ANALYSIS-BY-SYNTHESIS-BASED QUANTIZATION OF COMPRESSED SENSING MEASUREMENTS

Amirpasha Shirazinia, Saikat Chatterjee, Mikael Skoglund

Communication Theory Lab, ACCESS Linnaeus Centre
KTH Royal Institute of Technology, Stockholm, Sweden
Email: amishi@kth.se, sach@kth.se, skoglund@kth.se

ABSTRACT

We consider a resource-constrained scenario where a compressed sensing- (CS) based sensor has a low number of measurements which are quantized at a low rate followed by transmission or storage. Applying this scenario, we develop a new quantizer design which aims to attain a high-quality reconstruction performance of a sparse source signal based on analysis-by-synthesis framework. Through simulations, we compare the performance of the proposed quantization algorithm vis-a-vis existing quantization methods.

Index Terms— Quantization, compressed sensing, analysis-by-synthesis, sparsity, mean-square error

1. INTRODUCTION

Using a model of under-determined linear set of equations, compressed sensing (CS) [1] aims to reconstruct a high-dimensional sparse source vector (where most of coefficients are zero) from an under-sampled low-dimensional measurement vector. With a limited number of measurements (or a limited resource of sampling), CS has emerged as a new powerful tool for sparse signal acquisition, compression and reconstruction. In many practical applications, CS measurements need to be quantized into a finite resolution representation, and then transmitted to a destination point for sparse signal reconstruction followed by other inference tasks. In this paper, we consider application scenarios where both measurement (or sampling) and transmission resources are constrained. For transmission resource, we mean that the available bits to quantize the CS measurements are limited. Considering availability of limited number of measurements and quantization bits, we design new quantization algorithms where our goal is to achieve high quality sparse signal reconstruction from the quantized CS measurements.

CS with quantized measurements has recently started to gain significant attention in literature, and most commonly, the focus in this area is on three main categories: (1) extensions to existing CS reconstruction algorithms while quantization schemes remain unchanged [2–8]. (2) In the second category, trade-offs between the aspects of quantization (e.g., quantization rate) and CS (e.g., number of measurements and loss in sparse reconstruction) have been considered [5, 9, 10]. (3) In another important class, the main concentration is on quantizer design for CS measurements while CS reconstruction methods are fixed [11–15].

The main contribution of this work is in the third category mentioned above, i.e., quantizer design for CS measurements while CS reconstruction methods are fixed. We develop a new framework for scalar quantization of CS measurements with the objective of achieving a lower *end-to-end reconstruction distortion* rather than *quantization distortion* for CS measurements. Technically, given

a fixed quantizer look-up table and a CS reconstruction algorithm, our proposed algorithm strategically employs a two-stage mechanism in a closed-loop: (1) the synthesis step uses a sparse signal reconstruction technique for measuring the direct effect of quantization of CS measurements on the final sparse signal reconstruction quality, and (2) the analysis step decides appropriate quantized values to maximize the final sparse signal reconstruction quality. This closed-loop strategy is known as *analysis-by-synthesis* (AbS) which has been widely used in multi-media coding [16–18]. To the best of our knowledge, the AbS approach has not been used for quntization of CS measurements, where we show by exploiting this framework, a significantly better reconstruction performance is provided compared to the schemes which only consider quantization distortion, but at the expense of a higher computation. We analyze computational complexity of the proposed algorithm, where it is shown that the complexity depends upon the availability of two compression resources, i.e., quantization bit rate, and number of CS measurements.

Notations: Scalar random variables (RV's) will be denoted by upper-case letters and their instants by the respective lower-case letters. Random vectors will be represented by boldface characters. Further, a set is shown by a calligraphic character and its cardinality by $|\cdot|$. We will also denote the transpose of a vector by by $(\cdot)^T$. We will use $\mathbb{E}[\cdot]$ to denote the expectation operator. The ℓ_p-norm ($p \geq 0$) of a vector will be denoted by $\|\cdot\|_p$.

2. PROBLEM STATEMENT

2.1. Preliminaries of CS Framework

Formally, we let a random sparse (in a fixed basis) signal $\mathbf{X} \in \mathbb{R}^M$ be linearly encoded using a known deterministic sensing matrix $\mathbf{\Phi} \in \mathbb{R}^{N \times M}$ ($N < M$) representing measuring (sampling) system which results in an under-determined set of linear measurements $\mathbf{Y} = \mathbf{\Phi X} \in \mathbb{R}^N$. We let $\mathbf{X}$ be a K-sparse vector, i.e., it has at most K ($K < N$) non-zero coefficients, where the location and magnitude of the non-zero coefficients are drawn randomly from known distributions. We also note that the sparsity level K is known in advance. We define the support set of the sparse vector $\mathbf{X} = [X_1, \ldots, X_M]^T$ by $\mathcal{S} \triangleq \{m : X_m \neq 0\} \subset \{1, \ldots, M\}$ with $|\mathcal{S}| = \|\mathbf{X}\|_0 \leq K$.

In order to estimate a sparse source vector from under-determined linear measurements, several efficient techniques have been developed based on convex optimization (see e.g. [2, 19]), iterative greedy search (see e.g. [20–23]) and Bayesian estimation approaches (see e.g. [24–27]). The results of this paper are generic, and we do not use a specific CS reconstruction algorithm. Denoting a sparse reconstruction function by R, it is defined by a mapping $\mathsf{R} : \mathbb{R}^N \to \mathbb{R}^M$ which takes a (possibly corrupted) measurement vector in N-

dimensional space, and produces an estimate of the sparse source vector in M-dimensional space ($N < M$).

2.2. Quantization of CS Measurements

We consider scalar quantization of the random CS measurements Y_n's ($n = 1, \ldots, N$). For this purpose, quantization is divided into *encoding* and *decoding* tasks. We consider a scalar *quantizer encoder* which maps each measurement to an appropriate index in a finite integer set in order for a *quantizer decoder* to make an estimate of the measurements based on the received index and a known decoding look-up table. We assume that the total bit budget (rate) allocated for quantization is $R_x \triangleq M r_x$ bits per vector $\mathbf{X}$ in which $r_x \in \mathbb{R}^+$ is the assigned quantization rate to a scalar component of $\mathbf{X}$. Having the observations $\mathbf{Y} = \boldsymbol{\Phi} \mathbf{X}$, each entry of the measurement vector, Y_n ($n = 1, \ldots, N$), is encoded via $r_y \triangleq M r_x / N$ bits. For each entry Y_n, a quantizer encoder is defined by a mapping $\mathsf{E} : \mathbb{R} \to \mathcal{I}$, where $\mathcal{I}$ denotes the index set defined as $\mathcal{I} \triangleq \{0, 1, \ldots, 2^{r_y} - 1\}$ with $|\mathcal{I}| = 2^{r_y}$. Denoting the quantized index by the RV I_n ($n = 1, \ldots, N$), the encoder acts according to $Y_n \in \mathcal{R}^{i_n} \Rightarrow I_n = i_n$, where the sets $\{\mathcal{R}^{i_n}\}_{i_n=0}^{2^{r_y}-1}$ are called encoder regions and $\bigcup_{i_n=0}^{2^{r_y}-1} \mathcal{R}^{i_n} = \mathbb{R}$. In words, when Y_n belongs to the region $\mathcal{R}^{i_n}$, the encoder picks the index $i_n \in \mathcal{I}$. Next, we define quantizer decoder which is characterized by a mapping $\mathsf{D} : \mathcal{I} \to \mathcal{C}_n$. The quantizer decoder takes the index I_n, and performs according to an available look-up table; $I_n = i_n \Rightarrow \widehat{Y}_n = c_{i_n}$ such that when the received index is i_n, the decoder outputs the codepoint c_{i_n}. Note that $\widehat{Y}_n$ is the quantized measurement RV associated with the entry Y_n, and the set of all reproduction *codepoints* $\mathcal{C}_n \triangleq \{c_{i_n}\}_{i_n=0}^{2^{r_y}-1}$ associated with this entry is called a *codebook*. We denote by $\widehat{\mathbf{X}} = \mathsf{R}\left([c_{I_1}, \ldots, c_{I_N}]^T\right) \in \mathbb{R}^M$ the estimation of the source from the quantized measurements using a CS reconstruction function R.

2.3. Objective and Performance Criterion

In this paper, we are interested in addressing the following quantizer design problem: Given a CS measurement vector $\mathbf{Y} = \boldsymbol{\Phi} \mathbf{X} \in \mathbb{R}^N$, a CS reconstruction function R and codebook sets $\mathcal{C}_n = \{c_{i_n}\}_{i_n=0}^{2^{r_y}-1}$ ($n = 1, \ldots, N$) for a fixed bit budget R_x, the objective is to find encoding indexes $i_n \in \mathcal{I}$ ($n = 1, \ldots, N$), such that the end-to-end MSE of the estimated vector $\widehat{\mathbf{X}} \in \mathbb{R}^M$, i.e. $\mathbb{E}[\|\mathbf{X} - \widehat{\mathbf{X}}\|_2^2]$, is minimum. In other words, we address the optimization problem

$$\{i_1^\star, \ldots, i_N^\star\} = \underset{\{i_n \in \mathcal{I}\}_{n=1}^N}{\arg\min} \ \mathbb{E}[\|\mathbf{X} - \widehat{\mathbf{X}}\|_2^2], \tag{1}$$

where $\{i_n^\star\}_{n=1}^N$ are the optimal encoding indexes (w.r.t. to minimizing the end-to-end MSE given codebook sets) for quantization of the measurement vector $\mathbf{Y} = [Y_1, \ldots, Y_N]^T$. Also, note that the end-to-end distortion $\mathbb{E}[\|\mathbf{X} - \widehat{\mathbf{X}}\|_2^2]$ depends upon CS reconstruction distortion as well as quantization distortion.

In this paper, our aim is higher than just minimizing the *quantization distortion* $\mathbb{E}[\|\mathbf{Y} - \widehat{\mathbf{Y}}\|_2^2]$ considered in the design of *nearest-neighbor coding* where each measurement entry is coded to its nearest codepoint. The nearest-neighbor coding does not necessarily guarantee that the *end-to-end distortion*, i.e., $\mathbb{E}[\|\mathbf{X} - \widehat{\mathbf{X}}\|_2^2]$, is also minimized subject to fixed codebook sets. This is due to nonlinear behavior of CS reconstruction algorithms and non-orthogonality of the CS system.

Unfortunately, solving (1) jointly for all encoding indexes is not analytically and practically feasible for a generic sparse reconstruction algorithm since it is performed by searching over all possible
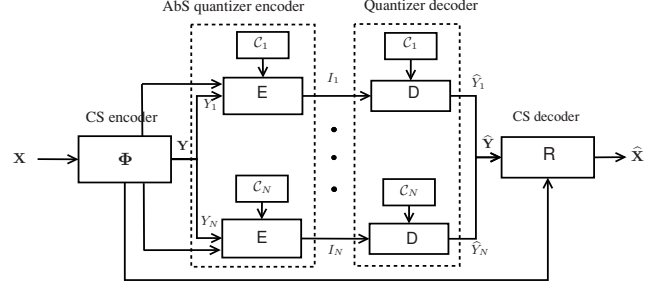

Fig. 1. AbS quantization of CS measurements.

$2^{M r_x}$ codepoints, leading to high complexity. Instead, in this work, we focus on a suboptimal technique for quantization of CS measurements which is computationally efficient and also provides a high-quality reconstruction performance.

3. ANALYSIS-BY-SYNTHESIS QUANTIZATION OF CS MEASUREMENTS

Our proposed AbS-based quantization system is illustrated in Figure 1. In order to feasibly solve (1), we consider optimizing one variable by fixing the others, that is, optimizing the index i_n by fixing the indexes $i_1, \ldots, i_{n-1}, i_{n+1}, \ldots, i_N$ using an *alternating optimization* approach which is potentially suboptimal compared to a joint optimization method, but provides a feasible solution. In what follows, we show how an MSE-minimizing transmission index can be chosen by fixing the others.

3.1. Optimizing Encoding Indexes

Let us first rewrite the end-to-end MSE, $\mathbb{E}[\|\mathbf{X} - \widehat{\mathbf{X}}\|_2^2]$, as (3), where (a) is followed by marginalization over I_n and $\mathbf{Y}$. Also, (b) follows from interchanging the integral and summation and the fact that $\Pr\{I_n = i_n | \mathbf{Y} = \mathbf{y}\} = 1, \forall y_n \in \mathcal{R}^{i_n}$, and otherwise the probability is zero. Moreover, $f(\mathbf{y})$ is the N-fold probability density function (p.d.f.) of the measurement vector. Note that by denoting $\widehat{\mathbf{X}}(I_n) \triangleq \mathsf{R}\left([c_{i_1}, \ldots, c_{I_n}, \ldots, c_{i_N}]\right)^T$, we imply that the reconstructed signal is dependent only upon the index (equivalently codepoint) associated with the n^{th} measurement entry. Now, let denote the *minimum mean square error* (MMSE) estimation of $\mathbf{X}$ given the measurements $\mathbf{Y} = \mathbf{y}$ by

$$\tilde{\mathbf{x}}(\mathbf{y}) \triangleq \mathbb{E}[\mathbf{X} | \mathbf{Y} = \mathbf{y}] \in \mathbb{R}^M, \tag{2}$$

then, given the fixed codebooks $\mathcal{C}_n$ ($n = 1, \ldots, N$), the MSE-minimizing index (assuming other indexes are fixed) is identical to finding the index that minimizes the term in the braces in the last expression of (3) since $f(y)$ is non-negative. The resulting index denoted by $i_n^\star \in \mathcal{I}$ is given by (4), where (a) follows from the fact that $\mathbf{X}$ is independent of I_n conditioned on $\mathbf{Y}$, hence, $\mathbb{E}\left[\|\mathbf{X}\|_2^2 | \mathbf{Y} = \mathbf{y}, I_n = i_n\right] = \mathbb{E}\left[\|\mathbf{X}\|_2^2 | \mathbf{Y} = \mathbf{y}\right]$ which is pulled out of the optimization. Also, (b) follows from a similar rationale, i.e., $\widehat{\mathbf{X}}(I_n)$ is independent of $\mathbf{Y}$ given I_n. Further, $\mathbf{X}$ and $\widehat{\mathbf{X}}(I_n)$ are independent given $\mathbf{Y}$ and I_n.

Following (2), the last expression in (4) can be rewritten as

$$i_n^\star = \arg\min_{i_n \in \mathcal{I}} \left\{ \|\widehat{\mathbf{x}}(i_n)\|_2^2 - 2\tilde{\mathbf{x}}(\mathbf{y})^T \widehat{\mathbf{x}}(i_n) \right\}. \tag{5}$$

One method to predict the MSE-minimizing encoding index is to find the codepoint which after passing through a sparse reconstruction algorithm reproduces a signal vector that is the best estimation to the current input signal vector. Interestingly, (5) implies such an

$$\mathbb{E}[\|\mathbf{X} - \widehat{\mathbf{X}}\|_2^2] = \mathbb{E}[\|\mathbf{X} - \widehat{\mathbf{X}}(I_n)\|_2^2] \stackrel{(a)}{=} \int_{\mathbf{y}} \sum_{i_n} \Pr\{I_n = i_n | \mathbf{Y} = \mathbf{y}\} \mathbb{E}[\|\mathbf{X} - \widehat{\mathbf{X}}(I_n)\|_2^2 | \mathbf{Y} = \mathbf{y}, I_n = i_n] f(\mathbf{y}) d\mathbf{y}$$

$$\stackrel{(b)}{=} \int_{y_1} \cdots \int_{y_{n-1}} \int_{y_{n+1}} \cdots \int_{y_N} \sum_{i_n} \int_{y_n \in \mathcal{R}^{i_n}} \left\{ \mathbb{E}[\|\mathbf{X} - \widehat{\mathbf{X}}(I_n)\|_2^2 | \mathbf{Y} = \mathbf{y}, I_n = i_n] \right\} f(\mathbf{y}) d\mathbf{y} \tag{3}$$

$$i_n^\star = \underset{i_n \in \mathcal{I}}{\arg\min} \, \mathbb{E}[\|\mathbf{X} - \widehat{\mathbf{X}}(I_n)\|_2^2 | \mathbf{Y} = \mathbf{y}, I_n = i_n] \stackrel{(a)}{=} \underset{i_n \in \mathcal{I}}{\arg\min} \left\{ \mathbb{E}[\|\widehat{\mathbf{X}}(I_n)\|_2^2 | \mathbf{Y} = \mathbf{y}, I_n = i_n] - 2\mathbb{E}[\mathbf{X}^T \widehat{\mathbf{X}}(I_n) | \mathbf{Y} = \mathbf{y}, I_n = i_n] \right\}$$

$$\stackrel{(b)}{=} \underset{i_n \in \mathcal{I}}{\arg\min} \left\{ \mathbb{E}[\|\widehat{\mathbf{X}}(I_n)\|_2^2 | I_n = i_n] - 2\mathbb{E}[\mathbf{X}^T | \mathbf{Y} = \mathbf{y}]\mathbb{E}[\widehat{\mathbf{X}}(I_n) | I_n = i_n] \right\} \tag{4}$$

analysis-by-synthesis (AbS) method. We use this principle to first find the optimized encoding index for each measurement entry separately (while others are fixed given the codepoints), and then combine them in an alternate-iterate procedure which will be described in details in the next section. Note that we also assume that the codebook sets are available at the encoder as well as the decoder.

3.2. Proposed Quantization Algorithm

We first describe the framework of the proposed quantization method summarized in Algorithm 1. Suppose that the codebook sets $\mathcal{C}_n$ ($n = 1, \ldots, N$) are designed offline, and let the quantizer encoder have access to the sensing matrix $\mathbf{\Phi}$ and sparsity level K as well as the measurements $\mathbf{y}$ (step (1)). In our formulations (e.g. (5)), the MMSE estimator is required, however, in practice, implementing the MMSE estimator may not be feasible. Therefore, in order to obtain a locally reconstructed vector $\tilde{\mathbf{x}}(\mathbf{y})$, we will approximate the MMSE estimator by the output of the low-complexity greedy *orthogonal matching pursuit* (OMP) [20, 21] reconstruction algorithm (step (2)). Now, we define a dummy vector $\mathbf{z} \in \mathbb{R}^N$ where at the first iteration, its n^{th} component is chosen uniformly at random from the set $\mathcal{C}_n$ ($\forall n$) (step (3)). Indeed, the vector $\mathbf{z}$ implies the predicted quantized measurement which is synthesized at the encoder. Throughout iterations, the entries of the vector $\mathbf{z}$ are adjusted towards the directions of the codepoints (in a sequential manner) which give the minimum reconstruction MSE when retrieved using a CS reconstruction algorithm. Now, we describe the subroutine AbS_seq($\cdot$) executed in Algorithm 1.

Algorithm 1 : AbS-based Quantization

1: **input:** $\mathcal{C}_n = \{c_{i_n}\}_{i_n=0}^{2^{r_y}-1}$ ($\forall n = 1, \ldots, N$) and $\mathbf{\Phi}, \mathbf{y}, K, \gamma$ (stopping threshold)
2: **compute:** $\tilde{\mathbf{x}}(\mathbf{y})$ in (2)
3: **initialize** $\mathbf{z}^{(0)} \in \mathbb{R}^N$, where $z_n^{(0)} \in \mathcal{C}_n, \forall n$.
4: Set $l \leftarrow 0$ (iteration counter)
5: **repeat**
6: $\quad [i_n^\star, \widehat{\mathbf{x}}^{(l+1)}(i_n^\star), \mathbf{z}^{(l+1)}] = \text{AbS\_seq}(\mathcal{C}_n, \tilde{\mathbf{x}}(\mathbf{y}), \mathbf{z}^{(l)}), \forall n$
7: $\quad l \leftarrow l + 1$
8: **until** $\Big| \|\widehat{\mathbf{x}}^{(l)}(i_n^\star)\|_2^2 - 2\tilde{\mathbf{x}}(\mathbf{y})^T \widehat{\mathbf{x}}^{(l)}(i_n^\star)$
$\qquad - \|\widehat{\mathbf{x}}^{(l-1)}(i_n^\star)\|_2^2 + 2\tilde{\mathbf{x}}(\mathbf{y})^T \widehat{\mathbf{x}}^{(l-1)}(i_n^\star) \Big| > \gamma, \forall n$
9: **output:** $I_n = i_n^\star, \widehat{Y}_n = c_{i_n^\star}, \forall n$

Sequential AbS quantization The proposed AbS-based quantization method is summarized in the subroutine AbS_seq($\cdot$) where the main idea is that each measurement entry is sequentially adjusted towards the direction of its MSE-minimizing codepoint at each iteration. Using Algorithm 1, the function AbS_seq($\cdot$) accepts the codebooks $\mathcal{C}_n, \forall n$, the locally reconstructed vector $\tilde{\mathbf{x}}(\mathbf{y})$

and the dummy vector $\mathbf{z}$. At iteration l, the n^{th} ($n = 1, \ldots, N$) entry of $\mathbf{z}^{(l)}$, denoted by $z_n^{(l)}$, is replaced by all 2^{r_y} codepoints from the set $\mathcal{C}_n$ (step (3)) while the other entries are fixed, and the reconstructed vectors, denoted by $\widehat{\mathbf{x}}^{(l)}(i_n) = \mathsf{R}(\mathbf{z}^{(l)})$ ($i_n \in \mathcal{I} = \{0, \ldots, 2^{r_y} - 1\}$), are synthesized corresponding to each vector (step (4)). Then, an optimization is carried out by solving $\arg\min_{i_n \in \mathcal{I}} \|\widehat{\mathbf{x}}^{(l)}(i_n)\|_2^2 - 2\tilde{\mathbf{x}}(\mathbf{y})^T \widehat{\mathbf{x}}^{(l)}(i_n)$ so as to find the wining MSE-minimizing encoding index $i_n^\star$ (step (6)). Now, the n^{th} entry of the vector $\mathbf{z}^{(l)}$ is updated by the codepoint associated with the analyzed index, i.e., $c_{i_n^\star}$ (step (7)). This procedure continues for each entry of $\mathbf{z}^{(l)}$ sequentially, and the subroutine produces the optimized transmission index $i_n^\star$, and the reconstructed vector $\widehat{\mathbf{x}}^{(l)}(i_n^\star)$ as well as the updated quantized vector $\mathbf{z}^{(l)}$ which will be used by the function at the next iteration of Algorithm 1 (step (9)).

Subroutine: AbS_seq $\left(\mathcal{C}_n, \tilde{\mathbf{x}}(\mathbf{y}), \mathbf{z}^{(l)} \right)$

1: **for** $n = 1 : N$ **do**
2: $\quad$ **for** $i = 0 : 2^{r_y} - 1$ **do**
3: $\qquad z_n^{(l)} \leftarrow c_{i_n}$
4: $\qquad$ **compute:** $\widehat{\mathbf{x}}^{(l)}(i_n) = \mathsf{R}(\mathbf{z}^{(l)})$
5: $\quad$ **end for**
6: $\quad i_n^\star = \underset{i_n \in \mathcal{I}}{\arg\min} \|\widehat{\mathbf{x}}^{(l)}(i_n)\|_2^2 - 2\tilde{\mathbf{x}}(\mathbf{y})^T \widehat{\mathbf{x}}^{(l)}(i_n)$
7: $\quad$ **update:** $z_n^{(l)} \leftarrow c_{i_n^\star}(i_n)$
8: **end for**
9: **output:** $i_n^\star, \widehat{\mathbf{x}}^{(l)}(i_n^\star), \mathbf{z}^{(l)}$

Algorithm 1 iterates until convergence where the stopping criterion is that reconstruction improvement at two consecutive iterations is smaller than a predefined threshold $\gamma > 0$. After convergence, the algorithm outputs the transmission indexes I_n's and the quantized CS measurements $\widehat{Y}_n$'s, $\forall n = 1, \ldots, N$, (step (9)).

Now, we analyze the computational complexity of the proposed quantization method. We quantify how many times a CS reconstruction algorithm is invoked throughout the procedures. First, recall from (1) that an exhaustive search for the joint optimization requires $\mathcal{O}(2^{Mr_x})$, or $\mathcal{O}(2^{Nr_y})$ (since $Mr_x = Nr_y$), computations of a CS reconstruction algorithm which is not permissible in practice. Next, let us consider the sequential AbS-based quantization (Subroutine AbS_seq) at one iteration of Algorithm 1. The operations for calculating the transmission indexes increase at most like $\mathcal{O}(2^{Mr_x/N} N)$. This implies that for a fixed bit budget $R_x = Mr_x$, by increasing the number of measurements, first the complexity decreases sharply, and then at some point it starts increasing with a small slope. This is due to the fact that the complexity depends on the compression resources, i.e., number of measurements through the linear term N

and quantization rate r_y through the exponential term $2^{R_x/N}$.

4. EXPERIMENTS AND RESULTS

4.1. Experimental Setups and Results

We quantify the performance using normalized MSE (NMSE) defined as $\text{NMSE} \triangleq \frac{\mathbb{E}[\|\mathbf{X}-\hat{\mathbf{X}}\|_2^2]}{\mathbb{E}[\|\mathbf{X}\|_2^2]}$. In order to measure level of undersampling, we define the measurement rate as $\alpha \triangleq N/M$ ($0 < \alpha \leq 1$). We choose α for given values of sparsity level K and input vector size M, and round the number of measurements N to its nearest integer. We randomly generate a set of K-sparse vector $\mathbf{X}$ where the support set $\mathcal{S}$ is chosen uniformly at random over the set $\{1, 2, \ldots, M\}$. Non-zero coefficients of $\mathbf{X}$ are drawn according to i.i.d. standard Gaussian random variables. We let the elements of the sensing matrix be $\mathbf{\Phi}_{ij} \overset{\text{iid}}{\sim} \mathcal{N}(0, 1/N)$, and then normalize the columns of $\mathbf{\Phi}$ to unit-norm. We apply the OMP reconstruction algorithm as a realization of the CS reconstruction function R.

Using simulation parameters $M = 512$, $K = 35$ (sparsity ratio $\approx 6.8\%$), $r_x = 0.75$ bit per component of $\mathbf{X}$, we have performed 1000 Monte-Carlo simulations to illustrate the performance (NMSE). For implementation of Algorithm 1 using the subroutine AbS_seq, we choose the stopping threshold $\gamma = 10^{-6}$, where we have observed that Algorithm 1 converges in at most 5 iterations. We have compared our proposed quantization method (labeled by *"sequential AbS quantization"* in the figures) with relevant methods such as *"nearest-neighbor coding"* of the measurement entries and *"support set coding"* [28]. Employing the support set coding, the K largest non-zero components of $\tilde{\mathbf{x}}$ (in estimated support set) can be represented by $K \log_2 M$ bits, and then their magnitudes are coded to nearest codepoints using $R_x - K \log_2 M$ bits. Another possible scheme is to quantize each component of the reconstructed signal directly using available bit budget. However, our simulations (not included here) have shown that this coding scheme provides a very poor performance which is excluded for the comparison. Further, we use the same codebooks for each individual scheme which are designed by the *Lloyd algorithm* [29, Chapter 6]. Further, we initialize Algorithm 1 with the same codebook sets for the nearest-neighbor coding. Using the support set coding, a codebook set is designed for a Gaussian scalar component. Hereafter, we make a convention that $\text{NMSE} = 1$ for the support set coding scheme if the total available bit-budget is $R_x < K \log_2 M$ bits.

The performance of the qunatization algorithms as a function of measurement rate α is reported in Figure 2. Let us first consider the nearest-neighbor coding and the proposed AbS-based quantization scheme. From the curves, it can be observed that given a very small values of α, the OMP algorithm fails to detect the sparsity pattern and reconstruct the source which results in a poor performance although the quantization rate per entry is high. As α increases to a certain amount, the reconstruction algorithm succeeds to reconstruct the sparse source precisely out of the measurements since the number of measurements is sufficient, and the quantization error is small enough. At this point ($\alpha = 0.25$ for Figure 2), the curves reach the best performance. However, for higher α's, due to the limited quantization rate r_y, the quantization error per entry increases which leads to a poorer performance. Next, we evaluate the performance of the support set coding where the quantization is performed on the locally reconstructed signal domain. Similarly, at small α's, the OMP reconstruction algorithm fails to reconstruct the locally sparse source, where the performance is insignificant. It can be seen that as α increases and the OMP is able to reconstruct the input signal vector, the performance improves slightly by further increasing measurement rate since the allocated quantization bits using this method are
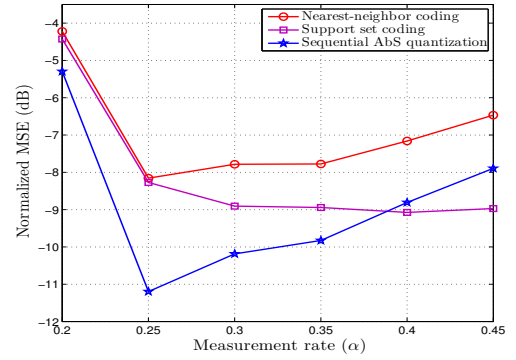

Fig. 2. NMSE as a function of measurement rate ($\alpha = N/M$).
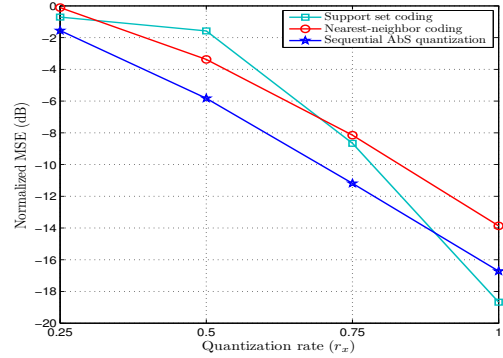

Fig. 3. NMSE as a function of quantization rate r_x.

independent of number of measurements. In the spirit of exploiting CS for practical applications, we are mainly interested in the lower ranges of α (e.g. $\alpha = 0.25$), where the AbS-based quantization scheme achieves a considerable 3 dB reduction in MSE.

Now, we show the performance (NMSE) as a function of quantization rate per entry of $\mathbf{X}$ (i.e., r_x) in Figure 3 using parameters $M = 512$, $K = 35$ at $\alpha = 0.25$. It can be observed that at low to moderate-high rates, the AbS quantization outperforms the other schemes, while at high rates the support set coding attains a slightly better performance. At very high rates, the performance of the support set coding is expected to saturate since, at $\alpha = 0.25$, the distortion due to recovery of the locally reconstructed vector $\tilde{\mathbf{x}}$, according to which the quantization is performed, remains constant. In total, the AbS-based quantization outperforms at operational ranges of measurement and quantization rates where the available resources (number of sensors and quantization rate) are constrained.

5. CONCLUSIONS

We have developed a new framework of quantizer design for CS measurements. We have considered a resource-constrained application where both measurements and transmission rates are limited. Using this scenario, we have addressed the problem of encoding CS measurements where inspired by the AbS framework, a new quantization algorithm has been proposed for coding of linear CS measurements. Numerical results have shown the promising performance gain obtained using this scheme.

5813

6. REFERENCES

[1] E. Candes and M. Wakin, "An introduction to compressive sampling," *IEEE Signal Processing Magazine*, vol. 25, no. 2, pp. 21 –30, Mar. 2008.

[2] E. Candes, J. Romberg, and T. Tao, "Stable signal recovery from incomplete and inaccurate measurements," *Comm. Pure Appl. Math*, vol. 59, no. 8, pp. 1207–1223, 2006.

[3] C. Güntürk, M. Lammers, A. Powell, R. Saab, and O. Ylmaz, "Sigma delta quantization for compressed sensing," in *Annual Conference Inf. Sciences and Systems*, March 2010, pp. 1 –6.

[4] A. Zymnis, S. Boyd, and E. Candes, "Compressed sensing with quantized measurements," *IEEE Sig. Proc. Lett.*, vol. 17, no. 2, pp. 149 –152, Feb. 2010.

[5] W. Dai and O. Milenkovic, "Information theoretical and algorithmic approaches to quantized compressive sensing," *IEEE Trans. Commun.*, vol. 59, no. 7, pp. 1857 –1866, Jul. 2011.

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

[7] M. Yan, Y. Yang, and S. Osher, "Robust 1-bit compressive sensing using adaptive outlier pursuit," *IEEE Trans. Sig. Proc.*, vol. 60, no. 7, pp. 3868 –3875, July 2012.

[8] U. Kamilov, V. K. Goyal, and S. Rangan, "Message-passing estimation from quantized samples," *CoRR*, vol. abs/1105.6368, 2011.

[9] V. Goyal, A. Fletcher, and S. Rangan, "Compressive sampling and lossy compression," *IEEE Sig. Proc. Mag.*, vol. 25, no. 2, pp. 48–56, March 2008.

[10] J. Laska and R. Baraniuk, "Regime change: Bit-depth versus measurement-rate in compressive sensing," *IEEE Trans. Sig. Proc.*, vol. 60, no. 7, pp. 3496 –3505, Jul. 2012.

[11] J. Sun and V. Goyal, "Optimal quantization of random measurements in compressed sensing," in *IEEE Int. Symp. Inf. Theory*, July 2009, pp. 6 –10.

[12] U. Kamilov, V. Goyal, and S. Rangan, "Optimal quantization for compressive sensing under message passing reconstruction," in *IEEE Int. Symp. Inf. Theory*, Aug. 2011, pp. 459 –463.

[13] P. Boufounos, "Universal rate-efficient scalar quantization," *IEEE Trans. Inf. Theory*, vol. 58, no. 3, pp. 1861 –1872, March 2012.

[14] A. Shirazinia, S. Chatterjee, and M. Skoglund, "Channel–optimized vector quantizer design for compressed sensing measurements," in *IEEE Int. Conf. Acoust., Speech, and Sig. Proc. (ICASSP)*, May 2013, to appear.

[15] ——, "Performance bounds for vector quantized compressive sensing," in *Int. Symp. Inf. Theory and App.*, Oct. 2012, pp. 289–293.

[16] P. Kroon and E. Deprettere, "A class of analysis-by-synthesis predictive coders for high quality speech coding at rates between 4.8 and 16 kbit/s," *IEEE Journal Select. Areas Commun.*, vol. 6, no. 2, pp. 353 –363, Feb 1988.

[17] K. Aizawa and T. Huang, "Model-based image coding advanced video coding techniques for very low bit-rate applications," *Proceedings of the IEEE*, vol. 83, no. 2, pp. 259 –271, Feb. 1995.

[18] E. George and M. Smith, "Speech analysis/synthesis and modification using an analysis-by-synthesis/overlap-add sinusoidal model," *IEEE Trans. Speech Audio Proc.*, vol. 5, no. 5, pp. 389 –406, Sep. 1997.

[19] E. Candes and T. Tao, "Rejoinder: the Dantzig selector: statistical estimation when p is much larger than n," *Annals of Statistics*, vol. 35, pp. 2392 – 2404, 2007.

[20] J. Tropp and A. Gilbert, "Signal recovery from random measurements via orthogonal matching pursuit," *IEEE Trans. Inf. Theory*, vol. 53, no. 12, pp. 4655 –4666, Dec. 2007.

[21] T. Blumensath and M. Davies, "Gradient pursuits," *IEEE Trans. Sig. Proc.*, vol. 56, no. 6, pp. 2370 –2382, June 2008.

[22] W. Dai and O. Milenkovic, "Subspace pursuit for compressive sensing signal reconstruction," *IEEE Trans. Inf. Theory*, vol. 55, no. 5, pp. 2230 –2249, May 2009.

[23] S. Chatterjee, D. Sundman, M. Vehkapera, and M. Skoglund, "Projection-based and look-ahead strategies for atom selection," *IEEE Trans. Sig. Proc.*, vol. 60, no. 2, pp. 634 –647, Feb. 2012.

[24] E. G. Larsson and Y. Selen, "Linear regression with a sparse parameter vector," *IEEE Trans. Sig. Proc.*, vol. 55, no. 2, pp. 451 –460, Feb. 2007.

[25] S. Ji, Y. Xue, and L. Carin, "Bayesian compressive sensing," *IEEE Trans. Sig. Proc.*, vol. 56, no. 6, pp. 2346 –2356, June 2008.

[26] M. Elad and I. Yavneh, "A plurality of sparse representations is better than the sparsest one alone," *IEEE Trans. Inf. Theory*, vol. 55, no. 10, pp. 4701 –4714, Oct. 2009.

[27] M. Protter, I. Yavneh, and M. Elad, "Closed-form MMSE estimation for signal denoising under sparse representation modeling over a unitary dictionary," *IEEE Trans. Sig. Proc.*, vol. 58, no. 7, pp. 3471 –3484, July 2010.

[28] A. Shirazinia, S. Chatterjee, and M. Skoglund, "Analysis-by-synthesis quantization of compressed sensing measurements," in *IEEE Trans. Sig. Proc.*, submitted.

[29] A. Gersho and R. M. Gray, *Vector Quantization and Signal Compression*. Kluwer Academic Publishers, 1991.