

Clock Synchronization Over Networks: Identifiability of the Sawtooth Model

POL DEL AGUILA PLA ^{1,2} (Member, IEEE), LISSY PELLACO³ (Student Member, IEEE), SATYAM DWIVEDI⁴ (Member, IEEE), PETER HÄNDEL ³ (Senior Member, IEEE), AND JOAKIM JALDÉN ³ (Senior Member, IEEE)

¹Center for Biomedical Imaging, Lausanne 1015, Switzerland

²Biomedical Imaging Group, École Polytechnique Fédérale de Lausanne, Lausanne 1015, Switzerland

³Division of Information Science and Engineering, School of Electrical Engineering and Computer Science, KTH Royal Institute of Technology, Stockholm 11428, Sweden

⁴Ericsson Research, Stockholm 164 83, Sweden

CORRESPONDING AUTHOR: POL DEL AGUILA PLA (e-mail: pol.delaguilapla@epfl.ch)

This work was supported by the SRA ICT TNG project Privacy-preserved Internet Traffic Analytics (PITA).

This article has supplementary downloadable material available at <https://ieeexplore.ieee.org>, provided by the authors.

ABSTRACT In this paper, we analyze the two-node joint clock synchronization and ranging problem. We focus on the case of nodes that employ time-to-digital converters to determine the range between them precisely. This specific design choice leads to a sawtooth model for the captured signal, which has not been studied before from an estimation theoretic standpoint. In the study of this model, we recover the basic conclusion of a well-known article by Freris, Graham, and Kumar in clock synchronization. More importantly, we discover a surprising identifiability result on the sawtooth signal model: noise improves the theoretical condition of the estimation of the phase and offset parameters. To complete our study, we provide performance references for joint clock synchronization and ranging using the sawtooth signal model by presenting an exhaustive simulation study on basic estimation strategies under different realistic conditions. With our contributions in this paper, we enable further research in the estimation of sawtooth signal models and pave the path towards their industrial use for clock synchronization and ranging.

INDEX TERMS Clock synchronization, ranging, identifiability, sawtooth model, sensor networks, round-trip time (RTT).

I. INTRODUCTION

Clock synchronization across a deployed network is a pervasive and long-standing challenge [1]–[6]. Furthermore, new-generation technologies each require more accurate synchronization. To name a few, i) in cellular communications, synchronization between base stations through a backhaul channel is fundamental to maintain frame alignment and permit handover among neighboring cells, and has been identified as a crucial requirement for distributed beamforming, interference alignment, and user positioning [7], [8], ii) in radio-imaging technology [9], accurate clock synchronization between the sparse chips that form an array is critical, and, in active-sensing 3-dimensional cases [10], [11], it results in low-cost wide-aperture ultra-short ultra-wideband

(UWB) pulses, increasing both the angular and depth resolutions of the captured images, iii) in wireless sensor networks [12], [13], synchronization is critical to data-fusion, channel-sharing, coordinated scheduling [14], [15], and distributed control [2] and iv) in distributed database solutions that provide external consistency, clock synchronization accuracy regulates latency, throughput, and performance [5].

Consequently with this wide range of application, theoretical insights on the fundamental limitations of clock synchronization over networks are likely to incite radical innovations in a number of fields. In [16], Freris, Graham, and Kumar established the fundamental limitations of the clock synchronization problem in an idealized scenario. Particularly, given a network of nodes with noise-less affine clocks

and fixed unknown link delays that exchange time-stamped messages, [16] i) showed that clock synchronization was only possible if the link delays were known to be symmetric, and ii) characterized the uncertainty regions of the clock synchronization parameters under different hypotheses. In this paper, we analyze the same problem from a perspective that is closer to real implementation. In short, we analyze the two-node joint clock synchronization and ranging problem [17]–[19] with noisy round-trip time (RTT) measurements without time-stamps [20], for a node design originally proposed in [21] to improve ranging accuracy. The resulting analysis has several advantages. First, because protocols without time stamps require only minimal transmissions of very short pulses carrying no information, the resulting technology minimizes communication overhead, and is beneficial in applications in which the data-rates are critically needed for other uses [6, p. 29]. Second, because we consider hardware specifically tailored to ranging accuracy, we reveal how applications that require this accuracy, such as cooperative localization [22], positioning [23], and control [2], can harness the same hardware and protocols for synchronization. Third, the analysis is more realistic, because it takes into account the real-world stochasticity of the measurements. In particular, in our analysis of the problem we i) unveil the need for symmetric delays in RTT-based protocols, in a direct parallel to the discovery in [16], ii) find novel results on the identifiability of sawtooth signal models under diverse conditions, which are of interest by their own right to chaotic system analysis [24], [25] and control, and iii) provide performance references to guide practitioners in their use of this technology.

In summary, in this paper we first derive from basic principles a model for RTT measurements between two nodes equipped with time-to-digital converters (TDC) in a network with fixed, unknown link delays (Theorem 1). Then, we shift our focus towards an encompassing family of signal models, i.e., sawtooth signal models, when one considers different stochastic effects. In this context, we provide results on the identifiability of these models, both negative (Lemma 1) and positive (Theorem 2), under different noise conditions. Here, we obtain the surprising result that the presence of noise inside a non-linear model term makes said model identifiable. We then shift the focus again towards clock synchronization and ranging, and we provide a thorough and verifiable empirical evaluation of the basic estimation techniques we propose in [26] to exploit the sawtooth model (Figs. 8–11, implementation accessible in [27]). These empirical results, together with the approximated Cramér-Rao lower bounds derived in [26], are clear and simple performance references for clock synchronization and ranging using sawtooth models. Such performance references are of use to both engineers that use this technology and to researchers aiming to develop estimation techniques for sawtooth signal models.

A. NOTATION

Discrete random processes will be in uppercase letters and square brackets, such as $Y[n]$, while deterministic sequences,

e.g., realizations of said processes, will be lowercase with square brackets, i.e., $y[n]$. For both these sequences, the notation will be simplified by omitting the discrete time index when it can be established by context. Vector random variables will be bold uppercase letters, e.g., \mathbf{Y} , while deterministic vectors will be bold lowercase letters, e.g., $\mathbf{y} \in \mathbb{R}^N$. Functions, on the other hand, will be non-italics lower-case letters, e.g., the probability density function (PDF) of a vector random variable \mathbf{Y} on a parametric family with vector parameter θ will be $f_{\mathbf{Y}}(\mathbf{y}; \theta)$. Through the paper, we view the modulus equivalence as a function, i.e., we note the mapping $x \mapsto y \in [0, a) \mid (y = x) \bmod(a)$ as $\bmod_a : \mathbb{R} \rightarrow [0, a)$.

II. THE SAWTOOTH MODEL

In applications in which high ranging accuracy with low communication overhead is desired, a low-cost solution in terms of both complexity and power consumption is using node designs that include TDCs to measure RTTs [21]. Indeed, such sensors were successfully incorporated in a prototype system to aid firefighters by providing on-site infrastructure-free indoor positioning [23]. TDCs, however, induce an asymmetry between the rate at which nodes can measure time and the rate at which they can act upon their environment. This asymmetry generates an unexpected waveform in the sequence of RTT measurements over time. This phenomenon was first reported by [19], where a sawtooth model was proposed and empirically validated, and possible applications to clock synchronization over networks were identified. In this section, we reintroduce the design of [21] and derive the sawtooth model from a few simple assumptions.

Consider now the design of [21], described in Fig. 1. Here, each node or sensor \mathcal{N} has a processing unit, a transceiver and a TDC. With this design, a sensor can measure RTTs at the resolution of the TDC, usually in the order of ps, much finer than the period of the processing unit’s clock, usually in the order of tens of ns. Besides the clear advantage of this design for ranging through RTT measurements, this creates an interesting asymmetric behavior of the node as an agent and as a measuring device. As we will show below (Theorem 1), this asymmetry produces a sawtooth waveform in the measured RTTs that depends on the synchronization parameters, and, under reasonable assumptions, leads to a viable system for clock synchronization over networks. As a consequence, such a design may also be considered for wired networks, where ranging information is usually not relevant. A final by-product of the inclusion of a TDC in the design in Fig. 1 is that we can consider that each node has perfect knowledge of its own clock period, which it measures directly with its TDC.

A. DETERMINISTIC MODEL

Consider two nodes, \mathcal{M} and \mathcal{S} , designed as \mathcal{N} in Fig. 1, in a network (wired or wireless). These two nodes execute the RTT measurement scheme illustrated in Fig. 2. In this scheme, \mathcal{M} measures the RTT between itself and \mathcal{S} by sending pulses (a.k.a. pings) to \mathcal{S} and using its TDC to accurately record when a response (a.k.a. pong) is received from \mathcal{S} . In particular,

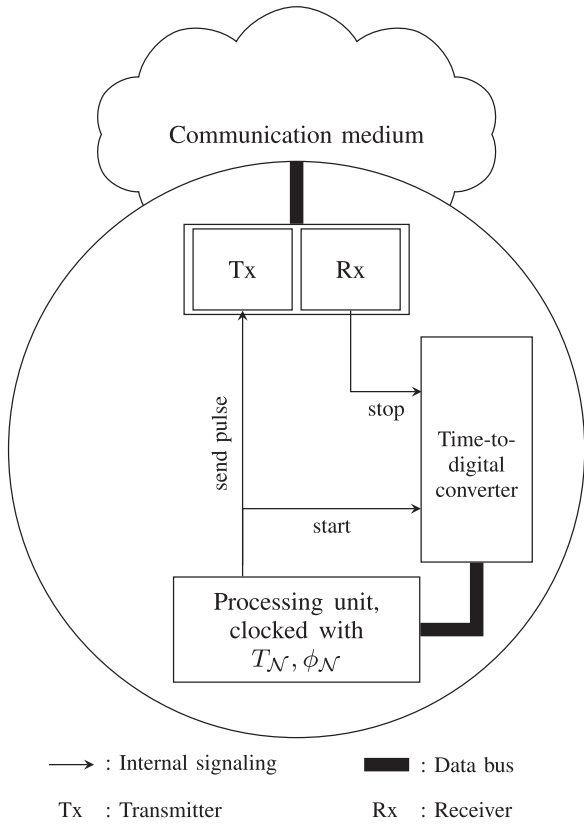


FIGURE 1. Internal design of any node \mathcal{N} considered throughout this paper. Initially proposed in [21], this design uses an independent time-to-digital converter (TDC) to accurately measure round-trip times (RTT). The resulting node \mathcal{N} can measure RTTs with a much finer time-resolution than it can react to incoming pulses. Indeed, in order to react to an incoming pulse, \mathcal{N} must first access the TDC's memory, process the reading, and decide to send a pulse, all of which require waiting until its next clock cycle.

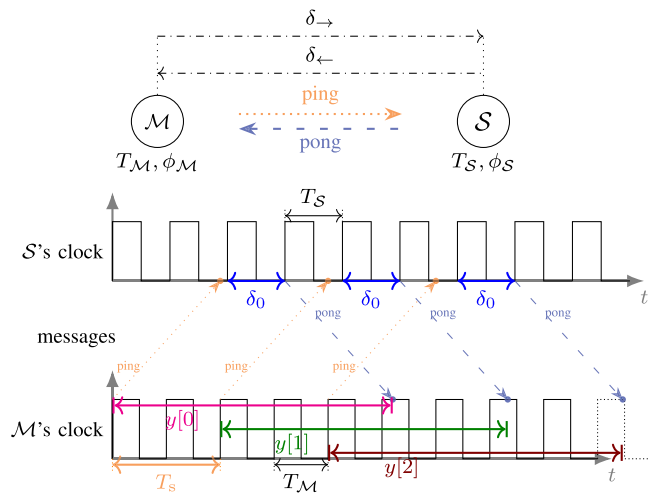


FIGURE 2. Example of the round-trip time (RTT) measurement scheme in Section II-A. \mathcal{M} sends ping pulses to \mathcal{S} at its clock upflanks every $T_s = 2T_M$. The ping is recorded in \mathcal{S} 's time-to-digital converter (TDC). At its next clock upflank, \mathcal{S} accesses its TDC and starts a delay of $\delta_0 = T_S$ before responding with a pong pulse. The pong is recorded in \mathcal{M} 's TDC as soon as it arrives.

\mathcal{M} sends a pulse at some of the times at which its clock has upflanks, i.e., at the times $t_n = KT_M n$. Here $T_M > 0$ [s] is \mathcal{M} 's clock period, the sampling factor $K \in \mathbb{N}$ is designed to determine the sampling period $T_s = KT_M$ [s], $n \in \mathbb{N}$ is a discrete-time index, and we assume without loss of generality that \mathcal{M} 's clock phase offset is zero, i.e., $\phi_M = 0$ [rad]. If we assume that the delays involved in the pulse traveling from \mathcal{M} to \mathcal{S} accumulate to a constant value δ_{\rightarrow} [s], the n -th pulse arrives at \mathcal{S} and is recorded in its TDC at time $t_n + \delta_{\rightarrow}$. Nonetheless, \mathcal{S} will not be able to access the TDC's memory before its next clock upflank, and consequently, any action by \mathcal{S} will be further delayed until $t_n + \delta_{\rightarrow} + \Delta_n$. Here, $\Delta_n \geq 0$ [s] is the time remaining until \mathcal{S} 's next clock upflank. If we consider that \mathcal{S} has a clock with period $T_S > 0$ [s] and phase ϕ_S [rad], i.e., an offset delay of $\varphi_S = T_S \phi_S / (2\pi)$ [s], then \mathcal{S} has its clock upflanks at those times that are at an integer number of periods away from φ_S , i.e., at the times $\tau \geq 0$ when $\text{mod}_{T_S}(\tau + \varphi_S) = 0$. Furthermore, we know that $\Delta_n \leq T_S$, as T_S is the time between consecutive upflanks. Consequently, to obtain a closed-form expression for Δ_n , we need to find

$$\Delta_n = \min \{ \tau \in (0, T_S) : \text{mod}_{T_S}(t_n + \delta_{\rightarrow} + \tau + \varphi_S) = 0 \}.$$

Because $\text{mod}_a(b + c) = \text{mod}_a(\text{mod}_a[b] + \text{mod}_a[c])$ for any $a \geq 0$ and $b, c \in \mathbb{R}$, the condition for $t_n + \delta_{\rightarrow} + \tau$ to be the time of one of \mathcal{S} 's clock upflanks can be rewritten as $\tau = QT_S - \text{mod}_{T_S}(t_n + \delta_{\rightarrow} + \varphi_S)$ for some $Q \in \mathbb{Z}$. Then, because $\text{mod}_{T_S}(t_n + \delta_{\rightarrow} + \varphi_S) < T_S$ and $\tau \in (0, T_S)$, we conclude that

$$\Delta_n = T_S - \text{mod}_{T_S}(t_n + \delta_{\rightarrow} + \varphi_S). \quad (1)$$

We now allow for a known delay δ_0 [s] to be introduced by \mathcal{S} , which can account for any processing required to read the TDC's state and prepare the new pulse, and will usually be an integer number of \mathcal{S} 's clock periods, i.e., $\delta_0 = K_0 T_S$. Finally, as we did for the ping pulse, we consider δ_{\leftarrow} to express the fixed delay for a pong pulse from \mathcal{S} to reach \mathcal{M} and be captured by the TDC. In conclusion, if we disregard the effect of the resolution of the TDC, which is usually four orders of magnitude finer than that of the nodes' clocks, the n -th RTT measurement will amount to

$$y_{\text{det}}[n] = \delta_0 + \delta_{\leftrightarrow} + \Delta_n, \quad (2)$$

where $\delta_{\leftrightarrow} = \delta_{\rightarrow} + \delta_{\leftarrow}$. We summarize our result in the following theorem.

Theorem 1 (Deterministic RTT measurement model): Consider two nodes \mathcal{M} and \mathcal{S} designed as specified in Fig. 1. Then, if \mathcal{M} and \mathcal{S} follow the RTT measurement protocol specified above, $f_d = 1/T_S - 1/T_M$ and $\delta_{\leftrightarrow} = \delta_{\rightarrow} + \delta_{\leftarrow}$, the n -th RTT measurement $y_{\text{det}}[n]$ can be expressed as

$$y_{\text{det}}[n] = \delta_{\leftrightarrow} + \delta_0 + T_S h[n], \quad \text{where} \quad h[n] = 1 - \text{mod}_1 \left(T_S f_d n + \frac{\delta_{\rightarrow}}{T_S} + \frac{\varphi_S}{2\pi} \right). \quad (3)$$




FIGURE 3. TDC measurement buffer, in units of the TDC’s clock, taken by \mathcal{M} and \mathcal{S} in a simulation of the protocol described by Fig. 2 and Section II-A. Unrealistic parameters were used to obtain a cheap-to-compute, simple, representative figure. For more details on how this simulation was performed see this project’s repository at [27].

Proof: From (1) we have that

$$\Delta_n = T_S - \text{mod}_{T_S}(T_S n + \delta_{\rightarrow} + \phi_S) \quad (4)$$

$$= T_S \left(1 - \text{mod}_1 \left(K \frac{T_M}{T_S} n + \frac{\delta_{\rightarrow}}{T_S} + \frac{\phi_S}{2\pi} \right) \right) \quad (5)$$

$$= T_S \left(1 - \text{mod}_1 \left(K \frac{T_M - T_S}{T_S} n + \frac{\delta_{\rightarrow}}{T_S} + \frac{\phi_S}{2\pi} \right) \right) \quad (6)$$

$$= T_S \left(1 - \text{mod}_1 \left(T_s f_d n + \frac{\delta_{\rightarrow}}{T_S} + \frac{\phi_S}{2\pi} \right) \right). \quad (7)$$

Here, we have used that $T_s = K T_M$ and $\phi_S = T_S \phi_S / (2\pi)$ in (5), that mod_1 is periodic with period one in (6), and that $f_d = 1/T_S - 1/T_M$ in (7). Finally, (3) follows from substituting (7) in (2). ■

In conclusion, by running the RTT protocol specified in Fig. 2, \mathcal{M} obtains data intimately related with the parameters it needs to predict \mathcal{S} ’s clock signal, i.e., to synchronize to \mathcal{S} . Indeed, having measured its own clock period T_M using its TDC, f_d reveals T_S , which together with ϕ_S characterizes \mathcal{S} ’s clock signal completely. The goal of our study is to establish under which conditions \mathcal{M} will be able to simultaneously estimate these parameters.

Other formulations of the model (3) in terms of the usual synchronization parameters for affine clocks, i.e. the clock skew $\alpha_S = T_S/T_M$ and the offset delay ϕ_S , including the general expression for when $\phi_M \neq 0$, can be found in the supplementary material to this paper. Nonetheless, the expression in (3) remains the most practical, because it expresses the compromise between the sampling period T_s and the frequency difference f_d of the system, which will prove to be relevant to our analysis.

Incidentally, under the simple assumption that $T_s \geq \delta_0 + T_S$, which can be guaranteed under any reasonable f_d if $K > K_0 + 1$, if we assume that \mathcal{S} ’s TDC starts measuring every

time \mathcal{S} sends a pong and stops measuring when the next ping is received, the n -th measurement $x[n]$ taken by \mathcal{S} ’s TDC can be expressed as

$$x_{\text{det}}[n] = T_s - \delta_0 - T_s h[n]. \quad (8)$$

As we will see, this will imply that even while \mathcal{M} is leading the RTT measurement protocol, \mathcal{S} could still perform frequency synchronization. Nonetheless, we will not consider \mathcal{S} ’s TDC measurements for most of the paper, and we will instead focus on determining the conditions under which \mathcal{M} can achieve full synchronization and ranging. More details on the derivations of (3) and (8) can be found in the supplementary material to this paper.

In this project’s repository, accessible at [27], we validate (3) by simulating an ideal physical system as described above and verifying the exact correspondence between the model and the obtained measurements. In Fig. 3, we show the fits of (3) and (8) on the TDC measurements of \mathcal{M} and \mathcal{S} throughout a simulated run with noisy clock periods and noisy transmission delays.

B. STOCHASTIC MODEL

In Fig. 4, we show real RTT data obtained in [19] from the ultra-wide band testbed of [21] using this RTT scheme, accompanied by an example model fit. Given the observed signal and its expected shape, a simple observation is that \mathcal{S} ’s clock was faster than that of \mathcal{M} in the specific experimental set-up, because the ramps in the sawtooth signal have negative slope, which implies that $f_d > 0$. Fig. 4 also exemplifies the two distinct effects that random deviations of the physical parameters can produce on the data. On one hand, large jumps of approximately T_S in the measured RTT are observed (effect i) if a random deviation influences the specific clock period at which \mathcal{S} reads the arrival of a ping pulse from its TDC, i.e., it changes which is the first up-flank in \mathcal{S} ’s clock after




FIGURE 4. RTT measurements from the ultra-wide band testbed from [21], compared to a fit of the deterministic model (3). In dash-dotted horizontal lines, the maximum and minimum values of the model.

the ping pulse arrives. On the other hand, if this does not happen, random deviations appear directly in the signal as additive noise (effect ii). From a modeling perspective, these two effects are not easily represented distinctively. Indeed, variations of the transmission time from \mathcal{M} to \mathcal{S} , δ_{\rightarrow} , or jitter in any of the two clock periods, $T_{\mathcal{M}}$ or $T_{\mathcal{S}}$, could lead to any of the two described effects, while variations of the transmission time from \mathcal{S} to \mathcal{M} , δ_{\leftarrow} , can only ever lead to effect ii]. In this paper, we will consider the effect of random variations on the physical parameters, as well as the quantization by the TDC, in the form of two additive white noise processes $W[n]$ and $V[n]$, inside and outside the nonlinearity, respectively. In short, our stochastic model for the RTT measurements taken by \mathcal{M} is

$$Y[n] = \delta_{\leftrightarrow} + \delta_0 + W[n] + T_{\mathcal{S}}H[n], \quad \text{where}$$

$$H[n] = 1 - \text{mod}_1 \left(T_{\mathcal{S}}f_d n + \frac{\delta_{\rightarrow}}{T_{\mathcal{S}}} + \frac{\phi_{\mathcal{S}}}{2\pi} + V[n] \right). \quad (9)$$

For simplicity, we will assume that $W[n]$ and $V[n]$ are zero-mean Gaussian processes with respective standard deviations σ_w and σ_v and we will consider them independent. Analyzing the effect of the existing dependence between them, or evaluating the magnitude or effect of this dependence, is outside of the scope of this paper. Only adding $V[n]$, i.e., setting $\sigma_w = 0$, could explain the two effects explained above for most sample indices n . Nonetheless, as shown by the indicators of the maximum and minimum of the model fit in Fig. 4, the experimental RTT measurements are not bounded in the range $[\delta_{\leftrightarrow} + \delta_0, \delta_{\leftrightarrow} + \delta_0 + T_{\mathcal{S}}]$, indicating that the noise term outside the non-linearity $W[n]$ is necessary. Furthermore, random variations in δ_{\leftarrow} or δ_0 cannot be meaningfully represented by $V[n]$, since variations of these parameters of any magnitude will never affect which upflank of \mathcal{S} detects the ping pulse. Fig. 5 exemplifies the effect of each of the noise terms by




FIGURE 5. Two realizations of our stochastic model (9), exemplifying the effects of additive noise inside and outside $\text{mod}_1(\cdot)$. On one hand, $\sigma_w = 0$ and $\sigma_v > 0$ leads to many high jumps around the wrapping points. On the other hand, $\sigma_w > 0$ and $\sigma_v = 0$ leads to a signal that is not bounded by the minimum and maximum values of the deterministic model (3) (shown in dash-dotted horizontal lines). For further examples of the effects of noise in measurements following our stochastic model (9), as well as the effects of randomness in the physical quantities described above, see this project's GitHub repository at [27].

showing two realizations of our stochastic model (9), one in which $\sigma_w = 0$ and $\sigma_v > 0$, and one in which $\sigma_w > 0$ and $\sigma_v = 0$.

In order to simplify the notation for the rest of the paper and abstract some of our theoretical results, we will express the stochastic sawtooth model in terms of four generic parameters, an offset $\alpha \in \mathbb{R}$, a non-zero amplitude $\psi \in \mathbb{R} \setminus \{0\}$ with known sign $\text{sign}(\psi)$, a normalized frequency $\beta \in [-1/2, 1/2)$, and a normalized phase offset $\gamma \in [0, 1)$. In other words, we will express the sawtooth signal model as

$$Y[n] = \alpha + W[n] + \psi \text{mod}_1(\beta n + \gamma + V[n]), \quad (10)$$

with $W[n]$ and $V[n]$ independent additive white Gaussian noise processes. An empirical analysis verifying this model (10) on real data from the testbed of [21] can be found in [19]. Here, the restriction of the parameters β and γ simply reflects the maximum ranges that we can expect to distinguish, given the periodicity of $\text{mod}_1(\cdot)$ as a function. Indeed, adding any integer factor of n inside the modulus, or any integer by itself, will not change $Y[n]$, and establishes an equivalence of period one for both β and γ . Here, we have chosen $\beta \in [-1/2, 1/2)$ and $\gamma \in [0, 1)$ to preserve their intuitive meanings as a normalized frequency and a phase term, respectively. Several initial insights can be drawn from the parallel between (9) and (10). First, the condition $|f_d| < 1/(2T_{\mathcal{S}})$, resembling the Nyquist sampling condition, arises from the restriction in β .

Second, if we consider this restriction and examine the relation between the parameters of both models, we observe that

$$\alpha = \delta_0 + \delta_{\leftrightarrow} + T_S, \psi = -T_S, \beta = f_d T_s, \quad \text{and}$$

$$\gamma = \text{mod}_1 \left(\frac{\delta_{\rightarrow}}{T_S} + \frac{\phi_S}{2\pi} \right),$$

and, incorporating that $T_S = T_M/(T_M f_d + 1)$ and that T_M, T_s and δ_0 are known,

$$\delta_{\leftrightarrow} = \alpha - \delta_0 - \frac{T_M}{T_M f_d + 1}, \quad (11)$$

$$f_d = \frac{\beta}{T_s} = - \left(\frac{1}{T_M} + \frac{1}{\psi} \right), \quad \text{and} \quad (12)$$

$$\phi_S = 2\pi \text{mod}_1 \left(\gamma - \text{mod}_1 \left(\delta_{\rightarrow} \frac{T_M f_d + 1}{T_M} \right) \right). \quad (13)$$

Clearly, then, unless further constraints relating $\delta_{\rightarrow}, \delta_{\leftarrow}$ and f_d are given, it is impossible to recover $\delta_{\rightarrow}, \delta_{\leftarrow}, f_d$, and ϕ_S from α, ψ, β and γ . In the context of clock synchronization over networks, this is equivalent to the impossibility result of [16], which studied the uncertainty sets where the synchronization parameters are known to lie given time-stamped message exchanges under different conditions. An analysis similar to that in [16] under idealized, noise-free conditions could be reproduced for (3), but is outside of the scope of this paper. In contrast, we will provide an analysis of identifiability when every physical parameter can be subject to noise. In fact, this analysis will reveal that synchronization with the sawtooth signal model requires a certain level of randomness, i.e., it is impossible without it. Consequently with the discussion above, then, we will assume that δ_{\rightarrow} is given when one knows δ_{\leftrightarrow} and f_d , as it happens in a number of applications. For example, in wireless sensor networks, one may generally consider that all nodes are equal and the channels between any two of them are symmetric, and thereby one can assume $\delta_{\rightarrow} = \delta_{\leftarrow} = \delta_{\leftrightarrow}/2 = \delta_1 + \rho/c$ where $\delta_1 > 0$ [s] is a known delay, $\rho > 0$ [m] is the unknown range between \mathcal{M} and \mathcal{S} in the communication medium and c [m/s] is the speed of light in the medium. Even in this context, true line-of-sight communication is not a requirement, and one only needs to assume that the multipath is not dense and the direct path is not fully blocked, as in such a case the TDC will trigger on the first pulse, corresponding to the shortest path. These assumptions are consistent with ultra-wide band pulses such as the ones used in [19]. When convenient in the paper, we will use this assumption combined with $\delta_1 \approx 0$, and consider the ranging problem of [21], [23] jointly with clock synchronization [19]. In the following section, we will characterize the model (10) statistically, providing conditions for its identifiability. Our aims in doing that are 1) to present novel results on the sawtooth signal model, and 2) to provide guarantees for practical synchronization systems that use nodes modeled by the design in Fig. 1.

III. IDENTIFIABILITY OF THE SAWTOOTH MODEL

Identifiability is a basic requirement on any statistical model that relates to the minimal conditions that make parameter estimation a reasonable goal [28].

Definition 1 (Identifiability): (From [29, Definition 11.2.2, p. 523]) Let \mathcal{Y}_θ be a statistical model with parameter $\theta \in \Omega$. Assume that if $\mathbf{Y} \sim \mathcal{Y}_\theta$ for some given θ , \mathbf{Y} has PDF $f_{\mathbf{Y}}(\mathbf{y}; \theta)$. Then, \mathcal{Y}_θ is an identifiable model, and θ is an identifiable parameter, if and only if

$$f_{\mathbf{Y}}(\mathbf{y}; \theta^{(1)}) = f_{\mathbf{Y}}(\mathbf{y}; \theta^{(2)}), \quad \forall \mathbf{y} \Leftrightarrow \theta^{(1)} = \theta^{(2)}. \quad (14)$$

That is, the mapping between the parameter θ and the distribution specified by \mathcal{Y}_θ is one-to-one.

If (14) is not met, the data observed when the parameter value is $\theta^{(1)}$ and the data observed when the parameter value is $\theta^{(2)}$ have the same distribution, and therefore, distinction between these two parameters from observed data is impossible. Unintuitively, even if (14) is not given, one could possibly design *good* estimators for θ . Specifically, as long as the selected metric in the space of parameters Ω does not assign much importance to the difference between the pairs $\theta^{(1)}$ and $\theta^{(2)}$ that do not fulfill (14), estimation could remain a sensible objective. In this paper, the data model \mathcal{Y}_θ is defined by (10), and the considered parameters are $\theta = [\alpha, \psi, \beta, \gamma]^T$. Hence, this section will be dedicated to establishing under which conditions, in terms of the values of σ_w and σ_v in (10), a one-to-one relation between θ and the distribution of the data $\mathbf{Y} = [Y[0], Y[1], \dots, Y[N-1]]^T$ can be ensured.

A. UNIDENTIFIABILITY WITHOUT INNER NOISE

In order to analyze the relation between θ and $f_{\mathbf{Y}}(\mathbf{y}; \theta)$, we will first consider the simplifying assumption $\sigma_v = 0$. This is an unrealistic assumption under most applicable uses of the sawtooth model (10), including that of clock synchronization, but it will be useful for our analysis. We will show that under this assumption, (10) yields an unidentifiable model \mathcal{Y}_θ in which the effect of α and γ cannot be fully distinguished in the observed data $\mathbf{Y} \sim \mathcal{Y}_\theta$.

Consider first (10) with $\sigma_v = 0$ and observe that then, $\mathcal{Y}_\theta = \mathcal{N}(\boldsymbol{\mu}_\theta, \sigma_w^2 \mathbf{I}_N)$, where \mathbf{I}_N is the $N \times N$ identity matrix and

$$\boldsymbol{\mu}_\theta = \alpha \mathbf{1}_N + \psi \text{mod}_1(\beta \mathbf{n} + \gamma \mathbf{1}_N), \quad (15)$$

with the modulus operation $\text{mod}_1(\cdot)$ applied component-wise, $\mathbf{1}_N = [1, 1, \dots, 1]^T \in \mathbb{R}^N$, and $\mathbf{n} = [0, 1, \dots, N-1]^T$. Because the normal distribution is fully characterized by its location and scale parameters, we know that changes in θ will only affect the distribution in terms of its location, controlled by its mean $\boldsymbol{\mu}_\theta$. Consequently, the condition for identifiability in Definition 1, i.e., (14), can be restated as $\boldsymbol{\mu}_{\theta^{(1)}} = \boldsymbol{\mu}_{\theta^{(2)}} \Leftrightarrow \theta^{(1)} = \theta^{(2)}$. Lemma 1 establishes that, when $\sigma_v = 0$, there are changes in α and γ that violate this condition.

Lemma 1 (Unidentifiability of (10) when $\sigma_v = 0$): Let \mathcal{Y}_θ express the model of the data $\mathbf{Y} = [Y[0], Y[1], \dots, Y[N-1]]^T$ given by (10) with $\theta = [\alpha, \psi, \beta, \gamma]^T$, $\alpha \in \mathbb{R}$, $|\psi| > 0$, $\beta \in [-1/2, 1/2]$ and $\gamma \in [0, 1)$, when the inner noise is disregarded, i.e., $\sigma_v = 0$. Then, \mathcal{Y}_θ is unidentifiable.




FIGURE 6. Examples of non-negative changes in the model parameters $\Delta\alpha \geq 0$ and $\Delta\gamma \geq 0$ (assuming $\psi > 0$) that lead to different (a) or the same (b) means $\mu_{\Delta\alpha}$ and $\mu_{\Delta\gamma}$ of the data according to model (10) with $\sigma_v = 0$, the mean values fully determine identifiability. For any change $\Delta\alpha > 0$ (if $\psi < 0$, $\Delta\alpha < 0$) such that $|\Delta\alpha| \leq \psi\epsilon_+$, there is a positive change in the phase $\Delta\gamma > 0$ that fulfills $\Delta\gamma \leq \epsilon_+$ and yields the same mean. The same is true vice versa.

In particular, there are different combinations of α and γ that yield the same distribution of \mathbf{Y} under \mathcal{Y}_θ .

Proof: We will show that $\mu_{\theta^{(1)}} = \mu_{\theta^{(2)}} \not\Rightarrow \theta^{(1)} = \theta^{(2)}$, i.e., the forward implication of (14) in Definition 1 is not fulfilled. In particular, given a parameter vector $\theta = [\alpha, \psi, \beta, \gamma]^T$, we will find $\theta^{(1)}, \theta^{(2)}$ such that $\theta^{(1)} \neq \theta^{(2)}$ and $\mu_{\theta^{(1)}} = \mu_{\theta^{(2)}}$.

Observe that, because $\text{mod}_1 : \mathbb{R} \rightarrow [0, 1)$,

$$\epsilon_+ = 1 - \max_{n \in \{0, 1, \dots, N-1\}} \{\text{mod}_1(\beta n + \gamma)\} > 0.$$

Then, $\forall n \in \{0, 1, \dots, N-1\}$ and $\forall \epsilon \in [0, \epsilon_+)$,

$$\alpha + \psi \text{mod}_1(\beta n + [\gamma + \epsilon]) = [\alpha + \psi\epsilon] + \psi \text{mod}_1(\beta n + \gamma).$$

Therefore, for any $\epsilon \in [0, \epsilon_+)$,

$$\theta^{(1)} = [\alpha + \psi\epsilon, \psi, \beta, \gamma]^T \text{ and } \theta^{(2)} = [\alpha, \psi, \beta, \gamma + \epsilon]^T,$$

yield $\mu_{\theta^{(1)}} = \mu_{\theta^{(2)}}$, i.e., $\mu_{\theta^{(1)}} = \mu_{\theta^{(2)}} \not\Rightarrow \theta^{(1)} = \theta^{(2)}$ and \mathcal{Y}_θ is unidentifiable. ■

Fig. 6 illustrates the idea of our proof of Lemma 1 when $\psi > 0$, and shows, in Fig. 6(b), changes in α and γ that cannot be distinguished in the mean of \mathbf{Y} for a simple example with

$N = 3$. Consequently, Fig. 6(b) serves as a straightforward counter-example to the identifiability of \mathcal{Y}_θ when $\sigma_v = 0$. Note that both the result in Lemma 1 and the counter-example of Fig. 6(b) are valid when $\sigma_w = 0$, i.e., when the model is deterministic as in (3). In contrast with our result in Theorem 2, this implies arbitrarily accurate synchronization could be impossible in an idealized noise-less scenario.

In our proof of Lemma 1, we exploit the formal definition of $\text{mod}_1(\cdot)$ to claim that its value will always be strictly less than one, and therefore, we obtain the margin ϵ_+ under which changes in α of the same sign as ψ and positive changes in γ are not distinguishable. However, the real limitation on identifiability is given by the points closest to the discontinuity from both sides, and, in most cases (i.e., $\gamma \neq 0$ and for most β s), a similar margin ϵ_- can be obtained under which changes in α of the sign opposite to ψ and negative changes in γ are not distinguishable. While our analysis is concerned with a fixed value of N , the lack of identifiability stated in Lemma 1 may be less problematic in an asymptotic regime. In particular, if increasing the sample size tends to reduce the segment at the left of the non-linearity without any sample, i.e., $\epsilon_+ \rightarrow 0$ when $N \rightarrow +\infty$, the size of the changes in α and γ that cannot be distinguished in the data would decrease with N , making the model identifiable in an asymptotic regime, or at least taking away importance from our proof of non-identifiability for large N s. In particular, if we consider the sequence of elements of the vector (15), μ_{θ_n} , we obtain what is known in dynamical systems as the orbit of a rotation of the circle. Then, if $\beta \in \mathbb{R} \setminus \mathbb{Q}$ we have that the orbit is minimal [30, ch. 1.3., proposition 1.3.3.], i.e., that the set $\{\mu_{\theta_n}\}_1^N$ when $N \rightarrow +\infty$ is dense in $[0, 1)$, and thus, $\epsilon_+ \rightarrow 0$. This contradicts the intuitive notion of finite-sample identifiability as a necessary condition for consistent estimators to exist, seen, for example, in [31, p. 62]. In contrast, when $\beta \in \mathbb{Q}$, (15) is periodic and hence $\epsilon_+ \rightarrow \epsilon_+^* > 0$. In particular, if $\beta = \pm M/Q$ with M and Q two co-prime naturals, then (15) is periodic with minimal period Q , and increasing the sample size beyond Q will not result in any further reduction of ϵ_+ , i.e., any further improvement from an identifiability perspective. In the case of clock synchronization, this corresponds to coherent sampling, in which $T_s f_d = \pm M/Q$. The effect of this specific case in the estimation error of a global grid search (GGS) strategy on the prediction mean square error of the model (3) (see [26]) when ψ and β are known is illustrated in Fig. 7.

Our analysis has assumed that α was part of the parameter vector θ , and that one wanted to recover it. Although this can be the prominent case in many applications of the sawtooth model, e.g., synchronization in wireless sensor networks or networks of autonomous vehicles, other applications may consider α to be known. Within synchronization, this would be the case of base-station synchronization in cellular networks, in which the backhaul links will most likely have a known and stable transmission delay. This would invalidate the identifiability analysis in Lemma 1, and under some additional conditions, \mathcal{Y}_θ could be shown to be identifiable. Regardless, in the next section we analyze the full model in the

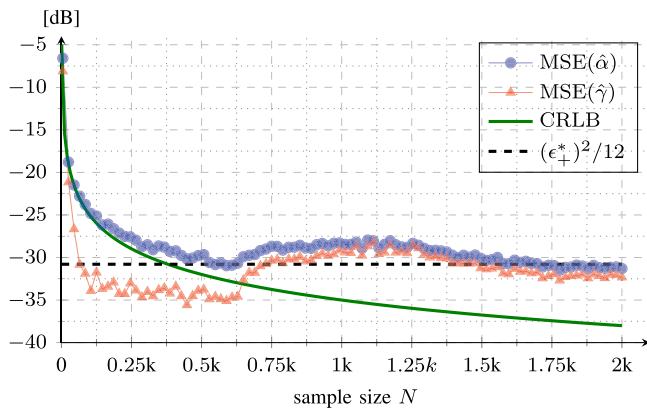


FIGURE 7. $MSE(\hat{\alpha}_{\text{GGS}})$ and $MSE(\hat{\gamma}_{\text{GGS}})$ against the sample size N when $\sigma_v = 0$, $\psi = 1$ and $\beta = M/Q$ with $M = 1$ and $Q = 10$. $\hat{\alpha}_{\text{GGS}}$ and $\hat{\gamma}_{\text{GGS}}$ are the result of a global grid search (GGS) on the prediction mean squared error with model (3) (see [26]) with 1000 discretization points for $\gamma \in [0, 1)$ when β and ψ are known. Results obtained from 2000 Monte Carlo repetitions and $\text{SNR}_{\text{out}} = 5$ dB (see Table 1). To access the fully reproducible code to generate this figure, see this project’s repository [27]. For comparison, the figure shows the Cramér-Rao lower bounds (CRLB) for estimating an offset in white noise, which here serves as a reference for the estimation of α and γ (see the supplementary material of this paper for details). Here, we observe that while the MSE for the estimate of α initially decays as predicted by the CRLB, it stabilizes around the variance of a uniform noise of width ϵ_+^* . Furthermore, the MSE for the estimate of γ becomes worse and stabilizes around the same value when the estimation of α reaches that level.

presence of noise inside the $\text{mod}_1(\cdot)$ non-linearity, i.e., with $\theta = [\alpha, \psi, \beta, \gamma]^T$ when $\sigma_v > 0$, and show its identifiability.

B. IDENTIFIABILITY WITH INNER NOISE

Theorem 2 (Identifiability of (10) when $\sigma_v > 0$): Let \mathcal{Y}_θ express the model of the data $\mathbf{Y} = [Y[0], Y[1], \dots, Y[N-1]]^T$ given by (10) when $\theta = [\alpha, \psi, \beta, \gamma]^T$, the parameters fulfill $\alpha \in \mathbb{R}$, $|\psi| > 0$ with $\text{sign}(\psi)$ known, $\beta \in [-1/2, 1/2)$ and $\gamma \in [0, 1)$, there is noise inside the $\text{mod}_1(\cdot)$ non-linearity, i.e., $\sigma_v > 0$ and $\sigma_w \geq 0$, and at least two RTT measurements have been taken, i.e., $N \geq 2$. Then, \mathcal{Y}_θ is an identifiable model and θ is an identifiable parameter.

Because the proof of Theorem 2 is rather technical, we place it in Appendix A. However, it is worthwhile to note here that it is not limited to the case in which $W[n]$ and $V[n]$ are Gaussian processes. Indeed, the statements in there apply *mutatis mutandis* under a wide variety of distributions for $W[n]$ and $V[n]$. In particular, any $W[n]$ consisting of independent and identically distributed (IID) samples from any location-scale family with some reference PDF $\varphi(w)$ with unbounded support will allow for the conclusion in (19). Furthermore, such a $W[n]$ together with any $V[n]$ consisting of IID samples from a location family that accepts a PDF and leads to a monomodal distribution after wrapping with mode equal to the location parameter, e.g., IID Cauchy distributed samples [32, p. 51], will also preserve all the statements therein. Nonetheless, to our knowledge, the literature mostly considers timing errors to be Gaussian (see, among others, [5],

[7], [8], [17], [18], [33], [34]), with little empirical incentive to consider other models.

The contrast between Lemma 1 and Theorem 2 is initially non-intuitive. Indeed, it implies that the presence of noise inside the non-linearity improves the theoretical condition of the estimation problem. This result recalls the popular theories of stochastic resonance for testing and estimation [35]–[39] and of dithering for improving the signal-independence of quantization noise [40], but is, in fact, less expected. In summary, both these theories delve into using noise to improve the performance of knowingly suboptimal strategies. In contrast, our identifiability result reveals how the inclusion of noise makes the data more informative with respect to the underlying parameters. To understand this result, one must first consider that it is very specific to sawtooth models, as it relies on that for any $x \in \mathbb{R} \setminus \mathbb{Z}$, there is an $\epsilon > 0$ such that $\text{mod}_1(x + \epsilon) = \epsilon + \text{mod}_1(x)$. In other words, small changes in phase and offset are equivalent around almost every x . In order to have any hope to distinguish them, one needs to guarantee that the points in which their effect differs play a role in shaping the distribution of the data. A phase noise $V[n]$ with long enough tails provides this guarantee, because changing the phase will alter the wrapping of the tails of the distribution through these non-linear points, while changing the offset will not.


IV. NUMERICAL RESULTS

In [26], we propose two estimators for the parameters of the sawtooth model (10). In particular, we first introduce an heuristic technique that is computationally light, intuitive, and surprisingly robust, which we call periodogram and correlation peaks (PCP). In this approach, one uses peaks in the power of the discrete Fourier transform to estimate the normalized frequency, and peaks in a circular correlation of the first estimated period to estimate the phase parameter. Using PCP as a starting point, we also introduce the local grid search (LGS), a computationally heavy method that improves the final performance by exploring a grid of normalized frequencies and phase parameters around the PCP estimate, and picks the combination that minimizes the model’s prediction mean square error, i.e., how well the model fits the observed signal. In both techniques, the offset parameter is chosen as the least squares estimate for each pair of frequency and phase estimates. For more detail on these techniques, and for the values of their parameters in the experimental results presented here, see [26], [27].


Here, we evaluate these estimators in a series of realistic simulation studies of clock synchronization and ranging. In the exposition of these results we intend to aid 1) engineers that aim to apply this technology, by providing reasonable expectations on its current possibilities, and 2) theoreticians that aim to develop estimators for the parameters of the sawtooth signal model (10), by revealing the strengths and pitfalls of the techniques that are currently available. With respect to 1), we include in all our results references that make it easier to identify different standard performance measures, regardless

TABLE 1. Default Values for the Physical and Simulation Parameters in our Numerical Results

Parameter	Interpretation	Value
N	sample size	2000
δ_0	delay introduced by \mathcal{S} [s]	$5 \cdot 10^{-6}$
$T_{\mathcal{M}}$	\mathcal{M} 's period [s]	10^{-8}
T_s	sampling period [s]	10^{-4}
ρ	range [m]	$\mathcal{U}[1, 3]$
f_d	frequency difference [Hz]	$\mathcal{U}[-200, 200]$
$\phi_{\mathcal{S}}$	\mathcal{S} 's phase [rad]	$\mathcal{U}[0, 2\pi]$
SNR_{in}	SNR for $V[n]$, $1/\sigma_v^2$ [dB]	40
SNR_{out}	SNR for $W[n]$, ψ^2/σ_w^2 [dB]	20


FIGURE 8. Result of 300 Monte Carlo repetitions in the conditions specified in Table 1 evaluating the MSE in the estimation of f_d by both PCP and LGS with respect to its actual value.

of the scale and aspect of each figure. In particular, a) in figures reporting ranging performance, i.e., $\text{MSE}(\hat{\rho})$, horizontal lines corresponding to standard errors of 1 cm or 0.1 cm are shown, b) in figures reporting frequency-difference estimation performance, i.e., $\text{MSE}(\hat{f}_d)$, horizontal lines corresponding to standard errors of 10 ppb (1 Hz) and 1 ppb (0.1 Hz) of \mathcal{M} 's frequency, $1/T_{\mathcal{M}} = 100$ MHz (see Table 1), are shown, and c) in figures reporting \mathcal{S} 's phase estimation performance, i.e., $\text{MSE}(\hat{\phi}_{\mathcal{S}})$, horizontal lines corresponding to standard errors in $\phi_{\mathcal{S}}$ that are one or two orders of magnitude below $T_{\mathcal{S}}$ are shown. In reference to c), we intentionally report performance on the estimation of the phase parameter $\phi_{\mathcal{S}}$ instead of the absolute time delay $\varphi_{\mathcal{S}} = T_{\mathcal{S}}\phi_{\mathcal{S}}/(2\pi)$. In our opinion, this is a better and fairer measure of how useful a specific clock synchronization technique is, because the scale of the errors in $\varphi_{\mathcal{S}}$ will always be dominated by the order of magnitude of $T_{\mathcal{S}}$. In other words, if $T_{\mathcal{S}} \approx 10$ ns, even guessing $\phi_{\mathcal{S}}$ at random between 0 and 2π achieves errors in $\varphi_{\mathcal{S}}$ that are on the order of ns. In order to streamline the exposition of this section and avoid unnecessary repetitions of the experimental conditions, we detail the default values for the physical and simulation parameters in Table 1. Note here that we chose to randomize the most critical parameters, so that the dependences shown in Figs. 8–11 can be considered the average behavior within the selected ranges.


FIGURE 9. Result of 300 Monte Carlo repetitions in the conditions specified in Table 1 evaluating the MSE in the estimation of $\phi_{\mathcal{S}}$ by both PCP and LGS with respect to its actual value.

A. SENSITIVITY TO THE PARAMETER VALUES

One of the weaknesses of the estimation approaches we present in [26] is low performance when f_d is small. In particular, because the PCP includes a mean-removal step before the DFT, the low frequencies are suppressed. If f_d is small, then, the peak we aim to detect in the DFT is most likely dampened and we detect noise instead. This is visualized in Fig. 8, which shows the average performance in the estimation of f_d by both PCP and LGS as a function of f_d , when all other parameters are randomized according to Table 1. Here, we see that the estimators fail when $f_d \in [-10, 10]$. Furthermore, we also observe the effect of the grid underlying the DFT on the PCP frequency estimate. In particular, for a given N , the PCP will only propose as estimates frequencies that are at one of the points in the DFT (e.g., -200 Hz and 200 Hz in Fig. 8), and so frequencies that are close to those will be better estimated than those that are far. This same phenomenon is observed in the experimental results we present in [26] for varying N and a fixed combination of physical parameters. Remarkably, the performance obtained with LGS remains below 1 ppb of $1/T_{\mathcal{M}}$ for most frequencies f_d outside the $[-10, 10]$ area.

A weakness of the measurement protocol described in Section II-A is that, due to the lack of time-stamps in the exchanged packets, the time origin shift between \mathcal{M} and \mathcal{S} 's clock appears only in terms of a phase term inside the $\text{mod}_1(\cdot)$ function, e.g., $\phi_{\mathcal{S}}$ if we assume that $\varphi_{\mathcal{M}} = 0$. If one desires absolute time synchronization, this can have dire consequences. While $\phi_{\mathcal{S}} = 2\pi(1 - \xi/2)$ and $\phi_{\mathcal{S}} = \pi\xi$ for some small $\xi > 0$ are only $2\pi\xi$ away under the non-linear wrapping, their difference implies errors in $\varphi_{\mathcal{S}}$ of $T_{\mathcal{S}}(1 - \xi)$. In order to take this effect into account, we use the conventional Euclidean distance to quantify the error for the phase term $\phi_{\mathcal{S}}$, without taking into account the wrapping effect. The ensuing increase of error around the extremes can be seen plainly in Fig. 9. This weakness is characteristic of protocols for clock synchronization without time-stamping. In turn, however, these protocols minimize the communication overhead and are more robust to malicious nodes [6, p. 29].




FIGURE 10. Result of 2000 Monte Carlo repetitions evaluating the MSE in the estimation of ρ , f_d and ϕ_S by both PCP and LGS with respect to the sample size. The range and phase terms ρ and ϕ_S were randomized as specified in Table 1, while $f_d \sim \mathcal{U}([-200, -10] \cup [10, 200])$, in order to avoid the wildly irregular behavior of the estimators when $f_d \in [-10, 10]$, shown in Fig. 8. For reference and comparison, the convergence rates given by the inverse Fisher information matrix we present in [26] are portrayed by lines with the corresponding slope adjusted to fit the data.

B. AVERAGE CONSISTENCY

An important property of estimators is consistency, i.e., the convergence of the MSE towards zero as the sample size increases. In order to characterize consistency beyond the example given in [26] with fixed physical parameters, in Fig. 10 we report the average performance for randomized parameter values for PCP and LGS in the estimation of the clock synchronization and ranging parameters as a function of the sample size. However, we sample frequency values in a reduced range, i.e., $f_d \in [-200, -10] \cup [10, 200]$, in order to avoid the instability of our estimators when $f_d \in [-10, 10]$. This provides an impression of the estimators' performance within their range of probable use, and significantly reduces

the amount of Monte Carlo repetitions necessary to obtain intelligible results.

Fig. 10 strengthens the conclusions from the analysis of the empirical results in [26]. Indeed, 1) the results suggest that both PCP and LGS are consistent for the estimation of ρ , f_d and ϕ_S , 2) LGS seems to have an asymptotically efficient rate of convergence with N for the estimation of both ρ and f_d , while the rate of convergence stagnates for ϕ_S . Here, we use as a reference the asymptotic rates of convergence for a linearized version of (9), which we derive from the Fisher information matrix in [26].

Finally, from a practical point of view, average estimation errors under 0.1 cm in the range parameter can be expected for $N \geq 1000$, and average estimation errors under 1 ppb in the frequency parameter can be expected for $N \geq 1500$. For the phase parameter ϕ_S , estimation errors below $2\pi/10$ can only be expected with LGS and for $N \geq 1000$.

C. SENSITIVITY TO THE INNER AND OUTER NOISES

All the results in Figs. 8 and 9 (as well as those in [26]) were obtained under the optimistic noise conditions $\text{SNR}_{\text{in}} = 40$ dB and $\text{SNR}_{\text{out}} = 20$ dB. In Fig. 11 we report the average performance for randomized parameter values for PCP and LGS in the estimation of the clock synchronization and ranging parameters as a function of either SNR_{in} or SNR_{out} , when $\text{SNR}_{\text{out}} = 30$ dB or $\text{SNR}_{\text{in}} = 40$ dB, respectively. Similarly as in Section IV-B, we sample f_d in the range $f_d \in [-200, -10] \cup [10, 200]$, in order to avoid the instability of our estimators when $f_d \in [-10, 10]$.

Fig. 11 reveals that range estimation performance decays progressively with the decrease of SNR_{out} throughout the investigated range for both PCP and LGS. In contrast, the decrease of SNR_{in} creates a progressive decay of performance only up to a breaking point around $\text{SNR}_{\text{in}} = 10$ dB. In this regime (very low SNR_{in}), one could consider modeling the non-linear term in (10) as a uniform noise term and employ techniques tailored to the estimation of offsets in Gaussian plus uniform noise [41]. Practically, both PCP and LGS achieve estimation accuracies below 0.1 cm only when $\text{SNR}_{\text{out}} \geq 10$ dB for $\text{SNR}_{\text{in}} = 40$ dB, and only when $\text{SNR}_{\text{in}} \geq 15$ dB for $\text{SNR}_{\text{out}} = 30$ dB.

For frequency estimation, we see that both PCP and LGS are very robust to a decrease of SNR_{out} up to a breaking point around $\text{SNR}_{\text{out}} = 10$ dB, and LGS maintains an improvement of 15 dB over PCP for any SNR_{out} above the breaking point. In contrast, when SNR_{in} degrades, the breaking point for both PCP and LGS is $\text{SNR}_{\text{in}} = 20$ dB, and the improvement of LGS over PCP increases gradually as SNR_{in} increases, only reaching 15 dB when $\text{SNR}_{\text{in}} = 40$ dB. Practically, on one hand, PCP achieves estimation accuracies below 10 ppb only when $\text{SNR}_{\text{out}} \geq 10$ dB for $\text{SNR}_{\text{in}} = 40$ dB, and only when $\text{SNR}_{\text{in}} \geq 20$ dB for $\text{SNR}_{\text{out}} = 30$ dB. On the other hand, LGS consistently improves on it, reaching estimation accuracies below 1 ppb only when $\text{SNR}_{\text{out}} \geq 10$ dB for $\text{SNR}_{\text{in}} = 40$ dB, and only when $\text{SNR}_{\text{in}} \geq 30$ dB for $\text{SNR}_{\text{out}} = 30$ dB. Promising directions for improving frequency estimation could come

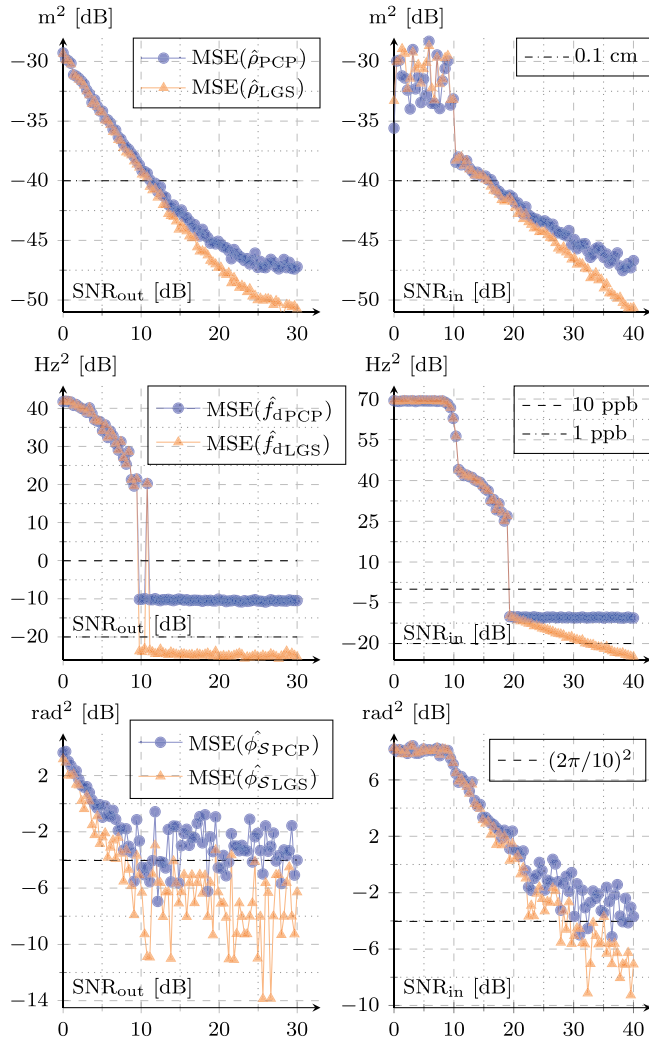


FIGURE 11. Result of 1000 Monte Carlo repetitions evaluating the MSE in the estimation of the clock synchronization and ranging parameters ρ , f_d , and ϕ_S , by both PCP and LGS with respect to the value of the SNRs. Each SNR was fixed to its maximum when the other was varied. The range and phase terms ρ and ϕ_S were randomized as specified in Table 1, while $f_d \sim \mathcal{U}([-200, -10] \cup [10, 200])$, in order to avoid the wildly irregular behavior of the estimators when $f_d \in [-10, 10]$, shown in Fig. 8.

from two different fronts. First, one could generalize the framework in [41] to admit frequency estimation, aiming to protect the resulting method for noises inside the non-linearity beyond the breaking points in SNR_{in} . Second, one could use the structure of the sawtooth signal in techniques similar to [42] to extract more information from the spectrum of the data.

For phase estimation, our results are not conclusive, because the randomization of both the frequency and phase parameters result in wide variability that would require further Monte Carlo averaging. Furthermore, while the unfavorable region of frequency parameters $f_d \in [-10, 10]$ has been avoided, the phase parameters are still sampled from the whole range $\phi_S \in [0, 2\pi)$, which includes the very challenging regions around the wrapping point (see Fig. 9). However,

the results seem to suggest a progressive decay of performance for both PCP and LGS when either SNR_{out} or SNR_{in} degrade. Furthermore, we see that phase estimation is much more sensitive to a degradation of SNR_{in} than that of SNR_{out} . Possible improvements of phase estimation could be expected using LGS-type techniques complimented with better frequency parameter estimates.

V. CONCLUSION

In this paper, we provide practical and theoretical insights on the fundamental limitations of clock synchronization over networks in applications that require high ranging accuracy and low communication overhead. From a practical standpoint, we show from first principles that using TDCs for measuring RTTs enables the use of a mathematical model that leads to very accurate ranging (e.g., accuracies beyond 0.1 cm for $N \geq 1000$ samples in Fig. 10) and remarkable frequency-synchronization performance (e.g., accuracies of 1 ppb for $N \geq 1500$ samples in the same figure), all with very simple estimation techniques. Furthermore, we point at promising directions of research that could hold the key to the improvement of these performance values, such as the extension of [41] to improve range estimation when SNR_{in} is very low (e.g., under 10 dB in Fig. 11), or the use of techniques similar to [42] to improve frequency synchronization. In fact, in [26] we provide a reference on the potential for improvement in the form of approximated Cramér-Rao lower bounds based on a linearization of the sawtooth model and standard results for Gaussian models. Additionally, we pinpoint the weaknesses of both the model and our estimation strategies. First, we acknowledge that more research is needed to consistently obtain accurate phase synchronization (phase estimation accuracies beyond $2\pi/100$). Second, we identify that schemes that rely on this mathematical model are not best suited for absolute time synchronization because the offset delay only appears as part of a phase term, which leads to wrapping errors.

From a theoretical standpoint, we establish that RTT-based schemes for clock synchronization are characterized by similar fundamental limitations as those relying on time-stamped message exchanges, previously discovered by [2]. Namely, synchronization is impossible with unknown path delays δ_{\rightarrow} and δ_{\leftarrow} if one cannot assume some relation between δ_{\rightarrow} , δ_{\leftarrow} and the frequency difference f_d , e.g., that they are symmetric, i.e., $\delta_{\rightarrow} = \delta_{\leftarrow}$. Furthermore, we have discovered a surprising property of sawtooth signal models, i.e., that adding noise inside the non-linearity allows for the joint identifiability of its offset and phase parameters (cf. Lemma 1 and Theorem 2). This result challenges the convention of random variations as a negative component of a model, and is of use beyond our application domain.

APPENDIX A PROOF OF IDENTIFIABILITY

Proof of Theorem 2, Identifiability of (10) when $\sigma_v > 0$: The backward implication of identifiability, i.e., that the same

parameters lead to the same data distribution (see Definition 1), is immediately clear from (10).

In the following, we will show that under the conditions above, the forward implication is also true, i.e., that

$$f_Y(\mathbf{y}; \boldsymbol{\theta}^{(1)}) - f_Y(\mathbf{y}; \boldsymbol{\theta}^{(2)}) = 0, \forall \mathbf{y} \in \mathbb{R}^N \quad (16)$$

implies that $\boldsymbol{\theta}^{(1)} = \boldsymbol{\theta}^{(2)}$, where $\boldsymbol{\theta}^{(1)}$ and $\boldsymbol{\theta}^{(2)}$ are two vectors of parameters, and we denote their respective components by the same super-index, i.e., $\alpha^{(i)}, \psi^{(i)}, \beta^{(i)}$, and $\gamma^{(i)}$ for $i \in \{1, 2\}$. Consequently, we assume (16), and will reach the conclusion that $\boldsymbol{\theta}^{(1)} = \boldsymbol{\theta}^{(2)}$. We start by defining the random process

$$P[n] = \alpha + \psi \bmod_1(\beta n + \gamma + V[n]), \quad (17)$$

and observe that (10) implies that $Y[n] = W[n] + P[n]$. Because $W[n]$ and $V[n]$ are assumed independent (see (10) and the discussion in Section II) and $W[n] \sim \mathcal{N}(0, \sigma_w^2)$, we have that, if $*$ denotes a convolution and $\varphi(\cdot)$ is the PDF of the standard normal distribution $\mathcal{N}(0, 1)$, the distribution of $Y[n]$ is, for any given $n \in \{0, 1, \dots, N-1\}$,

$$f_Y(\mathbf{y}; \boldsymbol{\theta}) = \left[\varphi \left(\frac{\tau}{\sigma_w} \right) * f_P(\tau; \boldsymbol{\theta}) \right](\mathbf{y}). \quad (18)$$

Then, consider the difference between the PDFs of $Y[n]$ that correspond to the parameter vectors $\boldsymbol{\theta}^{(1)}$ and $\boldsymbol{\theta}^{(2)}$, i.e.,

$$\begin{aligned} \Delta_Y(\mathbf{y}) &= f_Y(\mathbf{y}; \boldsymbol{\theta}^{(1)}) - f_Y(\mathbf{y}; \boldsymbol{\theta}^{(2)}) \\ &= \left[\varphi \left(\frac{\tau}{\sigma_w} \right) * \underbrace{(f_P(\tau; \boldsymbol{\theta}^{(1)}) - f_P(\tau; \boldsymbol{\theta}^{(2)}))}_{\Delta_P(\tau)} \right](\mathbf{y}), \end{aligned}$$

and observe that (16) implies that $\Delta_Y(\mathbf{y}) = 0, \forall \mathbf{y} \in \mathbb{R}$ and $\forall n \in \{0, 1, \dots, N-1\}$. Because $\varphi(\tau) > 0$ for any $\tau \in \mathbb{R}$, we have that $\Delta_Y(\mathbf{y}) = 0, \forall \mathbf{y} \in \mathbb{R}$ if and only if $\Delta_P(p) = 0, \forall p \in \mathbb{R}$. Because $P[n] \in [\alpha, \alpha + \psi]$ if $\psi > 0$ and $P[n] \in (\alpha + \psi, \alpha]$ if $\psi < 0$, this can only happen if $\alpha^{(1)} = \alpha^{(2)}$ and $\psi^{(1)} = \psi^{(2)}$, i.e.,

$$\Delta_Y(\mathbf{y}) = 0, \forall \mathbf{y} \in \mathbb{R} \Rightarrow \alpha^{(1)} = \alpha^{(2)} \text{ and } \psi^{(1)} = \psi^{(2)}. \quad (19)$$

Therefore, the ambiguity between α and γ that led to Lemma 1 has been resolved by the presence of the noise term $V[n]$ inside the non-linearity. This effect is exemplified by the simple example of Fig. 12, which studies the same exact situation as Fig. 6 (b), but with $\sigma_v > 0$.

Consider then the random process $Q[n]$ such that $Q[n] = \bmod_{2\pi}(2\pi\beta n + 2\pi\gamma + 2\pi V[n])$, and observe that

$$P[n] = \alpha + \frac{\psi}{2\pi} Q[n].$$

$Q[n]$ has a wrapped-normal distribution [32, p. 50] that is monomodal with mode $\bmod_{2\pi}(2\pi\beta n + 2\pi\gamma)$ and therefore, $P[n]$ is monomodal with mode $\alpha + \psi \bmod_1(\beta n + \gamma)$. Because $\Delta_Y(\mathbf{y}) = 0, \forall \mathbf{y} \in \mathbb{R}$ implies that $\Delta_P(p) = 0, \forall p \in \mathbb{R}$, we have that in particular, the mode of $P[n]$ under $\boldsymbol{\theta}^{(1)}$ and $\boldsymbol{\theta}^{(2)}$ must also be the same, which, because $\alpha^{(1)} = \alpha^{(2)}$ and $\psi^{(1)} = \psi^{(2)}$, implies that

$$\bmod_1(\beta^{(1)}n + \gamma^{(1)}) = \bmod_1(\beta^{(2)}n + \gamma^{(2)}). \quad (20)$$




FIGURE 12. Example on how the same changes of parameters $\Delta\alpha \geq 0$ and $\Delta\gamma \geq 0$ that exemplified in Fig. 6(b) that the model \mathcal{Y}_θ with $\sigma_v = 0$ was not identifiable yield different supports of $P[n]$, i.e., $\text{supp}_{\Delta\alpha}$ and $\text{supp}_{\Delta\gamma}$, when $\sigma_v > 0$. Accordingly, changes in α and γ , however small, will lead to different distributions of $Y[n]$.

Because $N \geq 2$, we know that (20) must be verified at least for $n \in \{0, 1\}$. For $n = 0$, because $\gamma^{(1)}, \gamma^{(2)} \in [0, 1]$, (20) implies that

$$\bmod_1(\gamma^{(1)}) = \bmod_1(\gamma^{(2)}) \Rightarrow \gamma^{(1)} = \gamma^{(2)}.$$

For $n = 1$, because $\gamma^{(1)} = \gamma^{(2)}$, (20) implies that there are $K_1, K_2 \in \mathbb{Z}$ with $Q = K_2 - K_1$ such that

$$\beta^{(1)} + K_1 = \beta^{(2)} + K_2 \Rightarrow \beta^{(1)} = Q + \beta^{(2)}.$$

Because $\beta^{(1)}, \beta^{(2)} \in [-\frac{1}{2}, \frac{1}{2}]$ this can only be fulfilled for $Q = 0$, and thus, $\beta^{(1)} = \beta^{(2)}$. In conclusion, (16) implies that $\alpha^{(1)} = \alpha^{(2)}$, $\psi^{(1)} = \psi^{(2)}$, $\gamma^{(1)} = \gamma^{(2)}$, and $\beta^{(1)} = \beta^{(2)}$, i.e., $\boldsymbol{\theta}^{(1)} = \boldsymbol{\theta}^{(2)}$. ■

ACKNOWLEDGMENT

The authors would like to thank Dr. H. Tullberg for referring them the existing results on circular statistics, which were fundamental to finalize the proof of Theorem 2. They also thank G. Farré for lengthy discussions on dynamical systems, the mod equivalence, and first directing them to the concept of the orbit of a rotation of the circle. Pol del Aguila Pla performed the work while at the KTH Royal Institute of Technology.

REFERENCES

- [1] L. Lamport, "Time, clocks, and the ordering of events in a distributed system," *Commun. Assoc. Comput. Machinery*, vol. 21, no. 7, pp. 558–565, Jul. 1978.
- [2] N. M. Freris, H. Kowshik, and P. R. Kumar, "Fundamentals of large sensor networks: Connectivity, capacity, clocks, and computation," *Proc. IEEE*, vol. 98, no. 11, pp. 1828–1846, Nov. 2010.
- [3] B. Etlzinger, H. Wymeersch, and A. Springer, "Cooperative synchronization in wireless networks," *IEEE Trans. Signal Process.*, vol. 62, no. 11, pp. 2837–2849, Jun. 2014.
- [4] W. Xia and M. Cao, "Determination of clock synchronization errors in distributed networks," *SIAM J. Control Optim.*, vol. 56, no. 2, pp. 610–632, 2018.

- [5] Y. Geng *et al.*, “Exploiting a natural network effect for scalable, fine-grained clock synchronization,” in *Proc. 15th USENIX Conf. Networked Syst. Des. Implementation*, 2018, pp. 81–94.
- [6] B. Etzlinger and H. Wymeersch, “Synchronization and localization in wireless networks,” *Found. Trends Signal Process.*, vol. 12, no. 1, pp. 1–106, 2018.
- [7] D. Zachariah, S. Dwivedi, P. Händel, and P. Stoica, “Scalable and passive wireless network clock synchronization in LoS environments,” *IEEE Trans. Wireless Commun.*, vol. 16, no. 6, pp. 3536–3546, Jun. 2017.
- [8] M. Koivisto *et al.*, “Joint device positioning and clock synchronization in 5 G ultra-dense networks,” *IEEE Trans. Wireless Commun.*, vol. 16, no. 5, pp. 2866–2881, May 2017.
- [9] Event Horizon Telescope Collaboration, “First M87 event horizon telescope results. I. The shadow of the supermassive black hole,” *Astrophysical J. Lett.*, vol. 875, no. 1, p. L1, 2019.
- [10] P. Chen and A. Babakhani, “3-D radar imaging based on a synthetic array of 30-GHz impulse radiators with on-chip antennas in 130-nm SiGe BiCMOS,” *IEEE Trans. Microw. Theory Techn.*, vol. 65, no. 11, pp. 4373–4384, Nov. 2017.
- [11] B. Jamali and A. Babakhani, “A self-mixing picosecond impulse receiver with an on-chip antenna for high-speed wireless clock synchronization,” *IEEE Trans. Microw. Theory Techn.*, vol. 66, no. 5, pp. 2313–2324, May 2018.
- [12] J. He, P. Cheng, L. Shi, J. Chen, and Y. Sun, “Time synchronization in WSNs: A maximum-value-based consensus approach,” *IEEE Trans. Autom. Control*, vol. 59, no. 3, pp. 660–675, Mar. 2014.
- [13] R. Carli and S. Zampieri, “Network clock synchronization based on the second-order linear consensus algorithm,” *IEEE Trans. Autom. Control*, vol. 59, no. 2, pp. 409–422, Feb. 2014.
- [14] K. S. Kim, S. Lee, and E. G. Lim, “Energy-efficient time synchronization based on asynchronous source clock frequency recovery and reverse two-way message exchanges in wireless sensor networks,” *IEEE Trans. Commun.*, vol. 65, no. 1, pp. 347–359, Jan. 2017.
- [15] S. Bolognani, R. Carli, E. Lovisari, and S. Zampieri, “A randomized linear algorithm for clock synchronization in multi-agent systems,” *IEEE Trans. Autom. Control*, vol. 61, no. 7, pp. 1711–1726, Jul. 2016.
- [16] N. M. Freris, S. R. Graham, and P. R. Kumar, “Fundamental limits on synchronizing clocks over networks,” *IEEE Trans. Autom. Control*, vol. 56, no. 6, pp. 1352–1364, Jun. 2011.
- [17] J. Zheng and Y.-C. Wu, “Joint time synchronization and localization of an unknown node in wireless sensor networks,” *IEEE Trans. Signal Process.*, vol. 58, no. 3, pp. 1309–1320, Mar. 2010.
- [18] S. P. Chepuri, R. T. Rajan, G. Leus, and A.-J. van der Veen, “Joint clock synchronization and ranging: Asymmetrical time-stamping and passive listening,” *IEEE Signal Process. Lett.*, vol. 20, no. 1, pp. 51–54, Jan. 2013.
- [19] S. Dwivedi, A. D. Angelis, D. Zachariah, and P. Händel, “Joint ranging and clock parameter estimation by wireless round trip time measurements,” *IEEE J. Sel. Areas Commun.*, vol. 33, no. 11, pp. 2379–2390, Nov. 2015.
- [20] M. R. Gholami, S. Dwivedi, M. Jansson, and P. Händel, “Ranging without time stamps exchanging,” in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, Apr. 2015, pp. 3981–3985.
- [21] A. D. Angelis, S. Dwivedi, and P. Händel, “Characterization of a flexible UWB sensor for indoor localization,” *IEEE Trans. Instrum. Meas.*, vol. 62, no. 5, pp. 905–913, May 2013.
- [22] N. Patwari, J. N. Ash, S. Kyperountas, A. O. Hero, R. L. Moses, and N. S. Correal, “Locating the nodes: Cooperative localization in wireless sensor networks,” *IEEE Signal Process. Mag.*, vol. 22, no. 4, pp. 54–69, Jul. 2005.
- [23] J.-O. Nilsson, J. Rantakokko, P. Händel, I. Skog, M. Ohlsson, and K. V. Hari, “Accurate indoor positioning of firefighters using dual foot-mounted inertial sensors and inter-agent ranging,” in *Proc. IEEE/ION Position, Location Navigation Symp.*, May 2014, pp. 631–636.
- [24] M. Ciftci and D. B. Williams, “Optimal estimation and sequential channel equalization algorithms for chaotic communications systems,” *EURASIP J. Appl. Signal Process.*, vol. 2001, no. 1, pp. 249–256, Jan. 2001.
- [25] D. F. Drake and D. B. Williams, “Linear, random representations of chaos,” *IEEE Trans. Signal Process.*, vol. 55, no. 4, pp. 1379–1389, Apr. 2007.
- [26] P. del Aguila Pla, L. Pellaco, S. Dwivedi, P. Händel, and J. Jaldén, “Clock synchronization over networks using sawtooth models,” in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.*, 2020.
- [27] P. del Aguila Pla and L. Pellaco, “Clock sync and range,” GitHub repository, 2018. [Online]. Available: https://github.com/poldap/clock_sync_and_range
- [28] E. Walter and L. Pronzato, *Identification of Parametric Models From Experimental Data*. ch. 2.6.1 Identifiability, Berlin, Germany: Springer-Verlag, 1997, pp. 20–32.
- [29] G. Casella and R. L. Berger, *Statistical Inference*. 2nd ed. Pacific Grove, CA, USA: Duxbury, 2002.
- [30] A. Katok and B. Hasselblatt, *Introduction to the Modern Theory of Dynamical Systems*. Cambridge, U.K.: Cambridge Univ. Press, 1995.
- [31] A. W. van der Vaart, *Asymptotic Statistics*. Cambridge, U.K.: Cambridge Univ. Press, 1998.
- [32] K. V. Mardia and P. E. Jupp, *Directional Statistics*. vol. 494, Hoboken, NJ, USA: Wiley, 2009.
- [33] K.-L. Noh, Q. M. Chaudhari, E. Serpedin, and B. W. Suter, “Novel clock phase offset and skew estimation using two-way timing message exchanges for wireless sensor networks,” *IEEE Trans. Commun.*, vol. 55, no. 4, pp. 766–777, Apr. 2007.
- [34] I. Skog and P. Händel, “Synchronization by two-way message exchanges: Cramér-Rao bounds, approximate maximum likelihood, and offshore submarine positioning,” *IEEE Trans. Signal Process.*, vol. 58, no. 4, pp. 2351–2362, Apr. 2010.
- [35] S. Kay, “Can detectability be improved by adding noise?” *IEEE Signal Process. Lett.*, vol. 7, no. 1, pp. 8–10, Jan. 2000.
- [36] H. Chen, P. K. Varshney, S. Kay, and J. H. Michels, “Theory of the stochastic resonance effect in signal detection: Part I—Fixed detectors,” *IEEE Trans. Signal Process.*, vol. 55, no. 7, pp. 3172–3184, Jul. 2007.
- [37] H. Chen and P. K. Varshney, “Theory of the stochastic resonance effect in signal detection—Part II: Variable detectors,” *IEEE Trans. Signal Process.*, vol. 56, no. 10, pp. 5031–5041, Oct. 2008.
- [38] S. Kay, “Noise enhanced detection as a special case of randomization,” *IEEE Signal Process. Lett.*, vol. 15, pp. 709–712, Nov. 2008.
- [39] H. Chen, L. R. Varshney, and P. K. Varshney, “Noise-enhanced information systems,” *Proc. IEEE*, vol. 102, no. 10, pp. 1607–1621, Oct. 2014.
- [40] L. Schuchman, “Dither signals and their effect on quantization noise,” *IEEE Trans. Commun. Technol.*, vol. 12, no. 4, pp. 162–165, Dec. 1964.
- [41] H. Lundin, M. Skoglund, and P. Händel, “On the estimation of quantizer reconstruction levels,” *IEEE Trans. Instrum. Meas.*, vol. 55, no. 6, pp. 2176–2182, Dec. 2006.
- [42] A. Camacho and J. G. Harris, “A sawtooth waveform inspired pitch estimator for speech and music,” *J. Acoustical Soc. Amer.*, vol. 124, no. 3, pp. 1638–1652, 2008.



POL DEL AGUILA PLA (Member, IEEE) received the double degree in telecommunications and electrical engineering from the Universitat Politècnica de Catalunya, Barcelona, and the KTH Royal Institute of Technology, Stockholm, Sweden, in 2014, and the Ph.D. degree in electrical engineering from the KTH Royal Institute of Technology in 2019. During his Ph.D., he investigated several inverse problems in signal processing applications such as biomedical imaging and clock synchronization over networks. Since October 2019, he has been a

Research Staff Scientist with the Center for Biomedical Imaging, Switzerland, and a Postdoctoral researcher with the EPFL’s Biomedical Imaging Group, Lausanne, Switzerland.

He is a Reviewer for the IEEE TRANSACTIONS ON SIGNAL PROCESSING, IEEE WIRELESS COMMUNICATIONS LETTERS, and *Elsevier Signal Processing*, and he is part of the local committee of ICASSP 2020, as well as a reviewer for ICML 2020.



LISSY PELLACO (Student Member, IEEE) received the M.Sc. degree (*summa cum laude*) in multimedia signal processing and telecommunication networks from the University of Genoa, Genoa, Italy, in 2016. She is currently working toward the Ph.D. degree in machine learning applied to radio networks with the KTH Royal Institute of Technology, Stockholm, Sweden under the supervision of Joakim Jaldén. She was at an international assignment on Industrial Automation Networks with ABB, Cleveland, OH, USA. She is

affiliated to the Wallenberg AI, Autonomous Systems and Software Program. In her Ph.D. program, she is investigating the merger of traditional parameterized and machine learning-based approaches to wireless communication, both satellite and terrestrial. She is also a reviewer for EUSIPCO 2019, and, since February 2019, she is part of the coordination group of the Female Ph.D. Student Network with the School of Electrical Engineering and Computer Science in KTH.

In December 2018, she was the recipient of the graduate award issued by the Italian Association of Electrical, Electronics, Automation, Information and Communication Technology to the most promising graduate in ICT Engineering. In April 2018, she was awarded the Excellence grant by the Executive Committee of the Doctoral Program council of the school of Electrical Engineering and Computer Science at KTH. Also, in 2013 she was the recipient of the ASING and IRIS INGEGNERIA awards issued by the University of Genoa to the most brilliant first-year B.Sc. students.



SATYAM DWIVEDI (Member, IEEE) received the M.S. and Ph.D. degrees from the Indian Institute of Science, Bangalore, India. He is a Senior Researcher with Ericsson Research, Stockholm, Sweden. He has been a Researcher and a Teacher with KTH, Stockholm. His research interests include wireless positioning, time synchronization, wireless propagation, and wireless testbeds.



PETER HÄNDEL (Senior Member, IEEE) received the M.Sc. degree in engineering physics and the Lic.Eng. and Ph.D. degrees in automatic control from the Department of Technology, Uppsala University, Uppsala, Sweden, in 1987, 1991, and 1993, respectively. He held a part-time position as an Associate Director of research with the Swedish Defense Research Agency from 2000 to 2006. In 2010, he joined the Indian Institute of Science, Bangalore, India, as a Guest Professor. He was a Guest Professor with the University of Gävle between 2007 and 2013. Since 1997, he has been with the KTH Royal Institute of Technology, Stockholm, Sweden, where he is currently a Professor of signal processing with the School of Electrical Engineering and Computer Science. He has authored more than 300 scientific publications. He was the recipient of a number of awards, including the Best Survey Paper Award in IEEE TRANSACTIONS ON INTELLIGENT TRANSPORTATION SYSTEMS. He is the former President of the IEEE Finland Joint Signal Processing and Circuits and Systems Chapter and the former President of the IEEE Sweden Signal Processing Chapter. He was an Associate Editor for the IEEE TRANSACTIONS ON SIGNAL PROCESSING.

He was an Associate Editor for the IEEE TRANSACTIONS ON SIGNAL PROCESSING.



JOAKIM JALDÉN (Senior Member, IEEE) received the M.Sc. and Ph.D. degrees in electrical engineering from the KTH Royal Institute of Technology, Stockholm, Sweden in 2002 and 2007, respectively. From July 2007 to June 2009, he held a post-doctoral research position with the Vienna University of Technology, Vienna, Austria. He also studied with Stanford University, CA, USA, from September 2000 to May 2002, and worked with ETH, Zürich, Switzerland, as a Visiting Researcher, from August to September, 2008. In July

2009, he returned to KTH, where he is currently a Professor of signal processing. His recent work includes work on signal processing for biomedical data analysis, including single-cell tracking for time-lapse microscopy and inverse diffusion for immunoassays. He was an Associate Editor for the IEEE COMMUNICATIONS LETTERS between 2009 and 2011, and an Associate Editor for the IEEE TRANSACTIONS ON SIGNAL PROCESSING between 2012 and 2016. He has been a member of the IEEE Signal Processing for Communications and Networking Technical Committee since 2013, where he serves as the Chair 2019–2020. Since 2016 he is also responsible for the five year B.Sc and M.Sc. Degree Program in Electrical Engineering at KTH.

For his work on MIMO communications, he has been awarded the IEEE Signal Processing Society's 2006 Young Author Best Paper Award, and the Best Student Conference Paper Award at IEEE ICASSP 2007. He was also the recipient of the Ingvar Carlsson Career Award issued in 2009 by the Swedish Foundation for Strategic Research.