# *Paper – V(B):* NoSQL Data Bases

[4 HPW :: 4 Credits :: 100 Marks (External:80, Internal:20)]

**Objective:** The main objective of this course is to cover core concepts of NoSQL databases, along with an example database for each ofthe key-value, document, column family, and graph databases

**Outcomes:**
At the end of the course the student will be able to

- Understand the need for NoSQL databases and their characteristics
- Understand the concepts of NoSQL databases
- Implement the concepts of NoSQL databases using four example databases: Redis for key-value databases, MongoDB for document databases, Cassandra for column-family databases, and Neo4J for graphdatabases.

## Unit-I

**Why NoSQL:** The Value of Relational Databases, Impedance Mismatch, Application and Integration Databases, Attack of the Clusters, The Emergence of NoSQL
**Aggregate Data Models:** Aggregates, Column-Family Stores, Summarizing Aggregate-Oriented Databases
**More Details on Data Models:** Relationships, Graph Databases, Schemaless Databases, Materialized Views, Modeling for Data Access

## Unit-II

**Distribution Models:** Single Server, Sharding, Master-Slave Replication, Peer-to-Peer Replication, Combining Sharding and Replication
**Consistency:** Update Consistency, Read Consistency, Relaxing Consistency, Relaxing Durability, Quorums
**Version Stamps:** Business and System Transactions, Version Stamps on Multiple Nodes
**Map-Reduce:** Basic Map-Reduce, Partitioning and Combining, Composing Map-Reduce Calculations

## Unit-III

**Key-Value Databases:** What Is a Key-Value Store, Key-Value Store Features, Suitable Use Cases, When Not to Use
**Document Databases:** What Is a Document Database, Features, Suitable Use Cases, When Not to Use

## Unit-IV

**Column-Family Stores:** What Is a Column-Family Data Store, Features, Suitable Use Cases, When Not to Use
**Graph Databases:** What Is a Graph Database, Features, Suitable Use Cases, When Not to Use

**Reference:**

1. Pramod J. Sadalage, Martin Fowler. NoSQL Distilled, Addison Wesley 2013

**Suggested Reading**

2. Luc Perkins, Eric Redmond, Jim R. Wilson. Seven Databases in Seven Weeks. The Pragmatic Bookshelf, 2018
3. Guy Harrison. Next GenerationDatabases: NoSQL, NewSQL, and Big Data. Apress, 2015

# *Practical – 5(B) :* NoSQL Data Bases (Lab)

[3 HPW :: 1 Credit :: 25 Marks]

**Objective:** The main objective of this lab is to become familiar with the four NoSQL databases: Redis for key-value databases, MongoDB for document databases, Cassandra for column-family databases, and Neo4J for graphdatabases

**NoSQL Databases:**
Redis (http://redis.io)
MongoDB (http://www.mongodb.org)
Cassandra (http://cassandra.apache.org)
Neo4j (http://neo4j.com)

**Exercises:**
1. Installation of NoSQL Databases: Redis,MongoDB, Cassandra, Neo4j on Windows & Linux
2. Practice CRUD (*Create*, *Read*, *Update*, and *Delete)* operations on the four databases: Redis,MongoDB, Cassandra, Neo4j
3. Usage of Where Clause equivalent in MongoDB
4. Usage of operations in MongoDB – AND in MongoDB, OR in MongoDB, Limit Records and Sort Records. Usage of operations in MongoDB – Indexing, Advanced Indexing, Aggregation and Map Reduce.
5. Practice with ' macdonalds ' collection data for document oriented database. Import restaurants collection and apply some queries to get specified output.
6. Write a program to count the number of occurrences of a word using MapReduce