



CLOUD COMPUTING AND DISTRIBUTED SYSTEMS

PROF. RAJIV MISRA

Department of Computer Science and Engineering
IIT Patna

PRE-REQUISITES : Minimum: Data Structures and Algorithms; Ideal: Computer Architecture, Basic OS and Networking concepts

INDUSTRIES SUPPORT : Companies like Amazon, Microsoft, Google, IBM, Facebook and start-ups working on this field.

COURSE OUTLINE :

Cloud computing is the on-demand delivery of computations, storage, applications, and other IT resources through a cloud services platform over the internet with pay-as-you-go business model. Today's Cloud computing systems are built using fundamental principles and models of distributed systems. This course provides an in-depth understanding of distributed computing "concepts", distributed algorithms, and the techniques, that underlie today's cloud computing technologies. The cloud computing and distributed systems concepts and models covered in course includes: virtualization, cloud storage: key-value/NoSQL stores, cloud networking, fault-tolerance cloud using PAXOS, peer-to-peer systems, classical distributed algorithms such as leader election, time, ordering in distributed systems, distributed mutual exclusion, distributed algorithms for failures and recovery approaches, emerging areas of big data and many more. And while discussing the concepts and techniques, we will also look at aspects of industry systems such as Apache Spark, Google's Chubby, Apache Zookeeper, HBase, MapReduce, Apache Cassandra, Google's B4, Microsoft's Swan and many others. Upon completing this course, students will have intimate knowledge about the internals of cloud computing and how the distributed systems concepts work inside clouds.

ABOUT INSTRUCTOR :

Prof. Rajiv Misra is an Associate Professor in Department of Computer Science and Engineering at Indian Institute of Technology Patna, India. He obtained his Ph.D degree from IIT Kharagpur, M.Tech degree in Computer Science and Engineering from the Indian Institute of Technology (IIT) Bombay, and Bachelors of engineering degree in Computer Science from MNIT Allahabad. His research interests spanned a design of distributed algorithms for Mobile, Adhoc and Sensor Networks, Cloud Computing and Wireless Networks. He has contributed significantly to these areas and published more than 70 papers in high quality journals and conferences, and 2 book chapters. His h-index is 10 with more than 590 citations. He has authored papers in IEEE Transactions on Mobile Computing, IEEE Transaction on Parallel and Distributed Systems, Adhoc Networks, Journal of Parallel and Distributed Computing.

COURSE PLAN :

Week 1: Introduction to Clouds, Virtualization and Virtual Machine

1. Introduction to Cloud Computing: Why Clouds, What is a Cloud, Whats new in today's Clouds, Cloud computing vs. Distributed computing, Utility computing, Features of today's Clouds: Massive scale, AAS Classification: HaaS, IaaS, PaaS, SaaS, Data-intensive Computing, New Cloud Paradigms, Categories of Clouds: Private clouds, Public clouds

2. Virtualization: What's virtualization, Benefits of Virtualization, Virtualization Models: Bare metal, Hosted hypervisor

3. Types of Virtualization: Processor virtualization, Memory virtualization, Full virtualization, Para virtualization, Device virtualization

4. Hotspot Mitigation for Virtual Machine Migration: Enterprise Data Centers, Data Center Workloads, Provisioning methods, Sandipiper Architecture, Resource provisioning, Black-box approach, Gray-box approach, Live VM Migration Stages, Hotspot Mitigation

Week 2: Network Virtualization and Geo-distributed Clouds

1. Server Virtualization: Methods of virtualization: Using Docker, Using Linux containers, Approaches for Networking of VMs: Hardware approach: Single-root I/O virtualization (SR-IOV), Software approach: Open vSwitch, Mininet and its applications
2. Software Defined Network: Key ideas of SDN, Evolution of SDN, SDN challenges, Multi-tenant Data Centers: The challenges, Network virtualization, Case Study: VL2, NVP
3. Geo-distributed Cloud Data Centers: Inter-Data Center Networking, Data center interconnection techniques: MPLS, Google's B4 and Microsoft's Swan

Week 3: Leader Election in Cloud, Distributed Systems and Industry Systems

1. Leader Election in Rings (Classical Distributed Algorithms): LeLann-Chang-Roberts (LCR) algorithm, The Hirschberg and Sinclair (HS) algorithm
2. Leader Election (Ring LE & Bully LE Algorithm): Leader Election Problem, Ring based leader election, Bully based leader election, Leader Election in Industry Systems: Google's Chubby and Apache Zookeeper
3. Design of Zookeeper: Race condition, Deadlock, Coordination, Zookeeper design goals, Data model, Zookeeper architecture, Sessions, States, Usecases, Operations, Access Control List (ACL), Zookeeper applications: Katta, Yahoo! Message Broker

Week 4: Classical Distributed Algorithms and the Industry Systems

1. Time and Clock Synchronization in Cloud Data Centers: Synchronization in the cloud, Key challenges, Clock Skew, Clock Drift, External and Internal clock synchronization, Christians algorithm, Error bounds, Network time protocol (NTP), Berkley's algorithm, Datacenter time protocol (DTP), Logical (or Lamport) ordering, Lamport timestamps, Vector timestamps
2. Global State and Snapshot Recording Algorithms: Global state, Issues in Recording a Global State, Model of Communication, Snapshot algorithm: Chandy-Lamport Algorithm
3. Distributed Mutual Exclusion: Mutual Exclusion in Cloud, Central algorithm, Ring-based Mutual Exclusion, Lamport's algorithm, Ricart-Agrawala's algorithm, Quorum-based Mutual Exclusion, Maekawa's algorithm, Problem of Deadlocks, Handling Deadlocks, Industry Mutual Exclusion : Chubby

Week 5: Consensus, Paxos and Recovery in Clouds

1. Consensus in Cloud Computing and Paxos: Issues in consensus, Consensus in synchronous and asynchronous system, Paxos Algorithm
2. Byzantine Agreement: Agreement, Faults, Tolerance, Measuring Reliability and Performance, SLIs, SLOs, SLAs, TLAs, Byzantine failure, Byzantine Generals Problem, Lamport-Shostak-Pease Algorithm, Fischer-Lynch-Paterson (FLP) Impossibility
3. Failures & Recovery Approaches in Distributed Systems: Local checkpoint, Consistent states, Interaction with outside world, Messages, Domino effect, Problem of Livelock, Rollback recovery schemes, Checkpointing and Recovery Algorithms: Koo-Toueg Coordinated Checkpointing Algorithm

Week 6: Cloud Storage: Key-value stores/NoSQL

1. Design of Key-Value Stores: Key-value Abstraction, Key-value/NoSQL Data Model, Design of Apache Cassandra, Data Placement Strategies, Snitches, Writes, Bloom Filter, Compaction, Deletes, Read, Membership, CAP Theorem, Eventual Consistency, Consistency levels in Cassandra, Consistency Solutions
2. Design of HBase: What is HBase, HBase Architecture, Components, Data model, Storage Hierarchy, Cross-Datacenter Replication, Auto Sharding and Distribution, Bloom Filter, Fold, Store, and Shift

Week 7: P2P Systems and their use in Industry Systems

1. Peer to Peer Systems in Cloud Computing: Napster, Gnutella, FastTrack, BitTorrent, DHT, Chord, Pastry and Kelips.

Week 8: Cloud Applications: MapReduce, Spark and Apache Kafka

1. MapReduce: Paradigm, Programming Model, Applications, Scheduling, Fault-Tolerance, Implementation Overview, Examples
2. Introduction to Spark: Resilient Distributed Datasets (RDDs), RDD Operations, Spark applications: Page Rank Algorithm, GraphX, GraphX API, GraphX working
3. Introduction to Kafka: What is Kafka, Use cases for Kafka, Data model, Architecture, Types of messaging systems, Importance of brokers