CS632A: Distributed Systems (3-0-0-9)

Course objectives:

The course is designed with objectives to provide the students, a foundations of Distributed Systems with focus on both middleware and theoretical issues that arise in building large scale distributed systems. The students are expected to have some knowledge of networks and operating systems for crediting the course. However some of the background materials will be reviewed in course of introductory lectures.

Course contents:

Introduction: Characteristics and goals of distributed systems, Difficulties in implementation of distributed systems, Architectural organizations of parallel and distributed systems, grid systems, cloud systems, evolution of pervasive, mobile, grid, cloud and distributed systems. [2]

Middleware tools: Socket programming, using RPC, RMI, MPI with examples. [2]

Internet: Basic concepts of Internet, Review of Internet protocols, TCP, UDP, Domain name service, Content delivery networks (youtube, yahoo-akamai). [3]

Concept of Clock in Distributed Systems: Notion of UTC, UT1 and GMT times, Protocols for synchronizing system clocks to wall clock times (Christian's clock, Berkeley clock, NTP), External and Internal synchronization mechanism, Temporal ordering of events in distributed systems, Causal dependencies, Logical clocks (Lamport's clock, vector clock), Birman-Schiper-Stephenson protocol, Schiper-Eggli-Sandoz protocol for unicast message ordering. Multicast message ordering [9]

Global State and Termination Detection: Notion of global states of distributed system, cuts and consistent cuts, recording of global state (Chandy-Lamport Algorithm), stable properties of a distributed systems, safe and nonsafe properties, termination and deadlock. [9]

Leader Election: Impossibility results, and a number of leader election algorithms.

Mutual Exclusion: Necessity of mutual exclusion, coordinator based solutions, distributed mutual exclusion algorithms. [6]

Agreement and Consensus: Consensus, Byzantine and Interactive consistency (problem definitions, system models). Lamport-Shostak-Pease algorithm, Phase King algorithm, Two-phase commit and Three-phase commit protocols, Consensus in asynchronous systems, Paxos algorithm. [6]

Memory Consistency Models: Data and client centric models, replica management [3]

Instructor: R. K. Ghosh

Class Room: KD 103
Time: 10:30-11:45AM

Course Organization: All relevant course material including discussion and meeting time will be notified via emails apart from class lectures.

Assignments, Quizzes and Exams:

1. There will be 5mts review questions on material taught in previous class during each lecture. 15 best scores of each student will considered as credit of the final score out of 15 marks.
2. One major quiz for 10 marks.
3. One midterm examination for 20 marks.
4. One end term examination for 20 marks.
5. One major project for 35 marks.

Review questions are meant for understanding problems related assimilation of lectures and also serves the purpose of getting student prepared for the exams, quiz and projects without exerting last minute pressure. So, students can get a decent grade with reasonable efforts.

Attendance: Not mandatory, but advisable as review question quizzes will be held during each lecture.

Books and course material:

1. Niranjan Shivaratri and Mukesh Singhal, *Concepts of Advanced Operating Systems*. McGraw Hill.
2. Nancy Lynch, *Distributed Algorithms*. Morgan Kaufman.
3. Nicolo Santoro, Design and Analysis of Distributed Algorithms. John Wiley.
4. Sukumar Ghosh, Distributed Systems. CRC Press.
5. Course Notes will be sent through Email.

R. K. Ghosh, Dept of CSE