG.TABLE1 table

9. Which of the following are introduced to reduce the overheads caused by the log-based recovery?
a. Indices b. Checkpoints c. Deadlocks d. Foreign keys

10. Which of the following statements is a set of one or more attributes taken collectively to uniquely identify a record?
a. Primary key b. Secondary key c. Foreign key d. Index key

11. When transaction Ti requests a data item currently held by Tj, Ti is allowed to wait only if it has a timestamp smaller than that of T_j (that is, T_i is older than T_j). Otherwise, T_i is rolled back (dies). Which scheme is this?
a. Wait-wound b. Wait-die c. Wound-wait d. Wait

12. A table is in 3NF if it is in 2NF and if it has no _____ dependencies
a. transitive b. functional c. trivial d. multi-valued

13. Which of the following operations do not preserve non-matched tuples?
a. Left outer join b. External join c. Inner join d. Right outer join

14. If a transaction has obtained a _____ mode lock, it can both read and write on the item
a. shared b. exclusive c. read-only d. write-only

15. In which of the following phase, the system replays updates of all transactions by scanning the log forward from the last checkpoint?
a. write b. read c. undo d. redo

16. What does the following Cypher query do?

```
CREATE (p:Person)-[:RATES]->(b:Book)
```

a. It creates one RATES node and two relationships of label Person and Book.
b. It creates one node of label Person and connects it to an existing node of label Technology and value t.
c. It creates two nodes and connects them.
d. It looks for the pattern consisting of two nodes connected through a RATES relationship.

17. Which characteristics do document-based NoSQL databases have?
    a. They store data in tabular format
    b. They support transactions and joins
    c. They are best suited for handling structured data
    d. They use a hierarchical, JSON-like data model

18. Which of the following is a benefit of using NoSQL databases?
    a. Strict data modeling
    b. Limited scalability
    c. Easy schema evolution
    d. Reduced data storage capacity

19. Consider the following set F of functional dependencies on schema (A, B, C, D, E). Which of the following functional dependencies cannot be logically implied using Armstrong's axioms?
    $F = \{X \rightarrow Y, Y \rightarrow Z, WX \rightarrow Y, UX \rightarrow W\}$
    a. $X \rightarrow Z$
    b. $WX \rightarrow WZ$
    c. $XYZ \rightarrow W$
    d. $UX \rightarrow Y$

20. What happens if a piece of data is stored in two places in the database?
    a. Concurrent access of data will be easier
    b. Multiple authorization mechanism will be required
    c. Atomicity and isolation will be ensured automatically.
    d. Changing the data in one spot will cause data inconsistency

KATHMANDU UNIVERSITY
End Semester Examination [C]
June/July 2024

Level : B.E./B.Sc.
Year : II
Time : 2 hrs. 30mins.

05 JUL 2024

Course : COMP 232
Semester : II
F. M. : 40

SECTION "B"
[6 Q. × 4 = 24 marks]

Attempt *ANY SIX* questions.

1. Write about different components included in a query processor. Explain major steps involved in query processing. [3+1=4]

2. Define JOIN operation in SQL. Explain different types of JOINs. [1+3=4]

3. Define conflict-serializability. Consider the following schedule with two transactions T_1 and T_2. Is this schedule conflict-serializable? If yes, provide the equivalent serial schedule(s). If no, explain why not. [1+3=4]

Time	T_1	T_2
1	read(A);	
2		read(A);
3		read(B);
4	read(B);	
5	A = A - 5;	
6	write(A);	
7		A = A + 5;
8		write(A);
9		B = B - 1;
10		write(B);
11	if B = B+ 1;	
12	write(B)	

4. Consider the following log records. During the recovery after system failure, which transactions will be redone and which transactions will be undone? What will be the value of the data items A, B, C, and D after recovery? [2+2=4]

<T_0 start>
<T_0, A, 100, 50>
<T_0, B, 20, 40>
<T_1 start>
<T_2 start>
<T_1, C, 100, 400>
<T_0, B, 40, 500>
<T_1 commit>
<T_2, C, 400, 100>
<checkpoint {T_0, T_2}>
<T_3 start>
<T_2 commit>
<T_3, D, 200, 2000>
<T_3, A, 100, 40>
<T_4 start>

5. What are the ACID properties in a database management system? Is it always mandatory to ensure ACID properties in a DBMS? [3+1=4]

6. What are the trade-offs between using relational databases and NoSQL databases for handling large-scale data? [4]

7. Write short notes on ANY TWO: [2+2=4]
   a. Schema-agnostic databases
   b. Data Definition Language (DDL)
   c. Different types of mapping cardinality in a binary relationship set

SECTION "C"

[2 Q. × 8 = 16 marks]

Attempt *ANY TWO* questions.

8. Answer the following questions:
   a. What is database normalization? Explain its importance. [2+2=4]
   b. When should we opt for denormalization? [2]
   c. Determine if the following relation is in BCNF. [2]

ID	Name	Dept_name	Building	Salary
223	Ryan	Physics	5	60000
214	Kim	History	8	75000
412	Tom	Finance	6	45000
645	Alina	Chemistry	4	55000
162	Xi	Computer Science	9	75000
111	John	Physics	5	45000
432	Anne	History	8	23000
125	Wu	Computer Science	9	60000
161	James	Chemistry	4	54000
754	Michael	Computer Science	9	51000

9. Construct an ER diagram and write an SQL script for creating the corresponding database schema for a hospital management system considering the following constraints:
   a. No two doctors can have the same license number.
   b. Patients cannot book multiple appointments for the same time slot.
   c. Appointment status must be one of the following: 'Pending', 'Accepted', 'Declined'
   d. When a patient is deleted, all related appointments, medical records, and prescriptions should be automatically deleted.
   e. Your ER model should have at least 3 entity sets. Clearly show mapping cardinalities, primary keys, and participation types. [4+4=8]

P.T.O.

10. Consider the following tables and answer the following questions:
job (id, title, city, description, min_qualification, min_salary, max_salary)
jobseeker (id, name, age, qualification, city)
job_application (job_id, jobseeker_id, application_date)

i. Give an expression in relational algebra for the following query: [2]
   a. Find the name of all jobseekers who live in Kathmandu.
   b. Find the title of the jobs with the salary range 10000-30000.

ii. Give an expression in SQL for the following queries: [2+2+2=6]
   a. Find all jobs that received more than 100 applications.
   b. Find the name of jobseekers for whom there is at least one job matching their qualification.
   c. List all jobs with the salary in the range 10000-30000.