ON QUANTIZED COMPRESSED SENSING WITH SATURATED MEASUREMENTS VIA GREEDY PURSUIT

Ines Elleuch, Fatma Abdelkefi*, Mohamed Siala*, Ridha Hamila†, Naofal Al-Dhahir‡*

* MEDIATRON Laboratory, SUP'COM, University of Carthage, Tunisia
† Electrical Engineering Department, Qatar University, Qatar
‡ Electrical Engineering Department, University of Texas at Dallas, USA

ABSTRACT

We consider the problem of signal recovery under a sparsity prior, from multi-bit quantized compressed measurements. Recently, it has been shown that allowing a small fraction of the quantized measurements to saturate, combined with a saturation consistency recovery approach, would enhance reconstruction performance. In this paper, by leveraging the potential sparsity of the corrupting saturation noise, we propose a model-based greedy pursuit approach, where a *cancel-then-recover* procedure is applied in each iteration to estimate the unbounded sign-constrained saturation noise and remove it from the measurements to enable a *clean* signal estimate. Simulation results show the performance improvements of our proposed method compared with state-of-the-art recovery approaches, in the noiseless and noisy settings.

***Index Terms*—** Multi-Bit Quantized Compressed Sensing, Saturation, Sparse Corruptions, Sign Constraint, Cancel-Then-Recover, Greedy Pursuit

1. INTRODUCTION

Quantized Compressed Sensing (QCS) is an emerging research field addressing the practical implementation of Compressed Sensing (CS) paradigm by considering measurements discretization arising from quantization. This discretization process implies two kinds of acquisition noise: *quantization noise* of bounded energy, due to the quantizer finite precision, and *saturation noise* of large and potentially unbounded amplitudes, arising from the quantizer finite dynamic range.

In addition to convex optimization based recovery algorithms [1–4], some papers tackle the sparse recovery problem from QCS measurements via a greedy approach. To address the extreme case of 1-bit CS, where all measurements saturate, the iterative algorithms CoSaMP [5] and IHT [6] have been tuned, through Matching Sign Pursuit (MSP) in [7] and Binary Iterative Hard Thresholding (BIHT) in [8], respectively. The key ingredient for these algorithm customizations is to incorporate a sign consistency objective function to penalize sign violation. MSP and BIHT consider one sided ℓ_2-norm and one sided ℓ_1-norm penalties, respectively.

The multi-bit quantization case is more subtle, because saturated and unsaturated measurements are contaminated with noise of different nature. In [9], the authors proposed Quantized IHT (QIHT) where a unified penalty function is considered to account for quantization and saturation inconsistency. In [4], the authors proposed a rejection approach and a saturation consistency approach, based on CoSaMP. In the rejection approach, saturated measurements, regarded as outliers that would significantly deteriorate CS reconstruction performances, are simply discarded, and a Least Squares (LS) objective function is applied on the remaining measurements. Alternatively, saturated and unsaturated measurements are decoupled within the Saturation Consistency SC-CoSaMP method, where two objective functions are used: an LS objective related to the unsaturated measurements and a one sided ℓ_2-norm objective to account for saturation violation.

In the aforementioned papers, soft consistency was preferred to hard consistency, because of its robustness to noise and support mis-identication. Moreover, it has been shown in [8], that different choices of penalty functions to account for inconsistency, provide different performances, depending on the noise level. These choices also rely on the existence of a tractable analytical expression of the penalty subgradient.

In this paper, we propose a different approach to solve the sparse recovery problem from quantized and possibly saturated measurements. By leveraging the sign characterization of the saturation errors, we implicitly promote their sparsity even within the saturated measurements. We consider a greedy approach that provides an appropriate framework to apply the *cancel-then-recover* approach of [10].

The rest of this paper is organized as follows. Section II presents the QCS model with saturation. Section III provides the framework and motivation behind this work. Section IV proposes a novel sparse recovery method that incorporates the *cancel-then-recover* approach within the greedy pursuit. Section V demonstrates the performance gain of the proposed method, via simulation results. Finally, Section VI concludes the paper.

This paper was made possible by NPRP grant # 6-070-2-024 from the Qatar National Research Fund (a member of Qatar Foundation). The statements made herein are solely the responsibility of the authors.

In the sequel, we will use bold-face lowercase letters for vectors and bold-face capital letters for matrices. For an ordered set T, $T(i)$ denotes its i^{th} element, $|T|$ denotes its cardinality, T^c denotes its complement. Notations $\mathbf{x}_T$ and $\mathbf{\Phi}_T$ stand for the sub-vector of elements of $\mathbf{x}$, and the sub-matrix of columns of $\mathbf{\Phi}$, indexed by T, respectively. The support of $\mathbf{x}$, denoted $\text{supp}(\mathbf{x})$, corresponds to the ordered set of indices of its nonzero entries and $\|\mathbf{x}\|_0 \triangleq |\text{supp}(\mathbf{x})|$. The ℓ_1-norm of $\mathbf{x} \in \mathbb{R}^N$, is defined as $\|\mathbf{x}\|_1 \triangleq \sum_{n=1}^N |x_n|$. $\mathbf{x}_{(K)}$ is the best K-term approximation of $\mathbf{x}$, obtained by keeping the K largest entries of $\mathbf{x}$ and zeroing the others. The sign operator sign is applied on vectors, componentwise. The hard thresholding operator $\mathcal{H}_\lambda$ sets the components of a vector, below λ, to 0. The symbols $\succeq$ and $\preceq$ denote entry-wise inequalities.

2. QUANTIZED COMPRESSED SENSING MODEL

We consider that compressed sensing measurements $\mathbf{z} \in \mathbb{R}^M$ are quantized using a b-bit uniform midrise quantizer operator $\mathcal{Q}_b$ with quantization interval δ, 2^b quantization levels and a saturation level $g = 2^{b-1}\delta$, such that:

$$\mathbf{y} = \mathcal{Q}_b(\mathbf{z}), \quad \mathbf{z} = \mathbf{\Phi x}, \tag{1}$$

where $\mathbf{x} \in \mathbb{R}^N$ is a K-sparse signal with $K \ll N$, $\mathbf{\Phi} \in \mathbb{R}^{M\times N}$ with $M < N$ is the measurement matrix, and $\mathbf{y} \in \mathbb{R}^M$ is the quantized measurements vector.

The model in (1) can be rewritten into

$$\mathbf{y} = \mathbf{\Phi x} + \mathbf{e} + \mathbf{n}, \tag{2}$$

where $\mathbf{e}$ and $\mathbf{n} \in \mathbb{R}^M$ account for the saturation noise and the quantization noise, respectively. At a low saturation rate, $\mathbf{e}$ is a sparse but large error term ($e_m = 0$ or $|e_m| > \frac{\delta}{2}$) corrupting a fraction of the measurements $\mathbf{z}$ and $\mathbf{n}$ as a dense but small error term ($|n_m| \le \frac{\delta}{2}$) affecting all its entries.
Let $\mathcal{S}$ denote the known support of the potentially saturated measurements corresponding to $|z_m| \ge g - \delta$

$$\mathcal{S} \triangleq \{m \in \{1,\dots,M\} : |y_m| = (g - \delta/2)\}, \quad |\mathcal{S}| = S.$$

Then, we have $\mathbf{e}_{\mathcal{S}^c} = \mathbf{0}$. By exploiting partial knowledge of the support of $\mathbf{e}$, the model in (2) can be further adjusted as

$$\mathbf{y} = \mathbf{\Phi x} + \mathbf{I}_{\mathcal{S}}\,\mathbf{e}_{\mathcal{S}} + \mathbf{n}, \tag{3}$$

where $\mathbf{I}$ is the identity matrix of order M. Let $\mathcal{E}$ denote the unknown support of the corruption term $\mathbf{e}$, i.e. that of effectively saturated measuremets

$$\mathcal{E} \triangleq \{m \in \{1,\dots,M\} : g \le |z_m|\}, \qquad |\mathcal{E}| = E.$$

Then, we have $\mathcal{E} \subseteq \mathcal{S}$, and $\mathbf{y}_{\mathcal{S}\cap\mathcal{E}^c}$ represents artificially saturated measurements, i.e. their pre-quantized measurements lie within the quantization regions of the extremum quantization levels. Formally, we have $\mathbf{n}_{\mathcal{E}} = \mathbf{0}$, $|\mathbf{n}_{\mathcal{E}^c}| \preceq \frac{\delta}{2}$, and

$$\mathbf{e}_{\mathcal{E}^c} = \mathbf{0}, \quad |\mathbf{e}_{\mathcal{E}}| \succeq \tfrac{\delta}{2}, \quad \text{sign}(\mathbf{e}_{\mathcal{E}}) = -\text{sign}(\mathbf{y}_{\mathcal{E}}). \tag{4}$$

Let $\mathbf{\Theta} = \mathbf{I}_{\mathcal{S}}\mathbf{\Lambda} \in \mathbb{R}^{M\times S}$, $\mathbf{s} = \mathbf{\Lambda e}_{\mathcal{S}} \in \mathbb{R}^S$ where $\mathbf{\Lambda} \in \mathbb{R}^{S\times S}$ is a diagonal matrix whose diagonal elements are given by $-\text{sign}(\mathbf{y}_{\mathcal{S}})$, and $\Sigma = \text{supp}(\mathbf{s}) = \{i \in \{1,\dots,S\} : \mathcal{S}(i) \in \mathcal{E}\}$. Then, by incorporating (4) into the model in (3), we get

$$\mathbf{y} = \mathbf{\Phi x} + \mathbf{\Theta s} + \mathbf{n}, \quad \text{where} \quad \mathbf{s}_{\text{supp}(\mathbf{s})} \succeq \tfrac{\delta}{2}. \tag{5}$$

3. GREEDY FRAMEWORK AND MOTIVATION

We propose to customize the iterative greedy algorithm CoSaMP [5] to perform sparse recovery from quantized and partially saturated measurements. The basic CoSaMP algorithm solves the following optimization problem:

$$\min_{\widetilde{\mathbf{x}}\in\mathbb{R}^N} \|\mathbf{y} - \mathbf{\Phi}\widetilde{\mathbf{x}}\|_2^2 \quad s.t. \quad \|\widetilde{\mathbf{x}}\|_0 \le K. \tag{6}$$

After initialization, each CoSaMP iteration repeats five steps:

1. Compute a proxy of the residual as $\mathbf{\Phi}^T\mathbf{r}$,
2. Identify a support candidate Ω by merging the best $2K$-term sparse approximation support of the proxy and the signal support from the previous iteration,
3. Compute a signal estimate using LS on candidate support $\widehat{\mathbf{x}}_\Omega \leftarrow \text{argmin}_{\widetilde{\mathbf{x}}\in\mathbb{R}^{|\Omega|}} \|\mathbf{y} - \mathbf{\Phi}_\Omega\widetilde{\mathbf{x}}\|_2^2$,
4. Prune the support via the best K-term approximation,
5. Update the residual $\mathbf{r} = \mathbf{y} - \mathbf{\Phi}\widehat{\mathbf{x}}$.

In the QCS framework, as the saturation noise is potentially unbounded, saturated measurements could be regarded as outliers. LS fitting is known to be sensitive to outliers, which would dramatically degrade the overall performance of CoSaMP. A simple solution to overcome this issue is to discard saturated measurements and reconstruct the signal using the remaining ones, as suggested in [4], via the so-called rejection approach. However, this approach is suboptimal in comparaison with the conventional setting where the saturation level is increased so that saturation never occurs.

Alternatively, the authors in [4], showed that allowing a fraction of measurements to saturate would be advantageous, provided that information held by saturated measurements is appropriately incorporated within the reconstruction process. They considered the saturation consistency constraint

$$-\mathbf{\Theta}^T\mathbf{\Phi x} \succeq g - \delta. \tag{7}$$

This constraint is softly incorporated into the signal estimation step of SC-CoSaMP, as follows:

$$\widehat{\mathbf{x}}_\Omega \leftarrow \underset{\widetilde{\mathbf{x}}\in\mathbb{R}^{|\Omega|}}{\text{argmin}} \|\mathbf{y}_{\mathcal{S}^c} - \mathbf{I}_{\mathcal{S}^c}^T\mathbf{\Phi}_\Omega\widetilde{\mathbf{x}}\|_2^2 + \|\big((g-\delta)\mathbf{1} + \mathbf{\Theta}^T\mathbf{\Phi}_\Omega\widetilde{\mathbf{x}}\big)_+\|_2^2, \tag{8}$$

where $\mathbf{1}$ denotes the unit vector and the function $(\cdot)_+$ acts by zeroing the negative elements. In this minimization subproblem, the first term accounts for the fitting error, with respect to $\mathbf{y}_{\mathcal{S}^c}$, due to quantization noise and the second one-sided ℓ_2-norm term penalizes saturation violation with respect to $\mathbf{y}_{\mathcal{S}}$.

Allowing saturation violation ensures feasible solutions even with erroneous signal support candidate or in the presence of noise. However, the saturation consistency approach of [4] remains suboptimal. Indeed, saturated measurements, even those artificially saturated, don't contribute in the LS objective and saturation is only addressed as a consistency constraint. Hence, apart from missing a fraction of the quantized measurements because they artificially saturate, the consistency approach fails to minimize the underlying saturation noise. Moreover, on the one hand, hard SC-CoSaMP is prone to infeasibility, especially in the presence of signal noise. On the other hand, in soft SC-CoSaMP, it is not clear whether the one-sided ℓ_2-norm penalty is optimal. Besides, the choice of the penalty parameter that balances the conventional objective and the saturation consistency violation, is not addressed.

4. PROPOSED SPARSE RECOVERY FROM QCS MEASUREMENTS VIA GREEDY PURSUIT

In order to retain all the potential of the saturated measurements, we build upon (5) and propose a sparse recovery approach that solves the following optimization problem:

$$\min_{\substack{\widetilde{\mathbf{x}}\in\mathbb{R}^N \\ \widetilde{\mathbf{s}}\in\mathbb{R}^S}} \|\mathbf{y}-\mathbf{\Theta}\widetilde{\mathbf{s}}-\mathbf{\Phi}\widetilde{\mathbf{x}}\|_2^2 \quad s.t. \quad \begin{cases} \|\widetilde{\mathbf{x}}\|_0 \le K, \\ \tilde{s}_i \in \{0\} \cup [\frac{\delta}{2}, +\infty[. \end{cases} \tag{9}$$

In this formulation, $\mathbf{y}-\mathbf{\Theta}\widetilde{\mathbf{s}}$ could be seen as a *clean* measurements vector obtained by removing the saturation noise producing annoying outliers. The key observation is that all M *cleaned* measurements contribute to the LS objective. Moreover, potential overfitting, caused by saturation noise support over-identification ($\mathcal{E} \subseteq \mathcal{S}$) is mitigated, by modeling the corruption term $\mathbf{s}$ as non-negative with possibly null elements.

4.1. Key Ingredients for the Proposed Method

We propose to solve problem (9) via a model-based CoSaMP, where the signal estimation and the residual computation steps are modified according to the model in (5). A key enabler for the proposed approach, is that at each iteration of the greedy pursuit, the signal estimation task is performed by considering a reduced cardinality support candidate Ω such that $|\Omega| \le 3K < M$. Hence, the underlying system is overdetermined. As the signal of interest is corrupted by a non-negative, presumably sparse and unbounded saturation noise term, we borrow from [10] and [11]. In the context of error correction coding, the authors of [10], proposed a method to recover a general (not sparse) signal from its noisy and corrupted encoded measurements. Formally, the received distorted codeword is modeled as $\mathbf{b} = \mathbf{A}\mathbf{x} + \mathbf{e} + \mathbf{n}$ where $\mathbf{A} \in \mathbb{R}^{M\times N}$ with $M > N$, is the coding matrix, $\mathbf{e}$ is a gross but sparse noise term and $\mathbf{n}$ is a small but dense noise term. By leveraging the sparsity of the corruption term, the following *cancel-then-recover* decoding approach was proposed in [10] to recover $\mathbf{x}$. First, the signal term is canceled from the measurements, by projecting them, using the matrix $\mathbf{Q}$, onto the orthogonal complement of the space spanned by the columns of $\mathbf{A}$, such that $\mathbf{Q}^T\mathbf{A} = \mathbf{0}$. Then, the sparse corruption term is estimated as

$$\hat{\mathbf{e}} = \underset{\tilde{\mathbf{e}}\in\mathbb{R}^M}{\operatorname{argmin}} \|\tilde{\mathbf{e}}\|_1 \quad s.t. \quad \|\mathbf{Q}^T(\mathbf{b}-\tilde{\mathbf{e}})\|_2^2 \le \epsilon. \tag{10}$$

Finally, $\mathbf{x}$ is recovered using a *clean* LS solution $\hat{\mathbf{x}} = (\mathbf{A}^T\mathbf{A})^{-1}\mathbf{A}^T(\mathbf{b}-\hat{\mathbf{e}})$.

This *cancel-then-recover* approach is not directly applicable in our context. Recall that partial support information is already incorporated into the observation model (5) by substituting the M-dimensional saturation noise $\mathbf{e}$ by the S-dimensional potential saturation noise $\mathbf{s}$. Given artificially saturated measurements, only $(S-E)$ elements of the saturation noise $\mathbf{s}$ are effectively zeros and $\mathbf{s}$ could be thought as a non-negative E-sparse vector with potentially high fraction of sparsity $\frac{E}{S}$. Hence, recovering the corruption term $\mathbf{s}$ using the ℓ_1-norm minimization in (10), is likely suboptimal.

Fortunately, incorporating a sign-constraint into the LS problem, leading to the Non-Negative Least Squares (NNLS) formulation, is shown to be an effective sparsity-promoting regularization [11]. If $\mathbf{b} = \mathbf{A}\mathbf{x} + \mathbf{n}$, and $\mathbf{x} \succeq 0$, then $\hat{\mathbf{x}}^{(\text{NNLS})} = \operatorname{argmin}_{\widetilde{\mathbf{x}}\succeq 0} \|\mathbf{b}-\mathbf{A}\widetilde{\mathbf{x}}\|_2$. Besides, NNLS does not require a tuning parameter. Subsequent hard thresholding of the NNLS solution, yielding the thresholded NNLS (tNNLS) solution $\hat{\mathbf{x}}^{(\text{tNNLS})} = \mathcal{H}_\lambda(\hat{\mathbf{x}}^{(\text{NNLS})})$, where λ is the minimum amplitude prior on sparse $\mathbf{x}$, shows a performance gain with regard to support recovery. Particularly, tNNLS outperforms the threshold non-negative ℓ_1-norm minimization, in the difficult regime characterized by a high sparsity level.

4.2. Proposed Greedy Pursuit Method

We propose to apply the *cancel-then-recover* approach of [10] within the signal estimation step of CoSaMP, by substituting the LS corruption estimation step in (10) by a tNNLS estimation with $\lambda = \frac{\delta}{2}$, where partial knowledge of the saturation noise sparsity pattern is incorporated into sparse $\mathbf{\Theta}$.

At each iteration, once a support candidate is identified, we pretend to cancel out the signal contribution from the measurements, using an orthobasis $\mathbf{Q}$ of the orthogonal complement of the column span of $\mathbf{\Phi}_\Omega$. Then, we aim to estimate the saturation noise, using a NNLS estimation instead of the ℓ_1-norm minimization of (10), as

$$\widehat{\mathbf{s}} \leftarrow \underset{\tilde{\mathbf{s}}\succeq 0}{\operatorname{argmin}} \|\mathbf{Q}^T\mathbf{y} - \mathbf{Q}^T\mathbf{\Theta}\tilde{\mathbf{s}}\|_2^2. \tag{11}$$

By leveraging the sign constraint on the saturation noise, NNLS is expected to perform better, albeit no explicit ℓ_1 regularization is employed, especially given that the corruption term $\mathbf{s}$ is not sufficiently sparse. Moreover, NNLS does not require any parameter tuning whose choice is problematic especially in the first iterations of the greedy pursuit for which the prior signal cancelation step is potentially erroneous. Upon convergence, we estimate the saturation noise support

as $\Sigma = \text{supp}\big(\mathcal{H}_{\delta/2}(\hat{\mathbf{s}}^{(\text{NNLS})})\big)$. Then, we refine the signal estimate, by considering a last run of the *cancel-then-recover* procedure using $\mathbf{\Theta}_\Sigma$ and $\Omega = \text{supp}(\hat{\mathbf{x}})$.
Algorithm 1 summarizes the proposed modified CoSaMP.

Algorithm 1 Proposed Modified CoSaMP

Require: CS matrix $\mathbf{\Phi}$, quantized measurements $\mathbf{y}$, sparsity level K

$\hat{\mathbf{x}} \leftarrow \mathbf{0}, \mathbf{r} \leftarrow \mathbf{y}$ {initialize}
while halting criterion false **do**
  $\mathbf{p} \leftarrow \mathbf{\Phi}^T\mathbf{r}$ {form signal proxy}
  $\Omega \leftarrow \text{supp}(\mathbf{p}_{(2K)}) \cup \text{supp}(\hat{\mathbf{x}})$ {identify signal support}
  $\{\mathbf{b}_\Omega, \hat{\mathbf{s}}\} \leftarrow$ CANCEL-THEN-RECOVER($\mathbf{\Phi}_\Omega, \mathbf{\Theta}, \mathbf{y}$)
  $\mathbf{b}_{\Omega^c} \leftarrow \mathbf{0}$ {estimate signal on candidate support}
  $\hat{\mathbf{x}} \leftarrow \mathbf{b}_{(K)}$ {prune signal support}
  $\mathbf{r} \leftarrow \mathbf{y} - \mathbf{\Theta}\hat{\mathbf{s}} - \mathbf{\Phi}\hat{\mathbf{x}}$ {update residual according to model}
end while
$\Omega \leftarrow \text{supp}(\hat{\mathbf{x}})$ {focus on signal support}
$\Sigma \leftarrow \text{supp}(\mathcal{H}_{\delta/2}(\hat{\mathbf{s}}))$ {prune saturation noise support via t-NNLS}
$\{\hat{\mathbf{x}}_\Omega, \hat{\mathbf{s}}_\Sigma\} \leftarrow$ CANCEL-THEN-RECOVER($\mathbf{\Phi}_\Omega, \mathbf{\Theta}_\Sigma, \mathbf{y}$) {refine estimates}

function CANCEL-THEN-RECOVER($\mathbf{A}, \mathbf{B}, \mathbf{y}$)
  $\mathbf{Q} \leftarrow \text{null}(\mathbf{A}^T)$ {form $\perp$ compl.}
  $\mathbf{s} \leftarrow \text{argmin}_{\tilde{\mathbf{s}} \succeq 0} \|\mathbf{Q}^T(\mathbf{y} - \mathbf{B}\tilde{\mathbf{s}})\|_2^2$ {estimate corruption via NNLS}
  $\mathbf{b} \leftarrow (\mathbf{A}^T\mathbf{A})^{-1}\mathbf{A}^T(\mathbf{y} - \mathbf{B}\mathbf{s})$ {estimate signal via clean LS}
  return $\mathbf{b}, \mathbf{s}$
end function

5. SIMULATION RESULTS AND DISCUSSIONS

In this section, we demonstrate the performance gain of our proposed CoSaMP based recovery method, in comparaison with hard and soft Saturation Consistency CoSaMP based methods [4], denoted here, hSC-CoSaMP and sSC-CoSaMP, respectively. We implement the signal estimation step within hSC-CoSaMP and the proposed method, using the general-purpose convex optimization package CVX [12] and the lsqnonneg MATLAB routine, respectively. In each trial, the $M \times N$ measurement matrix $\mathbf{\Phi}$ is generated from an i.i.d. Gaussian distribution with mean zero and variance $1/M$. The K-sparse signal $\mathbf{x}$, with support selected uniformly at random in $\{1, \dots, N\}$, is drawn from an i.i.d Gaussian distribution and then normalized to have unit ℓ_2-norm. For all experiments, we set $N = 1000$ and $K = 20$ and measure the reconstruction performance by the Recovery Signal-to-Noise-Ratio RSNR $\triangleq 10\log_{10}\left(\frac{\|\mathbf{x}\|_2^2}{\|\mathbf{x}-\hat{\mathbf{x}}\|_2^2}\right) = -20\log_{10}(\|\mathbf{x} - \hat{\mathbf{x}}\|_2)$ where $\hat{\mathbf{x}}$ is the reconstructed signal.

In the first experiment, we consider a noiseless setting, where the quantizer is the unique source of measurements distortion. Figure 1 depicts the RSNR of the three algorithms, averaged over 100 trials, as a function of the saturation level g varied over the range $[0, 0.4]$, for $b = 2, 4$ and $M = 200, 700$. Solid lines, dashed lines and dashed dot lines, follow the scale on the left vertical axis, while dot lines are associated with the right vertical axis. The saturation rate S/M, averaged over 1000 trials, depends on g, δ and also M through the matrix normalization of $\mathbf{\Phi}$. All RSNR

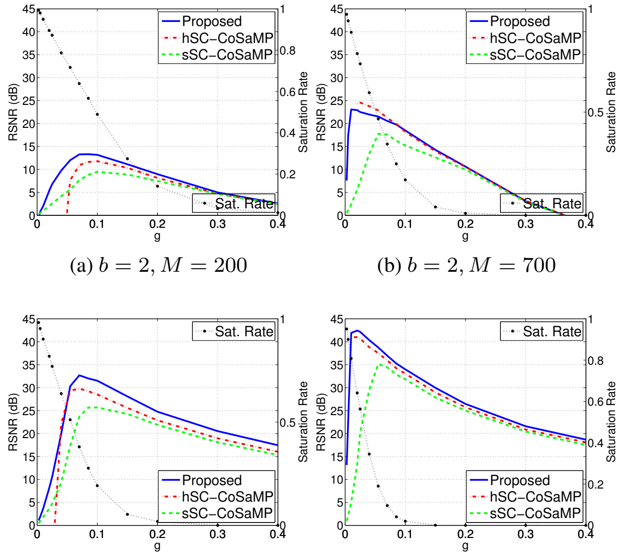

(a) $b = 2, M = 200$ (b) $b = 2, M = 700$ (c) $b = 4, M = 200$ (d) $b = 4, M = 700$

Fig. 1: RSNR versus g, with $N = 1000$, $K = 20$, $b = 4$, and different bit-depths b and measurement regimes M.

curves tend to meet as the saturation rate get closer to zero. Indeed, in the unsaturated quantizer regime, all three methods reduce to the basic CoSaMP method. The three methods achieve their optimal RSNR performances at a nonzero saturation rate, which confirms the benefit of saturated measurements for sparse recovery. The optimal saturation rate for each method, achieves the best tradeoff between model fitting and saturation consistency, depending on how saturated measurements are handled within each reconstruction procedure (model-based, penalty regularization and hard consistency constraint for the proposed method, sSC-CoSaMP and hSC-CoSaMP, respectively). Indeed, an increasing fraction of saturated measurements would not only help strengthening the regularization process in each method by reducing the feasible solution space, but it would also implies a reduced quantization noise for a given bit-depth. At the same time, the saturation rate should remain small enough i.e. the saturation errors should be *sparse* enough, in order to guarantee a sufficient number of the more precise measurements (i.e. those unsaturated) to guarantee more reliable reconstruction. Moreover, the optimal saturation rate increases with the number of measurements M. For instance, for $b = 4$, the proposed method archives a maximum RSNR of 32dB and 42dB, at a saturation rate of 38% and 64%, at the low and high measurement regimes, respectively. The proposed method outperforms the sSC-CoSaMP method with a higher maximum RSNR (around 5dB gain, for $b = 2$), at a higher optimal saturation rate. In other terms, the proposed method better leverages the potential of saturated measurements in the reconstruction process, and provides higher robustness to quantization noise. The proposed method achieves lower

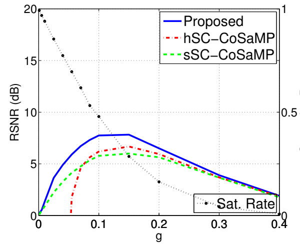

(a) $b = 2$, $M = 200$

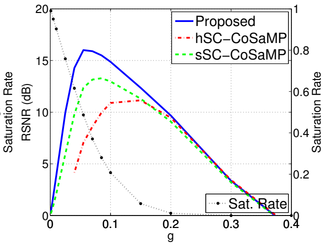

(b) $b = 2$, $M = 700$

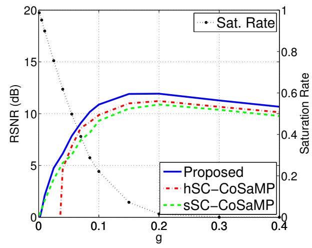

(c) $b = 4$, $M = 200$

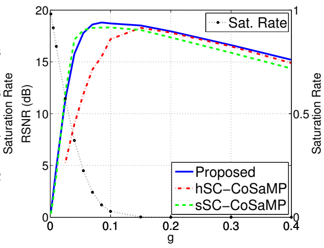

(d) $b = 4$, $M = 700$

Fig. 2: RSNR versus g, with ISNR $= 10$, $N = 1000$, $K = 20$, and different bit-depth b and measurement regimes M.

performance gain over the hSC-CoSaMP method, in terms of maximum RSNR. However, it provides a substantial improvement by maintaining feasibility. Indeed, hSC-CoSaMP is prone to severe failures, due to potential support candidate mis-identification. This is shown by broken curves reflecting potential infeasibility, near its optimal saturation rate, in the high measurement regime, and negative RSNR reflecting local minimum convergence, towards high saturation rates.

In the second experiment, we consider a noisy setting, where an interference $\mathbf{u} \in \mathbb{R}^N$ is present on the input signal. Formally, we acquire $\mathbf{y} = \mathcal{Q}_b(\mathbf{\Phi}(\mathbf{x} + \mathbf{u}))$. We generate the signal noise from an i.i.d. Gaussian distribution with mean zero and variance σ_u^2. We tune the signal noise variance σ_u^2, to obtain a desired Input Signal-to-Noise-Ratio, defined as $\mathrm{ISNR} \triangleq 10\log_{10}\left(\frac{\mathbb{E}[\|\mathbf{x}\|_2^2]}{\mathbb{E}[\|\mathbf{u}\|_2^2]}\right) = -10\log_{10}(N\sigma_u^2)$. Figure 2 depicts the average RSNR of the three algorithms, with the same experimental setup as Fig. 1, except that an ISNR of 10dB is considered. As expected, hSC-CoSaMP shows a dramatic performance decrease and becomes the less reliable method. The proposed method provides the best performances, especially at the high measurement regime and for $b = 2$, where it shows more robustness to coarse quantization, with 3dB and 5dB gain in terms of RSNR in comparaison with sSC-CoSaMP and hSC-CoSaMP, respectively.

6. CONCLUSION

We presented a novel approach to recover sparse signals from their partially saturated QCS measurements. We capitalized on the sign characterization of the saturation noise and its partial support information, to provide a model based greedy pursuit recovery method based on CoSaMP. We demonstrated, by numerical simulations, the performance gain of our proposed method in comparaison with the soft and hard saturation consistency reconstruction methods, in terms of recovery SNR, in the noisless and noisy setting.

REFERENCES

[1] Y. Plan and R. Vershynin, "One-bit compressed sensing by linear programming," *Comm. on Pure and Applied Mathematics*, vol. 66, no. 8, pp. 1275–1297, 2013.

[2] L. Jacques, D. Hammond, and J. Fadili, "Dequantizing compressed sensing: When oversampling and non-gaussian constraints combine," *IEEE Trans. on Information Theory*, vol. 57, no. 1, pp. 559–571, Jan. 2011.

[3] J. Laska and R. Baraniuk, "Regime change: Bit-depth versus measurement-rate in compressive sensing," *IEEE Trans. on Signal Processing,*, vol. 60, no. 7, pp. 3496–3505, July 2012.

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

[5] D. Needell and J. A. Tropp, "Cosamp: Iterative signal recovery from incomplete and inaccurate samples," *App. and Comp. Harm. Anal.*, vol. 26, no. 3, pp. 301–321, 2009.

[6] T. Blumensath and M. E. Davies, "Iterative hard thresholding for compressed sensing," *App. and Comp. Harm. Anal.*, vol. 27, no. 3, pp. 265–274, 2009.

[7] P. T. Boufounos, "Greedy sparse signal reconstruction from sign measurements," in *Proc. of Asilomar Conference on Signals, Systems and Computers*, Nov 2009.

[8] L. Jacques, J. Laska, P. Boufounos, and R. Baraniuk, "Robust 1-bit compressive sensing via binary stable embeddings of sparse vectors," *IEEE Trans. on Information Theory*, vol. 59, no. 4, pp. 2082–2102, April 2013.

[9] L. Jacques, L. K. Degraux, and C. D. Vleeschouwer, "Quantized iterative hard thresholding: Bridging 1-bit and high-resolution quantized compressed sensing," in *Proc. Intl. Conf. Sampling Theory and Applications, (SampTA)*, Germany, March 2013.

[10] E. J. Candes and P. A. Randall, "Highly robust error correction by convex programming," *IEEE Trans. on Information Theory*, vol. 54, no. 7, pp. 2829–2840, 2008.

[11] M. Slawski and M. Hein, "Sparse recovery by thresholded non-negative least squares," in *Proc. of Adv. in Neural Infor. Processing Systems (NIPS)*, Spain, 2011.

[12] M. Grant and S. Boyd, "CVX: Matlab software for disciplined convex programming, version 2.0," available online at http://cvxr.com/cvx, Aug. 2012.