

III Year II Semester

L T P C

Code: 20DS6659

4 0 0 4

NEXT GENERATION DATABASES (Honors)

Course Objectives:

1. To introduce about the new generation database management systems
2. To give a good formal foundation on the document databases
3. To introduce the concepts of use of databases in Graphs and Columns
4. To introduce the concepts of memory databases and object databases
5. To provide an overview of the databases techniques that are useful in future.

Course Outcomes:

After completion of the course, the students should be able to

1. Understand the first, second and third revolution of databases from file base databases to SQL to NoSQL to NewSQL.
2. Understand the document database overview, NoSQL databases, overview of MongoDB and CouchDB.
3. Understand the graph and column databases overview and overview of Neo4j, Gremlin, Sybase IQ, C-Store, and Vertica.
4. Understand features of in-memory databases and object databases overview and overview of Times Ten, Redis, VoltDB, Oracle 12c, SPARK architectures.
5. Understands about future database technologies like storage, block chain and quantum computing.

UNIT-I:

Database Revolution: Three Database Revolutions, Early Database Systems-The First Database Revolution, The Second Database Revolution- Relational theory, Transaction Models, The First Relational Databases, Client- server Computing, Object-oriented Programming and the OODBMS, The Relational Plateau, The Third Database Revolution, Google and Hadoop, The Rest of the Web, Cloud Computing, Document Database, The “NewSQL”, The Nonrelational Explosion. Google, Big Data, and Hadoop, The Big Data Revolution- Cloud, Mobile, Social, and Big Data, Google: Pioneer of Big Data, Google Hardware, The Google Software Stack, More about MapReduce, Hadoop: Open-Source Google Stack -Hadoop’s Origins, The Power of Hadoop, Hadoop’s Architecture, HBase ,Hive, Pig , The Hadoop Ecosystem.

UNIT-II:

Document Databases: What is a document database, NoSQL databases, Why choose NoSQL?, Performance overview of different databases, Why a document store, How does it work, Data storage ,Data querying and the map/reduce paradigm ,Inserting and Modifying, ACID , The different solutions -Open source solution, Proprietary solution. Examples - CouchDB, Why CouchDB , The storage, concurrency, Managing the database, Querying the database, Specificity of Couch DB. Examples - MongoDB, Why MongoDB?, The storage, concurrency, Managing the database, Querying the database, Specificity of Mongo DB.

UNIT-III:

Graph Databases & Column Databases: What is a Graph? RDBMS Patterns for Graphs, RDF and SPARQL, Property Graphs and Neo4j, Gremlin, Graph Database Internals, Graph Compute Engines. What is Column Databases, Why it used? The Columnar Alternative - Columnar Compression, Columnar Write Penalty, Sybase IQ, C-Store, and Vertica, Column Database Architectures -Projections, Columnar Technology in Other Databases

UNIT-IV:

In-Memory Databases & Object Databases: What is In-Memory Databases? The End of Disk? -Solid State Disk, The Economics of Disk, SSD- Enabled Databases, In-Memory Databases- Examples Times Ten, Redis, SAP HANA, VoltDB, Oracle 12c “in-Memory Database”, Berkeley Analytics Data Stack and Spark, Spark Architecture. Overview of object databases, Object Oriented Database, Object Relational Database, mapping of object relational mapping and standards of ODBMS, Examples: Object DB, Object Store, Objectivity DB, db4o and Gemstone features and advantages.

UNIT-V:

Databases of the future: The revolution revisited, counterrevolutionaries-have we come full circle? can we have it all? - consistency models, schema, database languages, storage, a vision for a converged database, other convergent databases, Disruptive database technologies-storage technologies, Blockchain-What it is?, Understanding Technologies, When it is used?, Quantum computing-Quantum Transaction, Quantum Search, Quantum Query Language.

Text Books:

1. Next Generation Databases — NoSQL, NewSQL and Bigdata, Guy Harrison, A press.
2. Couch DB, Document oriented Databases, Alain Issa, François Schiltz, ULB
3. Mongo DB Architecture Guide, MongoDB university, white paper
4. Graph Databases-neo4j, Ian Robinson, Jim Webber & Emil Eifrem, 2nd edition, O’Reilly
5. The Design and Implementation of Modern Column-Oriented Database Systems, Daniel Abadi, Foundations and Trends in Databases Vol. 5, No. 3 (2012) 197–280, 2013