

When graduate students work on a project, a professor must supervise their work on the project. Graduate students can work on multiple projects; in this case they can have a different supervisor for each one.

Departments have a department number, a department name and a main office. Departments have a professor who runs the department. Professors works in one or more departments, and for each department that they work in a time percentage is associated with their job. Graduate, students have a major department in which they are working on their degree. Each graduate student has another, more senior graduate student (known as student advisor) who advises him or her on what courses to take.

Design/Draw an ER diagram that captures the full information about the university. Use only the basic ER model, here, that is, entities, relationships, and attributes. Be sure to indicate any key and participation constraints with underlines and arrows.

Or

- (b) Consider the below schedule for three transactions. R indicates a read of a page, W indicates a write of a page, and COM indicates a commit.
- (i) Draw the arrows for the dependency graph for this schedule. (4)
 - (ii) Is the schedule conflict serializable? Explain Conflict Serializable. (6)
 - (iii) Which of the following changes to the above schedule will result in a schedule that is possible using strict two-phase locking? (5)

T1	R(A)						R(C)	W(B)	COM		
T2				R(A)	W(A)	COM					
T3		R(A)	R(C)							W(B)	COM

Reg. No. :

Question Paper Code : 20872

B.E./B.Tech. DEGREE EXAMINATIONS, NOVEMBER/DECEMBER 2023.

Third/Fourth Semester

Computer Science and Engineering

CS 3492 – DATABASE MANAGEMENT SYSTEMS

(Common to : Computer Science and Design/Computer and Communication Engineering/Computer Science and Business Systems and Information Technology)

(Regulations – 2021)


Time : Three hours

Maximum : 100 marks

Answer ALL questions.

PART A — (10 × 2 = 20 marks)

1. Relate the terms database, database management system.
2. What is the difference between logical data independence and physical data independence?
3. Define functional dependency.
4. Transform the following ER diagram into a relational schema diagram.



5. State ACID Properties.
6. What is the maximum number of superkeys for the relation schema R(E, F, G, H) with E as the key attribute? Write all such possible superkeys.
7. Define B+ trees.
8. Write the problems of executing two concurrent transactions.
9. State CAP Theorem.
10. List the different types of SQL injection attacks?

PART B — (5 × 13 = 65 marks)

11. (a) Discuss the main categories of data models. What are the basic differences between the relational model, the object model, and the XML model?

Or

- (b) Describe the three-schema architecture. Why do we need mappings between schema levels? How do different schema definition languages support this architecture?

12. (a) Consider the following schemas. The primary key for each relation is denoted by the underlined attribute.

LIVES (person-name, street, city)

WORKS (person-name, company-name, salary)

LOCATED-IN (company-name, city)

MANAGES (person-name, manager-name)

Write relational algebra expressions for the following queries:

- (i) Find the name of all employees (i.e., persons) who work for the City Bank company (which is a specific company in the database).
- (ii) Find the name and city of all employees who work for City Bank.
- (iii) Find the name, Street and city of all employees who work for City Bank and earn more than \$10,000.
- (iv) Find all employees who live in the same city as the company they work for.
- (v) Find all persons who do not work for City Bank.
- (vi) Find the second largest salary earned by the employee.

Or

- (b) Consider the relation for car dealership given below.

Dealership (Uniqueid, Customer-name, Purchase, Address, Newsfeed, Supplier, Price)

Normalize the table so that all resulting tables are in 3NF stating the definitions of various normal forms.

13. (a) Explain the concepts of serial, non-serial and conflict-serializable schedules with examples.

Or

- (b) What is the two-phase locking protocol? How does it guarantee serializability? Explain with an example.

14. (a) Construct a B+- tree for the following set of key values:

(2,3,5,7,11, 17, 19, 23, 29, 31)

- (i) Assume that the tree is initially empty and values are added in ascending order, Construct B+ - trees for the cases where the number of pointers that will fit in one node is as follows: (7)

(1) Four (2) Six (3) Eight

- (ii) Show the form of the tree after each of the following series of operations (6)

(1) Insert 9 (2) Insert 10 (3) Delete 23 (4) Delete 19

Or

- (b) (i) Brief on Static and Dynamic Hashing. (7)

- (ii) Suppose that we are using extendable hashing on a file that contains records with the following search-key values:

(2,3,5,7,11,17,19,23,29,31)

Show the extendable hash structure for this file if the hash function is $h(x) = x \text{ mod } 8$ and buckets can hold three records. (6)

15. (a) Describe in detail Query Processing and Optimization.

Or

- (b) Discuss on the following.

- (i) Access control based on privileges (7)

- (ii) Role Based access control (6)

PART C — (1 × 15 = 15 marks)

16. (a) Construct an E-R diagram for university database with the following information:

Professors have an SSN (Social Security Number), a name, an age, a rank, and a research speciality. Projects have a project number, a sponsor name, a starting date, an ending date and a budget. Graduate students have an SSN, a name, an age, and a degree program. Example: M.S or Ph.D. Each project is managed by one professor. Each project is worked by one or more professors, Professors can manage and/or work on multiple projects. Each project is worked by one or more graduate students (known as the projects research Assistants).