

CSE 550: Systems for all

Au 2021

Ratul Mahajan

What is a distributed system?

Separate time domains

- Processes have independent clocks
- Impossible to perfectly synchronize all clocks
 - Non-trivial message delays

Separate fault domains

- Processes can fail independently

Challenges of independent time

Event ordering is hard

- Key to reasoning about system behavior
- Program logic may depend on ordering -- e.g., Unix make

Consistent snapshot is hard

- Debugging bad system states – e.g., detect deadlocks
- Recovery after failures, migration


Two approaches to aligning times

Synchronize physical clocks

Use logical clocks

Synchronizing physical clocks (1/2)

One solution: Time server broadcasts time, clients set their own clocks




GPS uses a generalization of this idea

- Multiple satellites
- Triangulation yields location too

Min delay is more reliable
Set clock to received time + min delay

Synchronizing physical clocks (2/2)

Another solution: Clients interrogate time server, set clock to (server time + $RTT/2$)



NTP uses this idea

- Use multiple time servers

Asymmetric latencies cause problems

Logical clocks

Physical time is not needed for many tasks

Need ordering and causality, which logical clocks can provide

We read two versions

- Lamport clocks
- Vector clocks

Over to Jiarong