RECOVERY OF SPARSE SIGNALS FROM AMPLITUDE-LIMITED SAMPLE SETS

Adam C. Polak, Marco F. Duarte, Robert W. Jackson and Dennis L. Goeckel

Department of Electrical and Computer Engineering, University of Massachusetts Amherst

ABSTRACT

Motivated by the compelling application of interference mitigation at wideband receivers in wireless communication and sensing systems, we consider the recovery of a frequency-sparse signal from samples of small magnitude. The standard ℓ_1-norm minimization results in an inadequate signal-dependent recovery performance, and hence we introduce three techniques to improve the quality of recovery. The performance of each of these three techniques is characterized through numerical simulations, from which we conclude that each of the proposed techniques show the promise of substantially improving recovery performance.

Index Terms— interference rejection, nonlinear distortion, compressive sensing, nonuniform sampling, cognitive radio.

1. INTRODUCTION

In this paper, we consider the compressive sensing (CS) recovery of frequency-sparse signals from samples taken only when the signal amplitude is small; that is, we attempt to recover the signal using samples with values within a range $[-\tau, \tau]$.

For motivation, consider the reception of a low-power message signal in the presence of high-power interferers at a receiver in a wideband wireless communications or sensing system. For instance, cognitive radio receivers feature an analog front-end that starts with an initial bandwidth wide enough for all supported applications and then down-selects in a reconfigurable way to a set of one or more relatively narrow sub-bands. The low noise amplifier (LNA), which generally comes first in the analog receiver chain, needs to be wideband as the amplification takes place before the sub-band selection. Despite decades of significant efforts by the microwave circuits community, RF LNAs feature inherent nonlinearities that become apparent at high power levels. Because of the large bandwidths employed by cognitive radio systems, interference is nearly always present, and the power of the interfering signals can often be multiple orders of magnitude higher than the power of the signal of interest as the interfering transmitters can be located much closer to the receiver than the transmitter whose message needs to be received. Nonlinearities in the first stages of the hardware allow large interferers to corrupt lower-level signals before the interferers can be de-selected, even if the interferers occupy frequencies different than the message frequencies. Conventional filtering or projection methods (e.g., [1]) after the front-end cannot remove this distortion due to the presence of the nonlinearity. This problem arises in other wideband applications ranging from environmental sensing to vehicle surveillance, etc.

Combining observations on circuits (devices are linear for small amplitudes) and the fact that the signals of interest are often frequency-sparse, we propose reconstruction of a sparse signal from small-amplitude samples that preserve linearity of the receiver's front-end. The small-amplitude signal sampling approach introduced in this paper falls in the field of signal dependent non-uniform sampling. Early and significant work on signal-dependent sampling was done by Logan [2], who established sufficient conditions for the zero-crossings of a signal to uniquely determine it. Existing practical recovery algorithms from the zero-crossing information are however known to be unstable. Boufounos and Baraniuk [3] introduced an additional signal sparsity assumption to gain robustness in signal recovery from zero-crossings. Recovery of frequency-sparse signals from non-zero level crossings as well as from multiple level crossings has been addressed recently by Sharma and Sreenivas [4]. Our work is significantly different from [3, 4]: instead of sampling non-uniformly at the times when the signal crosses predefined levels, we consider sampling the signal uniformly at high sampling rates and then selecting only the samples whose amplitudes are below a given threshold τ, while discarding potentially nonlinearly distorted samples with values that exceed the threshold.

Perhaps the prior contribution most closely related to our work is the recent independent work of [5], which also considers the recovery of frequency-sparse signals from a reduced set of samples. The sample subselection in [5] is driven by signal clipping; the resulting algorithms that account for clipping are similar to those we discuss here. However, in contrast to [5], we study the performance of CS with the proposed algorithms as a function of the amplitude of the threshold that samples must meet to be deemed suitable for signal recovery. Such threshold τ controls the fraction of samples that are used in recovery, cf. Figure 4. Our results show the impact that different options to leverage sample se-

This paper is based in part upon work supported by the National Science Foundation under grants CNS-0905349 and ECCS-1201835.

lection information during recovery have on CS performance.

Parallel work by a subset of the authors [6] describes an application where measurements are judged according to their quality. In that work, a low-level message signal was recovered from the output of a low-cost LNA at the same time it is being saturated by a two-toned interferer. The message's complex amplitude was recovered by simple projection under the assumption that the interferer and message signal frequencies are known. In contrast, this work assumes no knowledge of the signal's frequency components and uses CS techniques to retrieve the signal from the below-threshold sample set. We use ℓ_1-norm minimization recovery algorithms and show that, for many cases, successful recovery is possible even if only a small subset of low-amplitude samples is available. We also propose three alternative measurement and recovery approaches for the cases when recovery from small-amplitude samples leads to erroneous results: ℓ_1-norm minimization with additional nonlinearity constraints, iterative ℓ_1-norm minimization, and injection of a known signal to be subtracted after measuring. All three approaches can help to recover successfully.

The remainder of the paper is organized as follows. Section 2 presents the formal problem statement. Our three proposed approaches for signal recovery from small-amplitude samples are described in Section 3. Numerical results that verify the utility of the proposed recovery approaches are presented in Section 4. Finally, Section 5 concludes the paper.

2. PROBLEM STATEMENT

Consider a frequency sparse signal s, with unknown support, captured at the receiver. s needs to be recovered from a set of its small-amplitude samples. For that purpose, we solve a linear program

$$\hat{S} = \arg\min_{\bar{S}} ||\bar{S}||_1 \text{ s.t. } A\bar{S} = AS, \quad (1)$$

where $S = \mathcal{F}s$ is the Fourier representation of s, $A = M\mathcal{F}^H$ is a transformation matrix, M is a measurement matrix, and $\mathcal{F}^H$ is the Hermitian conjugate of the Fourier matrix $\mathcal{F}$. The probability of recovery error P_{err} is defined as the probability of the normalized recovery error

$$NRE = \|s - \hat{s}\|_2 / \|s\|_2 \quad (2)$$

being above a target value ρ.

The characteristics of the measurement matrix depend on the measurement scheme. In order to reduce the sampling rate of analog-to-digital converters (ADCs), individual measurements can be built as a linear combination of multiple time samples [7]. We assume that the measurement matrix M is built out of rows of an identity matrix that correspond to the indices of small-amplitude samples.

It is well known [8] that if the time samples are taken uniformly at random, then the recovery guarantees of compressive sensing are independent from the support of the frequency-sparse signal that needs to be recovered. When the uniform randomness of the sampling scheme is violated, recovery performance can become support-dependent [9]. The signal-dependent sampling approach considered in this work clearly violates the randomness assumptions of compressive sensing. Thus, it is expected that the recovery method (1) will have varying performance for signals with different supports of the same size, even if the size of the set of small-amplitude samples used for the recovery is the same. The next section presents approaches for enhancing the recoverability of the sparse signals from the amplitude limited sample sets for the cases when the standard recovery (1) leads to erroneous results.

3. APPROACHES

We will demonstrate in Section 4 that an ℓ_1-norm minimization (1) fitting only the values of the samples with amplitude less than τ will encounter ambiguities for some signals (i.e., signal-dependent performance). In this section, we consider three approaches for improving performance of the recovery. Methods described in Subsections 3.1 and 3.2 have been considered in independent work of [5], where signal clipping was driving small sample selection.

3.1. ℓ_1-norm minimization with inequality constraints

The recoverability of a frequency-sparse signal s from a small-amplitude sample set can be enhanced by taking into account additional information about s that becomes available while discarding large-amplitude samples. In particular, the indices of samples whose amplitudes exceed the predefined threshold τ are known, and this information can be exploited via inequality constraints in the linear program (1). The resulting optimization problem becomes

$$\hat{S} = \arg\min_{\bar{S}} ||\bar{S}||_1 \text{ s.t. } A\bar{S} = AS, \ |s(\underline{\Gamma})| > \tau \quad (3)$$

where $\underline{\Gamma}$ is a vector of time stamps of samples that have been discarded. The incorporation of these inequality constraints into standard CS recovery was suggested in [10], where unbounded measurement quantization errors caused by the saturation of ADCs were considered. Because of the relatively easy implementation of threshold comparators at the receiver, the extension of the constraints from (3) to multiple thresholds $\tau_n > \tau_{n-1} > \cdots > \tau$ is worth considering for our application of interest. The advantage of adding a second threshold and additional constraints to (3) of the form

$$A\bar{S} = AS, \quad \tau_2 \geq |s(\underline{\Gamma})| > \tau, \quad |s(\underline{\Gamma}_2)| > \tau_2, \quad (4)$$

will be studied in Section 4.

3.2. Iterative ℓ_1-norm minimization

A second approach for performance enhancement of (1) when only small-amplitude samples are available is iterative

ℓ_1-norm minimization. The minimization problem from (1) is a relaxation of a computationally intractable combinatorial problem of ℓ_0-"norm" minimization

$$\hat{S} = \arg\min_{\bar{S}} ||\bar{S}||_0 \text{ s.t. } A\bar{S} = AS. \tag{5}$$

The problem (1) can be solved efficiently and a body of existing work has shown there exist conditions under which the combinatorial problem (5) and its relaxation (1) are equivalent [11]. However, with the signal-dependent sampling scheme considered in this work, the conventional assumptions of CS are violated, which often leads to non-equivalence of (1) and (5). In [12], the authors introduced an iterative recovery algorithm consisting of a sequence of weighted ℓ_1-norm minimizations that promotes the sparsity of the result of the computationally tractable ℓ_1-norm minimization for the cases when (1) and (5) are not equivalent:

$$\hat{S}_i = \arg\min_{\bar{S}} \|C_i\bar{S}\|_1 \text{ s.t. } A\bar{S} = AS. \tag{6}$$

The diagonal matrix C_i in (6) contains positive weights that are updated in every iteration i to be inversely proportional to the values of the solution of the previous iteration:

$$C_{i+1}(k,k) = \frac{1}{|\hat{S}_i(k)| + \epsilon}, \tag{7}$$

with ϵ being a positive constant used for stability; one can set C_1 to be the identity. The algorithm is robust with respect to the choice of ϵ, which, as found empirically [12], should be set to a value smaller than the expected amplitudes of coefficients of the solution. The weights (7) promote sparsity of the solution, as the coefficients with small amplitude values contribute strongly to the weighted ℓ_1-norm $\|C_i\bar{S}\|_1$ in consecutive iterations. Thus, the final solution tends to consist of a small number of coefficients of highest significance.

As will be shown in Section 4, the iterative recovery algorithm (6) can lead to successful recovery of signals from small-amplitude samples when (1) and (5) are not equivalent due to a violation of the assumptions in CS, which leads to an incorrect solution during the first iteration of (6).

3.3. Injection of artificial interferers

As a third approach to enhance the performance of (1) when only small-amplitude sample sets are available, we consider injection of a known interferer to the signal s. This corresponds to the addition of a known interferer to the received signal before the LNA in a wideband receiver. After injection of the interferer, the samples of the signal $s' = s + i_{add}$ for which the amplitude exceeds τ are discarded. Since the interferer is known, the values of i_{add} for the samples retained are subtracted from the respective samples of s' and the resulting signal is used for recovery. If the injected interferer is uncorrelated with the signal s and the power of s and i_{add} are similar, then the sampling times get decorrelated from the frequency content of the signal s. The level of the randomization is higher as the injected interferer becomes more unstructured. Since in practice the threshold τ is a fixed value specified by the nonlinearity of the LNA, the injection of the interferer implies a reduction of the number of samples retained, due to the increased power in s' with respect to s. However, as will be shown in Section 4, the injection of a known interferer can significantly enhance recoverability, despite the penalty (decrease) on the number of samples caused by the increase of the power of the sampled signal s'.

4. SIMULATIONS

In this section, we present simulation results for CS recovery from sets of small-amplitude samples of frequency-sparse signals. Consider a discrete signal s of length $N = 751$ that consists of 10 tones. We consider two cases: (i) the tones are randomly located on the frequency axis; and (ii) the tones are positioned adjacently to build a single frequency band, located randomly on the frequency axis. For both considered cases (i) and (ii), the amplitudes and phases of the tones are chosen uniformly at random from respective ranges: $[0, 1]$ and $[0, 2\pi]$. We let the threshold τ vary over the range $[0, s_{max}]$, where s_{max} is the maximal amplitude of the signal s. We then discard all samples whose amplitudes are above the threshold τ and preserve the remaining samples as measurements. These measurements are used to solve the minimization problem (1) and to find the estimate of the message signal $\hat{s}(t)$ as described in Section 2. Figure 1 shows the probability of recovery error P_{err} of (1), defined as the probability that NRE from (2) is above $\rho = 3\%$, calculated over 100 trials for both considered cases (i) and (ii) as a function of the threshold τ. The figure shows that signal recovery is possible from fewer low-amplitude samples of s for the case (i) as compared to the case (ii).

For the case (ii), ℓ_1-norm minimization with inequality constraints and iterative ℓ_1-norm minimization (cf. Sections 3.1 and 3.2) were applied to improve recovery performance from small-amplitude samples. Figure 2 shows the probability of error P_{err} calculated over 50 trials as a function of the threshold τ for the case (ii) when (1), (3) and (6) were used. Five iterations were used for method (6); increasing the number of iterations above five did not lead to meaningful performance improvements. Figure 2 also shows P_{err} for the case (ii) when (3) was used with the additional threshold constraint (4). The second threshold τ_2 was used only when $\tau < 0.7 \cdot s_{max}$ and was set to $\tau_2 = 0.75 \cdot s_{max}$.

Finally, we study the recovery performance improvement achieved via injection of known interferers. Figure 3 shows P_{err} of (1), calculated over 100 trials as a function of the threshold τ for the case (ii), when three different types of known interferers were injected: 1 and 5 randomly positioned tones and a random Gaussian noise. The average power of the injected interferer was set to be equal to the power of the sig-

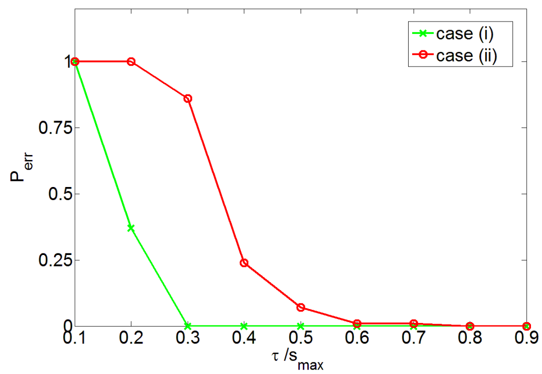

Fig. 1. Probability of recovery error for the cases (i) and (ii) for ℓ_1-norm minimization (1) as a function of the threshold τ.

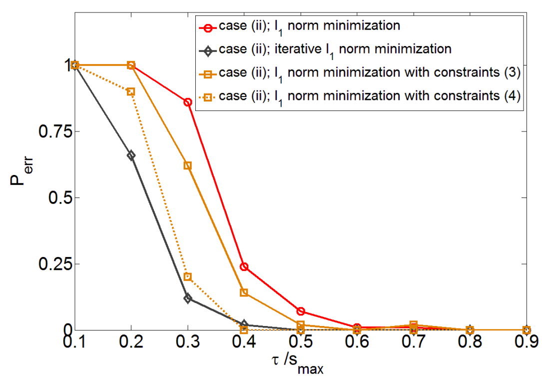

Fig. 2. Probability of recovery error for the case (ii) for ℓ_1-norm minimization (1), for iterative ℓ_1-norm minimization (6) and for constrained ℓ_1-norm minimization (3) and (4) as a function of the threshold τ.

nal s. For the considered frequency-sparse signal s, even a highly structured injected interferer (i.e., the sum of 5 tones) leads to significant decorrelation of the sampling times from the signal structure and thus significant recovery performance enhancement. Figure 3 shows that, even for a fixed threshold, which in practice is dictated by the characteristics of the nonlinearity of the LNA, adding an interferer i_{add} enhances recovery performance despite a reduction of the number of samples due to the average power increase of $s' = s + i_{add}$ with respect to s. Figure 4 shows the mapping between the threshold τ and the number of small-amplitude samples used for recovery, calculated over 100 trials, for different types of known injected interferers. It shows how a fraction of samples is lost due to the injected interferer, and how the choice of the interferer is causing only a small difference in the number of samples preserved.

5. CONCLUSIONS

Interference mitigation in wideband receivers is a critical component in many modern wireless communication and

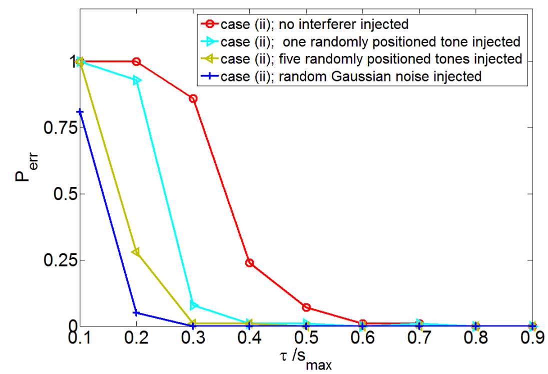

Fig. 3. Probability of recovery error for the case (ii) for ℓ_1-norm minimization (1) as a function of the threshold τ, for different types of known injected interferers.

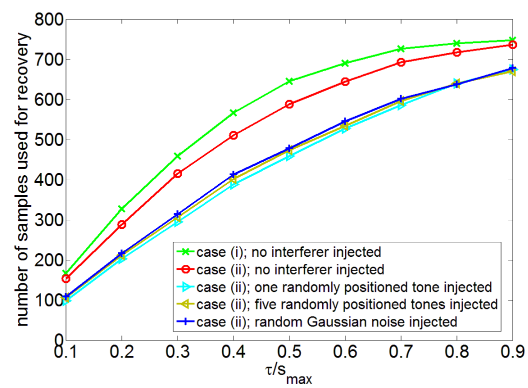

Fig. 4. Number of small-amplitude samples used for recovery as a function of the threshold τ for the case (i) and for the case (ii) for different types of known interferers injected.

sensing applications. Performing signal recovery that considers only samples of small magnitude (for which the RF front-end is linear) has been recently proposed by a subset of the authors. However, in this paper we have shown that the standard ℓ_1-norm minimization recovery performance becomes signal-dependent due to the correlation between the signal structure and the location of small-amplitude samples, thus motivating the exploration of enhanced CS sampling and recovery schemes. We have presented three such schemes that show significant improvement over the standard ℓ_1-norm minimization recovery from signal samples.

Future work will consider further algorithm development, the integration to existing CS analog-to-digital converters [7], and applications to interference mitigation in wideband receivers. Numerous challenges remain in this application, including the consideration of memory effects at the LNA output that can make sample timing matched to the linear region of the LNA challenging. Implementing the proposed interferer injection schemes also presents challenges such as the feasibility of accurate signal generation and removal.

6. REFERENCES

[1] M. Davenport, P. Boufounos, and R. Baraniuk, "Compressive domain interference cancellation," in *Workshop on Signal Processing with Adaptive Sparse Structured Representations*, April 2009.

[2] B.F. Logan Jr., "Information in the zero crossings of bandpass signals," *ATT Technical Journal*, vol. 56, pp. 487–510, Apr. 1977.

[3] P.T. Boufounos and R.G. Baraniuk, "Reconstructing sparse signals from their zero crossings," *IEEE International Conference on Acoustics, Speech and Signal Processing. ICASSP*, p. 3361 3364, Apr. 2008.

[4] N. Sharma and T.V. Sreenivas, "Sparse signal reconstruction based on signal dependent non-uniform samples," *IEEE International Conference on Acoustics, Speech and Signal Processing. ICASSP*, pp. 3453 – 3456, March 2012.

[5] A. J. Weinstein and M. B. Wakin, "Recovering a clipped signal in sparseland," *to appear in Sampling Theory in Signal and Image Processing.*, 2013.

[6] R. W. Jackson and M. Shusta, "Interference rejection by time selective sampling," in *submitted to 2013 European Microwave Conference*.

[7] J. A. Tropp, J. N. Laska, M. F. Duarte, J. K. Romberg, and R. G. Baraniuk, "Beyond Nyquist: Efficient sampling of sparse bandlimited signals," *IEEE Transactions on Information Theory*, vol. 56, no. 1, pp. 520–544, Jan. 2010.

[8] E. J. Candès and J. K. Romberg, "Sparsity and incoherence in compressive sampling," *Inverse Problems*, vol. 23, pp. 969–985, 2007.

[9] A.C. Polak, M. F. Duarte, and D. L. Goeckel, "Grouped incoherent measurements for compressive sensing," in *IEEE Statistical Signal Processing Workshop (SSP)*, Ann Arbor, MI, August 2012.

[10] J.N. Laska, P. T. Boufounos, M.A. Davenport, and R.G. Baraniuk, "Democracy in action: Quantization, saturation, and compressive sensing," *Appl. Comput. Harmon. Anal.*, vol. 31, no. 3, pp. 429–443, Nov. 2011.

[11] D. L. Donoho, "For most large underdetermined systems of linear equations the minimal ℓ_1 -norm solution is also the sparsest solution," *Commun. Pure Appl. Math.*, vol. 59, no. 6, pp. 797–829, 2006.

[12] E. Candès, M. B. Wakin, and S. P. Boyd, "Enhancing sparsity by reweighted l_1 minimization," *J. Fourier Anal. Appl.*, pp. 877–905, 2008.