Distributed Systems

ID2201

clocks

Johan Montelius

What is time?

- No absolute time:
  - but for most practical purposes it is
- Assuming absolute time:
  - how can we tell the time
- If we're not synchronized what can we do:
  - causal effects give us logical time that order some events
  - does our distributed system give us a consistent view of logical time

Why?

- Give me an example were it is important to know the time.

Synchronization problems

- Who has the correct time?
  - earths rotation - UT1
  - one “atomic” clock - UTC
- Even if wee all agree, how do we keep nodes synchronized?
  - it takes time to send a signal
  - in between signals nodes will drift
  - how often can we send signals

Synchronization

- External synchronization (accuracy)
  - each node in our network is synchronized with an external (global) source within a bound D
- Internal synchronization (precision)
  - every pair of nodes in our network are synchronized within a bound D

Drift and Monotonicity

- Drift is change in how well one clock can measure a time interval.
  - can drift change?
- Monotonicity is the property that time always moves forward.
  - can this be hard to guarantee?
- Correctness often means monotonicity and low drift.
- A correct clock might not be synchronized

How to synchronize

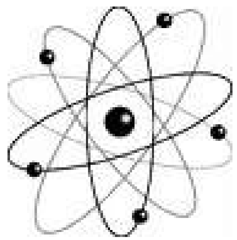

~300 ns

033-415783

~100 ms

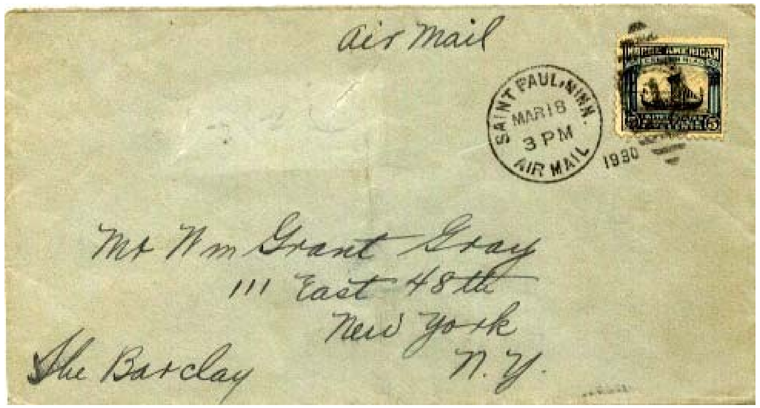

Asynchronous networks

- One server is connected to external source and used to synchronize other nodes in the network.
- Problem is of course that round-trip times are unknown and that they vary.
- A minimum propagation time can be known.
- How do we synchronize our nodes?

Christian's algorithm

$$T = t + r/2$$

$$diff = r/2$$

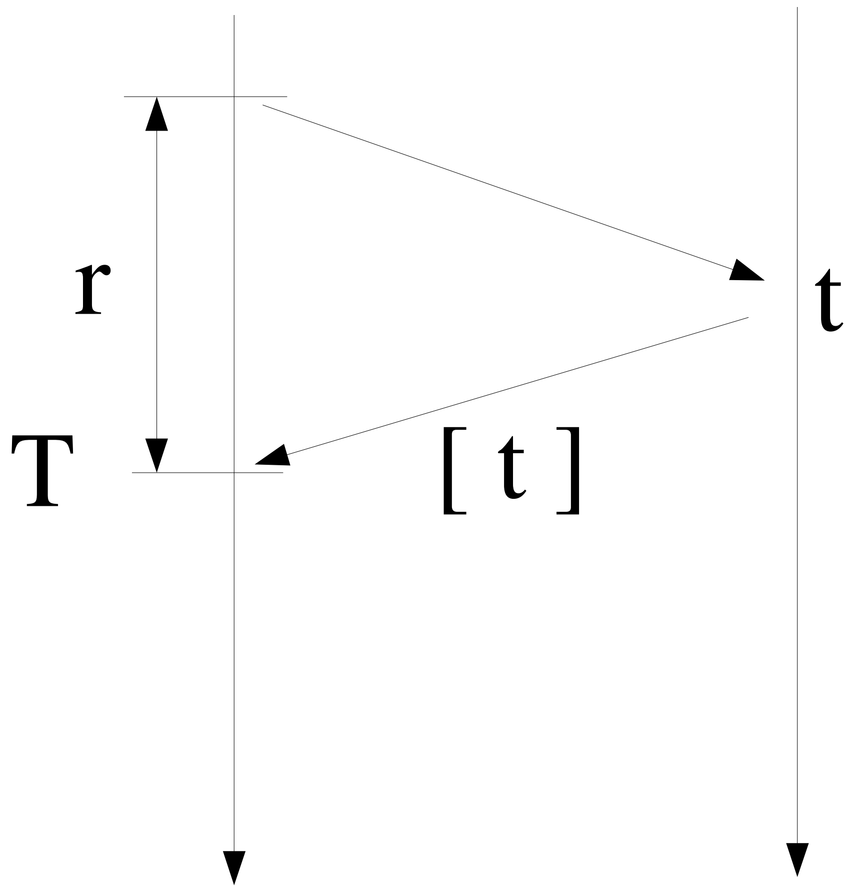

Christian's algorithm

$$T = t + r/2$$

$$diff = r/2 - m$$

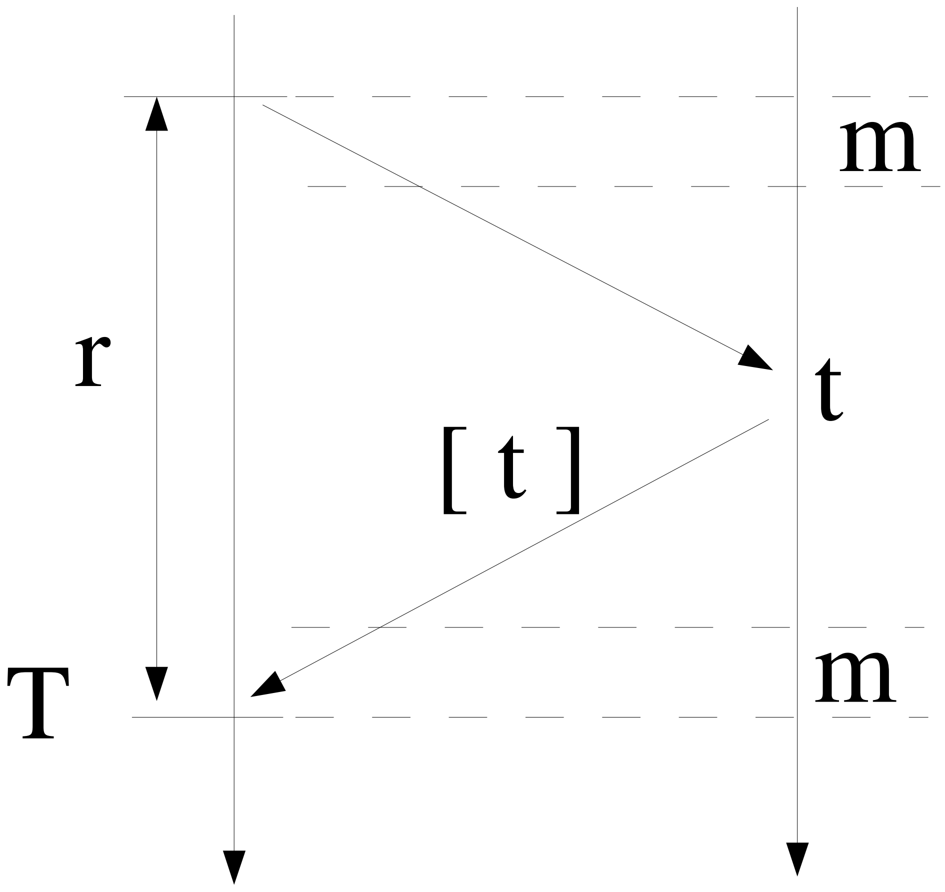

Berkeley algorithm

- Used to synchronize a network of nodes.
- Send requests to all nodes.
- Collect ti and calculate an “average” time T.
- Send out individual deltas to each node.

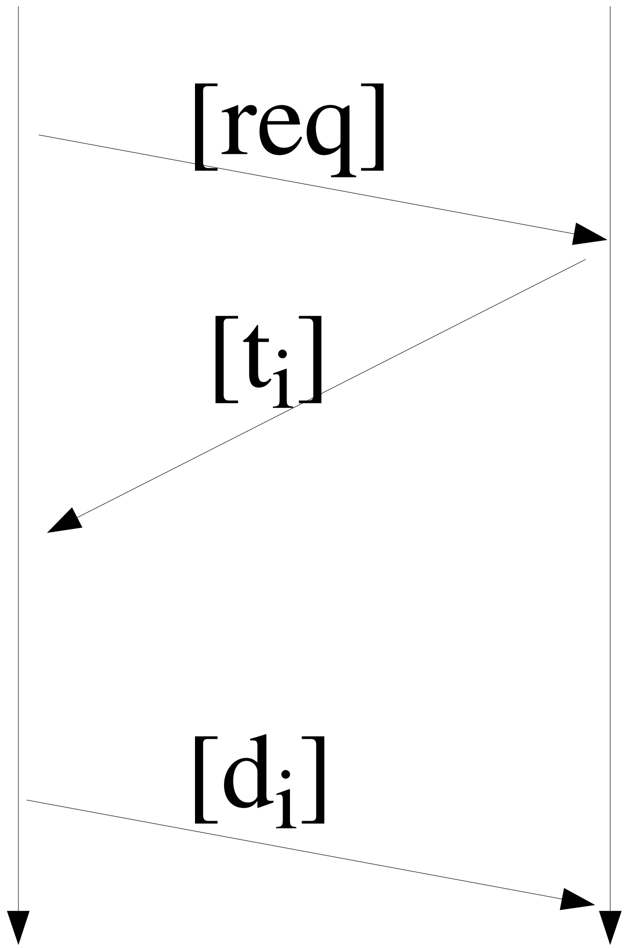

NTP – network time protocol

- An architecture targeting reliability and wide area networks.
- A hierarchy of servers: *stratum-1* connected to external sources.
- Fault tolerant: servers can be degraded to lower stratum if external source is lost, client can connect to secondary servers.
- Several synchronization protocols: LAN multicast, request and, synchronous.

NTP

- Client collects: t_0, t_1, t_2 and t_3.
- The two messages are independent.
- Similar to Christan's but with better estimate of delay.

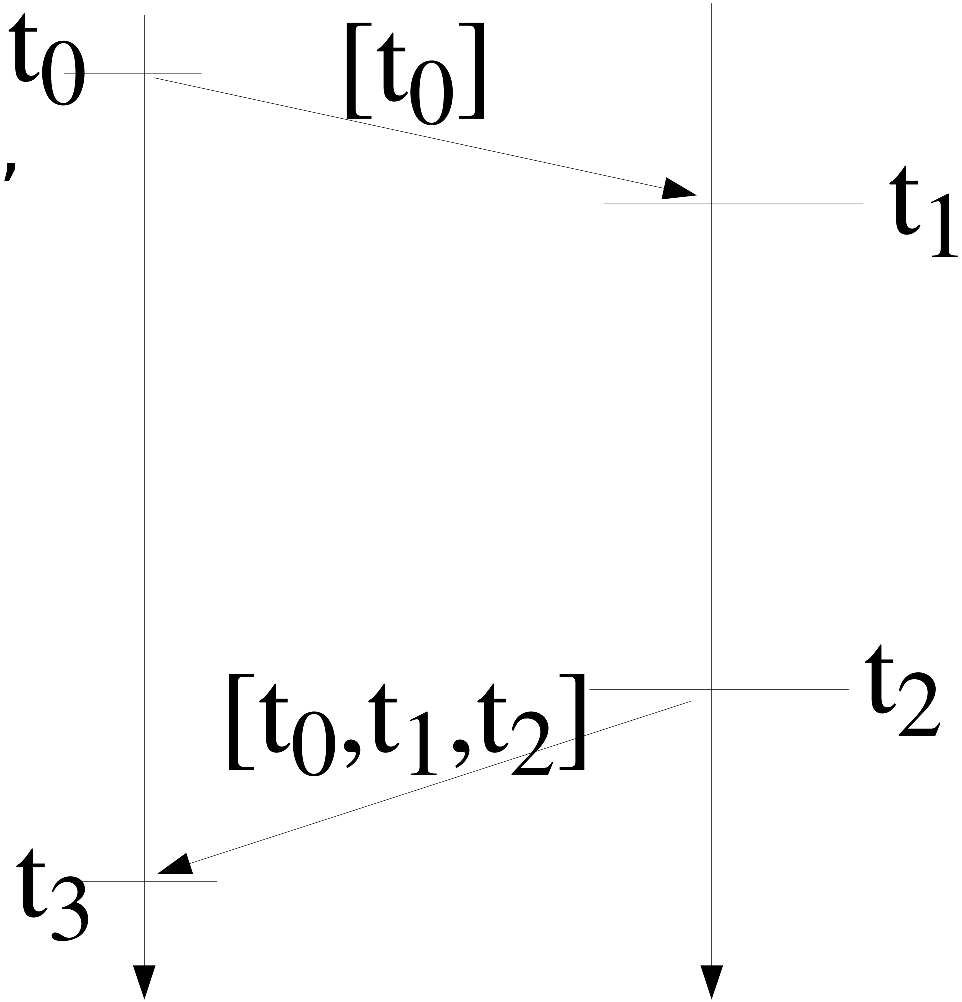

Summary

- Clocks give us a estimate on real time.
- Clocks can be synchronized internally and/or to an external source.
- Synchronization limited by network jitter and clock drift.
- Synchronize to UTC
  - NTP connected over Internet: a few 10 ms
  - local GPS clocks connected to LAN: < 1 ms
  - on board GPS clock: few µs to ns