CSE535 Asynchronous Systems

YoungMin kwon

Course Overview

- Logical Time
- Global State
- Basic Graph Algorithms
- Termination Detection
- Mutual Exclusion
- Deadlock Detection
- Global Predicate Detection
- Distributed Shared Memory
- Consensus and Agreement

Logical Time

- Issues
  - No global time
  - Network Time Protocol (NTP): a few tens of msec accuracy -> cannot capture the causality
- Logical clock
  - Maps an event to a time
  - Preserves the causal precedence
  - Scalar time, Vector time, Matrix time

Global State

- A State of a Distributed System
  - State of local processes (memory, register, …)
  - Messages in transit
- Issues
  - All processes cannot take a snapshot at the same time
- Consistent Global State
  - A sent message is either received or is in transit
  - No message is received or is in transit without being sent.
- We will cover several snapshot algorithms

Basic Graph Algorithms

- Spanning Tree
- Shortest Path Routing
- Flooding
- Synchronizers
- Leader Election

Termination Detection

- Termination of a distributed computation
  - Every process is locally terminated and no message is in transit.
- Distributed Snapshots
- Weight Throwing

Mutual Exclusion

- Only one process is allowed to execute the Critical Section (CS) at any given time
- Token-based
  - Only a process with a token is allowed to enter CS
- Non-token based
  - Rounds of messages are exchanged to determine who can enter CS next.
- Quorum-based
  - Quorums are formed such that when two processes request a CS, only one is granted.

Deadlock Detection

- Dealing with Deadlocks
  - Deadlock prevention, Deadlock Avoidance, Deadlock detection
- Models of Deadlocks
  - Single-Resource model, AND model, OR model, P out of Q model
- Algorithms
  - Path-pushing algorithms, Edge-chasing algorithms, Diffusing computation-based algorithm, Global state detection-based algorithm

Global Predicate Detection

- Stable Predicates
  - Predicates that remain true once they become true.
  - E.g. Deadlock, Termination
- Unstable Predicates
  - Predicates that may hold intermittently
  - Modality: Possibly(ϕ), Definitely(ϕ)

Distributed Shared Memory

- Abstraction:
  - Access the data across the network using only read and write.
- Memory Consistency Models:
  - Strict consistency (atomic consistency), linearizability,
  - Sequential consistency
  - Causal consistency
  - Processor consistency
  - Slow memory
- Mutual Exclusion
  - Lamport's bakery algorithm, Lamport's WRWR mechanism,…
  - Wait freedom

Consensus and Agreement

- Failure-free system
- Crash Failures
- Byzantine Failures