Brief Announcement: Gradient Clock Synchronization in Sensor Networks*

Lennart Meier
meier@tik.ee.ethz.ch

Lothar Thiele
thiele@tik.ee.ethz.ch

Computer Engineering and Networks Laboratory
Swiss Federal Institute of Technology (ETH) Zurich

Categories and Subject Descriptors: F.2.2 [Analysis of Algorithms and Problem Complexity]: Nonnumerical Algorithms and Problems

General Terms: Algorithms, Theory.

Keywords: lower bounds, sensor networks, synchronization.

We examine gradient clock synchronization [1], where the difference between any two network nodes' clocks is required to be upper-bounded by a non-decreasing function of their distance: A node has to be synchronized better to nearby nodes than to faraway nodes. We look at the gradient property in a system model that is typical for sensor networks, provide a lower bound for the achievable synchronization quality in our model, and discuss its relation to the bound in the model from [1].

Time information is degraded by clock drift, and it cannot be communicated without loss due to delay uncertainties. Upper bounds on drift and delay lead to lower bounds on the synchronization error. For instance, the worst-case error between two nodes with delay uncertainty D is in $\Omega(D)$. The central theorem of [1] states that in a network with maximal delay uncertainty D, the error between two nodes with constant delay uncertainty is in $\Omega(\frac{\log D}{\log\log D})$, i.e. the error grows with D although the delay uncertainty between the two nodes remains constant.

The system model in [1] assumes unbounded communication frequency. The upper-bounded drifts and delays are unknown to the nodes, and the lower bound is derived by letting an adversary modify them. We use a system model that we consider more appropriate for sensor networks. Clock drifts are still bounded, but we assume also the communication frequency to be bounded; in sensor networks, communication is expensive in terms of energy, and hence infrequent communication is desirable. As a consequence of infrequent communication, we neglect delay uncertainties and eliminate them from our analysis, i.e. communication occurs in zero time. This is reasonable: As the frequency of communication decreases, the uncertainty due to clock drift increases, while the uncertainty due to message delays remains constant. Two particular characteristics of sensor nodes further strenghten the case for the dominance of the drift:

On the one hand, time-stamping on sensor nodes can be done at a low level, such as in the MAC layer, leading to a small delay uncertainty. Recent algorithms reduce it to a few microseconds, e.g. by using packet streams or reference broadcasts. On the other hand, sensor nodes typically employ inexpensive oscillators with drifts of up to 100 ppm.

A numeric example: If the delay uncertainty is 1 μs and the clock drift's absolute value is bounded by 100 ppm, then after 5 ms, the drift's contribution to the uncertainty equals that of the delay. After one hour, it is 720000 times larger. Even for "optimistic" values of 1 ms (uncertainty) and 10 ppm (drift), drift and uncertainty have equal impact after 50 s.

In our model, the synchronization service uses only the communication that takes place anyway for achieving the overall goal of the sensor network; the time information is sent piggyback with the application data. Therefore, the communication pattern cannot be influenced by the synchronization algorithm. Examples of such sensor networks are those in which environmental data is collected periodically but communicated only sporadically, e.g. when time-critical data is recorded, when a master node explicitly requests the data, or when the solar cells of a node are providing sufficient energy for communication.

To derive any bounds, we have to make some quantitative assumption about the communication frequency. We assume that a node communicates at least once every d time units with each of its neighbors. In practice, this could be a bound required by the application. The adversary can modify the communication pattern by shifting the times at which communication events occur. He can also change the clock drifts.

Our main contribution is to show that an analogous lower bound as in [1] exists also in our model. This is not obvious: In [1], the bound is based on delay uncertainties, which do not exist in our model. In our model, the adversary that shifts the time at which an communication event occurs also has to modify the clock drifts of both nodes involved in the communication.

We require the synchronization algorithm to maintain at all times a maximum error $\Delta L := |L_i - L_j|$ between any two neighboring nodes' clocks L_i and L_j. Note that this implies the gradient property, as the maximum error between two nodes that are an arbitrary hop distance s apart is upper-bounded by $s \cdot \Delta L$.

We show that in a network consisting of a chain of n nodes with maximal clock drift $\hat{\rho}$ and maximal time d between communications, the smallest value ΔL that can be guaranteed by any algorithm is bounded by $\Delta L \geq \frac{\hat{\rho} d}{8(1+\hat{\rho})} \frac{\log(n-1)}{\log\left(\frac{8(1+\hat{\rho})}{\hat{\rho}} \log(n-1)\right)}$. The bound increases with increasing n, d, and $\hat{\rho}$.

*For a full version of this paper, see [2]. This work was supported (in part) by the National Competence Center in Research on Mobile Information and Communication Systems (NCCR MICS), a center supported by the Swiss National Science Foundation under grant number 5005-67322.

PODC'05, July 17–20, 2005, Las Vegas, Nevada, USA.
ACM 1-58113-994-2/05/0007.

REFERENCES

[1] R. Fan and N. Lynch. Gradient clock synchronization. In *Proceedings of PODC 2004*, pages 320–327, 2004.

[2] L. Meier and L. Thiele. Gradient clock synchronization in sensor networks. TR 219, Computer Engineering and Networks Laboratory, ETH Zurich, 2005.