Signal Processing 145 (2018) 37–47

Contents lists available at ScienceDirect

Signal Processing

journal homepage: www.elsevier.com/locate/sigpro

Multi-dimensional grid-less estimation of saturated signals☆

F. Elvander*, J. Swärd, A. Jakobsson

Department of Mathematical Statistics, Lund University, P.O. Box 118, Lund SE-221 00, Sweden

ARTICLE INFO

Article history:
Received 17 August 2017
Revised 10 October 2017
Accepted 15 November 2017
Available online 16 November 2017

Keywords:
Atomic norm
De-clipping
Gridless reconstruction

ABSTRACT

This work proposes a multi-dimensional frequency and amplitude estimator tailored for noise corrupted signals that have been clipped. Formulated as a sparse reconstruction problem, the proposed algorithm estimates the signal parameters by solving an atomic norm minimization problem. The estimator also exploits the waveform information provided by the clipped samples, incorporated in the form of linear constraints that have been augmented by slack variables as to provide robustness to noise. Numerical examples indicate that the algorithm offers preferable performance as compared to methods not exploiting the saturated samples.

1. Introduction

Many forms of practical measurements suffer from clipping, for instance due to limitations in the dynamic span of the analog-to-digital (AD) converter, possibly necessitated by needs of resolution, or by additive interference offsetting the signal unexpectedly. In such cases, the measured signal is occasionally saturated at its minimum or maximum values, typically requiring these samples to be treated as missing. One may attempt to reconstruct such samples using various forms of interpolation or by using estimators of the relevant signal information that allow for missing samples (see, e.g., [1–4]). There have also been methods proposed for using gain masks in the sampling stage as to mitigate the effects of clipping [5], as well as post-processing methods for countering the harmonic distortion induced by clipping [6].

More recently, several reconstruction schemes exploiting an assumed signal sparsity have been proposed. In [7], Adler et al. extend the concept of image inpainting (see, e.g., [8]) to audio signals in order to reconstruct the clipped samples. In [9], Defraene et al. utilize a compressed sensing formulation, as well as exploit features of the human auditory system, in order to increase the perceived signal quality. Other approaches include iterative hard thresholding [10], greedy methods [11], smooth regularization [12], social sparsity exploiting temporal dependence [13], and non-negative matrix factorization [14], whereas theoretical recovery guarantees have been studied in [15]. The related field of estimation and reconstruction of 1-bit signals is also attracting interest (see, e.g., [16,17]). Such signals only retain the sign of the sampled analog waveform, which can be seen as an extreme form of clipping. The problem of signal reconstruction of more generally quantized measurements has been explored in [18].

In this work, we propose an algorithm that exploits the assumed *a priori* structure of the signals of interest. This structure may, for instance, be that the signal can be well modelled as a sum of decaying sinusoids, as is common in areas such as spectroscopy, or by some other well structured signal. By formulating an estimator of the unknown parameters detailing the assumed signal structure, taking into account both the available and the saturated samples, we propose a sparse reconstruction algorithm that is able to exploit the information available in the saturated samples, while still being robust to the presence of additive noise. Robustness against noise is achieved by not enforcing hard clipping constraints, i.e., the proposed estimator does not constrain the reconstructed waveform to saturate at precisely the same samples as the observed signal, as this would make the estimator vulnerable to amplitude bias. Instead, the clipping information is taken into account by adding linear constraints, relaxed using slack variables, allowing also the noise to cause saturation.

Assuming that the measured signal consists of relatively few signal components, the algorithm may be constructed as a sparse reconstruction problem using a signal dictionary formed using the assumed signal waveforms, taking into account the saturation information of the clipped samples. In order to allow the signal of interest to be formed over a continuous parameter space, we express the resulting optimization as an atomic norm minimization. The atomic norm has previously been successfully exploited to de-

☆ This work was supported in part by the Swedish Research Council the grant number is VR 2015-04148. Parts of the material herein has been submitted to the Asilomar Conference on Signals, Systems, and Computers, Pacific Grove, CA, November 2017.

* Corresponding author.

E-mail addresses: filipelv@maths.lth.se, filip.elvander@matstat.lu.se (F. Elvander), js@maths.lth.se (J. Swärd), aj@maths.lth.se (A. Jakobsson).

https://doi.org/10.1016/j.sigpro.2017.11.008
0165-1684/

velop estimators allowing for off-grid components (see, e.g., [19–21]). Here, we propose a similar formulation to exploit the structure of the assumed signal, while incorporating information of the saturated samples. We note that an approach reminiscent of ours was recently proposed in [22] for line spectrum estimation from 1-bit samples, although that work considered only noise-free signals. If considering signals with further structure, the herein proposed framework may be extended correspondingly. For instance, in audio applications, signals may often be well modeled as a sum of harmonically related sinusoids. For such signals, one may also exploit the expected harmonic structure by instead using the atomic norm framework developed in [23,24].

In summary, the proposed algorithm allows for an efficient exploitation of the 1-bit information present in saturated periodic one- or multi-dimensional signals, allowing for both accurate parameter estimation and signal reconstruction. Further extensions to non-periodic or non-stationary signals may be formed by generalizing the atomic norm formulation to such signals.

This paper is organized as follows. In Section 2, we state the considered problem of estimating clipped signals and present the proposed estimator, where the one-dimensional and multi-dimensional cases are considered in Sections 2.1 and 2.2, respectively. Section 3 evaluates the performance of the proposed method using simulation studies, also considering the impact of the choice of regularization parameters. Section 4 concludes upon the work.

2. Proposed estimator

In this section, we present the proposed estimator. We begin by initially presenting the one-dimensional (1-D) version for real-valued sinusoidal data, and then generalize the formulation to allow for both complex and multi-dimensional data.

2.1. One-dimensional case

To illustrate the proposed algorithm, we assume that the signal of interest, $\mathbf{y}^{\text{unclipped}}$, consists of N samples of a sum of K real-valued sinusoids corrupted by an additive Gaussian noise, such that

$$\mathbf{y}^{\text{unclipped}} = \mathbf{A}\mathbf{d} + \mathbf{e} \tag{1}$$

where $\mathbf{d} \in \mathbb{R}^{K\times 1}$ denotes the amplitude vector, $\mathbf{e}$ the additive noise, and

$$\mathbf{A} = \begin{bmatrix} \mathbf{a}_1 & \dots & \mathbf{a}_K \end{bmatrix}$$

$$\mathbf{a}_k = \begin{bmatrix} \cos(2\pi f_k t_1 + \phi_k) & \dots & \cos(2\pi f_k t_N + \phi_k) \end{bmatrix}^T$$

with f_k and ϕ_k denoting the kth frequency and phase, respectively, and t_k the time index of the kth sample. Here, we assume that the measured signal is $\mathbf{y}$, and that $\mathbf{y}^{\text{unclipped}}$ is unavailable. For 1-D real-valued signals, we define clipping as follows.

Definition 2.1. Clipping of real-valued data.

The nth sample of a real-valued 1-D signal, $\mathbf{y}_n^{\text{unclipped}}$, is subjected to clipping if

$$|\mathbf{y}_n^{\text{unclipped}}| > \gamma \tag{2}$$

and the corresponding measured value will be $\mathbf{y}_n = \gamma\,\text{sign}(\mathbf{y}_n^{\text{unclipped}})$, where $\gamma \geq 0$ is referred to as the clipping level or clipping limit. □

Using this definition, let Ω^-, Ω^+, and Ω denote the indices of $\mathbf{y}$ that are clipped from below, from above, and all the non-clipped indices of $\mathbf{y}$, respectively. Correspondingly, for any vector $\boldsymbol{b}$, let $\boldsymbol{b}_\Omega$ denote the vector constructed from $\boldsymbol{b}$ using only the elements corresponding to the indices in Ω. Thus, $\mathbf{y}_\Omega = \mathbf{y}_\Omega^{\text{unclipped}}$, $\mathbf{y}_{\Omega^-} = -\gamma\mathbf{1}$, and $\mathbf{y}_{\Omega^+} = \gamma\mathbf{1}$, where $\mathbf{1}$ is a vector of ones, of appropriate dimension. In order to reconstruct the signal of interest successfully, one needs to estimate the signal parameters, here the frequencies, amplitudes, and phases, as well as the model order, K, all which are assumed to be unknown. The typical way of dealing with the clipped samples in $\mathbf{y}$ is to treat these as missing data points, and simply omit them from the measurement vector. The unknown parameters, and the model order, are then estimated using a technique that allows for missing samples, such as, e.g., [25].

It is well known that dictionary techniques using a predefined grid suffer when the true parameters are not on the grid. To alleviate this problem, and also account for the missing samples, we here make use of an atomic norm formulation. Defining an atom set as $\mathcal{A} = \{\mathbf{a}(f,\phi):\ f \in [0,1],\ \ \phi \in [0,2\pi)\}$, with atoms $[\mathbf{a}(f,\phi)]_t = \cos(2\pi f t + \phi)$, a signal containing a sum over K sinusoids may be expressed as

$$\mathbf{y}^\star = \sum_{k=1}^{K} d_k \mathbf{a}(f_k,\phi_k) \tag{3}$$

The atomic norm for a signal $\mathbf{y}$ is defined as

$$\begin{aligned}||\mathbf{y}||_{\mathcal{A}} &= \inf\{t > 0:\ \mathbf{y} \in t\ \text{conv}(\mathcal{A})\} \\ &= \inf_{d_k \geq 0, \phi_k \in [0,2\pi), f_k \in [0,1]} \left\{ \sum_k d_k :\ \mathbf{y} = \sum_k d_k \mathbf{a}(f_k,\phi_k) \right\}\end{aligned}$$

where conv($\mathcal{A}$) denotes the convex hull of $\mathcal{A}$. This formulation may be interpreted as finding the sparsest linear combination of atoms that constitutes the signal. In [20], it was shown that the atomic norm denoising, i.e., the analogous problem with additive noise, may be expressed as the (computationally tractable) semidefinite program (SDP)

$$\begin{aligned} \underset{x,\mathbf{z},\mathbf{u}}{\text{minimize}} \quad & x + u_1 + \frac{1}{2}\|\mathbf{y}_\Omega - \mathbf{z}_\Omega\|_2^2 \\ \text{subject to} \quad & \begin{bmatrix} x & \mathbf{z}^H \\ \mathbf{z} & \mathbf{T}(\mathbf{u}) \end{bmatrix} \succeq 0 \\ & \mathbf{T}(\mathbf{u}) \in \mathbb{T}^{N\times N} \end{aligned} \tag{4}$$

where $\mathbb{T}^{N\times N}$ denotes the set of all Hermitian $N\times N$ Toeplitz matrices, with $\mathbf{T}(\mathbf{u})$ denoting the Toeplitz matrix with $\mathbf{u}$ on its first column. Here, u_1 denotes the first element of the vector $\mathbf{u}$. Since the problem in (4) is an SDP, it is also convex, and may as a result be computed using solvers, such as, e.g., CVX [26], yielding a computational complexity of $\mathcal{O}(N^3)$. The third term in (4) penalizes the difference between the observed samples for the measured signal and the optimization variable, $\mathbf{z}$, corresponding to the noise-free, non-clipped signal. Solving this optimization problem will yield a signal, $\mathbf{z}$, where the missing values have been estimated, a scalar, x, corresponding to the sum of the absolute values of the amplitudes, and the vector $\mathbf{u}$ that determines the Toeplitz matrix $\mathbf{T}(\mathbf{u})$, from which, using, e.g., a Vandermonde decomposition, the resulting frequency estimates may be found. This approach has been shown to be very efficient in both retrieving the missing samples, as well as estimating the frequencies [20,27]. However, it should be noted that the approach treats the clipped samples as missing, and is thus wasteful in the sense that the information that the measured signal is above (or below) the clipping limit is not incorporated in the optimization problem.

To alleviate this, we proceed to extend the minimization to also incorporate this information in the saturated samples. Clearly, since a clipped sample may not always indicate that the true wave form should be clipped, this should be taken into consideration when forming the optimization problem. This discrepancy appears when the true wave form is inside the measurable region, but the noise pushes the sample over (under) the saturation limit. To incorporate this effect, we introduce the variables $\boldsymbol{\epsilon}^+$ and $\boldsymbol{\epsilon}^-$. These