Precise Timestamp-Free Network Synchronization

D. Richard Brown III and Andrew G. Klein
Worcester Polytechnic Institute
100 Institute Rd, Worcester, MA 01609
Email: {drb,klein@wpi.edu}

Abstract—**This paper describes an approach to master/slave network synchronization based on bidirectional message exchanges without the use of timestamps. Rather than the usual approach of exchanging digital timestamps through a dedicated synchronization protocol, an approach is described in which synchronization information is conveyed implicitly at the physical layer through the timing of the master node's responses to the slave nodes. This approach can reduce overhead and allow the embedding of synchronization functions in existing network traffic. A timestamp-free synchronization protocol is described and its performance is quantified in the presence of delay estimation error and stochastic local oscillator dynamics. A filtering framework is also developed to allow each slave node to accurately infer and correct local clock drifts from multiple noisy clock offset estimates. Based on fundamental delay estimation bounds for narrowband signals, numerical results show that synchronization among the slave nodes can be achieved quickly and that the resulting steady-state accuracy can be sufficient to support distributed transmission techniques requiring carrier phase alignment, e.g. distributed beamforming.**

Index Terms—**synchronization, delay estimation, oscillator dynamics, wireless communication, distributed communication systems**

I. Introduction

Synchronization is the process of establishing a common notion of time among two or more entities. In the context of communication networks, synchronization enables coordination among the nodes in the network and can facilitate scheduling of communication resources, interference avoidance, event detection/ordering, data fusion, and coordinated wake/sleep cycles [1]. Precise synchronization, to the order of a fraction of a carrier period, can also enable distributed transmission schemes such as retrodirective distributed beamforming [2].

A variety of synchronization protocols and systems have been developed over the last 30 years, including Network Time Protocol (NTP) [3], Precision Time Protocol (PTP) [4], and the Global Positioning System (GPS) [5], [6]. These systems have limited applicability in wireless networks with cost/battery/computationally-constrained devices [1]. While GPS can be used to discipline oscillators, its applicability is typically limited to infrastructure equipment since GPS receivers require line of sight channels to GPS satellites, are slow to acquire accurate timing estimates, are expensive, and have prohibitive power consumption. NTP and PTP are considered to be prohibitively complex for resource constrained devices like sensor nodes [7], [8] and have limited accuracy. Synchronization protocols developed in the last decade specifically for sensor networks include [1], [7], [9]–[15].

A common theme running through all of the cited synchronization protocols and systems is that they are based on application-layer or MAC-layer exchanges of digital timestamps between nodes in the network. Digital timestamps inherently have a minimum resolution, limiting accuracy, and add overhead to the network traffic. Timestamp-free synchronization techniques were originally studied in the context of natural phenomena, e.g. synchronization of the firing rates of fireflies, and led to formal mathematical models for systems of pulse-coupled oscillators in [16], [17] and application of these models to wireless sensor networks in [18]–[20]. While these studies represented an exciting paradigm shift with respect to the prior work, a limitation of the pulse-coupled oscillator literature is that it is based on unidirectional transmissions and assumes negligible propagation delays. This inherently limits the synchronization accuracy of these methods.

Precise timestamp-free synchronization techniques accounting for propagation delays were developed recently in the context of retrodirective distributed transmit beamforming systems [2], [21]–[24]. These techniques are based on serial exchanges of unmodulated beacons through a network with a ring architecture and achieve precise synchronization through estimation of RF carrier phase and frequency offsets. While these techniques were shown to be effective for carrier phase and frequency synchronization, they require some amount of overhead since they are based on a dedicated synchronization protocol. Also, due to the unmodulated nature of the beacons, the synchronization solution achieved with these schemes has an inherent carrier period ambiguity. To resolve this ambiguity, an additional time reference is necessary in these systems.

This paper develops an approach to network synchronization based on bidirectional message exchanges between a master node and N slave nodes. Rather than requiring a dedicated synchronization protocol and exchanging digital timestamps, our approach is to convey implicit timing information *in the physical layer* through the timing of the responses from the master node to the slave nodes. Since no timestamps are exchanged, synchronization functions can be embedded in existing network traffic. Also, since the technique is based on bidirectional message exchanges and delay estimation, it accounts for propagation delay and its accuracy is only limited by fundamental estimation bounds. Our focus is on the development of precise synchronization schemes in support of distributed transmission techniques, e.g. distributed beamforming and/or virtual MIMO, hence we explicitly consider the effects of local clock drift and stochastic oscillator dynamics and

develop a filtering framework to optimally track and correct local clock offsets and drifts given only occasional noisy estimates of the local clock offsets. Based on fundamental delay estimation bounds for narrowband signals, our numerical results for timestamp-free synchronization with local filtering show that synchronization among the slave nodes can be achieved quickly and that the resulting steady-state accuracy can be on the order of tens of picoseconds with realistic oscillator and signaling parameters.

II. SYSTEM MODEL

We consider a time-division duplexed (TDD) wireless network with $N + 1$ nodes with node 0 serving as the "master" node and nodes $n = 1, \ldots, N$ comprising the "slave" nodes. The master node could be an infrastructure device such as an access point or it could simply be a node in the network promoted to serve as the network's time server. It is assumed that the slave nodes can all communicate directly with the master node.

To ease exposition, we assume the channel between slave node n and the master node is time-invariant and single-path with propagation delay denoted as τ_n. Since the network is TDD, the propagation delay is assumed to be reciprocal. Basic electromagnetic principles have long established that channel reciprocity holds at the antennas when the channel is accessed at the same frequency in both directions [25]. Channel reciprocity can also be quite accurate at intermediate-frequency (IF) and/or baseband if a reciprocal transceiver architecture is used [26] and can be further improved through transceiver calibration techniques to remove I/Q imbalance effects [27], [28].

The nodes in the network do not possess a common notion of time. The following section presents a model of local and reference time that will be subsequently used in the description and analysis of the synchronization protocol.

A. Local Clock Dynamics

The focus of this paper is the description and analysis of a synchronization technique for devices in a wireless ad-hoc network. To support this focus, it is necessary to explicitly present a model of local time at each node and describe how the local time at each node relates to the reference time t at the master node. The notation t refers to the reference time, i.e. the "true" time, in the system. All time-based quantities such as propagation delays and/or frequencies are specified in reference time unless otherwise noted. The notation t_n corresponds to the local notion of time at node n with $t_0 = t$. The clock offset at node n at time t with respect to the master node is denoted as $\Delta_n(t) = t_n - t$. An event a that occurs at local time $t_n^{(a)}$ at node n occurs at local time $t^{(a)} = t_n^{(a)} - \Delta_n(t^{(a)})$ at the master node.

None of the slave nodes have knowledge of the reference time t. Based on the two-state models in [29], we define the discrete-time state of the n^{th} node's clock at time $t = kT$ as

$$\boldsymbol{x}_n[k] = \left[\Delta_n[k], \dot{\Delta}_n[k] \right]^\top$$

where $\Delta_n[k]$ corresponds to the clock offset at node n with respect to the master clock. The state update of the n^{th} transmit node's clock is governed by

$$\boldsymbol{x}_n[k+1] = \boldsymbol{f}(T)\boldsymbol{x}_n[k] + \boldsymbol{u}_n[k] \tag{1}$$

where T is the state update period, the state update matrix

$$\boldsymbol{f}(T) = \begin{bmatrix} 1 & T \\ 0 & 1 \end{bmatrix}, \tag{2}$$

and the process noise vector $\boldsymbol{u}_n[k] \overset{\text{i.i.d.}}{\sim} \mathcal{N}\left(\boldsymbol{0}, \boldsymbol{Q}_n(T)\right)$ captures the effect of stochastic phase noise in the n^{th} node's local oscillator. The covariance of the discrete-time process noise is derived from a continuous-time model in [29] and is given as

$$\boldsymbol{Q}_n(T) = T \begin{bmatrix} p_n + q_n \frac{T^2}{3} & q_n \frac{T}{2} \\ q_n \frac{T}{2} & q_n \end{bmatrix} \tag{3}$$

where p_n (units of seconds) and q_n (units of Hertz) are the process noise parameters corresponding to white frequency noise and random walk frequency noise, respectively. The process noise parameters p_n and q_n can be estimated by fitting the theoretical Allan variance $\sigma_y^2(\tau) = p_n/\tau + q_n\tau/3$ to experimental measurements of the Allan variance over a range of τ values. For notational convenience hereafter we will assume all of the transmitter oscillators to have the same process noise parameters such that $p_n = p$ and $q_n = q$ for all $n = 1, \ldots, N$.

III. TIMESTAMP-FREE SYNCHRONIZATION PROTOCOL

This section describes a timestamp-free synchronization protocol based on bidirectional message exchanges between a slave node and the master node. Figure 1 shows the interactions between node n and the master node. We assume that these interactions occur over a short interval during which the discrete-time state $\boldsymbol{x}_n[k]$ is approximately constant.

The protocol begins with node n transmitting a packet to the master node at arbitrary local time $t_n^{(a)}$. The packet arrives at the master node at local time

$$t^{(b)} = t_n^{(a)} - \Delta_n[k] + \tau_n$$

where τ_n is the propagation delay between node n and the master node. The master node then transmits a packet back to node n at time $t^{(c)}$ where $t^{(c)}$ is selected such that

$$\frac{t^{(b)} + t^{(c)}}{2} \pmod{T_0} = 0 \tag{4}$$

where T_0 is master node *clock tick period*. Note that, unlike the usual sender/receiver synchronization protocol, e.g. [7], no timestamps are exchanged between the nodes here. Implicit timing information is embedded in the master node's response to node n by selecting $t^{(c)}$ so that a local clock tick the master node is centered between $t^{(b)}$ and $t^{(c)}$. Since the channel is reciprocal, node n receives the reply packet from the master node at local time

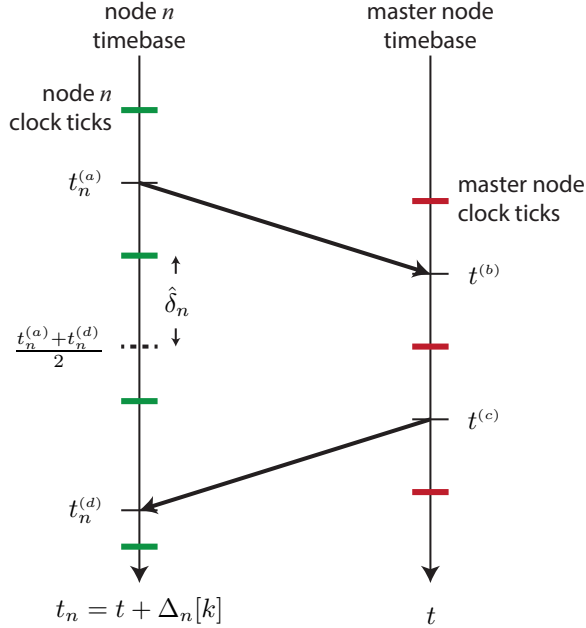$$t_n^{(d)} = t^{(c)} + \Delta_n[k] + \tau_n.$$

Fig. 1. Timestamp-free synchronization bidirectional message exchange.

Node n can now estimate its clock tick offset with respect to the master node by calculating

$$\hat{\delta}_n = \left(\frac{t_n^{(a)} + t_n^{(d)}}{2} \right)_{T_0} \tag{5}$$

where the notation $(z)_{T_0}$ corresponds to wrapping z to the interval $[-T_0/2, T_0/2)$. Note that, in the absence of estimation error and wrapping ambiguity, $\hat{\delta}_n = \boldsymbol{h}\boldsymbol{x}_n[k]$ with $\boldsymbol{h} = [1,0]$. The offset estimate in (5) can be used directly for immediate clock offset correction at node n or as an input to a filtering algorithm to correct both clock offsets and drifts.

This timestamp-free synchronization technique accounts for propagation delay, hence its accuracy is only limited by the fundamental bounds of delay estimation and the accuracy of the channel reciprocity assumption. Note that the technique also does not require single-path channels. As long as the *effective* propagation delay is reciprocal, e.g. both nodes use the strongest path to estimate arrival times, the propagation delays cancel in the calculation of $\hat{\delta}_n$. The overhead incurred in this approach is the delay caused by the master node in timing its response to node n so that a local clock tick at the master node is centered between $t^{(b)}$ and $t^{(c)}$. This overhead can be made small by making the clock ticks at the master node more frequent, at the expense of increased potential for wrapping ambiguity.

IV. DELAY ESTIMATION AND FILTERING

The timestamp-free synchronization protocol described in the previous section requires the master and slave nodes to both be able to accurately estimate time of arrival of the wirelessly transmitted packets. It is well-known, e.g. [30, p.55], that the Cramer-Rao lower bound (CRLB) for delay estimation of signals observed in additive white Gaussian noise

(AWGN) is inversely proportional to the signal-to-noise ratio (SNR) and the mean squared bandwidth of the signal. The analysis leading to this bound, however, assumes a baseband signal. Weiss and Weinstein [31], [32] showed that *narrowband* signal delay estimation performance also improves with carrier frequency when the SNR exceeds a certain threshold. Intuitively, above this SNR threshold, the presence of the carrier in the narrowband signal provides additional detail in the correlator output that can be used to refine delay estimates to a fraction of a carrier period.

From [31], [32], the CRLB for estimating the time of arrival of a *narrowband* signal in AWGN can be written as

$$\text{var}(\hat{\tau}) \geq \frac{\pi}{WT \cdot \omega_0^2 \cdot \text{SNR}} \tag{6}$$

where WT is the time-bandwidth product of the narrowband waveform, ω_0 is the radian carrier frequency of the narrowband waveform, and SNR is the pre-integration SNR at the receiver. It is assumed here that T is short enough so that the effect of any oscillator dynamics over the observation interval are negligible. As an example, in the case of a pre-integration SNR of 10 dB, bandwidth $B = 50$ MHz, waveform duration $T = 10$ μs, and carrier frequency $\omega_0 = 2\pi \cdot 1$ GHz, the narrowband CRLB implies that RMS delay estimation errors can be as small as approximately 4 ps. It was shown in [33] that, by using a two-step maximum-likelihood estimator, the narrowband CRLB could be achieved for these values of system parameters. The experimental results reported in [34] with similar signaling parameters in an outdoor environment and with off-the-shelf hardware yielded RMS delay estimation errors of less than 10 ps.

While the proposed timestamp-free synchronization protocol can be used by directly applying the offset estimates calculated in (5) for local clock correction at each slave node, better synchronization accuracy can be achieved through repeated bidirectional exchanges with the master node and filtering of the delay estimates. The filtering approach is also natural in the context of the Gauss-Markov state space model in (1). We assume that the observations at node n in timeslot k are given by

$$\hat{\delta}_n[k] = \Delta_n[k] + V[k] = \boldsymbol{h}\boldsymbol{x}_n[k] + V[n] \tag{7}$$

with $\boldsymbol{h} = [1,0]$ and where $V[k] \sim \mathcal{N}(0, R)$ represents the estimation error with R specified according to (6), assumed to be spatially and temporally independent, identically distributed, and independent of the process noise $\boldsymbol{u}_n[k]$ for all n and k. This observation model implicitly assumes that the observation interval is short so that any frequency estimate gleaned from carrier or symbol timing is not useful.

In timeslot k, the slave node n that performed a bidirectional message exchange with the master node updates its state predictions, state estimates, and the associated covariance matrices with the usual Kalman filter update [35]. Denoting state estimates and predictions at node n as $\hat{\boldsymbol{x}}_n[k|k]$ and $\hat{\boldsymbol{x}}_n[k+1|k]$, respectively, and the associated estimation and prediction covariance matrices $\boldsymbol{\Sigma}_n[k|k]$ and $\boldsymbol{\Sigma}_n[k+1|k]$, the

filter recursion is specified as

$$\hat{\boldsymbol{x}}_n[k|k] = \hat{\boldsymbol{x}}_n[k|k-1] + \boldsymbol{K}_n \left(\hat{\delta}_n[k] - \boldsymbol{h}\hat{\boldsymbol{x}}_n[k|k-1] \right)$$

$$\hat{\boldsymbol{x}}_n[k+1|k] = \boldsymbol{f}(T)\hat{\boldsymbol{x}}_n[k|k]$$

$$\boldsymbol{\Sigma}_n[k|k] = \boldsymbol{\Sigma}_n[k|k-1] - \boldsymbol{K}_n(\boldsymbol{h}\boldsymbol{\Sigma}_n[k+1|k]\boldsymbol{h}^\top + R)\boldsymbol{K}_n^\top$$

$$\boldsymbol{\Sigma}_n[k+1|k] = \boldsymbol{f}(T)\boldsymbol{\Sigma}_n[k|k](\boldsymbol{f}(T))^\top + \boldsymbol{Q}$$

with $\boldsymbol{K}_n = (\boldsymbol{\Sigma}_n[k|k-1]\boldsymbol{h}^\top)/(\boldsymbol{h}\boldsymbol{\Sigma}_n[k+1|k]\boldsymbol{h}^\top + R)$ and where $\boldsymbol{f}(T)$ and $\boldsymbol{Q}$ are defined in (2) and (3), respectively. The remaining $N-1$ slave nodes update their state estimates and predictions with $\boldsymbol{K}_m = 0$ for all $m \neq n$ since they do not receive an observation.

Remark: The modulo and wrapping operations in (4) and (5) imply that the observations received at node n are nonlinear in the state $\boldsymbol{x}_n[k]$, whereas we have assumed an observation model in (7) that is linear (affine) in the state. The linear model is reasonable, however, if the slave node clock offsets are small with respect to the master node clock tick period T_0. This can be achieved either by making T_0 large for initial message exchanges or by having the master node initially broadcast a single low-precision, e.g. 1 ms resolution, timestamp to all of the slave nodes prior to commencing the timestamp-free synchronization protocol.

V. NUMERICAL RESULTS

The numerical results in this section assume a network with $N = 10$ slave nodes, i.i.d. Gaussian distributed initial clock offsets $\Delta_n[0]$ with standard deviation 5 ms, and i.i.d. uniformly distributed initial drifts $\dot{\Delta}_n[0]$ with maximum deviation of ± 10 parts-per-million (ppm) for $i = 1, \ldots, N$. The stochastic oscillator parameters p and q were derived via a least-squares fit from a table of typical Allan variance characteristics of a "good" crystal oscillator in [36]. The master node clock period T_0 was set to 100 ms and the master/slave messaging period T was set to 250 ms. The measurement error standard deviation $\sqrt{R}$ was set to 20 ps, consistent with the system parameters in Section IV. In each iteration the slave node was chosen randomly, with equal probability for each slave node.

Figure 2 shows one realization of the uncorrected clock offsets at the 10 slave nodes in the system. In this case, since the stochastic oscillator parameters were chosen based on a "good" crystal oscillator and the duration of the simulation is only 250 seconds, the clock offset trajectories are dominated by the initial offsets and the ± 10 ppm oscillator drifts.

To demonstrate the efficacy of simple timestamp-free synchronization without filtering, Figure 3 shows corrected clock offsets at the 10 slave nodes in the system using the timestamp-free synchronization protocol. In this example, the clock offset estimates at each slave node were directly applied to correct the local clock without filtering. Each time a slave node performs a bidirectional message exchange with the master node, it accurately corrects is local offset but does not correct its local drift. The effect of uncorrected local drift can be seen in the approximately linear clock offset trajectories in Figure 3, where nodes with larger rate offsets tend to
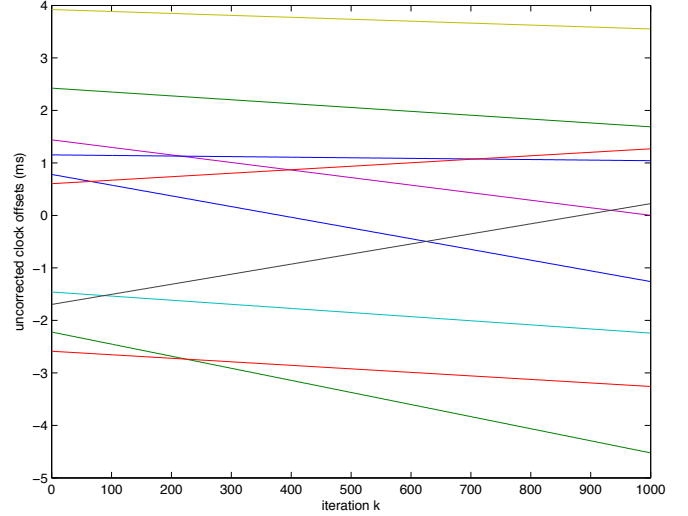

Fig. 2. Uncorrected clock offsets example with different colors corresponding to the clock offsets at different slave nodes.

have larger excursions from zero. Although the nodes are synchronized to a precision better than 0.1 ms in this example, the effect of uncorrected drift is evident in the steady-state performance and would be exacerbated in systems with less frequent master/slave message exchanges.
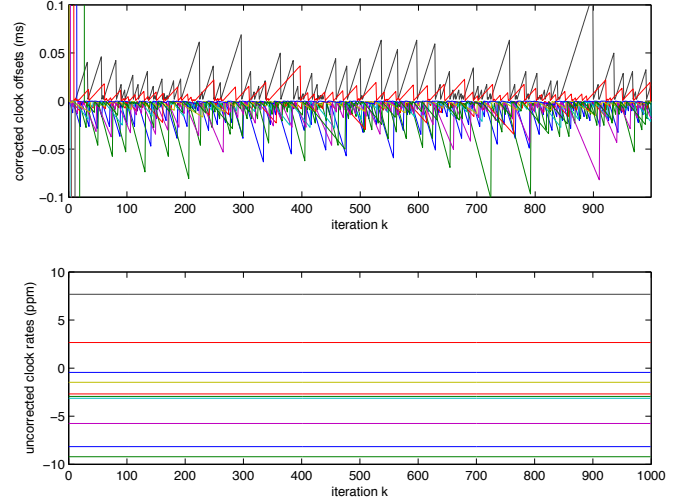

Fig. 3. Corrected clock offsets example (no filtering).

To demonstrate the performance improvement that can be achieved by filtering the offset estimates, Figure 4 shows corrected clock offsets and drifts at the 10 slave nodes in the system where each node locally filters the delay estimates to infer and correct local drifts. Note the change of axis units/limits with respect to Figure 3. Initially, the clock offsets at each slave node are dominated by the initial offsets and the uncompensated drifts, but after each slave node has received a few offset estimates, it is able to infer and accurately correct the initial offset and local drift. Offset correction in this case is significantly improved with respect to Figure 3 with a standard

deviation of approximately 11 ps over the last 500 iterations of the simulation. A small amount of residual drift is still evident in these results, but the performance tends to be dominated by the unpredictable stochastic dynamics of the local oscillators. Drift compensation is also quite effective in this case with corrected clock rates better than 0.01 parts-per-billion (ppb).
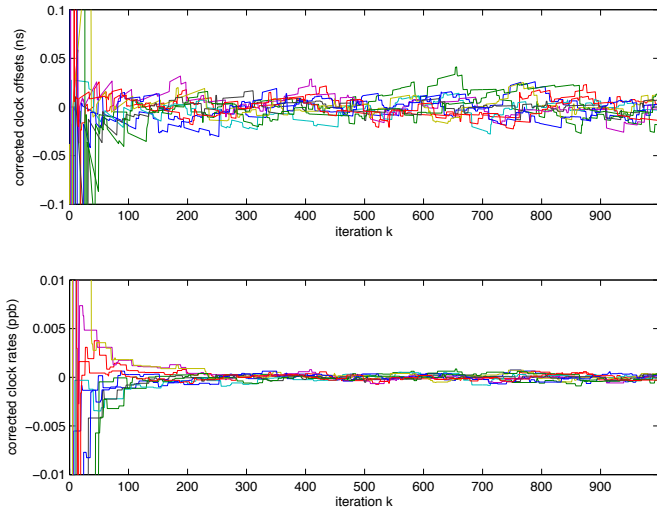

Fig. 4. Corrected clock offsets example with filtering.

While the numerical results in this section are for a specific realization of the clock process and measurement noises, they are typical of the results obtained in our testing.

VI. CONCLUSION

This paper describes a new approach to master/slave network synchronization that exploits the physical layer properties of narrowband wireless signals and a filtering framework motivated by the state-space description of local oscillator dynamics to achieve precise synchronization performance without the exchange of timestamps. Since no timestamps are exchanged, synchronization functions can be embedded in existing network traffic. Also, since the protocol is based on bidirectional message exchanges, it accounts for propagation delay and its accuracy is only limited by fundamental estimation bounds. Our numerical results for timestamp-free synchronization with local filtering show that synchronization among the slave nodes can be achieved quickly and that the resulting steady-state accuracy can be sufficient to support distributed transmission techniques requiring carrier phase alignment, e.g. virtual MIMO and/or distributed beamforming.

While this paper has shown that timestamp-free synchronization can be effective in master/slave networks, an important extension of this work is to develop timestamp-free synchronization protocols for general network architectures including networks with multi-hop communication to a time server. An analytical characterization of the achievable steady-state performance of timestamp-free synchronization as a function of the number of nodes in the network, local oscillator parameters, and measurement noise parameters would also be valuable toward understanding the effect of these parameters on synchronization accuracy. Asymptotic results for large networks are also of interest to understand the scaling laws of these synchronization techniques.

REFERENCES

[1] M. Maggs, S. O'Keefe, and D. Thiel, "Consensus clock synchronization for wireless sensor networks," *IEEE Sensors J.*, vol. 12, no. 6, pp. 2269–2277, Jun. 2012.

[2] R. Preuss and D.R. Brown III, "Two-way synchronization for coordinated multi-cell retrodirective downlink beamforming," *IEEE Trans. Signal Process.*, vol. 59, no. 11, pp. 5415–5427, Nov. 2011.

[3] D. Mills, "Internet time synchronization: the network time protocol," *IEEE Trans. Commun.*, vol. 39, no. 10, pp. 1482–1493, Oct. 1991.

[4] *IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems*, IEEE Std. 1588, 2008.

[5] W. Lewandowski, J. Azoubib, and W. Klepczynski, "GPS: primary tool for time transfer," *Proc. IEEE*, vol. 87, no. 1, pp. 163–172, Jan. 1999.

[6] J. Hahn and E. Powers, "Implementation of the GPS to Galileo time offset (GGTO)," in *Proc. of the 2005 Joint IEEE International Frequency Control Symposium and Precise Time and Time Interval (PTTI) Systems and Applications Meeting*, Aug. 2005, pp. 33–212.

[7] S. Ganeriwal, R. Kumar, and M. Srivastava, "Timing-sync protocol for sensor networks," in *Proceedings ACM SenSys 2003*. ACM New York, NY, USA, Nov. 2003, pp. 138–149.

[8] F. Sivrikaya and B. Yener, "Time synchronization in sensor networks: a survey," *IEEE Netw.*, vol. 18, no. 4, pp. 45–50, Jul.–Aug. 2004.

[9] J. van Greunen and J. Rabaey, "Lightweight time synchronization for sensor networks," in *Proc. 2nd ACM International Workshop on Wireless Sensor Networks and Applications (WSNA 2003)*, Sep. 2003.

[10] M. Sichitiu and C. Veerarittiphan, "Simple, accurate time synchronization for wireless sensor networks," in *Proc. IEEE Wireless Communications and Networking (WCNC'03)*, vol. 2, Mar. 2003, pp. 1266–1273.

[11] M. Maroti, B. Kusy, G. Simon, and A. Ledeczi, "The flooding time synchronization protocol," in *Proc. 2nd Intl. Conf. on Embedded Networked Sensor Systems*, Nov. 2004.

[12] W. Su and I. Akyildiz, "Time-diffusion synchronization protocol for wireless sensor networks," *IEEE/ACM Trans. Netw.*, vol. 13, no. 2, pp. 384–397, Apr. 2005.

[13] Q. Li and D. Rus, "Global clock synchronization in sensor networks," *IEEE Trans. Comput.*, vol. 55, no. 2, pp. 214–226, Feb. 2006.

[14] R. Solis, V. Borkar, and P. Kumar, "A new distributed time synchronization protocol for multihop wireless networks," in *IEEE Conf. on Decision and Control*, Dec. 2006, pp. 2734–2739.

[15] L. Schenato and G. Gamba, "A distributed consensus protocol for clock synchronization in wireless sensor network," in *IEEE Conf. on Decision and Control*, Dec. 2007, pp. 2289–2294.

[16] R. Mirollo and S. Strogatz, "Synchronization of pulse-coupled biological oscillators," *SIAM J. on Appl. Math.*, pp. 1645–1662, 1990.

[17] D. Lucarelli and I. Wang, "Decentralized synchronization protocols with nearest neighbor communication," in *Proceedings of the 2nd international conference on Embedded networked sensor systems (ACM SenSys'04)*, November 3-5 2004, pp. 62–68.

[18] Y.-W. Hong and A. Scaglione, "A scalable synchronization protocol for large scale sensor networks and its applications," *IEEE Journal on Selected Areas in Communications*, vol. 23, no. 5, pp. 1085–1099, May 2005.

[19] G. Werner-Allen, G. Tewari, A. Patel, M. Welsh, and R. Nagpal, "Firefly-inspired sensor network synchronicity with realistic radio effects," in *Proceedings of the 3rd international conference on Embedded networked sensor systems (ACM SenSys'05)*, November 2-4 2005, pp. 142–153.

[20] A.-S. Hu and S. Servetto, "On the scalability of cooperative time synchronization in pulse-connected networks," *IEEE Trans. Inf. Theory*, vol. 52, no. 6, pp. 2725–2748, Jun. 2006.

[21] D.R. Brown III, G. Prince, and J. McNeill, "A method for carrier frequency and phase synchronization of two autonomous cooperative transmitters," in *Proceedings of the 6th IEEE Signal Processing Advances in Wireless Communications*, New York, NY, June 5-8 2005, pp. 260–264.

[22] I. Ozil and D.R. Brown III, "Time-slotted round-trip carrier synchronization," in *Proceedings of the 41st Asilomar Conference on Signals, Systems, and Computers*, Pacific Grove, CA, November 4-7 2007, pp. 1781–1785.

[23] D.R. Brown III and H.V. Poor, "Time-slotted round-trip carrier synchronization for distributed beamforming," *IEEE Trans. on Signal Processing*, vol. 56, no. 11, pp. 5630–5643, November 2008.

[24] R. Preuss and D.R. Brown III, "Retrodirective distributed transmit beamforming with two-way source synchronization," in *Proceedings of the Conference of Information Sciences and Systems (CISS 2010)*, Princeton, NJ, March 2010, pp. 1–6.

[25] G. Fettweis, E. Zimmermann, V. Jungnickel, and E. Jorswieck, "Challenges in future short range wireless systems," *IEEE Veh. Technol. Mag.*, vol. 1, no. 2, pp. 24–31, Jun. 2006.

[26] D. Parish, F. Farzaneh, and C. Barrat, "Methods and apparatus for calibrating radio frequency base stations using antenna arrays," U.S. Patent 6,037,898, Oct., 1997.

[27] A. Bourdoux, B. Come, and N. Khaled, "Non-reciprocal transceivers in ofdm/sdma systems: impact and mitigation," in *Proc. Radio and Wireless Conference (RAWCON'03)*, Aug. 2003, pp. 183–186.

[28] M. Guillaud, D. Slock, and R. Knopp, "A practical method for wireless channel reciprocity exploitation through relative calibration," in *Proc. Intl. Symp. on Signal Processing and Its Applications*, vol. 1, Aug. 2005, pp. 403–406.

[29] L. Galleani, "A tutorial on the 2-state model of the atomic clock noise," *Metrologia*, vol. 45, no. 6, pp. S175–S182, Dec. 2008.

[30] S. M. Kay, *Fundamentals of Statistical Signal Processing Volume I: Estimation Theory*. Upper Saddle River, NJ: Prentice Hall, 1993.

[31] A. Weiss and E. Weinstein, "Fundamental limitations in passive time delay estimation–part i: Narrow-band systems," *Acoustics, Speech and IEEE Transactions on Signal Processing*, vol. 31, no. 2, pp. 472–486, Apr. 1983.

[32] E. Weinstein and A. Weiss, "Fundamental limitations in passive time-delay estimation–part ii: Wide-band systems," *Acoustics, Speech and IEEE Transactions on Signal Processing*, vol. 32, no. 5, pp. 1064–1078, Oct. 1984.

[33] P. Bidigare, U. Madhow, R. Mudumbai, and D. Scherber, "Attaining fundamental bounds on timing synchronization," in *Acoustics, Speech and Signal Processing (ICASSP), 2012 IEEE International Conference on*, Mar. 2012, pp. 5229–5232.

[34] P. Bidigare, S. Pruessing, D. Raeman, D. Scherber, U. Madhow, and R. Mudumbai, "Initial over-the-air performance assessment of ranging and clock synchronization using radio frequency signal exchange," in *IEEE Stat. Signal Processing Workshop (SSP)*, Aug. 2012, pp. 273–276.

[35] B. Anderson and J. Moore, *Optimal Filtering*. Dover, 2005.

[36] W. Klepczynski and P. Ward, "Frequency stability requirements for narrow band receivers," in *32nd Annual Precise Time and Time Interval Meeting*, November 2000.