Course Code & Name	26CAPC7002 Advanced Database Management Systems		
Teaching Hours /week (L: T:P: R)	3:1:0:0	Credits	4
CIE Marks	40	ESE Marks	60
Course Type	Theory (PCC)	Exam Hours	3
Prerequisites (if any)	26CAPC2003		

Course Objectives:

- To design efficient and reliable databases using advanced normalization, transaction management, and data storage techniques.
- To understand and evaluate modern database systems such as distributed, object oriented, XML, and NoSQL databases.

Course Outcomes (COs)

At the end of the course students should be able to:

CO#	Course Outcomes	Bloom's Knowledge Level (KL)
CO 1	Apply and evaluate advanced normalization concepts, including higher normal forms and functional dependencies, for efficient and consistent database design.	K3
CO 2	Analyse and implement transaction management and concurrency control mechanisms, including locking, timestamping, and recovery methods, to maintain database consistency and reliability.	K2
CO 3	Apply appropriate data storage, indexing, and hashing techniques to optimize data retrieval and query performance in database systems.	K3
CO 4	Analyze and evaluate distributed, object-oriented, XML, and NoSQL databases to identify their features, architectures and use cases for modern data management.	K3

Note: *K1- Remember, K2- Understand, K3- Apply, K4- Analyse, K5- Evaluate, K6- Create*

CO-PO Mapping Table (Mapping of Course Outcomes to Program Outcomes)

CO #	PO 1	PO 2	PO 3	PO 4	PO 5	PO 6	PO 7	PO 8
CO 1	3	3	3	2	1	-	1	2
CO 2	3	3	3	3	2	1	2	2
CO 3	3	2	3	3	1	1	1	2
CO 4	3	2	3	3	2	1	2	3

Note: *1: Slight (Low), 2: Moderate (Medium), 3: Substantial (High), - : No Correlation*

Course Assessment Method

(CIE: 40 marks, ESE: 60 marks)

Continuous Internal Evaluation Marks (CIE)

Attendance	Assignment/ Micro-project	Internal Examination-1 (Written)	Internal Examination- 2 (Written)	Total
8	12	10	10	40 Marks

End Semester Examination Marks (ESE)

In Part A, all questions need to be answered and in Part B, each student can choose any one full question out of two questions.

PART A	PART B	TOTAL
• 2 Questions from each module. • Total of 8 Questions, each carrying 3 marks **(8x3 =24 marks)**	• Each question carries 9 marks. • Two questions will be given from each module, out of which 1 question should be answered. • Each question can have a maximum of 3 subdivisions. **(4x9 = 36 marks)**	**60 Marks**

SYLLABUS

Module #	Syllabus Description	Hours
1	**Advanced Database Design**: Review of Normal Forms Based on Primary Keys- (1NF, 2NF, 3NF), Inference Rules for Functional Dependencies (proof not needed) - Minimal set of Functional Dependencies - Improving the Design - Surrogate Key Considerations - Higher Level Normal Forms: BCNF, Multi-valued Dependencies, 4NF, Join dependencies, 5NF, Domain Key Normal Form.	**10**
2	**Transaction Management and concurrency :** Transaction: Transaction Concepts, Transaction Properties, Transaction Management with SQL, The Transaction Log – Concurrency Control: Lost Updates, Uncommitted Data, Inconsistent Retrievals, The Scheduler – Concurrency Control with Locking Methods: Lock Granularity, Lock Types, Two Phase Locking to Ensure Serializability, Deadlocks – Concurrency Control with Timestamping Methods: Wait/Die and Wait/Wound Schemes – Concurrency Control with Optimistic Methods - Multiversion Concurrency Control (MVCC)- Database Recovery Management: Transaction Recovery.	**12**
3	**Data Storage and Querying:-** RAID – File Organization – Organization of Records in Files – Indexing and Hashing: Basic concept, Ordered Indices, B+ tree Index Files: Structure of a B+-Tree (structure only, algorithms not needed) - B tree index files – Static Hashing – Dynamic Hashing – Query Processing: Overview	**12**
4	**System Architecture, Object Oriented Databases, XML and NoSQL:-** Distributed Databases: Homogeneous and Heterogeneous Databases, Distributed Data Storage, Distributed Transactions, Google Spanner- Object Based Databases: Overview, Complex Data types, Structured types and inheritance in SQL, Table Inheritance, Array and Multiset types in SQL, Object identity	**12**

	and reference types in SQL - XML: DTD and XML Schema, XML presentation, XML Applications - Next Generation Databases: Distributed Relational Databases - Nonrelational Distributed Databases - MongoDB Sharding and Replication - Hbase - Cassandra - CAP Theorem.	

Text Books				
Sl. No	**Title of the Book**	**Name of the Author/s**	**Name of the Publisher**	**Edition & Year**
1	Database Principles fundamentals of design, implementation and management	Carlos Coronal, Steven Morris, Keeley Crocket, Craig Blewett	Cengage	3rd Edition 2020
2	Next Generation Databases NoSQL, NewSQL, and Big Data	Guy Harrison	Apress	2015

Reference Books				
Sl. No	**Title of the Book**	**Name of the Author/s**	**Name of the Publisher**	**Edition & Year**
1	Database Concepts	Silberschatz, Korth and Sudarshan	Tata Mc Graw Hill	6th Edition 2010
2	Database Systems: The Complete Book	Hector Gracia, Jeff Ullman, Jennifer Widom	Pearson	1st Edition 2011
3	MongoDB and Python	Niall O' Higgins	O'reilly	2011

Web References/ Video Links (NPTEL, SWAYAM…)	
Sl.No	**Link ID**
1	https://www.coursera.org/specializations/google-cloud-database-engineer
2	https://onlinecourses.nptel.ac.in/noc21_cs04

Course Level Assessment Questions:

Course Outcome 1 (CO1):

1. Define *functional dependency* and list the different normal forms used in database normalization.
2. Explain how *partial* and *transitive* dependencies cause anomalies in a relational database design.
3. Given a relation schema R(A, B, C, D, E) with FDs: A→B, B→C, AC→D, D→E, identify the highest normal form and justify.
4. Differentiate between *BCNF* and *3NF* with suitable examples.
5. Evaluate the role of *surrogate keys* in improving database design compared to natural keys.

Course Outcome 2 (CO2):

1. Define a *database transaction* and list the *ACID properties* that ensure reliable transaction processing.
2. Explain the problems caused by *lost updates*, *uncommitted data*, and *inconsistent retrievals* in concurrent transaction environments.
3. Illustrate the working of the *Two-Phase Locking (2PL)* protocol with a suitable example showing how serializability is achieved.
4. Compare *timestamp-based* concurrency control with *locking-based* methods, highlighting advantages and disadvantages of each.
5. Evaluate how *deadlocks* occur in a transaction system and discuss strategies such as *Wait/Die* and *Wait/Wound* schemes for deadlock handling.

Course Outcome 3 (CO3):

1. Define RAID and list its different levels used in database storage systems.
2. Explain how file organization and record organization affect data access performance in a database.

3. Illustrate the structure of a B+ tree index and explain how it supports efficient data retrieval.
4. Differentiate between static hashing and dynamic hashing with suitable examples.
5. Evaluate the advantages of using B+ tree indices over B tree indices in query processing.

Course Outcome 4 (CO4):

1. Define a distributed database and differentiate between homogeneous and heterogeneous distributed systems.
2. Explain object-oriented features in databases, such as inheritance, complex data types, and object identity, with examples.
3. Apply XML schema and DTD concepts to represent data for a Library Information System.
4. Analyze how distributed transactions ensure consistency across nodes in a distributed database environment.
5. In the context of NoSQL databases and explain how systems like MongoDB and Cassandra address trade-offs between consistency, availability, and partition tolerance.

Model Question Paper

Reg No.:_____ Name:_____

FEDERAL INSTITUTE OF SCIENCE AND TECHNOLOGY

MCA S1 (R,S) Examination ---------- (2026 Scheme)

Course Code: 26CAPC7002

Course Name: Advanced Database Management Systems

Max. Marks: 60 Duration: 3 Hours

PART A

Answer all questions, each carries 3 marks.

		CO	Marks
1	Explain the need of applying the normalization process.	CO1	(3)
2	Differentiate single and multi valued dependency with suitable examples.	CO1	(3)
3	Discuss the ACID properties of a transaction.	CO2	(3)
4	What is a transaction log? Why is it used for?	CO2	(3)

5	Explain two file organization techniques with suitable examples.	CO3	(3)
6	Explain the steps in query processing.	CO3	(3)
7	How is shading different from replication in the context of MongoDB?	CO4	(3)
8	Explain the use of gossip protocol in Cassandra.	CO4	(3)

PART B

(Answer any one full question from each module, each question carries 9 marks)

Module I

9	a. Explain the axioms of functional dependency.	CO1	
	b. Define normalization? Apply the normalization technique to the following attributes to arrive at a database design – Emp_ID, Proj_ID, Proj_Name, Emp_Name, Emp_Designation, Charge_Per_Hour, Total_Hours.		(3) (6)

OR

10	a. Explain the steps in the Minimal Cover algorithm.	CO1	(4)
	b. Given a relation M(P,Q,R,S,T,U) with FDs, F={P → R, PQ → R, R → SU, RS → U, TR→ PQ}, compute the minimal cover of F.		(5)

Module II

11	a. Illustrate the following problems with suitable examples: i) Lost-update problem (ii) Inconsistent Retrieval Problem (iii) Uncommitted Data Problem.	CO2	(6) (3)
	b. Explain the transaction recovery process. Differentiate the deferred-write and write-through transaction recovery procedures.		

OR

12	a. Explain why concurrency control mechanisms are needed in transaction management.	CO2	(5)
	b. Briefly discuss the two-phase locking protocol used in concurrency control. How does it ensure concurrency?		(4)

Module III

13	a. Define RAID and list its different levels used in database storage systems.	CO3	(6) (3)
	b. How is dense index different from sparse index?		

OR

14	a.	Illustrate the structure of a B+ tree index and explain how it supports efficient data retrieval.	CO3	(5)
	b.	Differentiate between static hashing and dynamic hashing with suitable examples.		(4)

Module IV

15	a.	What are the desirable properties of distributed databases with replication?	CO4	(4)
	b.	Define Fragmentation. Consider a College database that keeps records of all registered students in a Student table with the schema. Student (RegNo, Name, Course, Address, Semester, Fees, Marks). Create a fragment for the accounts section of the college to store the fees details of the registered students.		(5)

OR

16	a.	Explain with the help of suitable examples, how multi-set and array types are created in object-based databases.	CO4	(3)
	b.	Illustrate, with the help of a suitable example, the structure of a DTD. Write any two sample XML documents that satisfy your example DTD.		(6)

*****

FISAT
Focus on Excellence