RECONSTRUCTING SPARSE SIGNALS FROM THEIR ZERO CROSSINGS

Petros T. Boufounos and Richard G. Baraniuk

Rice University, ECE Department

ABSTRACT

Classical sampling records the signal level at pre-determined time instances, usually uniformly spaced. An alternative implicit sampling model is to record the timing of pre-determined level crossings. Thus the signal dictates the sampling times but not the sampling levels. Logan's theorem provides sufficient conditions for a signal to be recoverable, within a scaling factor, from only the timing of its zero crossings. Unfortunately, recovery from noisy observations of the timings is not robust and usually fails to reproduce the original signal. To make the reconstruction robust this paper introduces the additional assumption that the signal is sparse in some basis. We reformulate the reconstruction problem as a minimization of a sparsity inducing cost function on the unit sphere and provide an algorithm to compute the solution. While the problem is not convex, simulation studies indicate that the algorithm converges in typical cases and produces the correct solution with very high probability.

Index Terms— Level-crossing problems, implicit sampling, sparse reconstruction

1. INTRODUCTION

Classical signal acquisition systems usually sample signals at specified time instances and acquire the value of the signal amplitude. The defining characteristic of such systems is that the sampling times are in general signal-independent. Most acquisition systems sample uniformly in time or use a non-uniform but periodic sampling pattern. Both cases are well established theoretically and in practical implementations.

In contrast, in the case of sampling using level crossings—also referred to as implicit sampling—the signal dictates the sampling times. Specifically, the sampling device continuously compares the amplitude of the signal to pre-determined discrete levels and records the time of each level crossing together with the amplitude of the level crossed and, often, the direction of the crossing [1]. In such a scheme, the signal dictates the timing of the samples, but not the amplitude of the levels acquired. Although subsequent digital processing usually implies that the timing information is also discretized, recent work has developed hardware to perform analog-time, discrete-amplitude processing triggered by such level crossings (for examples, see [2, 3] and the references within).

Signal acquisition from level crossings is appealing, especially in cases in which the amplitude of a signal cannot be accurately measured but can be compared to a few, predefined levels. For example, a hardware implementation of a sampling system that detects the timing of zeros crossings is extremely simple since it only requires a simple comparator and an accurate timer. Furthermore, the timing of the level crossings can be a parsimonious signal representation for certain signal classes, as compared with uniform sampling [4, 5].

Despite its potential appeal in hardware design and signal representation, the problem has not, to date, been theoretically characterized in the general case. The most significant result is Logan's theorem [6], which states sufficient conditions for the zero crossings to uniquely represent the signal. However, a more general result on the recovery of a one-dimensional signal from the timings of arbitrary level crossings has remained elusive, despite significant results on multidimensional signals [7].

Unfortunately, recovery of the signal from the level crossing information is not always robust to sampling noise and timing quantization. Logan's theorem provides no robustness guarantee for the reconstruction performance. Indeed, practical algorithms are either unstable, or require the measurement of additional information on the signal to stabilize the representation [1, 8, 9]. This paper addresses this problem. Our approach is to stabilize the reconstruction by assuming the signal is sparse in some basis.

Recent results in signal acquisition, processing, and representation demonstrate that several classes of interesting signals can be sparsely represented in some basis. This information is used, for example, to compressively sample and reconstruct [10, 11] or denoise a signal [12]. The assumption of sparsity serves as a prior model to guide the signal reconstruction method and select the signal that is sparsest in the chosen basis among all the possible reconstruction candidates. The ability of the sparsity prior to describe signals and guide the reconstruction process is what makes it a good candidate to stabilize the reconstruction of signals from their zero crossings.

The goal of this paper is to provide a new formulation for the problem of reconstructing the signal from the timing of its zero crossings under the assumption that the acquired signal is sparse in some basis and to demonstrate experimentally that under this formulation it is possible to reconstruct the sparse signal with very high probability. Specifically, we assume the signal is sparse in the Fourier basis. Under this assumption it is possible to formulate the reconstruction problem as an optimization constrained on the unit sphere. Although the reconstruction is a non-convex problem, we present a simple algorithm that converges experimentally to the global optimum with high probability and, as verified by our experimental results, recovers the signal. Even though the algorithm works well in practice, the focus of this paper is neither on the efficiency of the algorithm nor on theoretical guarantees on its performance. Instead, the emphasis is on the new formulation and its applicability in signal recovery. Note that sparsity in the Fourier basis is not necessary; it is straightforward to modify the algorithm to use any basis of interest.

Sections 2.1 and 2.2 present an overview of implicit sampling and sparse signal recovery, respectively. They establish the notation and serve as a quick reference for the remainder of the paper. Section 3.1 uses the sparsity prior to reformulate the problem as an

Rice University, ECE Department, MS 380, P.O. Box 1892, Houston, TX 77251-1892, E-mail:{petrosb,richb}@rice.edu, Web: dsp.rice.edu. This work was supported by the grants DARPA/ONR N66001-06-1-2011 and N00014-06-1-0610, NSF CCF-0431150, ONR N00014-07-1-0936, AFOSR FA9550-07-1-0301, ARO W911NF-07-1-0502, ARO MURI W311NF-07-1-0185, and the Texas Instruments Leadership University Program.

optimization constrained on the unit sphere. An algorithm to solve the new formulation is introduced in Sec. 3.2, where it is also argued that it converges to the global optimum with very high probability. Section 4 provides experimental results that verify this claim and some discussion on future directions.

2. BACKGROUND

2.1. Signal Recovery From Zero Crossings

The recovery of a one-dimensional (1-D) signal from its zero crossings has not been fully explored to date. In [6] it is shown that the zero crossings of a real bandpass signal uniquely determine the signal within a scalar factor as long as its bandwidth is restricted to $[B, 2B - \epsilon]$ for any $0 < \epsilon < B$. A practical algorithm, based on the singular value decomposition (SVD), is developed in [9] but, as the authors acknowledge, it often fails to correctly reconstruct the signal. Although we follow their initial development closely, we reformulate the problem to exploit the assumption of sparsity, and we provide an algorithm that uses repeated random initializations to ensure that it recovers the signal with high probability.

Logan's theorem [6] applies to infinite dimensional spaces, but in this paper we restrict our attention to finite dimensional spaces of periodic bandpass signals. Under this assumption, the Fourier series coefficients are sufficient to represent the signal. Without loss of generality, we assume the signal has unit period, and therefore it contains only harmonic frequencies of the form $f_n = 2\pi n$, for $n \in \mathbb{Z}, B/2\pi \leq n \leq B/\pi$. Thus

$$x(t) = \sum_{n \in \mathcal{B}} a_n \cos(2\pi n t) + b_n \sin(2\pi n t), \tag{1}$$

where $\mathcal{B} = \{\lceil B/2\pi \rceil, \dots, \lfloor B/\pi \rfloor\} \equiv \{n_1, \dots, n_{N/2}\}$ denote the indices of the Fourier Series coefficients a_n and b_n sufficient to represent the signal, and $N = 2\lfloor B/2\pi \rfloor$ denotes the total number of those coefficients. For the remainder of this paper, we use $\mathbf{x} \in \mathbb{R}^N$ to denote the vector of coefficients a_n and b_n. Note that the ordering of a_n and b_n in $\mathbf{x}$ is not important, as long as it is consistent with all the other vectors and operators introduced later. Assuming an ordering of the cosine coefficients followed by the sine coefficients, the signal at time t satisfies:

$$x(t) = \langle \phi_t, \mathbf{x} \rangle, \tag{2}$$

where ϕ_t is the measurement vector at time t:

$$\phi_t = \big[\cos(2\pi n_1 t) \quad \dots \quad \cos\left(2\pi n_{N/2} t\right) \quad \sin(2\pi n_1 t) \quad \dots \quad \sin\left(2\pi n_{N/2} t\right) \big]^T. \tag{3}$$

Using Logan's theorem on this signal model, it follows that the zero crossings over a single signal period are sufficient to represent the signal and, therefore, the coefficients in $\mathbf{x}$. Denoting the set of zero crossings using $\mathcal{T} = \{t_k \in [0, 1), k = 1, \dots, Z\}$:

$$\langle \mathbf{x}, \phi_{t_k + l} \rangle = x(t_k + l) = 0, \text{ for all } t_k \in \mathcal{T} \text{ and } l \in \mathbb{Z}, \tag{4}$$

with Z denoting the number of zero crossings of the signal in the period $[0, 1)$. Any set of Z time instances, $t_k, k = 1, \dots, Z$ defines a sampling operator $\mathbf{\Phi}_{\{t_k\}}$ for all signals in the model:

$$\mathbf{\Phi}_{\{t_k\}}(\mathbf{x}) = [x(t_1), \dots, x(t_Z)]^T, \tag{5}$$

which is the matrix that transforms $\mathbf{x} \in \mathbb{R}^N$ to a vector in $\mathbb{R}^Z$:

$$\mathbf{\Phi}_{\{t_k\}} = [\phi_{t_1} \dots \phi_{t_Z}]^T. \tag{6}$$

In contrast to classical sampling systems the $\mathbf{\Phi}_{\mathcal{T}}$ is signal dependent and defined by the zero crossings of the sampled signal.

The sampled signal lies in the nullspace null$(\mathbf{\Phi}_{\mathcal{T}})$. It follows from Logan's theorem that $\mathbf{\Phi}_{\mathcal{T}}$ has rank $N-1$ and that its nullspace has dimension exactly equal to 1. Recovering any non-zero vector in this nullspace reconstructs $\mathbf{x}$ within a scalar factor, which cannot be identified since $x(t_k) = 0$ implies that $cx(t) = 0$ for all constants c. In principle, the 1-D nullspace can be identified using any nullspace identification method, such as the SVD suggested in [9].

Unfortunately the reconstruction is not robust to perturbations in the recorded timing of the zeros crossings. Given a set of noisy measurements of the zero crossings $\{t_k\}$, there is no guarantee that a bandpass signal with such zero crossings exists or, if it does, that it is close in some meaningful sense to the original signal. Experimental results demonstrate that the lack of such guarantees makes reconstruction difficult and unreliable. Specifically, errors in the measurement of the zero crossings set $\mathcal{T}$, due to measurement noise or quantization of the timings, usually make the resulting sampling operator $\mathbf{\Phi}_{\mathcal{T}}$ full rank. Thus, a one-dimensional null-space cannot be immediately recovered.

One approach to resolve this issue is to use the SVD to decompose the domain of $\mathbf{\Phi}_{\mathcal{T}}$ into singular vectors and then use the vector corresponding to the smallest singular value as the reconstructed signal. In practice, however, this method is very sensitive to measurement noise and quantization. Experimental results demonstrate that $\mathbf{\Phi}_{\mathcal{T}}$ has a number of small singular values with similar magnitude and the sampled signal is not always reconstructed using the singular vector corresponding to the smallest one. Although the signal of interest does lie in the space spanned by the vectors corresponding to the few smallest singular values, it is difficult to identify a single vector in the space to represent the signal. To guide the reconstruction we propose an extra constraint on the signal of interest, namely that it is sparse in some basis.

2.2. Sparse Signal Recovery

Recent advances in signal theory have demonstrated that the class of sparse signals is a good signal model for several kinds of interesting signals, often encountered in communications, radar, and image processing applications. The assumption in such a model is that the signal, when expressed in some sparsifying basis or dictionary, has very few significant coefficients, and the remaining ones are zero or approximately zero. For example, communications signals are often sparse in the short-time Fourier domain, and radar signals are sparse in the chirplet domain.

Specifically, a vector $\mathbf{x}$ is sparse in some basis if its basis expansion has a small ℓ_p-norm, for $p \leq 1$. The ℓ_p-norm (technically a pseudonorm if $p < 1$) of a vector is defined as:

$$\|\mathbf{x}\|_p = \left(\sum_i |(\mathbf{x})_i|^p \right)^{\frac{1}{p}}, \ p > 0, \tag{7}$$

where $(\cdot)_i$ denotes the i^{th} coefficient of the vector. For $p = 0$, the norm is defined as the number of nonzero coefficients, and is equal to the limit of $\|\mathbf{x}\|_p^p$ as p tends to 0.

The assumption of sparsity has proven very useful for resolving ambiguities in signal reconstruction. In most practical cases it is necessary to allow for some deviation, in an ℓ_2 sense, from the measurements or the signal to accommodate measurement noise. For example, in signal denoising applications, the noisy signal implicitly defines a set of signals that are close in the ℓ_2 sense. The denoised signal is the one in this set that is the sparsest when expressed in an appropriate dictionary [12]. Similarly, in compressive sensing applications, the compressive sensing measurements define an affine subspace in which the sampled signal may lie. The signal chosen by the reconstruction algorithm is the one with the sparsest representation in some basis that is close to the measurements [10, 11].

The general formulation of these problems often takes the form of an optimization that balances the sparsity of the signal with the deviation from the measurements, such as:

$$\widehat{\mathbf{x}} = \arg\min_{\mathbf{x}} \|\mathbf{x}\|_p^p + \frac{\lambda}{2}\|\mathbf{\Phi}\mathbf{x} - \mathbf{y}\|_2^2, \tag{8}$$

where $\mathbf{y}$ is the data (for example, the noisy signal or the compressive measurements), $\mathbf{\Phi}$ is the usually signal-independent sparsity dictionary, and $\mathbf{x}$ is the sparse representation of the signal. The parameter λ balances the sparsity of the solution with the fidelity to the data. Although the ℓ_0 norm is a strict measure of the signal sparsity, the ℓ_1 norm is often used in practice in (8) because it is continuous, convex and easy to optimize. In certain cases, depending on $\mathbf{\Phi}$, the solution of (8) for any $0 \leq p \leq 1$ has the same sparsity structure as for $p = 0$ [13, 14]. This minimization has been well studied, with a variety of algorithms to solve it (for some examples, see [15] and references within). For the remainder of this paper we focus on the case of $p = 1$, although the development in subsequent sections can be immediately generalized to any $0 < p \leq 1$.

3. SPARSE RECOVERY FROM ZERO CROSSINGS

3.1. Sparsity on the Unit Sphere

Although sparsity is a very effective model to resolve reconstruction ambiguities, the application of this model to the problem of reconstruction from zero crossings is not straightforward. A sparsity model on the reconstruction from the zero crossings can be imposed using the minimization in (8) with the signal-dependent sampling operator $\mathbf{\Phi}_{\mathcal{T}}$, as defined in (6) and setting the data vector to $\mathbf{y} = \mathbf{0}$. Unfortunately, in this case the optimal solution is $\mathbf{x} = \mathbf{0}$. This is due to the ambiguity of resolving $\mathbf{x}$ only within a scalar multiple of itself inherent in the problem formulation. Thus, we impose the additional constraint that the solution has unit energy, i.e., $\|\mathbf{x}\|_2 = 1$:

$$\widehat{\mathbf{x}} = \arg\min_{\mathbf{x}} \|\mathbf{x}\|_1 + \frac{\lambda}{2}\|\mathbf{\Phi}_{\mathcal{T}}\mathbf{x}\|_2^2, \text{ s.t. } \|\mathbf{x}\|_2 = 1. \tag{9}$$

The constraint in (9) places the solution on the surface of the unit sphere in $\mathbb{R}^N$ — a non-convex set. Thus, any convergence guarantees due to the convexity of the cost function in (8) do not hold for the solution of (9). Note that as λ tends to infinity, the optimum tends to the singular vector corresponding to the smallest singular value. Furthermore, as λ tends to 0 the solution tends to a vector $\mathbf{x}$ with a single non-zero component at the location corresponding to the column of $\mathbf{\Phi}_{\mathcal{T}}$ with the smallest norm.[1] In that sense, the problem is parallel to the overdetermined case of (8): as λ tends to infinity $\widehat{\mathbf{x}}$ converges to the least-squares solution $\widehat{\mathbf{x}} = \mathbf{\Phi}^\dagger\mathbf{y}$, where $(\cdot)^\dagger$ denotes the pseudoinverse. Similarly as λ tends to zero, (8) converges to $\widehat{\mathbf{x}} = \mathbf{0}$.

The use of the ℓ_2 norm of $\mathbf{\Phi}_{\mathcal{T}}\mathbf{x}$ as a measure of closeness to the samples seems reasonable and is also justified experimentally. Specifically, our preliminary experiments, not presented here due to space limitations, demonstrate that the empirical distribution of the signal amplitude at the quantized time of the zero crossing is close to a Gaussian with variance that decreases as the quantization becomes finer. The weight λ in (9) should also be set inversely proportional to the square root of this variance.

[1] This solution is reached only as the limit of the solution path as λ tends to zero. If the problem is solved with $\lambda = 0$, the solution can be any vector that has a single non-zero component in any location.

Algorithm 1 Renormalized Fixed Point Iteration

1. **Initialize:**
   $\widehat{\mathbf{x}}_{\text{best}} \leftarrow \arg\min_{\mathbf{x}} \|\mathbf{\Phi}_{\mathcal{T}}\mathbf{x}\|_2$, s.t. $\|\mathbf{x}\|_2 = 1$
2. **Seed Inner Iteration:**
   $\widehat{\mathbf{x}}_0 \leftarrow$ random vector $\in \mathbb{R}^N$ s.t. $\|\widehat{\mathbf{x}}_0\|_2 = 1$,
   $k \leftarrow 0$
   - (a) **Inner Loop Counter:**
     $k \leftarrow k + 1$
   - (b) ℓ_2 **Gradient:**
     $\mathbf{g}_k \leftarrow \mathbf{\Phi}_{\mathcal{T}}^T\mathbf{\Phi}_{\mathcal{T}}\mathbf{x}_{k-1}$
   - (c) ℓ_2 **Gradient Projection on Sphere Surface:**
     $\widetilde{\mathbf{g}}_k \leftarrow \mathbf{g}_k - \langle \mathbf{g}_k, \mathbf{x}_{k-1}\rangle \mathbf{x}_{k-1}$
   - (d) ℓ_2 **Gradient Descent:**
     $\mathbf{h} \leftarrow \widehat{\mathbf{x}}_{k-1} - \delta\widetilde{\mathbf{g}}_k$
   - (e) **Shrinkage:**
     $(\mathbf{u})_i \leftarrow \operatorname{sgn}\left((\mathbf{h})_i\right) \max\left\{|\,(\mathbf{h})_i\,| - \frac{\delta}{\lambda}, 0\right\}$, for all i,
   - (f) **Normalization:**
     $\mathbf{x}_k \leftarrow \mathbf{u}/\|\mathbf{u}\|$
   - (g) **Inner Iteration:** Repeat from (a) until convergence.
3. **Update Best Estimate:**
   If $\|\widehat{\mathbf{x}}_k\|_1 < \|\widehat{\mathbf{x}}_{\text{best}}\|_1$, set $\widehat{\mathbf{x}}_{\text{best}} \leftarrow \widehat{\mathbf{x}}_k$.
4. **Outer Loop:** Repeat from 2, L times.

3.2. Renormalized Fixed Point Iteration

One of the simplest approaches to solve (9) is to modify an existing iterative algorithm that solves (8) such that the solution is constrained by the algorithm to be on a sphere. Since the constraint is not convex, convergence to the global minimum is not guaranteed. However, the experimental results in Sec. 4 demonstrate that global convergence can be achieved in practice with very high probability.

This paper presents four modifications to the fixed point algorithm described in [15]. First, the constraint that the solution has unit norm is enforced by re-normalizing the estimate after each iteration. Second, the gradient descent step is modified such that the gradient is followed only in a direction tangent to the sphere. Third, since speed of convergence is not a primary concern of this paper, the continuation strategy discussed in [15] is not applied. Fourth, to combat the non-convexity of the search space, the algorithm is enclosed in an outer loop that performs several trials, each initialized with a new random seed. Algorithm 1 summarizes the modifications.

The outer loop of the algorithm seeds the inner iteration loop randomly and keeps track of the result with the smallest ℓ_1 norm. Since the search space is not convex, the inner loop is not guaranteed to converge to the global minimum. However, if each trial initialized with a new random seed converges to the global minimum with positive probability P, then it follows that the probability that at least one of the L trials produce the global minimum approaches 1 exponentially fast as a function of L:

$$P_{\text{global}} = 1 - (1 - P)^L. \tag{10}$$

The experimental results in Sec. 4 demonstrate that this is indeed often the case in practice. The outer loop is initialized in Step 1 using the smallest singular vector. This initialization can also be used as the seed for one of the trials in the inner loop.

The inner iteration loop follows [15] closely. Step (a) keeps track of the iteration counter in the inner loop. Step (b) computes

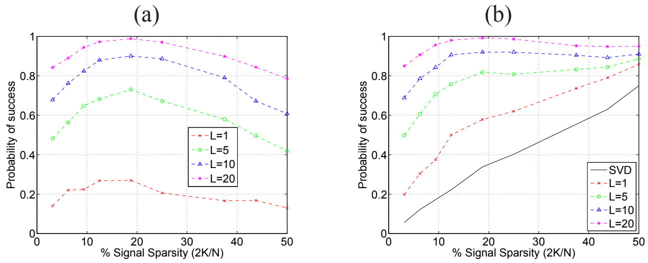

Fig. 1. Success rate after L iterations (a) if all L iterations are randomly initialized and (b) if $(L-1)$ iterations are randomly initialized while one is initialized using the SVD . The solid line plots the performance of the SVD method [9].

the gradient of $\|\mathbf{\Phi}_{\mathcal{T}}\mathbf{x}\|_2^2/2$ at the current solution estimate $\widehat{\mathbf{x}}_{k-1}$. Step (c) projects that gradient along the surface of the unit sphere by removing the component of the gradient parallel to the current solution estimate. That step is not necessary since the solution estimate is renormalized at the end of each iteration, but our experiments demonstrated that it improves the global convergence properties of the algorithm. Step (d) performs a gradient descent along the projected gradient and stores the result temporarily in $\mathbf{h}$. The descent step size is δ. Step (e) performs a shrinkage operation on $\mathbf{h}$ and stores the result temporarily in $\mathbf{u}$. The shrinkage, as described in [15], can be considered as a gradient descent with step size δ along the $\|\mathbf{x}\|_1/\lambda$ part of the cost function. Step (f) normalizes the result to have unit power and the current solution estimate is updated. Finally, Step (g) compares the new solution to the previous one and declares the algorithm converged or iterates from Step (a).

4. RESULTS AND DISCUSSION

This section presents experimental results for signal recovery from zero crossings using Alg. 1. In the experiments we set $\lfloor B/2\pi \rfloor = 128$, i.e., $N = 256$. The timing information in $\mathcal{T}$ is represented using 20 bits per zero crossing, with each value rounded towards 1 during quantization. To generate the sparse vector $\mathbf{x}$, a set of K random frequencies are selected, for which the sine and the cosine Fourier coefficients are populated with values drawn from a normal distribution; the remaining coefficients are set to 0. The resulting vector, which has sparsity $2K/N$, is subsequently normalized to have unit magnitude and transformed to the time domain to determine its zero crossings. These are used in Alg. 1 to recover the signal.

As the algorithm is executed the number of successes is recorded, where success is defined as recovering the signal with a signal-to-noise ratio of more than 20dB. The experiments are executed with $L = 1, 5, 10$, and 20. For each value of K we execute 500 experiments for a total of 10000 iterations of the inner loop. The results are reported in Fig. 1, in which we report the probability of recovering the signal within the first L iterations of the inner loop.

Figure 1(a) demonstrates the performance of the algorithm when all trials are seeded randomly, using vectors with coefficients drawn from an i.i.d. Bernoulli ± 1 distribution, normalized to be on the unit sphere. As predicted by (10) increasing L improves the performance of the algorithm. As the signal becomes less sparse (larger $2K/N$), the sparsity model does not accurately describe it, and the performance deteriorates. As the signal becomes more sparse (smaller $2K/N$) the algorithm on average encounters more cases with only low frequencies and therefore fewer zero crossings. The reconstruction of these signals is more difficult, which is reflected in the performance. Note that setting $L = 75$ in our experiments was sufficient to always recover the signal.

A simple modification of the algorithm uses the singular vector corresponding to the smallest singular value to seed one of the trials, and seeds the remaining $(L-1)$ trials randomly, as above. The results from this modification are reported in Fig. 1(b). For comparison we also plot the results using the SVD method in [9]. It is evident from the figure that this modification significantly improves the probability that the algorithm converges in few iterations, especially as K increases and the SVD method provided better initial estimates.

In summary, formulating the problem of sparse signal reconstruction from zero crossings as a sparsity inducing optimization on the unit sphere stabilizes the reconstruction. This formulation can be extended beyond Fourier sparsity to accommodate other sparsity bases by appropriately modifying the sampling operator $\mathbf{\Phi}_{\mathcal{T}}$, but the performance of such a modification needs to be evaluated. Furthermore, although the algorithm described provides the solution with high empirical probability, further research on reconstruction algorithms and their performance is necessary.

5. REFERENCES

[1] R. Hummel and R. Moniot, "Reconstructions from zero crossings in scale space," *IEEE Trans. Acoustics, Speech, and Sig. Proc.*, vol. 37, no. 12, pp. 2111–2130, 1989.

[2] B. Schell and Y. Tsividis, "Analysis of continuous-time digital signal processors," in *Proc. ISCAS 2007*, May 2007, pp. 2232–2235.

[3] Y. W. Li, K. L. Shepard, and Y. P. Tsividis, "A continuous-time programmable digital fir filter," *IEEE J. Solid-State Circuits*, vol. 41, no. 11, pp. 2512–2520, Nov. 2006.

[4] K. Guan and A. C. Singer, "A level-crossing sampling scheme for bursty signals," in *Proc. 40th Annual Conf. Information Sciences and Systems, 2006*, 2006, pp. 1357–1359.

[5] K. Guan and A. C. Singer, "A level-crossing sampling scheme for non-bandlimited signals," in *Proc. ICASSP 2006*, May 2006, vol. 3, pp. III–381–III–383.

[6] B. F. Logan, Jr., "Information in the zero crossings of bandpass signals," *AT&T Technical Journal*, vol. 56, pp. 487–510, Apr. 1977.

[7] A. Zakhor and A. V. Oppenheim, "Reconstruction of two-dimensional signals from level crossings," *Proc. IEEE*, vol. 78, no. 1, pp. 31–55, Jan. 1990.

[8] S. Mallat, "Zero-crossings of a wavelet transform," *IEEE Trans. Info. Theory*, vol. 37, no. 4, pp. 1019–1033, 1991.

[9] S. Roweis, S. Mahajan, and J. Hopfield, "Signal reconstruction from zero-crossings," Draft (Accessed Sept. '07), http://citeseer.ist.psu.edu/roweis98signal.html, Aug. 1998.

[10] E. J. Candès, "Compressive sampling," in *Proc. Int. Congress of Mathematics*, Madrid, Spain, 2006, pp. 1433–1452.

[11] D. Donoho, "Compressed sensing," *IEEE Trans. Info. Theory*, vol. 52, no. 4, pp. 1289–1306, Sept. 2006.

[12] S. S. Chen, D. L. Donoho, and M. A. Saunders, "Atomic decomposition by basis pursuit," *SIAM J. Scientific Computing*, vol. 20, no. 1, pp. 33–61, 1999.

[13] D. L. Donoho, "For most large underdetermined systems of linear equations the minimal ℓ_1-norm solution is also the sparsest solution," *Comm. Pure and Applied Math.*, vol. 59, no. 6, pp. 797–829, 2006.

[14] D. L. Donoho, "For most large underdetermined systems of equations, the minimal ℓ_1-norm near-solution approximates the sparsest near-solution," *Comm. Pure and Applied Math.*, vol. 59, no. 7, pp. 907–934, 2006.

[15] E. T. Hale, W. Yin, and Y. Zhang, "A fixed-point continuation method for ℓ_1-regularized minimization with applications to compressed sensing.," CAAM Tech. Report TR07-07, Rice University, July 2007.