

CLOCK SYNCHRONIZATION OVER NETWORKS USING SAWTOOTH MODELS

Pol del Aguila Pla[†], Lissy Pellaco[‡], Satyam Dwivedi^{††}, Peter Händel[‡], and Joakim Jaldén[‡]

[†] Biomedical Imaging Group, EPFL, Lausanne, Switzerland
[‡] Division of Information Science and Engineering, School of EECS
 KTH Royal Institute of Technology, Stockholm, Sweden
^{††} Ericsson Research, Stockholm, Sweden

ABSTRACT

Clock synchronization and ranging over a wireless network with low communication overhead is a challenging goal with tremendous impact. In this paper, we study the use of time-to-digital converters in wireless sensors, which provides clock synchronization and ranging at negligible communication overhead through a sawtooth signal model for round trip times. In particular, we derive Cramér-Rao lower bounds for a linearization of the sawtooth signal model, and we thoroughly evaluate simple estimation techniques by simulation, giving clear and concise performance references for this technology.

Index Terms— Clock synchronization, ranging, wireless sensor networks (WSN), round-trip time.

1. INTRODUCTION

Time-to-digital converters (TDC) are independently clocked, low-power, highly accurate time measurement devices. Incorporating TDCs in the design of wireless sensors provides very accurate ranging information from basic round trip time (RTT) measurement protocols [1]. Such a scheme has been used to devise reliable and cost-efficient systems for indoor localization [2]. A similar scheme, introduced in [3], uses an improved RTT protocol to address clock synchronization across a deployed network. This approach is extensively analyzed in [4], both practically and theoretically. Clock synchronization becomes possible due to the presence of two different clock speeds within each wireless sensor, i.e., that of the sensor and that of its TDC. The resulting RTT measurements follow a sawtooth signal model [3], which, under realistic assumptions, leads to the identifiability of the clock synchronization and ranging parameters [4]. In this paper, we provide performance references for the use of this technology, which will benefit both engineers that use it and researchers studying the estimation of sawtooth signal models.

Clock synchronization in wireless sensor networks has been studied extensively from a variety of perspectives [5–13]. Most studies focus on global synchronization performance through a network based on some form of time-stamped message exchange. Although some target specific objectives, e.g., a fast consensus across the network [6, 11] or energy efficiency [9, 10], communication overhead due to the arguably unnecessary exchange of time stamps is usually disregarded. However, several studies [3, 4, 8, 12, 14] have reported that two-way message exchanges without time stamps

have the potential to substantially lower communication overhead while still providing accurate synchronization. Our study provides performance references on the synchronization accuracy of two TDC-equipped sensors in a WSN that exchange messages without time stamps, reducing communication overhead but still obtaining remarkable performance in ranging and frequency synchronization (under 0.1 cm and 1 ppb of the clock frequency).

2. SAWTOOTH MODEL AND CRAMÉR-RAO LOWER BOUNDS

An empirical study run by our group [3] revealed that with a specific measurement protocol (see [3] and [4]), the RTTs $Y[n]$ measured between two sensors with TDCs, which we name \mathcal{M} and \mathcal{S} , follow the sawtooth signal model, i.e.,

$$Y[n] = \alpha + W[n] + \psi \bmod_1(\beta n + \gamma + V[n]), \quad (1)$$

where $W[n]$ and $V[n]$ are noise processes, which are assumed to be white, independent, zero-mean Gaussian processes with standard deviations σ_v and σ_w , respectively. Here, α , ψ , β and γ are the generic sawtooth model parameters, respectively, for offset, amplitude, normalized frequency, and phase. In [4], we show that under simple modeling assumptions, when \mathcal{M} measures RTTs to and from \mathcal{S} , it obtains

$$Y[n] = \delta_{\leftrightarrow} + \delta_0 + W[n] + T_S H[n], \text{ where} \quad (2)$$

$$H[n] = 1 - \bmod_1 \left[T_s f_d n + \frac{\delta_{\rightarrow}}{T_S} + \frac{\phi_S}{2\pi} + V[n] \right].$$

Here, δ_0 [s] is a known delay introduced by \mathcal{S} , $\delta_{\leftrightarrow} \approx 2\rho/c$ [s] is the transmission time of each message back and forth, which we assume to be the result of two identical delays, $\delta_{\rightarrow} = \delta_{\leftarrow}$, and where ρ [m] is the range between \mathcal{M} and \mathcal{S} and c [m/s] is the speed of light in the communication medium. Further, T_S [s] (unknown by \mathcal{M}) and $T_{\mathcal{M}}$ [s] (known by \mathcal{M} , measured through its TDC) are, respectively, the clock periods of \mathcal{S} and \mathcal{M} , while $f_d = 1/T_S - 1/T_{\mathcal{M}}$ [Hz] is the difference between their frequencies, and $T_s = KT_{\mathcal{M}}$ is the known time between two consecutive measurements. Finally, ϕ_S [rad] is the unknown phase of \mathcal{S} 's clock when $\phi_{\mathcal{M}} = 0$ rad is assumed.

In [4], we show that (2) is an identifiable model, i.e., that the distribution of the data contains enough information to singularly identify these parameters. Nonetheless, the likelihood function is not differentiable everywhere. This violates the assumptions of the Cramér-Rao lower bound (CRLB) for the mean square error (MSE) of unbiased estimators, hindering our objective of providing performance references for the estimation of the model's parameters. Instead, we analyze a linear model that results from assuming that

The first author performed the work while at the KTH Royal Institute of Technology. Contact authors, pol.delaguilapla@epfl.ch and pellaco@kth.se.

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

an oracle has removed the effect of the nonlinearity (phase unwrapping). The model then becomes

$$Z[n] = \delta_0 + \frac{\delta_{\leftrightarrow}}{2} + T_S \left(1 - \frac{\phi_S}{2\pi} \right) - T_S T_s f_d n + U[n], \quad (3)$$

with $U[n]$ a white Gaussian process such that $U[n] \sim \mathcal{N}(0, \sigma^2)$ with $\sigma^2 = \sigma_w^2 + T_S^2 \sigma_v^2$. The resulting model (3) is not without complications. First, ϕ_S and δ_{\leftrightarrow} are not jointly identifiable, because only their weighted sum affects the distribution of $Z[n]$. Second, the variance of the noise now depends on T_S , i.e., on f_d , one of the parameters to estimate. Therefore, we analyze first a general linear model with slope-dependent noise power, i.e., the model

$$\mathbf{Z} = [\mathbf{1}_N, \mathbf{n}] \boldsymbol{\omega} + \mathbf{U}, \text{ with } \mathbf{U} \sim \mathcal{N}(0, \sigma^2 \mathbf{I}_N) \quad (4)$$

with $\sigma^2 = \sigma_0^2 + (\sigma_1 + \tilde{\beta} \sigma_2)^2$, $\mathbf{Z} = [Z[0], Z[1], \dots, Z[N-1]]^T$, $\boldsymbol{\omega} = [\tilde{\alpha}, \tilde{\beta}]^T$, and where $\sigma_0 \geq 0$, $\sigma_1 \geq 0$ and $\sigma_2 \geq 0$ are known. This model is equivalent to (3) when $\tilde{\alpha} = \delta_0 + \delta_{\leftrightarrow}/2 + T_S(1 - \phi_S/2\pi)$, $\tilde{\beta} = -T_S T_s f_d$, $\sigma_0 = \sigma_w$, $\sigma_1 = T_M \sigma_v$, and $\sigma_2 = \sigma_v/K$. Here, recall that $K = T_s/T_M$. The advantages of (4) with respect to (3) are that i) it is an identifiable model, and ii) it can be analyzed using standard results for the Fisher information matrix of Gaussian models [15, ch. 3.9, p. 47]. Furthermore, given the Fisher information matrix \mathbf{I}_ω for (4), one can obtain CRLBs for f_d , for ϕ_S when δ_{\leftrightarrow} is known, and for δ_{\leftrightarrow} when ϕ_S is known, by using the CRLB on functions of vector parameters [16, corollary 5.23, p. 306], i.e.,

$$\text{MSE}(\hat{\mathbf{g}}(\mathbf{Z})) \geq \text{CRLB}_u(\mathbf{g}(\boldsymbol{\omega})) = (\nabla_{\boldsymbol{\omega}} \mathbf{g})^T \mathbf{I}_\omega^{-1} (\nabla_{\boldsymbol{\omega}} \mathbf{g}), \quad (5)$$

where $\hat{\mathbf{g}}(\mathbf{Z})$ is an unbiased estimator of $\mathbf{g}(\boldsymbol{\omega})$, a bounded function, and $\nabla_{\boldsymbol{\omega}} \mathbf{g}$ is its gradient. The derivation and statement of the inverse Fisher information matrix necessary for (5) can be found in Section 5. Then, from the relation between (3) and (4), one obtains that

$$\begin{aligned} f_d &= \mathbf{g}_{f_d}(\boldsymbol{\omega}) = -\frac{\tilde{\beta}}{T_M (KT_M + \tilde{\beta})}, \\ \phi_S &= \mathbf{g}_{\phi_S}(\boldsymbol{\omega}) = 2\pi + \frac{2\pi}{T_M + \frac{\tilde{\beta}}{K}} \left(\frac{\delta_{\leftrightarrow}}{2} + \delta_0 - \tilde{\alpha} \right), \text{ and} \\ \delta_{\leftrightarrow} &= \mathbf{g}_{\delta_{\leftrightarrow}}(\boldsymbol{\omega}) = 2 \left(\tilde{\alpha} - \delta_0 - \left(T_M + \frac{\tilde{\beta}}{K} \right) \left(1 - \frac{\phi_S}{2\pi} \right) \right). \end{aligned} \quad (6)$$

The expressions for ϕ_S or δ_{\leftrightarrow} assume that the respective other is known. This circumvents the joint identifiability problem stated above, but the resulting CRLBs will disregard that both parameters need to be estimated simultaneously. Nonetheless, our purpose in deriving these bounds is to use them as a plausible reference for the performance one can obtain using (2), for which we proved identifiability in [4]. In order to establish the CRLBs using (5) we obtain

$$\begin{aligned} \nabla_{\boldsymbol{\omega}} \mathbf{g}_{f_d}(\boldsymbol{\omega}) &= -\frac{1}{T_S^2 K} [0, 1]^T, \\ \nabla_{\boldsymbol{\omega}} \mathbf{g}_{\delta_{\leftrightarrow}}(\boldsymbol{\omega}) &= 2 \left[1, \frac{1}{K} \left(\frac{\phi_S}{2\pi} - 1 \right) \right]^T, \text{ and} \\ \nabla_{\boldsymbol{\omega}} \mathbf{g}_{\phi_S}(\boldsymbol{\omega}) &= \frac{-2\pi}{T_S} \left[1, \frac{1}{K} \left(\frac{\phi_S}{2\pi} - 1 \right) \right]^T. \end{aligned} \quad (7)$$

The obtained CRLBs are valid for unbiased estimators from data \mathbf{Z} generated according to (3), but they are not guaranteed to hold for unbiased estimators from data \mathbf{Y} generated from (2). Furthermore, they are not valid bounds on the MSE of biased estimators from either model. Nonetheless, we believe they provide a linear intuition that, as our experimental results confirm, is practically relevant.

3. BASIC ESTIMATION STRATEGIES

We present simple estimators for the parameters of a sawtooth signal model (1) based on the techniques proposed in [3]. In their simplicity, they show remarkable robustness for the ranges of parameters α , β , γ and ψ that arise in practical clock synchronization and ranging scenarios. Consequently, we consider them to be a good reference on the minimum expected performance that can be obtained from systems that use the proposed technology. For the sake of reproducibility and direct impact onto the interested communities, we provide thoroughly documented Jupyter notebooks that contain the implementation of all the presented techniques in this project's repository [17].

We expose our estimation methods in the more general notation of (1). However, we will consider that given β or ψ , the other is fully determined. This parallels clock synchronization, in which $\beta = T_s f_d$ and $\psi = -T_S = -T_M/(T_M f_d + 1)$. For practical application of these techniques to clock synchronization, it suffices to transform the estimators of α , β , γ and ψ to suitable estimators of ρ , f_d and ϕ_S through the comparison between (1) and (2) (for details, see [4]).

3.1. Periodogram and correlation peaks (PCP), a fast and simple solution

Deliberately developed to be computationally cheap, PCP uses only very simple and efficient operations such as discrete Fourier transforms (DFTs), sorting algorithms, and sample means. The estimator is divided in three steps, and relies on the assumption that the sign of the amplitude ψ is known. First, one uses a periodogram of the $L-1$ -times zero-padded centered data to estimate the absolute value of the frequency parameter β , i.e., $|\hat{\beta}| = \arg \max_{k \in \mathcal{K}} \{|\text{DFT}_{NL}(\tilde{y}[n])|^2\} / (NL)$ where $\tilde{y}[n]$ is a length NL signal such that

$$\tilde{y}[n] = \begin{cases} y[n] - \frac{1}{N} \sum_{m=0}^{N-1} y[m] & \text{if } n < N, \\ 0 & \text{if } N \leq n \leq NL - 1, \end{cases}$$

and $\mathcal{K} = \{0, 1, \dots, \lfloor NL/2 \rfloor\}$.

Second, one uses this unsigned frequency estimate to build two length $\lfloor 1/|\hat{\beta}| \rfloor$ signals $p_+[n]$ and $p_-[n]$ such that $p_{\pm}[n] = \text{sign}(\psi) \text{mod}_1(\pm|\hat{\beta}|n)$ for $0 \leq n < \lfloor 1/|\hat{\beta}| \rfloor$. These two reference signals and the first estimated period of the data, i.e., the length $\lfloor 1/|\hat{\beta}| \rfloor$ signal $\tilde{y}[n]$ such that $\tilde{y}[n] = y[n]$ for $0 \leq n < \lfloor 1/|\hat{\beta}| \rfloor$, are centered, max-normalized, and circularly correlated using length $\lfloor 1/|\hat{\beta}| \rfloor$ DFTs to estimate the sign of β and the value of γ . In particular, if $\tilde{y}[n]$, $\hat{p}_+[n]$, and $\hat{p}_-[n]$ are the centered and successively max-normalized signals, one computes two numbers l_+ and l_- as $l_{\pm} = \max_{0 \leq n < \lfloor 1/|\hat{\beta}| \rfloor} \{|\text{IDFT}[\text{DFT}(\hat{p}_{\pm}[n]) \text{DFT}(\tilde{y}[n])^*]|\}$, where $*$ represents complex conjugation. Here, one also stores at which indices $n_{\pm}^{\text{opt}} \in \{0, 1, \dots, \lfloor 1/|\hat{\beta}| \rfloor - 1\}$ the maxima l_{\pm} are achieved. Then, if $l_{\pm} > l_{\mp}$, one estimates $\hat{\beta} = \pm|\hat{\beta}|$ and $\hat{\gamma} = \text{mod}_1(\hat{\beta} n_{\pm}^{\text{opt}})$ with $n^{\text{opt}} = n_{\pm}^{\text{opt}}$, and the amplitude of the signal is considered estimated as $\hat{\psi}_{\hat{\beta}}$ through its relation with the frequency β .

Third, one employs the closed-form solution for the minimum mean square error estimator for the offset parameter α assuming that $\hat{\beta}$, $\hat{\gamma}$ and $\hat{\psi}_{\hat{\beta}}$ are correct, i.e.,

$$\hat{\alpha}_{\hat{\beta}, \hat{\gamma}} = \sum_{n=0}^{N-1} y[n] - \sum_{m=0}^{N-1} \hat{\psi}_{\hat{\beta}} \text{mod}_1[\hat{\beta} m + \hat{\gamma}]. \quad (8)$$

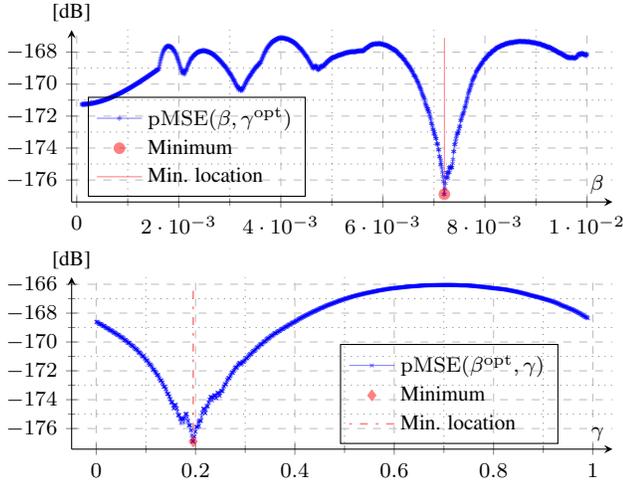


Fig. 1. Example of the prediction mean squared error (PMSE, (9)) of the model (1) obtained in a global grid search procedure (as described in Section 3.2) with a $10^3 \times 10^3$ grid with $\mathcal{B} = [0, 10^{-2}]$ and $\mathcal{G} = [0, 1)$, when either β or γ are fixed to their minimizing values β^{opt} and γ^{opt} . In this example, $N = 500$, $T_{\mathcal{M}} = 10$ ns, $T_s = 100 \mu\text{s}$, $\delta_0 = 5 \mu\text{s}$, $f_d = 73$ Hz, $\phi_S = \pi$ rad, and $\rho = 2$ m. For more details about the example and our implementation, as well as the image representation of the PMSE jointly over β and γ , see this project's repository [17].

Although this three-step estimator is heuristic, its computational cost is very low, and it can be implemented in lightweight hardware. Furthermore, while some of its steps are rather counter-intuitive, they show remarkable robustness. For example, using only the first estimated period of the data $\hat{y}[n]$ to estimate the phase parameter γ is clearly not an optimal strategy, but shows unparalleled robustness to errors in the estimation of the unsigned frequency parameter $|\beta|$, while steeply reducing the computational burden.

3.2. Local or global grid search (LGS or GGS), an exhaustive and costly solution

In contrast to PCP, the second technique we propose is computationally heavy. Nonetheless, our simulation study in Section 4 will suggest that it exhibits desirable statistical properties. In particular, we propose to minimize the prediction MSE (PMSE), i.e.,

$$\min_{(\beta, \gamma) \in \mathcal{G} \times \mathcal{B}} \left\{ \sum_{n=0}^{N-1} \left(y[n] - \hat{\alpha}_{\beta, \gamma} - \hat{\psi}_{\beta} \bmod_1[\beta n + \gamma] \right)^2 \right\}, \quad (9)$$

In (9), $\hat{\psi}_{\beta}$ is the implied estimator of ψ for a given β we mentioned at the start of Section 3, and $\hat{\alpha}_{\beta, \gamma}$ is the α that minimizes the cost function in (9), parametrized by β and γ and given *mutatis mutandis* by the expression in (8). Regrettably, the solution to (9) has to be approximated, because the PMSE over β and γ is neither convex nor unimodal, which implies that current iterative solvers are unable to find its global minimum efficiently. Example cuts of the profile of the PMSE over β and γ are reported in Fig. 1. We propose to approximately solve (9) by grid search, i.e., build a grid over some given ranges $\mathcal{G} \subset [0, 1)$ for γ and $\mathcal{B} \subset [-1/2, 1/2)$ for β and pick the parameters (β, γ) in the grid that yield the smallest value of the cost function. We call this technique either global grid search (GGS) when \mathcal{B} and \mathcal{G} contemplate all possible values, and local grid

Table 1. Values for the parameters of PCP, LGS and GGS throughout the paper, unless otherwise stated.

Parameter	Interpretation	Default value
L	zero-padding factor	5
\mathcal{B}_{LGS}	range for β in LGS	$\hat{\beta}_{\text{PCP}} + [-5, 5] \cdot 10^{-4}$
\mathcal{G}_{LGS}	range for γ in LGS	$\hat{\gamma}_{\text{PCP}} + [-28, 28] \cdot 10^{-3}$
$(N_{\mathcal{B}}, N_{\mathcal{G}})$	gridpoints for LGS	$(10^2, 10^3)$
\mathcal{B}_{GGS}	range for β in GGS	$[10^{-4}, 10^{-2}]$
\mathcal{G}_{GGS}	range for γ in GGS	$[0, 1)$
$(N_{\mathcal{B}}, N_{\mathcal{G}})$	gridpoints for GGS	$(10^3, 10^3)$

search (LGS) when they are defined as small neighborhoods around the PCP estimates. The performance of these methods will critically depend on the number and location of the grid points, which are design parameters that set the compromise between accuracy and computational complexity. The simplest distribution of these grid points is uniformly across $\mathcal{G} \times \mathcal{B}$, with $N_{\mathcal{G}}$ possible values for γ and $N_{\mathcal{B}}$ possible values for β .

4. EMPIRICAL RESULTS

In Fig. 2, we illustrate the convergence of the MSE for the PCP and LGS estimators proposed in Section 3 with the sample size N and compare it with the CRLBs for the unwrapped model derived in 2. The results we report were obtained from 300 Monte Carlo repetitions for specific physical parameters, i.e., $\delta_0 = 5 \mu\text{s}$, $T_{\mathcal{M}} = 10$ ns, $f_d = 73$ Hz, $T_s = 100 \mu\text{s}$, $\rho = 2$ m, and $\phi_S = 3\pi/4$ rad. Furthermore, the noise conditions were quite benign ($\text{SNR}_{\text{in}} = 1/\sigma_v^2 = 40$ dB and $\text{SNR}_{\text{out}} = \Psi^2/\sigma_w^2 = 20$ dB) and the algorithm's parameters were set as in Table 1. The results suggest that both estimators are consistent for these specific values of the parameters, in the sense that their overall error tends to decrease with increasing sample size, i.e., $\text{MSE} \rightarrow 0$ with $N \rightarrow +\infty$. This is coherent with the results we report in [4], where we evaluate these algorithms with randomized physical parameters and under varying noise conditions.

For PCP, the convergence of the MSE is clearly inefficient, and one observes it only by the decay of the envelope of the error. The regular bumps observed in the graphs of $\text{MSE}(\hat{f}_{d\text{PCP}})$ and $\text{MSE}(\hat{\phi}_{S\text{PCP}})$ are caused by the resolution of the underlying periodogram estimate. On one hand, if $\beta = k/NL$ for some $k \in \{0, 1, \dots, \lfloor NL/2 \rfloor\}$, β will be included in the periodogram's grid and the PCP will be biased towards it and achieve very low MSE. On the other hand, if β is between two such points, the PCP's bias will increase the MSE instead.

For LGS, the error seems to follow the decay of the CRLB of the unwrapped model in the estimation of ρ and f_d . However, the convergence of $\text{MSE}(\hat{\phi}_{S\text{LGS}})$ is much slower than that predicted by the CRLB of the unwrapped model. This is to be expected, since the bounds in (5) do not take into account the non-linearity of the model, and therefore, the wrapping effect of the phase term. Although this non-linear behavior is what makes the joint estimation of ϕ_S and ρ possible, it also makes ϕ_S much harder to estimate than a simple offset. Furthermore, one must consider that the MSE in the estimation of ϕ_S only plays a role when one aims to obtain time synchronization. If only phase synchronization is desired, however, consistency and efficiency may be defined using more appropriate evaluation metrics [18, p. 84]. The evaluation with respect to these metrics are outside the scope of this paper.

For both PCP and LGS, the error in the estimation of the range

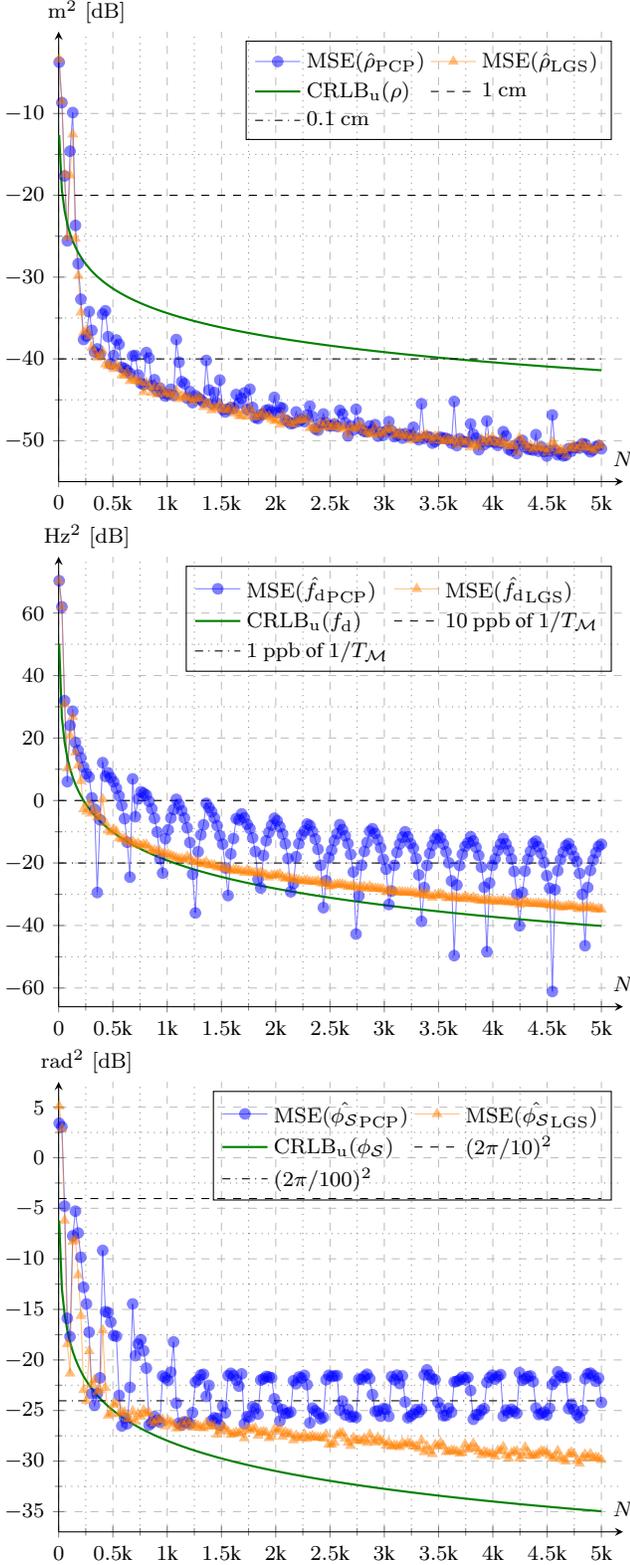


Fig. 2. Result of 300 Monte Carlo repetitions for the physical parameters specified in Section 4, evaluating the MSE in the estimation of ρ , f_d and ϕ_S by both PCP and LGS with respect to the sample size. For reference and comparison, we include the CRLBs for the unwrapped model derived in Section 2, and given by (5) and (7).

ρ is well below the CRLB, and for $N \geq 500$, it is mostly below 0.1 cm. Similarly, for PCP, $N \geq 500$ leads to average frequency estimation errors below 10 ppb of $1/T_M$ and average phase estimation errors well below $2\pi/10$. For LGS, $N \geq 500$ leads to average phase estimation errors below $2\pi/100$, and $N \geq 1500$ to frequency estimation errors of less than 1 ppb of $1/T_M$.

In conclusion, incorporating TDCs in wireless nodes to benefit from sawtooth modeling of RTT measurements is a promising strategy to simultaneously achieve remarkable ranging and frequency synchronization accuracy (errors under 0.1 cm and 1 ppb, respectively) and drastically decrease communication overhead. On the other hand, absolute time synchronization seems to be less suited to the sawtooth model, at least without more complex techniques (see the extended discussion we present in [4]).

5. APPENDIX: FISHER INFORMATION MATRIX FOR THE LINEAR MODEL WITH SLOPE-DEPENDENT NOISE POWER

Consider the model for \mathbf{Z} in (4) and recall that $\boldsymbol{\omega} = [\tilde{\alpha}, \tilde{\beta}]^T$. [15, ch. 3.9, p. 47] conveniently provides the expression for the Fisher information matrix of a generic Gaussian model in which $\mathbf{Z} \sim \mathcal{N}(\boldsymbol{\mu}_\omega, C_\omega)$ as

$$\mathbf{I}_\omega = \left[\begin{aligned} & \left[\frac{\partial}{\partial \omega_i} \boldsymbol{\mu}_\omega \right]^T C_\omega^{-1} \left[\frac{\partial}{\partial \omega_j} \boldsymbol{\mu}_\omega \right] + \\ & + \frac{1}{2} \text{Tr} \left[C_\omega^{-1} \frac{\partial C_\omega(\boldsymbol{\omega})}{\partial \omega_i} C_\omega^{-1} \frac{\partial C_\omega(\boldsymbol{\omega})}{\partial \omega_j} \right] \end{aligned} \right]_{i,j \in \{1,2\}}.$$

For (4), $\partial \boldsymbol{\mu}_\omega / \partial \tilde{\alpha} = \mathbf{1}_N$, $\partial \boldsymbol{\mu}_\omega / \partial \tilde{\beta} = \mathbf{n}$, $\partial C_\omega / \partial \tilde{\alpha} = 0 \mathbf{I}_N$ and $\partial C_\omega / \partial \tilde{\beta} = 2\sigma_2^2 (\sigma_1 + \tilde{\beta}\sigma_2) \mathbf{I}_N$. Considering that $\mathbf{1}_N^T \mathbf{1}_N = N$, $\mathbf{n}^T \mathbf{1}_N = N(N-1)/2$ and $\mathbf{n}^T \mathbf{n} = N(N-1)(2N-1)/6$, we obtain the Fisher information matrix for (4), i.e.,

$$\mathbf{I}_\omega = \frac{N}{\sigma^2} \begin{pmatrix} 1 & \frac{N-1}{2} \\ \frac{N-1}{2} & \frac{(N-1)(2N-1)}{6} + \frac{2\sigma_2^2(\sigma_1 + \tilde{\beta}\sigma_2)^2}{\sigma^2} \end{pmatrix}. \quad (10)$$

Inverting (10) leads to

$$\mathbf{I}_\omega^{-1} = \frac{\sigma^2/N}{\frac{N+1}{12} + \frac{2\sigma_2^2(\sigma_1 + \tilde{\beta}\sigma_2)^2}{\sigma^2(N-1)}} \begin{pmatrix} \frac{2N-1}{6} + \frac{2\sigma_2^2(\sigma_1 + \tilde{\beta}\sigma_2)^2}{\sigma^2(N-1)} & -\frac{1}{2} \\ -\frac{1}{2} & \frac{1}{N-1} \end{pmatrix} \quad (11)$$

which allows for the computation of the CRLBs for the estimation of $\tilde{\alpha}$ and $\tilde{\beta}$, and, through the relations (6) and their gradients (7), the CRLBs for the estimation of f_d , δ_{\leftrightarrow} when φ_S is known, and φ_S when δ_{\leftrightarrow} is known. In terms of the rates of convergence for the variance of efficient estimators, one can see that

$$\mathbf{I}_\omega^{-1} = \begin{pmatrix} \frac{\sigma^2}{\frac{N(N+1)}{2(2N-1)} + \mathcal{O}(N^{-1})} + \frac{2\sigma_2^2(\sigma_1 + \tilde{\beta}\sigma_2)^2}{N(N^2-1) + \mathcal{O}(1)} & -\frac{\sigma^2}{\frac{N(N+1)}{6} + \mathcal{O}(1)} \\ -\frac{\sigma^2}{\frac{N(N+1)}{6} + \mathcal{O}(1)} & \frac{\sigma^2}{\frac{N(N^2-1)}{12} + \mathcal{O}(N)} \end{pmatrix},$$

i.e., the efficient estimators of the offset $\tilde{\alpha}$ and the slope $\tilde{\beta}$ still have the same rates of convergence as in a standard linear model, with additions of only non-dominating terms.

6. REFERENCES

- [1] Alessio De Angelis, Satyam Dwivedi, and Peter Händel, “Characterization of a flexible UWB sensor for indoor localization,” *IEEE Transactions on Instrumentation and Measurement*, vol. 62, no. 5, pp. 905–913, May 2013.
- [2] J.-O. Nilsson, J. Rantakokko, Peter Händel, Isaac Skog, M. Ohlsson, and K. V. S. Hari, “Accurate indoor positioning of firefighters using dual foot-mounted inertial sensors and inter-agent ranging,” in *2014 IEEE/ION Position, Location and Navigation Symposium (PLANS 2014)*, May 2014, pp. 631–636.
- [3] Satyam Dwivedi, Alessio De Angelis, Dave Zachariah, and Peter Händel, “Joint ranging and clock parameter estimation by wireless round trip time measurements,” *IEEE Journal on Selected Areas in Communications*, vol. 33, no. 11, pp. 2379–2390, Nov. 2015.
- [4] Pol del Aguila Pla, Lissy Pellaco, Satyam Dwivedi, Peter Händel, and Joakim Jaldén, “Clock synchronization over networks — Identifiability of the sawtooth model,” Submitted to the IEEE Open Journal of Signal Processing, 2020.
- [5] Nikolaos M. Freris, Hemant Kowshik, and P. R. Kumar, “Fundamentals of large sensor networks: Connectivity, capacity, clocks, and computation,” *Proceedings of the IEEE*, vol. 98, no. 11, pp. 1828–1846, Nov. 2010.
- [6] Jianping He, Peng Cheng, Ling Shi, Jiming Chen, and Youxian Sun, “Time synchronization in WSNs: a maximum-value-based consensus approach,” *IEEE Transactions on Automatic Control*, vol. 59, no. 3, pp. 660–675, Mar. 2014.
- [7] Bernhard Etlzinger, Henk Wymeersch, and Andreas Springer, “Cooperative synchronization in wireless networks,” *IEEE Transactions on Signal Processing*, vol. 62, no. 11, pp. 2837–2849, June 2014.
- [8] Mohammad Reza Gholami, Satyam Dwivedi, Magnus Jansson, and Peter Händel, “Ranging without time stamps exchanging,” in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Apr. 2015, pp. 3981–3985.
- [9] Kyeong Soo Kim, Sanghyuk Lee, and Eng Gee Lim, “Energy-efficient time synchronization based on asynchronous source clock frequency recovery and reverse two-way message exchanges in wireless sensor networks,” *IEEE Transactions on Communications*, vol. 65, no. 1, pp. 347–359, Jan. 2017.
- [10] Dave Zachariah, Satyam Dwivedi, Peter Händel, and Petre Stolica, “Scalable and passive wireless network clock synchronization in LoS environments,” *IEEE Transactions on Wireless Communications*, vol. 16, no. 6, pp. 3536–3546, June 2017.
- [11] Jianping He, Xiaoming Duan, Peng Cheng, Ling Shi, and Lin Cai, “Accurate clock synchronization in wireless sensor networks with bounded noise,” *Automatica*, vol. 81, pp. 350–358, 2017.
- [12] Bernhard Etlzinger and Henk Wymeersch, “Synchronization and localization in wireless networks,” *Foundations and Trends® in Signal Processing*, vol. 12, no. 1, pp. 1–106, 2018.
- [13] Weiguo Xia and Ming Cao, “Determination of clock synchronization errors in distributed networks,” *SIAM Journal on Control and Optimization*, vol. 56, no. 2, pp. 610–632, 2018.
- [14] Dave Zachariah, Alessio De Angelis, Satyam Dwivedi, and Peter Händel, “Schedule-based sequential localization in asynchronous wireless networks,” *EURASIP Journal on Advances in Signal Processing*, vol. 16, pp. 1–12, 2014.
- [15] Steven Kay, *Fundamentals of statistical signal processing: Estimation theory*, Prentice-Hall, Inc., 1993.
- [16] Mark J. Schervish, *Theory of statistics*, Springer Science & Business Media, 1995.
- [17] Pol del Aguila Pla and Lissy Pellaco, “clock sync and range,” GitHub repository, https://github.com/poldap/clock_sync_and_range, 2018.
- [18] Kantilal Varich Mardia and Peter E. Jupp, *Directional statistics*, vol. 494, John Wiley & Sons, 2009.