Analog to digital conversion for finite frames

John J. Benedetto[a], Alexander M. Powell[b], and Özgür Yılmaz[c]

[a]Department of Mathematics, University of Maryland, College Park, MD, USA;
[b]Applied and Computational Mathematics, Princeton University, Princeton, NJ, USA;
[c]Department of Mathematics, The University of British Columbia, Vancouver, B.C., Canada

ABSTRACT

Sigma-Delta ($\Sigma\Delta$) schemes are shown to be an effective approach for quantizing finite frame expansions. Basic error estimates show that first order $\Sigma\Delta$ schemes can achieve quantization error of order $1/N$, where N is the frame size. Under certain technical assumptions, improved quantization error estimates of order $(\log N)/N^{1.25}$ are obtained. For the second order $\Sigma\Delta$ scheme with linear quantization rule, error estimates of order $1/N^2$ can be achieved in certain circumstances. Such estimates rely critically on being able to construct sufficiently small invariant sets for the scheme. New experimental results indicate a connection between the orbits of state variables in $\Sigma\Delta$ schemes and the structure of constant input invariant sets.

Keywords: Finite frames, Sigma-Delta ($\Sigma\Delta$) quantization, stability, quantization error.

1. INTRODUCTION

Given a signal x of interest, a first step towards a digital representation is to obtain an *atomic decomposition* for x. One expands x with respect to a set $\{e_n\}_{n\in\Lambda}$ of vectors to obtain the discrete representation

$$x = \sum_{n\in\Lambda} c_n e_n, \tag{1}$$

where the c_n are real or complex numbers, and the index set Λ is finite or countably infinite. Such an expansion is *redundant* if the choice of c_n in (1) is not unique. The discrete decomposition (1) is not digital since the coefficients c_n may take on a continuum of values. *Quantization* is the intrinsically lossy and nonlinear process of reducing the continuous range of the coefficients to a finite set. Given a finite set $\mathcal{A}$ of numbers, called a *quantization alphabet*, quantization approximates the decomposition (1) by a digital decomposition

$$\widetilde{x} = \sum_{n\in\Lambda} q_n e_n,$$

where $q_n \in \mathcal{A}$.

There are two thematically different approaches to practical quantization, *fine quantization* and *coarse quantization*. In fine quantization one approximates the individual coefficients c_n in (1) with high precision. In coarse quantization, one approximates each c_n with less precision, for example by imposing $c_n \in \{-1, 1\}$ in the (extreme) 1-bit case, but compensates for this by exploiting the redundancy present in (1).

We shall discuss the mathematical theory of quantization for the particular class of atomic decompositions given by finite frames for $\mathbb{R}^d$, with one area of focus on rigorous approximation error estimates and another on stability theorems. Many existing methods require nonrigorous statistical assumptions, and still give only relatively poor approximations. Our results greatly improve on this.

Further author information: (Send correspondence to J.J.B.)
J.J.B.: E-mail: jjb@math.umd.edu
A.M.P.: E-mail: apowell@math.princeton.edu
Ö.Y.: E-mail: oyilmaz@math.ubc.ca

2. FINITE FRAMES

A set $\{e_n\}_{n=1}^N \subseteq \mathbb{R}^d$ of vectors is a *finite frame* for $\mathbb{R}^d$ with *frame constants* $0 < A \leq B < \infty$ if

$$\forall x \in \mathbb{R}^d, \quad A||x||^2 \leq \sum_{n=1}^{N} |\langle x, e_n \rangle|^2 \leq B||x||^2. \tag{2}$$

Here $||\cdot||$ is the d-dimensional Euclidean norm. A frame is *unit norm* if each $||e_n|| = 1$, and it is *tight* if the frame constants are equal, i.e., $A = B$. Although frames can be defined in separable Hilbert spaces we shall only consider finite frames for $\mathbb{R}^d$ here.

If $\{e_n\}_{n=1}^N$ is frame for $\mathbb{R}^d$ then there exists a *dual frame* $\{\widetilde{e}_n\}_{n=1}^N$ for $\mathbb{R}^d$ which gives the canonical frame decompositions

$$\forall x \in \mathbb{R}^d, \quad x = \sum_{n=1}^{N} \langle x, e_n \rangle \widetilde{e}_n = \sum_{n=1}^{N} \langle x, \widetilde{e}_n \rangle e_n. \tag{3}$$

Moreover, if $\{e_n\}_{n=1}^N$ is a tight frame with frame constant A, then the dual frame is $\{\frac{1}{A} e_n\}_{n=1}^N$, and the frame expansions in (3) give

$$\forall x \in \mathbb{R}^d, \quad x = \frac{1}{A} \sum_{n=1}^{N} x_n e_n, \quad x_n = \langle x, e_n \rangle. \tag{4}$$

For unit norm tight frames for $\mathbb{R}^d$, the frame constant A is solely determined by the size N of the frame and the dimension d. In this case, one can show[1–4] that $A = N/d$, which directly reflects the frame's redundancy.

It is straightforward to construct *generic* finite frames; in fact, any finite set of vectors is a frame for its span. However, useful frames demand a high degree of *geometric structure*. This is analogous to the $L^2(\mathbb{R}^d)$ setting where structured systems such as wavelet and Gabor frames have set the bar for implementability and performance. Therefore, the problem of building structured finite frames has emerged as an important and extremely active area.

There are many constructions of unit norm tight finite frames. A simple example in $\mathbb{R}^2$ is given by

$$F_N = \{(\cos(2\pi n/N), \sin(2\pi n/N))\}_{n=1}^N, \tag{5}$$

for any fixed $N > 2$. This corresponds to the Nth roots of unity. Vertices of the platonic solids[1] provide examples in $\mathbb{R}^3$. The classical examples in $\mathbb{R}^d$ are *harmonic frames.*[2,4] Harmonic frames are built using discrete Fourier matrices, and can be made arbitrarily redundant. More advanced examples include Grassmanian frames[5] for wireless communication and multiple description coding, geometrically uniform (GU) frames,[6] ellipsoidal frames,[7] and unions of orthonormal bases for code division multiple access (CDMA) systems.[8] There are frames which especially provide robustness with respect to erasures,[9,10] there are approaches which are rooted in group theory,[11] and others which are closely linked with sphere and line packing problems.[5]

3. QUANTIZATION OF FINITE FRAME EXPANSIONS

3.1. The quantization problem

Let us begin by stating the general *quantization problem* for unit norm tight finite frame expansions. If K is a positive integer and $\delta > 0$, then the 2K level *midrise quantization alphabet* with step size δ is the finite set of numbers

$$\mathcal{A}_K^\delta = \{-(K-1/2)\delta, -(K-3/2)\delta, \cdots, -\delta/2, \delta/2, \cdots, (K-1/2)\delta\}.$$

Given $x \in \mathbb{R}^d$, a unit norm tight finite frame $\{e_n\}_{n=1}^N$for $\mathbb{R}^d$, and the corresponding frame expansion (4), design low complexity algorithms to find a *quantized frame expansion*

$$\widetilde{x} = \frac{d}{N} \sum_{n=1}^{N} q_n e_n, \quad q_n \in \mathcal{A}_K^\delta \tag{6}$$

such that the *approximation error* $||x - \widetilde{x}||$ is small.

The reconstruction (6) from the quantized coefficients q_n is called *linear reconstruction.*

3.2. Background on quantization of finite frames

The *scalar quantizer* is a basic component of quantization algorithms for finite frame expansions. Given the quantization alphabet $\mathcal{A}_K^\delta$, the associated scalar quantizer is the function $Q : \mathbb{R} \to \mathcal{A}_K^\delta$ defined by

$$Q(x) = \arg\min_{a \in \mathcal{A}_K^\delta} |x - a|. \tag{7}$$

In other words, Q quantizes real numbers by rounding them to the nearest element of the quantization alphabet.

The classical and most common approach to quantizing the finite frame expansion (3) is first to quantize each frame coefficient $\langle x, e_n \rangle$ by $Q(\langle x, e_n \rangle)$. This step, often referred to as *pulse code modulation* (PCM), is then followed by linear reconstruction

$$\widetilde{x} = \frac{d}{N} \sum_{n=1}^{N} q_n e_n, \qquad \text{where} \qquad q_n = Q(\langle x, e_n \rangle). \tag{8}$$

This approach, while simple, gives highly non-optimal estimates for the approximation error $||x - \widetilde{x}||$ when the frame has even moderate amounts of redundancy. The basic deterministic estimate for this approach is $||x - \widetilde{x}|| \leq d\delta/2$, and it does not utilize any of the frame's redundancy. Practical analysis usually makes nonrigorous probabilistic assumptions on quantization noise[2] and leads to the *mean square error (MSE)* estimate

$$MSE = E(||x - \widetilde{x}||^2) = \frac{d^2 \delta^2}{N}.$$

Here, $E(\cdot)$ denotes expectation with respect to the associated probabilistic assumptions.[2]

Other existing approaches to finite frame quantization improve the quantization error by using more advanced reconstruction strategies. Given the frame expansion (3), such schemes still encode the frame coefficients using PCM, but do not require the linear reconstruction rule (8). Consistent reconstruction is one especially important class of nonlinear reconstruction which can be achieved in different ways.[2, 12] For example, projection onto convex sets (POCS) and linear programming were used to this end.[2, 13, 14] PCM with consistent reconstruction[2] outperforms PCM with linear reconstruction and achieves an improved MSE of order $1/N^2$. A drawback of consistent reconstruction is the higher computational complexity associated with it.

A differently motivated approach[15, 16] proceeds by producing equivalent vector quantizers (EVQ) with periodic structure. While this approach allows for simple reconstruction, it is primarily suited for low to moderate dimensions, and for frames with low redundancy. There is also a different method based on predictive quantization.[17] This noise shaping approach is related to our work, and was shown to be effective for subband coding applications.

3.3. First order ΣΔ quantization

In[18, 19] the authors investigated the use of first order Sigma-Delta (ΣΔ) schemes for quantizing finite frame expansions. In contrast to PCM, ΣΔ schemes are iterative algorithms and thus require one to specify an order in which frame coefficients are quantized. Let p be a permutation of $\{1, \cdots, N\}$ which denotes the quantization ordering, let Q be the scalar quantizer associated to the alphabet $\mathcal{A}_K^\delta$, and let $\{x_n\}_{n=1}^N$ be a sequence of frame coefficients corresponding to some $x \in \mathbb{R}^d$ and a frame $F = \{e_n\}_{n=1}^N$.

Definition 3.1. *The* first order ΣΔ scheme with ordering p and alphabet $\mathcal{A}_K^\delta$ *is defined by the iteration:*

$$\begin{aligned} u_n &= u_{n-1} + x_{p(n)} - q_n, \\ q_n &= Q(u_{n-1} + x_{p(n)}), \end{aligned} \tag{9}$$

for $n = 1, \cdots, N$*, where* $u_0 = 0$.

The u_n are internal state variables of the scheme, and the q_n are the quantized coefficients from which one reconstructs $\widetilde{x}$ by the linear reconstruction

$$\widetilde{x} = \frac{d}{N} \sum_{n=1}^{N} q_n e_{p(n)}.$$

The choice of the permutation p should depend solely on the frame being used, and not on the specific signal x being quantized. Roughly speaking, a good choice of quantization ordering p allows one to use the redundancy of the frame, i.e., interdependencies between the frame vectors, to iteratively compensate for the errors incurred when each frame coefficient x_n is replaced by some $q_n \in \mathcal{A}_K^\delta$. We define the notion of *frame variation* to quantify the role of the frame and permutation p in the $\Sigma\Delta$ algorithm.

DEFINITION 3.2. *Let $F = \{e_n\}_{n=1}^N$ be a unit-norm tight frame for $\mathbb{R}^d$ and let p be a permutation of $\{1, 2, \cdots, N\}$. The* frame variation of F with respect to p *is defined by*

$$\sigma(F,p) = \sum_{n=1}^{N-1} ||e_{p(n)} - e_{p(n+1)}||.$$

The basic error estimate for first order $\Sigma\Delta$ quantization of finite frame expansions may now be stated as follows.[18, 19]

THEOREM 3.3. *Consider the first order $\Sigma\Delta$ scheme with ordering p and alphabet $\mathcal{A}_K^\delta$. Let $F = \{e_n\}_{n=1}^N$ be a unit-norm tight frame for $\mathbb{R}^d$ and let $x \in \mathbb{R}^d$ satisfy $||x|| < (K - 1/2)\delta$. If $\widetilde{x}$ is the quantized output of the first order $\Sigma\Delta$ scheme then*

$$||x - \widetilde{x}|| \leq \frac{\delta d}{2N}(\sigma(F,p) + 1). \tag{10}$$

For certain infinite families of frames and associated permutations it is possible to obtain uniform bounds on the frame variation independent of the frame size. For example, if p_N is the identity permutation of $\{1, \cdots, N\}$ and if F_N is the Nth roots of unity frame, as in (5), then[18]

$$\sigma(F_N, p) \leq 2\pi.$$

Likewise, we have shown[18] that for a fixed dimension d the harmonic frames H_N^d are an infinite family of frames for which the frame variation has a relatively small uniform bound indepedent of N for simple choices of permutations p_N. For such families of frames, (10) implies that the approximation error $||x - \widetilde{x}||$ of the $\Sigma\Delta$ scheme (9) is at most of order $1/N$.

The $\Sigma\Delta$ error estimates can be improved beyond (10). We have proven[18] that under certain technical assumptions the $\Sigma\Delta$ scheme satisfies the refined bound

$$||x - \widetilde{x}|| \leq C_x \frac{\log N}{N^{5/4}}, \tag{11}$$

but with the constant C_x depending on x.

The improved estimate (11) illustrates some of the difficulties in the finite frame setting, since results for $\Sigma\Delta$ quantization of bandlimited signals[20] lead one to expect error of order $(\log N)/N^{\frac{4}{3}}$. Refined estimates in the bandlimited setting are valid away from points where the signal has a vanishing derivative. The non-local nature of estimates in the finite frame setting causes technical problems when the sequence of frame coefficients behaves like a function with points of vanishing derivative, as $N \to \infty$. The specific difficulties arise when we use stationary phase methods to estimate certain exponential sums associated with the $\Sigma\Delta$ scheme.[18]

EXAMPLE 3.4. *Let F_N be the Nth roots of unity frame for $\mathbb{R}^2$, as in (5), and let $x = (.157, 1/\sqrt{\pi})$. Let $\widetilde{x_N}$ be the quantized output of the first order $\Sigma\Delta$ algorithm with alphabet $\mathcal{A}_1^2 = \{-1, 1\}$, when the Nth roots of unity frame is used. Figure 1 shows a log-log plot of the approximation error $||x - \widetilde{x_N}||$ as a function of N.*

3.4. Second order $\Sigma\Delta$ quantization

Although (10) gives an error estimate whose utilization of redundancy is of order $1/N$, it is natural to seek even better utilization of the frame redundancy. Higher order schemes make this possible. Work on using higher order $\Sigma\Delta$ schemes to quantize finite frame expansions was initiated in.[21]

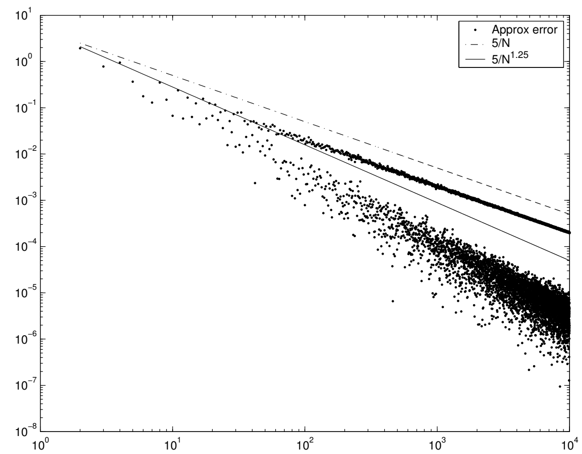

Figure 1. The frame coefficients of $x = (.157, 1/\sqrt{\pi})$ with respect to the Nth roots of unity tight frame are quantized using the first order $\Sigma\Delta$ scheme (9) with quantization alphabet $\mathcal{A}_1^2 = \{-1, 1\}$. The figure shows a log-log plot of the approximation error $||x - \widetilde{x}||$ as a function of the frame size N, compared with $5/N$ and $5/N^{1.25}$.

Let $\{x_n\}_{n=1}^N$ be a sequence of frame coefficients and let p be a permutation of $\{1, 2, \cdots, N\}$. We shall consider the particular 1-bit second order $\Sigma\Delta$ scheme[21–23] defined by the iteration:

$$\begin{aligned} u_n &= u_{n-1} + x_{p(n)} - q_n, \\ v_n &= v_{n-1} + u_n, \\ q_n &= \frac{\delta}{2}\text{sign}(u_{n-1} + \frac{v_{n-1}}{2}), \end{aligned} \tag{12}$$

for $n = 1, \cdots, N$, where $u_0 = v_0 = 0$.

The following theorem[21] is a representative error estimate for the second order $\Sigma\Delta$ scheme (12) in the setting of finite frames.

THEOREM 3.5. *Let $F_N = \{e_n\}_{n=1}^N$ be the Nth roots of unity frame for $\mathbb{R}^2$, and let p be the identity permutation of $\{1, 2, \cdots, N\}$. Suppose that $x \in \mathbb{R}^2, ||x|| \leq \frac{\delta}{2}\alpha$, where $0 < \alpha$ is a sufficiently small fixed constant. Let $\{q_n\}_{n=1}^N$ be the quantized bits produced by (12), and suppose that the input to the scheme is given by the frame coefficients $\{x_{p(n)}\}_{n=1}^N$ of x. Then if N is even, we have*

$$\|x - \widetilde{x}\| \leq C\frac{d\delta}{N^2}, \tag{13}$$

and if N is odd, we have

$$C_1\frac{d\delta}{N} \leq ||x - \widetilde{x}|| \leq C_2\frac{d\delta}{N}. \tag{14}$$

One surprising point of Theorem 3.5 is that $\Sigma\Delta$ quantization behaves quite differently when used to quantize finite frame expansions than in the original setting of quantizing sampling expansions. In particular, when a stable second order $\Sigma\Delta$ scheme is used to quantize sampling expansions one has the approximation error estimate[24]

$$||f - \widetilde{f}||_{L^\infty(\mathbb{R})} \leq \frac{C}{\lambda^2},$$

where λ is the sampling rate. Theorem 3.5 shows that analagous approximation error estimates are generally not true in the setting of finite frames. In particular, for N odd the approximation error is at best of order $1/N$ as N tends to infinity, see (14). A key issue is that the finite nature of the problem for finite frame expansions

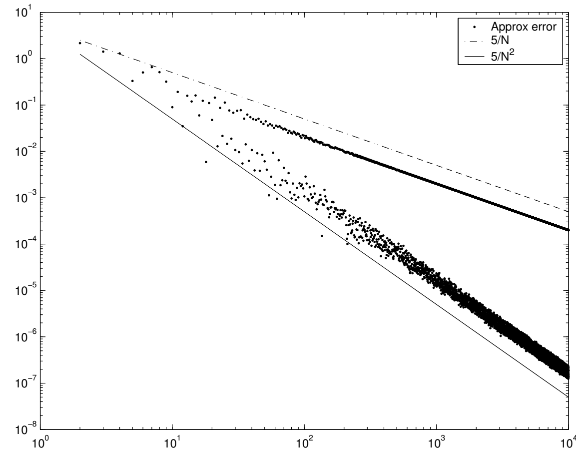

Figure 2. The frame coefficients of $x = (1/\sqrt{37}, 1/100)$ with respect to the Nth roots of unity tight frame are quantized using the second order $\Sigma\Delta$ scheme (12). The figure shows a log-log plot of the approximation error $||x - \widetilde{x}||$ as a function of the frame size N, compared with $5/N$ and $5/N^2$.

gives rise to non-zero boundary terms in certain situations, and that these boundary terms may negatively affect error estimates.

An important component in the proof of Theorem 3.5 is to prove that the scheme (12) is *stable*. In other words, one must show that there exists $0 < \alpha$ and a bounded set $S \subset \mathbb{R}^2$ such

$$\forall n = 1, \ldots, N, \quad |x_{p(n)}| \leq \alpha \implies (u_n, v_n) \in S.$$

Stability can be studied by rewriting (12) in the form of the following piecewise affine map:

$$\begin{pmatrix} u_n \\ v_n \end{pmatrix} = T_{x_n}(u_{n-1}, v_{n-1}) \equiv \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix} \begin{pmatrix} u_{n-1} \\ v_{n-1} \end{pmatrix} + \begin{pmatrix} x_n - \text{sign}(u_{n-1} + .5v_{n-1}) \\ x_n - \text{sign}(u_{n-1} + .5v_{n-1}) \end{pmatrix}.$$

The following theorem[21] shows that the second order $\Sigma\Delta$ scheme (12) is stable.

THEOREM 3.6. *There exists $0 < \alpha$ and $S \subset (-2, 2) \times \mathbb{R}$ such that if $|x_n| \leq \alpha$ for all n, and if u_n, v_n are the state variables of the second order linear $\Sigma\Delta$ scheme then*

$$\forall n, \quad (u_n, v_n) \in S.$$

In fact,

$$|x| < \alpha \implies T_x(S) \subseteq S.$$

Although the scheme (12) was previously known to be stable,[22, 23] Theorem 3.6 contains an important improvement because it shows that the invariant set S can be chosen to be sufficiently small, i.e., bounded inside $(-2, 2) \times \mathbb{R}$. The proof of Theorem 3.5 depends heavily on this fact.

EXAMPLE 3.7. *Let F_N be the Nth roots of unity frame for $\mathbb{R}^2$, as in (5), and let $x = (1/\sqrt{37}, .01)$. Let $\widetilde{x_N}$ be the quantized output of second order $\Sigma\Delta$ scheme (12) when the Nth roots of unity frame is used. Figure 2 shows a log-log plot of the approximation error $||x - \widetilde{x_N}||$ as a function of the frame size N.*

EXAMPLE 3.8. *Theorem 3.6 shows that for any sequence of sufficiently small input coefficients, x_n, the state variables of the second order $\Sigma\Delta$ scheme (12) stay bounded inside the set S from Theorem 3.6. As such, the stability theorem assumes no structure on the x_n, merely smallness. However, in our setting the coefficients have*

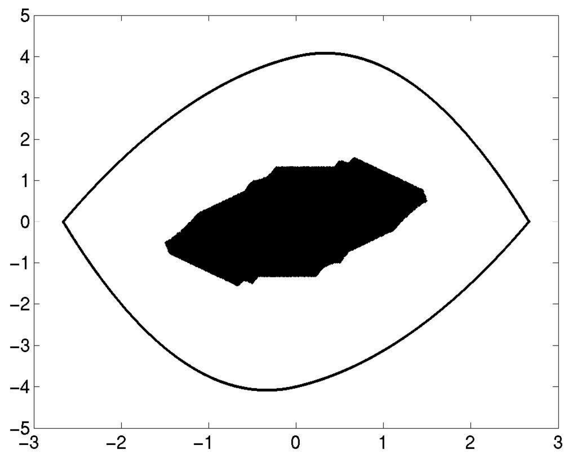

Figure 3. The orbit of the state variables (u_n, v_n) in (12) using $x = (1/3\pi, 1/100)$ and frame coefficients x_n given by the 1,000,000th roots of unity tight frame.

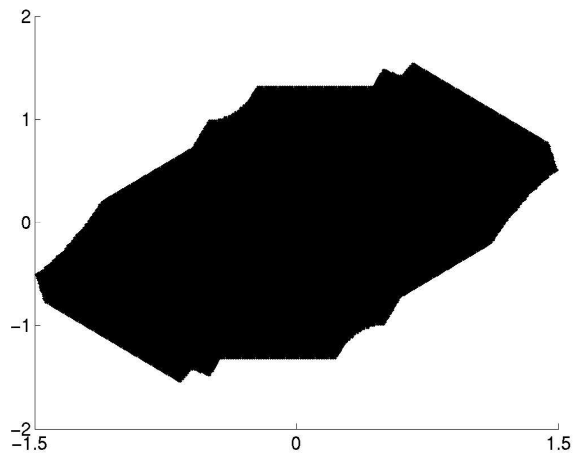

Figure 4. The union of the invariant sets Γ_y for $-||x|| \leq y \leq ||x||$, where x is as in Example 3.8.

the additional property of being frame coefficients $x_n = \langle x, e_n \rangle$. *In practice, this additional structure can lead to invariant sets which are smaller than those constructed by Theorem 3.6.*

For example, Figure 3 shows the orbit of the state variables (u_n, v_n) *in (12), using* $x = (1/3\pi, 1/100)$ *and the frame coefficients* x_n *determined by the 1,000,000th roots of unity tight frame for* $\mathbb{R}^2$. *For perspective, the orbit is shown inside of one of the invariant sets obtained in*[22] *(not the smaller set*[21] *of Theorem 3.6). Note that the orbit appears to be contained in a set whose width is noticeably smaller than* $(-2, 2) \times \mathbb{R}$.

It has been shown[25] *that for each fixed* x *there exists an invariant set* Γ_x *so that* $T_x(\Gamma_x) = \Gamma_x$ *and such that* Γ_x *tiles* $\mathbb{R}^2$ *by* $2\mathbb{Z} \times 2\mathbb{Z}$. *A proof of this phenomenon for general, stable, arbitrary order* $\Sigma\Delta$ *rules was recently obtained in.*[26] *In Figure 4 we plot the union of all the constant input invariant sets* Γ_y *for* $-||x|| \leq y \leq ||x||$, *where* $x = (1/3\pi, 1/100)$ *is as above. Observe that Figures 3 and 4 appear to be very similar.*

Example 3.8 leads us pose the following problem.

PROBLEM 3.9. *Mathematically explain the observations in Example 3.8. That is, prove a link between 1) orbits of the state variables* (u_n, v_n) *in (12) for "nicely" varying input sequences* x_n, *and 2) unions of constant invariant sets* Γ_y *for (12).*

Acknowledgments

The authors are grateful to Ingrid Daubechies, Sinan Güntürk, Mark Lammers, Aram Tangboondouangjit, and Nguyen Thao for valuable discussions related to the material.

This work was partially funded by NSF DMS Grant 0139759, NSF DMS Grant 0219233, and ONR Grant N0001-4021-0398. The second author was also supported by an Erwin Schrödinger Institute (ESI) Junior Research Fellowship, and all three authors are very grateful to Hans Feichtinger for having helped to arrange visits to ESI, where a portion of this work was completed.

REFERENCES

1. J. Benedetto and M. Fickus, "Finite normalized tight frames," *Advances in Computational Mathematics* **18**, pp. 357–385, February 2003.
2. V. Goyal, M. Vetterli, and N. Thao, "Quantized overcomplete expansions in $\mathbb{R}^n$: Analysis, synthesis, and algorithms," *IEEE Transactions on Information Theory* **44**, pp. 16–31, January 1998.
3. V. Goyal, J. Kovačević, and J. Kelner, "Quantized frame expansions with erasures," *Applied and Computational Harmonic Analysis* **10**, pp. 203–233, 2001.
4. G. Zimmermann, "Normalized tight frames in finite dimensions," in *Recent Progress in Multivariate Approximation*, K. Jetter, W. Haußmann, and M. Reimer, eds., Birkhäuser, 2001.
5. T. Strohmer and R. H. Jr., "Grassmannian frames with applications to coding and communications," *Applied and Computational Harmonic Analysis* **14**(3), pp. 257–275, 2003.
6. H. Bölcskei and Y. Eldar, "Geometrically uniform frames," *IEEE Transactions on Information Theory* **49**, pp. 993–1006, April 2003.
7. K. Dykema, D. Freeman, K. Kornelson, D. Larson, M. Ordower, and E. Weber, "Ellipsoidal tight frames and projection decomposition of operators," *Illinois Journal of Mathematics* **48**(2), pp. 477–489, 2004.
8. R. H. Jr., T. Strohmer, and A. Paulraj, "Grassmannian signatures for CDMA systems," *Proceedings IEEE Global Telecommunications Conference, San Francisco* , 2003.
9. R. Holmes and V. Paulson, "Optimal frames for erasures," *Lin. Alg. Appl.* **377**, pp. 31–51, January 2004.
10. P. Casazza and J. Kovačević, "Equal-norm tight frames with erasures," *Advances in Computational Mathematics* **18**, pp. 387–430, February 2003.
11. R. Vale and S. Waldron, "Tight frames and their symmetries," *Const. Approx.* **21**(1), pp. 83–112, 2005.
12. Z. Cvetković, "Resilience properties of redundant expansions under additive noise and quantization," *IEEE Transactions on Information Theory* **49**, pp. 644–656, March 2003.
13. D. Youla, "Mathematical theory of image restoration by the method of convex projections," in *Image recovery: theory and applications*, H. Stark, ed., Academic, 1987.
14. G. Strang, *Introduction to applied mathematics*, Wellesley-Cambridge, Cambridge, MA, 1986.
15. B. Beferull-Lozano and A. Ortega, "Efficient quantization for overcomplete expansions in $\mathbb{R}^n$," *IEEE Transactions on Information Theory* **49**, pp. 129–150, January 2003.
16. B. Beferull-Lozano and N. Sloan, "Quantizing using lattice intersections," *Discrete and Computational Geometry* **25**, pp. 799–824, 2003.
17. H. Bölcskei, "Noise reduction in oversampled filter banks using predictive quantization," *IEEE Transactions on Information Theory* **47**, pp. 155–172, January 2001.
18. J. Benedetto, A. Powell, and Ö. Yılmaz, "Sigma-delta ($\Sigma\Delta$) quantization and finite frames," *Submitted to IEEE Transactions on Information Theory* , 2004.
19. J. Benedetto, A. Powell, and Ö. Yılmaz, "Sigma-delta quantization and finite frames," *Proceedings IEEE ICASSP, Montreal* , 2004.
20. C. Güntürk, "Approximating a bandlimited function using very coarsely quantized data: Improved error estimates in sigma-delta modulation," *Journal of the American Mathematical Society* **17**(1), pp. 229–242, 2004.
21. J. Benedetto, A. Powell, and Ö. Yılmaz, "Second order sigma-delta ($\Sigma\Delta$) quantization of finite frame expansions," *To appear in Applied and Computational Harmonic Analysis* , 2004.

22. O. Yılmaz, “Stability analysis for several sigma-delta methods of coarse quantization of bandlimited functions,” *Constructive Approximation* **18**, pp. 599–623, 2002.
23. S. Pinault and P. Lopresti, “On the behavior of the double-loop sigma-delta modulator,” *IEEE Transactions on Circuits and Systems* **40**, pp. 467–479, August 1993.
24. I. Daubechies and R. DeVore, “Reconstructing a bandlimited function from very coarsely quantized data: A family of stable sigma-delta modulators of arbitrary order,” *Annals of Mathematics* **158**(2), pp. 679–710, 2003.
25. C. Güntürk, *Harmonic Analysis of Two Problems in Signal Quantization and Compression.* PhD thesis, Princeton University, 2000.
26. C. Güntürk and T. Nguyen, “Ergodic dynamics in $\Sigma\Delta$ quantization: tiling invariant sets and spectral analysis of error,” *Advances in Applied Mathematics* **34**(3), pp. 523–560, 2005.