Generalized Approximate Message Passing Estimation from Quantized Samples

Ulugbek Kamilov, Vivek K Goyal, and Sundeep Rangan

Abstract—**Estimation of a vector from quantized linear measurements is a common problem for which simple linear techniques are sometimes greatly suboptimal. This paper summarizes the development of generalized approximate message passing (GAMP) algorithms for minimum mean-squared error estimation of a random vector from quantized linear measurements, notably allowing the linear expansion to be overcomplete or undercomplete and the scalar quantization to be regular or non-regular. GAMP is a recently-developed class of algorithms that uses Gaussian approximations in belief propagation and allows arbitrary separable input and output channels. Scalar quantization of measurements is incorporated into the output channel formalism, leading to the first tractable and effective method for high-dimensional estimation problems involving non-regular scalar quantization. Non-regular quantization is empirically demonstrated to greatly improve rate–distortion performance in some problems with oversampling or with undersampling combined with a sparsity-inducing prior. Under the assumption of a Gaussian measurement matrix with i.i.d. entries, the asymptotic error performance of GAMP can be accurately predicted and tracked through the state evolution formalism.**

I. Introduction

Estimation of a signal from quantized samples arises both from the discretization in digital acquisition devices and the quantization performed for lossy compression. In some settings, much can be gained from treating quantization with care. A key example is analog-to-digital conversion (ADC), where the advantage from oversampling is increased by replacing conventional linear estimation with nonlinear estimation procedures [1]–[6]. Sophisticated approaches are also helpful when using sparsity or compressibility to reconstruct an undersampled signal [7]–[9].

This paper focuses on using a simple message-passing algorithm based on belief propagation (BP). Implementation of BP for estimation of a continuous-valued quantity requires discretization of densities; this is inexact and leads to high computational complexity. To handle quantization effects without any heuristic additive noise model and with low complexity, we use a recently-developed Gaussian-approximated BP algorithm, called *generalized approximate message passing* (GAMP) [10] or *relaxed belief propagation* [11], which extends earlier methods [12], [13] to nonlinear output channels.

GAMP provides significantly-improved performance over traditional methods for estimating from quantized samples. Additionally, when quantizer outputs are used as inputs to a nonlinear estimation algorithm, minimizing the mean-squared error (MSE) between quantizer inputs and outputs is generally not equivalent to minimizing the MSE of the final reconstruction [14]. To optimize the quantizer for the GAMP algorithm, one can use that the MSE under large random mixing matrices $\mathbf{A}$ can be predicted accurately from a set of simple state evolution (SE) equations [10], [15]. Then, by modeling the quantizer as a part of the measurement channel, the SE formalism can be used to optimize the quantizer to asymptotically minimize distortions after the reconstruction by GAMP. Note that our use of random $\mathbf{A}$ is for rigor of the SE formalism; the effectiveness of GAMP does not depend on this. Additional details, especially on SE analysis and quantizer optimization, appear in [16].

II. Quantized Linear Expansions

This paper focuses on the general quantized measurement abstraction of

$$\mathbf{y} = Q(\mathbf{A}\mathbf{x}), \tag{1}$$

where $\mathbf{x} \in \mathbb{R}^n$ is a signal of interest, $\mathbf{A} \in \mathbb{R}^{m \times n}$ is a linear *mixing matrix*, and $Q : \mathbb{R}^m \to \mathbb{R}^m$ is a scalar quantizer. We will be primarily interested in (per-component) MSE $n^{-1}\mathbb{E}[\|\mathbf{x}-\widehat{\mathbf{x}}\|^2]$ for various estimators $\widehat{\mathbf{x}}$ that depend on $\mathbf{y}$, $\mathbf{A}$, and Q. The cases of $m \geq n$ and $m < n$ are both of interest. We sometimes use $\mathbf{z} = \mathbf{A}\mathbf{x}$ to simplify expressions.

A. Overcomplete Expansions

Let $\mathbf{A} \in \mathbb{R}^{m \times n}$ with $m \geq n$ have rank n. Then $\mathbf{A}\mathbf{x}$ is an *overcomplete expansion* of $\mathbf{x}$ and $Q(\mathbf{A}\mathbf{x})$ is a *quantized overcomplete expansion*. Under several reasonable models, linear estimation of $\mathbf{x}$ has MSE inversely proportional to m. More sophisticated algorithms have focused on enforcing *consistency* of an estimate with the quantized samples to yield MSE inversely proportional to m^2, which is the best possible dependence on m [17]. Full consistency is not necessary for optimal MSE dependence on m [4], [6].

A nonlinear estimate may exploit boundedness of the sets

$$\mathcal{S}_i(\mathbf{y}_i) = \{\mathbf{x} \in \mathbb{R}^n \mid q_i(\mathbf{z}_i) = \mathbf{y}_i\}, \qquad i = 1, 2, \ldots, m,$$

which we call *single-sample consistent sets*. Assuming for now that scalar quantizer q_i is regular and its cells are bounded, the boundary of $\mathcal{S}_i(\mathbf{y}_i)$ is two parallel hyperplanes. The full set of hyperplanes obtained for one index i by varying $\mathbf{y}_i$ over the output levels of q_i is called a hyperplane wave partition [17], as illustrated for a uniform quantizer in Fig. 1(a). The set enclosed by two neighboring hyperplanes in a hyperplane wave partition is called a *slab*; one slab is shaded in Fig. 1(a).

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.

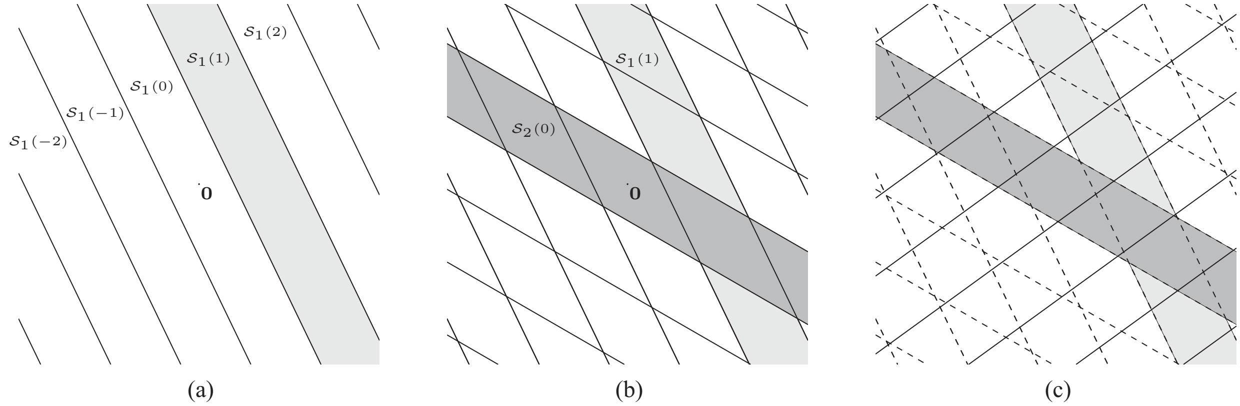

Fig. 1: Visualizing the information present in a quantized overcomplete expansion of $\mathbf{x} \in \mathbb{R}^2$ when each q_i is a regular quantizer. (a) A single hyperplane wave partition with one single-sample consistent set shaded. (b) Partition boundaries from two hyperplane waves; $\mathbf{x}$ is specified to the intersection of two single-sample consistent sets, which is a bounded convex cell. (c) Partition from part (b) in dashed lines with a third hyperplane wave added in solid lines.

Intersecting $\mathcal{S}_i(\mathbf{y}_i)$ for n distinct indexes specifies an n-dimensional parallelotope as illustrated in Fig. 1(b). Using more than n of these single-sample consistent sets restricts $\mathbf{x}$ to a finer partition, as illustrated in Fig. 1(c). The intersection $\mathcal{S}(\mathbf{y}) = \bigcap_{i=1}^{m} \mathcal{S}_i(\mathbf{y}_i)$ is called the *consistent set*.

Quantized overcomplete expansions arise naturally in acquisition subsystems such as ADCs, where m/n represents oversampling factor relative to Nyquist rate. In such systems, high oversampling factor may be motivated by a trade-off between MSE and power consumption or manufacturing cost: within certain bounds, faster sampling is cheaper than a higher number of quantization bits per sample. Ordinarily, the bit-rate inefficiency of the raw output is made irrelevant by recoding, at or near Nyquist rate, soon after acquisition or within the ADC. An alternative explored in this paper is to combat this bit-rate inefficiency through the use of non-regular quantization.

B. Non-Regular Quantization

The bit-rate inefficiency of the raw output with regular quantization is easily understood with reference to Fig. 1(c). With $\mathbf{y}_1$ and $\mathbf{y}_2$ fixed, $\mathbf{x}$ must lie in the intersection of the shaded strips. Only four values of $\mathbf{y}_3$ are possible (the solid hyperplane wave breaks $\mathcal{S}_1(1) \cap \mathcal{S}_2(0)$ into four cells); bits are wasted if this is not exploited in the representation of $\mathbf{y}_3$.

Binning does not change the boundaries of the single-sample consistent sets, but it makes these sets unions of slabs that may not even be connected. Thus, while binning reduces the quantization rate, in the absence of side information that specifies which slab contains $\mathbf{x}$ (at least with moderately high probability), it increases distortion significantly. The increase in distortion is due to *ambiguity* among slabs. Taking $m > n$ quantized samples together may provide adequate information to disambiguate among slabs, thus removing the distortion penalty.

The key concepts in the use of non-regular quantization are illustrated in Fig. 2. Suppose one quantized sample $\mathbf{y}_1$ specifies a single-sample consistent set $\mathcal{S}_1(\mathbf{y}_1)$ composed of two slabs, such as the shaded region in Fig. 2(a). A second quantized sample $\mathbf{y}_2$ will not disambiguate between the two slabs. In the example shown in Fig. 2(b), $\mathcal{S}_2(\mathbf{y}_2)$ is composed of two slabs, and $\mathcal{S}_1(\mathbf{y}_1) \cap \mathcal{S}_2(\mathbf{y}_2)$ is the union of four connected sets. A third quantized sample $\mathbf{y}_3$ may now completely disambiguate; the particular example of $\mathcal{S}_3(\mathbf{y}_3)$ shown in Fig. 2(c) makes $\mathcal{S} = \mathcal{S}_1(\mathbf{y}_1) \cap \mathcal{S}_2(\mathbf{y}_2) \cap \mathcal{S}_3(\mathbf{y}_3)$ a single convex set.

When the quantized samples together completely disambiguate the slabs as in the example, the rate reduction from binning comes with no increase in distortion. The price to pay comes in complexity of estimation. This paper provides a tractable and effective method for reconstruction from a quantized linear expansion with non-regular quantizers.

C. Undercomplete Expansions

Maintaining the quantized measurement model (1), let us turn to the case of $m < n$. We now call $Q(\mathbf{Ax})$ a *quantized undercomplete expansion* of $\mathbf{x}$.

Since the rank of $\mathbf{A}$ is less than n, $\mathbf{A}$ is a many-to-one mapping. Thus, even without quantization, one cannot recover $\mathbf{x}$ from $\mathbf{Ax}$. Rather, $\mathbf{Ax}$ specifies a proper subspace of $\mathbb{R}^n$ containing $\mathbf{x}$; when $\mathbf{A}$ is in general position, the subspace is of dimension $n - m$. Quantization increases the ambiguity in the value of $\mathbf{x}$, yielding consist sets similar to those depicted in Fig. 1(a) and 2(a). However, knowledge that $\mathbf{x}$ is sparse or approximately sparse could be exploited to enable accurate estimation of $\mathbf{x}$ from $Q(\mathbf{Ax})$. See [16] for an illustration.

III. Estimation from Quantized Samples

Generalizing (1), let

$$\mathbf{y} = Q(\mathbf{z} + \mathbf{w}) \qquad \text{where} \qquad \mathbf{z} = \mathbf{Ax}. \tag{2}$$

The input vector $\mathbf{x} \in \mathbb{R}^n$ is random with i.i.d. entries with prior p.d.f. $p_{\mathbf{x}}$. The linear mixing matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$ is random with i.i.d. entries $a_{ij} \sim \mathcal{N}(0, 1/m)$. The (pre-quantization) additive noise $\mathbf{w} \in \mathbb{R}^m$ is random with i.i.d. entries $w_i \sim \mathcal{N}(0, \sigma^2)$.

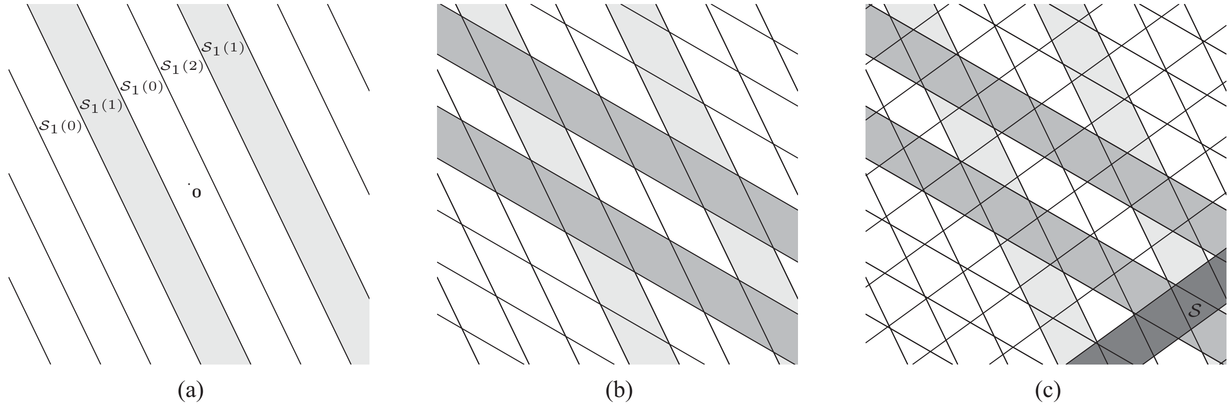

Fig. 2: Visualizing the information present in a quantized overcomplete expansion of $\mathbf{x} \in \mathbb{R}^2$ when using non-regular (binned) quantizers. (a) A single hyperplane wave partition with one single-sample consistent set shaded. Note that binning makes the shaded set not connected. (b) Partition boundaries from two hyperplane waves; $\mathbf{x}$ is specified to the intersection of two single-sample consistent sets, which is now the union of four convex cells. (c) A third sample now specifies $\mathbf{x}$ to within a consistent set $\mathcal{S}$ that is convex.

The scalar quantizer Q has identical component quantizers q_i, each with K output levels.

The estimator $\widehat{\mathbf{x}}$ is a function of $\mathbf{A}$, $\mathbf{y}$, Q, and σ^2. We wish to minimize the MSE $n^{-1}\mathbb{E}[\|\mathbf{x} - \widehat{\mathbf{x}}\|^2]$.

Our primary interest is in the case of $\sigma^2 = 0$, but allowing a nontrivial distribution for $\mathbf{w}$ is not only more general but also makes the derivations more clear.

IV. GAMP for a Quantizer Output Channel

The acquisition model (2) is suitable for GAMP estimation under the conditions in [10] after one simple observation: the mapping from $\mathbf{z}$ to $\mathbf{y}$ is a separable probabilistic mapping with identical marginals. Specifically, quantized measurement $\mathbf{y}_i$ indicates $\mathbf{s}_i \in q_i^{-1}(\mathbf{y}_i)$, so each component *output channel* can be characterized as

$$p_{\mathbf{y}|\mathbf{z}}(y \mid z) = \int_{q_i^{-1}(y)} \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{(t-z)^2}{2\sigma^2}\right) dt.$$

Explicit update rules are given in [16].

V. Experimental Results

Overcomplete Expansions: We first consider overcomplete expansion of $\mathbf{x}$ as discussed in Section II-A. We generate the signal $\mathbf{x}$ with i.i.d. elements from the standard Gaussian distribution $\mathbf{x}_j \sim \mathcal{N}(0, 1)$. We form $\mathbf{A}$ from i.i.d. zero-mean Gaussian random variables. To concentrate on the degradation due to quantization we assume noiseless measurement model (1); i.e., $\sigma^2 = 0$ in (2).

Fig. 3 presents squared-error performance of three estimation algorithm while varying the oversampling ratio m/n and holding $n = 100$. To generate the plot we considered estimation from measurements discretized by a 16-level regular uniform quantizer. For each value of m/n, 100 random realizations of the problem were generated; the curves show the median squared error over these 100 Monte Carlo trials. For comparison to GAMP, we also plot the performance for two other common reconstruction methods: linear MMSE and maximum a posteriori probability (MAP). The MAP estimator was implemented using quadratic programming (QP) and obtains the minimum Euclidean-norm estimate consistent with the quantized samples.

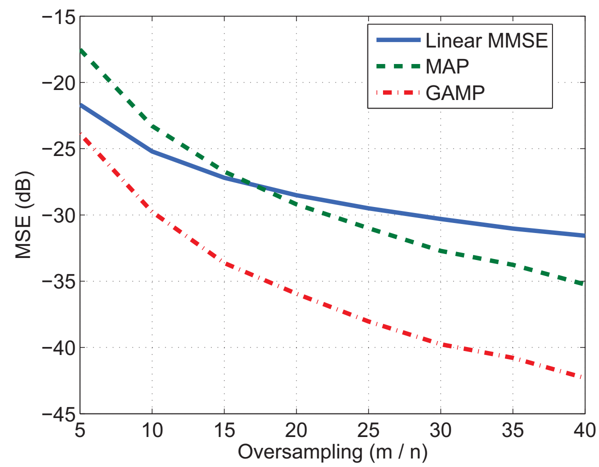

Fig. 3: Performance comparison for oversampled observation of a jointly Gaussian signal vector (no sparsity). GAMP outperforms linear MMSE and MAP estimators.

We see that GAMP offers significantly better performance with more than 5 dB improvement for many values of m/n. MAP provides poor performance compared to GAMP because it finds a corner of the consistent set, which is suboptimal as compared to the centroid of the consistent set.

Compressive Sensing with Quantized Measurements: We now would like to estimate a sparse signal $\mathbf{x}$ from $m < n$ random measurements. We assume that the signal $\mathbf{x}$ is generated with i.i.d. elements from the Gauss–Bernoulli

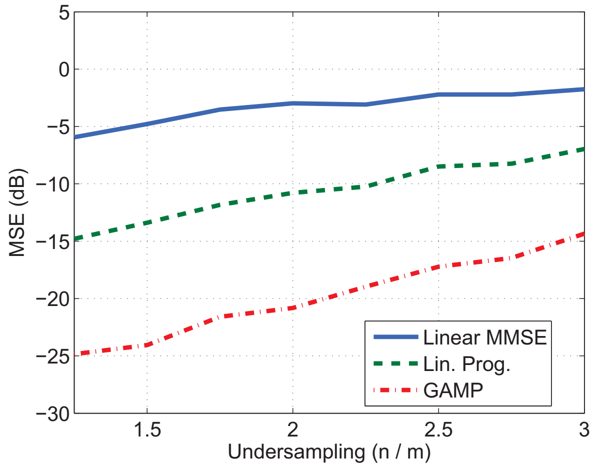

Fig. 4: Performance comparison of GAMP with other sparse estimation methods.

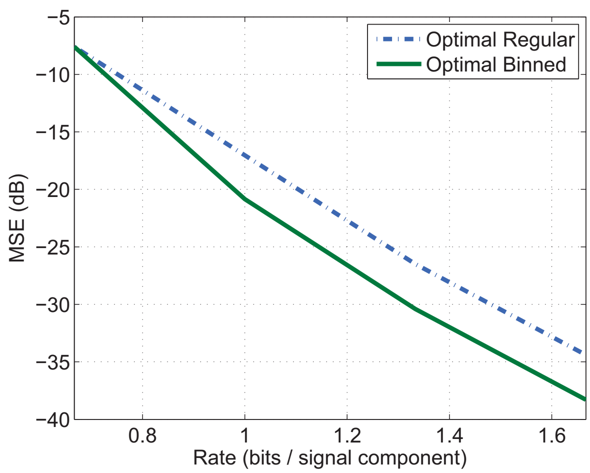

Fig. 5: Performance comparison of GAMP with regular and binned quantizers under sparse Gauss-Bernoulli prior.

distribution

$$\mathbf{x}_j \sim \begin{cases} \mathcal{N}(0, 1/\rho), & \text{with probability } \rho; \\ 0, & \text{with probability } 1-\rho, \end{cases} \quad (3)$$

where ρ is the sparsity ratio that represents the average fraction of nonzero components of $\mathbf{x}$. In the following experiments we assume $\rho = 0.1$. Similarly to overcomplete case, we form the measurement matrix $\mathbf{A}$ from i.i.d. Gaussian random variables, i.e., $A_{ij} \sim \mathcal{N}(0, 1/m)$; and we assume no additive noise ($\sigma^2 = 0$ in (2)).

Fig. 4 compares squared-error performance of GAMP with two other sparse estimation methods. We obtain the curves by varying the undersampling ratio $\beta = n/m$ and holding $m = 100$. We perform estimation from measurements obtained from a 16-level regular uniform quantizer. The figure plots the median of squared error values from 100 Monte Carlo trials for each value of β. The top curve (worst performance) is for linear MMSE estimation; and the middle curve is for the basis pursuit estimator, $\widehat{\mathbf{x}} = \arg\min_{\mathbf{x}\,:\,\mathbf{y}=Q(\mathbf{Ax})} \|\mathbf{x}\|_1$, which can be cast and solved as a linear program (LP). Again GAMP offers substantial improvement over the other methods.

Non-Regular Quantization: We now repeat the case of the undersampled sparse signal by again using $\mathbf{x}$ with an i.i.d. Gauss–Bernoulli distribution with $\rho = 0.1$. To study the effect of non-regular quantization, we introduce a binning function that reduces the rate of the quantizer by 1 bit per sample (halving the number of output levels) by performing reduction modulo 2.

Fig. 5 plots the MSE performance of GAMP under regular and non-regular quantization. We obtain the plot by varying the number of bits per component of $\mathbf{x}$ and holding $\beta = 3$. We see that, in comparison to regular quantizers, binned quantizers with GAMP estimation achieve much lower distortions for the same rates.

References

[1] N. T. Thao and M. Vetterli, "Reduction of the MSE in R-times oversampled A/D conversion from $O(1/R)$ to $O(1/R^2)$," *IEEE Trans. Signal Process.*, vol. 42, no. 1, pp. 200–203, Jan. 1994.

[2] ——, "Deterministic analysis of oversampled A/D conversion and decoding improvement based on consistent estimates," *IEEE Trans. Signal Process.*, vol. 42, no. 3, pp. 519–531, Mar. 1994.

[3] V. K. Goyal, M. Vetterli, and N. T. Thao, "Quantized overcomplete expansions in $\mathbb{R}^N$: Analysis, synthesis, and algorithms," *IEEE Trans. Inform. Theory*, vol. 44, no. 1, pp. 16–31, Jan. 1998.

[4] S. Rangan and V. K. Goyal, "Recursive consistent estimation with bounded noise," *IEEE Trans. Inform. Theory*, vol. 47, no. 1, pp. 457–464, Jan. 2001.

[5] J. J. Benedetto, A. M. Powell, and Ö. Yilmaz, "Sigma–Delta ($\Sigma\Delta$) quantization and finite frames," *IEEE Trans. Inform. Theory*, vol. 52, no. 5, pp. 1990–2005, May 2006.

[6] A. M. Powell, "Mean squared error bounds for the Rangan–Goyal soft thresholding algorithm," *Appl. Comput. Harm. Anal.*, vol. 29, no. 3, pp. 251–271, Nov. 2010.

[7] A. Zymnis, S. Boyd, and E. Candès, "Compressed sensing with quantized measurements," *IEEE Signal Process. Lett.*, vol. 17, no. 2, pp. 149–152, Feb. 2010.

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

[9] 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.

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

[11] ——, "Estimation with random linear mixing, belief propagation and compressed sensing," in *Proc. Conf. on Inform. Sci. & Sys.*, Princeton, NJ, Mar. 2010, pp. 1–6.

[12] 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, Oct. 2006, pp. 194–198.

[13] M. Bayati and A. Montanari, "The dynamics of message passing on dense graphs, with applications to compressed sensing," *IEEE Trans. Inform. Theory*, vol. 57, no. 2, pp. 764–785, Feb. 2011.

[14] 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, no. 8, pp. 5298–5325, Aug. 2011.

[15] 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, Jun. 2007, pp. 946–950.

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

[17] N. T. Thao and M. Vetterli, "Lower bound on the mean-squared error in oversampled quantization of periodic signals using vector quantization analysis," *IEEE Trans. Inform. Theory*, vol. 42, no. 2, pp. 469–479, Mar. 1996.