CS 755 – System and Network Architectures and Implementation

Module 6 – Fault Tolerance

Martin Karsten

mkarsten@uwaterloo.ca

State Machine Replication

- multiple replicas execute *identical* operations
- *identical* vs. *equivalent*?
  - depends on program logic
  - hidden channels (dependencies)?
- consensus -> global ordering
  - in the presence of transient failures

System Parameters

- synchronous vs. asynchronous communication
  - validity of timeouts?
- reliable communication?
  - ordering, loss, duplication, corruption
- failures: stop vs. recover vs. byzantine
- group semantics
  - open vs. closed, static vs. dynamic

System Objectives

- Performance
- Availability
- Ordering Requirements
- Uniformity of Replication
  - proper replication for *correct* nodes
  - proper replication for *all* nodes
    - including crash/recovered nodes

Ongoing Research Topic

- still early on the maturity curve
  - fault-tolerant distributed systems
  - state machine replication
  - consensus
- many subtleties in parameters & requirements
  - difficult to compare approaches
- gap: theory vs. practice
  - see recent paper about Raft system