

Quantization of Sparse Representations

Petros Boufounos, Richard Baraniuk*

{petrosb,richb}@rice.edu

Rice University, ECE Dept., MS 380

P.O. Box 1892, Houston, TX 77251-1892

Abstract

Compressive sensing (CS) is a new signal acquisition technique for sparse and compressible signals. Rather than uniformly sampling the signal, CS computes inner products with randomized basis functions; the signal is then recovered by a convex optimization. Random CS measurements are universal in the sense that the same acquisition system is sufficient for signals sparse in any representation. This paper examines the effect of quantization of CS measurements. A careful study of strictly sparse, power-limited signals concludes that CS with scalar quantization does not use its allocated rate efficiently. The inefficiency, which is quantified, can be interpreted as the price that must be paid for the universality of the encoding system. The results in this paper complement and extend recent results on the quantization of compressive sensing measurements of compressible signals.

1 Introduction

Sparse representations have received significant attention in the signal processing literature. Their importance is highlighted in applications to transform coding and signal compression. In these applications, the signal is acquired using a high-rate sampling process. The samples are subsequently transformed to a sparse linear representation, in which most of the representation coefficients are zero or close to zero. This sparsity is exploited in the encoding process to reduce the rate required to store and transmit the signal.

Recent results in the area of compressive sensing demonstrate that signals with sparse or approximately sparse representations can also be efficiently acquired [1–4]. Classical sampling theory suggests that sampling and reconstruction of a signal in

*Rice ECE Dept. Technical Report 0701; Summary to appear in Proc. DCC '07, Snowbird, UT. Research supported by ONR grants N00014-06-1-0768 and N00014-06-1-0829; AFOSR grant FA9550-04-0148; DARPA grants N66001-06-1-2011 and N00014-06-1-0610; NSF grants CCF-0431150, CNS-0435425, and CNS-0520280; and the Texas Instruments Leadership University Program.

a vector space requires at least as many linear measurements of the signal as the dimensionality of the space. Compressive sensing significantly reduces the number of measurements required to capture and subsequently reconstruct the signal, as long as it has a sparse representation or can be well approximated by one. The results also guarantee that the signals can be recovered sufficiently well, even in the presence of measurement error and distortion such as quantization. One of the advantages of compressive sensing is the universality property, which guarantees that with proper random selection of the measurement space, the signal can be recovered with very high probability, even if the representation in which the signal is sparse is not known in advance of the measurements [2].

The sparsity models used in the compressive sensing literature vary, depending on the application. A sparsity dictionary is always assumed, which is a collection of vectors—usually an orthonormal basis—in which the signal can be sparsely represented. One common signal model assumes that the signal is power limited and sparse, i.e. has very few non-zero coefficients when represented using the dictionary. Other common models assume that the representation is compressible, i.e. that very few coefficients are sufficient to approximate the signal well enough. The most common model for compressible signals assumes that the representation lies in an ℓ_p ball, i.e. has bounded ℓ_p norm for some $p \leq 1$. Under this assumption, which guarantees compressibility, the quality of the approximation depends on p and the number of representation coefficients used to approximate the signal.

The quantization of compressible signals with bounded ℓ_p norm in the context of compressive sensing has been studied in [5]. Specifically, lower bounds for the quantization error are presented, applicable to any quantized representation. These bounds are compared with the performance of compressive sensing followed by scalar quantization and non-linear reconstruction. It is thus demonstrated that the performance of a compressive sensing universal encoding method is within a logarithmic factor of the optimal encoding. This is a small price to pay for the universality of the algorithm.

The main contribution of this paper is to examine the quantization of strictly sparse, power limited signals in terms of the quantization error and the rate-use efficiency of compressive sensing. Although the results presented extend and complement the work on the quantization error presented in [5], sparse signals behave very differently than compressible signals with bounded ℓ_p norm. The strictly sparse signal model is better suited to describe man-made signals, such as radar and communication signals. It is also a good model for signals that have already been encoded and subsequently decoded using classical transform coding methods, and have thus become sparse in the process.

After presenting some background and establishing the notation in Sec. 2, it is shown in Sec. 3.1 that compressive sensing measurements of sparse signals use only a fraction of the available quantization points to encode the signal. This implies that the encoding bits are inefficiently used. This inefficiency, which is examined and quantified in Sec. 3.2, can be considered as the cost to achieve the universality

of compressive sensing encoding, or a target for elimination by subsequent lossless entropy coding.

Using the bounds on the quantization efficiency, a lower bound for the quantization error is derived in Sec. 3.3. It thus follows that compressive sensing and quantization of sparse signals incur a penalty in the error, compared to the optimal error achievable using transform coding. This penalty can also be considered as the cost of universality. It is also shown that reconstruction using linear programming incurs a further penalty in the achievable error. The results in this paper are used in Sec. 3.4 to improve the lower bounds presented in [5], specifically in the case of low rate quantization and highly compressible signals.

2 Background

This section presents some background on sampling, quantization, and compressive sensing. The main purpose is to establish the notation and serve as a quick reference.

2.1 Sampling and Quantization

For the purposes of this paper we assume a signal, denoted using \mathbf{x} , that lies in an N dimensional space, denoted using \mathcal{S} , and referred to as the signal space. Sampling is performed by taking measurements using inner products of \mathbf{x} with an orthonormal sampling basis $\{\phi_k | k = 1, \dots, M\}$. The sampling basis spans the M -dimensional sampling space \mathcal{W} , which is a subspace of \mathcal{S} . The sampling operator is defined as:

$$\Phi : \mathcal{S} \rightarrow \mathbb{R}^M, \text{ s.t. } (\Phi \mathbf{x})_k = \langle \mathbf{x}, \phi_k \rangle, \quad (1)$$

in which $(\cdot)_k$ is the k^{th} component of a vector in \mathbb{R}^M . The vector of sampling coefficients is denoted using $\mathbf{y} = \Phi \mathbf{x}$. Correspondingly, the reconstruction function is denoted using:

$$\mathbf{R}(\cdot), \text{ s.t. } \mathbf{R} : \mathbb{R}^M \rightarrow \mathcal{S}. \quad (2)$$

The parenthetical notation is used as a reminder that the reconstruction is not necessarily linear.

In order to process and transmit the sampling coefficients, they are quantized to a finite precision. It is assumed that each coefficient is mapped to one of L quantization levels, in which L is finite. Unless otherwise noted, most of the results require no assumption on the quantization intervals. In some cases uniform quantization is assumed and Δ is used to denote the quantization interval length.

The quantization of the coefficients using an L -level quantizer provides a possible of L^M combinations of levels for each coefficients to represent the vector, which corresponds to L^M possible quantization points. Assuming no subsequent entropy coding, such a representation requires at least $B \geq M \log_2 L$ bits to represent the coefficients. Alternatively, given a fixed budget of B bits, assuming that all the bits are equally distributed to all the coefficients, then the quantizer can have at most $L \leq 2^{B/M}$ levels per coefficient. For exposition simplicity, we assume that B and L are integers such that $L = 2^{B/M}$.

Scalar quantization of the sampling coefficients partitions \mathbb{R}^M —and, correspondingly, \mathcal{S} and \mathcal{W} —into L^M cells which contain the quantization points. The edges of each cell in \mathbb{R}^M and \mathcal{W} are orthogonal to each other, with dimensions that are implied by the quantization intervals of the scalar quantizer. The quantization levels of the quantizer determine the quantization points inside each cell. Any vector \mathbf{y} contained in a cell is quantized to the quantization point $\hat{\mathbf{y}}$ corresponding to that particular cell. For example, the cells of a uniform quantizer are all hypercubes of size Δ^M , with the corresponding quantization points at the center of each hypercube. The metric of interest is the ℓ_2 norm of the reconstruction error using the quantized coefficients, i.e. the magnitude $\|\mathbf{x} - \mathbf{R}(\hat{\mathbf{y}})\|_2$. A two-dimensional example is shown in Fig. 1(a).

Depending on the formulation, the performance of the quantizer is commonly measured using either the expected value or the worst case quantization error over all the vectors of interest. The signal is assumed to lie in a K -dimensional subspace and, without loss of generality, it is power-limited by $\|\mathbf{x}\|_2 \leq 1$. Under these assumptions the minimum worst-case error, and the minimum mean square error of a vector quantizer can be derived for any quantizer that uses P cells to partition the space. Using the minimax approach as in [5, 6] the worst case error magnitude can be shown to be greater than:

$$\max_{\mathbf{x}} \|\epsilon\|_2 \geq P^{-1/K}, \quad (3)$$

for any choice of vector quantizer. Similarly, using Zador’s formula [7, 8], the minimum mean squared error for any quantizer can be shown to be:

$$E\{\|\epsilon\|^2\} = \Omega(P^{2/K}), \quad (4)$$

in which the constants only depend on the distribution of the signal in the subspace. Within the constant of proportionality, both results are essentially the same. Thus, the minimax approach is used in the remainder of this paper. It should be noted that an assumption in this work, consistent with the minimax approach, is that the locations of the sparse coefficients are uniformly distributed. Often correlations in the locations of the sparse coefficients can reduce the rates required to encode the signals. In such cases, more general methods, such as the information theoretic rate-distortion analysis introduced in [9] can be used to analyze the system performance.

2.2 Sparse Representations and Compressive Sensing

Signals of interest in signal processing can often be represented or well approximated using sparse representations in a particular basis. In this paper it is assumed that the signals of interest, $\mathbf{x} \in \mathcal{S}$ are all the K -sparse signals in a known orthonormal basis $\{\psi_1, \dots, \psi_N\}$, such that the sum in the representation $\mathbf{x} = \sum_i \langle \mathbf{x}, \psi_i \rangle \psi_i$ only has K non-zero terms. It follows that the signals of interest lie in a collection of K -dimensional subspaces.

Recent advances in compressive sensing [1–4] demonstrate that sparse signals can be sampled and exactly reconstructed using fewer than N linear measurements. The signal \mathbf{x} is sampled using M measurements with the orthogonal basis $\{\phi_1, \dots, \phi_M\}$ to




Figure 1: Example of uniform scalar quantization in a two dimensional sampling space. (a) Any vector \mathbf{y} in a quantization cell is quantized to the center of the cell $\hat{\mathbf{y}}$. (b) The one-dimensional subspace \mathcal{H} intersects very few of the quantization cells.

produce the vector $\mathbf{y} = \Phi \mathbf{x}$, as described in Sec. 2.1. It is possible to recover \mathbf{x} using a convex program if the sampling operator Φ satisfies a restricted isometry property (RIP), which is satisfied with very high probability if $M = cK \log_2(N/K)$ random measurement vectors are used [4]. If the RIP holds, then the vector $\hat{\mathbf{x}}$ that satisfies:

$$\hat{\mathbf{x}} := \min \|\tilde{\mathbf{x}}\|_{\ell_1} \text{ s.t. } \Phi \tilde{\mathbf{x}} = \mathbf{y}, \quad (5)$$

is equal to the original vector \mathbf{x} . Measurements can also be perturbed, for example due to measurement noise or quantization, by an error vector \mathbf{e} , in which case the measured signal is $\mathbf{y} = \Phi \mathbf{x} + \mathbf{e}$. If the error vector is bounded by $\|\mathbf{e}\|_{\ell_2} \leq \delta$ then the signal can be approximately recovered using the following convex program:

$$\hat{\mathbf{x}} := \min \|\tilde{\mathbf{x}}\|_{\ell_1} \text{ s.t. } \|\Phi \tilde{\mathbf{x}} - \mathbf{y}\|_{\ell_2} \leq \delta. \quad (6)$$

In this case, the recovered signal $\hat{\mathbf{x}}$ satisfies $\|\mathbf{x} - \hat{\mathbf{x}}\| \leq C\delta$ for some constant C that only depends on the sampling operator and not on the signal \mathbf{x} . For the remainder of this paper the measurement error is attributed to quantization. The performance in the presence of noise can only be worse. A study of the performance in the presence of noise, although interesting, is beyond the scope of this paper.

3 Quantization and Compressive Sensing

This section examines the scalar quantization of compressive sensing measurements. Section 3.1 presents an upper bound on the number of quantization cells intersected in the sampling process. Sections 3.2 and 3.3 use the bound to evaluate the rate efficiency and the error performance, respectively. Section 3.4 applies the error bounds to the case of compressible signals.

3.1 Sampling and Quantization of Subspaces

The sampling and quantization of low dimensional signal subspaces using higher dimensional sampling spaces has been studied in the context of frame representations

and oversampling [6, 8]. The bounds derived in both references are similar, but the approach in [6] can be immediately applied to the quantization of sparse signals. Scalar quantization of a basis expansion in an M -dimensional space \mathcal{W} using an L -level scalar quantizer generates an orthogonal quantization lattice with L^M vector quantization cells. It is shown in [6] that a K -dimensional subspace of \mathcal{W} intersects a limited number of cells, denoted using $I_o(M, K, L)$, which is upper bounded by:

$$I_o(M, K, L) \leq (2L)^K \binom{M}{K}. \quad (7)$$

It thus follows that any cells outside the ones intersected are never used by the quantizer. Figure 1(b) presents an example for the case of $M = 2$ and $K = 1$, in which the one-dimensional subspace \mathcal{H} only intersects very few quantizations cells. The bound in Eq. (7) can be immediately used in Eqs. (3) and (4) to derive a lower bound on the reconstruction error for any signal in the subspace.

In the case of sparse signals, this result is immediately applicable by realizing that K -sparse signals occupy a set of $\binom{N}{K}$ subspaces in the signal space \mathcal{S} . It follows that the projection implied in the sampling operation produces at most $\binom{N}{K}$ of subspaces of dimension at most K in the sampling space \mathcal{W} . Thus the number of cells intersected, denoted using $I_N(M, K, L)$, is at most:

$$I_N(M, K, L) \leq (2L)^K \binom{M}{K} \binom{N}{K}. \quad (8)$$

To minimize the maximum error, these quantization levels are equally distributed to each of the subspaces, and therefore at most $I_o(M, K, L)$ quantization points should be used to encode each of the subspaces. It should be noted that this is equivalent to the assumption that the K components of the sparse signal are uniformly chosen from the N possible components, and the use of Zador's formula in Eq. (4) to lower bound the achievable mean squared error.

3.2 Rate Penalty for Compressive Sensing

A consequence of the discussion above is that there is a rate penalty associated with scalar quantization of samples of sparse signals. This section quantifies that penalty and examines the asymptotic behavior in a variety of asymptotic assumptions.

Assuming a fixed number of quantization levels L , the rate used for transmission is at least $B = M \log_2 L$. On the other hand, only $I_N(M, K, L)$ cells are intersected, which implies that the rate required is at most $\log_2(I_N(M, K, L))$. Thus, an upper bound on the effective bit use of the quantizer—which is also a measure of the quantizer efficiency—can be derived from Eq (8):

$$\frac{\log_2(I_N(M, K, L))}{M \log_2 L} \leq \frac{\log_2((2L)^K \binom{M}{K} \binom{N}{K})}{M \log_2 L} \leq \frac{K \log_2(2L \frac{Me}{K} \frac{Ne}{K})}{M \log_2 L} \quad (9)$$

Specifically in the case of compressive sensing, reconstruction robustness requires that the number of samples is equal to $M = cK \log_2(N/K)$ for some constant $c > 1$. In this case, the bit use is upper bounded by:

$$\frac{\log_2(I_N(M, K, L))}{M \log_2 L} \leq \frac{\log_2 L + \log_2(\frac{N}{K}) + \log_2(\log_2(\frac{N}{K})) + \log_2(e^2 c) + 1}{c \log_2(\frac{N}{K}) \log_2 L} \quad (10)$$

$$= O\left(\frac{1}{c \log_2(\frac{N}{K})} + \frac{1}{c \log_2 L}\right) \quad (11)$$

As the quantization becomes finer, L grows to infinity and the fraction of bits used converges to $1/(c \log_2(\frac{N}{K}))$. Similarly, as the signal becomes sparser, the fraction $\rho = N/K$ grows to infinity, and the fraction of bits used converges to $1/(c \log_2 L)$.

Alternatively, the bounds above could be derived assuming a fixed rate of B bits, i.e. $L = 2^{B/M}$ quantization levels. In this case:

$$\frac{\log_2(I_N(M, K, L))}{B} \leq \frac{K}{M} + \frac{1 + eK(\log_2 M + \log_2 N - 2 \log_2 K)}{B}. \quad (12)$$

Again, when obtaining $M = cK \log_2(N/K)$ compressive sensing measurements, the bit use is upper bounded by:

$$\frac{\log_2(I_N(M, K, L))}{B} \leq \frac{1}{c \log_2(\frac{N}{K})} + \frac{1 + eK(\log_2(\frac{N}{K}) + \log_2(c \log_2(\frac{N}{K})))}{B}. \quad (13)$$

The bounds in Eq. (11) and (13) describe the gap between the number of bits used by the quantized representations and the number of bits that are necessary to represent the same quantization points. In principle, subsequent lossless entropy coding could be used to eliminate the gap between the bits necessary and the bits used. However such a coding method is not currently available for compressive sensing measurements. Some of the difficulty in implementing a lossless entropy coding stage is a consequence of the randomness of the sampling dictionary, which is used to achieve the universality property. The gap in bit-use efficiency is the price to pay for the universality of the currently state of the art universal encoding strategies.

3.3 Error Bounds

Coupled with the rate efficiency of a sampling and quantization algorithm for sparse representations is a bound on the maximum decoding error. Specifically, this section uses the number of quantization cells intersected by the sampling method together with Eq. (3) to determine a lower bound on the worst case quantization error, assuming the best possible distribution of these points in the volume of interest. That is compared to the best case lower bound, which assumes the use of all the possible quantization points at that rate. Furthermore, the efficiency of the classical compressive sensing paradigm is considered, in which sampling and scalar quantization is followed by non-linear reconstruction using linear programming algorithms.

As a benchmark for further comparison, a lower bound on the worst case error is derived and it is shown that it is achievable. Specifically, the lowest worst case error can be achieved by equally distributing the available quantization points to all $\binom{N}{K}$ subspaces that constitute the signal set. Using a pre-determined bit-rate B , this corresponds to $2^B/\binom{N}{K}$ quantization points per subspace. From Eq. (3) and Stirling's formula it follows that the worst case error is lower bounded by:

$$\epsilon_B \geq \frac{N}{2^{B/K}K}. \quad (14)$$

This bound is tight since it can be achieved using $\log_2 \binom{N}{K} \approx K \log_2 \left(\frac{N}{K}\right)$ bits to encode the locations of the sparse coefficients and allocating the remaining bits to encode the implied K -dimensional subspace using a vector quantizer. It should be noted that if, instead, a scalar quantizer is used to encode each of the K sparse coefficients independently, there is a $K^{1/2}$ penalty to the achievable worst case error:

$$\epsilon_{B-Sc} \leq \frac{N}{2^{B/K}K^{1/2}}. \quad (15)$$

At the same bit-rate B , sampling followed by scalar quantization using $L = 2^{B/M}$ levels intersects only $I_N(M, K, 2^{B/M})$ quantization cells. If these cells and the corresponding quantization points are efficiently placed in the space, the worst case quantization error is lower bounded by:

$$\epsilon_S \geq I_o(M, K, 2^{B/M})^{-1/K} \geq \frac{K}{2^{1+B/M}Me}. \quad (16)$$

This bound, applied to $M = cK \log_2(N/K)$ compressive sensing measurements, yields the following lower bound for any reconstruction algorithm:

$$\epsilon_{CS} \geq \frac{1}{2^{1+B/(cK \log_2(N/K))}ce \log_2 \left(\frac{N}{K}\right)} = \Omega \left(\frac{1}{2^{B/(cK \log_2(N/K))} \log_2 \left(\frac{N}{K}\right)} \right) \quad (17)$$

Assuming reconstruction using the linear program in Eq. (6), the error can be guaranteed to be:

$$\epsilon_{CS-LP} \leq O(M^{1/2}\Delta/2) = O \left(\frac{(K \log_2(N/K))^{1/2}}{2^{B/cK \log_2(N/K)}} \right), \quad (18)$$

in which $\Delta = 2^{1-B/M}$ is the uniform quantization interval achievable at rate B .

The lower error bound in Eq. (17) is worse than the lower error bound in Eq. (14). The logarithmic term in the denominator of the exponent can be explained by considering that the optimal vector quantization method only needs to quantize a K -dimensional space instead of the $M = cK \log_2(N/K)$ -dimensional one used for compressive sensing. On the other hand compressive sensing gains an N/K factor since

the location of the sparse components does not need to be encoded. The remaining penalty of $1/\log_2(N/K)$ should be attributed to the inefficient use of the M -dimensional space, since the union of the projected subspaces does not intersect all the quantization cells. It is the price for the universality of compressive sensing.

There is a further gap between the lower bound in Eq. (17) and the achievable error upper bound in Eq. (18). A factor of $M^{1/2} = (K \log_2(N/K))^{1/2}$ in this gap can be attributed to the use of scalar quantization instead of vector quantization in the M dimensional space. However, there is still a factor of $\log_2(N/K)$ that is not accounted for. This factor is due to the inefficiency of the recovery process and the inefficient placement of the quantization cells on the subspaces. A similar logarithmic gap is demonstrated in [5] in the case of compressible signal models.

3.4 Compressible Signals

As mentioned in the Introduction, another commonly used sparsity model is the set of vectors \mathbf{x} with bounded ℓ_p norm for some $0 < p < 1$. This section uses the results developed above to improve on the bounds presented in [5], specifically in the case of low bit-rates and highly compressible signals.

Without loss of generality it is assumed that the signals lie in the unit ℓ_p ball, $\mathcal{B}_p = \{\mathbf{x} \mid \|\mathbf{x}\|_{\ell_p} \leq 1\}$. In [5] the following error bound is presented for the quantization of such signals using any encoding algorithm:

$$\epsilon_{\ell_p} = \Omega \left(\frac{B}{\log_2(N/B) + 1} \right)^{-1/p+1/2}. \quad (19)$$

As p tends to 0, this bound also tends to 0 and becomes loose. At the limit of $p = 0$, it can be shown that \mathcal{B}_p converges to the set of signals that are 1-sparse¹ and power limited to 1. In this case, Eq. (14) provides a tighter bound than Eq. (19):

$$\epsilon_{\ell_p} \geq \frac{N}{2^B} = 2^{-B+\log_2 N}. \quad (20)$$

Using the nesting property of ℓ_p balls, namely that $\mathcal{B}_p \subset \mathcal{B}_{p'}$ if $p < p'$, it follows that the error bound is non-decreasing in p . Therefore, Eq. (20) is a lower bound of the quantization error for all p . This bound is tighter than the one in Eq. (19) when p is small. In the case of very low rate quantization (i.e. B close to $\log_2 N$) the bound in Eq. (20) can be dominant even for large $p < 1$. However, as the rate B increases, the bound decreases exponentially and the bound in Eq. (19) dominates, even for p close to 0.

4 Conclusions

This work demonstrates that compressive sensing of sparse signals followed by scalar quantization is inefficient both in terms of the rate utilization and the error performance. Some of the inefficiency may be recovered using vector instead of scalar

¹In general, $\lim_{p \rightarrow 0} \{\mathbf{x} \mid \|\mathbf{x}\|_{\ell_p} \leq C\}$ is the set of 1-sparse signals that are power limited to \sqrt{C} .

quantization. Still, the universality of the encoding algorithm requires that the quantization cells uniformly cover the sampling space. On the other hand the samples of a sparse signal form low dimensional subspaces in the sampling space. These subspaces do not intersect all the available quantization cells. The resulting gap is the cost of universality. The results suggest some further inefficiency in reconstruction using linear programming. This might be recoverable or partly recoverable using reconstruction methods that ensure reconstruction consistency with the quantized measurements.

References

- [1] D. Donoho, “Compressed Sensing,” *IEEE Trans. Info. Theory*, vol. 52, no. 4, pp. 1289–1306, September 2006.
- [2] E. J. Candès and T. Tao, “Near Optimal Signal Recovery From Random Projections: Universal Encoding Strategies?,” *IEEE Trans. Info. Theory*, 2004, (Submitted. Preprint available at <http://www.acm.caltech.edu/emmanuel/papers/OptimalRecovery.pdf>).
- [3] E. J. Candès, J. Romberg, and T. Tao, “Robust Uncertainty Principles: Exact Signal Reconstruction from Highly Incomplete Frequency Information,” *IEEE Trans. Info. Theory*, vol. 52, no. 2, pp. 489–509, Feb. 2006.
- [4] E. Candès, J. Romberg, and T. Tao, “Stable Signal Recovery from Incomplete and Inaccurate Measurements,” *Communications on Pure and Applied Mathematics*, vol. 59, no. 8, pp. 1207–1223, Aug. 2006.
- [5] E. J. Candès and J. Romberg, “Encoding the ℓ_p Ball From Limited Measurements,” in *Data Compression Conference, 2006*, March 2006, pp. 33–42.
- [6] P. Boufounos, *Quantization and Erasures in Frame Representations*, D.sc., MIT, Cambridge, MA, January 2006.
- [7] R. M. Gray and D. L. Neuhoff, “Quantization,” *IEEE Transactions on Information Theory*, vol. 44, no. 6, pp. 2325–2383, October 1998.
- [8] 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. Info. Theory*, vol. 42, no. 2, pp. 469–479, March 1996.
- [9] S. Sarvotham, D. Baron, and R. G. Baraniuk, “Measurements vs. Bits: Compressed Sensing meets Information Theory,” in *Proc. 44th Allerton Conference on Communication, Control, and Computing*, Monticello, IL, September 2006.