

Optimal Quantization for Compressive Sensing under Message Passing Reconstruction

Ulugbek Kamilov^{†,‡}, Vivek K Goyal[†], and Sundeep Rangan^{*}

[†]Research Laboratory of Electronics, Massachusetts Institute of Technology

[‡]École Polytechnique Fédérale de Lausanne

^{*}Polytechnic Institute of New York University

kamilov@ieee.org, v.goyal@ieee.org, srangan@poly.edu

Abstract—We consider the optimal quantization of compressive sensing measurements along with estimation from quantized samples using generalized approximate message passing (GAMP). GAMP is an iterative reconstruction scheme inspired by the belief propagation algorithm on bipartite graphs which generalizes approximate message passing (AMP) for arbitrary measurement channels. Its asymptotic error performance can be accurately predicted and tracked through the state evolution formalism. We utilize these results to design mean-square optimal scalar quantizers for GAMP signal reconstruction and empirically demonstrate the superior error performance of the resulting quantizers.

I. INTRODUCTION

By exploiting signal sparsity and smart reconstruction schemes, compressive sensing (CS) [1], [2] can enable signal acquisition with fewer measurements than traditional sampling. In CS, an n -dimensional signal x is measured through m random linear measurements. Although the signal may be undersampled ($m < n$), it may be possible to recover x assuming some sparsity structure.

So far, most of the CS literature has considered signal recovery directly from linear measurements. However, in many practical applications, measurements have to be discretized to a finite number of bits. The effect of such quantized measurements on the performance of the CS reconstruction has been studied in [3], [4]. In [5]–[7] the authors adapt CS reconstruction algorithms to mitigate quantization effects. In [8], high-resolution functional scalar quantization theory was used to design quantizers for lasso reconstruction [9].

The contribution of this paper to the quantized CS problem is twofold: First, for quantized measurements, we propose reconstruction algorithms based on Gaussian approximations of belief propagation (BP). BP is a graphical model-based estimation algorithm widely used in machine learning and channel coding [10], [11] that has also received significant recent attention in compressed sensing [12]. Although exact implementation of BP for dense measurement matrices is generally computationally difficult, Gaussian approximations of BP have been effective in a range of applications [13]–[18]. We consider a recently developed Gaussian-approximated

BP algorithm, called *generalized approximate message passing* [16], [19], that extends earlier methods [15], [18] to nonlinear output channels. We show that the GAMP method is computationally simple and, with quantized measurements, provides significantly improved performance over traditional CS reconstruction based on convex relaxations.

Our second contribution concerns the quantizer design. With linear reconstruction and mean-squared error distortion, the optimal quantizer simply minimizes the mean squared error (MSE) of the transform outputs. Thus, the quantizer can be optimized independently of the reconstruction method. However, when the quantizer outputs are used as an input to a nonlinear estimation algorithm, minimizing the MSE between quantizer input and output is not necessarily equivalent to minimizing the MSE of the final reconstruction [20]. To optimize the quantizer for the GAMP algorithm, we use the fact that the MSE under large random transforms can be predicted accurately from a set of simple state evolution (SE) equations [19], [21]. Then, by modeling the quantizer as a part of the measurement channel, we use the SE formalism to optimize the quantizer to asymptotically minimize distortions after the reconstruction by GAMP.

A longer document discusses both undersampling and oversampling and both regular and non-regular quantization [22]. Code for GAMP is available online [23].

II. NOTATIONS

We use two sets of indices, with $i, j \in \{1, \dots, n\}$ denoting signal components and $a, b \in \{1, \dots, m\}$ measurement components. Bold faced variables \mathbf{x}, \mathbf{y} represent random quantities. The (a, i) entry of the matrix A is denoted by A_{ai} . The Gaussian probability density function with mean μ and variance σ^2 is denoted by $\phi(\cdot; \mu, \sigma^2)$.

III. BACKGROUND

A. Compressive Sensing

In a noiseless CS setting the signal $x \in \mathbb{R}^n$ is acquired via $m < n$ linear measurements of the type

$$z = Ax, \quad (1)$$

where $A \in \mathbb{R}^{m \times n}$ is the *measurement matrix*. The objective is to recover x from (z, A) . Although the system of equations

This material is based upon work supported by the National Science Foundation under Grant No. 0729069 and by the DARPA InPho program through the US Army Research Office award W911-NF-10-1-0404.

formed is underdetermined, the signal is still recoverable if some favorable conditions on x and A are satisfied and exploited. Generally, in CS the common assumption is that the signal is exactly or approximately sparse in some orthonormal basis Ψ , i.e., there is a vector $u = \Psi^{-1}x \in \mathbb{R}^n$ with most of its elements equal or close to zero. Additionally, for certain guarantees on the recoverability of the signal to hold, the matrix A must satisfy the *restricted isometry property* (RIP) [24]. Some families of random matrices, like appropriately-dimensional matrices with i.i.d. Gaussian elements, have been demonstrated to satisfy the RIP with high probability.

A common method for recovering the signal from the measurements is basis pursuit. Although it is possible to solve basis pursuit in polynomial time by casting it as a linear program, its complexity has motivated researchers to look for even cheaper alternatives like numerous recently-proposed iterative methods [12], [16], [17], [25], [26]. Moreover, in real applications a CS reconstruction scheme must be able to mitigate imperfect measurements, due to noise or limited precision [3], [5], [6].

B. Scalar Quantization

A quantizer is a function that discretizes its input by performing a mapping from a continuous set to some discrete set. More specifically, consider K -point regular scalar quantizer Q , defined by its output levels $C = \{c_i; i = 1, 2, \dots, K\}$, decision boundaries $\{(b_{i-1}, b_i) \subset \mathbb{R}; i = 1, 2, \dots, K\}$, and a mapping $c_i = Q(s)$ when $s \in [b_{i-1}, b_i]$ [27]. Additionally define the inverse image of the output level c_i under Q as a *cell* $Q^{-1}(c_i) = [b_{i-1}, b_i)$. For $i = 1$, if $b_0 = -\infty$ we replace the closed interval $[b_0, b_1)$ by an open interval (b_0, b_1) .

Typically quantizers are optimized by selecting decision boundaries and output levels in order to minimize the distortion between the random vector $s \in \mathbb{R}^m$ and its quantized representation $\hat{s} = Q(s)$. For example, for a given vector s and the MSE distortion metric, optimization is performed by solving

$$Q^\# = \arg \min_Q \mathbb{E} \left[\|s - Q(s)\|^2 \right], \quad (2)$$

where minimization is done over all K -level regular scalar quantizers. One standard way of optimizing Q is via the *Lloyd algorithm*, which iteratively updates the decision boundaries and output levels by applying necessary conditions for quantizer optimality [27].

However, for the CS framework finding the quantizer that minimizes MSE between s and \hat{s} is not necessarily equivalent to minimizing MSE between the sparse vector x and its CS reconstruction from quantized measurements \hat{x} [8], [20]. This is due to the nonlinear effect added by any particular CS reconstruction function. Hence, instead of solving (2), it is more interesting to solve

$$Q^* = \arg \min_Q \mathbb{E} \left[\|x - \hat{x}\|^2 \right], \quad (3)$$

where minimization is performed over all K -level regular scalar quantizers and \hat{x} is obtained through a CS reconstruction

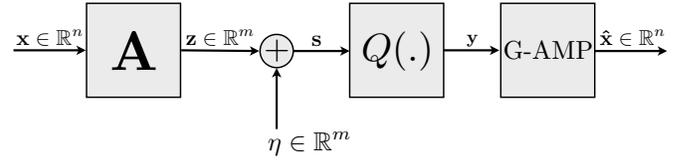


Fig. 1: Compressive sensing set up with quantization of noisy measurements s . The vector z denotes noiseless random measurements.

method like basis pursuit or GAMP. This is the approach taken in this work.

C. Message Passing Algorithms

Consider the problem of estimating a random vector $x \in \mathbb{R}^n$ from noisy measurements $y \in \mathbb{R}^m$, where the noise is described by a measurement channel $p_{y|z}(y_a | z_a)$, which acts identically on each measurement z_a of the vector z obtained via (1). Moreover suppose that elements in the vector x are distributed i.i.d. according to $p_x(x_i)$. Then we can construct the following conditional probability distribution over random vector x given the measurements y :

$$p_{x|y}(x | y) = \frac{1}{Z} \prod_{i=1}^n p_x(x_i) \prod_{a=1}^m p_{y|z}(y_a | z_a), \quad (4)$$

where Z is the normalization constant and $z_a = (Ax)_a$. By marginalizing this distribution it is possible to estimate each x_i . Although direct marginalization of $p_{x|y}(x | y)$ is computationally intractable in practice, we approximate marginals through BP [12], [16], [17]. BP is an iterative algorithm commonly used for decoding of LDPC codes [11]. We apply BP by constructing a bipartite factor graph $G = (V, F, E)$ from (4) and passing the following messages along the edges E of the graph:

$$p_{i \rightarrow a}^{t+1}(x_i) \propto p_x(x_i) \prod_{b \neq a} \hat{p}_{b \rightarrow i}^t(x_i), \quad (5)$$

$$\hat{p}_{a \rightarrow i}^t(x_i) \propto \int p_{y|z}(y_a | z_a) \prod_{j \neq i} p_{j \rightarrow a}^t(x_j) dx, \quad (6)$$

where \propto means that the distribution is to be normalized so that it has unit integral and integration is over all the elements of x except x_i . We refer to messages $\{\hat{p}_{i \rightarrow a}\}_{(i,a) \in E}$ as variable updates and to messages $\{\hat{p}_{a \rightarrow i}\}_{(i,a) \in E}$ as measurement updates. We initialize BP by setting $p_{i \rightarrow a}^0(x_i) = p_x(x_i)$.

Earlier works on BP reconstruction have shown that it is asymptotically MSE optimal under certain verifiable conditions. These conditions involve simple single-dimensional recursive equations called *state evolution* (SE), which predicts that BP is optimal when the corresponding SE admits a unique fixed point [15], [21]. Nonetheless, direct implementation of BP is still impractical due to the dense structure of A , which implies that the algorithm must compute the marginal of a high-dimensional distribution at each measurement node. However, as mentioned in Section I, BP can be simplified

through various Gaussian approximations, including the *relaxed BP* method [15], [16] and *approximate message passing (AMP)* [17], [19]. Recent theoretical work and extensive numerical experiments have demonstrated that, in the case of certain large random measurement matrices, the error performance of both relaxed BP and AMP can also be accurately predicted by SE. Hence the optimal quantizers can be obtained in parallel for both of the methods, however in this paper we concentrate on design for Generalized AMP, while keeping in mind that identical work can be done for Relaxed BP as well.

Due to space limitations, in this paper we will limit our presentation of GAMP and SE equations to the setting in Figure 1. See [19] for more general and detailed analysis.

IV. QUANTIZED GAMP

Consider the CS setting in Figure 1, where without loss of generality we assumed that $\Psi = I_n$. The vector $\mathbf{x} \in \mathbb{R}^n$ is measured through the random matrix \mathbf{A} to result in $\mathbf{z} \in \mathbb{R}^m$, which is further perturbed by some additive white Gaussian noise (AWGN). The resulting vector \mathbf{s} can be written as

$$\mathbf{s} = \mathbf{z} + \boldsymbol{\eta} = \mathbf{A}\mathbf{x} + \boldsymbol{\eta}, \quad (7)$$

where $\{\eta_a\}$ are i.i.d. random variables distributed as $\mathcal{N}(0, \sigma^2)$. These noisy measurements are then quantized by the K -level scalar quantizer Q to give the CS measurements $\mathbf{y} \in \mathbb{R}^m$. The GAMP algorithm is used to estimate the signal \mathbf{x} from the corrupted measurements \mathbf{y} , given the matrix \mathbf{A} , noise variance $\sigma^2 > 0$, and the quantizer mapping Q . Note that under this model each quantized measurement y_a indicates that $s_a \in Q^{-1}(y_a)$, hence our measurement channel can be characterized as

$$p_{\mathbf{y}|\mathbf{z}}(y_a | z_a) = \int_{Q^{-1}(y_a)} \phi(t; z_a, \sigma^2) dt, \quad (8)$$

for $a = 1, 2, \dots, m$ and where $\phi(\cdot)$ is the Gaussian function.

GAMP can be derived by approximating the updates in (5) and (6) by two scalar parameters each and introducing some first-order approximations, as discussed in [19]. Then given the functions F_{in} , \mathcal{E}_{in} , D_1 , and D_2 described below, for each iteration $t = 0, 1, 2, \dots$, the GAMP algorithm produces estimates \hat{x}_i^t of the variables x_i . The estimation is performed as follows:

$$\hat{x}_i^{t+1} \equiv F_{\text{in}} \left(\hat{x}_i^t + \frac{\sum_{a=1}^m A_{ai} u_a^t}{\sum_{a=1}^m A_{ai}^2 \tau_a^t}, \frac{1}{\sum_{a=1}^m A_{ai}^2 \tau_a^t} \right), \quad (9)$$

$$\hat{\tau}_i^{t+1} \equiv \mathcal{E}_{\text{in}} \left(\hat{x}_i^t + \frac{\sum_{a=1}^m A_{ai} u_a^t}{\sum_{a=1}^m A_{ai}^2 \tau_a^t}, \frac{1}{\sum_{a=1}^m A_{ai}^2 \tau_a^t} \right), \quad (10)$$

$$u_a^t \equiv D_1 \left(y_a, \sum_{i=1}^n A_{ai} \hat{x}_i^t - u_a^{t-1} \sum_{i=1}^n A_{ai}^2 \hat{\tau}_i^t, \sum_{i=1}^n A_{ai}^2 \hat{\tau}_i^t + \sigma^2 \right), \quad (11)$$

$$\tau_a^t \equiv D_2 \left(y_a, \sum_{i=1}^n A_{ai} \hat{x}_i^t - u_a^{t-1} \sum_{i=1}^n A_{ai}^2 \hat{\tau}_i^t, \sum_{i=1}^n A_{ai}^2 \hat{\tau}_i^t + \sigma^2 \right), \quad (12)$$

where σ^2 is the variance of the components η_a .

We refer to messages $\{\hat{x}_i, \hat{\tau}_i\}_{i \in V}$ as variable updates and to messages $\{u_a, \tau_a\}_{a \in F}$ as measurement updates. The algorithm is initialized by setting $\hat{x}_i^0 = \hat{x}_{\text{init}}$, $\hat{\tau}_i^0 = \hat{\tau}_{\text{init}}$, and $u_a^{-1} = 0$, where \hat{x}_{init} and $\hat{\tau}_{\text{init}}$ are the mean and variance of the prior $p_{\mathbf{x}}(x_i)$. The nonlinear functions F_{in} and \mathcal{E}_{in} are the conditional mean and variance

$$F_{\text{in}}(q, \nu) \equiv \mathbb{E}[\mathbf{x} | \mathbf{q} = q], \quad (13)$$

$$\mathcal{E}_{\text{in}}(q, \nu) \equiv \text{var}(\mathbf{x} | \mathbf{q} = q), \quad (14)$$

where $\mathbf{q} = \mathbf{x} + \mathbf{v}$, $\mathbf{x} \sim p_{\mathbf{x}}(x_i)$, and $\mathbf{v} \sim \mathcal{N}(0, \nu)$. Note that these functions can easily be evaluated numerically for the given values of q and ν . Similarly, the functions D_1 and D_2 can be computed via

$$D_1(y, \hat{z}, \nu) \equiv \frac{1}{\nu} (F_{\text{out}}(y, \hat{z}, \nu) - \hat{z}), \quad (15)$$

$$D_2(y, \hat{z}, \nu) \equiv \frac{1}{\nu} \left(1 - \frac{\mathcal{E}_{\text{out}}(y, \hat{z}, \nu)}{\nu} \right), \quad (16)$$

where the functions F_{out} and \mathcal{E}_{out} are the conditional mean and variance

$$F_{\text{out}}(y, \hat{z}, \nu) \equiv \mathbb{E}[\mathbf{z} | \mathbf{z} \in Q^{-1}(y)], \quad (17)$$

$$\mathcal{E}_{\text{out}}(y, \hat{z}, \nu) \equiv \text{var}(\mathbf{z} | \mathbf{z} \in Q^{-1}(y)), \quad (18)$$

of the random variable $\mathbf{z} \sim \mathcal{N}(\hat{z}, \nu)$. These functions admit closed-form expressions in terms of $\text{erf}(z) = \frac{2}{\sqrt{\pi}} \int_0^z e^{-t^2} dt$.

V. STATE EVOLUTION FOR GAMP

The equations (9)–(12) are easy to implement, however they provide us no insight into the performance of the algorithm. The goal of SE equations is to describe the asymptotic behavior of GAMP under large measurement matrices. The SE for our setting in Figure 1 is given by the recursion

$$\bar{\tau}_{t+1} = \bar{\mathcal{E}}_{\text{in}} \left(\frac{1}{\bar{D}_2(\beta \bar{\tau}_t, \sigma^2)} \right), \quad (19)$$

where $t \geq 0$ is the iteration number, $\beta = n/m$ is a fixed number denoting the measurement ratio, and σ^2 is the variance of the AWGN components which is also fixed. We initialize the recursion by setting $\bar{\tau}_0 = \hat{\tau}_{\text{init}}$, where $\hat{\tau}_{\text{init}}$ is the variance of x_i according to the prior $p_{\mathbf{x}}(x_i)$. We define the function $\bar{\mathcal{E}}_{\text{in}}$ as

$$\bar{\mathcal{E}}_{\text{in}}(\nu) = \mathbb{E}[\mathcal{E}_{\text{in}}(q, \nu)], \quad (20)$$

where the expectation is taken over the scalar random variable $\mathbf{q} = \mathbf{x} + \mathbf{v}$, with $\mathbf{x} \sim p_{\mathbf{x}}(x_i)$, and $\mathbf{v} \sim \mathcal{N}(0, \nu)$. Similarly, the function \bar{D}_2 is defined as

$$\bar{D}_2(\nu, \sigma^2) = \mathbb{E}[D_2(y, \hat{z}, \nu + \sigma^2)], \quad (21)$$

where D_2 is given by (16) and the expectation is taken over $p_{\mathbf{y}|\mathbf{z}}(y_a | z_a)$ and $(\mathbf{z}, \hat{\mathbf{z}}) \sim \mathcal{N}(0, P_z(\nu))$, with the covariance matrix

$$P_z(\nu) = \begin{pmatrix} \beta \hat{\tau}_{\text{init}} & \beta \hat{\tau}_{\text{init}} - \nu \\ \beta \hat{\tau}_{\text{init}} - \nu & \beta \hat{\tau}_{\text{init}} - \nu \end{pmatrix}. \quad (22)$$

One of the main results of [19] was to demonstrate the convergence of the error performance of the GAMP algorithm to the SE equations under large measurement matrices. The

following theorem is a corollary of this result; see Section VI-D of [19].

Theorem 1. Consider the GAMP algorithm under assumptions in Sections V-B and V-D of [19]. Then for any fixed iteration number t and fixed variable index i , the error variances satisfy the limit

$$\lim_{n \rightarrow \infty} \mathbb{E} \left[|\mathbf{x}_i - \hat{\mathbf{x}}_i^t|^2 \right] = \bar{\tau}_t, \quad (23)$$

where $\bar{\tau}_t$ is the output of the SE equation (19).

Another important result regarding SE recursion in (19) is that it admits at least one fixed point. It has been shown that as $t \rightarrow \infty$ the recursion decreases monotonically to its largest fixed point [16].

It is important to note that the results on convergence of GAMP to SE have been demonstrated for large i.i.d. Gaussian matrices \mathbf{A} [18], [19]. Prior analyses of approximate BP algorithms, like Relaxed BP, were based on certain large sparse assumptions for the matrix \mathbf{A} , which are rarely satisfied in the CS setting [16].

VI. OPTIMAL QUANTIZATION

We now return to the problem of designing MSE-optimal quantizers under GAMP presented in (3). By modeling the quantizer as part of the channel and working out the resulting equations for GAMP and SE, we can make use of the convergence results to recast our optimization problem to

$$Q^{\text{SE}} = \arg \min_Q \{ \bar{\tau}^* \}, \quad (24)$$

where $\bar{\tau}^*$ is the largest fixed point of the SE equations given by

$$\bar{\tau}^* = \lim_{t \rightarrow \infty} \bar{\tau}_t. \quad (25)$$

Minimization is done over all K -level regular scalar quantizers. In practice, about 10 to 20 iterations are sufficient to reach the fixed point of $\bar{\tau}_t$. Then by applying Theorem 1, we know that the asymptotic performance of Q^* will be identical to that of Q^{SE} . It is important to note that the SE recursion behaves well under quantizer optimization. This is due to the fact that SE is independent of actual output levels and small changes in the quantizer boundaries result in only minor change in the recursion (see (18)). Although closed-form expressions for the derivatives of $\bar{\tau}_t$ for large t 's are difficult to obtain, we can approximate them by using finite difference methods. Finally, the recursion itself is fast to evaluate, which makes the scheme in (24) practically realizable using standard optimization methods.

VII. EXPERIMENTAL RESULTS

We now present experimental validation for our results. Assume that the signal \mathbf{x} is generated with i.i.d. elements from the Gauss–Bernoulli distribution

$$\mathbf{x}_i \sim \begin{cases} \mathcal{N}(0, 1/\rho), & \text{with probability } \rho; \\ 0, & \text{with probability } 1 - \rho, \end{cases} \quad (26)$$

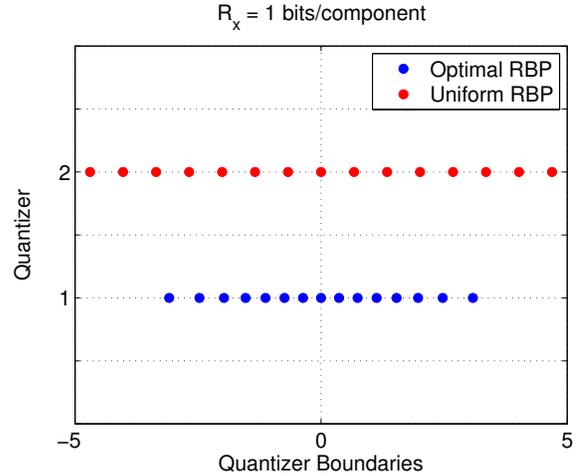


Fig. 2: Optimized quantizer boundaries for 1 bit/component of \mathbf{x} . Optimal quantizer is found by optimizing boundaries for each β and then picking the result with smallest distortion.

where ρ is the sparsity ratio that represents the average fraction of nonzero components of \mathbf{x} . In the following experiments we assume $\rho = 0.1$. We form the measurement matrix \mathbf{A} from i.i.d. Gaussian random variables, i.e., $A_{ai} \sim \mathcal{N}(0, 1/m)$; and we assume that AWGN with variance $\sigma^2 = 10^{-5}$ perturbs measurements before quantization. Related results with no additive noise appear in [22].

Now, we can formulate the SE equation (19) and perform optimization (24). We compare two CS-optimized quantizers: *Uniform* and *Optimal*. We fix boundary points $b_0 = -\infty$ and $b_K = +\infty$, and compute the former quantizer through optimization of type (2). In particular, by applying the central limit theorem we approximate elements s_a of \mathbf{s} to be Gaussian and determine the *Uniform* quantizer by solving (2), but with an additional constraint of equally-spaced output levels. To determine *Optimal* quantizer, we perform (24) by using a standard SQP optimization algorithm for nonlinear continuous optimization.

Figure 2 presents an example of quantization boundaries. For the given bit rate R_x over the components of the input vector \mathbf{x} , we can express the rate over the measurements \mathbf{s} as $R_s = \beta R_x$, where $\beta = n/m$ is the measurement ratio. To determine the optimal quantizer for the given rate R_x we perform optimization for all β s and return the quantizer with the least MSE. As we can see, in comparison with the uniform quantizer obtained by merely minimizing the distortion between the quantizer input and output, the one obtained via SE minimization is very different; in fact, it looks more concentrated around zero. This is due to the fact that by minimizing SE we are in fact searching for quantizers that asymptotically minimize the MSE of the GAMP reconstruction by taking into consideration the nonlinear effects due to the method. The trend of having more quantizer points near zero is opposite to the trend shown in [8] for quantizers optimized for lasso reconstruction.

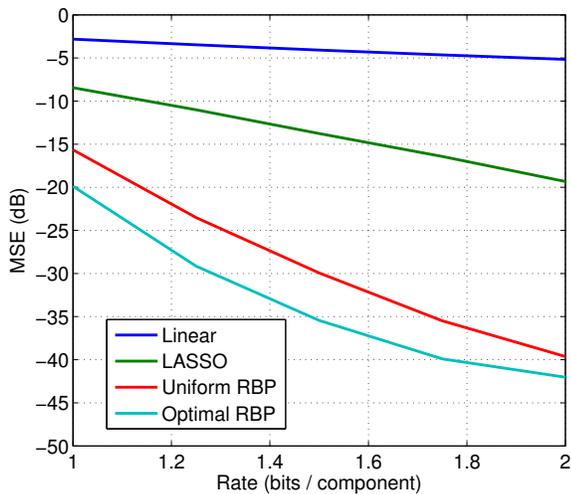


Fig. 3: Performance comparison of GAMP with other sparse estimation methods.

Figure 3 presents a comparison of reconstruction distortions for our two quantizers and confirms the advantage of using quantizers optimized via (19). To obtain the results we vary the quantization rate from 1 to 2 bits per component of \mathbf{x} , and for each quantization rate, we optimize quantizers using the methods discussed above. For comparison, the figure also plots the MSE performance for two other reconstruction methods: linear MMSE estimation and the widely-used lasso method [9], both assuming a bounded uniform quantizer. The lasso performance was predicted by state evolution equations in [19], with the thresholding parameter optimized by the iterative approach in [28]. It can be seen that the proposed GAMP algorithm offers dramatically better performance—more than 10 dB improvement. At higher rates, the gap is still larger although GAMP performance saturates due to the AWGN at the quantizer input. Similarly we can see that the MSE of the quantizer optimized for the GAMP reconstruction is much smaller than the MSE of the standard one, with more than 4 dB difference for many rates.

VIII. CONCLUSIONS

We present generalized approximate message passing as an efficient algorithm for compressive sensing reconstruction from the quantized measurements. We integrate ideas from recent generalization of the algorithm for arbitrary measurement channels to design a method for determining optimal quantizers under GAMP reconstruction. Although computationally simpler, experimental results show that under quantized measurements GAMP offers significantly improved performance over traditional reconstruction schemes. Additionally, performance of the algorithm is further improved by using the state evolution framework to optimize the quantizers.

REFERENCES

[1] E. J. Candès, J. Romberg, and T. Tao, “Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information,” *IEEE Trans. Inform. Theory*, vol. 52, pp. 489–509, Feb. 2006.

[2] D. L. Donoho, “Compressed sensing,” *IEEE Trans. Inform. Theory*, vol. 52, pp. 1289–1306, Apr. 2006.

[3] E. J. Candès and J. Romberg, “Encoding the ℓ_p ball from limited measurements,” in *Proc. IEEE Data Compression Conf.*, (Snowbird, UT), pp. 33–42, Mar. 2006.

[4] V. K. Goyal, A. K. Fletcher, and S. Rangan, “Compressive sampling and lossy compression,” *IEEE Signal Process. Mag.*, vol. 25, pp. 48–56, Mar. 2008.

[5] W. Dai, H. V. Pham, and O. Milenkovic, “A comparative study of quantized compressive sensing schemes,” in *Proc. IEEE Int. Symp. Inform. Theory*, (Seoul, Korea), pp. 11–15, June–July 2009.

[6] A. Zymnis, S. Boyd, and E. Candès, “Compressed sensing with quantized measurements,” *IEEE Signal Process. Lett.*, vol. 17, pp. 149–152, Feb. 2010.

[7] J. N. Laska, P. T. Boufounos, M. A. Davenport, and R. G. Baraniuk, “Democracy in action: Quantization, saturation, and compressive sensing,” *Appl. Comput. Harm. Anal.*, vol. 31, 2011.

[8] J. Z. Sun and V. K. Goyal, “Optimal quantization of random measurements in compressed sensing,” in *Proc. IEEE Int. Symp. Inform. Theory*, (Seoul, Korea), pp. 6–10, June–July 2009.

[9] R. Tibshirani, “Regression shrinkage and selection via the lasso,” *J. Royal Stat. Soc., Ser. B*, vol. 58, no. 1, pp. 267–288, 1996.

[10] J. Pearl, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. San Mateo, CA: Morgan Kaufmann Publ., 1988.

[11] T. J. Richardson and R. L. Urbanke, “The capacity of low-density parity check codes under message-passing decoding,” *IEEE Trans. Inform. Theory*, vol. 47, pp. 599–618, Feb. 2001.

[12] D. Baron, S. Sarvotham, and R. G. Baraniuk, “Bayesian compressive sensing via belief propagation,” *IEEE Trans. Signal Process.*, vol. 58, pp. 269–280, Jan. 2010.

[13] J. Boutros and G. Caire, “Iterative multiuser joint decoding: Unified framework and asymptotic analysis,” *IEEE Trans. Inform. Theory*, vol. 48, pp. 1772–1793, July 2002.

[14] T. Tanaka and M. Okada, “Approximate belief propagation, density evolution, and neurodynamics for CDMA multiuser detection,” *IEEE Trans. Inform. Theory*, vol. 51, pp. 700–706, Feb. 2005.

[15] D. Guo and C.-C. Wang, “Asymptotic mean-square optimality of belief propagation for sparse linear systems,” in *Proc. IEEE Inform. Theory Workshop*, (Chengdu, China), pp. 194–198, Oct. 2006.

[16] S. Rangan, “Estimation with random linear mixing, belief propagation and compressed sensing,” in *Proc. Conf. on Inform. Sci. & Sys.*, (Princeton, NJ), pp. 1–6, Mar. 2010.

[17] D. L. Donoho, A. Maleki, and A. Montanari, “Message-passing algorithms for compressed sensing,” *Proc. Nat. Acad. Sci.*, vol. 106, pp. 18914–18919, Nov. 2009.

[18] M. Bayati and A. Montanari, “The dynamics of message passing on dense graphs, with applications to compressed sensing,” *IEEE Trans. Inform. Theory*, vol. 57, pp. 764–785, Feb. 2011.

[19] S. Rangan, “Generalized approximate message passing for estimation with random linear mixing,” arXiv:1010.5141v1 [cs.IT], Oct. 2010.

[20] V. Misra, V. K. Goyal, and L. R. Varshney, “Distributed scalar quantization for computing: High-resolution analysis and extensions,” *IEEE Trans. Inform. Theory*, vol. 57, Aug. 2011. To appear.

[21] D. Guo and C.-C. Wang, “Random sparse linear systems observed via arbitrary channels: A decoupling principle,” in *Proc. IEEE Int. Symp. Inform. Theory*, (Nice, France), pp. 946–950, June 2007.

[22] U. Kamilov, V. K. Goyal, and S. Rangan, “Message-passing estimation from quantized samples,” arXiv:1105.6368v1 [cs.IT], May 2011.

[23] S. Rangan *et al.*, “Generalized approximate message passing.” SourceForge.net project gampmatlab. Available on-line at <http://gampmatlab.sourceforge.net/>.

[24] E. J. Candès and T. Tao, “Decoding by linear programming,” *IEEE Trans. Inform. Theory*, vol. 51, pp. 4203–4215, Dec. 2005.

[25] A. Maleki and D. L. Donoho, “Optimally tuned iterative reconstruction algorithms for compressed sensing,” *IEEE J. Sel. Topics Signal Process.*, vol. 4, pp. 330–341, Apr. 2010.

[26] J. A. Tropp and S. J. Wright, “Computational methods for sparse solution of linear inverse problems,” *Proc. IEEE*, vol. 98, pp. 948–958, June 2010.

[27] R. M. Gray and D. L. Neuhoff, “Quantization,” *IEEE Trans. Inform. Theory*, vol. 44, pp. 2325–2383, Oct. 1998.

[28] S. Rangan, A. K. Fletcher, and V. K. Goyal, “Asymptotic analysis of MAP estimation via the replica method and applications to compressed sensing,” arXiv:0906.3234v1 [cs.IT], June 2009.