

Greedy Sparse Signal Reconstruction From Sign Measurements

Petros T. Boufounos

Mitsubishi Electric Research Laboratories, Cambridge, MA 02139, *petrosb@merl.com*

Abstract—This paper presents Matched Sign Pursuit (MSP), a new greedy algorithm to perform sparse signal reconstruction from signs of signal measurements, i.e., measurements quantized to 1-bit. The algorithm combines the principle of consistent reconstruction with greedy sparse reconstruction. The resulting MSP algorithm has several advantages, both theoretical and practical, over previous approaches. Although the problem is not convex, the experimental performance of the algorithm is significantly better compared to reconstructing the signal by treating the quantized measurement as values. Our results demonstrate that combining the principle of consistency with a sparsity prior outperforms approaches that use only consistency or only sparsity priors.

I. INTRODUCTION

Compressive Sensing (CS) is a recently emerged signal acquisition technology for signals that are sparse or compressible. Specifically, compressive sensing leverages the structure that exists in signals of interest to enable acquisition at rates significantly lower than the Nyquist rate. To do so, CS uses randomized, linear, non-adaptive measurements, followed by non-linear reconstruction [1].

The non-linear processing to recover the signal from the measurements is a key hallmark of CS. Reducing the number of measurements below the Nyquist rate results to significant reconstruction ambiguity. Non-linear reconstruction is the workhorse of CS that takes into account the signal sparsity model and resolves this ambiguity. Until recently, however, very little attention has been paid in leveraging the power of non-linear reconstruction algorithms to incorporate accurate models of the measurement system in addition to models of the signal. Most of the literature to-date assumes a simple linear measurement system with additive noise.

Incorporating an accurate measurement system model can provide significant reconstruction improvements. This was first demonstrated in [2], [3] for quantized frame expansions. Imposing a correct model of coefficient quantization and enforcing reconstruction consistent with the model significantly outperformed linear reconstruction, which approximately models quantization as an additive uncorrelated noise process. In the context of CS, an explicit model of quantization first appeared in [4]. More recent work on quantization [5]–[7], saturation [8], and non-linear distortion [9], demonstrates the advantages of explicitly incorporating an accurate model of the measurement system during the reconstruction stage.

This paper revisits [4] and provides a new greedy CS reconstruction algorithm to recover signals from heavily quantized measurements. The present work capitalizes on a recent flurry

of research in greedy reconstruction algorithms, designed explicitly for CS reconstruction problems, which demonstrate fast performance and provable guarantees (e.g., [10], [11]). In particular, we modify CoSaMP [10], one of the most recent examples of such greedy algorithms, to produce the Matching Sign Pursuit (MSP), an algorithm designed to reconstruct signals from the sign of their measurements.

The MSP algorithm’s design addresses two issues specific to reconstruction from sign measurements, namely the loss of global amplitude information and the desire for consistent reconstruction. The former is a natural consequence of recording only the signs of measurements; any positive scalar multiple of the acquired signal will have measurements with the same sign. The latter is a consequence of incorporating an accurate acquisition and quantization system model. The new algorithm has similar goals as the ℓ_1 minimization on the unit ℓ_2 sphere, proposed in [4] to solve the same problem.

The MSP algorithm has several advantages over ℓ_1 minimization on the unit ℓ_2 sphere. First, it is straightforward to prove that, as the algorithm iterates and updates a signal estimate, any sparse enough estimate consistent with the measurements terminates the algorithm. This implies that if the signal estimate is equal to the signal that produced the measurements, the algorithm terminates with the correct solution. Second, the algorithm always returns a sparse solution by construction. Third, the MSP measures sparsity explicitly, using the number of non-zero coefficients, and not by a relaxation using the ℓ_1 norm. Thus, in contrast to [4], it allows the use of the ℓ_1 norm as a signal magnitude constraint instead of the ℓ_2 norm. Fourth, although the problem is non-convex, MSP has experimentally showed more consistent and robust performance compared to the ℓ_1 based algorithm in [4].

The remainder of this paper is structured as follows. The next section presents a brief background on CS, consistent reconstruction and 1-bit quantization. The goal of this section is to serve as a quick reference, setup the problem and establish the notation. Section III presents the MSP algorithm and reviews its features. Section IV presents some experimental results that validate our approach. Section V concludes with a brief discussion on the algorithm and the results.

II. BACKGROUND AND PROBLEM FORMULATION

A. Compressive Sensing

Compressive (or Compressed) Sensing (CS) is a recently emerged field in signal processing that enables signal acquisition using very few measurements compared to the signal

dimension, as long as the signal is sparse in some basis. Using CS, a signal $\mathbf{x} \in \mathbb{R}^N$ with only K non-zero coefficients can be recovered from only $M = O(K \log(N/K))$ linear non-adaptive measurements, compactly represented using

$$\mathbf{y} = \mathbf{A}\mathbf{x}, \mathbf{y} \in \mathbb{R}^M, \quad (1)$$

where $\mathbf{A} \in \mathbb{R}^{M \times N}$ is the measurement matrix, modeling the measurement system. To be able to recover all K -sparse signals \mathbf{x} it is necessary that \mathbf{A} does not map two distinct K -sparse signals to the same measurement vector, i.e., that for all K -sparse vectors $\mathbf{x}_1 \neq \mathbf{x}_2$

$$\mathbf{A}\mathbf{x}_1 - \mathbf{A}\mathbf{x}_2 = \mathbf{A}(\mathbf{x}_1 - \mathbf{x}_2) \neq \mathbf{0}. \quad (2)$$

Assuming (2), we can recover \mathbf{x} by solving

$$\hat{\mathbf{x}} = \arg \min_{\mathbf{x} \in \mathbb{R}^N} \|\mathbf{x}\|_0 \text{ subject to } \mathbf{y} = \mathbf{A}\mathbf{x}, \quad (3)$$

where the $\|\cdot\|_0$ counts the number of non-zero coefficients. This is an NP-hard problem [12] in general and becomes infeasible in high dimensions.

Computationally efficient exact signal recovery can be guaranteed if \mathbf{A} obeys a restricted isometry property (RIP) of order $2K$, i.e., if there exists a universal constant δ_{2K} such that for all $2K$ -sparse signals \mathbf{z}

$$(1 - \delta_{2K})\|\mathbf{z}\|_2^2 \leq \|\mathbf{A}\mathbf{z}\|_2^2 \leq (1 + \delta_{2K})\|\mathbf{z}\|_2^2. \quad (4)$$

If \mathbf{A} has a small RIP constant δ_{2k} , it approximately maintains ℓ_2 distances between K -sparse signals. In addition to efficient exact recovery, the RIP guarantees robustness to measurement noise and robustness to deviation from the strictly sparse signal model. Although verifying the RIP also has combinatorial complexity, a surprising result is that random matrices with sufficient number of rows can achieve small coherence and small RIP constants with overwhelming probability [1].

There are two fundamental approaches to reconstruction from CS measurements: convex optimization and greedy search algorithms. If the measurement matrix obeys the RIP with sufficiently small constant δ_{2K} and there is no measurement noise, it is possible to exactly recover signals from the measurement vector \mathbf{y} using the convex optimization [1]

$$\hat{\mathbf{x}} = \arg \min_{\mathbf{x} \in \mathbb{R}^N} \|\mathbf{x}\|_1 \text{ subject to } \mathbf{y} = \mathbf{A}\mathbf{x}. \quad (5)$$

The RIP further guarantees robustness to noise and stable recovery of compressible signals. Similarly, the RIP guarantees that greedy sparse reconstruction algorithms—which have origins at the Matching Pursuit [13]—robustly recover the signal. A recently emerged body of literature provides a variety of greedy algorithms with such guarantees, of which the Compressive Sampling Matching Pursuit (CoSaMP) [10] and the Subspace Pursuit [11] are the most recent examples.

B. Quantization and Consistent Reconstruction

Quantization—a necessary step for digitization—is usually modeled as additive measurement noise, denoted using \mathbf{n} :

$$\mathbf{y} = Q(\mathbf{A}\mathbf{x}) = \mathbf{A}\mathbf{x} + \mathbf{n}, \quad (6)$$

where $Q(\cdot)$ is the quantizer and \mathbf{n} is energy-limited to some ϵ depending on the quantization accuracy:

$$\|\mathbf{n}\|_2 \leq \epsilon. \quad (7)$$

For a uniform linear quantizer with quantization interval Δ , $\epsilon = \sqrt{M}\Delta/2$. Often the noise is further modeled as a zero-mean uniform independent and identically distributed random process, with standard deviation $E\{\|\mathbf{n}\|_2\} = \sqrt{M}\Delta^2/12$.

In the presence of norm-limited measurement noise such as quantization, it has been shown that robust reconstruction can be achieved by solving

$$\hat{\mathbf{x}} = \arg \min_{\mathbf{x}} \|\mathbf{x}\|_1 \text{ s.t. } \|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2 \leq \epsilon, \quad (8)$$

or by appropriately modifying the stopping criteria of greedy algorithms to accommodate for the uncertainty. In this case, the reconstruction error norm is bounded by $\|\mathbf{x} - \hat{\mathbf{x}}\|_2 \leq C\epsilon$, where the constant C depends on the properties of the measurement system \mathbf{A} and the reconstruction approach chosen, but not on the signal [1], [10].

An alternative to treating quantization as norm-limited noise is to enforce consistent reconstruction, which requires that the solution is consistent with all our knowledge about the signal and the measurement process. In the case of quantized measurements, this implies that if the reconstructed signal is re-measured using the measurement system \mathbf{A} and quantized at the same accuracy then the quantized measurements should be exactly the same as the original measurements used to reconstruct the signal. In [2], [3] it is shown that consistent reconstruction significantly improves the reconstruction performance in quantized frame representations.

Although in the general case of norm-limited measurement noise the reconstruction in (8) is consistent with the measurements, this is not the case if the measurement noise is due to quantization. Specifically, in the case of uniform linear quantization, all noise components have magnitude $|n_i| \leq \Delta/2$, which implies that consistent reconstruction should produce a signal that satisfies

$$\|\mathbf{A}\hat{\mathbf{x}} - \mathbf{y}\|_\infty \leq \frac{\Delta}{2}, \quad (9)$$

where $\|\mathbf{y}\|_\infty$ denotes the maximum magnitude of any element in \mathbf{y} . Several reconstruction approaches designed for quantized measurements have appeared in the literature [4]–[7], enforcing consistent reconstruction either explicitly [6] or implicitly by approximating the constraint with an ℓ_p , $p > 2$ norm [7].

C. 1-bit Compressive Sensing

In the case of 1-bit quantization, the quantizer is most often implemented as a comparator to zero. In particular, we consider linear measurements of a signal $\mathbf{x} \in \mathbb{R}^N$ using inner products with the rows of a measurement matrix $\mathbf{A} \in \mathbb{R}^{M \times N}$. Each measurement is quantized to 1-bit by recording only its sign:

$$\mathbf{y} = \text{sign}(\mathbf{A}\mathbf{x}), \quad (10)$$

where the function $\text{sign}(y_i) = y_i/|y_i|$ is applied element-wise to the measurement vector. Since each measurement only

uses one bit, the number of measurements is the same as the number of bits used for the signal acquisition. Without loss of generality we assume the vector \mathbf{x} is K -sparse in the canonical basis, i.e., it has at most K non-zero elements.

In this case, the measurement error can be unbounded. Reconstruction using standard CS approaches, treating the signs of the measurements are just ± 1 values, does not perform as well as consistent reconstruction [4].

Consistent reconstruction requires that measurements of the reconstructed signal \hat{x} should have the same sign as the measurements obtained from the measurement system:

$$(\mathbf{A}\hat{\mathbf{x}})_i \begin{matrix} y_i=+1 \\ \geq \\ y_i=-1 \end{matrix} 0 \quad (11)$$

$$\Leftrightarrow y_i \times \text{sign}(\mathbf{A}\hat{\mathbf{x}})_i \geq 0. \quad (12)$$

where $y_i = \pm 1$ is the quantized measurement

If \mathbf{x} is consistent with the measurements then so is $a\mathbf{x}$ for all $a \geq 0$. Since 0 is the sparsest vector possible, a reconstruction algorithm that only requires consistency with the measurements will drive the solution to $\hat{\mathbf{x}} = 0$. To enforce reconstruction at a non-trivial solution we artificially resolve this fundamental amplitude ambiguity. Specifically, we impose an energy constraint that the reconstructed signal lies on the unit ℓ_2 -sphere, i.e., that

$$\|\hat{\mathbf{x}}\|_2 = 1. \quad (13)$$

Note that this constraint significantly reduces the search space. This reduction plays an important role in improving the reconstruction performance.

Sparse recovery from sign measurements was first examined in [4], combining the principle of consistent reconstruction with ℓ_1 minimization on the unit ℓ_2 sphere to reconstruct the signal. The present work, instead, proposes an iterative greedy algorithm, the Matching Sign Pursuit (MSP). The MSP also uses the principle of consistent reconstruction but instead of an ℓ_1 minimization it performs a greedy search similar to the Compressive Sampling Matching Pursuit (CoSaMP) [10] and the Subspace Pursuit [11] approaches.

Consistent reconstruction from sign measurements can also be considered as the limit of recovering from saturated measurements, considered in [8]. As the signal gain increases, so does the saturation rate. The majority of the measurements then saturate. At the limit, the only information that needs to be encoded is if the saturation is positive or negative. The connections between [8] and the MSP algorithm are deferred to a future publication.

III. GREEDY SIGNAL RECONSTRUCTION

A. The MSP Algorithm

The MSP algorithm is described in Algorithm 1. Similarly to most greedy algorithms, after each iteration l the algorithm produces an estimate $\hat{\mathbf{x}}^l$ of the measured signal and its sparse support. Each iteration refines the estimate until the algorithm converges. The refinement first attempts to identify which constraints are violated and how severely. It then attempts to identify the support of the signal that is most effective in

reducing the sign violations. Finally it reconstructs the signal on that support to be as consistent with the data as possible and iterates with the new signal estimate.

The steps of MSP are very similar to CoSaMP. Specifically, in steps 3 and 4 of each iteration the MSP computes a set of measurements using the previous signal estimate $\hat{\mathbf{x}}^{l-1}$ and determines which measurements have signs inconsistent with the data. Based on the inconsistent measurements, step 5 computes a proxy, \mathbf{s}^l , for the part of the signal not explained by $\hat{\mathbf{x}}^{l-1}$. Step 6 combines the support of the $2K$ largest components of \mathbf{s}^l with the support of $\hat{\mathbf{x}}^{l-1}$ to produce a candidate support T over which step 7 performs consistent reconstruction. The reconstructed signal \mathbf{b}^l has support at most $3K$ components, which is truncated down to K components in step 8 and normalized to produce the signal estimate $\hat{\mathbf{x}}^l$. The algorithm terminates once a convergence criterion is satisfied. In our simulations we terminate once the change of \hat{x}^l from iteration to iteration becomes very small.

B. Normalization

As noted in [4] and Sec. II-C, the measurement process in (10) eliminates any magnitude information and a global amplitude constraint should be imposed. Similarly to [4], the MSP enforces an ℓ_2 constraint. The signal is reconstructed such that $\|\hat{\mathbf{x}}\|_2 = 1$ by re-normalizing the signal estimate after each iteration.

Of course, alternative normalizations are possible. Depending on the application or additional information on the signal, an ℓ_1 or an ℓ_∞ normalization might be more appropriate. Alternatively, it might be necessary to normalize the signal such that its projection to a specific measurement vector has a predetermined value. These alternatives can be implemented by appropriately modifying steps 7 and 8 of the algorithm. Note that in the algorithm presented in [4] it is not possible to enforce an ℓ_1 amplitude constraint; the ℓ_1 norm of the reconstructed signal is used as a proxy for the signal sparsity.

C. Enforcing Consistency

Consistency is enforced in two separate places in the algorithm. First, in steps 4 and 5 the sign violations are identified and the correction signal proxy is computed based on those violations and their amplitude. Second, step 7 performs consistent reconstruction over the candidate support, as opposed to the standard least squares inversion performed by CoSaMP.

To enforce consistency, the MSP attempts to minimize the norm of the sign violations in the measurement domain, i.e.,

$$\|(\text{diag}(\mathbf{y}) \mathbf{A}\hat{\mathbf{x}})^-\|_2, \quad (14)$$

where $\text{diag}(\mathbf{y})$ denotes a matrix with the measurement signs in the main diagonal, and $(x)^- = (x - |x|)/2$ is the ‘‘negative part’’ function, applied element-wise to each vector element (i.e., sets all positive elements of the vector to zero). This is equivalent to penalizing only the sign violations using a quadratic penalty, similar to [4], [8]. When the reconstruction is exactly consistent, the penalty is zero. The deviation away from zero provides a measure of inconsistency.

Algorithm 1 The Matching Sign Pursuit (MSP) Algorithm

Inputs: Measurement signs $\mathbf{y} \in \{\pm 1\}^N$, Measurement matrix \mathbf{A} , Signal sparsity K .

Initialization: Iteration count $l = 0$, Initial K -sparse estimate $\hat{\mathbf{x}}^0$

1: **while** not converged **do**

2: *Increase iteration count:*

$$l \leftarrow l + 1$$

3: *Compute measurements estimate:*

$$\tilde{\mathbf{y}}^l = \text{diag}(\mathbf{y}) \mathbf{A} \hat{\mathbf{x}}^{l-1},$$

where $\text{diag}(\mathbf{y})$ is a diagonal matrix with the measured signs. Consistency is achieved when $\tilde{\mathbf{y}}$ has positive elements.

4: *Identify sign violations:*

$$\mathbf{r}^l = (\tilde{\mathbf{y}}^l)^-,$$

where the function $(\cdot)^-$ is the element-wise negative part of a vector (i.e., sets all the positive elements to 0).

5: *Compute correction signal proxy:*

$$\mathbf{s}^l = \mathbf{A}^T \text{diag}(\mathbf{y}) \mathbf{r}^l$$

6: *Identify correction support:*

$$T^l = \text{supp}(\mathbf{s}^l|_{2K}) \cup \text{supp}(\hat{\mathbf{x}}^{l-1}),$$

where $\text{supp}(\cdot)$ determines the support set of a vector, and $\mathbf{x}|_K$ selects the K components of \mathbf{x} with the largest magnitude, setting all others to 0.

7: *Perform Consistent Reconstruction over chosen support:*

$$\mathbf{b}^l|_T = \arg \min_{\mathbf{x}} \left\| (\text{diag}(\mathbf{y}) \mathbf{A} \mathbf{x})^- \right\|_2^2, \text{ s.t. } \|\mathbf{x}\|_2 = 1 \text{ and } \mathbf{x}|_{T^c} = 0$$

i.e., minimize the squared norm of the negative (inconsistent) part of the vector subject to the normalization constraint.

The optimization is only over the support T . This (also non-convex) optimization is performed by iterating a gradient descent step followed by a normalization step until convergence, as described in Sec. III-C.

8: *Truncate, normalize, and update the estimate:*

$$\hat{\mathbf{x}}^l = \frac{\mathbf{b}^l|_K}{\|\mathbf{b}^l|_K\|_2}$$

Output: $\hat{\mathbf{x}}^l$

In step 7 of the algorithm we make the same choice for enforcing consistency: the cost function (14) is minimized on the unit ℓ_2 sphere, but only over the support T . In our implementation we perform this minimization by iterating a simple gradient descent step on the cost function followed by renormalization of the estimate to be on the unit sphere. Since the optimization on the unit sphere is not convex, the initial seed is important. We chose the initial seed to be the current estimate $\hat{\mathbf{x}}^{l-1}$ from the previous iteration of the MSP. Note that if the seed is consistent with the measurements, the cost is zero and step 7 terminates without modifying the estimate.

This choice of quadratic penalty provides a soft consistency metric. The alternative is to enforce hard consistency constraints, as in [2], [3], [6], i.e. not allow any constraint violations and eliminate all signals that violate constraints from the solution space. Soft consistency is preferred because of its robustness to noise and support mis-identification. If hard consistent reconstruction is enforced, the reconstruction constraints can produce an infeasible set in the presence of measurement noise. Furthermore, since there is no guarantee that the support has been correctly identified at step 7 of the MSP (otherwise there would be no need for iteration) there is no guarantee that hard consistency is feasible even in the absence of noise. Using the soft consistency penalty on unfeasible problems, the MSP attempts to reduce that penalty as close to zero as possible. The penalty functions chosen

in [5], [7] have similar soft-threshold properties, although the exact form is different.

IV. EXPERIMENTAL RESULTS

To tease out the effects of only imposing consistency, only imposing sparsity, and imposing both, our experiments compare three distinct approaches: (i) consistent reconstruction on the unit sphere, (ii) using the signs as measurement values in CoSaMP, and (iii) sparse consistent reconstruction using the MSP. The consistent reconstruction approach (i) attempts to reconstruct \mathbf{x} using:

$$\hat{\mathbf{x}} = \arg \min_{\mathbf{x}} \left\| (\text{diag}(\mathbf{y}) \mathbf{A} \mathbf{x})^- \right\|_2^2 \text{ s.t. } \|\mathbf{x}\|_2 = 1.$$

This program is optimized using the same algorithm as in step 7 of the MSP, but without explicit support constraints.

The CoSaMP approach (ii) treats each measurement as a value of ± 1 , depending on the sign of the measurement. This has the potential to produce a signal significantly away from $\|\mathbf{x}\|_2 = 1$, depending on the measurement matrix \mathbf{A} , which would unfairly penalize CoSaMP. To ensure a fair comparison, the CoSaMP output was normalized to $\|\mathbf{x}\|_2 = 1$ before evaluating its performance. Not normalizing the output significantly degraded CoSaMP's performance.

The experimental performance of MSP compared to the other two approaches is shown in Figure 1. Specifically, the

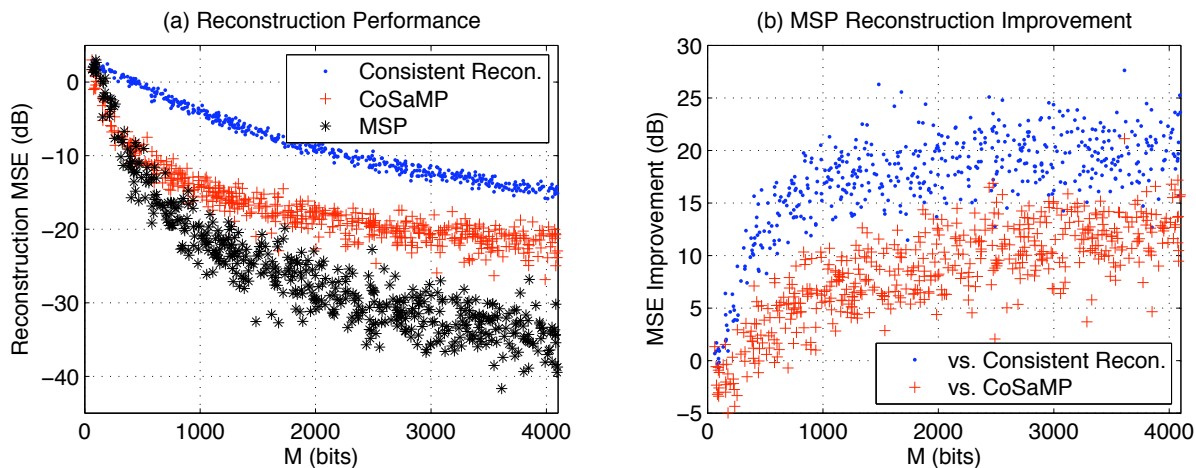


Fig. 1: Experimental performance of MSP vs. Consistent Reconstruction and CoSaMP

figure plots experimental results for $K = 20$, $N = 256$, and randomly generated $M \leq 4096$ (i.e. at a rate of up to 16 bits per signal coefficient). Our experiments demonstrated similar behavior for other values of K and N .

Figure 1(a) demonstrates the reconstruction performance of the three approaches. The figure plots the reconstruction MSE versus the number of measurements for each individual experiment. As shown in the figure, the MSP algorithm outperforms both consistent reconstruction without sparsity constraints and CoSaMP. It is also evident that the effect of sparsity is significant enough to ensure that CoSaMP outperforms consistent reconstruction. The effect is easier to quantify in Figure 1(b), in which the improvement of MSP over consistent reconstruction and CoSaMP is plotted. It is evident in the figure that at high bit rates MSP can have a gain of 12dB over CoSaMP and 20dB over Consistent Reconstruction.

Both the MSP and the ℓ_1 optimization on the unit ℓ_2 sphere [4] solve non-convex problems. Thus, initialization is important. We found that initializing the MSP with a few iterations of standard CoSaMP, treating the measurements as values would eliminate any issue of local convergence to an undesired optimum. Even with similar initialization on the ℓ_1 optimization, we found that the MSP has significantly better convergence properties.

V. DISCUSSION

In conclusion, it is evident from the results in this paper and in the CS literature that carefully exploiting the information provided in the acquisition system model provides significant improvements in the reconstruction performance. In the particular case of 1-bit quantization, this information is only about the sign of the measurements and nothing more. Since the measurements are unbounded, treating each quantized measurement as a value and attempting to reconstruct causes over-fitting and distortion to the reconstructed signal.

The MSP algorithm we developed explicitly incorporates a quantization model appropriate for 1-bit quantization. Using that model it significantly outperforms standard CS approaches. The same algorithm has significant applications beyond one-bit quantization—an example is sparse signal reconstruction from non-linearly distorted measurements [9]—which we defer to later publications. The MSP is an important tool in any system designer toolbox.

REFERENCES

- [1] E. J. Candès, “Compressive sampling,” in *Proc. International Congress of Mathematicians*, vol. 3, Madrid, Spain, 2006, pp. 1433–1452.
- [2] N. Thao and M. Vetterli, “Reduction of the MSE in R-times oversampled A/D conversion $O(1/R)$ to $O(1/R^2)$,” *Signal Processing, IEEE Transactions on*, vol. 42, no. 1, pp. 200–203, Jan 1994.
- [3] V. K. Goyal, M. Vetterli, and N. T. Thao, “Quantized overcomplete expansions in \mathbb{R}^N : Analysis, synthesis, and algorithms,” *IEEE Trans. Info. Theory*, vol. 44, no. 1, pp. 16–31, Jan. 1998.
- [4] P. Boufounos and R. G. Baraniuk, “One-Bit Compressive Sensing,” in *Proc. 42nd annual Conference on Information Sciences and Systems (CISS)*, Princeton, NJ, Mar 19–21 2008.
- [5] A. Zymnis, S. Boyd, and E. Candès, “Compressed sensing with quantized measurements,” *Preprint*, 2009.
- [6] W. Dai, H. Pham, and O. Milenkovic, “Distortion-rate functions for quantized compressive sensing,” *Preprint*, 2009.
- [7] L. Jacques, D. Hammond, and M. Fadili, “Dequantizing compressed sensing: When oversampling and non-gaussian constraints combine,” *Preprint*, 2009.
- [8] J. Laska, P. Boufounos, M. Davenport, and R. Baraniuk, “Democracy in action: Quantization, saturation, and compressive sensing,” *Preprint*, 2009.
- [9] P. Boufounos, “Reconstruction of sparse signals from distorted randomized measurements,” (*submitted*), 2009.
- [10] D. Needell and J. A. Tropp, “CoSaMP: Iterative signal recovery from incomplete and inaccurate samples,” *Appl. Comp. Harmonic Anal.*, vol. 26, pp. 301–321, 2008.
- [11] W. Dai and O. Milenkovic, “Subspace pursuit for compressive sensing signal reconstruction,” submitted, 2008. [Online]. Available: <http://arxiv.org/abs/0803.0811>
- [12] B. K. Natarajan, “Sparse approximate solutions to linear systems,” *SIAM J. Comput.*, vol. 24, pp. 227–234, 1995.
- [13] S. G. Mallat and Z. Zhang, “Matching pursuits with time-frequency dictionaries,” *IEEE Trans. Signal Processing*, vol. 41, no. 12, pp. 3397–3415, Dec. 1993.