18CS632 Distributed Systems 3-0-1-4

Course Objective:

This course is an introduction to the design of distributed systems and algorithms that support distributed computing. It aims to provide a practical exposure into the design and functioning of existing distributed systems

Prerequisites:

Operating Systems, Networks

Course Syllabus

Unit 1:

Introduction and types of distributed systems – Architecture of DS - Overview of processes - Taxonomy of Distributed Systems - scalable performance - load balancing and availability. Models of computation - shared memory and message passing system - synchronous and asynchronous systems. Communication in Distributed Systems - Remote Procedure Calls and Message Oriented Communications and implementation - High-level communication and publish-subscribe in Map reduce.

Unit 2:

Logical time and event ordering - Global state and snapshot algorithms - distributed snapshots in VMs - clock synchronization - Distributed mutual exclusion - Group based Mutual Exclusion - leader election - deadlock detection - termination detection - Distributed Databases - implementations over a simple distributed system and case studies of distributed databases and systems - Distributed file systems: scalable performance, load balancing, and availability. Examples from Dropbox, Google FS (GFS)/ Hadoop Distributed FS (HDFS), Bigtable/HBase MapReduce, RDD

Unit 3:

Consistency control: Data Centric Consistency - Client Centric Consistency - Replica Management - Consistency Protocols. Fault tolerance and recovery: basic concepts - fault models - agreement problems and its applications - commit protocols - voting protocols - check pointing and recovery. Case Studies from Apache Spark, Google Spanner, Amazon Aurora, Block Chain Systems etc.

References:

1. M. Van Steen, A.S. Tanenbaum, Distributed Systems, Third Edition, CreateSpace Independent Publishing Platform, 2017.
2. Ajay D. Kshemkalyani and Mukesh Singhal, Distributed Computing: Principles, Algorithms and Systems, Cambridge University Press, 2011.
3. Garg VK. Elements of distributed computing. John Wiley & Sons, 2002.
4. George Coulouris, Jean Dollimore, Tim Kindberg and Gordon Blair, Distributed Systems: Concepts and Design, Fifth Edition, Pearson Education, 2017.

5. Fokkink W. Distributed algorithms: an intuitive approach, Second Edition, MIT Press, 2018.

Course Outcomes:

At the end of the course the students will be able to

	Course Outcome	**BTL**
CO 1	Understand the design principles in distributed systems and the architectures for distributed systems.	3
CO 2	Apply various distributed algorithms related to clock synchronization, concurrency control, deadlock detection, load balancing, voting etc.	4
CO 3	Analyze fault tolerance and recovery in distributed systems and algorithms for the same.	4
CO 4	Analyze the design and functioning of existing distributed systems and file systems.	4
CO 5	Implement different distributed algorithms over current distributed platforms	5

CO-PO-PSO mapping:

CO	**PO1**	**PO2**	**PO3**	**PO4**	**PO5**	**PSO1**	**PSO2**	**PSO3**
1	2		2	2		2	3	3
2	2		2	2		2	3	3
3	2		2	2		2	3	3
4	3	3	3	3		2	3	3
5	3		3	2	3	2	3	3

Evaluation Pattern: 4B

Component	Weight (%)
Assignments	15
Lab Exercises	15
Case Study	10
Periodical 1	15
Periodical 2	15
End semester	30

Programme Objectives (PSO)

Hone the skill of computer science professionals in areas of research and innovation.

Develop experts with high professional competence in recent and futuristic technologies.

Create man power with technical competency in computer science to design and develop solutions for the societal problems.

Program Outcomes (PO)

Ability to independently carry out research investigation and development work to solve practical problems

Ability to write and present a substantial technical report/document

Students should be able to demonstrate a degree of mastery over the area

Ability to design and develop computing solutions using emerging computing paradigms to interdisciplinary problems following standard practices, tools and technologies

Ability to demonstrate commitment to professional ethics