

CSE 452

Distributed Systems

What is a Distributed System?

- Multiple machines
 - machines are faulty
 - Concurrency!
- Connected by a network
 - the network is faulty

Leslie Lamport

Why?

- harness the power of multiple machines
- horizontal scaling
- redundancy/replication
- fault tolerance
- Placing data near users

How hard is it to build a DS?

- hard to maintain "coherence"
 - replication makes updates hard
- multiple machines working — one could fail
 - partial failure
- concurrency

Fault Model

★ = Standard fault model

- list of failures we plan to tolerate
 - tolerate automatically
- what failures are possible?
 - power goes out — machines crash ★
 - network
 - reordering ★
 - dropped ★
 - corruption
 - duplicate ★
 - delay ★
 - msg injection
 - unplugging the cable — drop