

Estimating the underlying signal waveform, noise covariance and synchronization jitter from unsynchronized measurements

J Martinsson

EISLAB, Department of Computer Science and Electrical Engineering, Luleå University of Technology, SE-971 87 Luleå, Sweden

E-mail: Jesper.Martinsson@ltu.se

Received 17 October 2007, in final form 27 November 2007

Published 14 January 2008

Online at stacks.iop.org/MST/19/025101

Abstract

In this paper a new synchronization technique is presented for applications using repeated measurements or experiments with periodically excited signals. The objective with repeated or periodic measurements is often to retrieve an estimate of the (noise reduced) signal and its uncertainties. However, these measurements need to be synchronized to obtain accurate estimates. Existing synchronization techniques are limited to specific signal and noise conditions, such as white Gaussian noise or narrowband signals, to achieve good performance. The proposed method, not limited by these conditions, extracts statistical information regarding the underlying signal and the noise contained in the measurements, to obtain good synchronization (asymptotically optimal). The Cramér–Rao lower bound (CRLB) is derived for the synchronization problem, including bounds for the underlying signal waveform and the covariance of the measurement noise, both considered unknown. The method, which is shown to be the maximum-likelihood estimator (MLE) in both white and colored Gaussian noise, is compared with the CRLB along with standard sub-sample estimation and aligning techniques using Monte Carlo simulations. The results show significant mean square error (MSE) improvements compared to standard synchronization techniques. Synchronization results using the proposed technique are presented for repeated ultrasonic measurements, to validate the method in a real measurement situation, and to experimentally support theoretical results.

Keywords: synchronization, signal estimation, trigger, variations, trigger jitter, repeated measurements, periodic measurements, unsynchronized measurements

1. Introduction

In many measurement situations, and in particular ultrasonics, it is possible to conduct either repeated measurements of the same underlying signal, or to design experiments using periodic excitation. In both cases the objective is often to obtain repeated or periodic redundancy so that information regarding the underlying signal and the noise can be extracted from the data.

Given that the measurements are perfectly synchronized, an estimate of the underlying signal waveform is easily obtained by averaging over the repeated measurements or over

the measured periods. By subtracting this average from each measurement or each period, an estimate of the measurement noise and its statistical properties can be obtained.

However, perfect synchronization is often difficult to obtain in practice. In the case of repeated measurements, the problem is generated by trigger variations, resulting in a slight misalignment of the measurement window. In the case of periodic excitation, the problem occurs if the signal generation part is unsynchronized with the data collection part, creating a small drift between generation and sampling.

Assume that the retrieved data are a collection of M unsynchronized repeated measurements or periods of a signal.

The signal model can be written as

$$y_m[n] = x(n \cdot T_s - \tau_m) + w_m[n], \quad (1)$$

for $m = 1, 2, \dots, M$ repeated measurements or periods, and $n = 1, 2, \dots, N$ time samples. In equation (1),

- $y_m[n]$ represent the m th measurement,
- $x(n \cdot T_s)$ is the unknown underlying signal (bandlimited),
- T_s is the sampling time,
- τ_m is the unknown synchronization jitter,
- $w_m[n]$ is the m th measurement noise.

Consider the synchronization jitter, τ , to be a random variable, independent of the measurement noise, described by its probability density function (PDF) $p_\tau(\tau)$. The expected value of $y[n]$, with respect to τ and w , is

$$\begin{aligned} E_{\tau,w}\{y[n]\} &= E_\tau\{x(n \cdot T_s - \tau)\} + E_w\{w[n]\} \\ &= \int_{-\infty}^{\infty} x(n \cdot T_s - \tau) p_\tau(\tau) d\tau, \end{aligned} \quad (2)$$

assuming zero mean measurement noise. For unsynchronized measurements, the expected value (the usual estimate of the underlying signal) results in a convolution of the underlying signal with the PDF of the synchronization jitter, destroying the estimated signal's shape and the covariance estimate of the measurement noise. A more detailed study of these effects can be found in section 2.

Accurate estimates of the measured data's first- and second-order statistics are important and often required in many applications. The filtering effect described above usually results in low-pass filtered first moment (the underlying signal). This aggravates the possibility of retrieving information from high frequency regions and it also destroys second-order estimates. The second-order statistics (covariance matrices) are needed to optimally weight objective functions when fitting models or parameters to measured data, obtaining uncertainty bounds for estimated models and parameters, detecting model errors, and for experimental design [1].

The problem of compensating for synchronization jitter is often solved by estimating the sub-sample time delay with respect to one of the received signals, and then pre-aligning the whole set before estimating the underlying signal and the noise covariance. The aligning usually involves interpolation in the time-domain, and the sub-sample time delay estimation often involves some interpolation of the cross-correlation function [2]. Currently available methods assume either white noise or narrowband signals; see, e.g., [2–4].

In this paper, a different procedure is proposed¹ which takes advantage of the information contained in the entire set of unsynchronized measurements. Instead of using one noisy measurement as reference (under limited signal and noise conditions), the procedure simultaneously estimates the underlying signal waveform, the covariance matrix of the noise and the sub-sample synchronization jitters from the available

¹ An open source synchronization package for Matlab, containing the derived estimators, is available at <http://www.ltu.se/csee/research/eislab/areas/sensorsystems/toolboxes> (or upon request).

data. In addition to these, the proposed method also yields asymptotic estimates of their corresponding uncertainties. The proposed estimator is shown to be the maximum likelihood estimator (MLE) for the underlying signal, the covariance and the jitter, in white or colored Gaussian noise, independent of the signal bandwidth. This paper is an extension of previous work, presented in [5].

The estimator is derived and implemented using a frequency-domain approach, which means that time-domain interpolation and finite difference approximations of derivatives are avoided.

The Fisher information matrix (FIM) and the Cramér–Rao lower bound (CRLB) [6] are derived for the synchronization problem, including bounds on the underlying signal waveform, the covariance matrix, and the corresponding synchronization jitter. A closed-form expression for the CRLB of the synchronization jitter is found, providing a useful insight on how the shape of the signal waveform and how the strength and color of the noise influence the estimation result. A closed-form expression is also valuable from an optimization point of view, as it increases the stability and reduces the numerical complexity of the maximization algorithm.

The proposed estimator is compared with the CRLB along with standard sub-sample estimation and aligning techniques using Monte Carlo simulations, under different noise color, signal-to-noise ratio and data size conditions. The results show that the mean square error (MSE) of the estimates is considerably smaller using the proposed method compared to standard synchronization techniques, already for a fairly small number of repeated measurements.

The remainder of the paper is organized as follows: in section 2, a detailed study of the convolution effect, caused by unsynchronized measurements, is presented. In section 3, the log-likelihood function is given, the estimators are derived and optimization issues are addressed. In section 4, theoretical results are presented and interpreted. The estimator's MSE performance is evaluated in section 5. Estimation results from repeated ultrasonic measurements are presented in section 6. A discussion is found in section 7, covering aspects concerning experiments and signal model extensions. Finally, conclusions are drawn in section 8.

2. The effect of unsynchronized measurements

The importance of accurate estimates of the measured data's first- and second-order statistics was briefly discussed in section 1. The expected effect, if synchronization is ignored, was given by the convolution in (2).

An informative way to investigate how the convolution affects the estimates is to express it in the frequency-domain and visualize it for a given PDF. In the frequency-domain, the expected value and the covariance of the measurements are given by

$$E_{\tau,\tilde{w}}\{\tilde{y}(\omega)\} = \tilde{p}_\tau(\omega)\tilde{x}(\omega) \quad (3)$$

and

$$\begin{aligned} \text{cov}_{\tau,\tilde{w}}\{\tilde{y}(\omega), \tilde{y}(\xi)\} &= (\tilde{p}_\tau(\omega - \xi) - \tilde{p}_\tau(\omega)\tilde{p}_\tau(-\xi))\tilde{x}(\omega)\tilde{x}(-\xi) \\ &\quad + E_{\tilde{w}}\{\tilde{w}(\omega)\tilde{w}(-\xi)\}, \end{aligned} \quad (4)$$

where ω and ξ represent angular frequencies. A complete derivation of (3) and (4) can be found in appendix A.1. In this paper, the tilde notation $\tilde{\cdot}$ will be used to denote frequency-domain description.

The convolution in the time-domain (2) results in a multiplication (3) of the Fourier transformed signal $\tilde{x}(\omega)$ with the characteristic function $\tilde{p}_\tau(\omega)$ for the PDF (the Fourier transform of $p_\tau(\tau)$). The characteristic function of the jitter acts as a linear filter on the estimated signal. The covariance function in (4) is also affected by the characteristic function, where autocorrelation components of the signal $\tilde{x}(\omega)\tilde{x}(-\xi)$ leaks into the covariance function by the factor $(\tilde{p}_\tau(\omega-\xi) - \tilde{p}_\tau(\omega)\tilde{p}_\tau(-\xi))$. For synchronized measurements, the characteristic function is equal to one for all ω and no filtering effect is present, i.e. $p_\tau(\tau) = \delta(\tau) \leftrightarrow \tilde{p}_\tau(\omega) = 1$, where $\delta(\tau)$ represents the Dirac's delta function.

2.1. Jitter distributions

In the case of periodic excitation, the convolution occurs if there is a small drift between signal generation and sampling. This drift is normally lagging behind or sampling slightly faster than the signal generation part, meaning that each estimated period contains a bit more or a bit less than the generated period. For this type of cumulative drift, each τ_m is either greater or less than the previous one. A good model to describe this behavior is to represent τ as dependent and uniformly distributed over the total cumulative drift occurred during the measurement time.

In the case of repeated measurement, the convolution effect is obtained due to trigger variations. This can be modeled either as uniform or Gaussian distributed, depending on the trigger design and the measurement situation. Note that neither the convolution effect described here, nor the estimators derived in this paper, are affected by the dependence between each τ_m .

Unsynchronized measurements can also be generated by small non-stationary environmental effects, such as temperature or pressure fluctuations, during the measurement time. These effects are common in acoustic measurements, as the speed of sound (time-of-flight) is very sensitive to non-stationary fluctuations.

2.2. The convolution effect

In figure 1, the effect of using unsynchronized measurements is shown for different jitter distributions. The left column represents the effect if the jitter is uniformly distributed as $\tau \sim \mathcal{U}[-\alpha T_s, \alpha T_s]$, for the varying distribution intervals $\alpha = 1, 1/2, 1/4, 1/8$. The right column represents the effect if the jitter is normally distributed as $\tau \sim \mathcal{N}(0, (\alpha T_s)^2)$, for the varying standard deviations αT_s where $\alpha = 1, 1/2, 1/4, 1/8$.

In the first row, figures 1(a)–(b), we can see the different distributions of the synchronization jitter. Here $\tilde{\tau} = \tau/T_s$ is shown instead of τ to exclude the dependence on T_s .

The second row shows the characteristic functions for the uniform distributed jitter in figure 1(c), and for the Gaussian distributed jitter in figure 1(d). The curves show the suppression of high frequency information of the underlying

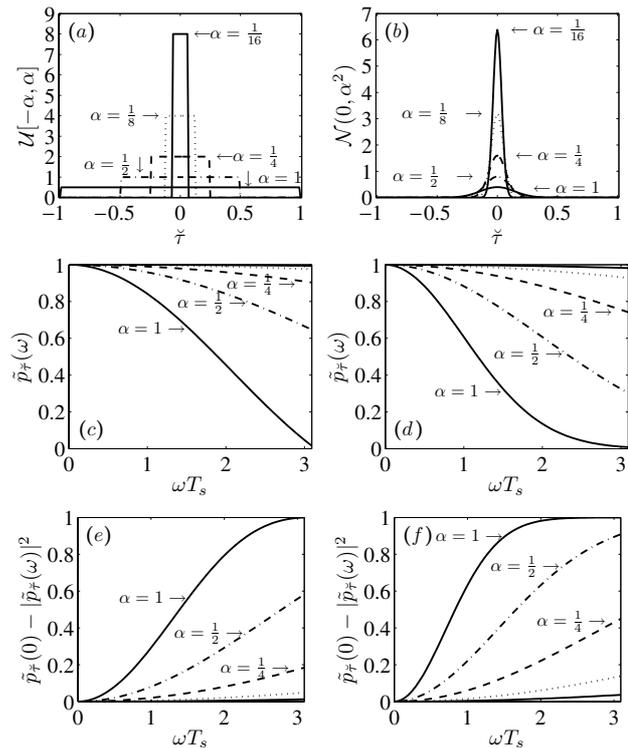


Figure 1. The effect of using unsynchronized measurements, where $\tilde{\tau} = \tau/T_s$ is represented to obtain independence of T_s : (a) uniform distribution, $\tilde{\tau} \sim \mathcal{U}[-\alpha, \alpha]$, for the varying intervals $\alpha = 1, 1/2, 1/4, 1/8$; (b) Gaussian distribution, $\tilde{\tau} \sim \mathcal{N}(0, \alpha^2)$, for the varying standard deviations $\alpha = 1, 1/2, 1/4, 1/8$; (c) characteristic function $\tilde{p}_\tau(\omega)$ for the uniform PDF, approaches one for small α ; (d) characteristic function $\tilde{p}_\tau(\omega)$ for the Gaussian PDF, approaches one for small α ; (e) the error factor $(\tilde{p}_\tau(0) - |\tilde{p}_\tau(\omega)|^2)$ in (5) for the uniform PDF, approaches zero for small α ; (f) the error factor $(\tilde{p}_\tau(0) - |\tilde{p}_\tau(\omega)|^2)$ in (5) for the Gaussian PDF, approaches zero for small α .

signal in (3), if standard averaging is applied to estimate the signal. For the ideal synchronized case, α is very small, no filtering effect is present, i.e., the characteristic function is equal to one for all ω . For increasing α the high frequency suppression becomes visible, destroying the underlying signal's shape. For a jitter between $\pm T_s$ or $\alpha = 1$, most high frequencies components in the signal are lost.

The covariance function of the measurement, given by the expression in (4), indicates that signal components which normally should not exist, are present in both the diagonal elements ($\omega_k = \xi_k$) and off diagonal elements ($\omega_k \neq \xi_k$) in the covariance matrix. These components are introduced by the factor $(\tilde{p}_\tau(\omega - \xi) - \tilde{p}_\tau(\omega)\tilde{p}_\tau(-\xi))$, which for synchronized measurements is zero ($\tilde{p}_\tau(\omega) = 1$). Assuming that the asymptotic diagonal property of the covariance matrix applies, discussed in section 3.1, the covariance matrix is diagonal and an expression for the effect can be investigated along the diagonal elements ($\omega_k = \xi_k$). The diagonal effect is given by the variance of the measurement as

$$\text{var}_{\tau, \tilde{\omega}}\{\tilde{y}(\omega)\} = (\tilde{p}_\tau(0) - |\tilde{p}_\tau(\omega)|^2)|\tilde{x}(\omega)|^2 + \text{E}_{\tilde{\omega}}\{|\tilde{w}(\omega)|^2\}. \quad (5)$$

In the third row, the error factor $(\tilde{p}_\tau(0) - |\tilde{p}_\tau(\omega)|^2)$ is shown for the uniform distribution in figure 1(e) and for the Gaussian distribution in figure 1(f). As mentioned earlier, and given by (5), this factor should be equal to zero for all ω to obtain an accurate estimate of the covariance matrix. If synchronization is ignored the high frequency components of the signal, which were suppressed earlier, leaks into the covariance estimate in the error formation step, $\tilde{y}(\omega) - E_{\tau, \tilde{w}}\{\tilde{y}(\omega)\}$, appendix A.1, destroying the covariance estimate in high frequency regions.

If this covariance matrix is used to weight objective functions, create uncertainty bounds or detect model errors, it will damage the estimates and create false bounds and errors.

Note that averaging over the measurements is an observation of the expected value, and hence (2) or (3) describes the average effect, or the effect when M is very large. For a small number of measurements M , the effect varies with each realization around the expected effect. By the central limit theorem (and also by the asymptotic properties discussed in section 3.1), the average estimator

$$\hat{x}(\omega) = \frac{1}{M} \sum_{m=1}^M \tilde{y}_m(\omega) \quad (6)$$

is asymptotically complex Gaussian distributed with the expected value given by (3) and with a covariance given in (4) scaled with the factor M^{-1} .

3. The maximum likelihood estimator

This section begins with a short summary of the main benefits obtained by using a frequency-domain approach to address the synchronization problem. The problem is rewritten in the discrete frequency domain, and then the log-likelihood function is stated, required to form the maximum likelihood estimator (MLE). The MLE is derived for two cases, one where the covariance matrix is known *a priori*, and one where it is unknown and needs to be estimated. Also, an additional iterative procedure for the second case is presented, that reduces the computational complexity and the sensitivity against local maxima. Finally, comments concerning the numerical optimization, required to estimate the parameters, are given for the derived methods.

3.1. Frequency-domain representation

Transferring the problem into the discrete frequency-domain results in a number of advantages compared to standard time-domain techniques.

- (i) The covariance matrix of the discrete Fourier transformed noise is asymptotically (as the number of data points goes to infinity) diagonal [7–9], making it easily invertible and mathematically tractable for large data sets.
- (ii) If the measurement noise is non-Gaussian distributed in the time-domain, the discrete Fourier transformed noise is asymptotically complex Gaussian distributed [7, 10], and the asymptotic MLE results in minimizing a weighted least-squares (WLS) cost-function (identical to the Gaussian case).

- (iii) Operations using derivatives of the cost-function, e.g. to find critical points or theoretical bounds, are easily implemented in the frequency domain (discrete difference approximations or sub-sample interpolation is avoided).

In the discrete frequency domain, (1) can be written as

$$\tilde{y}_m[k] = \exp\{-j\omega_k \tau_m\} \tilde{x}[k] + \tilde{w}_m[k], \quad (7)$$

for $k = 1, 2, \dots, N$ frequency points, where $\tilde{y}_m[k]$, $\tilde{x}[k]$ and $\tilde{w}_m[k]$ represent the discrete Fourier transform (DFT) of the sequences $y_m[n]$, $x[n]$ and $w_m[n]$, respectively, and ω_k is the k th corresponding angular frequency. Here $j = \sqrt{-1}$.

For aperiodic burst type of signals or pulses, the continuous signal $x(t)$ can be considered in practice both band- and time limited, if the sampling frequency and window size are correctly chosen. For the right choice, the effect of leakage- and aliasing errors occurring when truncating and sampling it can be neglected. This means that the signal can be periodically extended, with practically negligible artifacts, at the transition from one period to another. For periodic (bandlimited) signals the DFT coefficients gives a perfect representation of the true spectrum of the signal at the discrete frequencies.

Note that for unsynchronized measurements of a periodic signal, each estimated period contains either slightly more or slightly less than one period, due to sampling drifts. In theory, this introduces artifacts if the estimated period is periodically extended. However, if the drift is moderate, i.e. a fraction of a sample over an estimated period, this effect can be neglected. For moderate drifts, the DFT coefficients gives a close to perfect representation of the underlying periodic signal and the model in (7) is applicable.

3.2. The log-likelihood function

The log-likelihood function for M repeated measurements, assuming that $\tilde{\mathbf{w}}_m$ is complex Gaussian distributed, i.e. $\tilde{\mathbf{w}}_m \sim \mathcal{CN}(\mathbf{0}, \mathbf{C}_{\tilde{\mathbf{w}}})$, and $\tilde{\mathbf{w}}_m$ is independent of $\tilde{\mathbf{w}}_l$ for $m \neq l$, is

$$\ln(p(\tilde{\mathbf{Y}}; \tilde{\mathbf{x}}, \tau)) = A - \sum_{m=1}^M \tilde{\mathbf{e}}_m^H \mathbf{C}_{\tilde{\mathbf{w}}}^{-1} \tilde{\mathbf{e}}_m, \quad (8)$$

where

$$\begin{aligned} \tilde{\mathbf{Y}} &= [\tilde{\mathbf{y}}_1, \dots, \tilde{\mathbf{y}}_M], \\ \tilde{\mathbf{y}}_m &= [\tilde{y}_m[1], \dots, \tilde{y}_m[N]]^T, \\ \tilde{\mathbf{e}}_m &= \tilde{\mathbf{y}}_m - \mathbf{G}(\tau_m)\tilde{\mathbf{x}}, \\ \mathbf{G}(\tau_m) &= \text{diag}\{\exp\{-j\omega_1 \tau_m\}, \dots, \exp\{-j\omega_N \tau_m\}\}^T, \\ \tau &= [\tau_1, \dots, \tau_M]^T, \\ \tilde{\mathbf{x}} &= [\tilde{x}[1], \dots, \tilde{x}[N]]^T, \\ A &= -M(N \ln(\pi) + \ln(\det(\mathbf{C}_{\tilde{\mathbf{w}})})), \end{aligned}$$

where T denotes transposed, and H represents transposed and complex conjugated (Hermitian). $\mathcal{CN}(\mathbf{0}, \mathbf{C}_{\tilde{\mathbf{w}}})$ denotes a multivariate complex Gaussian distribution with mean $\mathbf{0}$ and covariance matrix $\mathbf{C}_{\tilde{\mathbf{w}}}$. $\text{diag}\{\mathbf{v}\}$ represents a diagonal matrix with diagonal elements given by the elements in the vector \mathbf{v} .

3.3. The MLE with known covariance matrix

Assuming that the covariance matrix $\mathbf{C}_{\tilde{\mathbf{w}}}$ is known *a priori*, the unknown quantities that need to be estimated are the vectors τ and $\tilde{\mathbf{x}}$. Maximizing (8) with respect to τ and $\tilde{\mathbf{x}}$, is a separable nonlinear weighted least-squares problem. The problem is linear in $\tilde{\mathbf{x}}$ but nonlinear in the synchronization jitter τ . For a given τ a closed form solution exist for maximizing (8) with respect to $\tilde{\mathbf{x}}$. Setting the derivative

$$\frac{\partial \ln(p(\tilde{\mathbf{Y}}; \tilde{\mathbf{x}}, \tau))}{\partial \tilde{\mathbf{x}}} = \sum_{m=1}^M (\mathbf{G}^H(\tau_m) \mathbf{C}_{\tilde{\mathbf{w}}}^{-1} (\tilde{\mathbf{y}}_m - \mathbf{G}(\tau_m) \tilde{\mathbf{x}}))^* \quad (9)$$

to zero, the maximum is given by

$$\hat{\tilde{\mathbf{x}}} = \left(\sum_{m=1}^M \mathbf{G}^H(\tau_m) \mathbf{C}_{\tilde{\mathbf{w}}}^{-1} \mathbf{G}(\tau_m) \right)^{-1} \sum_{m=1}^M \mathbf{G}^H(\tau_m) \mathbf{C}_{\tilde{\mathbf{w}}}^{-1} \tilde{\mathbf{y}}_m, \quad (10)$$

where $\hat{\cdot}$ denotes estimate and $*$ complex conjugate. By inserting the closed-form expression for $\hat{\tilde{\mathbf{x}}}$ into (8), the dimensionality of the optimization problem is reduced significantly to a maximization of

$$\ln(p(\tilde{\mathbf{Y}}; \hat{\tilde{\mathbf{x}}}, \tau)) = A - \sum_{m=1}^M \tilde{\varepsilon}_m^H \mathbf{C}_{\tilde{\mathbf{w}}}^{-1} \tilde{\varepsilon}_m \quad (11)$$

with respect to τ only, where

$$\tilde{\varepsilon}_m = \tilde{\mathbf{y}}_m - \mathbf{G}(\tau_m) \hat{\tilde{\mathbf{x}}}. \quad (12)$$

3.4. The MLE with unknown covariance matrix

The derivations in section 3.3 relies on *a priori* knowledge of the covariance matrix of the noise. If this information is unknown *a priori*, which is often the case, it can be estimated from the repeated measurements. The unknown quantities that need to be estimated are τ , $\tilde{\mathbf{x}}$ and $\mathbf{C}_{\tilde{\mathbf{w}}}$.

Assuming that the diagonal properties of the (frequency domain) covariance matrix applies, section 3.1, i.e. $\mathbf{C}_{\tilde{\mathbf{w}}} = \text{diag}([\sigma_{\tilde{w}}^2[1], \dots, \sigma_{\tilde{w}}^2[N]]^T)$, the log-likelihood function in (11) can be simplified to a maximization of

$$\begin{aligned} \ln(p(\tilde{\mathbf{Y}}; \hat{\tilde{\mathbf{x}}}, \sigma_{\tilde{w}}^2, \tau)) \\ = - \sum_{m=1}^M \sum_{k=1}^N \frac{\tilde{\varepsilon}_m^*[k] \tilde{\varepsilon}_m[k]}{\sigma_{\tilde{w}}^2[k]} + \ln(\sigma_{\tilde{w}}^2[k]) + \ln(\tau), \end{aligned} \quad (13)$$

with respect to $\sigma_{\tilde{w}}^2$ and τ . Setting the partial derivative of (13) with respect to $\sigma_{\tilde{w}}^2[k]$

$$\frac{\partial \ln(p(\tilde{\mathbf{Y}}; \hat{\tilde{\mathbf{x}}}, \sigma_{\tilde{w}}^2, \tau))}{\partial \sigma_{\tilde{w}}^2[k]} = \sum_{m=1}^M \frac{\tilde{\varepsilon}_m^*[k] \tilde{\varepsilon}_m[k]}{\sigma_{\tilde{w}}^4[k]} - \frac{1}{\sigma_{\tilde{w}}^2[k]} \quad (14)$$

to zero, the maximum is given by

$$\hat{\sigma}_{\tilde{w}}^2[k] = \frac{1}{M} \sum_{m=1}^M \tilde{\varepsilon}_m^*[k] \tilde{\varepsilon}_m[k], \quad (15)$$

assuming that $\hat{\sigma}_{\tilde{w}}^2[k] \neq 0, \forall k$. Inserting this closed-form expression into (13), the dimensionality reduces, analogous as before, to maximization of

$$\begin{aligned} \ln(p(\tilde{\mathbf{Y}}; \hat{\tilde{\mathbf{x}}}, \hat{\sigma}_{\tilde{w}}^2, \tau)) \\ = - \sum_{m=1}^M \sum_{k=1}^N \frac{\tilde{\varepsilon}_m^*[k] \tilde{\varepsilon}_m[k]}{\hat{\sigma}_{\tilde{w}}^2[k]} + \ln(\hat{\sigma}_{\tilde{w}}^2[k]) + \ln(\tau) \end{aligned} \quad (16)$$

with respect to τ only. Note that for a diagonal covariance matrix, the signal estimate in (10) reduces to

$$\hat{\tilde{\mathbf{x}}} = \frac{1}{M} \sum_{m=1}^M \mathbf{G}^H(\tau_m) \tilde{\mathbf{y}}_m, \quad (17)$$

from the diagonal and unitary property of $\mathbf{G}(\tau_m)$.

3.5. The iterative maximum likelihood estimator (IMLE)

In noisy measurement situations with few repetitions or periods M , the log-likelihood function in (16) is very sensitive to small changes in the parameters, and a good initial guess is necessary to avoid local maxima. Another problem is encountered for the opposite situation with a large number of repetitions or periods M . The Jacobian matrices, necessary for the optimization algorithm in section 3.6, are of dimension $\mathbb{C}^{M \times NM}$ and for large data records these matrices become unmanageable.

For these situations we can resort to sub-optimal methods, where the dependence of the parameter is less nonlinear and the log-likelihood surface more tractable. The idea with an iterative approximation is to reduce the parameter dependence, by replacing the unknown underlying signal waveform and covariance matrix, with estimates obtained from the previous iteration step. This reduces sensitivity of the log-likelihood function with respect to the parameter, and it reduces the dimension of the Jacobian matrices by a factor M to a more manageable size.

The IMLE can be summarized as follows:

- (1) Start with an initial guess of τ ,
- (2) Estimate the covariance matrix $\hat{\mathbf{C}}_{\tilde{\mathbf{w}}}$ and the underlying waveform $\hat{\tilde{\mathbf{x}}}$, using (15), (17), and the current τ ,
- (3) Maximize (11) with respect to τ , considering $\hat{\mathbf{C}}_{\tilde{\mathbf{w}}}$ and $\hat{\tilde{\mathbf{x}}}$ as known quantities (obtained from the previous step),
- (4) If not converged, go back to step 2.

This method produces a more malleable likelihood surface around its global maximum and it reduces the computational complexity of the optimization procedure by a factor M , compared to (16). Another advantage is that the full sample covariance matrix can be used (without any complications) if N is small and the assumption of diagonal covariance, discussed in section 3.1, is questionable. However, the benefits of an iterative approximation generally comes at the cost of losing some efficiency [10].

3.6. Optimization aspects

Estimating the synchronization jitter results in a nonlinear maximization problem with respect to the parameter τ , and numerical optimization methods must be used to find the maximum. Most numerical methods use local information, such as a truncated Taylor series expansion of the log-likelihood function, around a point to find the next step. If the initial starting point is well chosen and the function is well behaved in the neighborhood of its global maxima, the iteration will converge.

In this paper, the Gauss–Newton method is used to maximize log-likelihood functions. This choice is based on

the method's simplicity and local performance in solving least-squares problems. The method can be summarized by the following steps:

- (1) Create the measurement matrix $\tilde{\mathbf{Y}}$ in (8).
- (2) Find $\hat{\tau}$, $\hat{\mathbf{x}}$ and $\hat{\sigma}_w^2$ by maximizing the log-likelihood function, i.e. maximize (11) or (16) with respect to τ . The Newton maximization steps are:
 - (2.1) Find an initial guess $\tau^{(p)}$ for $p = 0$.
 - (2.2) Form the Hessian and gradient of the log-likelihood function, with respect to $\tau^{(p)}$.
 - (2.3) Find $\tau^{(p+1)}$ using (A.5) and (A.6) in appendix A.2.
 - (2.4) If $\|\Delta\tau^{(p)}\|$ in (A.6) and $|L(\tilde{\mathbf{Y}}; \tau^{(p+1)}) - L(\tilde{\mathbf{Y}}; \tau^{(p)})|$ are large then increase p , i.e. $p = p + 1$, and go back to step 2.2, otherwise continue.
- (3) The maximum likelihood is found and the estimates are given by (15), (17) and $\hat{\tau} = \tau^{(p+1)}$.

If the measurements are highly unsynchronized, i.e. $|\tau_m| > T_s$, a useful step is to pre-align the measurements to the closest whole sample using standard cross-correlation techniques, to avoid local likelihood maxima.

Note that the complex valued vectors in (8) originate from a DFT of real valued sequences. This means that the same information can be found in both halves of the vectors due to the symmetry property $\tilde{u}[k] = \tilde{u}^*[N+2-k]$, for $1 < k \leq N$, if $\tilde{\mathbf{u}} = \text{DFT}\{\mathbf{u}\}$ and $\mathbf{u} \sim \mathbb{R}^N$. Hence, the estimation problem can be confined to one half of the vectors, e.g. to $k = 1, \dots, K$, where $K = N/2 + 1$ if N is even, or $K = (N + 1)/2$ if N is odd.

4. Theoretical results

In this section the main theoretical results, obtained from the derivations of the estimators and bounds, are presented and interpreted.

4.1. Estimation with singular matrices

The Fisher information matrix (FIM) [6] in appendix A.3, and the Hessian matrix in appendix A.2 does not have full rank, if $\tau \sim \mathbb{R}^M$. The reason for this is that τ and $\tau + a$ for any scalar a produces the same likelihood. In other words the Hessian and the FIM have a null space. This is not a surprise since the perfect reference point is unknown, and the objective is only to minimize the synchronization jitter differences relative to each other. Hence, adding an arbitrary constant to all the jitters (shifting all measurements by an equal amount) does not change the synchronization problem and thereby not the likelihood. This rank deficiency is common in estimation problems with over-parameterized structures [11] and can be solved using various approaches.

One alternative is to reduce the model size by excluding parameters, in our case one parameter. This can be achieved by 'grounding' one synchronization jitter, i.e., set $\tau_k = 0$ for some k and make this the reference point (similar to an electric circuit). This approach excludes a column and a row in the FIM and the Hessian which produces full rank, i.e. $\tau \sim \mathbb{R}^{M-1}$.

However, if the particular realization τ_k lies far from its expected value (far from the average of all others), it will slow down the convergence process, since the other measurements or periods have to be aligned iteratively to this unlucky fixed reference point. To fasten up the convergence, an alternative is to choose a reference close the average, which brings us to the next approach. To solve this with respect to an average reference point, the system of equations can be solved using the Moore–Penrose pseudoinverse [11] to cope with the rank deficiency and produce the minimum norm solution.

4.2. Interpreting the CRLB for τ

In appendix A.4, expressions for the CRLB for τ are given. Assuming that grounding is applied (section 4.1) and $\tau \sim \mathbb{R}^{M-1}$, the variance of the estimated parameter is bounded by

$$\text{var}(\hat{\tau}) \geq \left(\sum_{k=1}^K \text{DNR}[k] \right)^{-1}, \quad (18)$$

where

$$\text{DNR}[k] = \frac{|j\omega_k \tilde{x}[k]|^2}{\sigma_w^2[k]}, \quad (19)$$

given that $\hat{\tau}$ is obtained using an unbiased estimator. Equation (19) can be interpreted as the signal's derivative-to-noise ratio (DNR) (at frequency ω_k), and the bound in (18) is given by the inverse of the sum of the DNR over all frequencies. The performance of an estimator is limited by the frequency content of the signal's derivative, in comparison to the frequency content of the noise.

The bound reveals some of the differences between the two implementation approaches described in section 4.1. An approximate expression for the CRLB using the pseudoinverse approach is given by (A.28). Comparing the diagonal (the variance) of this expression with the diagonal of (A.26) obtained using grounding, it differs with a factor $(1 - 1/M)$ compared to $(1 + 1)$. The pseudoinverse approach can be interpreted as aligning with respect to the average displacement (minimum norm solution), thus decreasing the variance of the parameter vector with $1/M$ (losing one degree of freedom). Grounding, on the other hand, aligns with respect to one synchronization jitter assumed known (which of course it is not), increasing the variance twice as much. Note that, even though the CRLB differs between the approaches, similar synchronization results are obtained. This means that given an estimate $\hat{\tau}$ obtained using one approach it can be converted to the other, either by subtracting its mean value (pseudoinverse) or by subtracting $\hat{\tau}_k$ for some k (grounding).

5. Simulation results

Four different methods are compared using computer simulations:

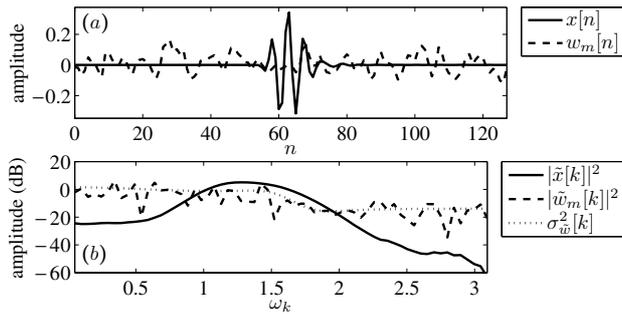


Figure 2. Time- and frequency-domain representation of the signal and noise, with a SNR of 0 dB: (a) discrete-time domain representation of the simulated signal $x[n]$ (solid line) and a realization of one noise sequence $w_m[n]$ (dashed line); (b) discrete frequency domain representation of the signal (solid line), noise (dashed line) and variance (dotted line).

SE	The ‘simple’ estimator, which is the most commonly used method to synchronize measurements. The measurements are aligned against one of the measured signals (used as reference). Neither the underlying signal waveform nor the noise covariance are estimated or used. The time complexity is of the order $O(NM)$, where O represents the <i>Big O notation</i> [12].
LSE	The (nonlinear) least-squares estimator, which is the MLE in additive white Gaussian noise or if the covariance is known <i>a priori</i> , described in section 3.3. The underlying signal waveform is estimated and used. The time complexity is of the order $O(NM^2)$.
MLE	The maximum likelihood estimator, described in section 3.4. The underlying signal waveform and the noise covariance are estimated and used. The time complexity is of the order $O(NM^2)$.
IMLE	The iterative maximum likelihood estimator, described in section 3.5, which is an iterative approximation of the MLE. Both the underlying signal waveform and the noise covariance are estimated and used. The time complexity is of the order $O(NM)$, the same as for the SE.

All methods are implemented in the frequency domain. To obtain fast convergence, the pseudoinverse approach discussed in section 4.1 is used. However, grounding is still applied (after the algorithm has converged, see last part of section 4.2) to give a fair comparison with the SE.

5.1. The simulated signal

In the simulations, a measured ultrasonic echo is used to represent the underlying signal waveform (see figure 2(a)). The signal is retrieved from a low-noise measurement of a broadband ultrasonic pulse (see section 6).

Since signals of this type have a close to finite support, the calculation of the signal-to-noise ratio (SNR) in this paper involves the energy of the signal (ignoring T_s), instead of the

power (or variance). The SNR is defined here as

$$\text{SNR} = \frac{\sum_{k=1}^K |\tilde{x}[k]|^2}{\sum_{k=1}^K \sigma_w^2[k]}. \quad (20)$$

5.2. The noise model

The additive noise $w_m[n]$ is modeled as Gaussian and colored using discrete ARMA filters as

$$\tilde{w}_m[k] = H[k]\tilde{\epsilon}_m[k] + B\tilde{\epsilon}_m[k], \quad (21)$$

where

$$H[k] = \frac{\sum_{l=1}^{N_b} b_l \exp\{-j2\pi(k-1)(l-1)/N\}}{\sum_{l=1}^{N_a} a_l \exp\{-j2\pi(k-1)(l-1)/N\}} \quad (22)$$

is the ARMA filter, B represents a white noise floor and $\tilde{\epsilon}_m[k]$ is the DFT of a white Gaussian noise sequence. The ARMA filter $H[k]$ is used to model:

- Applied pre-filters before sampling that affect the noise characteristics. Common pre-filters are high-pass, band-pass, and low-pass filters to remove trends, disturbances, and aliasing effects respectively.
- Additive (nonwhite) disturbances from the surrounding.

The noise floor B is used to model additive white noise at the analog-to-digital converter.

For the model in (21), the (co)variance is calculated as $\sigma_w^2[k] = |H[k] + B|^2$.

5.3. Additive colored Gaussian noise

In figure 2(a), the signal waveform is seen together with a realization of a colored noise sequence, using the signal and noise models described earlier. The noise is colored using (21), with a third-order Chebyshev type I filter with filter coefficients $\mathbf{b} = [0.1059, 0.3177, 0.3177, 0.1059]^T$, $\mathbf{a} = [1.0000, -0.5620, 0.7194, -0.3102]^T$, a white noise floor at $B = 0.2$ (-10 dB), and a total SNR of 0 dB. The specific parameters are chosen to resemble the experimental situation encountered in section 6, where the build in low-pass pre-filter in the pulser/receiver is used to reduce aliasing effects.

In the simulations, the synchronization jitter τ is modeled using independent realizations from a uniform distribution within the interval $[-T_s, T_s]$, i.e. $\tau_m \sim \mathcal{U}[-T_s, T_s]$. In figure 3(a), the mean square error (MSE) for the different methods can be seen together with the Cramér–Rao lower bound (CRLB), versus the number of measurements M , for a constant SNR of 0 dB.

For a small number of measurements, the MSEs are comparable between the methods. As M increases a big improvement can be seen as the IMLE and the MLE approaches the CRLB. A small difference is visible between the SE and the LSE for this particular SNR. To observe larger differences, the SNR has to be decreased further.

In figure 3(b), the MSE versus the SNR is shown for $M = 20$ measurements. As the SNR drops below 2 dB, a significant difference is observable between the SE and the LSE. The visible improvement is the result of estimating the underlying signal waveform (LSE), compared to using

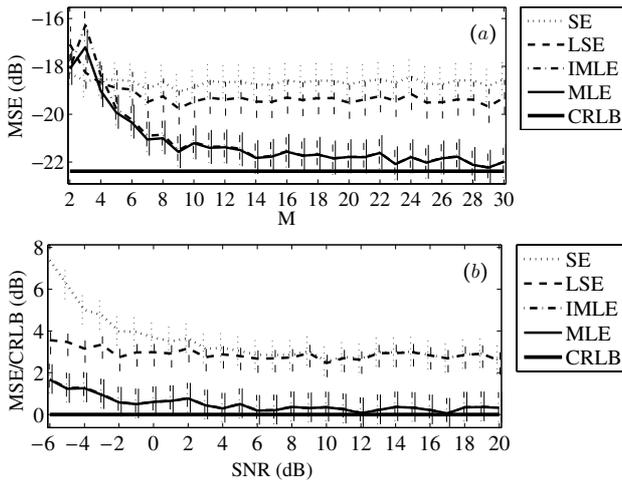


Figure 3. Mean square error (MSE) comparison of $\hat{\tau}_m$ using 500 realizations. The vertical lines mark the 99% confidence intervals. (a) MSE versus the number of measurements M , with the noise conditions described in figure 2; (b) MSE/CRLB versus SNR, with $M = 20$ and noise colored as in figure 2. To get a better view of the closeness to the CRLB, the ratio between the MSE and the CRLB is investigated in (b) to avoid the otherwise strong dependence on the SNR.

one noisy measurement as reference (SE). Nevertheless, the maximum-likelihood based methods are consistently much better, due to the use of second-order statistics.

Note that the CRLB is constant for increasing number of measurements M . The reason for this is that the parameter vector τ is increased proportionally. A necessary condition for the CRLB to decrease is that the dimension of the parameter vector remains constant while the amount of data is increased [6].

5.4. Additive white Gaussian noise

For white Gaussian noise the MSE for the LSE coincides with the IMLE, the MLE and the CRLB for large M . However, a small improvement can be observed for the LSE for small M , in figure 4(a). The improvement is due to the *a priori* information (white Gaussian noise) used in the LSE, and not considered by the IMLE or the MLE. Note that the LSE for this situation is the true MLE given this *a priori* information. For this particular situation, the LSE correctly guesses the noise color as white, hence the increase in performance.

For a small number of measurements, $M < 5$, the IMLE and the MLE give poor performance, due to the high uncertainties added by estimating the covariance in (15) from a small number of observations (M).

6. Experimental results

In this section, synchronization and estimation results are presented using the MLE (section 3) on repeated pulse-echo measurements.

6.1. Experimental setup

The experimental setup consists of an ultrasonic transducer mounted on a custom built measurement cell consisting of a

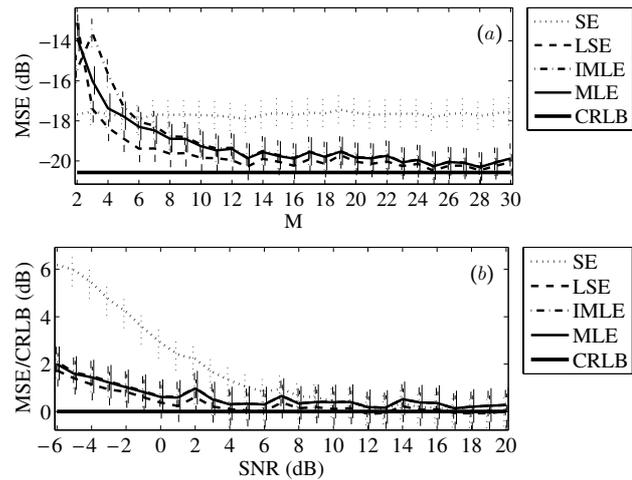


Figure 4. Mean square error (MSE) comparison of $\hat{\tau}_m$ using 500 realizations. The vertical lines mark the 99% confidence intervals. (a) MSE versus the number of measurements M , with the white Gaussian noise, i.e. $\mathbf{a} = 1$, $\mathbf{b} = 1$, $B = 0$ and SNR = 0 dB; (b) MSE/CRLB versus SNR, with $M = 20$ and white noise. To get a better view of the closeness to the CRLB, the ratio between the MSE and the CRLB is investigated in (b) to avoid the otherwise strong dependence on the SNR.

water region and plexiglass plate. A pulser/receiver 5073PR from Panametrics was used to excite the transducer and amplify the received signal. A built-in low-pass filter, with a cutoff frequency of 20 MHz, was used to avoid high frequency signal and noise components from folding down during sampling. The signal was then digitized using a CompuScope 12400 oscilloscope card, by Gage Applied Technologies Inc., Lachine, QC, Canada, at 12-bit resolution, using a sampling rate of 100 MHz. The oscilloscope was triggered using the synchronization output of the pulser/receiver.

Repeated pulse-echo tests were performed using a broadband piezoelectric transducer V317, manufactured by Panametrics, Waltham, MA, USA.

6.2. Synchronization and estimation results

A possible way to detect poor synchronization from experimental data, is to investigate the variance where the derivative of the signal is large (e.g. zero crossings) and compared it against the variance where the derivative is smaller [3]. Figure 5 shows a comparison between measured (unsynchronized) pulse-echo signals and synchronized signals using the MLE. The figure shows 20 measurements with an estimated SNR of 14 dB, using the experimental setup described in section 6.1. Notice the decrease in horizontal variation for the synchronized signals.

In figure 6, the mean value $\hat{x}[n]$ and the estimated standard deviation $\hat{\sigma}_{\hat{x}}[n]$ of the 20 measurements are shown. In figure 6(a) we can see that the estimated signals (the mean values) differ between the unsynchronized and synchronized case in both shape and size. This is a consequence of the convolution effect given by (2) in section 1, indicating that the signal shape is affected if standard averaging is applied

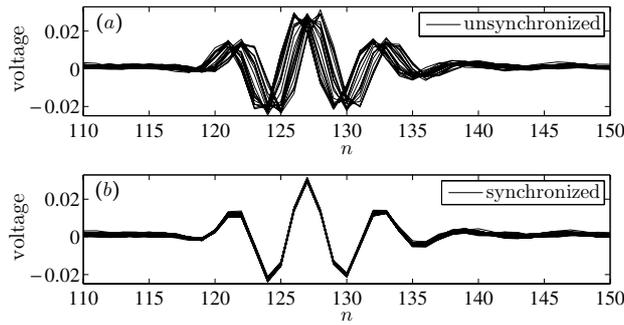


Figure 5. Time-domain plot: (a) 20 measured ultrasonic signals; (b) the same signals as in (a) but synchronized using the MLE.

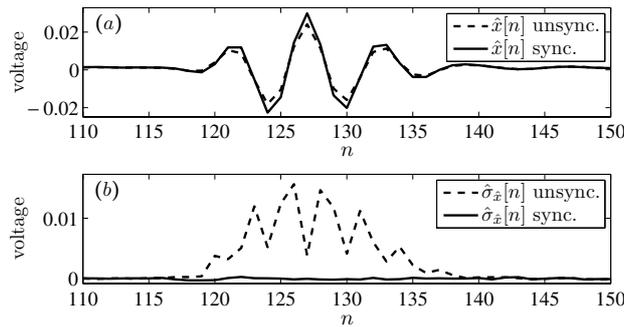


Figure 6. Time-domain plot: (a) estimated mean values of the measured signals in figure 5; (b) estimated standard deviations of the measured signals in figure 5.

to estimate the signal. The standard deviation curves in figure 6(b) clearly shows the increase in standard deviation for the unsynchronized case at regions where the derivative is large. For the synchronized case, no significant change is visible.

The theoretical investigation in section 2 showed that high frequency information of the underlying signal was suppressed. In figure 7(a), the (scaled) periodogram of the mean value $|\hat{x}[k]|^2$ is visible for the unsynchronized and synchronized case together with the periodogram of one single measurement $|\hat{y}_m[k]|^2$. A good agreement is visible between the synchronized mean value and one single measurement, indicating that a good synchronization is obtained. However, comparing the unsynchronized periodogram with the synchronized and measured one, we can observe that high frequency information of the estimated signal is suppressed. The difference is approximately -2.3 dB at 20 MHz (or $0.4 \cdot F_s/2$), which agrees with the calculated effect of $|\hat{p}_\tau(\omega)|^2 = -2.5$ dB at $\omega = 0.4\pi F_s$ for a uniform distribution with $\alpha = 1$; see (3) and figure 1(c). Assuming a uniform distribution for this experiment, with $\alpha = 1$, is validated by the histogram of the estimated jitters in figure 8.

For unsynchronized measurements, the variance is also higher at areas where the signal has high frequency components. This conclusion is given by the error factor $(\hat{p}_\tau(0) - |\hat{p}_\tau(\omega)|^2)$ in (5), and visible in figure 1(e)–1(f). In figure 7(b), estimates of the variance (noise color) are shown for both cases. The estimated variance obtained from the unsynchronized case clearly shows the presence of signal

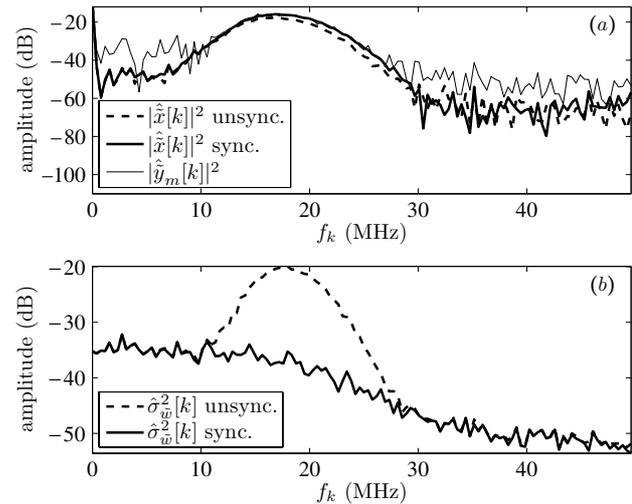


Figure 7. Frequency-domain plot: (a) periodogram the mean values in figure 6(a), shown together with a periodogram of a single measurement; (b) (co)variance comparison of raw measurements versus synchronized measurements.

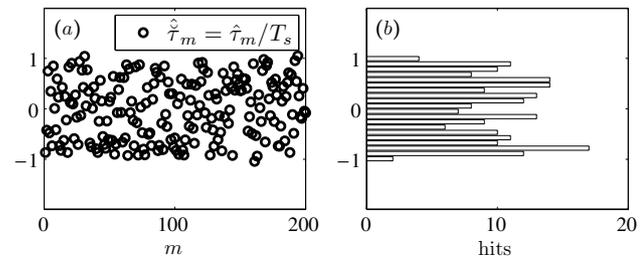


Figure 8. Estimated jitter distribution: (a) estimated jitter, normalized as in figure 1(a); (b) histogram.

components destroying the variance estimate. This agrees with the conclusions drawn in section 2. The estimated variance using the MLE indicates no signal components, and the shape of the periodogram agrees with the effect introduced by the applied low-pass anti-aliasing filter, see section 6.1. The calculated effect of $(\hat{p}_\tau(0) - |\hat{p}_\tau(\omega)|^2) = -3.7$ dB at 20 MHz (or $0.4 \cdot F_s/2$) agrees with an estimated effect of $(\hat{\sigma}_y^2[k] - \hat{\sigma}_w^2[k]) / |\hat{x}[k]|^2 = -3.5$ dB, using the relationship in (5), where $\hat{\sigma}_y^2[k]$ and $\hat{\sigma}_w^2[k]$ represents the estimated unsynchronized and synchronized variance, respectively, in figure 7(b).

7. Discussion

7.1. Real experiments

The proposed estimator has been used on real experiments of repeated measurements of ultrasonic signals. However, for real experiments, the problem of not knowing the truth, i.e. not knowing the underlying signal waveform, the covariance, nor the true synchronization jitter, prevent us from comparing the true performance. Higher variance at high frequencies or at areas where the signal has a large derivative is an indication

of poor synchronization and could be used as measure of the performance. This is, however, only an indication of poor synchronization, since other factors such as sample jitter [13] and noise color produce similar effects.

7.2. Asymptotic properties

The asymptotic properties of the MLE (the estimated parameter is asymptotically unbiased, normally distributed and with a covariance given by the CRLB) are only valid if the errors are small. Small errors means either that the data size is large enough, compared to the dimension of the estimated parameters, or that the SNR is so high that statistical linearity applies [6].

Note also that the CRLB is a lower bound for unbiased estimators, which none of the estimators in this paper are. However, the LSE, IMLE and MLE are asymptotically unbiased, which means that a comparison is valid against the bound for a sufficiently large data size or SNR.

7.3. Extensions

The truncation effect on aperiodic signals that do not have finite support in the time domain is not investigated in this paper. Simulations using truncated signals in moderate synchronization jitter indicate similar results as for those with finite support. For situations with highly unsynchronized measurements, adequate windowing centered at τ should be included in the signal model to preserve similarity between the measurements. This prevents signal components, obtained from the non-overlapping regions, from destroying the estimates.

The estimators derived in this paper are applicable to other types of repeated or periodically excited signals. An interesting extension would be to expand the model in dimension to cope with image processing applications, e.g., synchronization of repeated unsynchronized (shaky) image sequences.

8. Conclusions

In this paper a maximum likelihood estimator (MLE) is derived for the underlying signal waveform, the covariance matrix and the synchronization jitter. The estimator applies for repeated unsynchronized measurements, or unsynchronized measurements of periodically excited signals. The effect of poor synchronization on the estimated signal and covariance is studied, indicating that even a small sub-sample misalignment produces unwanted filtering effects.

The estimator is compared with other sub-sample synchronization techniques, for wideband signals in the presence of white and colored noise. We demonstrated with simulations that the mean square error (MSE) of the estimates is significantly lower using the proposed method than for related techniques. The estimator's MSE was also compared with the theoretical lower bound (CRLB), showing asymptotic efficiency for large data size or signal-to-noise ratio conditions. Estimation results are presented for repeated

measurements to validate the estimator for real measurements, and to experimentally support the theoretical results.

Acknowledgments

The author would like to express his sincerest gratitude to Anders Grennberg, Fredrik Häggglund, Johan E Carlson and Veronica J Nilsson for valuable comments and discussions.

Appendix

A.1. Derivation of the effect of unsynchronized measurements

Assuming that the synchronization jitter τ is independent of the measurement noise \tilde{w} , then the expected value of the measurements can be expressed in the frequency domain as

$$\begin{aligned} E_{\tau, \tilde{w}}\{\tilde{y}(\omega)\} &= E_{\tau, \tilde{w}}\{\exp\{-j\omega\tau\}\tilde{x}(\omega) + \tilde{w}(\omega)\} \\ &= E_{\tau}\{\exp\{-j\omega\tau\}\tilde{x}(\omega)\} + E_{\tilde{w}}\{\tilde{w}(\omega)\} \\ &= \tilde{p}_{\tau}(\omega)\tilde{x}(\omega), \end{aligned} \quad (\text{A.1})$$

where $\tilde{p}_{\tau}(\omega) = \int_{-\infty}^{\infty} p_{\tau}(\tau) \exp\{-j\omega\tau\} d\tau$ is the characteristic function, i.e. the Fourier transform of the PDF for τ .

The covariance function of the measurements in the frequency domain can be expressed as

$$\begin{aligned} \text{cov}_{\tau, \tilde{w}}\{\tilde{y}(\omega), \tilde{y}(\xi)\} &= E_{\tau, \tilde{w}}\{(\tilde{y}(\omega) - E_{\tau, \tilde{w}}\{\tilde{y}(\omega)\}) \\ &\quad \cdot (\tilde{y}(\xi) - E_{\tau, \tilde{w}}\{\tilde{y}(\xi)\})^*\} \\ &= E_{\tau, \tilde{w}}\{\tilde{y}(\omega)\tilde{y}(-\xi)\} - E_{\tau, \tilde{w}}\{\tilde{y}(\omega)\}E_{\tau, \tilde{w}}\{\tilde{y}(-\xi)\}, \end{aligned} \quad (\text{A.2})$$

where the property $\tilde{y}^*(\xi) = \tilde{y}(-\xi)$ for $y(t) \sim \mathbb{R}$ has been used. The last two expressions in (A.2) can be expanded as

$$\begin{aligned} E_{\tau, \tilde{w}}\{\tilde{y}(\omega)\tilde{y}(-\xi)\} &= E_{\tau, \tilde{w}}\{(\exp\{-j\omega\tau\}\tilde{x}(\omega) + \tilde{w}(\omega)) \\ &\quad \cdot (\exp\{j\xi\tau\}\tilde{x}(-\xi) + \tilde{w}(-\xi))\} \\ &= E_{\tau}\{\exp\{-j(\omega - \xi)\tau\}\} \tilde{x}(\omega)\tilde{x}(-\xi) + E_{\tilde{w}}\{\tilde{w}(\omega)\tilde{w}(-\xi)\} \\ &= \tilde{p}_{\tau}(\omega - \xi)\tilde{x}(\omega)\tilde{x}(-\xi) + E_{\tilde{w}}\{\tilde{w}(\omega)\tilde{w}(-\xi)\} \end{aligned} \quad (\text{A.3})$$

and

$$E_{\tau, \tilde{w}}\{\tilde{y}(\omega)\}E_{\tau, \tilde{w}}\{\tilde{y}(-\xi)\} = \tilde{p}_{\tau}(\omega)\tilde{p}_{\tau}(-\xi)\tilde{x}(\omega)\tilde{x}(-\xi). \quad (\text{A.4})$$

Inserting (A.3) and (A.4) into (A.2) gives (4).

A.2. Newton maximization

The optimization problem is to numerically maximize the log-likelihood function in equation (11) or equation (16) with respect to τ . The $(p+1)$ th iteration step is

$$\tau^{(p+1)} = \tau^{(p)} + \Delta\tau^{(p)}, \quad (\text{A.5})$$

where

$$\Delta\tau^{(p)} = - \left(\frac{\partial^2 L(\tilde{\mathbf{Y}}; \tau)}{\partial \tau \partial \tau^T} \Big|_{\tau=\tau^{(p)}} \right)^{\dagger} \frac{\partial L(\tilde{\mathbf{Y}}; \tau)}{\partial \tau} \Big|_{\tau=\tau^{(p)}}, \quad (\text{A.6})$$

and $L(\tilde{\mathbf{Y}}; \tau)$ is the shorthand notation for the log-likelihood function. Note that the inverse is replaced by the pseudoinverse [14], denoted \dagger , to cope with possible singularities, discussed in section 4.1.

To increase the stability and decrease the computational cost of the iteration in (A.6), the *method of scoring* is

applied [6]. Scoring improves the stability by taking the expected value of the Hessian in (A.6), i.e., the Hessian is replaced by the negative of the Fisher information matrix (FIM) [6]. Since analytical expressions for the FIM and its (pseudo)inverse can be derived, see appendix A.3, the computational cost of solving the linear system in (A.6) is reduced significantly. By applying the method of scoring, the Hessian is approximated using (A.24) as

$$\left. \frac{\partial^2 L(\tilde{\mathbf{Y}}; \boldsymbol{\tau})}{\partial \boldsymbol{\tau} \partial \boldsymbol{\tau}^T} \right|_{\boldsymbol{\tau}=\boldsymbol{\tau}^{(p)}} \approx -2 \sum_{k=1}^K \frac{|j\omega_k \hat{x}[k]|^2}{\hat{\sigma}_w^2[k]} \left(\mathbf{I} - \frac{1}{M} \right), \quad (\text{A.7})$$

where $\hat{x}[k]$ and $\hat{\sigma}_w^2[k]$ are obtained using (17) and (15) with $\boldsymbol{\tau} = \boldsymbol{\tau}^{(p)}$. The (pseudo)inverse is given by (A.26) or (A.28).

A.3. Calculation of Cramér–Rao lower bound (CRLB)

When deriving the CRLB in this section, the approach of setting one τ_k to zero (grounding) is assumed to be applied (discussed in section 4.1) unless otherwise stated. The expression for the Fisher information matrix (FIM) using complex valued signals and additive complex Gaussian noise is given in [6]. Applied to the measurements in $\tilde{\mathbf{Y}}$, where $\tilde{\mathbf{y}}_m \sim \mathcal{CN}(\tilde{\mathbf{s}}_m, \text{diag}(\boldsymbol{\sigma}_w^2))$ and $\tilde{\mathbf{w}}_m$ is independent of $\tilde{\mathbf{w}}_l$ for $m \neq l$, the FIM is given by

$$\begin{aligned} [\mathbf{F}(\boldsymbol{\theta})]_{a,b} &= \sum_{m=1}^M \sum_{k=1}^K 2\Re \left\{ \frac{\partial \tilde{s}_m^*[k]}{\partial \theta_a} \frac{1}{\sigma_w^2[k]} \frac{\partial \tilde{s}_m[k]}{\partial \theta_b} \right\} \\ &+ \frac{\partial \sigma_w^2[k]}{\partial \theta_a} \frac{\partial \sigma_w^2[k]}{\partial \theta_b} \sigma_w^{-4}[k], \end{aligned} \quad (\text{A.8})$$

where

$$\tilde{s}_m[k] = \exp\{-j\omega_k \tau_m\} \tilde{x}[k], \quad (\text{A.9})$$

$$\boldsymbol{\theta} = [\Re\{\tilde{\mathbf{x}}\}^T, \Im\{\tilde{\mathbf{x}}\}^T, \boldsymbol{\sigma}_w^2, \boldsymbol{\tau}^T]^T. \quad (\text{A.10})$$

Inserting the derivatives

$$\frac{\partial \tilde{s}_m[k]}{\partial \theta} = \begin{cases} \exp\{-j\omega_k \tau_m\} \delta_{k,a}, & \theta = \Re\{\tilde{x}[a]\} \\ j \exp\{-j\omega_k \tau_m\} \delta_{k,a}, & \theta = \Im\{\tilde{x}[a]\} \\ 0, & \theta = \sigma_w^2[a] \\ -j\omega_k \exp\{-j\omega_k \tau_m\} \delta_{m,a}, & \theta = \tau_a \end{cases} \quad (\text{A.11})$$

and

$$\frac{\partial \sigma_w^2[k]}{\partial \theta} = \begin{cases} 0, & \theta = \Re\{\tilde{x}[a]\} \\ 0, & \theta = \Im\{\tilde{x}[a]\} \\ \delta_{k,a}, & \theta = \sigma_w^2[a] \\ 0, & \theta = \tau_a \end{cases} \quad (\text{A.12})$$

in (A.8), the FIM can be expressed as

$$\mathbf{F}(\boldsymbol{\theta}) = \begin{bmatrix} \mathbf{A} & \mathbf{0} & \mathbf{0} & \mathbf{D} \\ \mathbf{0} & \mathbf{A} & \mathbf{0} & \mathbf{E} \\ \mathbf{0} & \mathbf{0} & \mathbf{B} & \mathbf{0} \\ \mathbf{D}^T & \mathbf{E}^T & \mathbf{0} & \mathbf{C} \end{bmatrix}, \quad (\text{A.13})$$

where

$$A_{a,b} = \delta_{a,b} \frac{2M}{\sigma_w^2[a]} (K \times K), \quad (\text{A.14})$$

$$B_{a,b} = \delta_{a,b} \frac{M}{\sigma_w^4[a]} (K \times K), \quad (\text{A.15})$$

$$C_{a,b} = \delta_{a,b} 2 \sum_{k=1}^K \frac{|j\omega_k \tilde{x}[k]|^2}{\sigma_w^2[k]} (M-1 \times M-1), \quad (\text{A.16})$$

$$D_{a,b} = 2 \frac{\omega_a \Im\{\tilde{x}[a]\}}{\sigma_w^2[a]} (K \times M-1), \quad (\text{A.17})$$

$$E_{a,b} = -2 \frac{\omega_a \Re\{\tilde{x}[a]\}}{\sigma_w^2[a]} (K \times M-1). \quad (\text{A.18})$$

The CRLB is given by the inverse of the FIM [6] as

$$\mathbf{CRLB}(\boldsymbol{\theta}) = \mathbf{F}^{-1}(\boldsymbol{\theta}). \quad (\text{A.19})$$

The lower bound for the covariance matrix of any unbiased estimate of $\boldsymbol{\theta}$ [6] is

$$\text{cov}(\hat{\boldsymbol{\theta}}) \geq \mathbf{CRLB}(\boldsymbol{\theta}). \quad (\text{A.20})$$

Note that if σ_w^2 is known *a priori*, the FIM is obtained by excluding the third block row and the third block column (those intersecting the matrix \mathbf{B}) in (A.13).

If the pseudoinverse approach is applied, discussed in section 4.1, the FIM for this situation is obtained by expanding (A.16)–(A.18) in dimension to include all M jitters. The CRLB for this case is given by the pseudoinverse [11] as

$$\mathbf{CRLB}(\boldsymbol{\theta}) = \mathbf{F}^\dagger(\boldsymbol{\theta}). \quad (\text{A.21})$$

A.4. The Cramér–Rao lower bound for τ

Let

$$\mathbf{F}(\boldsymbol{\theta}) = \begin{bmatrix} \hat{\mathbf{A}} & \hat{\mathbf{B}} \\ \hat{\mathbf{B}}^T & \hat{\mathbf{C}} \end{bmatrix}, \quad \hat{\mathbf{A}} = \begin{bmatrix} \mathbf{A} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{A} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{B} \end{bmatrix}, \quad \hat{\mathbf{B}} = \begin{bmatrix} \mathbf{D} \\ \mathbf{E} \\ \mathbf{0} \end{bmatrix}. \quad (\text{A.22})$$

If the grounding approach is applied, discussed in section 4.1, to produce nonsingular matrices then $\mathbf{F}(\boldsymbol{\theta})$ is nonsingular and the inverse can be expressed as

$$\mathbf{F}^{-1}(\boldsymbol{\theta}) = \begin{bmatrix} (\hat{\mathbf{A}} - \hat{\mathbf{B}}\mathbf{C}^{-1}\hat{\mathbf{B}}^T)^{-1} & -\hat{\mathbf{A}}^{-1}\hat{\mathbf{B}}(\mathbf{C} - \hat{\mathbf{B}}^T\hat{\mathbf{A}}^{-1}\hat{\mathbf{B}})^{-1} \\ -(\mathbf{C} - \hat{\mathbf{B}}^T\hat{\mathbf{A}}^{-1}\hat{\mathbf{B}})^{-1}\hat{\mathbf{B}}^T\hat{\mathbf{A}}^{-1} & (\mathbf{C} - \hat{\mathbf{B}}^T\hat{\mathbf{A}}^{-1}\hat{\mathbf{B}})^{-1} \end{bmatrix}, \quad (\text{A.23})$$

where the last block diagonal entity represent the CRLB for $\boldsymbol{\tau}$. Using (A.14)–(A.18), the Schur complement [14] of \mathbf{C} can be expressed as

$$\mathbf{C} - \hat{\mathbf{B}}^T\hat{\mathbf{A}}^{-1}\hat{\mathbf{B}} = c \left(\mathbf{I} - \frac{1}{M} \right), \quad (\text{A.24})$$

where

$$c = 2 \sum_{k=1}^K \frac{|j\omega_k \tilde{x}[k]|^2}{\sigma_w^2[k]}, \quad (\text{A.25})$$

and CRLB for $\boldsymbol{\tau}$ is given by

$$\mathbf{CRLB}(\boldsymbol{\tau}) = \frac{1}{c} (\mathbf{I} + 1). \quad (\text{A.26})$$

If $\hat{\tau}$ is obtained from an unbiased estimator, then

$$\text{var}(\hat{\tau}) \geq \frac{2}{c} = \left(\sum_{k=1}^K \frac{|\text{j}\omega_k \tilde{x}[k]|^2}{\sigma_w^2[k]} \right)^{-1}. \quad (\text{A.27})$$

If the pseudoinverse approach is applied, discussed in section 4.1, the matrix in (A.24) is singular, and an approximation of CRLB for τ is given by

$$\text{CRLB}(\tau) \approx (\mathbf{C} - \hat{\mathbf{B}}^T \hat{\mathbf{A}}^{-1} \hat{\mathbf{B}}^{-1})^\dagger = \frac{1}{c} \left(\mathbf{I} - \frac{1}{M} \right). \quad (\text{A.28})$$

References

- [1] Schoukens J and Pintelon R 2006 Estimation of nonparametric noise models *IEEE Instrumentation and Measurement Conf. (Sorrento, Italy)* pp 102–6
- [2] Björklund S and Ljung L 2004 A review of time-delay estimation techniques *Tech. Rep.* ISSN: 1400-3902, Linköping University
- [3] Grennberg A and Sandell M 1994 Estimation of subsample time delay differences in narrowband ultrasonic echoes using the Hilbert transform correlation *IEEE Trans. Ultrason. Ferroelectr. Freq. Control* **41** 588–95
- [4] Fischer B R and Medvedev A 1999 Laguerre domain estimation of time delays in narrowband ultrasonic echoes *14th Triennial IFAC World Congress (Beijing, China)* pp 361–6
- [5] Martinsson J, Häggglund F and Carlson J E 2006 Estimating the underlying signal waveform and synchronization jitter from repeated measurements *Proc. IEEE Int. Ultrason. Symp. (Vancouver, Canada)* pp 440–3
- [6] Kay S M 1993 *Fundamentals of Statistical Signal Processing: Estimation Theory* vol 1 (Englewood Cliffs, NJ: Prentice-Hall)
- [7] Brillinger D R 1981 *Time Series: Data Analysis and Theory (Holden-Day Series in Time Series Analysis)* (San Francisco: Holden-Day)
- [8] Schoukens J and Renneboog J 1986 Modeling the noise influence on the Fourier coefficients after a discrete Fourier transform *IEEE Trans. Instrum. Meas.* **20** 278–86
- [9] Ljung L 1985 On the estimation of transfer functions *Automatica* **21** 677–96
- [10] Pintelon R and Schoukens J 2001 *System Identification: A Frequency Domain Approach* (Piscataway, NJ: IEEE)
- [11] Stoica P and Marzetta T L 2001 Parameter estimation problems with singular information matrices *IEEE Trans. Signal Process.* **49** 87–90
- [12] Black P E 2007 ‘big O notation’ *Dictionary of Algorithms and Data Structures* <http://www.nist.gov/dads/HTML/bigOnotation.html>
- [13] Vandersteen G and Pintelon R 2000 Maximum likelihood estimator for jitter noise models *IEEE Trans. Instrum. Meas.* **49** 1282–4
- [14] Strang G 1986 *Linear Algebra and its Applications* 3rd edn (San Diego, CA: Harcourt Brace Jovanovich)