22DS1T3: ADVANCED DATABASE MANAGEMENT SYSTEMS

Course Name	Advanced Database Management Systems		L	T	P	C	CIA	SEE	TM
Course Code	22DS1T3		4	0	0	4	30	70	100
Year of Introduction: 2021	Year of Offering: 2022	Year of Revision: 2022	Percentage of Revision: 7%						
L-Lecture, T-Tutorial, **P**-Practical, **C**-Credits, **CIA**-Internal Marks, **SEE**-External Marks, **TM**-Total Marks									

Course Description and Purpose:
Advanced Database Management Systems is a course that illustrates basic concepts of *Structured Query Language*, *Relational Algebra and Relational Calculus*, *Functional Dependencies and Normalization for Relational Databases*, *Transaction Processing Concepts*, *Concurrency Control Techniques*, *Data Models*, *Distribution Models & Consistency of NoSQL*, *Querying and Creating*, *Updating & Deleting Documents in Mongo DB*, *Data Lakes*.

Course Objectives:
This course will help enable the students t o understand, learn and develop a various *Data Models and Basic Querying*, *Transaction Processing*, *Concurrency Control*, *Distributed Databases*, *Data Lakes* also apply *Creating, Querying, Updating & Deleting Documents in Mongo DB.*

Specific objectives include:
- ✓ To understand basic concepts of *Structured Query Language*, *Relational Algebra and Relational Calculus*.
- ✓ To learn the basics of *Functional Dependencies and Normalization* for Relational Databases and *Transaction Processing* Concepts.
- ✓ To learn *Concurrency Control Techniques* and *Distributed Database Concepts*.
- ✓ To understand the *Data Models*, *Distribution Models & Consistency of NoSQL*.
- ✓ To know *Querying*, *Creating*, *Updating & Deleting Documents in Mongo DB*, *Data Lakes*.

Course Learning Outcomes:
Upon successful completion of the course, the student will be able to:
CO1: Understand basic concepts *of Structured Query Language & Relational Algebra and Relational Calculus*.
CO2: Learn the basics of *Functional Dependencies and Normalization for Relational Databases* and *Transaction Processing Concepts*.
CO3: Learn *Concurrency Control Techniques* and *Distributed Database Concepts*.
CO4: To understand the *Data Models*, *Distribution Models & Consistency of NoSQL*.
CO5: To know *Querying*, *Creating*, *Updating & Deleting documents in Mongo DB*, *Data Lakes*.

UNIT I (12 Hours)

Basic & More SQL: SQL Data Definition and Data Types - Specifying Constraints in SQL - Basic Retrieval Queries in SQL - INSERT, DELETE, and UPDATE Statements in SQL - Additional Features of SQL - More Complex SQL Retrieval Queries - Specifying Constraints as Assertions and Actions as Triggers - Views (Virtual Tables) in SQL - Schema Change Statements in SQL.
Data Modeling Using the ER Model: Entity - Entity Types - Entity Sets - Attributes and Keys - Relationship Types - Relationship Sets - Roles and Structural Constraints - Weak Entity Types - Relationship Types of Degree Higher than Two - Refining the ER Design for the COMPANY Database.
The Relational Algebra and Relational Calculus: Unary Relational Operations: SELECT and PROJECT-Relational Algebra
Operations from Set Theory - Binary Relational Operations: JOIN and DIVISION - Additional Relational Operations - Examples of Queries in Relational Algebra - The Tuple Relational Calculus - The Domain Relational Calculus.

UNIT II (12 Hours)

Basics of Functional Dependencies and Normalization for Relational Databases: Functional Dependencies - Normal Forms Based on Primary Keys - General Definitions of Second and Third Normal Forms - Boyce Codd Normal Form - Multivalued Dependency and Fourth Normal Form - Join Dependencies and Fifth Normal Form.

Introduction to Transaction Processing Concepts and Theory: Introduction to Transaction Processing - Transaction and System Concepts - Desirable Properties of Transactions - Characterizing Schedules Based on Recoverability - Characterizing Schedules Based on Serializability -Transaction Support in SQL.

UNIT III (12 Hours)

Concurrency Control Techniques: Two-Phase Locking Techniques for Concurrency Control - Concurrency Control Based on Timestamp Ordering - Multiversion Concurrency Control Techniques - Validation (Optimistic) Techniques and Snapshot Isolation Concurrency Control - Granularity of Data Items and Multiple Granularity Locking - Using Locks for Concurrency Control in Indexes - Other Concurrency Control Issues.

Distributed Database Concepts: Data Fragmentation, Replication, and Allocation Techniques for Distributed Database Design - Overview of Concurrency Control and Recovery in Distributed Databases - Overview of Transaction Management in Distributed Databases - Query Processing and Optimization in Distributed Databases - Types of Distributed Database Systems - Distributed Database Architectures - Distributed Catalog Management.

UNIT IV (12 Hours)

Why NoSQL: The Value of Relational Database - Emergence of NoSQL .

Aggregate Data Models: Aggregates - Keyvalue and Document Data Models - Column Family Stores.

More Details on Data Models: Relationships - Graphs DB - Schemaless DB - Materialized Views.

Distribution Models: Single Server - Shrading - Master Slave Replication.

Consistency: Update - Read - Relax Consistency.

UNIT V (12 Hours)

Getting Started: Documents - Collections - Databases - Data Types.

Creating, Updating & Deleting Documents: Inserting & Saving Documents - Removing Documents - Updating Documents.

Querying: Introduction to Find - Query Criteria - Type Specific Queries - Where Queries - Cursors.

Data Lakes: Introduction - What is Data Lake? - The value of the Data Lake to ING - The 5 Level Model of Governance Maturity.

Reference Text Books:

1. Ramez Elmasri & Shamkant B. Navathe, Fundamentals of Database Systems, Pearson, Seventh Edition, 2016
2. Pramod J.Sadalage & Martin Fowler, No SQL Distilled, Addison-Wesley, Second Edition, 2013.
3. Kristina Chodorow, Mongo DB, O'Reilly, Second Edition, 2013
4. Mandy Chessell, Ferd Scheepers, Maryna Strelchuk, Ron van der Starre, Seth Dobrin, Daniel Hernandez, From Data Lake to Data Driven Organization, IBM-Red Guide,2018, https://www.redbooks.ibm.com/redpapers/pdfs/redp5486.pdf
5. Shashank Tiwari, Professional NoSQL, Wiley, 2011, Second Edition, 2011
6. Abraham Silberschatz, Henry F Korth , S Sudarshan, Database System Concepts, McGraw-Hill International Edition, Sixth Edition,2011

PARVATHANENI BRAHMAYYA SIDDHARTHA COLLEGE OF ARTS & SCIENCE

(An Autonomous College in the jurisdiction of Krishna University)
M.Sc.(Computer Science), First Semester
Course Name: Advanced Database Management Systems
Course Code: 22DS1T3
(w.e.f admitted batch 2022-23)

Time: 3 Hours **Max Marks: 70**

SECTION-A
Answer ALL questions. All Questions Carry Equal Marks. (5×4 = 20 Marks)

1.(a) Explain Trigger in SQL with example. (CO1,L2)
(or)
(b) Illustrate *DDL Commands* in SQL. (CO1,L2)
2. (a) Analyze *Second Normal Form*. (CO2,L4)
(or)
(b) List the *Properties of Transactions*. (CO2,L4)
3. (a) Explain *Multiple Granularity Locking*. (CO3,L2)
(or)
(b) Explain *Query Processing and Optimization* in Distributed Databases. (CO3,L2)
4. (a) Explain important features of NoSQL databases. (CO4,L5)
(or)
(b) Explain *Aggregate Data Models*. (CO4,L5)
5. (a) Apply Triggers in MongoDb. (CO5,L3)
(or)
(b) Construct where query in MongoDb. (CO5,L3)

SECTION-B
Answer ALL questions. All Questions Carry Equal Marks. (5×10 = 50 Marks)

6. (a) Explain *various constraints* of Relational Model. (CO1,L2)
(or)
(c) Illustrate *Select* & *Project* operations of *Relational Algebra*. (CO1,L2)

7 .(a) Explain *Fifth* Normal Forms in detail. (CO2,L5)
(or)
(b) Prove whether the transactions T1 & T2 ensure *serializability*. (CO2,L5)

T1	T2
read_item(x); X:=X – N;	
	read_item(x); X:=X + M;
write_item(X); read_item(Y);	
	Write_item(x);
Y:=Y+N; Write_item(Y);	

8. (a) Build *Concurrency Control* based on *Timestamp Ordering*. (CO3,L3)

(or)

(d) Identify the usage of *Data Fragmentation*, *Replication*, and *Allocation Techniques* for *Distributed Database Design*. (CO3,L3)

9. (a) Analyze *Graphs DB* and *Schemaless DB* in detail. (CO4,L4)

(or)

(b) Compare *Shrading* and *Master Slave Replication* in detail. (CO4,L4)

10. (a) Demonstrate *CRED Operation* in *MongoDB*? (CO5,L2)

(or)

(b) Explain the concepts of *Data Lake* in detail? (CO5,L2)