On one-stage recovery for $\Sigma\Delta$-quantized compressed sensing

Arman Ahmadieh
Mathematics Department
University of British Columbia
Vancouver, BC
arman@math.ubc.ca

Özgür Yılmaz
Mathematics Department
University of British Columbia
Vancouver, BC
oyilmaz@math.ubc.ca

***Abstract*—Compressed sensing (CS) is a signal acquisition paradigm to simultaneously acquire and reduce dimension of signals that admit sparse representations. When such a signal is acquired according to the principles of CS, the measurements still take on values in the continuum. In today's "digital" world, a subsequent quantization step, where these measurements are replaced with elements from a finite set is crucial. We focus on one of the approaches that yield efficient quantizers for CS: $\Sigma\Delta$ quantization, followed by a one-stage tractable reconstruction method, which was developed in [20] with theoretical error guarantees in the case of sub-Gaussian matrices. We propose two alternative approaches that extend the results of [20] to a wider class of measurement matrices including (certain unitary transforms of) partial bounded orthonormal systems and deterministic constructions based on chirp sensing matrices.**

***Index Terms*—compressed sensing, quantization, noise-shaping, $\Sigma\Delta$ quantization, one-stage reconstruction**

I. Introduction

Compressed sensing (CS) has recently emerged as a revolutionary sampling theory. This new theory is based on the empirical observation that various important classes of signals, such as audio and images, admit (nearly) sparse approximations when expanded with respect to an appropriate basis or frame, such as a wavelet basis or a Gabor frame. CS theory shows that one can recover such signals from only a few linear, non-adaptive measurements. As such, CS provides a dimension reduction paradigm. However, in todays digitally driven world, every sampling theory needs to be accompanied by a quantization theory. Next, we discuss this aspect of CS.

Formally, a signal is a vector x in $\mathbb{R}^n$, where n is potentially large. We say that x is k-sparse if $\|x\|_0 \leq k$ where $\|x\|_0$ is the cardinality of the support of $x = [x_1, \ldots, x_n]^T$ defined as $\text{supp}(x) := \{j : \ x_j \neq 0\}$. The set of all k-sparse signals in $\mathbb{R}^n$ is denoted by Σ_k^n.

Suppose $x \in \Sigma_k^n$ or it is *compressible*, i.e., it can be well approximated in Σ_k^n such that $\sigma_k(x) := \min_{v\in\Sigma_k^n} \|x - v\|_1$ is small. Compressed measurements of x are linear, non-adaptive measurements given by $y = \Phi x + \eta$. Here Φ is an $m \times n$ CS measurement matrix with $m \ll n$ and η is additive noise. Consequently, the "compressed" measurement vector y is still real valued, this time in $\mathbb{R}^m$, with $m \ll n$. As mentioned earlier, in the classical signal processing paradigm, such an acquisition or *sampling* stage is followed by *quantization* where the sample values are mapped from the continuum to a finite set. While quantization was mostly omitted in the early CS literature, there has been several recent papers that address this problem. The approaches in the literature focus mostly on either "memoryless scalar quantizers" (MSQ) or "noise-shaping quantizers".

A. Memoryless scalar quantization for CS

Suppose that $x \in \mathbb{R}^n$ and $y \in \mathbb{R}^m$ are as above. An MSQ with alphabet $\mathcal{A}$ rounds off each entry of y (independently) to the closest element of $\mathcal{A}$ [6], [9], [19]. A special case of MSQ is the 1-bit quantizers, where each measurement is replaced by its sign [7], [13], [17], [18], i.e., $\mathcal{A} = \{\pm 1\}$.

One way to analyze the error associated with MSQ is by interpreting the quantization error as additive noise. Such an approach shows that one can obtain an approximation $\tilde{x}$ using, for example, Basis Pursuit Denoise [8], [10]. In that case, we get an approximation error bound $\|x - \tilde{x}\|$ that is proportional to the quantizer resolution, say δ. This theoretical upper bound as well as the empirical performance –see [11]– does not improve by increasing the number of measurements m. On the other hand, it was observed in [11] that in a two-stage recovery method where the Penrose-Moore pseudo-inverse is used in the second stage (after support recovery), the error $\|x - \tilde{x}\|$ is empirically $\mathcal{O}(\frac{1}{\sqrt{m}})$. Motivated by this, [16] shows that $\|x - \tilde{x}\|$ is bounded by the sum of two terms: one that is independent on m but unobservably small in any realistic setting, and another that is indeed $\mathcal{O}(\frac{1}{\sqrt{m}})$, at least for a wide class of sub-Gaussian matrices with high probability. Similarly, it was also shown in the 1-bit CS context in [18] that for a fixed level of sparsity, the error in approximation using a specific convex minimization program decays as $\mathcal{O}(\frac{1}{m^{1/5}})$ up to a logarithmic factor.

While these improved results show some decay as a function of m, this decay is mild, suggesting that MSQ does not utilize extra measurements efficiently. This leads us to noise-shaping quantizers.

B. Noise-shaping quantizers for CS

Noise-shaping quantizers were originally introduced in the context of analogue-to-digital (A/D) conversion of bandlimited

Ö. Yılmaz is funded in part by an NSERC Discovery Grant (22R82411) and PIMS CRG 33: HDDA.

signals [12]. These A/D convertors, called $\Sigma\Delta$ quantizers, became popular [21] as they can be implemented using low-accuracy circuit elements and still produce high-accuracy approximations by oversampling. For many classes of signals it is much easier to oversample on circuitry compared to using high-accuracy circuit elements, for example scalar quantizers Q_δ with very small δ.

Motivated by their efficiency in exploiting redundancy, $\Sigma\Delta$ quantizers were considered in the context of frame expansions (which are inherently redundant). Indeed, they were shown to yield approximations that improve as the redundancy increases in the contexts of Gabor frames [23], [24], finite frames in $\mathbb{R}^d$ with certain regularity assumptions [3]–[5], Gaussian random frames [11], and sub-Gaussian random frames [14], [15].

These results in frame theory were instrumental in early work that proposed $\Sigma\Delta$ quantization in the setting of CS. In a nutshell, suppose $x \in \Sigma_k^n$, $\Phi \in \mathbb{R}^{m\times n}$ be an appropriate CS measurement matrix, and $y = \Phi x$ be the noise free compressive measurements. Also, let q be obtained by quantizing y using an rth order $\Sigma\Delta$ scheme and let D be the difference matrix as in [20, Section 2.1]. In [11] a *two-stage recovery algorithm* was proposed: first, the support set $T = \text{supp}(x)$ is recovered or estimated. Then, the reconstruction vector $\hat{x}$ is given by $\hat{x}_{\Sigma\Delta} = Fq$ with $F = (D^{-r}\Phi_T)^\dagger D^{-r}$, where Φ_T denotes the restriction of Φ to its columns indexed by T. While this two-stage reconstruction approach yields superior decay in approximation error as the number of measurements m increases –see [11], [14] – there are two major caveats: The two-stage approach is not robust with respect to additive noise, and it imposes size requirements on the smallest non-zero entry of the sparse signal.

II. One-stage recovery for $\Sigma\Delta$-quantized CS

As a remedy to the issues mentioned above, [20] proposed a *one-stage reconstruction method* which computes the approximation $\tilde{x}$ to x by solving the convex optimization problem

$$(\hat{x},\hat{\nu}) := \arg\min_{(z,\nu)} \|z\|_1 \text{ s.t. } \|D^{-r}(\Phi z + \nu - q)\|_2 \le C_r\delta\sqrt{m}, \text{ and } \|\nu\|_2 \le \epsilon\sqrt{m}. \quad (1)$$

Fix, now, any ℓ that is sufficiently large so that ℓ measurements suffice to recover x from Φx in the non-quantized CS setting. Then the approximation $\hat{x}$ obtained as above satisfies

$$\|\hat{x} - x\|_2 \le C\Big((\frac{m}{\ell})^{-r+1/2} + \sqrt{\frac{m}{\ell}}\epsilon\Big) \quad (2)$$

where c, C are constants that do not depend on m, ℓ, n.

Indeed, this method solves the issues mentioned in the previous section when the CS measurements are obtained via sub-Gaussian matrices and certain Fourier matrices [22]. On the other hand, it is not known if this one-stage recovery method enjoys recovery guarantees when we use other important classes of measurement matrices, e.g., random restrictions of discrete Fourier transform matrices or bounded orthonormal systems (BOS), or various classes of deterministic measurement matrices.

A. Generalizing to other measurement systems

In order to generalize the results of [20] to other classes of random matrices and also certain deterministic matrices, we isolate one main property, which we call *(P1)*, that the measurement matrices must satisfy for such a generalization.

Property (P1). Suppose that Φ is an $m \times n$ unnormalized CS measurement matrix, with (expected) column norm of $\sqrt{m}$. We say that Φ satisfies the property (P1) of order (k,ℓ) if the RIP constant of $\frac{1}{\sqrt{\ell}}(\Phi)_\ell$—where $(\Phi)_\ell$ is the restriction of Φ to its first ℓ rows—satisfies $\delta_{2k} < 1/9$.

Note that sub-Gaussian matrices, and random restrictions of BOS (including the DFT matrix) satisfy this property with high probability for appropriate choices of k and ℓ.

Let $y = \Phi x + \eta$, and $\|\eta\|_\infty \le \epsilon$. Set $H := [C_r D^r \ \frac{\epsilon}{\delta} I]$. Here C_r is a constant that can depend on the order r and δ and in the specific case of an rth order greedy $\Sigma\Delta$ quantizer, $C_r = 1/2$ [20]. Next, let $H = U\Sigma V^T$ be the singular value decomposition of H. With this notation, the approach used in [20] is to show that $U^T\Phi$ satisfies *(P1)*. It is well-known that sub-Gaussian matrices satisfy (P1) and this is leveraged in [20] to show that $U^T\Phi$ satisfies *(P1)* as well. Yet, this implication is non-trivial and not necessarily true, for example, when Φ is a partial BOS.

Here, we propose two ways to circumvent this issue. Specifically, we will devise two novel approaches where it will be sufficient that Φ (instead of $U^T\Phi$) satisfies *(P1)*.

III. Two novel approaches

A. Approach 1: Using a modified measurement matrix

It can be shown (similar to the proof of Theorem 1 in [20]) that one-stage reconstruction following $\Sigma\Delta$ quantization can be performed if

(i) Φ satisfies *(P1)*, and
(ii) measurements are obtained using $U\Phi$ as opposed to Φ.

In particular, under this condition, the reconstruction error is as in [20, Theorem 1]. Specifically, the following holds.

Theorem 1. *Suppose that Φ is an $m \times n$ CS matrix, $x \in \mathbb{R}^n$, and $k < \ell \le m$ is such that Φ satisfies (P1) of order (k,ℓ). Suppose the measurements of x are given by $y = \tilde{\Phi}x$, where $\tilde{\Phi} = U\Phi$ with U as above, and quantized by an rth-order $\Sigma\Delta$ scheme. Then $\hat{x}$, obtained via* (1) *after replacing Φ with $\tilde{\Phi}$ satisfies*

$$\|x - \hat{x}\|_2 \le C\Big((\frac{m}{\ell})^{-r+1/2}\delta + \frac{\sigma_k(x)}{\sqrt{k}} + \sqrt{\frac{m}{\ell}}\epsilon\Big) \quad (3)$$

where C is a constant that does not depend on m, ℓ, n.

Remark 1. *Theorem 1 shows that the one-stage reconstruction scheme can be applied when the measurement matrix is, e.g., $U\mathcal{F}^{m,n}$, $U\mathcal{C}^{m,n}$, or $U\mathcal{S}^{m,n}$, for large enough m. Here, $\mathcal{F}^{m,n}$, $\mathcal{C}^{m,n}$, and $\mathcal{S}^{m,n}$ denote random restrictions of a discrete Fourier transform, discrete cosine transform, and discrete sine transform, respectively. In fact, one can use $U\mathcal{B}^{m,n}$ where $\mathcal{B}^{m,n}$ is any partial BOS. Furthermore, multiplying any of*

these matrices with U can be implemented using FFT, thus fast, at least when $r = 1$.

Remark 2. *Alternatively, one can apply U after collecting the measurements. This would require all m analog measurements be stored until we apply U, and thus is not practicable when m is large.*

B. Approach 2: Using a digital buffer

Aside from the issues raised in Remark 2, the above approach is not ideal because the measurement matrix $U\Phi$ (specifically U) depends on m. This means that we must use a different measurement matrix if we wish to increase the number of measurements m, i.e., we can not "reuse" the measurements already collected. This problem would be resolved if we could modify the scheme so that we first collect $y = \Phi x$; quantize y; and then apply U digitally (which has a fast implementation) on the quantized measurements.

To that end, we propose the following scheme.

(1) Given a standard CS measurement matrix Φ, collect $y = \Phi x + \eta$, where the noise η satisfies $\|\eta\|_\infty \leq \epsilon$.
(2) Fix a small δ' (much smaller than the desired final accuracy) and quantize y using an MSQ with step size δ' resulting in y_{MSQ}. This is a high bit-budget representation of y and will be discarded after the next stages so, it is just kept in a buffer (with sufficiently large memory).
(3) Compute Uy_{MSQ}, which finely approximates $Uy = U\Phi x$ as U is unitary.
(4) Use a $\Sigma\Delta$ quantizer (of order r that matches the matrix U in step (3)) with step size δ to quantize Uy_{MSQ}. This will be the digital representation of x that we will keep.

Remark 3. *The high-resolution representation y_{MSQ} is used only during data collection and discarded once we obtain the $\Sigma\Delta$-quantized representation at the end of Step (4).*

Finally, we reconstruct an approximation to x by solving

$$(\hat{x}, \hat{\nu}) := \arg\min_{(z,\nu)} \|z\|_1 \text{ s.t. } \|D^{-r}(\Phi z + \nu - q)\|_2 \leq C_r \delta \sqrt{m} \text{ and } \|\nu\|_2 \leq \delta'' \sqrt{m} \tag{4}$$

with $\delta'' := \epsilon + \delta'/2$. Note that this method will be successful provided δ' in step (2) is sufficiently small to match the quantization error corresponding to the $\Sigma\Delta$ quantization of step (4). Thus, we will have to ensure that $m \leq m_{max}$ where δ' will be chosen depending on m_{max} (or vice versa). Collecting all these, we have the following Theorem:

Theorem 2. *Let $x \in \mathbb{R}^n$, Φ be an $m \times n$ CS measurement matrix, and let k and ℓ be such that Φ satisfies (P1) of order (k, ℓ). Suppose that q is obtained from x following the scheme suggested above where $\delta' := \frac{\delta}{(6r)^r m_{max}^r}$ and $\|\eta\|_\infty \leq \epsilon$. If $\hat{x}$ is obtained via* (4)*, the approximation error satisfies*

$$\|x - \hat{x}\|_2 \leq C\Big((\frac{m}{\ell})^{-r+1/2}\delta + \frac{\sigma_k(x)}{\sqrt{k}} + \sqrt{\frac{m}{\ell}}\epsilon\Big) \tag{5}$$

whenever $m \leq m_{max}$. Here C is a constant that does not depend on m, ℓ, n.

Note that partial BOS satisfy the conditions of Theorem 2. Further, Uy_{MSQ} can be computed fast, at least when $r = 1$.

IV. One-stage recovery for $\Sigma\Delta$ quantization with deterministic matrices

Chirp sensing matrices constitute an important class of deterministic matrices, first introduced by Applebaum et al. [2] in the context of CS. For a prime number p and $\omega := e^{i\frac{2\pi}{p}}$, the columns of a $p \times p^2$ chirp sensing matrix Φ are defined via

$$\Phi_{rp+m+1} = \Big[\omega^{r\cdot 0^2 + m\cdot 0}, \omega^{r\cdot 1^2 + m\cdot 1}, \ldots, \omega^{r\cdot (p-1)^2 + m\cdot (p-1)}\Big]^T \tag{6}$$

where r and m range between 0 and $p-1$. As in the case of random measurement matrices, it is natural to ask whether $\Sigma\Delta$ schemes can be used to quantize CS measurements obtained using chirp sensing matrices. Unfortunately, *(P1)* does not hold for these matrices. Next we show that this issue can be resolved by using a certain submatrix of the chirp sensing matrix Φ defined as follows.

Definition 1. *Let p be a prime and Φ the $p \times p^2$ chirp sensing matrix whose columns are indexed by r and m in $\mathbb{Z}_p$ and given as in* (6)*. Define $\bar{\Phi}$ as the $p \times p\lfloor\sqrt{p}\rfloor$ submatrix of Φ with columns as in* (6) *with r and m restricted to $\{0, 1, \ldots, p-1\}$, and $\{\lfloor\sqrt{p}\rfloor, 2\lfloor\sqrt{p}\rfloor, ..., (\lfloor\sqrt{p}\rfloor)^2\}$ respectively.*

Let $x \in \Sigma_k^n$, fix a quantization order r and let U be as defined in Section II-A. We propose to use $U\bar{\Phi}$ to obtain CS measurements $y = U\bar{\Phi}$ which we will then quantize using an rth-order $\Sigma\Delta$ scheme to obtain q. We will analyze the corresponding approximation error for a fixed sparsity level k as the number of measurements grows. In the case of $U\bar{\Phi}$, as we increase the number of measurements p, we must also increase the ambient dimension – as the measurement matrix is p-by-$p\lfloor\sqrt{p}\rfloor$. To deal with this, we embed x, originally in $\mathbb{R}^n$ into the higher dimensional space by padding it with zeros.

The following theorem shows that $\bar{\Phi}$ satisfies the property *(P1)* of order (k, ℓ) for appropriate choices of k and ℓ.

Theorem 3. *Consider the $p \times p\lfloor\sqrt{p}\rfloor$ matrix $\bar{\Phi}$ as defined in Definition 1. Then, there exists a prime number p_0 such that for $p \geq p_0$, the matrix Φ satisfies the property (P1) of order (k, ℓ) for $k \leq \sqrt[4]{p}\log p$ and $\ell = \lfloor p^{3/4}\log^2 p\rfloor$.*

Remark 4. *Theorem 3 shows that the restriction of $\bar{\Phi}$ to its top ℓ rows satisfies the RIP of order k such that $\delta_{2k} < 1/9$; such a conclusion cannot be obtained using the RIP bounds for chirp sensing matrices. The proof – given in [1] – relies on estimating certain Gauss-type exponential sums and will be omitted here due to lack of space.*

Corollary 1. *Let $x \in \Sigma_k^n$, let p_0 be as defined in Theorem 3, and suppose that $p_1 > p_0$ is a prime number such that $k \leq \sqrt[4]{p_1}\log p_1$. Then, for any $p \geq p_1$, x can be approximated by $\hat{x}$, the solution to* (1)*, if*

(i) the measurement matrix is $U\bar{\Phi}$, where $\bar{\Phi}$ is the $p \times p\lfloor\sqrt{p}\rfloor$ matrix defined as in Definition 1, and
(ii) q is obtained by quantizing $U\bar{\Phi}$ using an rth order $\Sigma\Delta$ scheme.

In the noise-free case, as we increase the number of measurements p, the approximation error satisfies

$$\|x-\hat{x}\|_2 \leq C(\log p)^{2r-1} p^{-\frac{1}{4}(r-\frac{1}{2})} \tag{7}$$

where C is a constant that does not depend on p_0, p, and p_1.

Remark 5. *Note that the error decay rate $O(p^{-\frac{1}{4}(r-\frac{1}{2})})$ (up to a factor logarithmic in p) given in Corollary 1 is inferior to $O(p^{-(r-\frac{1}{2})})$ which we obtain with random matrices (with $m=p$ measurements). This behaviour is due to the fact that the both dimensions of $\bar{\Phi}$ increase as we increase p. One way to circumvent this issue is to restrict the maximum number of measurements to some $p_{\max}$. Indeed, that way one can show that the approximation error behaves like $p^{-(r-\frac{1}{2})}$, similar to the case with random matrices. The experiments we present in the next section illustrate this behaviour.*

V. Numerical experiments

Here we present outcomes of numerical experiments within the deterministic CS setup we explored in the previous section. We consider the noise-free setting with exactly sparse signals. Specifically, we consider $p = 61$, 137, 223, 307, 397, 487, 593, 677, 787—these correspond to the $(18+15r)$th prime for $r = 0, 1, \cdots, 8$—and for each prime p, we draw 20 signals, each of which is a 4-sparse signal with a random support chosen from the set $\{1, 2, \cdots, 61\lfloor\sqrt{61}\rfloor\}$, and whose entries are chosen independently from a standard Gaussian distribution. In other words, the actual ambient dimension of signals that are considered is $61\lfloor\sqrt{61}\rfloor = 427$. For each such signal, we compute the CS measurements $y = U\bar{\Phi}$ which we subsequently quantize using a stable rth-order $\Sigma\Delta$ scheme to obtain q with $r = 1$. Next, we reconstruct an approximation $\hat{x}$ of x using (1) where we set $\Phi = U\bar{\Phi}$, $\delta = 0.1$, $r = 1$, and $\epsilon = 0$. Finally, for each p, we compute the average $\|x-\hat{x}\|_2$. We repeat this experiment with a second-order $\Sigma\Delta$ quantizer, i.e., $r = 2$. The results are reported in Figure 1. Observe that the approximation error behaves like $p^{-(r-1/2)}$, similar to the case with random matrices as discussed in Remark 5 as opposed to order $p^{-\frac{1}{4}(r-\frac{1}{2})}$ of Corollary 1.

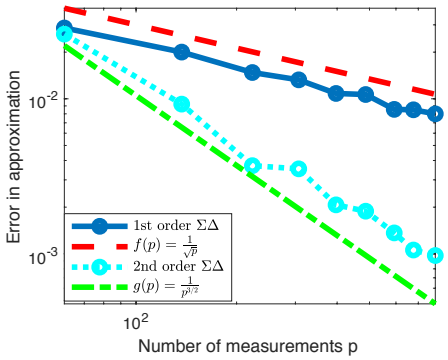

Fig. 1. Approximation error after using $\Sigma\Delta$ quantization order 1 and 2 to quantize CS measurements of 4-sparse signals and reconstructing via the one-stage recovery method with 4-sparse signals. We also plot the graphs of $f(p) = c_1 p^{-1/2}$ and $g(p) = c_2 p^{-3/2}$ for comparison.

References

[1] A. Ahmadieh and Ö. Yılmaz. One stage recovery for $\Sigma\Delta$ quantization in compressed sensing with chirp sensing matrices. Technical report, The University of British Columbia, Mathematics Department, March 2019.

[2] L. Applebaum, S. D. Howard, S. Searle, and R. Calderbank. Chirp sensing codes: Deterministic compressed sensing for fast recovery. *Applied and Computational Harmonic Analysis*, 26(2):283–290, 2009.

[3] J. Benedetto, A. Powell, and Ö. Yılmaz. Sigma-delta quantization and finite frames. *IEEE International Conference on Acoustics*, 52(5):1990–2005, 2004.

[4] J. Benedetto, A. Powell, and Ö. Yılmaz. Second-order sigma-delta quantization of finite frame expansions. *Applied and Computational Harmonic Analysis*, 20(1), 2006.

[5] B. Bodmann and V. Paulsen. Frames, graphs, and erasures. *Linear Algebra and its Applications*, 404, 2005.

[6] P. Boufounos, L. Jacques, F. Krahmer, and R. Saab. *Quantization and compressed sensing. Chapter in "Compressed sensing and its applications" (edited by H. Boche, R. Calderbank, G. Kutyniuk, J. Vybiral)*. Springer, 2015.

[7] P. T. Boufounos and R. G. Baraniuk. One-bit compressive sensing. *Proc. of Conf. on Information Sciences and Systems*, March 2008.

[8] E. Candes and T. Tao. Decoding by linear programming. *Information Theory, IEEE transactions on*, 51(12), December 2005.

[9] E. Chou, S. Gunturk, F. Krahmer, R. Saab, and Ö. Yılmaz. *Noise-shaping quantization methods for frame-based and compressive sampling systems, chapter 4 in "Sampling Theory, A Renaissance" (edited by G. Pfander)*. Birkhauser, Boston, 2015.

[10] D. Donoho. Compressed sensing. *IEEE Transactions on Signal Processing*, 52(4):1289–1306, 2006.

[11] C.S. Gunturk, M. Lammers, A. Powell, R. Saab, and Ö. Yılmaz. Sobolev duals for random frames and sigma-delta quantization of compressed sensing measurements. *Foundation of Computational Mathematics*, 13(1):1–36, 2013.

[12] H. Inose and Y. Yasuda. A unity bit coding method by negative feedback. *Proc. IEEE*, 51(11):1524–1535, 1963.

[13] L. Jacques, J. N. Laska, P. T. Boufounos, and R. G. Baraniuk. Robust 1-bit compressive sensing via binary stable embeddings of sparse vectors. *IEEE Transactions on Information Theory*, 59(4), April 2013.

[14] F. Krahmer, R. Saab, and Ö. Yılmaz. Sigma-delta quantization of sub-gaussian frame expamsions and its application to compressed sensing. *Information and Inference*, 3(1):40–58, 2014.

[15] F. Krahmer and R. Ward. Lower bounds for the error decay incurred by coarse quantization schemes. *Applied and Computational Harmonic Analysis*, 32(1):131–138, 2012.

[16] K. Melnykova and Ö. Yılmaz. Memoryless scalar quantization for random frames. *ArXiv preprint ArXiv: 1804.02839*, 2018.

[17] Y. Plan and R. Vershynin. Robust 1-bit compressed sensing and sparse logistic regression: a convex programming approach. *Information Theory, IEEE transactions on*, 59(1):482–494, December 2012.

[18] Y. Plan and R. Vershynin. One-bit compressed sensing by linear programming. *Communications on pure and Applied Mathematics*, 66(8):1275–1297, 2013.

[19] A. Powell, R. Saab, and Ö. Yılmaz. *Quantization and finite frames, Chapter 8 in "Finite frames: Theory and Applications" (edited by P. Casazza and G. Kutyniok)*. Birkhauser, Boston, 2012.

[20] R. Saab, R. Wang, and Ö. Yılmaz. Quantization of compressive samples with stable and robust recovery. *Applied and Computational Harmonic Analysis*, 44(1):123–143, 2018.

[21] R. Schreier and G. Temes. *Understanding delta-sigma data converters*. Wiley, Piscataway, New Jersey, 2004.

[22] R. Wang. Sigma Delta quantization with harmonic frames and partial Fourier ensembles. *Journal of Fourier Analysis and Applications*, 24(6):1460–1490, December 2018.

[23] Ö. Yılmaz. Coarse quantization of highly redundant time-frequency representations of square-integrable functions. *Applied and Computational Harmonic Analysis*, 14(2):107–132, 2003.

[24] Ö. Yılmaz. On coarse quantization of tight gabor frame expansions. *International Journal of Wavelets*, 3(2):283–299, 2005.