Topics in Distributed Computer Systems
Comprehensive Examination Syllabus
Posted Spring 2011

Reference: Distributed Computing: Principles, Algorithms, and Systems, by A.D. Kshemkalyani and M. Singhal (Cambridge University Press 2008).

1. Overview of computing and computer systems
   - Different forms of computing
   - Different types of parallel systems
   - Characteristics of and motivation for distributed systems
   - Interconnection networks for parallel systems
     - For shared memory multiprocessor systems: multistage switch networks (e.g., Omega, Butterfly)
     - For multicomputer systems: torus, hypercube
   - Communication primitives: blocking/non-blocking, synchronous/asynchronous
   - Challenges in designing and building distributed systems

2. Model of distributed computations
   - Distributed programs and distributed executions
   - Space-time diagram of a distributed execution
   - Causal precedence relation of events
   - Logical vs. physical concurrency
   - Communication models relevant in distributed computing (FIFO, non-FIFO, CO)
   - Global state and cut
   - Past and future cones of event

3. Logical time
   - Causality of events, foundation of logical time
   - Implementing logical clocks
     - Scalar time: Data structures and rules for clock updates
     - Vector time: Data structures and rules for clock updates, differential technique to reduce communication overhead, Singhal-Kshemkalyani's implementation of the differential technique to reduce storage overhead
     - Matrix time: Data structures and rules for clock updates
     - Basic properties of all above
   - NTP for physical clocks synchronization

4. Global state and snapshot recording algorithms
   - Consistent global state
   - Issues in recording a global state
   - Snapshot recording algorithms
     - For FIFO channels: Chandy-Lamport, Spezialetti-Kearns
     - For non-FIFO channels: Lai-Yang, Mattern's
     - For CO: Acharya-Badrinath, Alagar-Venkatesan
     - Comparison of snapshot recording algorithms

5. Termination detection (Chapter 7)
   - Definition and system model
   - Termination detection using distributed snapshots
   - Termination detection by weight throwing
   - Topor's spanning-tree based termination detection
   - Chandrasekaran and Venkatesan's message-optimal termination detection algorithm

6. Reasoning with knowledge (Chapter 8)
   - The muddy children puzzle, what does it reveal?
   - Kripke structure and its use for the muddy children puzzle
   - The challenges in establishing common knowledge in asynchronous systems
   - Concurrent common knowledge (CCK) and achieving CCK with a snapshot-based algorithm
   - Knowledge and clocks (scalar, vector and matrix clocks)

7. Distributed mutual exclusions (Chapter 9)
   - Why mutual exclusion mechanisms developed for centralized, shared-memory systems fail in distributed systems
   - System model for distributed mutual exclusion
   - Requirements and performance metrics
   - Approaches for distributed mutual exclusion (token-based, non-token-based, and quorum-based)
     - Non-token-based algorithms: Lamport's algorithm, Richart-Agrwala algorithm, Singhal's dynamic information structure algorithm
     - Quorum-based algorithms: Maekawa's algorithm, Agarwal-El Abbadi algorithm
     - Token-based algorithms: Suzuki-Kasami's broadcast algorithm, Raymond's tree-based algorithm
     - Performance analysis of algorithms

8. Consensus and Agreement (Chapter 14)
   - Problem definition and assumptions
   - The Byzantine agreement problem, and its equivalent problems (the consensus problem, and the interactive consistency problem)
   - Agreement in a failure-free system
   - Agreement in synchronous systems with failure
     - Exponential algorithm
     - Polynomial algorithm
   - Impossibility of consensus in an asynchronous system