LUND UNIVERSITY

Multi-dimensional Grid-less Estimation of Saturated Signals

F. ELVANDER, J. SWÄRD, AND A. JAKOBSSON

Published in: Elsevier Signal Processing
doi:10.1016/j.sigpro.2017.11.008

Lund 2017

Mathematical Statistics
Centre for Mathematical Sciences
Lund University

Multi-dimensional Grid-less Estimation of Saturated Signals

Filip Elvander*, Johan Swärd*, and Andreas Jakobsson*

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.

Index Terms

Atomic norm, de-clipping, gridless reconstruction

I. 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], the authors extend the concept of image inpainting (see, e.g., [8]) to audio signals in order to reconstruct the clipped samples. In [9], the authors 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

This work was supported in part by the Swedish Research Council and the Royal Physiographic Society in Lund.

Parts of the material herein has been submitted to the Asilomar Conference on Signals, Systems, and Computers, Pacific Grove, CA, November 2017

*Dept. of Mathematical Statistics, Lund University, P.O. Box 118, SE-221 00 Lund, Sweden, email: {filipelv,js,aj}@maths.lth.se.

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 develop 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 II, 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 II-A and II-B, respectively. Section III evaluates the performance of the proposed method using simulation studies, also considering the impact of the choice of regularization parameters. Section V concludes upon the work.

II. 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.

A. 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 II.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. 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 capture the discrepancy between the observed and the true signal waveform for the samples saturated due to the additive noise, and should preferably be as small as possible in order to reduce the influence of this problem. More specifically, as the number of samples corresponding to noise-free signal values that lie within the measurable region $[-\gamma, \gamma]$ but still are clipped due to the additive noise is expected to be small, we expect $\boldsymbol{\epsilon}^+$ and $\boldsymbol{\epsilon}^-$ to be sparse, i.e., these will have few non-zero components. Incorporating these changes, the minimization may be expressed as

$$\begin{aligned} \underset{x,\mathbf{z},\boldsymbol{\epsilon},\mathbf{u}}{\text{minimize}} \quad & \mu\left(x + u_1\right) + \lambda \left\|\boldsymbol{\epsilon}\right\|_1 + \frac{1}{2}\left\|\mathbf{y}_\Omega - \mathbf{z}_\Omega\right\|_2^2 \\ \text{subject to} \quad & \begin{bmatrix} x & \mathbf{z}^H \\ \mathbf{z} & \mathbf{T}(\mathbf{u}) \end{bmatrix} \succeq 0 \\ & \mathbf{T}(\mathbf{u}) \in \mathbb{T}^{N\times N} \\ & \mathbf{z}_{\Omega^+} + \boldsymbol{\epsilon}^+ \geq \gamma \\ & \mathbf{z}_{\Omega^-} + \boldsymbol{\epsilon}^- \leq -\gamma \end{aligned} \tag{5}$$

where

$$\boldsymbol{\epsilon} = \begin{bmatrix} \boldsymbol{\epsilon}^{+T} & \boldsymbol{\epsilon}^{-T} \end{bmatrix}^T \tag{6}$$

and with μ and λ denoting user-defined parameters governing the denoising and the regularization of the $\boldsymbol{\epsilon}$, respectively. This formulation is a generalization of the atomic norm denoising problem found in [27], and, thus, the standard atomic norm formulation for the missing data case is retrieved by setting $\lambda = 0$. It may also be noted that by letting $\lambda \to \infty$, the reconstructed waveform is constrained to be consistent the clipping observed in the measured signal. This is reminiscent of the approach in [9], wherein a grid-based method incorporating hard clipping constraints was proposed. It is worth noting that the formulation used in (5) utilizes the ℓ_1-norm as a convex relaxation of the ℓ_0-pseudonorm in order to promote sparsity in $\boldsymbol{\epsilon}$. As an alternative, one might consider using the ℓ_2-norm, but this would then not promote sparsity, but rather slack variables $\boldsymbol{\epsilon}$ whose elements have small magnitudes. This would in turn result in a worse solution,

Algorithm 1 The proposed estimator

1: Require γ. Set λ and μ.
2: Solve (5) to obtain Toeplitz matrix $\mathbf{T}$.
3: Obtain frequency estimates from Vandermonde decomposition of $\mathbf{T}$.
4: Obtain amplitude estimates by solving (7).

as the use of an ℓ_2-norm would cause an upward amplitude bias for the reconstructed waveform $\mathbf{z}$, due to a trade-off between the terms of the objective function penalizing the magnitude of $\boldsymbol{\epsilon}$ and the difference $\mathbf{y}_\Omega - \mathbf{z}_\Omega$.

As before, the resulting frequency estimates are then found by using, e.g., a Vandermonde decomposition on $\mathbf{T}(\mathbf{u})$. Although estimates of the amplitudes, d_k, can also be obtained from such a Vandermonde decomposition, these estimates will be biased towards zero due to the regularization parameter μ. In order to refine the amplitude estimates, we therefore propose to additionally solve

$$\begin{aligned} \underset{\boldsymbol{\alpha},\boldsymbol{\beta},\boldsymbol{\epsilon}}{\text{minimize}} \quad & \frac{1}{2}\left\|\mathbf{y}_\Omega - \mathbf{Z}_\Omega\big(\hat{\mathbf{f}}\big)\mathbf{r}\right\|_2^2 + \lambda\left\|\boldsymbol{\epsilon}\right\|_1 \\ \text{subject to} \quad & \mathbf{Z}_{\Omega^+}\big(\hat{\mathbf{f}}\big)\mathbf{r} + \boldsymbol{\epsilon}^+ \geq \gamma \\ & \mathbf{Z}_{\Omega^-}\big(\hat{\mathbf{f}}\big)\mathbf{r} + \boldsymbol{\epsilon}^- \leq -\gamma \end{aligned} \tag{7}$$

where

$$\mathbf{Z}\big(\hat{\mathbf{f}}\big) = \begin{bmatrix} \mathbf{c}_1 & \dots & \mathbf{c}_K & \mathbf{s}_1 & \dots & \mathbf{s}_K \end{bmatrix} \tag{8}$$

$$\mathbf{c}_k = \Big[\cos(2\pi\hat{f}_k t_1) \ \dots \ \cos(2\pi\hat{f}_k t_N)\Big]^T \tag{9}$$

$$\mathbf{s}_k = \Big[\sin(2\pi\hat{f}_k t_1) \ \dots \ \sin(2\pi\hat{f}_k t_N)\Big]^T \tag{10}$$

$$\mathbf{r} = \begin{bmatrix} \boldsymbol{\alpha}^T & \boldsymbol{\beta}^T \end{bmatrix}^T . \tag{11}$$

Here, $\hat{\mathbf{f}}$ denotes the vector of frequency estimates obtained from the Vandermonde decomposition of $\mathbf{T}(\mathbf{u})$, and $\mathbf{Z}\big(\hat{\mathbf{f}}\big)$ is the dictionary matrix of cosine and sine atoms corresponding to these frequencies. Note that the use of both cosine and sine atoms allows for the representation of the non-linear phase parameter by combining two such atoms. Thus, the resulting optimization is a least squares (LS) problem, constrained to satisfy the clipping conditions of the observed signal, where the slack variable $\boldsymbol{\epsilon}$ is again exploited in order to provide robustness against noise. Using trigonometric identities, each amplitude estimate, $\hat{d}_k$, is then constructed from the minimizing vectors $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ as

$$\hat{d}_k = \sqrt{\alpha_k^2 + \beta_k^2}\,. \tag{12}$$

Similarly, the phase, ϕ_k is estimated as

$$\hat{\phi}_k = -\arctan\left(\frac{\beta_k}{\alpha_k}\right) . \tag{13}$$

The full algorithm is outlined in Algorithm 1.

B. D-dimensional case

In this section, we generalize the proposed estimator to allow for complex-valued data as well as expand the optimization problem to also be able to deal with D-dimensional data. We begin by defining what we in this paper mean by complex clipping. As before, let γ denote the clipping level. The real and the imaginary parts of the signal are typically treated separately, resulting in the following definition of complex clipping:

Definition II.2. Clipping of complex-valued data.
Sample n in a complex signal $\mathbf{y}^{\text{unclipped}}$ is subjected to clipping if either
(i) $|\Im\{\mathbf{y}_n^{\text{unclipped}}\}| > \gamma$
and will assume the value $\Im\{\mathbf{y}_n\} = \gamma \text{sign}(\Im\{\mathbf{y}_n^{\text{unclipped}}\})$, and/or
(ii) $|\Re\{\mathbf{y}_n^{\text{unclipped}}\}| > \gamma$
and will assume the value $\Re\{\mathbf{y}_n\} = \gamma \text{sign}(\Re\{\mathbf{y}_n^{\text{unclipped}}\})$, where $\Re$ and $\Im$ denote the real and the imaginary part, respectively. □

It is worth noting that Definition II.2 allows the real part of a sample to be correctly recorded, whereas the imaginary part is clipped, or vice versa. It also allows both the real- and the imaginary parts of the sample to be clipped, as well as being within $[-\gamma, \gamma]$ in both dimensions. It may be noted that no assumption of circularity is required for either of the cases.

In [28], the atomic norm framework was expanded to allow for two-dimensional data, and this was further generalized in [29] for the multi-dimensional case, where it also was shown that the Vandermonde decomposition that is used in the one-dimensional case to retrieve the frequency estimates has a multi-dimensional counterpart, and may thus be used for frequency retrieval for multi-dimensional data. We now present the D-dimensional version of the proposed estimation algorithm for clipped complex data.

Let $\boldsymbol{\mathcal{Y}}$ be the $N_1 \times N_2 \times \cdots \times N_D$ data tensor and let $\mathbf{y}$ be the vectorized version of $\boldsymbol{\mathcal{Y}}$ with size $N \times 1$, where $N = \prod_{n=1}^{D} N_n$. We here define the vectorization as being operated on the mode-1 matricization, or unfolding (see also [30]). Thus, in the two-dimensional (2-D) case, the vectorization reduces to stacking the columns of the 2-D data matrix. The order of the vectorization is not important as long as it is consistent. The atomic norm minimization problem taking the clipping information into account may then be formulated as

$$
\begin{aligned}
\underset{x,\mathbf{z},\mathbf{T},\boldsymbol{\epsilon}}{\text{minimize}} \quad & \mu\left(x + \text{tr}\{\mathbf{T}\}\right) + \lambda \left\|\boldsymbol{\epsilon}\right\|_1 \\
& + \frac{1}{2}\left\|\Re(\mathbf{y}_{\Omega_\Re}) - \Re(\mathbf{z}_{\Omega_\Re})\right\|_2^2 \\
& + \frac{1}{2}\left\|\Im(\mathbf{y}_{\Omega_\Im}) - \Im(\mathbf{z}_{\Omega_\Im})\right\|_2^2 \\
\text{subject to} \quad & \begin{bmatrix} x & \mathbf{z}^H \\ \mathbf{z} & \mathbf{T} \end{bmatrix} \succeq 0 \\
& \Re(\mathbf{z}_{\Omega_\Re^+}) + \boldsymbol{\epsilon}_\Re^+ \geq \gamma \\
& \Im(\mathbf{z}_{\Omega_\Im^+}) + \boldsymbol{\epsilon}_\Im^+ \geq \gamma \\
& \Re(\mathbf{z}_{\Omega_\Re^-}) + \boldsymbol{\epsilon}_\Re^- \leq -\gamma \\
& \Im(\mathbf{z}_{\Omega_\Im^-}) + \boldsymbol{\epsilon}_\Im^- \leq -\gamma
\end{aligned}
\tag{14}
$$

where

$$
\boldsymbol{\epsilon} = \begin{bmatrix} \boldsymbol{\epsilon}_\Re^{+T} & \boldsymbol{\epsilon}_\Im^{+T} & \boldsymbol{\epsilon}_\Re^{-T} & \boldsymbol{\epsilon}_\Im^{-T} \end{bmatrix}^T \tag{15}
$$

with $\Omega_{\Re}$ and $\Omega_{\Im}$ denoting the subset of the elements corresponding to the samples in $\mathbf{y}$ that have not been clipped in their real and imaginary parts, respectively. The notation $\Omega_{\Re}^{+}$ and $\Omega_{\Re}^{-}$ denote for the subset of elements corresponding to the samples in $\mathbf{y}$ that have their real part clipped with positive sign and negative sign, respectively, and similar for $\Omega_{\Im}^{+}$ and $\Omega_{\Im}^{-}$. As before, $\boldsymbol{\epsilon}$ acts as a slack-variable, allowing the clipping of the real and imaginary parts to be considered caused by the noise and not the true waveform. Furthermore, $\mathbf{T}$ is a D-level Toeplitz matrix (see [29] for a detailed definition). In the 2-D case, the 2-level $N_1N_2 \times N_1N_2$ Toeplitz matrix becomes

$$\mathbf{T} = \begin{bmatrix} \mathbf{T}_0 & \mathbf{T}_{-1} & \dots & \mathbf{T}_{-(N_1-1)} \\ \mathbf{T}_1 & \mathbf{T}_0 & \dots & \mathbf{T}_{-(N_1-2)} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{T}_{N_1-1} & \mathbf{T}_{N_1-2} & \dots & \mathbf{T}_0 \end{bmatrix} \tag{16}$$

where each

$$\mathbf{T}_{n_1} = \begin{bmatrix} z_{n_1,0} & z_{n_1,-1} & \dots & z_{n_1,-(N_2-1)} \\ z_{n_1,1} & z_{n_1,0} & \dots & z_{n_1,-(N_2-2)} \\ \vdots & \vdots & \ddots & \vdots \\ z_{n_1,N_2-1} & z_{n_1,N_2-2} & \dots & z_{n_1,0} \end{bmatrix} \tag{17}$$

for $n_1 = -(N_1-1), \dots, N_1-1$, is an $N_2 \times N_2$ Toeplitz matrix.

Similar to (5), the first term in (14) minimizes $\mathbf{z}^H\mathbf{T}^{(-1)}\mathbf{z}$ and controls the size of $\mathbf{T}$. The second term regularizes the slack variables $\boldsymbol{\epsilon}$ using the ℓ_1-norm. This corresponds to letting only a few of the elements in $\boldsymbol{\epsilon}$ to be active in the solution. The third and fourth terms bound the variable $\mathbf{z}$ to be close to the noisy signal, in a ℓ_2-norm sense, corresponding to a data fitting term. It is worth noting that the two last terms correspond to the proposed denoising term in [27]. Having obtained estimates of the set of frequencies from the optimal $\mathbf{T}$ estimate, one may then estimate the amplitudes by solving a LS problem analogous to that in (7), namely

$$\begin{aligned} \underset{\mathbf{d},\boldsymbol{\epsilon}}{\text{minimize}} \quad & \lambda \left\|\boldsymbol{\epsilon}\right\|_1 + \frac{1}{2}\left\|\Re(\mathbf{y}_{\Omega_{\Re}}) - \Re\left(\mathbf{Z}_{\Omega_{\Re}}(\hat{\mathbf{f}})\mathbf{d}\right)\right\|_2^2 \\ & + \frac{1}{2}\left\|\Im(\mathbf{y}_{\Omega_{\Im}}) - \Im\left(\mathbf{Z}_{\Omega_{\Im}}(\hat{\mathbf{f}})\mathbf{d}\right)\right\|_2^2 \\ \text{subject to} \quad & \Re\left(\mathbf{Z}_{\Omega_{\Re}^{+}}(\hat{\mathbf{f}})\mathbf{d}\right) + \boldsymbol{\epsilon}_{\Re}^{+} \geq \gamma \\ & \Im\left(\mathbf{Z}_{\Omega_{\Im}^{+}}(\hat{\mathbf{f}})\mathbf{d}\right) + \boldsymbol{\epsilon}_{\Im}^{+} \geq \gamma \\ & \Re\left(\mathbf{Z}_{\Omega_{\Re}^{-}}(\hat{\mathbf{f}})\mathbf{d}\right) + \boldsymbol{\epsilon}_{\Re}^{-} \leq -\gamma \\ & \Im\left(\mathbf{Z}_{\Omega_{\Im}^{-}}(\hat{\mathbf{f}})\mathbf{d}\right) + \boldsymbol{\epsilon}_{\Im}^{-} \leq -\gamma \end{aligned} \tag{18}$$

where $\mathbf{d}$ is the vector of amplitudes of the K D-dimensional sinusoids and $\mathbf{Z}(\hat{\mathbf{f}})$ is the $N \times K$ matrix defined as

$$\mathbf{Z}(\hat{\mathbf{f}}) = \mathbf{A}^{(D)}(\hat{\mathbf{f}}) \otimes \ldots \otimes \mathbf{A}^{(1)}(\hat{\mathbf{f}}) \tag{19}$$

where $\otimes$ denotes the Kronecker product and

$$\mathbf{A}^{(d)}(\hat{\mathbf{f}}) = \begin{bmatrix} \mathbf{a}_1^{(d)}(\hat{\mathbf{f}}) & \dots & \mathbf{a}_K^{(d)}(\hat{\mathbf{f}}) \end{bmatrix} \tag{20}$$

$$\mathbf{a}_k^{(d)}(\hat{\mathbf{f}}) = \begin{bmatrix} \mathrm{e}^{2i\pi\hat{f}_k^{(d)}t_1^{(d)}} & \dots & \mathrm{e}^{2i\pi\hat{f}_k^{(d)}t_N^{(d)}} \end{bmatrix}^T \tag{21}$$

III. Numerical evaluation

We proceed to examine the performance of the proposed algorithm, initially striving to estimate the frequencies and amplitudes of a clipped 1-D signal consisting of $K = 2$ sinusoids using (5) and (7), respectively. This is done by forming 500 Monte Carlo (MC) simulations of $N = 100$ samples, where in each simulation the (normalized) frequencies, f_1 and f_2, are drawn uniformly on the intervals $[0.08, 0.1]$ and $[0.11, 0.13]$, respectively. Furthermore, the amplitudes and phases are drawn uniformly on $[0.8, 1.2]$ and $[0, 2\pi)$, respectively. We examine the performance for two cases, the first varying the number of clipped samples while keeping the signal to noise ratio (SNR) fixed at 15 dB, where SNR is defined as

$$\text{SNR} = 10 \log_{10} \left(\frac{P_y}{\sigma^2} \right) \tag{22}$$

with P_y denoting the power of the true signal, and σ^2 denoting the noise variance. In the second case, the SNR is instead varying for the case of 30% clipped samples. The performance is measured using the sum of the root mean squared error (RMSE) for the frequencies, f_1 and f_2, as well as for the amplitudes, d_1 and d_2, where the RMSE for each component is defined as

$$\text{RMSE} = \sqrt{\frac{1}{M} \sum_{m=1}^{M} |\hat{\theta}_m - \theta|^2} \tag{23}$$

where θ_m is the true parameter, $\hat{\theta}_m$ is the mth MC estimate of that parameter, and M is the number of MC simulations. In order to illustrate the performance of the proposed method, we use the regularization parameters $\lambda = \mu = 1$ in all simulations. It may be noted that these values are in no way optimal in the sense of minimizing the estimation errors, as is shown in Figures 5 and 6 below, and are here rather selected to show that the algorithm is robust to the choice of parameters. As comparison, we also include the performance of the atomic norm minimization which only considers the unclipped samples, i.e., where the estimates are obtained by solving (5) without including the constraints, or equivalently, by setting $\lambda = 0$. For the obtained set of frequency estimates, the amplitudes are then estimated using (7). Also, for the amplitude estimates, we include comparisons with three least squares (LS) estimators that have been given oracle knowledge of the frequencies. The first of these LS estimators estimates $\mathbf{d}$ by solving (7) using the true frequencies. The second considers only the unclipped samples, i.e., solves (7) using $\lambda = 0$. Lastly, the third estimators uses hard clipping constraint, i.e., $\boldsymbol{\epsilon} = \mathbf{0}$, or equivalently, (7) is solved using $\lambda = +\infty$.

For the scenarios considering varying fractions of non-clipped samples, the RMSE for the frequency and amplitude estimates are presented in Figures 1 and 2, respectively. As can be seen from Figure 1, the proposed estimator is robust to the occurrence of clipped samples, and produces estimates whose accuracy is close to unaffected by the fraction of clipped samples. By comparison, the alternative estimator that only considers the non-clipped samples breaks down as the fraction of non-clipped samples decreases. As can be seen in Figure 2, the robustness of the proposed estimator then translates into improved amplitude estimates. In the figure, it can be seen that the three estimators utilizing (7) perform the best, as they can salvage information contained in the clipped samples; the LS estimator operating on only non-clipped samples suffers from the smaller samples size, whereas the constrained LS estimator using no slack variables suffers from an upward amplitude bias induced by the corrupting noise component, $\mathbf{e}$.

Similar conclusions may be drawn from Figures 3 and 4, showing the RMSE for the frequency and amplitude estimates for the scenario with varying SNR. Also in this case, the proposed

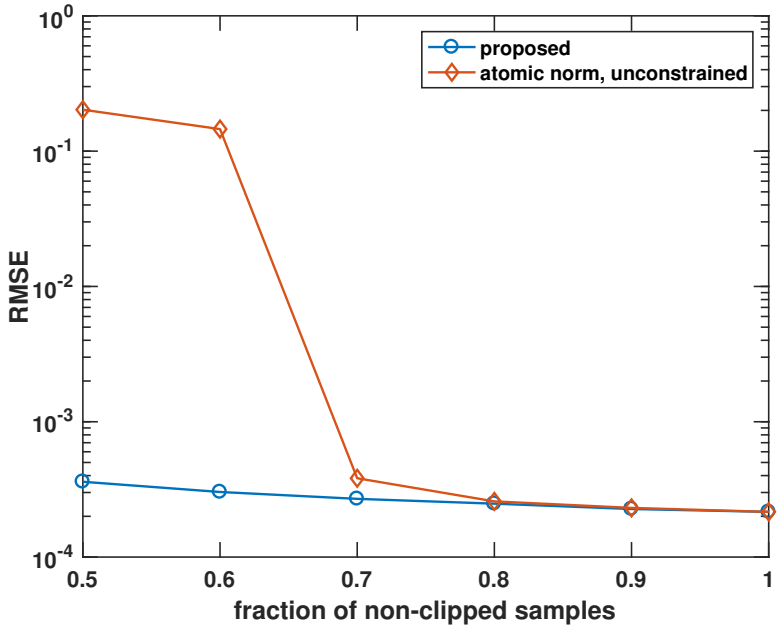

Fig. 1. RMSE for the frequency estimates produced by the proposed estimator, as well as the unconstrained atomic norm estimator, as a function of the fraction of unclipped data.

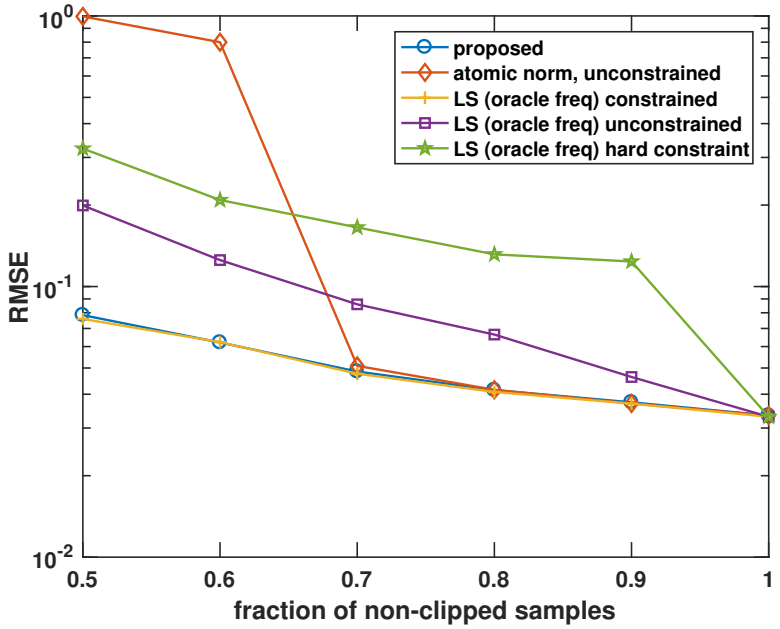

Fig. 2. RMSE for the amplitude estimates produced by the proposed estimator, the unconstrained atomic norm estimator, as well as the LS estimators, as a function of the fraction of unclipped data.

estimator displays greater robustness, and is less sensitive to noise than the estimators utilizing only the non-clipped samples. It may be noted from the figure that the RMSEs of the two estimators do not converge as the SNR increases; the proposed estimator consistently outperforms the estimator using only the non-clipped samples. Interestingly, for the highest SNR considered, i.e., 50 dB, the three LS estimators have identical performance. This is to be expected, as such a low noise setting renders the slack variable ϵ superfluous as the constraints are satisfied by

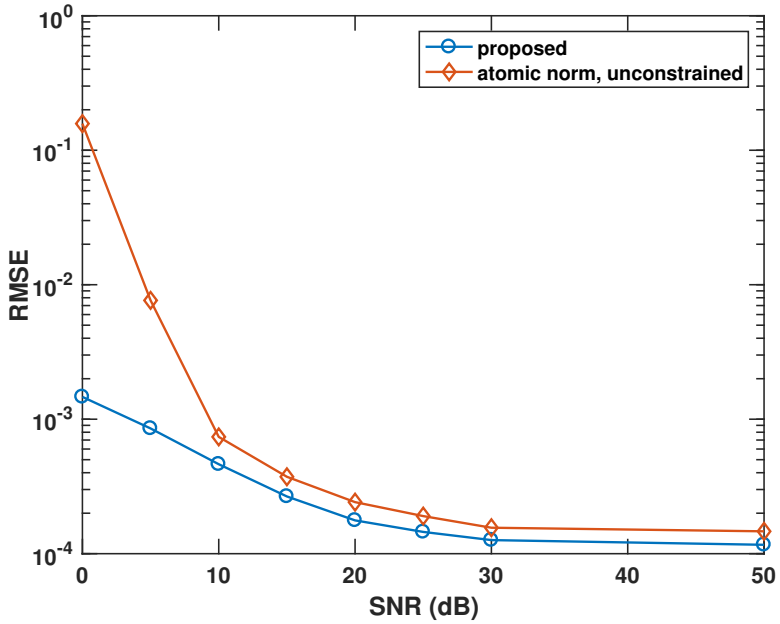

Fig. 3. RMSE for the frequency estimates produced by the proposed estimator, as well as the unconstrained atomic norm estimator, as a function of the SNR.

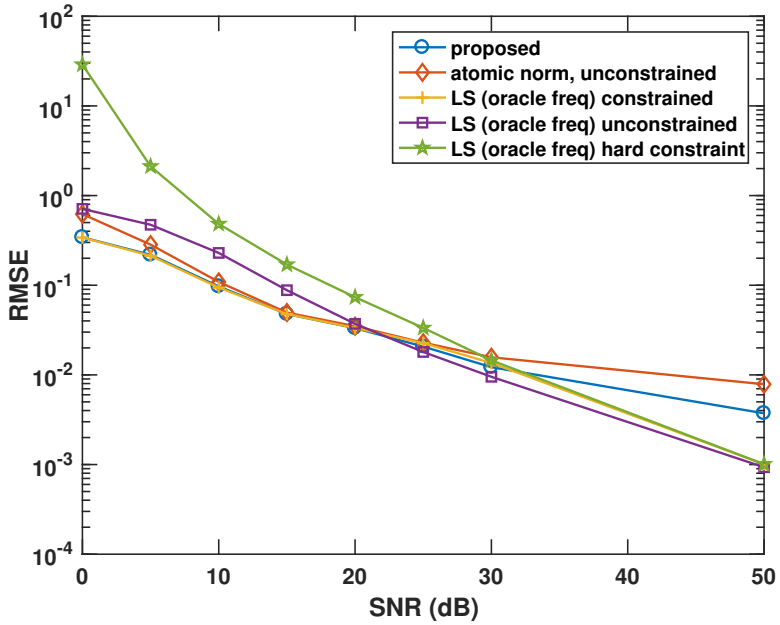

Fig. 4. RMSE for the amplitude estimates produced by the proposed estimator, the unconstrained atomic norm estimator, as well as the LS estimators, as a function of the fraction of unclipped data.

the uncorrupted waveform themselves. Also, as can be seen from the figure, for SNRs 25 and 30 dB, the LS estimator considering only non-clipped samples actually performs better than the estimators using (7), which is probably due to the slack variable ϵ introducing degrees of freedom that are not beneficial in such high, but not extreme, SNR settings. Furthermore, one may note that the atomic norm-based estimators perform worse than the LS estimators for the highest SNR settings, as they also have to estimate the frequencies.

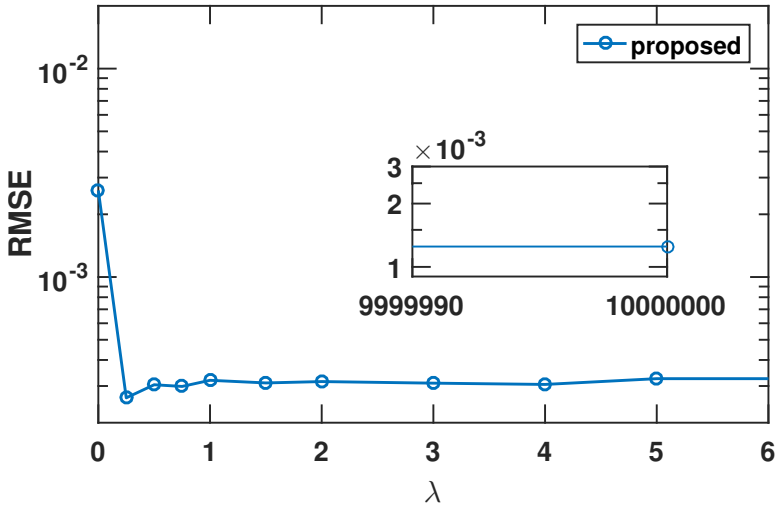

Fig. 5. RMSE for the frequency estimates produced by the proposed estimator, as a function of the regularization parameter λ for SNR 15 dB and 30% clipped samples.

Figures 5 and 6 show the effect on the RMSE for the frequency estimates when varying the regularization parameters λ and μ for the same simulation scenario as above, with SNR 15 dB and 30% clipped samples. When varying λ, we fix $\mu = 1$ and when varying μ, we fix $\lambda = 1$. As can be seen from Figure 5, setting $\lambda = 0$, corresponding to considering the clipped samples as missing, yields quite bad performance, whereas the accuracy of the estimates directly improve for $\lambda > 0$. Furthermore, the performance seems to be quite unaffected by the specific choice of λ. However, really large values of λ lead to performance degradation as $\lambda \to \infty$ corresponds to forcing the slack variable $\boldsymbol{\epsilon}$ to be zero, which in the case of noisy measurements negatively impacts the accuracy of the estimates. Indeed, for $\lambda = 10^7$, the RMSE for the frequency estimates was on the order of 10^{-3}, i.e., similar to the RMSE for $\lambda = 0$ and an order of magnitude larger than for $\lambda \in [1, 5]$. As can be seen from Figure 6, the performance of the proposed estimator is also fairly robust to varying values of μ, although setting $\mu = 0$ results in a rapid decrease in accuracy as this corresponds to not considering the sparsity of the solution, i.e., there is no penalty for non-sparse solutions. Conversely, increasing μ corresponds to increasing the sparsity of the solution at the expense of the fit to the observed waveform. As can be seen from the figure, performance indeed degrades when increasing μ to 4 and larger.

An illustration of the performance of the proposed estimator for signal reconstruction i shown in Figure 7. As can be seen, the proposed estimator is able to accurately reconstruct the unclipped and undistorted waveform, whereas the atomic norm estimator that consider clipped samples as missing incurs both a frequency and an amplitude bias.

We proceed by examining the performance of the estimator for multi-dimensional complex data using (14). All tests were done using 2-D data, with $N_1 = N_2 = 8$, containing two 2-D sinusoids with random phases, frequencies, and magnitudes. The data was on the form

$$\mathbf{y} = \mathbf{A}^{(2)} \otimes \mathbf{A}^{(1)} \boldsymbol{\alpha} + \mathbf{e} \tag{24}$$

where $\boldsymbol{\alpha}$ denotes the $K \times 1$ complex vector corresponding to the amplitudes, $\mathbf{e}$ the additive noise, $\otimes$ the Kronecker product, and the superscript $(\cdot)^{(d)}$ denotes the dimension d. The matrix

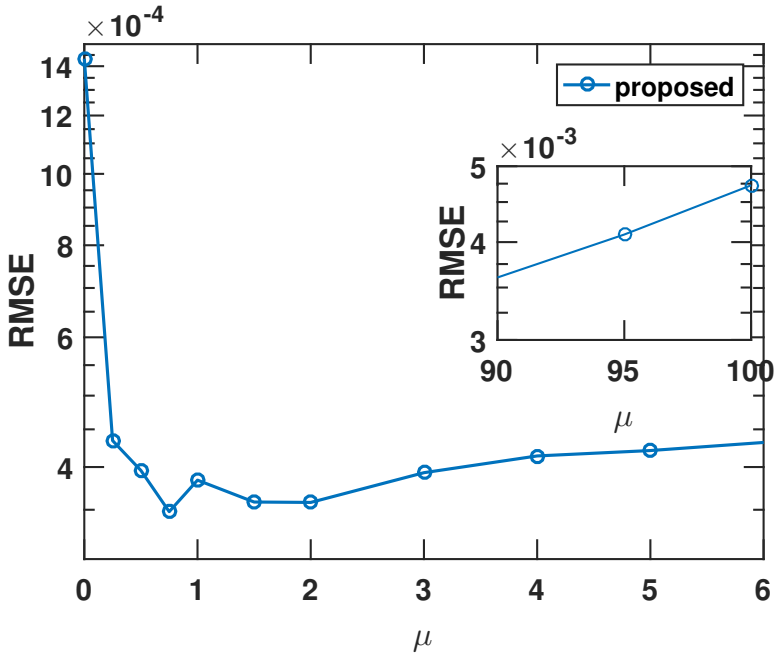

Fig. 6. RMSE for the frequency estimates produced by the proposed estimator, as a function of the regularization parameter μ for SNR 15 dB and 30% clipped samples.

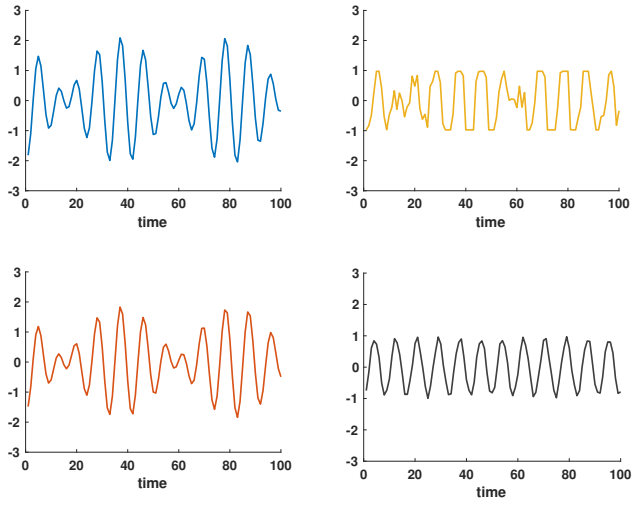

Fig. 7. Top left: the unclipped, undistorted waveform. Top right: the noise contaminated and clipped waveform. Bottom left: the reconstructed waveform as given by the proposed estimator. Bottom right: the reconstructed waveform as given by the atomic norm estimator, considering the clipped samples as missing.

$\mathbf{A}^{(d)}$, for $d = 1, 2$, are constructed as

$$\mathbf{A}^{(d)} = \left[\begin{array}{ccc} \mathbf{a}_1^{(d)} & \dots & \mathbf{a}_K^{(d)} \end{array} \right] \tag{25}$$

$$\mathbf{a}_k^{(d)} = \left[\begin{array}{ccc} \mathrm{e}^{2i\pi f_k^{(d)} t_1^{(d)}} & \dots & \mathrm{e}^{2i\pi f_k^{(d)} t_N^{(d)}} \end{array} \right]^T \tag{26}$$

The phase was sampled from $[0, 2\pi)$, whereas the (normalized) frequencies were selected uniformly from $[0, 1]$, but separated $1/N_d$ in each dimension. The magnitudes were uniformly selected between $[0.8, 1.2]$. In the first example, we investigate the RMSE on the frequency

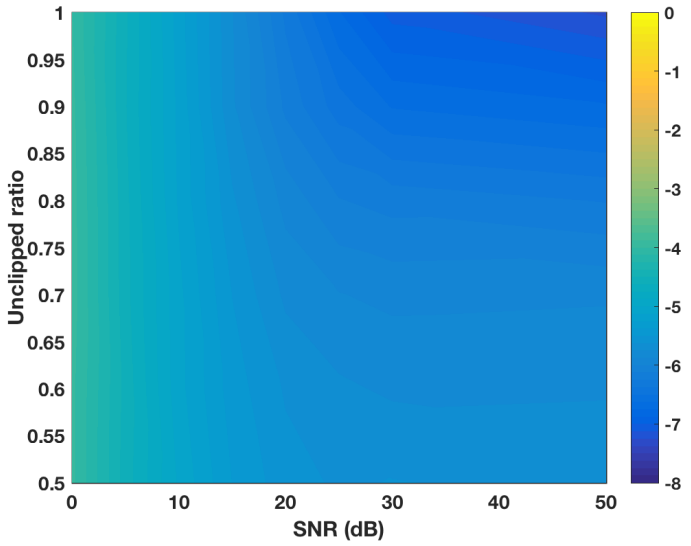

Fig. 8. RMSE for the frequency estimates produced by the proposed multi-dimensional estimator, as a function of the fraction of unclipped data and the SNR level. For clarity of presentation, the magnitudes are here given in log-scale.

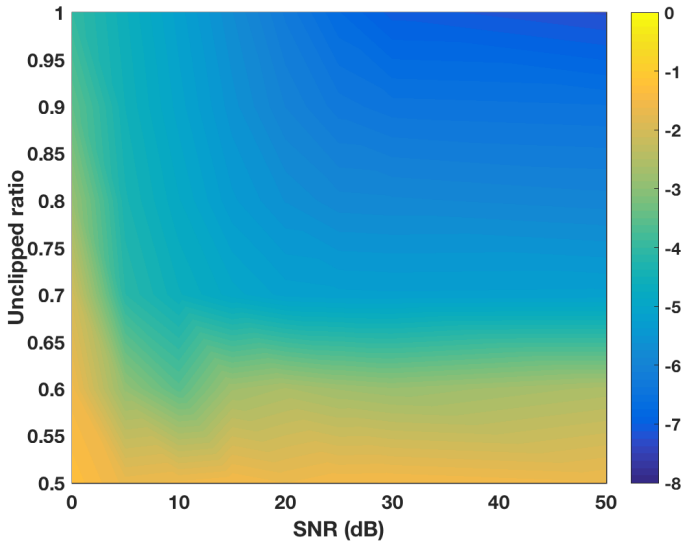

Fig. 9. RMSE for the frequency estimates produced by the unconstrained 2-D atomic norm estimator, as a function of the fraction of unclipped data and the SNR level. For clarity of presentation, the magnitudes are here given in log-scale.

estimation of the proposed method compared with the 2-D atomic norm proposed in [28], which treats the clipped samples as unknown. The frequency estimates were all retrieved using the MaPP estimator from [29]. The RMSE was evaluated for a range of different SNR levels and clipping ratios. The SNR ranged from 0 to 50, and the clipping ratios varied from 0.5 to 1. For each setting of SNR and clipping ratio, 500 Monte Carlo simulations were done and the user parameters were set to $\mu = \lambda = 1$.

Figures 8 and 9 show the performance of the proposed multi-dimensional estimator and the 2-D atomic norm, respectively. It can be seen from the figures that the results for the two estimators differ on two key points. First, the proposed estimator seems relatively unaffected by the clipping ratio; it is first when the SNR level drops to about 0 dB that any degradation

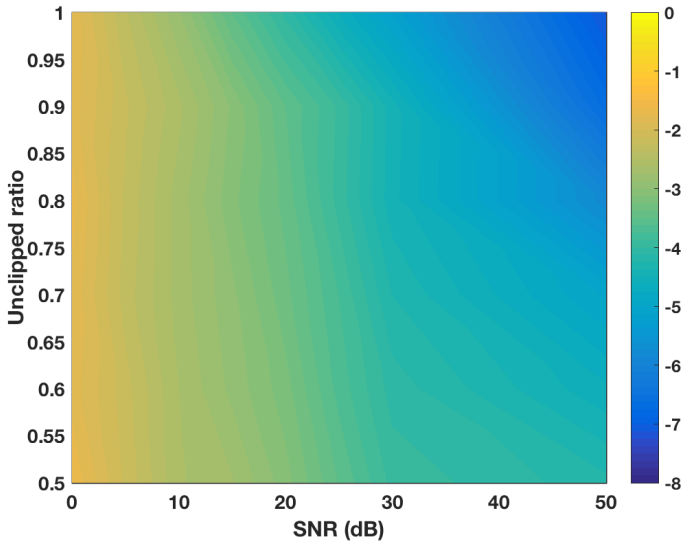

Fig. 10. RMSE for the amplitude estimates produced by the proposed multi-dimensional estimator as a function of the fraction of unclipped data and the SNR level. For clarity of presentation, the magnitudes are here given in log-scale.

starts to be noticeable. For the 2-D atomic norm estimator, the performance degrades both for low SNR levels and when the number of clipped samples increases. This figure corresponds well with the results for the 1-D case shown in Figures 1 and 3. Incorporating the information that the clipped sample should be above (below) the clipping threshold, as well as including the noise effect using ϵ, clearly shows its benefits.

As in the 1-D case, we proceed to examine the resulting RMSE for the amplitude estimates, comparing the proposed method to the 2-D atomic norm estimator, as well as the three LS estimators described above, which are given full knowledge about the true frequencies. Figures 10 and 11 show the RMSE of the amplitude estimates produced by solving (18) using the frequency estimate from the proposed method and the 2-D atomic norm estimator, respectively. Similar to the frequency estimation, the proposed method manages to better estimate the amplitudes. This is not surprising, as it also produced better frequency estimates. Figures 12, 13, and 14 show the three LS estimates with total knowledge of the true frequencies, corresponding to solving (18), with $\lambda = 1$, $\lambda = 0$, and $\lambda \to +\infty$, respectively. As can be seen from the figures, the proposed method in Figure 10 can be seen to produce similar results as the oracle LS estimator with $\lambda = 1$ in Figure 12. The unconstrained LS estimate, shown in Figure 13, seems to be more sensitive to the SNR, especially when the number of clipped samples is large. The final LS estimator, corresponding to the case when $\lambda \to +\infty$, seems even more sensitive to low SNR levels. We can conclude that the proposed method provides a better frequency estimate than the traditional 2-D atomic norm estimator. Furthermore, given these estimates, it is shown that the resulting amplitude estimates are almost as good as the one obtained if allowed full knowledge about the true frequencies.

IV. Acknowledgments

The authors would like to thank Dr. Zai Yang for providing the code for the MaPP method, as well as making us aware of several interesting references.

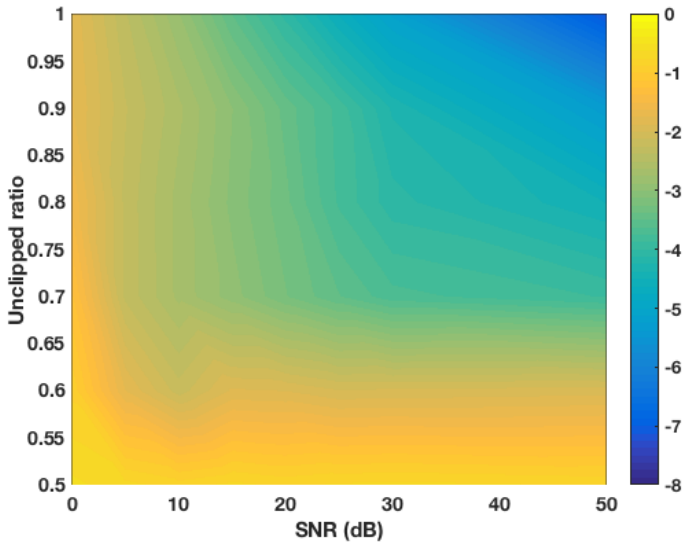

Fig. 11. RMSE for the amplitude estimates produced by the 2-D atomic norm estimator as a function of the fraction of unclipped data and the SNR level. For clarity of presentation, the magnitudes are here given in log-scale.

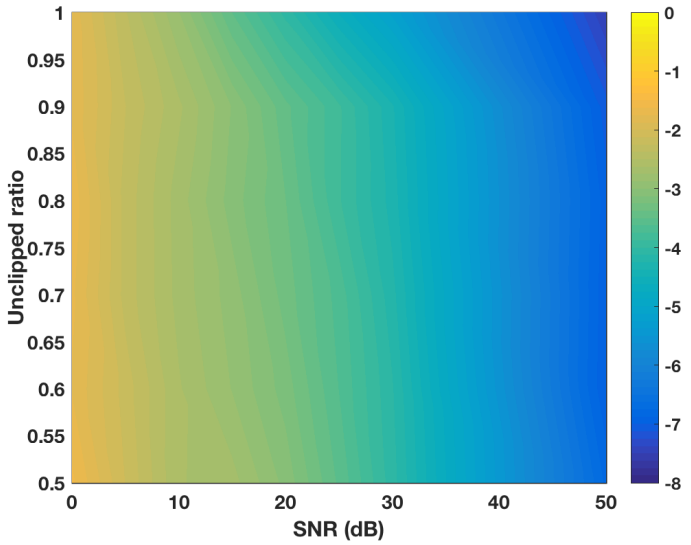

Fig. 12. RMSE for the amplitude estimates produced by the constrained ($\lambda = 1$) LS estimator, given oracle knowledge of the true frequencies, as a function of the fraction of unclipped data and the SNR level. For clarity of presentation, the magnitudes are here given in log-scale.

V. Conclusions

In this work, we have introduced a sparse reconstruction technique allowing for saturated signal samples. By exploiting the 1-bit information of the saturated samples, as well as allowing for the possibility that the noise causes the saturation of signals close to the saturation limits, the proposed estimator is shown to outperform alternative estimators not exploiting such information. The proposed estimator is formed using an atomic norm formulation allowing for a continuous parameter space, and does thus not suffer from the off-grid effects that often deteriorates dictionary based techniques.

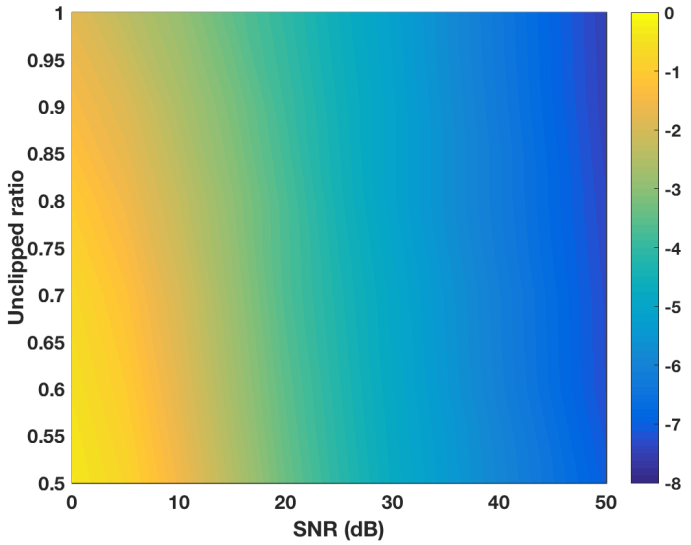

Fig. 13. RMSE for the amplitude estimates produced by the unconstrained ($\lambda = 0$) LS estimator, given oracle knowledge of the true frequencies, as a function of the fraction of unclipped data and the SNR level. For clarity of presentation, the magnitudes are here given in log-scale.

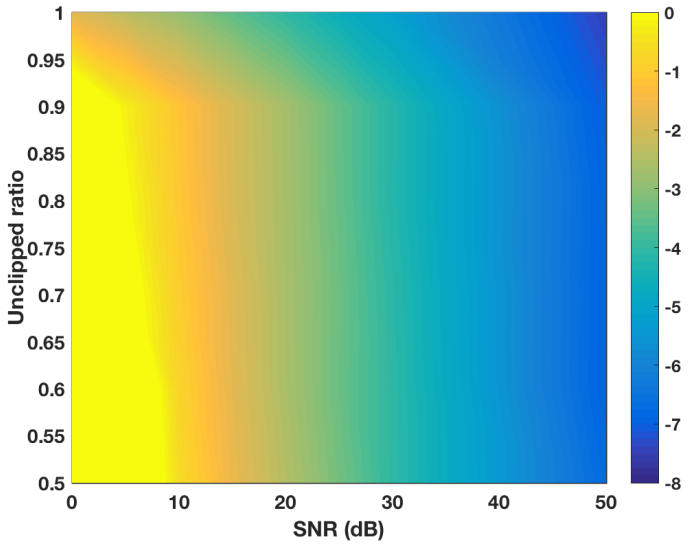

Fig. 14. RMSE for the amplitude estimates produced by the hard constrained ($\lambda \to +\infty$) LS estimator, given oracle knowledge of the true frequencies, as a function of the fraction of unclipped data and the SNR level. For clarity of presentation, the magnitudes are here given in log-scale.

REFERENCES

[1] J. S. Abel and J. O. Smith, "Restoring a Clipped Signal," in *16th IEEE Int. Conf. on Acoustics, Speech and Signal Processing*, Toronto, Canada, Apr. 14-17 1991, pp. 1745–1748.

[2] B. Porat and B. Friedlander, "ARMA spectral estimation of time series with missing observations," *IEEE Trans. Inform. Theory*, vol. 30, no. 4, pp. 601–602, July 1986.

[3] Y. Wang, J. Li, and P. Stoica, *Spectral Analysis of Signals - The Missing Data Case*, Morgan & Claypool, 2005.

[4] E. Gudmundson, P. Stoica, J. Li, A. Jakobsson, M. D. Rowe, J. A. S. Smith, and J. Ling, "Spectral Estimation of Irregularly Sampled Exponentially Decaying Signals with Applications to RF Spectroscopy," *J. Magn. Reson.*, vol. 203, no. 1, pp. 167–176, March 2010.

[5] P. Smaragdis, "Dynamic Range Extension Using Interleaved Gains," *IEEE Transactions on Acoustics Speech and Signal Processing*, vol. 17, no. 5, pp. 966–973, Jul. 2009.

[6] F. Esqueda, S. Bilbao, and V. Välimäki, "Aliasing Reduction in Clipped Signals," *IEEE Trans. Signa*, vol. 64, no. 20, pp. 5255–5267, Oct. 2016.

[7] A. Adler, V. Emiya, M. G. Jafari, M. Elad, R. Gribonval, and M. D. Plumbley, "Audio Inpainting," *IEEE Trans. Signal Process.*, vol. 20, no. 3, pp. 922–932, Mar. 2012.

[8] M. Elad, *Sparse and Redundant Representations*, Springer, 2010.

[9] B. Defraene, N. Mansour, S. De Hertogh, T. van Waterschoot, M. Diehl, and M. Moonen, "Declipping of Audio Signals Using Perceptual Compressed Sensing," *IEEE Trans. Acoust., Speech, Language Process.*, vol. 21, no. 12, pp. 2627–2637, Dec. 2013.

[10] S. Kitic, L. Jacques, N. Madhu, M. P. Hopwood, A. Spriet, and C. De Vleeschouwer, "Consistent Iterative Hard Thresholding for Signal Declipping," in *38th IEEE Int. Conf. on Acoustics, Speech and Signal Processing*, Vancouver, Canada, May 26-31 2013, pp. 5939–5943.

[11] A. J. Weinstein and M. B. Wakin, "Recovering a Clipped Signal in Sparseland," *Sampling Theory in Signal and Image Processing*, vol. 12, no. 1, pp. 55–69, Jan. 2013.

[12] M. J. Harvilla and R. M. Stern, "Efficient Audio Declipping using Regularized Least Squares," in *40th IEEE Int. Conf. on Acoustics, Speech, and Signal Processing*, Brisbane, Australia, Apr. 19-24 2015, pp. 221–225.

[13] K. Siedenburg, M. Kowalski, and M. Dörfler, "Audio Declipping with Social Sparsity," in *39th IEEE Int. Conf. on Acoustics, Speech and Signal Processing*, Florence, Italy, May 4-9 2014, pp. 1577–1581.

[14] Ç. Bilen, A. Ozerov, and P. Pérez, "Audio Declipping via Nonnegative Matrix Factorization," in *IEEE Workshop on Applications of Signal Processing to Audio and Acoustics*, Oct. 2015, pp. 1–5.

[15] C. Studer, P. Kuppinger, G. Pope, and H. Bolcskei, "Recovery of Sparsely Corrupted Signals," *IEEE Trans. Inf. Theor.*, vol. 58, no. 5, pp. 3115–3130, May 2012.

[16] L. Jacques, J. N. Laska, P. T. Boufounos, and R. G. Baraniuk, "Robust 1-Bit Compressive Sensing via Binary Stable Embeddings of Sparse Vectors," *IEEE Trans. Inf. Theor.*, vol. 59, no. 4, pp. 2082–2102, Apr. 2013.

[17] Y. Plan and R. Vershynin, "Robust 1-bit Compressed Sensing and Sparse Logistic Regression: A Convex Programming Approach," *IEEE Trans. Inf. Theor.*, vol. 59, no. 1, pp. 482–494, Jan. 2013.

[18] A. Zymnis, S. Boyd, and E. Candes, "Compressed Sensing With Quantized Measurements," *IEEE Signal Process. L*, vol. 17, no. 2, pp. 149–152, Feb. 2010.

[19] V. Chandrasekaran, B. Recht, P. A. Parrilo, and A. S. Willsky, "The Convex Geometry of Linear Inverse Problems," *Foundations of Computational Mathematics*, vol. 12, no. 6, pp. 805–849, Dec 2012.

[20] G. Tang, B. N. Bhaskar, P. Shah, and B. Recht, "Compressed Sensing Off the Grid," *IEEE Trans. Inform. Theory*, vol. 59, no. 11, pp. 7465–4790, Nov 2013.

[21] J. Swärd, S. I. Adalbjörnsson, and A. Jakobsson, "Generalized Sparse Covariance-based Estimation," *Elsevier Signal Processing*, 2017, Accepted for publication.

[22] C. Zhou, Z. Zhang, F. Liu, and B. Li, "Gridless compressive sensing method for line spectral estimation from 1-bit measurements," *Digit. Signal Process.*, vol. 60, pp. 152–162, Jan. 2017.

[23] T. L. Jensen and L. Vandenberghe, "Multi-pitch Estimation using Semidefinite Programming," in *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing*, New Orleans, LA, USA, March 2017.

[24] Z. Yang, "A Gridless Sparse Method for Super-Resolution of Harmonics," in *25th European Signal Processing Conference*, Aug 28 - Sep 2 2017.

[25] P. Stoica, J. Li, and J. Ling, "Missing Data Recovery via a Nonparametric Iterative Adaptive Approach," *IEEE Signal Process. Lett.*, vol. 16, no. 4, pp. 241–244, April 2009.

[26] Inc. CVX Research, "CVX: Matlab Software for Disciplined Convex Programming, version 2.0 beta," `http://cvxr.com/cvx`, Sept. 2012.

[27] Z. Yang and L. Xie, "On Gridless Sparse Methods for Line Spectral Estimation From Complete and Incomplete Data," *IEEE Trans. Signal Process.*, vol. 63, no. 12, pp. 3139–3153, June 2015.

[28] Y. Chi and Y. Chen, "Compressive Two-Dimensional Harmonic Retrieval via Atomic Norm Minimization," *IEEE Trans. Signal Process.*, vol. 63, no. 4, pp. 1030–1042, Feb 2015.

[29] Z. Yang, L. Xie, and P. Stoica, "Vandermonde Decomposition of Multilevel Toeplitz Matrices With Application to Multidimensional Super-Resolution," *IEEE Trans. Inf. Theor.*, vol. 62, no. 6, pp. 3685–3701, June 2016.

[30] T. G. Kolda and B. W. Bader, "Tensor Decompositions and Applications," *SIAM review*, vol. 51, no. 3, pp. 455–500, 2009.