MITSUBISHI ELECTRIC RESEARCH LABORATORIES
http://www.merl.com

Reconstruction of Sparse Signals from Distorted Randomized Measurements

Petros Boufounos

TR2010-014 April 2010

Abstract

In this paper we show that, surprisingly, it is possible to recover sparse signals from nonlinearly distorted measurements, even if the nonlinearity is unknown. Assuming just that the nonlinearity is monotonic, we use the only reliable information in the distorted measurements: their ordering. We demonstrate that this information is sufficient to recover the signal with high precision and present two approaches to do so. The first uses order statistics to compute the minimum mean square (MMSE) estimate of the undistorted measurements and use it with standard compressive sensing (CS) reconstruction algorithms. The second uses the principle of consistent reconstruction to develop a deterministic nonlinear reconstruction algorithm that ensures that measurements of the reconstructed signal have ordering consistent with the ordering of the distorted measurements. Our experiments demonstrate the superior performance of both approaches compared to standard CS methods.

International Conference on Acoustics Speech and Signal Processing

RECONSTRUCTION OF SPARSE SIGNALS FROM DISTORTED RANDOMIZED MEASUREMENTS

Petros T. Boufounos

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

ABSTRACT

In this paper we show that, surprisingly, it is possible to recover sparse signals from nonlinearly distorted measurements, even if the nonlinearity is unknown. Assuming just that the nonlinearity is monotonic, we use the only reliable information in the distorted measurements: their ordering. We demonstrate that this information is sufficient to recover the signal with high precision and present two approaches to do so. The first uses order statistics to compute the minimum mean square (MMSE) estimate of the undistorted measurements and use it with standard compressive sensing (CS) reconstruction algorithms. The second uses the principle of consistent reconstruction to develop a deterministic nonlinear reconstruction algorithm that ensures that measurements of the reconstructed signal have ordering consistent with the ordering of the distorted measurements. Our experiments demonstrate the superior performance of both approaches compared to standard CS methods.

Index Terms— Compressive Sensing, randomized sampling, consistent reconstruction, order statistics

1. INTRODUCTION

Compressive Sensing (CS) is a recently emerged signal acquisition technology that leverages the structure in most signals of interest to enable acquisition at rates significantly lower than previously thought possible. To do so, CS combines randomized linear measurements—which guarantee that the whole signal is observed—with nonlinear reconstruction to incorporate the signal structure in the recovery [1].

This paper extends the CS framework to recover signals from measurements observed through an unknown monotonic nonlinear distortion. We demonstrate that, surprisingly, the distortion maintains sufficient information to recover the signal. The key insight is that the relative ordering of the signal values is preserved because the distortion is monotonic. This ordering by itself preserves sufficient information for the signal reconstruction.

We contribute two significantly different recovery frameworks. The first is a statistical framework that estimates the undistorted measurements using the order statistics of the distorted measurements. This estimate can be used as the input to any reconstruction algorithm to recover the signal. The second is a deterministic framework that directly incorporates the ordering information in the reconstruction algorithm. In that framework, we present a new greedy reconstruction algorithm to produce a signal estimate consistent with the information in the measurement ordering. Both approaches outperform classical CS reconstruction on the distorted data.

The first CS hallmark we exploit is the randomized measurement process which makes individual measurements normally (or approximately normally) distributed i.i.d. random variables. We exploit the randomization by combining standard estimation theory with order statistics to compute a minimum mean squared error (MMSE) estimate of the undistorted measurements. The estimate is based only on the ordering of the distorted measurements. Our methodology does not rely on the signal structure or the reconstruction algorithm, only on the measurement process. Thus, although we present it in the context of CS, it can recover a variety of signals from their randomized distorted measurements using the appropriate recovery algorithm.

The second CS hallmark we exploit is nonlinear reconstruction which allows us to incorporate the ordering of the measurements as a constraint in the reconstruction process. Thus, in addition to prior knowledge of the signal structure we capitalize on our knowledge of the sampling system. The second part of this paper presents a general framework as well as a specific greedy reconstruction algorithm that imposes consistency with the measurement ordering. The algorithm harnesses recent results on consistent reconstruction from the sign of CS measurements [2, 3]. As in the first part, the consistent reconstruction principle can be used in non-sparse problems.

The applications of this work are numerous. Drift and nonlinear variations of sampling devices are very common in modern acquisition systems, both due to manufacturing and run-time conditions. For example, in optical acquisition systems such as [4], the operating temperature and the ambient light in the scene can make the device drift to a nonlinear acquisition region. The framework presented in this work is highly suitable to such acquisition systems.

The next section provides some background on CS with emphasis on 1-bit CS and a brief overview of the order statistics results used in this paper. Section 3 presents the measurement model and our assumptions. Section 4 develops the statistical approach to substitute the distorted measurements with their undistorted MMSE estimate. Section 5 develops the deterministic consistent reconstruction approach. Section 6 presents experimental results comparing the two approaches with each other and with standard practice.

2. BACKGROUND

2.1. Compressive Sensing

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

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

where $\mathbf{A} \in \mathbb{R}^{M \times N}$ models the measurement system.

Exact recovery is guaranteed if the measurement matrix $\mathbf{A}$ 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 $\mathbf{z}$

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

If δ_{2k} is small, $\mathbf{A}$ approximately maintains ℓ_2 distances between K-sparse signals. In this case, exact recovery is possible using a convex optimization [1] or a greedy algorithm with provable guarantees, such as the recently emerged Compressive Sampling Matching Pursuit (CoSaMP) [6] and the Subspace Pursuit [7].

An important result is that random matrices with sufficient number of rows can achieve small RIP constants with overwhelming probability. Thus, random matrices are commonly used for CS acquisition and reconstruction. The randomness of the acquisition matrix further ensures a nice statistical distribution of the results. Specifically, if the matrix is composed of i.i.d. random entries, then the measurements $\mathbf{y}$ also follow a normal distribution asymptotically [8]. We rely on this property in the first part of the paper.

It is important to note that the randomness of the measurements is with respect to the sample space of measurement matrices. Therefore, any results based on this distribution, such as the results in the first part of this paper, are for the typical case, not the worst case. This is in contrast to a significant part of the CS literature which considers worst case (adversarial) selection of the measured signal given the sampling matrix $\mathbf{A}$. The RIP is such an example property: it requires that the matrix is well-conditioned even for the worst case $\mathbf{x}$ to be sampled. Although this difference is often not consequential in practice, it can be an important distinction for certain applications.

2.2. 1-bit Compressive Sensing

A significant step in the CS literature that enables the development in this paper is 1-bit Compressive Sensing [2], in which a signal is acquired and reconstructed using only the signs of linear measurements. Specifically, a signal is sampled using

$$\mathbf{y} = \operatorname{sign}(\mathbf{A}\mathbf{x}), \tag{3}$$

where $\operatorname{sign}(\cdot) = \pm 1$ according to the sign of the measurement, and is applied element-wise. The signal is reconstructed to be consistent with the signs of the measurements, i.e., such that $\operatorname{sign}(\mathbf{A}\widehat{\mathbf{x}}) = \mathbf{y}$ and $\widehat{\mathbf{x}}$ is sparse.

The signs of the measurements eliminate any information about the amplitude of the signal. Thus, the constraint $\|\widehat{\mathbf{x}}\|_2 = 1$ is imposed in the reconstruction, i.e., the reconstruction is performed on the unit sphere. Sparsity is enforced by minimizing the ℓ_1 norm on the sphere (otherwise $\widehat{\mathbf{x}} = 0$ is the minimizer).

Consistency with the measurements is imposed by relaxing the strict constraints and introducing a one-sided quadratic penalty if a constraint is violated. This can be expressed as the squared norm of the measurements that violate the constraint.Specifically, we let $(\cdot)^-$ denote the negative part of a scalar, i.e.,

$$(x)^- = -\min(x, 0) = \frac{|x| - x}{2} = \begin{cases} 0, & \text{if } x \geq 0 \\ -x & \text{otherwise.} \end{cases} \tag{4}$$

The penalty is then equal to

$$c(\widehat{\mathbf{x}}) = \left\| \left(\operatorname{diag}(\mathbf{y}) \mathbf{A}\widehat{\mathbf{x}} \right)^- \right\|_2^2, \tag{5}$$

where $\operatorname{diag}(\mathbf{y})$ is a matrix with the signs of the measurements in its diagonal. The negative part operator $(\cdot)^-$ is applied element-wise to identify the constraint violations and the magnitude of the violation. A signal estimate consistent with the samples will produce no constraint violations and the penalty function $c(\widehat{\mathbf{x}})$ will equal zero.

Using (5) the reconstruction problem is relaxed to

$$\widehat{\mathbf{x}} = \underset{\mathbf{x},\, \|\mathbf{x}\|_2 = 1}{\arg\min} \|\mathbf{x}\|_1 + \frac{\lambda}{2} \left\| \left(\operatorname{diag}(\mathbf{y}) \mathbf{A}\widehat{\mathbf{x}} \right)^- \right\|_2^2. \tag{6}$$

Of course, (6) is non-convex and convergence to the global optimum cannot be guaranteed.

Another recent approach to recover the sparse signal is the Matching Sign Pursuit (MSP)—a greedy algorithm that computes a sparse minimum to the penalty function (5) [3]. The MSP performs an iterative greedy search inspired by CoSaMP [6]. Specifically, the MSP, described extensively in [3], updates a sparse estimate of the signal $\widehat{\mathbf{x}}$ by iterating the following until convergence:
(*i*) Identify which sign constraints are violated
(*ii*) Identify the signal components mostly effective in minimizing the cost function and reducing the sign violations
(*iii*) Minimize the cost function over those components
(*iv*) Truncate the signal to the desired sparsity, normalize to $\|\widehat{\mathbf{x}}\|_2 = 1$, and update the estimate.

It is important to highlight two distinct advantages of the MSP over the minimization in (6). First, it is straightforward to show that the original signal terminates the algorithm. So does any K-sparse consistent solution, as desired. Second, even though both attempt an optimization over a non-convex space, MSP experimentally performs better in avoiding the local minima. Thus, in Section 5 we prefer the MSP to perform consistent reconstruction.

2.3. Order Statistics

We assume a set of M random variables $y_i \in \mathbb{R}$, also denoted as a vector $\mathbf{y} \in \mathbb{R}^M$. These are independently drawn from some distribution $f(y)$, with cumulative distribution function $F(y)$. For the remaining of this paper we use $y_{(i)}$ to denote the set of variables sorted in ascending order, $y_{(1)} \leq y_{(2)} \leq \ldots \leq y_{(M)}$, and k_i to denote the index of the i^{th} sorted measurement, i.e.,

$$y_{(i)} = y_{k_i},\ i = 1, \ldots, M. \tag{7}$$

The sorted observations $y_{(i)}$ form the order statistics of the observations. We further define $p_i = \frac{i}{M+1}$ and $q_i = 1 - p_i$, which asymptotically counts the fraction of measurements less than and greater than $y_{(i)}$, respectively.

The moments of order statistics generally do not have a closed form. An asymptotically accurate unbiased approximation is [9]

$$E\left(y_{(i)}\right) = Q\left(p_i\right), \tag{8}$$

$$E\left(y_{(i)} y_{(j)}\right) = \frac{p_i q_j}{M+2} Q'(p_i) Q'(q_j) \tag{9}$$

where $Q(x) = F^{-1}(x)$ is the inverse of the CDF, often referred to as the quantile function, and $Q'(x)$ is its derivative evaluated at x. This should not be confused with the $Q(\cdot)$ function often denoting the tail integral of the normal distribution—not used in this paper.

3. MEASUREMENT MODEL

For the remainder of this paper we consider linear measurements of a signal $\mathbf{x} \in \mathbb{R}^N$ using inner products with the rows $\mathbf{a}_i$ of a measurement matrix $\mathbf{A} \in \mathbb{R}^{M \times N}$, which is random but known at the reconstruction. Without loss of generality, the matrix is assumed to have random i.i.d. elements drawn from the standard normal distribution $\mathcal{N}(0, 1)$.[1] Each measurement is observed through the same nonlinear function $g(x)$:

$$\mathbf{y} = g(\mathbf{A}\mathbf{x}), \tag{10}$$

[1] In most of the CS literature the entries of $\mathbf{A}$ are drawn with variance $\sigma^2 = 1/M$. This choice is inconsequential but simplifies further notation.

where $g(x)$ is applied element-wise to each of the vector coefficients. Although we do not assume knowledge of $g(x)$, we do assume it is strictly increasing, i.e., $g(x_1) > g(x_2) \Leftrightarrow x_1 > x_2$. We briefly discuss relaxing this assumption at the end of Section 6. Due to lack of space we focus on the fundamental concepts in this paper and we assume noiseless sampling. Still, in the experimental section we demonstrate that our methods are robust in the presence of noise.

Since the nonlinear distortion is unknown, only limited information is conveyed through $g(\cdot)$. For example, the unknown distortion $g(x)$ eliminates any magnitude information on $\mathbf{x}$: for any $c > 0$, $\mathbf{y} = g(\mathbf{A}\mathbf{x}) = \widehat{g}(\mathbf{A}c\mathbf{x})$, where $\widehat{g}(x) = g(x/c)$ is also monotonic. Either of the two functions could be the nonlinearity distorting the signal. Thus, $\mathbf{x}$ can only be recovered within a positive scalar factor. Furthermore, any other monotonic distortion $\widetilde{g}(\cdot)$ of the measurements may originate from the same signal since the composition of two monotonic functions is also monotonic.

On the other hand, it is trivial to show that the nonlinearity maintains the sorting order of the measurements. This is a property we exploit in our development. From (7) it follows that

$$\operatorname{sign}\left(y_{(i)} - y_{(j)}\right) = \operatorname{sign}\left(i - j\right). \tag{11}$$

The index sequence $\{k_1, \ldots, k_N\}$ is preserved among all monotonic distortions $g(x)$, including the identity. Furthermore, once the sequence $\{k_i\}$ is known, the exact values of $y_{(i)}$ provide no further information and cannot be useful in the reconstruction. This is because a nonlinear monotonic distortion can always be constructed that maps $y_{(i)}$ to any other $y'_{(i)}$ that has the same ordering $\{k_i\}$.

Since $\mathbf{A}$ is a random matrix with i.i.d. normally distributed entries, the undistorted measurements—denoted $\overline{\mathbf{y}} = \mathbf{A}\mathbf{x}$—are also normally distributed i.i.d. variables with respect to the sample space of $\mathbf{A}$. Asymptotically, this is true even if the matrix entries are i.i.d. but not normally distributed, due to the central limit theorem [8].

4. MEASUREMENT SUBSTITUTION

This section uses the measurement order statistics to reconstruct the signal. Since the nonlinearity $g(\cdot)$ is unknown, the ordering of the measurements is the only reliable information we obtain from $\mathbf{y}$. Using the normality of the undistorted measurements $\overline{\mathbf{y}}$ we develop a robust estimator for $\overline{\mathbf{y}}$ using its order statistics.

Specifically, instead of using the distorted measurements we replace them with the MMSE estimate of the undistorted values, conditioned only on the measurement ordering. The estimator is a function of the measurement ordering denoted using $\widehat{\mathbf{y}}\left(\{k_i\}\right)$. From standard estimation theory, it is the conditional expectation

$$\widehat{\mathbf{y}}\left(\{k_i\}\right) = E(\overline{\mathbf{y}}|\{k_i\}) \tag{12}$$

As argued in Section 3, the measurement process removes any information on the signal amplitude. Thus, the signal can only be identified within a positive scaling factor. As with [2], we normalize the recovered signal to have unit ℓ_2 norm. Under this assumption and the normality of the measurement matrix, the measurements follow the standard normal distribution $\mathcal{N}(0, 1)$. Using the asymptotic approximation (8) into (12), the estimator for $\widehat{y}_{(i)}$ follows

$$\widehat{y}_{(i)} = \widehat{y}_{k_i} = \Phi^{-1}\left(p_i\right), \tag{13}$$

where $\Phi(\cdot)$ denotes the CDF of the standard normal distribution

$$\Phi(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{x} e^{-u^2/2} du. \tag{14}$$

The estimated measurements can be used as input to any reconstruction algorithm to reconstruct the signal.

5. CONSISTENT RECONSTRUCTION

In this section instead of using the measurement ordering to estimate the undistorted measurements, we impose consistent measurement ordering as a constraint in the reconstruction algorithm. In other words, we use a reconstruction algorithm that ensures that the measurements of the reconstructed signal have the same ordering as the measurements of the original signal.

To impose consistency we capitalize on the the MSP algorithm briefly described in Section 2.2. We use the MSP with the measurement model in Section 3 to derive an implicit sampling matrix $\widetilde{\mathbf{A}}$ such that $\widetilde{\mathbf{y}} = \operatorname{sign}\left(\widetilde{\mathbf{A}}\mathbf{x}\right)$ can be derived from the ordering of the measurements. If we let $\mathbf{a}_k$ denote the k^{th} row of $\mathbf{A}$, then

$$y_{(i)} > y_{(j)} \Leftrightarrow y_{k_i} > y_{k_j} \Leftrightarrow \langle \mathbf{a}_{k_i}, \mathbf{x} \rangle > \langle \mathbf{a}_{k_j}, \mathbf{x} \rangle \tag{15}$$

$$\Leftrightarrow \langle \mathbf{a}_{k_i} - \mathbf{a}_{k_j}, \mathbf{x} \rangle > 0 \tag{16}$$

$$\Leftrightarrow \operatorname{sign}\left(\langle \mathbf{a}_{k_i} - \mathbf{a}_{k_j}, \mathbf{x} \rangle\right) = \operatorname{sign}\left(i - j\right), \tag{17}$$

where (15) follows from the monotonicity of the nonlinear distortion in (10), and (17) follows from (11), i.e. from the properties of the sorting index sequence $\{k_i\}$. In other words we can construct the matrix $\widetilde{\mathbf{A}}$ using rows of the form $\mathbf{a}_{k_i} - \mathbf{a}_{k_j}$ such that

$$\operatorname{sign}\left(\widetilde{\mathbf{A}}\mathbf{x}\right) = \operatorname{sign}\begin{pmatrix} \vdots \\ i - j \\ \vdots \end{pmatrix} \triangleq \widetilde{\mathbf{y}}. \tag{18}$$

This $\widetilde{\mathbf{A}}$ and the corresponding sign measurements $\widetilde{\mathbf{y}}$ are provided as input to the MSP algorithm to estimate $\mathbf{x}$. Eq. (18) above holds for any choice of index pairs (k_i, k_j) and, therefore, for any choice of vector pairs $(\mathbf{a}_{k_i}, \mathbf{a}_{k_j})$ chosen to construct $\widetilde{\mathbf{A}}$. This is a design choice in our approach. In this paper we use the $(M - 1)$ pairs $(k_{i+1}, k_i), i = 1, \ldots, M - 1$ which, in principle, guarantee that the reconstruction is consistent with every pair (k_i, k_j). In our experiments, incorporating more pairs increased the computational cost of the algorithm without any reconstruction benefits.

Since this approach attempts to solve a non-convex problem the choice of initial seed is also important to facilitate convergence to the global optimum. Even though the MSP algorithm has significantly better convergence performance than the ℓ_1 optimization on the sphere [2], we still observe convergence issues, especially with a small number of measurements M. Fortunately we already have a very good initial seed: we use measurement substitution and a few iterations of standard CS decoding (for example CoSaMP [6]). This provides an inexpensive way to 'warm start' the algorithm.

6. EXPERIMENTAL RESULTS AND DISCUSSION

To validate both approaches we performed a series of simulations, which we report here. Our results were robust to variations in parameters, so in the interest of space we only present a small snapshot. The experiments use random matrices with standard normal i.i.d. elements, of varying dimensions M and N. The sparse signals are have $K = 32$ random non-zero coefficients with amplitude selected from a normal distribution and normalized to have unit norm. The signal is sampled through a $\tanh(\cdot)$ nonlinearity to which white Gaussian noise $\mathbf{n}$ is added with standard deviation $\sigma_{\mathbf{n}} = 0.01$ and 0.001 per measurement. To summarize, the measurement process implemented the following equation

$$\mathbf{y} = \tanh\left(\mathbf{A}\mathbf{x}\right) + \mathbf{n}. \tag{19}$$

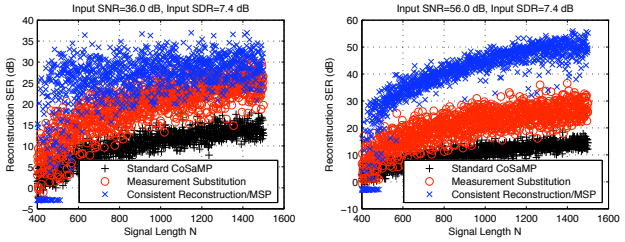

Fig. 1. Reconstruction Signal-to-Error ratio, for varying signal length N, with fixed sparsity $K = 32$ and undersampling ratio $M/N = 1/4$. Results for two different input SNR are displayed.

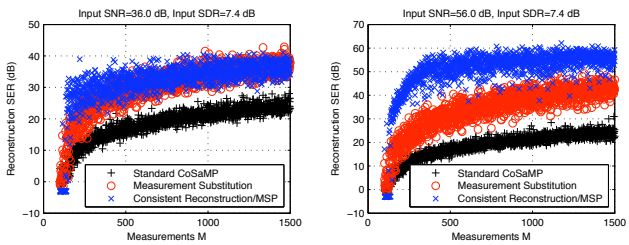

Fig. 2. Reconstruction Signal-to-Error ratio, for varying number of measurements M, with fixed sparsity $K = 32$ and signal length $N = 1000$. Results for two different input SNR are displayed.

The measured SNR of the measurements after the nonlinearity (i.e. only due to the noise) was

$$\text{SNR} \triangleq 20\log_{10}\left(\|\tanh\left(\mathbf{Ax}\right)\|_2/\|\mathbf{n}\|_2\right) \approx 36 \text{ and } 56 \text{ dB}, \quad (20)$$

for $\sigma_{\mathbf{n}} = 0.01$ and 0.001 respectively, while the total measured Signal-to-Distortion-Ratio (SDR), which includes the effect of the nonlinearity was

$$\text{SDR} \triangleq 20\log_{10}\left(\|\mathbf{Ax}\|_2/\|\mathbf{Ax}-\mathbf{y}\|_2\right) \approx 7.4 \text{ dB} \quad (21)$$

for both $\sigma_{\mathbf{n}} = 0.01$ and 0.001. In other words, the effect of the nonlinearity was significant and dominated the distortion, while the noise mostly affected the input SNR.

The reconstruction is performed using three different approaches. The first is a reference reconstruction using the standard CoSaMP algorithm on the distorted measurements $\mathbf{y}$. The second is using the measurement substitution combined with the standard CoSaMP, as described in the first part of the paper. The third is the consistent reconstruction using the MSP algorithm, 'warm started' using a few iterations of the measurement substitution with the standard CoSaMP. For a fair comparison we normalized all outputs to have unit ℓ_2 norm, i.e., $\|\widehat{\mathbf{x}}\|_2 = 1$ (without the normalization the performance of standard CoSaMP was worse).

We performed two different sets of experiments. In the first set the the sparsity level is constant $K = 32$ and we let N and M vary, keeping their ratio fixed at $M/N = 1/4$. For each choice of M and N a different set of random $\mathbf{A}$, $\mathbf{x}$, and $\mathbf{n}$ was generated.

In the second set of experiments the sparsity and the signal length are constant $K = 32$ and $N = 1000$. The number of measurements varies in the range $M = 100, \ldots, 1500$. Of course, as M increases the measurements are less "compressive" but we include these results because they demonstrate clear trends. Furthermore, as mentioned in the previous sections, the results are applicable even when the number of measurements is not the main cost in the system. The sparsity model and the non-linear reconstruction for $M \geq 1000$ serves more as a denoising model rather than a CS one.

Figures 1 and 2 plot the experimental reconstruction Signal-to-Error Ratio (SER)

$$\text{SER}_{\text{recon}} \triangleq 20\log_{10}\left(\|\mathbf{x}\|_2/\|\mathbf{x}-\widehat{\mathbf{x}}\|_2\right). \quad (22)$$

The figures demonstrate that the consistent reconstruction approach using the MSP outperforms the other approaches and is able to reconstruct the signal roughly at the sampling SNR, especially as M increases. It is also notable that at lower sampling SNRs and larger number of measurements, measurement substitution performs as well as consistent reconstruction. Both methods outperform standard reconstruction using CoSaMP. It is also important to point out that with very few measurements, in the range $M \approx 100, \ldots, 200$, consistent reconstruction was often not able to converge to the global optimum and did not perform well. Failure to converge was not encountered in any of our data points above $M \approx 200$. Another interesting observation is that as the number of measurements increase, the performance gap between measurement substitution and consistent reconstruction decreases.

We should note that noise under this distortion model can make the $\tanh(\cdot)$ distortion non-invertible, even when the distortion function is known and the recovery algorithm attempts to invert it. The noise might make the measurement magnitude exceed 1, the range of $\tanh(\cdot)$. This is robustly handled by the two approaches described in our paper, since neither attempts to invert the nonlinearity.

Each of the two methods has particular strengths and weaknesses. Incorporating nonlinearities that are not strictly monotonic, such as quantization and saturation, is straightforward with consistent reconstruction but not obvious in the measurement substitution framework. Furthermore, deterministic matrix constructions can be used with consistent reconstruction but not with measurement substitution. On the other hand, measurement substitution uses significantly less computation and well-established convex reconstruction algorithms with guaranteed performance. Both are significant additions to the system designer's toolkit with remarkable performance.

7. REFERENCES

[1] E. J. Candès, "Compressive sampling," in *Proc. International Congress of Mathematicians*, Madrid, Spain, 2006, vol. 3, pp. 1433–1452.

[2] 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.

[3] P. Boufounos, "Greedy sparse signal reconstruction from sign measurements," in *Asilomar Conference On Signals Systems and Computers*, 2009, accepted, to apear.

[4] M. F. Duarte, M. Davenport, D. Takhar, J. N. Laska, T. Sun, K. Kelly, and R. G. Baraniuk, "Single-pixel imaging via compressive sampling," in *IEEE Signal Processing Magazine*, March 2008, vol. 25, pp. pp. 83–91.

[5] S.G. Mallat and Zhifeng Zhang, "Matching pursuits with time-frequency dictionaries," *IEEE Transactions on Signal Processing*, vol. 41, no. 12, pp. 3397–3415, Dec 1993.

[6] 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.

[7] Wei Dai and Olgica Milenkovic, "Subspace pursuit for compressive sensing signal reconstruction," submitted, 2008.

[8] J. N. Laska, P. Boufounos, and R. G. Baraniuk, "Finite-range scalar quantization for compressive sensing," in *Conf. on Sampling Theory and Applications (SampTA)*, 2009.

[9] H.A David, *Order Statistics*, John Wiley & Sons Inc., 2nd edition, 1981.