

25UCST401	DATABASE MANAGEMENT SYSTEMS	Category	L	T	P	R	Credit	Year of Introduction
		PCC	3	1	0	0	4	2025

Preamble

This course covers database architecture, data modeling, SQL, indexing, and transaction management. It equips students to design and manage reliable database systems for real-world applications such as banking, e-commerce, and enterprise systems.

Prerequisite: 25UCST301 Data Structures, DCM205 Discrete Mathematics

Course Outcomes: After the completion of the course, the student will be able to:

CO	Course Outcomes	Knowledge Level
CO1	Describe DBMS architecture, data models, and ER modeling.	K2 Understand
CO2	Apply SQL, ER-to-relational mapping, and normalization techniques for effective database design and querying.	K3 Apply
CO3	Describe storage structures, indexing techniques (B-tree, hashing), and the fundamentals of query processing.	K2 Understand
CO4	Analyze ACID properties, concurrency control, and recovery techniques to ensure database consistency and reliability.	K4 Analyze

Mapping of course outcomes with program outcomes and program specific outcomes

CO/PO	PO1	PO2	PO3	PO4	PO5	PO6	PO7	PO8	PO9	PO10	PO11	PSO1	PSO2
CO1	2	1	1	-	1	-	-	-	-	-	-	1	1
CO2	2	2	2	-	2	-	-	-	-	-	1	1	2
CO3	2	1	-	1	-	-	-	-	-	-	1	1	1
CO4	2	2	1	1	1	-	-	-	-	-	1	1	1
AVG	2	1.5	1.3	1.0	1.3	-	-	-	-	-	1.0	1.0	1.3

Mapping of course outcomes with Knowledge and Attitude Profile (WK)

	WK1	WK2	WK3	WK4	WK5	WK6	WK7	WK8	WK9
--	-----	-----	-----	-----	-----	-----	-----	-----	-----

CO1	*	*	*	*		*		*	
CO2	*	*	*	*		*		*	
CO3	*	*	*	*	*			*	
CO4	*	*	*	*	*	*		*	

Wks Mapped with the course

WK1	WK2	WK3	WK4	WK5	WK6	WK7	WK8	WK9
*	*	*	*	*	*		*	

Mark Distribution:

Total Marks	CIE Marks	ESE Marks	ESE Duration
100	40	60	2.5 hrs

Continuous Internal Evaluation Pattern:

Attendance	: 5 marks
Continuous Assessment Test	: 20 marks
Continuous Assessment Assignment	: 15 marks

SYLLABUS

Module 1 (12 hrs)

Introduction to databases: File-based vs DBMS, advantages of DBMS, users and DBA roles, DBMS architecture. Data models: ER model, object-based and record-based models. Entity-Relationship (ER) modeling: entities, relationships, attributes, aggregation, specialization/generalization, weak entities; design guidelines. Enhanced ER (EER) concepts: super/subclasses, constraints, ternary relationships. Relational model basics: relations, keys, constraints, relational algebra (selection, projection, join, set operations).

Module 2 (15 hrs)

SQL fundamentals: DDL, DML, DCL, TCL; CREATE, ALTER, DROP, INSERT, UPDATE, DELETE. Advanced SQL: JOINS, nested queries, correlated subqueries, views, materialized views, set operations. Constraints, triggers, stored procedures, functions. Relational DB design: Functional dependencies (FDs), closure of attributes, minimal cover. Normalization: 1NF, 2NF, 3NF, BCNF (definitions, tests, decomposition algorithms), multi-valued dependencies and 4NF, 5NF overview.

Module 3 (14 hrs)

Secondary storage devices & file organization: record blocking, record formats, heap files, sorted files. Indexing fundamentals: primary & secondary indices, dense vs sparse indices. B+ tree indexes: structure and operations (search/insert/delete). Hash-based indexing: static hashing, extendible/hash partitioning basics. Query processing & Query optimization.

Module 4 (15 hrs)

Transactions and properties: ACID, transaction states, serializability (conflict & view). Concurrency control: locks, lock compatibility, two-phase locking (2PL), deadlock detection and prevention. Recovery techniques: log-based recovery, deferred and immediate update, checkpoints. Advanced / recent topics (overview): distributed databases, NoSQL systems (document/key-value/columnar).

Text Books:

1. Elmasri, R., & Navathe, S. B. (2016). *Fundamentals of database systems* (7th ed.). Pearson Education.
2. Silberschatz, A., Korth, H. F., & Sudarshan, S. (2021). *Database system concepts* (7th ed.). McGraw-Hill Education.
3. Ramakrishnan, R., & Gehrke, J. (2003). *Database management systems* (3rd ed.). McGraw-Hill Education..

Reference Books:

1. Date, C. J., Kannan, A., & Swaminathan, S. (2006). *An introduction to database systems* (8th ed.). Pearson Education.
2. Deshpande, P. S. (2014). *Database management systems*. Dreamtech Press.
3. Connolly, T., & Begg, C. E. (2020). *Database systems: A practical approach to design, implementation, and management* (6th ed.). Pearson Education

END SEMESTER EXAMINATION – MODEL QUESTION PAPER

Course Code	25UCST401
Course Name	DATABASE MANAGEMENT SYSTEMS
Max. Marks: 60	Duration: 2.5 Hours

Part A (Answer all questions. Each question carries 3 marks)				
Qn No.	Question	Marks	CO	BL
1.	Differentiate DBMS and File-Based systems. List any four advantages of DBMS.	3	CO1	L2
2.	Explain the components of an ER diagram for a Library Management System by identifying at least four entities, their attributes, and the relationships among them.	3	CO1	L2
3.	Write SQL queries to: (i) Create a table STUDENT with appropriate constraints. (ii) Retrieve names of students who scored above 80 using nested query.	3	CO2	L3
4.	Explain 3NF and BCNF with examples. Given a relation R(A,B,C,D) with FDs: AB→C, C→D, D→A, check whether R is in BCNF.	3	CO2	L2

5.	Compare dense and sparse indices. When is a sparse index preferred over a dense index?	3	CO3	L2
6.	Explain the structure and characteristics of a B+ tree. Describe the roles of internal nodes and leaf nodes in organizing and retrieving data efficiently.	3	CO3	L2
7.	Explain the ACID properties of transactions with suitable examples.	3	CO4	L2
8.	Analyze how the Two-Phase Locking (2PL) protocol enforces conflict serializability in a multi-user database system. Illustrate your analysis using a transaction schedule and explain the consequences if the growing and shrinking phases are violated.	3	CO4	L4

Part B (Answer any one full question from each Module. Each carry 9 marks)				
Module I				
Qn No.	Question	Marks	CO	BL
9.a	Explain the three-schema architecture of DBMS. How does it achieve data independence?	5	CO1	L2
9.b	Explain the concepts of specialization and generalization in Enhanced Entity-Relationship (EER) modeling with a suitable example.	4	CO1	L2
OR				
10.a	Describe relational algebra operations: selection, projection, join, and set operations with examples.	5	CO1	L2
10.b	Explain the entities, relationships, and cardinalities involved in a University database consisting of Student, Faculty, Course, and Department entities.	4	CO1	L2
Module II				
Qn No.	Question	Marks	CO	BL
11.a	Given the tables Student(StudentID, Name, DeptID) and Department(DeptID, DeptName), write SQL queries using appropriate JOIN operations to: a) Display the names of students along with their department names. b) Display all departments, including those that have no students assigned. Identify the JOIN type used in each query and justify your choice.	5	CO2	L3

11.b	Given relation R(A,B,C,D,E) with FDs: $A \rightarrow BC$, $CD \rightarrow E$, $B \rightarrow D$, $E \rightarrow A$. Find the canonical cover and decompose R into 3NF.	4	CO2	L3
OR				
12	An organization maintains an EMPLOYEE table and an AUDIT_LOG table. Write an SQL trigger that records the employee ID, old salary, new salary, and update time whenever an employee's salary is modified. Explain how the trigger is executed during the update operation.	9	CO2	L3
Module III				
Qn No.	Question	Marks	CO	BL
13	Explain query processing steps in detail. Describe how the query optimizer selects an efficient execution plan using heuristics.	9	CO3	L2
OR				
14.a	Explain static hashing and extendible hashing for database indexing with examples.	5	CO3	L2
14.b	Compare heap file organization and sorted file organization. In what scenarios is each preferred?	4	CO3	L2
Module IV				
Qn No.	Question	Marks	CO	BL
15.a	Analyze the deferred update and immediate update log-based recovery techniques used in database systems. Compare their recovery procedures after a system failure and explain the roles of REDO and UNDO operations with suitable examples.	4	CO4	L4
15.b	Analyze the document, key-value, and columnar NoSQL data models. Compare their scalability and consistency characteristics with those of relational databases, and determine the most suitable model for a large-scale e-commerce application. Justify your answer.	5	CO4	L4
OR				
16	Analyze the deadlock detection and prevention techniques used in database concurrency control. Given a wait-for graph of concurrent transactions, determine whether a deadlock exists, justify your conclusion, and recommend an appropriate resolution strategy.	9	CO4	L4

