Recovery of Quantized Compressed Sensing Measurements

Grigorios Tsagkatakis[a], Panagiotis Tsakalides[a,b]

[a] ICS - Foundation for Research & Technology - Hellas (FORTH), Crete, Greece
[b] Dept. Computer Science, University of Crete, Greece

ABSTRACT

Compressed Sensing (CS) is a novel mathematical framework that has revolutionized modern signal and image acquisition architectures ranging from one-pixel cameras, to range imaging and medical ultrasound imaging. According to CS, a sparse signal, or a signal that can be sparsely represented in an appropriate collection of elementary examples, can be recovered from a small number of random linear measurements. However, real life systems may introduce non-linearities in the encoding in order to achieve a particular goal. Quantization of the acquired measurements is an example of such a non-linearity introduced in order to reduce storage and communications requirements. In this work, we consider the case of scalar quantization of CS measurements and propose a novel recovery mechanism that enforces the constraints associated with the quantization processes during recovery. The proposed recovery mechanism, termed Quantized Orthogonal Matching Pursuit (Q-OMP) is based on a modification of the OMP greedy sparsity seeking algorithm where the process of quantization is explicit considered during decoding. Simulation results on the recovery of images acquired by a CS approach reveal that the modified framework is able to achieve significantly higher reconstruction performance compared to its naive counterpart under a wide range of sampling rates and sensing parameters, at a minimum cost in computational complexity.

Keywords: Compressed Sensing, Scalar Quantization, Image compression

1. INTRODUCTION

Compressed Sensing (CS)[1,2] is a novel mathematical framework that has revolutionized modern signal and image acquisition architectures ranging from Single Pixel Cameras[3] and CS based spectral imaging[4] , to range sensing[5] and Ultrasound imaging[6] . In imaging scenarios where hard constrains are imposed, including spatial, temporal and spectral resolution, power consumption, and robustness, CS based architectures outperform traditional approaches, both from a theoretical as well as a practical standpoint. However, despite their clear superiority as a sensing paradigm, the success of CS as an end-to-end sensing and encoding strategy has been hindered by various actors, including the effects of signal quantization. In practice, quantization is imposed by the real-life requirement that acquired signals have to be represented using a finite collection of codewords in order to transform the real valued measurements to binary vectors. Such binary represented measurements can then be stored and transmitted before a decoder employs a recovery algorithm.

Most of the effort involved around CS has been primarily focused on identifying the scenarios where a sufficiently large number of randomly encoded measurements are available at the decoder for signal reconstruction. Unfortunately, issues related to the performance of CS in the presence of quantization have hampered the widespread adoption of CS as a valid, generic, compression scheme. The key issue lays in the many-to-one mapping of quantization, which dictates that signal encoding is lossy even if sufficient measurements are acquired.[7] The non-linear nature of quantization leads to a complex interaction with CS recovery where one must address the tradeoff between number of measurements that must be acquired and the number of bits that will be allocated for representing them.

The problem has been predominantly studied in light of two scenarios, high resolution analysis[8,9] and 1-bit CS[10,11] . In the former case, the distance between consecutive quantization thresholds is small compared to the dynamic range of the signal, thus allowing the modeling of quantization as a noise source[12] . Under this

Further author information: (Send correspondence to G.T.) E-mails: G.T.: greg@ics.forth.gr, P.T.: tsakalid@ics.forth.gr

modelling assumption, signal approximation is constrained by the quantization-equivalent noise and recovery follows the same mentality as the case of recovery from noisy measurements. In the other extreme, 1-bit CS corresponds to the case of severe quantization where only the sign of each measurement is maintained. The motivation behind the employment of such extreme quantizers lays in the assumption that modern analog-to-digital converters are congested by the quantization process, leading to lower effective sampling rates, while utilizing high resolution quantizers can cause higher power consumption. A novel line of research regarding the interaction of quantization and CS sampling and reconstruction aims are producing a new generation of recovery algorithms that consider the effects of quantization during the signal estimation process. A prominent example of this direction is the quantized iterative hard thresholding algorithm[13] , which is a modified version of the well known iterative hard thresholding algorithm and has shown very good performance in generic signal recovery from quantized CS measurements[14] .

In this work we propose a novel recovery mechanism for recovering images from quantized CS measurements. Such an imaging architecture can be either supported by the appropriate hardware infrastructure or be employed by an encoder as a generic image compression scheme. Recovery is driven by a novel sparse reconstruction scheme, termed Quantized Orthogonal Matching Pursuit (Q-OMP) which is an extension of the Orthogonal Matching Pursuit (OMP) algorithm for signal recovery from compressed measurements. The reconstruction capabilities of Q-OMP are attributed to both the estimation capacity of OMP and the introduction of the concept of quantization consistency during the recovery process. As a result, our scheme can achieve superior image recovery performance, in a wide range of quantization settings, outperforming the recovery of standard OMP, at a minimal increase in recovery complexity.

2. CS BASED IMAGING

The theory of CS[15] suggests that one can achieve perfect reconstruction of a signal $\mathbf{x} \in \mathbb{R}^N$ from a small number of random measurements $\mathbf{y} \in \mathbb{R}^M$, far below the typical Shannon-Nyquist sampling limit, provided certain conditions are met. Formally, instead of directly sampling the signal $\mathbf{x}$, in CS one acquires a low-dimensional representation $\mathbf{y} = \mathbf{\Psi x} \in \mathbb{R}^M$, where $\mathbf{\Psi}$ is the $M \times N$ sensing (or measurements) matrix with $M << N$. Recovery from the compressed measurements $\mathbf{y}$ is possible by exploiting the sparsity of the signal, either in the original domain, or when the signal is expressed in an appropriate dictionary $\mathbf{D} \in \mathbb{R}^{M \times J}$ according to $\mathbf{x} = \mathbf{Ds}$, where J is the number of atoms in the dictionary and $\mathbf{s} \in \mathbb{R}^J$ is the representation coefficients vector. The full resolution signal can be recovered from the low dimensional embedding by solving a regularized optimization problem where the sparsity of the signal, measured by the l_0 non-zero counting pseudo norm is minimized subject to fidelity conditions:

$$\min ||\mathbf{s}||_0 \;\; \text{subject to} \;\; \mathbf{y} = \mathbf{\Psi D s} \tag{1}$$

A natural variation of this scheme allows for the approximate recovery of the signal by incorporating an error tolerance ϵ, in which case the optimization in Eq. (1) is given by:

$$\min ||\mathbf{s}||_0 \;\; \text{subject to} \;\; \|\mathbf{y} - \mathbf{\Psi D s}\| \leqslant \epsilon \tag{2}$$

An important issue regarding the formulation in Eq.(1) and Eq.(2) is that the l_0 minimization is an NP-hard problem and therefore inefficient to solve for even moderate sized problems.

To address this issue, greedy approaches such as the OMP[16] algorithm have been proposed. OMP greedily tries to identify the elements that contain most of the signal energy by iteratively selecting the dictionary element that best matches the signal and estimating the residual error by projecting the input signal to the linear span of the selected elements, until the residual error is within an acceptable approximation limit. Alternately, the CS theory suggests that for sufficiently sparse signals and for measurements matrices whose elements are drawn from appropriate distributions that satisfy the RIP property[17] , reconstructing the original signal $\mathbf{s}$ from the measurements $\mathbf{y}$ can be achieved by replacing the l_0 norm with the more tractable l_1 norm, defined as the sum of the absolutes values of the vector elements.

The sensing matrix plays a key role in the ability to recover sparse signal from CS measurements. In general, the deterministic construction of a universally optimal sensing matrix is a formidable task, that has been intensively investigated by the research community[18, 19] . Fortunately, the theory of CS has shown that recovery is possible when the entries of the sensing matrix are randomly drown from an appropriate distribution, two of the most prominent ones being the Gaussian and the Binary. While the former offers stronger theoretical guarantees, the latter is more attractive from an implementation perspective.

3. QUANTIZATION OF CS MEASUREMENTS

A scalar quantization process is defined by an ordered set of thresholds $\mathcal{T} = \{t_1, ..t_B | t_i < t_j, \forall i, j\}$, generating a codebook $\mathcal{C}$ containing $|\mathcal{C}| = B = 2^R$ codewords. Typically each codeword corresponds to the value in the middle between two consecutive thresholds. Each acquired measurement y is represented by the closest codeword such that $\mathcal{Q}(y) = \min_{c \in \mathcal{C}} \|c - y\|_2$. After quantization, the acquired signal is given by $\hat{\mathbf{y}} = \mathcal{Q}_B(\mathbf{y})$, where the nonlinear operator $\mathcal{Q} : \mathbb{R} \to 2^R$ models the process of mapping the set of real numbers to a set of specific elements indexed by R bits. As a consequence of the quantization, the measurements are compactly represented, requiring a total bit budget equal to RM. The most straightforward allocation of the bit budget is the uniform scalar quantization.

In the uniform scalar quantization, for a given bit resolution B per measurement, the quantizer thresholds the real line uniformly such that $|t_i - t_j| = \Delta$ for signals with bounded magnitude between t_1 and t_B. Values outsize the bounds are saturated and get truncated to the limiting bins. The output of a uniform scalar quantizer for an input uniformly distributed signal x is given by $\mathcal{Q}(x) = sgn(x) \cdot \Delta \cdot \lfloor \frac{|x|}{\Delta} + 0.5 \rfloor$ where sgn is the sign function.[20] In addition to the non-linear mapping of measurements, uniform scalar quantization also imposes constrains on the magnitude of the measurements by saturating values outside the predefined region.

While the uniform scalar quantization is a computational efficient approach, which works for any signal with bounded magnitude, the under-utilization of prior knowledge may lead to inefficient designs. Due to the importance of the allocation of this bit budget, different approaches have been presented. A more efficient quantization approach is the non-uniform scalar quantizer which is based on the Lloyd-Max optimal quantization approach[21, 22] . Unlike the uniform scalar quantization, in optimal scalar quantization one assumes knowledge of statistical properties of the signals under investigation extracted from training examples. These examples are employed in a iterative optimization approach that searches for the partition scheme that minimizes the representation error of the training samples.

In order to account for the effect of quantization, we propose a novel formulation of OMP based recovery. Formally, we consider the non-linear mapping function of the sparse signal $\mathbf{s}$ given by $\mathcal{Q}(\mathbf{s}; \mathbf{GD}) = \mathcal{Q}(\mathbf{GDs})$. In this case, the recovery program in Eq.(2) can be expressed according to:

$$\min \|\mathbf{s}\|_0 \;\; \text{subject to} \;\; \|\hat{\mathbf{y}} - \mathcal{Q}(\mathbf{s})\|_2 \leqslant \epsilon \; . \tag{3}$$

Recovery in the presence of the non-linear function $\mathcal{Q}$ can be achieved by solving a modified version of the greedy OMP approach, with additional constraints on the consistency of the recovered signal with respect to its quantized counterpart. The algorithmic steps of the proposed Quantized OMP (Q-OMP) is presented in Algorithm 1. Q-OMP, much like OMP, is an iterative process where in each iteration the algorithm first performs support identification by selecting the dictionary element that best matches (in an l_2 norm) the residual and then updates the residual error.

For the case of uniform scalar quantization, the approximation error can be bounded by the quantization equivalent noise, while in the case of the optimal quantization the error is related to the statistics of the input signals. With respect to the typical OMP, Q-OMP imposes a small increase in computational complexity due to the introduction of the quantization of the estimates. However, this increase is associated with a significantly more robust behavior compared to OMP as we will see in the next section.

4. EXPERIMENTAL RESULTS

In our system model, we follow a block coding approach where groups of pixels, 8×8 in our case, are multiplexed by a specific sampling pattern ϕ_i producing a single measurement y_i for that group. The process is simultaneously applied to all non-overlapping groups of image pixels and repeated with different sensing matrices in order to acquire a number of measurements. Since images are not naturally sparse signal, one has to resort to the use of

Algorithm 1: Quantized Orthogonal Matching Pursuit (Q-OMP)

Input: The measurements $\mathbf{y}$,
The sensing matrix $\mathbf{\Phi}$,
The dictionary of examples $\mathbf{D}$,
The error tolerance *threshold* and/or maximum number of iterations k.
Output: The sparse representation coefficients $\hat{\mathbf{s}}$.

initialization $T^0 = \varnothing, \mathbf{r}^0 = \mathbf{y}$
while $error \geqslant threshold$ or $k \leqslant iterationslimit$ **do**
  $T^k = T^{k-1} \cup \arg\max_j |\mathcal{Q}(< \mathbf{r}^{k-1}, (\mathbf{\Phi D})_j >)|$.
  $\hat{\mathbf{s}}_{T^k} = \arg\min_s \|y - \mathcal{Q}((\mathbf{\Phi D}_{T^k})\mathbf{s})\|_2$.
  $\mathbf{r}^k = \mathbf{y} - \mathbf{\Phi}_{T^k}\hat{\mathbf{s}}_{T^k}$.
  set $k \longleftarrow k + 1$
end while

a dictionary that can sparsely represent such signals. Although a large number of presentation basis have been proposed, ranging from wavelet-like basis to learned dictionaries, we consider the Discrete Cosine Transform (DCT) in our experiential section, motivated by the extremely wide use of DCT as a image transform in the JPEG image compression algorithm. To compare the performance of each method, we consider the PSNR of the full 8 bits recovered images as the error metric. The training of the quantizers and the testing on the reconstruction was performed on disjoint sets of grayscale versions of benchmark images*. We consider two cases where we employ either a binary sensing matrix or a Gaussian one. For the former case, we consider binary matrices where $p(\Phi_{i,j} = 1) = 0.5$ and for the former case we consider zero mean Gaussian distributions with variance $1/M^{23}$.

In the first set of experiments, we considered the uniform scalar quantization with saturation of measurements acquired with normalized Gaussian and Binary sensing matrices. Figure 1 presents the recovery performance with Gaussian sensing matrices at (a) 4 bits and (c) 2 bits per measurement and from Binary sensing matrices at (b) 4 bits and (d) 2 bits per measurement. Examining all four cases presented in Figure 1, we observe that the *ordering* of the algorithms with respect to performance is maintained accross sensing schemes and quantization resolutions, however, reconstruction quality is significantly affected by quantization. More specifically, one can notice there is always a significant performance reduction when the reconstruction algorithm is presented with the full resolution measurements (labeled unquantized) versus the quantized ones (labeled OMP). Furthermore, this performance gap is not constant but increases at higher sampling rates. In the moderate quantization of 4 bits per measurement, the reconstruction quality is monotonically increased at increasing sampling rates, following the monotonic increase in reconstruction quality observed in the unquantized case, yet, achieving a lower performance bound. However, in the higher compression case of 2 bits per measurement, the recovery of OMP in the presence of quantization appears to deteriorate at higher sampling rates. This phenomenon can be attributed to the fact that increasing the number of measurements, leads to higher dynamic range of values, which in the cases of extreme quantization such as the 2 bit per sample, causes the quantization equivalent noise to out-power the underlying true signal power. Regarding the recovery under the two sensing schemes, namely the Gaussian and the Binary, the theoretical assumptions regarding encoding capabilities are manifested in the experimental results, where recovery from Gaussian based measurements always achieves higher quality compared to recovery from the Binary ones.

Focusing on the comparative performance of OMP and Q-OMP reconstruction, simulation results indicate that Q-OMP is more robust to the quantization effects, typically achieving monotonic increase at higher sampling rates. In the 4 bits per sample case, we observe that both OMP and Q-OMP exhibit similar performance up to around 20 measurements, while Q-OMP is better at higher sampling rates. In the challenging case of 2 bits per measurement, there is a significant difference between the reconstruction of Gaussian sensed verses Binary sensed measurements, where in the Binary case, Q-OMP achieves almost 10dB gain in performance at high sampling rates. Furthermore, while OMP's performance in the Binary 2 bit case is extremely poor, leading to situations

*http://sipi.usc.edu/database/

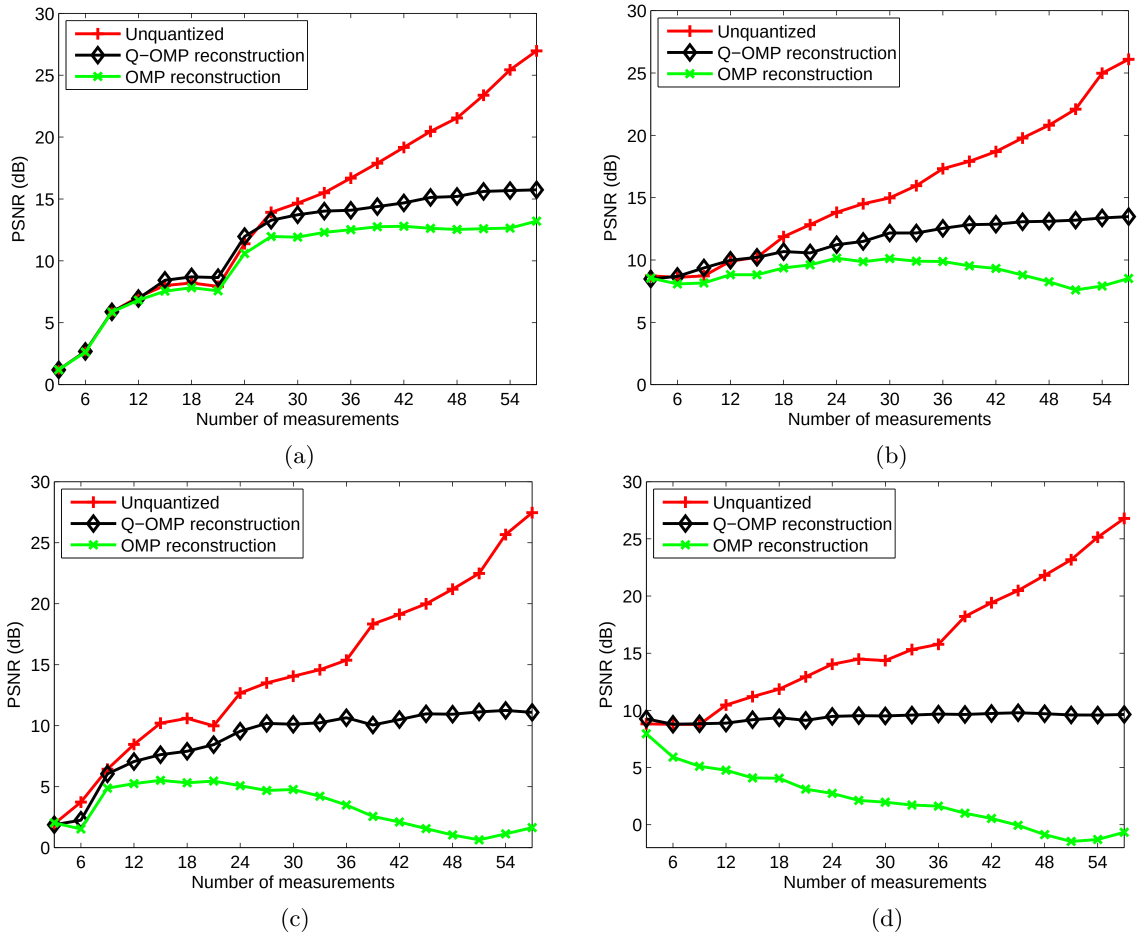

Figure 1: Recovery of compressible signals from unquantized and *uniformly* quantized measurements by employing the traditional OMP and the Quantization aware OMP (Q-OMP). We consider Gaussian (left column) and Binary (right column) sensing matrices and uniform scalar quantization with quantization at (a,b) 4 bits and (c,d) 2 bits per measurement.

where the noise out-powers the true signal, Q-OMP maintains a stable performance across sampling rates. Note that even in the scenario of extreme quantization where increasing the sampling rate without increasing the encoding resolution leads to worse performance for typical recovery approaches, the proposed Q-OMP algorithm is able to maintain a non-decreasing performance.

In the second groups of experiments, we consider the recovery performance when an optimal scalar quantizer is employed. Figure 2 presents the four cases regarding bit resolution and sensing matrices as Figure 1, however the quantization thresholds are found by the Lloyd-Max algorithm. Comparing the overall performance between the uniform and the optimal quantization, we observe that recovery is better for both OMP and Q-OMP in the optimal quantization case, as it would be expected. More specifically, OMP in the 4 bits per measurement achieves up to 5 dB increase in recovery performance from Gaussian measurements by going from uniform to optimal quantization, while little performance gain is observed for the binary measurements. In the 2 bit per measurement case, again there is a noticeable increase in performance for Gaussian measurements while for Binary sensing, actually increasing the sampling rate lead to performance drop due to the increase in the dynamic range of measurements.

The performance of Q-OMP in the 4 bits per measurement case when an optimal quantization is employed

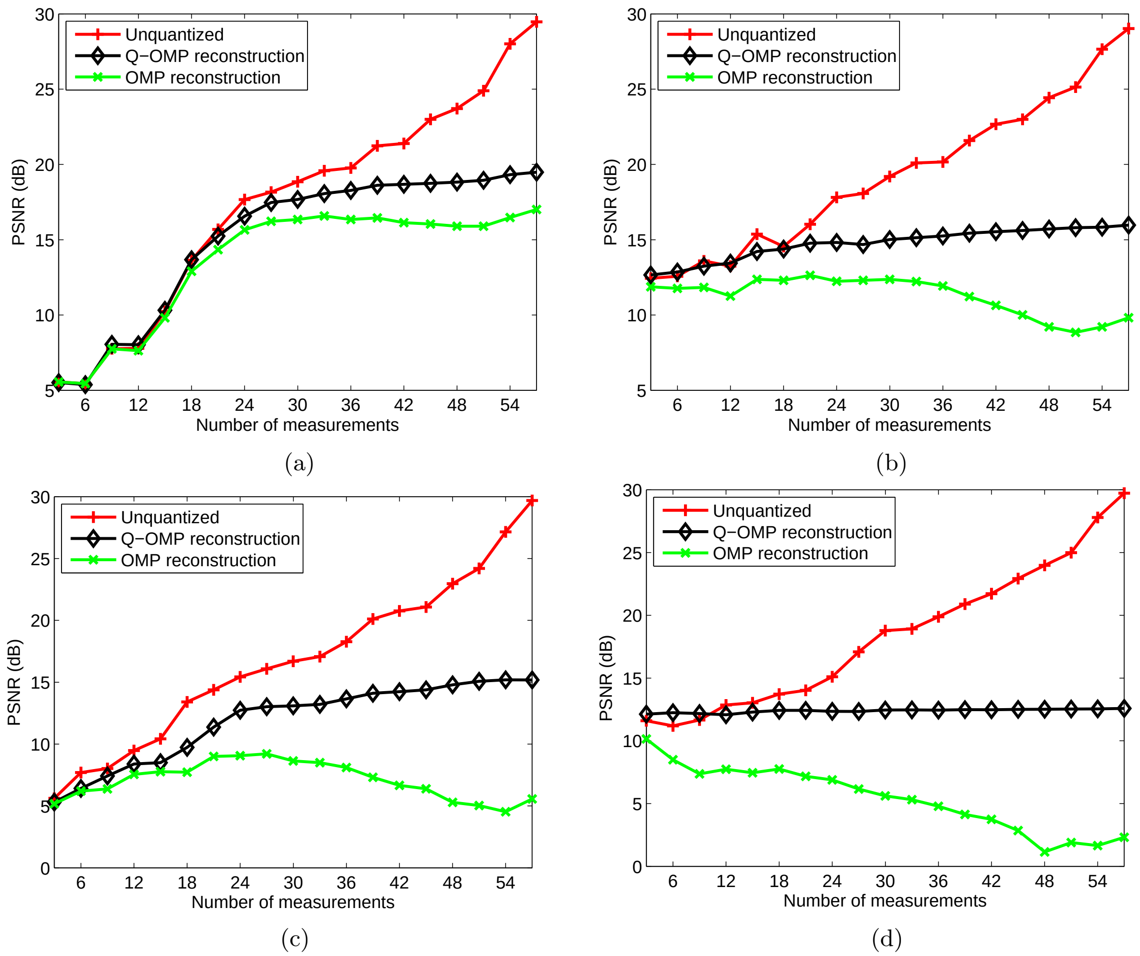

Figure 2: Recovery of compressible signals from unquantized and *optimally* quantized measurements by employing the traditional OMP and the Quantization aware OMP (Q-OMP). We consider Gaussian (left column) and a Binary (right column) sensing matrices and uniform scalar quantization with quantization step equal to 0.01 (a and b), 0.1 (c and d) and 0.2 (e and f).

follows the same behavior as OMP, *i.e* significant performance increase in for Gaussian sensing matrices and marginal for Binary. The gain observed for 4 bits per measurement in the Gaussian case is transfer to the 2 bit case, where unlike OMP, the performance of Q-OMP is comparable to the 4 bits per measurement of uniform quantization. Even in the challenging case of Binary sensing at 2 bits per measurement, the performance of Q-OMP is non decreasing.

Drowning some conclusion from both experimental setups, results suggest that although Binary sensing matrices are more implementation friendly, they cannot achieve the recovery performance when Gaussian sensing matrices are employed. In the case of extreme quantization (2 bits per measurement), the noise introduced by qunatization can even lead to performance drop for Binary sensing matrices at higher sampling rates. Regarding recovery capabilities, Q-OMP outperforms OMP is all cases achieving performance comparable to the unquantized case at low sampling rates.

To understand how recovery error is translated to visual artifacts, we present an example of images recovered from the typical and the proposed scheme in a scenario where an optimal scalar quantizer at 4 bits per measurement was employed in conjunction with a Gaussian sensing matrix which collected 32 measurements in total. One easily observes that even in the case of unquantized measurements, a realistic scenario such as the one discussed here significantly hinders the recovery capabilities. However, although the recovered image exhibits strong signs of blocking effects, a significant portion of details has been recovered. With further processing steps including the application of a de-blocking algorithm, the quality of the final image can be dramatically increased. Regarding the performance of recovery in the quantized measurements case (c), the image exhibit a considerable amount of complex structured noise. As a consequence, we observe that both low as well as high frequency details are lost due to white nature of quantization. The quantization aware Q-OMP achieves a better visual quality than OMP, sharing an important range of image features with the unquantized case.

(a)

(b)

(c)

(d)

Figure 3: Illustration of (a) original and recovered images with (b) OMP from original measurements, (c) OMP from quantized measurements and (d) Q-OMP from quantized measurements, from 32 measurements obtained with Gaussian sensing matrices, quantized at 4 bits per measurements with the Lloyd-Max quantization. The PSNR achieved by the methods are (b) 15.6 dB, (c) 14.4 dB and (d) 15.8 dB.

5. DISCUSSION

While the concept of CS has been introduced in numerous imaging architectures, quantization of compressibly sampled measurements is a necessary evil associated with real-life systems that has limited to adoption of CS as an end-to-end image acquisition and storage scheme. In this work, we investigate the effects of scalar quantization on images either captured by a CS architectures or encoded using a CS based image compression scheme. Simulation results verify the assumption that indeed quantization can significantly reduce the recovery capabilities of standard recovery algorithms such as the OMP. To address this issues, we propose a modified reconstruction algorithm, the Q-OMP, which considers the quantization process during the iterative recovery steps. By introducing the quantization consistency constraint, a significant portion of the reconstruction error is adsorbed by the proposed recovery algorithm, outperforming the typically approach under all scenarios under investigation.

6. ACKS

This work was funded by the IAPP CS-ORION (PIAP-GA-2009-251605) grant within 7th Framework Program of the European Community and co-financed by the European Union and Greek national funds through the National Strategic Reference Framework (NSRF), Research Funding Program: "Cooperation-2011", Project "SeNSE", grant number 11_6_1381 and GSRT O.P. Competitiveness and Entrepreneurship PEFYKA project within the KRIPIS ction of the GSRT.

REFERENCES

1. Donoho, D., "Compressed sensing," *Information Theory, IEEE Transactions on* **52**(4), 1289–1306 (2006).
2. Candès, E. J., Romberg, J., and Tao, T., "Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information," *Information Theory, IEEE Transactions on* **52**(2), 489–509 (2006).
3. Baraniuk, R. G., "Single-pixel imaging via compressive sampling," *IEEE Signal Processing Magazine* (2008).
4. Wagadarikar, A. A., Pitsianis, N. P., Sun, X., and Brady, D. J., "Video rate spectral imaging using a coded aperture snapshot spectral imager," *Optics Express* **17**(8), 6368–6388 (2009).
5. Tsagkatakis, G., Woiselle, A., Tzagkarakis, G., Bousquet, M., Starck, J. L., and Tsakalides, P., "Compressed gated range sensing," in [*SPIE Optical Engineering+ Applications*], 88581B–88581B, International Society for Optics and Photonics (2013).
6. Tzagkarakis, G., Achim, A., Tsakalides, P., and Starck, J.-L., "Joint reconstruction of compressively sensed ultrasound rf echoes by exploiting temporal correlations," in [*Biomedical Imaging (ISBI), 2013 IEEE 10th International Symposium on*], 632–635, IEEE (2013).
7. Boufounos, P. T., Jacques, L., Krahmer, F., and Saab, R., "Quantization and compressive sensing," *arXiv preprint arXiv:1405.1194* (2014).
8. Laska, J. N., Boufounos, P. T., Davenport, M. A., and Baraniuk, R. G., "Democracy in action: Quantization, saturation, and compressive sensing," *Applied and Computational Harmonic Analysis* **31**(3), 429–443 (2011).
9. Sarvotham, S., Baron, D., and Baraniuk, R. G., "Measurements vs. bits: Compressed sensing meets information theory," in [*Proceedings of 44th Allerton Conf. Comm., Ctrl., Computing*], (2006).
10. Boufounos, P. T. and Baraniuk, R. G., "1-bit compressive sensing," in [*Information Sciences and Systems, 2008. CISS 2008. 42nd Annual Conference on*], 16–21, IEEE (2008).
11. Jacques, L., Laska, J. N., Boufounos, P. T., and Baraniuk, R. G., "Robust 1-bit compressive sensing via binary stable embeddings of sparse vectors," *IEEE Transactions on Information Theory* **59** (April 2013).
12. Widrow, B. and Kollár, I., "Quantization noise," *Cambridge University Press* **2**, 5 (2008).
13. Jacques, L., Degraux, K., and De Vleeschouwer, C., "Quantized iterative hard thresholding: Bridging 1-bit and high-resolution quantized compressed sensing," *arXiv preprint arXiv:1305.1786* (2013).
14. Blumensath, T. and Davies, M. E., "Iterative hard thresholding for compressed sensing," *Applied and Computational Harmonic Analysis* **27**(3), 265–274 (2009).
15. Candes, E., Eldar, Y., Needell, D., and Randall, P., "Compressed sensing with coherent and redundant dictionaries," *Applied and Computational Harmonic Analysis* **31**(1), 59–73 (2011).

16. Tropp, J. A. and Gilbert, A. C., "Signal recovery from random measurements via orthogonal matching pursuit," *Information Theory, IEEE Transactions on* **53**(12), 4655–4666 (2007).
17. Baraniuk, R., Davenport, M., DeVore, R., and Wakin, M., "A simple proof of the restricted isometry property for random matrices," *Constructive Approximation* **28**(3), 253–263 (2008).
18. DeVore, R. A., "Deterministic constructions of compressed sensing matrices," *Journal of Complexity* **23**(4), 918–925 (2007).
19. Rauhut, H., "Compressive sensing and structured random matrices," *Theoretical foundations and numerical methods for sparse recovery* **9**, 1–92 (2010).
20. Gray, R. M., [*Source coding theory*], vol. 83, Springer (1990).
21. Lloyd, S., "Least squares quantization in pcm," *Information Theory, IEEE Transactions on* **28**(2), 129–137 (1982).
22. Max, J., "Quantizing for minimum distortion," *Information Theory, IRE Transactions on* **6**(1), 7–12 (1960).
23. Berinde, R., Gilbert, A. C., Indyk, P., Karloff, H., and Strauss, M. J., "Combining geometry and combinatorics: A unified approach to sparse signal recovery," in [*Communication, Control, and Computing, 2008 46th Annual Allerton Conference on*], 798–805, IEEE (2008).