Spectral Compressive Sensing

Marco F. Duarte, *Member, IEEE,* and Richard G. Baraniuk, *Fellow, IEEE*

Abstract

Compressive sensing (CS) is a new approach to simultaneous sensing and compression of sparse and compressible signals. A great many applications feature smooth or modulated signals that can be modeled as a linear combination of a small number of sinusoids; such signals are sparse in the frequency domain. In practical applications, the standard frequency domain signal representation is the discrete Fourier transform (DFT). Unfortunately, the DFT coefficients of a frequency-sparse signal are themselves sparse only in the contrived case where the sinusoid frequencies are integer multiples of the DFT's fundamental frequency. As a result, practical DFT-based CS acquisition and recovery of smooth signals does not perform nearly as well as one might expect. In this paper, we develop a new *spectral compressive sensing* (SCS) theory for general frequency-sparse signals. The key ingredients are an over-sampled DFT frame, a signal model that inhibits closely spaced sinusoids, and classical sinusoid parameter estimation algorithms from the field of spectrum estimation. Using peridogram and eigen-analysis based spectrum estimates (e.g., MUSIC), our new SCS algorithms significantly outperform the current state-of-the-art CS algorithms while providing provable bounds on the number of measurements required for stable recovery.

I. Introduction

The emerging theory of *compressive sensing* (CS) combines digital data acquisition with digital data compression to enable a new generation of signal acquisition systems that operate at sub-Nyquist rates. Rather than acquiring N samples $\mathbf{x} = [\mathbf{x}[1]\ \mathbf{x}[2]\ \ldots\ \mathbf{x}[N]]^T$ of an analog signal at the Nyquist rate, a CS system acquires $M < N$ measurements via the linear dimensionality reduction $\mathbf{y} = \Phi\mathbf{x}$, where Φ

Date: February 4, 2010. Rice University Technical Report TREE-1005. MFD is with the Program in Applied and Computational Mathematics, Princeton University, Princeton, NJ, 08544. RGB is with the Department of Electrical and Computer Engineering, Rice University, Houston, TX 77005. Email: mduarte@princeton.edu, richb@rice.edu. MFD and RGB were supported by grants NSF CCF-0431150 and CCF-0728867, DARPA/ONR N66001-08-1-2065, ONR N00014-07-1-0936 and N00014-08-1-1112, AFOSR FA9550-07-1-0301, ARO MURIs W911NF-07-1-0185 and W911NF-09-1-0383, and the Texas Instruments Leadership Program. MFD was also supported by NSF Supplemental Funding DMS-0439872 to UCLA-IPAM, P.I. R. Caflisch.

is an $M \times N$ measurement matrix. When the signal $\mathbf{x}$ has a *sparse* representation $\mathbf{x} = \Psi\theta$ in terms of an $N \times N$ orthonormal basis matrix Ψ, meaning that only $K \ll N$ out of N signal coefficients θ are nonzero, then the number of measurements required to ensure that $\mathbf{y}$ retains all of the information in $\mathbf{x}$ is just $M = O(K \log(N/K))$ [1–3]. Moreover, a sparse signal $\mathbf{x}$ can be recovered from its compressive measurements $\mathbf{y}$ via a convex optimization or iterative greedy algorithm. Random matrices play a central rôle as universal measurements, since they are suitable for signals sparse in any fixed basis with high probability. The theory also extends to noisy signals as well as to so-called *compressible* signals that are not exactly sparse but can be closely approximated as such. Sparse signals have coefficients θ that, when sorted, decay according to a power law: $|\theta[i]| < Ci^{-1/p}$ for some $p \leq 1$; the smaller the decay exponent p, the faster the decay and the better the recovery performance we can expect from CS.

A great many applications feature smooth or modulated signals that can be modeled as a linear combination of K sinusoids [4–7]:

$$\mathbf{x}[n] = \sum_{k=1}^{K} a_k \, e^{-j\omega_k n}, \tag{1}$$

where $\omega_k \in [0, 2\pi]$ are the sinusoid frequencies. When the sinusoids are of infinite extent, such signals have a K-sparse representation in terms of the discrete-time Fourier transform (DTFT),[1] since

$$X(\omega) = \sum_{k=1}^{K} a_k \, \delta(\omega - \omega_k), \tag{2}$$

where δ is the Dirac delta function. We will refer to such signals as *frequency-sparse*.

Practical applications feature signals of finite length N. In this case, the frequency domain tool of choice for analysis and CS recovery has been the discrete Fourier transform (DFT).[2] The DFT $\mathbf{X}[l]$ of N consecutive samples from the smooth signal model (1) can be obtained from the DTFT (2) by first convolving with a sinc function and then sampling:

$$\mathbf{X}[l] = \sum_{k=1}^{K} a_k \, \mathrm{sinc}\left(\frac{2\pi(l - l_k)}{N} \right), \tag{3}$$

where $\mathrm{sinc}(\omega) := \frac{\sin(\pi\omega)}{\pi\omega}$, and $l_k = \frac{N\omega_k}{2\pi}$.

Unfortunately, the DFT coefficients in (3) do not share the same sparsity property as the DTFT coefficients in (2), except in the (contrived) case when the sinusoid frequencies in (1) are *integral*, that is,

[1]Recall that the DTFT of a signal $\mathbf{x}$ is defined as $X(\omega) = \sum_{n=-\infty}^{\infty} \mathbf{x}[n] e^{-j\omega n}$ with inverse transformation $\mathbf{x}[n] = \frac{1}{2\pi} \int_{-\pi}^{\pi} X(\omega) e^{j\omega n} d\omega$.

[2]Recall that the DFT of a length-N signal $\mathbf{x}$ is defined as $\mathbf{X}[l] = \sum_{n=1}^{N} \mathbf{x}[n] e^{-j2\pi ln/N}$, $1 \leq l \leq N$, with inverse transformation $\mathbf{x}[n] = \frac{1}{N} \sum_{l=1}^{N} \mathbf{X}[l] e^{j2\pi ln/N}$, $1 \leq n \leq N$.

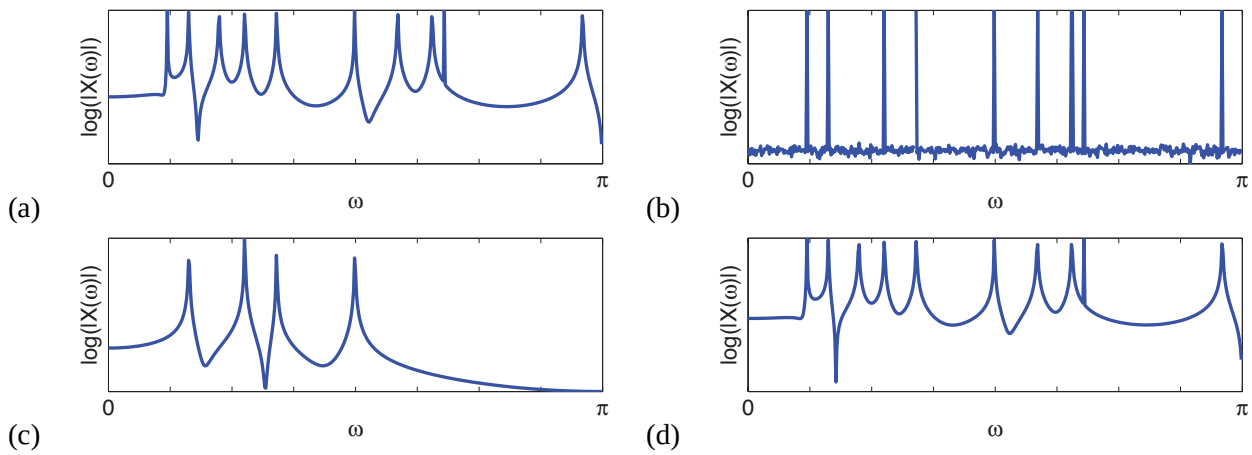

Fig. 1. *Compressive sensing (CS) sparse signal recovery from $M = 300$ noiseless random measurements of a signal of length $N = 1024$ composed of $K = 20$ complex-valued sinusoids with arbitrary real-valued frequencies. We compare the frequency spectra obtained from oversampled periodograms of (a) the original signal and its recovery using (b) standard CS using the orthonormal DFT basis (SNR = 5.35dB), (c) standard CS using a $10\times$ zero-padded, redundant DFT frame (SNR = -4.45dB), and (d) spectral CS using Algorithm 2 (SNR = 34.78 dB).*

when each and every l_k is equal to an integer. On closer inspection, we see that not only are most smooth signals *not sparse* in the DFT domain, but, owing to the slow asymptotic decay of the sinc function, they are *just barely compressible*, with a decay exponent of $p = 1$. As a result, practical CS acquisition and recovery of smooth signals does not perform nearly as well as one might expect (see Fig. 1(b) and the discussion in [7], for example).

The goal of this paper is to develop new CS recovery algorithms for practical smooth signals (with non-integral frequencies). The naïve first step is to change the signal representation to a zero-padded DFT, which provides samples from the signal's DTFT at a higher rate than the standard DFT. This is equivalent to replacing the DFT basis with a redundant frame [8] of sinusoids that we will call a *DFT frame*. Unfortunately, there exists a tradeoff in the use of these redundant frames for sparse approximation and CS recovery: if we increase the amount of zero-padding / size of the frame, then signals with non-integral frequency components become more compressible, which increases recovery performance. However, simultaneously, the frame becomes increasingly *coherent* [9, 10], which decreases recovery performance (see Fig. 1(c), for example). In order to optimize this tradeoff, we will leverage the last few decades of progress on sinusoid parameter estimation from the field of spectrum estimation [11–13] plus recent progress on model-based CS [14] and marry these techniques with a class of greedy

CS recovery algorithms. We will refer to our general approach as *spectral compressive sensing* (SCS).

The primary novelty of SCS is the concept of taming the coherence of the redundant DFT frame using an inhibition model that ensures the sinusoid frequencies ω_k of (1) are not too closely spaced. We will provide an analytical characterization of the number of measurements M required for stable SCS signal recovery under this model and will study the performance of the framework under parameter variations. As we see from Fig. 1, the performance improvement of SCS over standard DFT-based CS can be substantial.

While this paper focuses on frequency-sparse signals, the SCS concept generalizes to other settings featuring signals that are sparse in a parameterized redundant frame. Examples include the frames underlying localization problems [15–18], radar imaging [19–21], and manifold-based signal models [22].

This paper is organized as follows. Section II provides the usual background on CS and model-based CS, while Section III summarizes existing schemes for line spectrum estimation. Section IV develops our proposed SCS recovery algorithms, and Section V presents our experimental results. Section VI summarizes related work in this area, and Section VII closes with conclusions and suggestions for future work.

II. Background

A. Sparse approximation

A signal $\mathbf{x} \in \mathbb{R}^N$ is *K-sparse* ($K \ll N$) in a basis or frame[3] Ψ if there exists a vector θ with $\|\theta\|_0 = K$ such that $\mathbf{x} = \Psi\theta$. Here $\|\cdot\|_0$ denotes the ℓ_0 pseudo-norm, which simply counts the number of nonzero entries in the vector. Signal compression often relies on the existence of a known basis or frame Ψ such that for the signal of interest $\mathbf{x}$ there exists a *K-sparse approximation* $\mathbf{x}_K$ in Ψ that yields small approximation error $\|\mathbf{x} - \mathbf{x}_K\|_2$. When Ψ is a basis, the optimal K-sparse approximation of $\mathbf{x}$ in Ψ is trivially found through hard thresholding: we preserve only the entries of θ with the K largest magnitudes and set all other entries to zero. While thresholding is suboptimal when Ψ is a frame, there exist a bevy of *sparse approximation algorithms* that aim to find a good sparse approximation to the signal of interest. Such algorithms include basis pursuit [23], CoSaMP [24], matching pursuit [9], and iterative thresholding [25–28]. The approximation performance of the latter two algorithms is directly

[3]Recall that a frame is a matrix $\Psi \in \mathbb{R}^{D\times N}$, $D < N$, such that for all vectors $\mathbf{x} \in \mathbb{R}^D$, $A\|\mathbf{x}\|_2 \leq \|\Psi^T\mathbf{x}\|_2 \leq B\|\mathbf{x}\|_2$ with $0 < A \leq B < \infty$. A frame is a generalization of the concept of a basis to sets of possibly linearly dependent vectors [8].

tied to the *coherence* of the frame Ψ, defined as

$$\mu(\Psi) = \arg \min_{1 \leq i,j \leq N} |\langle \psi_i, \psi_j \rangle| ,$$

where ψ_i denotes the i^{th} column of Υ. For example, orthogonal matching pursuit (OMP) successfully obtains a K-sparse signal representation if $\mu(\Psi) \leq \frac{1}{16(K-1)}$ [9, 10].

B. Compressive sensing

Compressive Sensing (CS) is an efficient acquisition framework for signals that are sparse or compressible in a basis or frame Ψ. Rather than uniformly sampling the signal $\mathbf{x}$, we measure inner products of the signal against a set of measurement vectors $\{\phi_1, \ldots, \phi_M\}$; when $M < N$, we effectively compress the signal. By collecting the measurement vectors as rows of a measurement matrix $\Phi \in \mathbb{R}^{M \times N}$, this procedure can be written as $\mathbf{y} = \Phi\mathbf{x} = \Phi\Psi\theta$, with the vector $\mathbf{y} \in \mathbb{R}^M$ containing the CS measurements. We then aim to recover the signal $\mathbf{x}$ from the fewest possible measurements $\mathbf{y}$. Since $\Phi\Psi$ is a dimensionality reduction, it has a null space, and so infinitely many vectors $\mathbf{x}'$ yield the same recorded measurements $\mathbf{y}$. Fortunately, standard sparse approximation algorithms can be employed to recover the signal representation θ by finding a sparse approximation of $\mathbf{y}$ using the frame $\Upsilon = \Phi\Psi$.

The Restricted Isometry Property (RIP) has been proposed to measure the fitness of a matrix Υ for CS [1].

Definition 1: The *K-restricted isometry constant* for the matrix Υ, denoted by δ_K, is the smallest nonnegative number such that, for all $\theta \in \mathbb{R}^N$ with $\|\theta\|_0 = K$,

$$(1 - \delta_K)\|\theta\|_2^2 \leq \|\Upsilon\theta\|_2^2 \leq (1 + \delta_K)\|\theta\|_2^2. \tag{4}$$

A matrix has the RIP if $\delta_K > 0$. Since calculating δ_K for a given matrix requires a combinatorial amount of computation, random matrices have been advocated. For example, a matrix of size $M \times N$ with independent and identically distributed (i.i.d.) Gaussian entries with variance $1/M$ will have the RIP with very high probability if $K \leq M/\log(N/M)$. The same is true of matrices following Rademacher (± 1) or more general subgaussian distributions. Revisiting our previous example, OMP can recover a K-sparse representation θ from its measurements $\mathbf{y} = \Upsilon\theta$ if the restricted isometry constant $\delta_{K+1} < \frac{1}{3\sqrt{K}}$ [29].

C. Frequency-sparse signals

Recall from the introduction that frequency-sparse signals of the form (1) have a sparse DTFT (2). However, to exploit sparsity in CS, we require a discrete signal representation; thus, the DFT has been

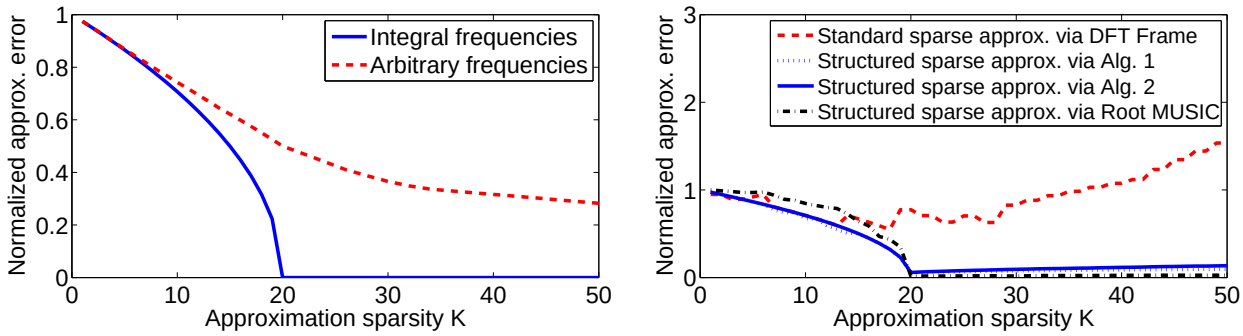

(a) (b)

Fig. 2. *Performance of K-term sparse approximation algorithms on signals of length $N = 1024$ containing 20 complex sinusoids of arbitrary integral and non-integral frequencies. We plot the approximation error as a function of the approximation sparsity K. (a) Orthonormal DFT basis approximation performance is perfect for signals with exclusively integral frequencies and atrocious for signals with non-integral frequencies. (b) Three potential approximation strategies for sinusoids with non-intgeral frequencies. Standard sparse approximation using the DFT frame $\Psi(c)$, $c = 10$, performs even worse than using the orthonormal DFT basis. Structured sparse approximations based on periodogram and Root MUSIC spectrum estimates perform much better.*

the tool of choice for frequency-sparse signals. Additionally, the fast Fourier transform (FFT) provides a very efficient algorithm to calculate the DFT coefficients of a signal. The DFT of an length-N signal $\mathbf{x}$ can be obtained as its representation in the orthonormal *DFT basis* $\mathbf{F}$, which has entries $\mathbf{F}[p,q] = e^{-j2\pi pq/N}/\sqrt{N}$, $1 \leq p, q \leq N$.

The DFT basis preserves the sparsity of the DTFT for frequency-sparse signals only when the signal components have *integral frequencies* of the form $2\pi n/N$, where n is an integer. Otherwise, the situation is decidedly more complicated due to the spectral leakage induced by windowing (sinc convolution). To illustrate the difficulty, Fig. 2(a) plots the sparse approximation error of signals of length $N = 1024$ containing 20 complex sinusoids of both integral and non-integral frequencies using the DFT basis. As expected, sparse approximation using a DFT basis fails miserably for signals with non-integral frequencies.

The naïve way to combat spectral leakage is to employ a redundant frame that we term a *DFT frame*. The DFT frame representation provides a finer sampling of the DTFT coefficients for the signal $\mathbf{x}$ observed. We let $c \in \mathbb{N}$ denote the frequency oversampling factor for the DFT frame, and define the

frequency sampling interval $\Delta := 2\pi/cN \in (0, 2\pi/N]$. We also let

$$\mathbf{e}(\omega) := \frac{1}{\sqrt{N}}[1 \; e^{j\omega/N} \; e^{j2\omega/N} \; \ldots \; e^{j\omega(N-1)/N}]^T$$

denote a normalized vector containing regular samples of a complex sinusoid with angular frequency $\omega \in [0, 2\pi)$. The DFT frame with oversampling factor c is then defined as $\Psi(c) := [\mathbf{e}(0) \; \mathbf{e}(\Delta) \; \mathbf{e}(2\Delta) \; \ldots \; \mathbf{e}(2\pi - \Delta)]^T$, and the corresponding signal representation $\theta = \Psi(c)^T\mathbf{x}$ provides us with cN equispaced samples of the signal's DTFT. Note that $\Psi(1) = \mathbf{F}$, the usual orthonormal DFT basis.

We can use the DFT frame $\Psi(c)$ to obtain sparse approximations for frequency-sparse signals with components at arbitrary frequencies; as the frequency oversampling factor c increases, the K-sparse approximation provided by $\Psi(c)$ becomes increasingly accurate. The proof of the following lemma is given in the Appendix.

Lemma 1: Let $\mathbf{x} = \sum_{k=1}^{K} a_k\mathbf{e}(\omega_k)$ be a K-frequency-sparse signal, and let $\mathbf{x}_K = \Psi(c)\theta_K$ be its best K-sparse approximation in the frame $\Psi(c)$, with $\|\theta_K\|_0 = K$. Then the corresponding best K-term approximation error for $\mathbf{x}$ obeys

$$\|\mathbf{x} - \mathbf{x}_K\|_2 \leq \|\mathbf{a}\|_1\sqrt{1 - \text{sinc}^2\frac{1}{2c}}, \tag{5}$$

where $\mathbf{a} = [a_1 \; \ldots \; a_K]^T$.

Unfortunately, standard algorithms that aim to find the sparse approximation of $\mathbf{x}$ in the frame $\Psi(c)$ do not perform well when c increases due to the high coherence between the frame vectors, particularly for large values of c:

$$\mu(\Psi(c)) = \frac{cN\sin(\pi/cN)}{\pi} \to 1 \text{ as } c \to \infty.$$

Due to this tradeoff, the maximum frequency oversampling factor that still allows for sparse representation of K-sparse signals is

$$c \leq \frac{1}{N \, \text{asinc}\left(\frac{1}{16(K-1)}\right)},$$

where $\text{asinc}(\cdot)$ denotes the inverse of the sinc function within the interval $[0, 1]$. In words, the sparsity K of the signal limits the maximum size of the redundant DFT frame that we can employ, and vice-versa. Fig. 2(b) demonstrates the performance of standard sparse approximation of the same signal with arbitrary frequencies as in Fig. 2(a), but using the redundant frame $\Psi(c)$ instead, with $c = 10$. Due to the high coherence of the frame $\Psi(c)$, the algorithm cannot obtain an accurate sparse approximation of the signal.

D. Model-based compressive sensing

While many natural and manmade signals and images can be described to first-order as sparse or compressible, the support of their large coefficients often has an underlying second-order inter-dependency structure. This structure can often be captured by a *union-of-subspaces* model that enables an algorithmic *model-based CS* framework to exploit signal structure during recovery [14, 30]. We provide a brief review of model-based CS below; in Section IV, we will use this framework to overcome the issues of sparse approximation and CS using coherent frames.

The set Σ_K of all length-N, K-sparse signals is the union of the $\binom{N}{K}$, K-dimensional subspaces aligned with the coordinate axes in $\mathbb{R}^N$. A *structured sparsity model* endows the K-sparse signal $\mathbf{x}$ with additional structure that allows only certain K-dimensional subspaces from Σ_K and disallows others. The signal model $\mathcal{M}_K$ is defined by the set of m_K allowed supports $\{\Omega_1, \ldots, \Omega_{m_K}\}$. Signals from $\mathcal{M}_K$ are called K-structured sparse. Signals that are well-approximated as K-structured sparse are called structured compressible.

If we know that the signal $\mathbf{x}$ being acquired is K-structured sparse or structured compressible, then we can relax the RIP constraint on the CS measurement matrix Υ to require isometry only for those signals in $\mathcal{M}_K$ and still achieve stable recovery from the compressive measurements $\mathbf{y} = \Upsilon\theta$. The *model-based* RIP requires that (4) holds *only* for signals with sparse representations $\theta \in \mathcal{M}_K$ [30, 31]; we denote this new property as $\mathcal{M}_K$-RIP to specify the dependence on the chosen signal model and change the model-based RIP constant from δ_K to $\delta_{\mathcal{M}_K}$ for clarity. This *a priori* knowledge reduces the number of random measurements required for model-based RIP with high probability to $M = \mathcal{O}(\log m_K)$ [30]. For some models, the reduction from $M = O(K \log(N/K))$ can be significant [14].

The $\mathcal{M}_K$-RIP property is sufficient for robust recovery of structured sparse signals using algorithms such as model-based CoSaMP and IHT [14]. These model-based CS recovery algorithms replace the standard optimal K-sparse approximation performed by thresholding with a *structured sparse approximation* algorithm $\mathbb{M}(\mathbf{x}, K)$ that returns the best K-term approximation of the signal $\mathbf{x}$ belonging in the signal model $\mathcal{M}_K$.

To summarize, the combination of a structured signal model and a structured sparse approximation algorithm enables us to design a model-based recovery algorithm that achieves a substantial reduction in the number of measurements required for stable recovery.

III. Parameter Estimation for Frequency-Sparse Signals

The goal of CS is to identify the values and locations of the nonzero / large coefficients of a sparse / compressible signal from a small set of linear measurements. For frequency-sparse signals, such an identification can be interpreted as a parameter estimation problem, since each coefficient index corresponds to a sinusoid of a certain frequency. Thus, in this case, CS aims to estimate the frequencies and amplitudes of the largest sinusoids present in the signal. In practice, most CS recovery algorithms iterate through a sequence of increasing-quality estimates of the signal coefficients by differentiating the signal's actual nonzero coefficients from spurious estimates; such spurious coefficients are often modeled as recovery noise.

Thus, we now briefly review the extensive prior work in parameter estimation for frequency-sparse signals embedded in noise [11, 12]. We start with the simple sinusoid signal model, expressed as $\mathbf{x} = A\mathbf{e}(\omega) + \mathbf{n}$, where $\mathbf{n} \sim \mathcal{N}(0, \sigma^2\mathbf{I})$ denotes a white noise vector with i.i.d. entries. The model parameters are A and ω, the complex amplitude and frequency of the sinusoid, respectively.

A. Periodogram-based methods

The maximum likelihood estimator (MLE) of the amplitude A when the frequency ω is known is given by the DTFT of $\overline{\mathbf{x}}$, the zero-padded, infinite length version of the signal $\mathbf{x}$, at frequency ω: $\widehat{A} = \frac{1}{N}\overline{\mathbf{X}}(\omega) = \langle \mathbf{e}(\omega), \mathbf{x} \rangle$. [11, 12] Furthermore, since only a single sinusoid is present, the MLE for the frequency ω is given by the frequency of the largest-magnitude DTFT coefficient of $\overline{\mathbf{x}}$: $\widehat{\omega} = \arg\max_\omega |\overline{\mathbf{X}}(\omega)| = \arg\max_\omega |\langle \mathbf{e}(\omega), \mathbf{x} \rangle|$ [11, 12]. This approach is often described as the *periodogram method* for parameter estimation [12]. This simple estimator can be extended to the multiple sinusoid setting by performing combinatorial hypothesis testing [12].

For frequency-sparse signals with components at integral frequencies, the signal's representation in the DFT basis provides the information needed by the MLEs above; in this case, the parameter estimation problem is equivalent to sparse approximation in the DFT basis. This equivalence can also be extended to frequency-sparse signals whose component frequencies are included in an oversampled DTFT sampling grid by using a DFT frame instead.

B. Window-based methods

From the spectral analysis point of view, we can argue that the coherence of the DFT frame $\Psi(c)$ is simply another manifestation of the spectral leakage problem. The classical way to combat spectral

leakage is to apply a tapered window function to the signal before computing the DFT [12, 13]. However, windowing can also hamper the spectral analysis resolution, making it more difficult to identify frequency-sparse signal components with similar frequencies.

An improved approach to spectral estimation proposed by Thomson [32] forms a weighted average of windowed DFTs using a set of windows known as prolate spheroidal wave functions (PSWF). The PSWF windows are orthogonal and optimally concentrated in frequency; hence, they optimize the resolution of the frequency analysis [32]. Thomson's method aims for a balance in the tradeoff between low spectral leakage and high spectral analysis resolution. The estimators of Section III-A can be adapted to employ the Thomson frequency analysis estimate instead of the signal's DTFT.

C. Eigenanalysis-based methods

A modern alternative to classical periodogram-based spectrum estimates are line spectrum estimation algorithms based on eigenanalysis of the signal's correlation matrix [12]. Such algorithms provide improved resolution of the parameters of a frequency-sparse signal by estimating the principal components of the signal's autocorrelation matrix in order to find the dominant signal modes in the frequency domain. Example algorithms include Pisarenko's method, multiple signal classification (MUSIC), and estimation of signal parameters via rotationally invariant techniques (ESPRIT). A line spectrum estimation algorithm $\mathbb{L}(\mathbf{x}, K)$ returns a set of dominant K frequencies for the input signal $\mathbf{x}$, with K being a controllable parameter.

As a concrete example, we describe the MUSIC algorithm, which estimates the parameters of a frequency-sparse signal embedded in noise. We revisit the model of Section III: $\mathbf{x} = \mathbf{s} + \mathbf{n}$, where $\mathbf{s}$ is now of the form (1) and $\mathbf{n} \sim \mathcal{N}(0, \sigma_n^2 \mathbf{I})$ denotes a noise vector. MUSIC operates on the autocorrelation matrix $\mathbf{R}_{\mathbf{xx}}$ of $\mathbf{x}$ of size $P \times P$; we obtain its eigendecomposition into the eigenvalues $\lambda_1, \ldots, \lambda_P$, sorted by decreasing magnitude, and the corresponding eigenvectors $\mathbf{v}_1, \ldots, \mathbf{v}_P$. The algorithm relies on a score function

$$P_{\mathsf{MUSIC}}(\omega) = \frac{1}{\sum_{p=K+1}^{P} |\mathbf{e}(\omega)^T \mathbf{v}_p|^2} \tag{6}$$

and returns the locations of the K largest score function peaks as the frequencies present in the signal. A modification known as Root MUSIC formulates a polynomial that depends on the noise subspace eigenvectors; the polynomial's zeros help determine the locations of the peaks of (6).

We can interpret the line spectrum estimation process $\mathbb{L}$ as a K-sparse approximation algorithm $\mathbb{T}'(\mathbf{x}, K)$ in the frequency domain: first, we obtain the K frequencies $\{\widehat{\omega}_k\}_{k=1}^{K} = \mathbb{L}(\mathbf{x}, K)$; second, we

estimate the values of the corresponding DTFT coefficients for the signal as shown in Section III-A. We note that most line spectrum estimation algorithms provide a tradeoff between estimation accuracy and computational complexity in the selection of the window size used to estimate the autocorrelation matrix $\mathbf{R_{xx}}$.

IV. Spectral Compressive Sensing

We are now in a position to develop new SCS recovery algorithms that are especially tailored to frequency-sparse signals of arbitrary frequencies. We will develop two sets of algorithms based on the periodogram and line spectrum estimation algorithms from Section III.

A. SCS via periodogram

To alleviate the performance-sapping coherence of the redundant DFT frame, we marry it with the model-based CS framework of Section II-D that forces the signal approximation to contain linear combinations of only incoherent frame elements. We assume initially that the components of the frequency-sparse signal $\mathbf{x}$ have frequencies in the oversampled grid of the frame $\Psi(c)$; we will then extend our analysis to signals with components at arbitrary frequencies at the end of the subsection.

1) Structured signal model: We begin by defining a structured signal model for frequency-sparse signals requiring that the components of the signal are incoherent with each other. Our structured signal model is defined as

$$\mathcal{T}_{K,c,\mu} = \left\{ \sum_{i=1}^{K} a_i \mathbf{e}(d_i \Delta) \text{ s. t. } d_i \in \{0, \ldots, cN-1\}, |\langle \mathbf{e}(d_i\Delta), \mathbf{e}(d_j\Delta)\rangle| \leq \mu,\ 1 \leq i,j \leq K \right\}, \quad (7)$$

where $\mu \in [0,1]$ is the maximal coherence allowed and $\Delta = 2\pi/cN$ as before. The union of subspaces contained in $\mathcal{T}_{K,c,\mu}$ corresponds to linear combinations of K incoherent elements from the DFT frame $\Psi(c)$, and $\mathcal{T}_{K,c,\mu} \subseteq \Sigma_K$. The coherence restriction in (7) imposes a lower limit to the frequency spacing between any two sinusoids present in a recoverable signal, which echoes the frequency resolution issue of spectrum estimators from Section III. Note that we require the signal model $\mathcal{T}_{K,\mathbf{c},\mu}$ to obey $\mu \leq \frac{1}{16(K-1)}$ — matching the requirements of sparse approximation algorithms mentioned in Section II-A — even though the actual coherence of the dictionary $\Psi(c)$ can be significantly larger [10].

2) Structured sparse approximation algorithm: Following the incoherent component model above, we modify a standard sparse approximation algorithm to avoid selecting highly coherent pairs of elements of the DFT frame $\Psi(c)$. Our structured sparse approximation algorithm is an adaptation of the refractory model-based algorithm of [33] and can be implemented as an integer program.

The algorithm $\mathbb{T}(\theta, K, c, \mu)$ takes as inputs the coefficient vector $\theta = \Psi(c)^T \mathbf{x} \in \mathbb{R}^{cN}$, the approximation sparsity K, and the maximum coherence allowed μ, and finds the best approximation for $\mathbf{x}$ under the model $\mathcal{T}_{K,c,\mu}$. The integer program implementation employs a constraint matrix $\mathbf{D}_\mu \in \mathbb{R}^{cN \times cN}$, which has binary entries that indicate whether each pair of elements from the DFT frame $\Psi(c)$ are coherent:

$$\mathbf{D}_\mu[i,j] := \begin{cases} 1 & \text{if } |\langle \mathbf{e}(i\Delta), \mathbf{e}(j\Delta) \rangle| \geq \mu, \\ 0 & \text{if } |\langle \mathbf{e}(i\Delta), \mathbf{e}(j\Delta) \rangle| < \mu. \end{cases}$$

The implementation also employs the cost vector $\mathbf{c}_\theta \in \mathbb{R}^{cN}$ defined as $\mathbf{c}_\theta(i) = \theta(i)^2,\ i = 0, \ldots, cN-1$. We proceed by solving the integer program

$$\mathbf{s}_{K,c,\mu} = \arg \min_{\mathbf{s} \in \{0,1\}^{cN}} \mathbf{c}_\theta^T \mathbf{s} \text{ such that } \mathbf{D}_\mu \mathbf{s} \leq \mathbf{1}, \mathbf{s}^T \mathbf{1} \leq K, \tag{8}$$

where $\mathbf{1}$ denotes a vector of ones of appropriate length, and then set the vector entry-wise product $\mathbb{T}(\theta, K, c, \mu) = \theta \cdot \mathbf{s}_{K,c,\mu}$ as the structured sparse approximation output.

When the matrix $\mathbf{D}_\mu$ is totally unimodular, the integer program (8) has the same solution as its noninteger relaxation, which is a linear program [34] whose complexity is cubic on the number of variables, $O(c^3 N^3)$. One example of totally unimodular matrices are *interval matrices*, which are binary matrices in which the ones appear consecutively in each row. While the matrix $\mathbf{D}_\mu$ we use in our case is not an interval matrix — as each row of $\mathbf{D}_\mu$ contains several intervals — it is possible to relax the integer program by using a modified matrix $\overline{\mathbf{D}_\mu}$. To obtain this new matrix, we decompose each row $\mathbf{s}_n$ of $\mathbf{D}_\mu$ into a set of rows $\mathbf{s}_{n,1},\ \mathbf{s}_{n,2}, \ldots$ that contain only one interval each and for which $\sum_i \mathbf{s}_{n,i} = \mathbf{s}_n$. The number of rows of $\overline{\mathbf{D}_\mu}$ is then at most $\frac{cN}{\pi\mu}$.

To reduce the computational complexity of this new structured sparse approximation algorithm, we propose a heuristic that relies on *frequency inhibition*. To obtain the model-based sparse approximation in the DFT frame $\theta = \Psi(c)^T \mathbf{x}$, we search for the coefficient $\theta(d)$ with the largest magnitude. Once a coefficient is selected, the algorithm inhibits all coefficients for coherent sinusoids (i.e., indices d' for which $|\langle \mathbf{e}(d\Delta), \mathbf{e}(d'\Delta) \rangle| > \mu$) by setting those coefficients to zero. This will include all coefficients for frequencies within $\kappa = 2\pi\, \mathrm{asinc}(1/\mu)/N$ radians/sample of the one selected. We then repeat the process by searching for the next largest coefficient in magnitude until K coefficients are selected or all coefficients are zero. This heuristic has complexity $O(cKN \log(cN))$ and offers very good performance for sparse approximation of arbitrary frequency-sparse signals, as shown in Fig. 2(b). Our experimental results in Section V employ this heuristic structured sparse approximation algorithm.

Algorithm 1 Spectral Iterative Hard Thresholding (SIHT) via periodogram

Inputs: CS Matrix Φ, DFT frame $\Psi := \Psi(c)$, measurements $\mathbf{y}$,

structured sparse approximation algorithm $\mathbb{T}(\cdot, K, c, \mu)$.

Outputs: K-sparse approximation $\widehat{\theta}$, signal estimate $\widehat{\mathbf{x}}$.

initialize: $\widehat{\theta}_0 = 0$, $\mathbf{r} = \mathbf{y}$, $i = 0$

while halting criterion false **do**

$i \leftarrow i + 1$

$\mathbf{b} \leftarrow \widehat{\theta}_{i-1} + \mathbb{T}(\Psi^T \Phi^T \mathbf{r}, N, c, \mu)$ {form signal estimate}

$\widehat{\theta}_i \leftarrow \mathbb{T}(\Psi^T \Psi \mathbf{b}, K, c, \mu)$ {prune signal estimate according to structured sparsity model}

$\mathbf{r} \leftarrow \mathbf{y} - \Phi\Psi\widehat{\theta}_i$ {update measurement residual}

end while

return $\widehat{\theta} \leftarrow \widehat{\theta}_i$, $\widehat{\mathbf{x}} \leftarrow \Psi\widehat{\theta}$

3) Recovery algorithm: The model-based IHT algorithm of [14] is particularly amenable to modification to incorporate our heuristic frequency-sparse approximation algorithm. Due to the redundancy of the frame $\Psi(c)$, we perform some minor surgery on the algorithm: we replace the matrix Φ by the matrix product $\Phi\Psi$ and multiply the signal estimate $\mathbf{b}$ by the Gramian matrix of the frame each time so that coherent frame elements featuring coefficients of opposing signs can cancel each other. The modified algorithm, which we dub *spectral iterative hard thresholding* (SIHT), is unfurled in Algorithm 1.

SIHT inherits a strong performance guarantee from standard IHT: If the matrix Φ has $\mathcal{T}_{2K,c,\mu}$-RIP with $\delta_{\mathcal{T}_{2K,c,\mu}} < \sqrt{2} - 1$, then we have

$$\|\mathbf{x} - \widehat{\mathbf{x}}\|_2 \leq C_0 \frac{\|\mathbf{x} - \Psi\mathbb{T}(\Psi^T\mathbf{x}, K, \mu)\|_2}{K^{1/2}} + C_1\epsilon; \tag{9}$$

where $\mathbb{T}(\Psi^T\mathbf{x}, K, c, \mu)$ denotes the structured sparse approximation of the input signal $\mathbf{x}$. We note that for signals that are frequency-sparse and composed of incoherent sinusoids with frequencies of the form $l\Delta$, we have $\Psi\mathbb{T}(\Psi^T\mathbf{x}, K, c, \mu) = \mathbf{x}$, meaning that recovery from noiseless measurements is exact.

4) Required number of CS measurements: To calculate the number of random CS measurements needed for signal recovery using Algorithm 1, we count the number of subspaces t_K that compose the signal model $\mathcal{T}_{K,c,\mu}$. We can obtain a loose measurement bound by counting the number of K-dimensional subspaces generated by subsets of the frame $\Psi(c)$ where no two vectors in a subspace have frequencies

closer than κ. From [33], we know the number of subspaces to be

$$t_K < \binom{cN - (K-1)(c\,\mathrm{asinc}(1/\mu) - 1)}{K};$$

this provides us with the following loose measurement bound for a random matrix to have the $\mathcal{T}_{2K,c,\mu}$-RIP:

$$M = \mathcal{O}\left(K \log\left(\frac{c(N - K\,\mathrm{asinc}(1/\mu))}{K}\right)\right). \tag{10}$$

The measurement bound (10) differs from that of CS with the orthonormal DFT basis only in the numerator inside the logarithm. While there is a reduction of up to cK in the numerator due to the $\mathrm{asinc}(1/\mu)$ term, this reduction is significantly smaller than the penalty of $(c-1)N$ due to frequency oversampling. If we ignore the small reduction, then the number of measurements needed corresponds to that required for a K-sparse, cN-dimensional signal; in other words, SIHT enjoys the benefits of a sparsifying coherent frame without the penalty on the number of measurements required for stable signal recovery. We will demonstrate below in Section V that, in practice, SCS offers significant reductions in the number of measurements needed for accurate recovery of frequency-sparse signals when compared against standard CS using both the orthonormal DFT basis and DFT frames.

5) Frequency interpolation: We now address the case where ω/Δ is non-integer, that is, where the frequency-sparse signal has components outside of the frequencies sampled by the DFT frame. We will modify the structured sparse approximation algorithm to include frequency and magnitude estimation steps. In this case, the approximation algorithm will find the frequency in the grid that is closest to each component's frequency. In other words, the standard structured sparse approximation of (8) of the signal estimate enables us to identify the grid frequencies closest to the frequencies of the signal components. It is then possible to estimate the component frequencies by performing a least squares fit: for each index selected by the structured sparse approximation, we fit the frame coefficients for a set of neighboring indices to a functional form for the sinc-shaped frequency response of a windowed sinusoid. For example, a quadratic fit works well for frequencies within the main lobe of the translated sinc function $\mathrm{sinc}(\omega - \overline{\omega})$ [32]. To improve the performance of the estimators for the frequency and amplitude of the sinusoids, we can use the multiple-window spectrum estimator of [32], which offers improved resolution on the frequency estimates and lower bias on the amplitude estimates.

B. SCS via line spectrum estimation

While the combination of a redundant frame and the coherence-inhibiting structured sparsity model yields an improvement in the performance of SIHT, the algorithm still suffers from a limitation in the

Algorithm 2 SIHT via Line Spectrum Estimation

Inputs: CS Matrix Φ, measurements $\mathbf{y}$, structured sparse approximation algorithm $\mathbb{T}'(\cdot, K)$.

Outputs: K-sparse approximation $\{\widehat{\omega}_k, \widehat{a}_k\}_{k=1}^K$, signal estimate $\widehat{\mathbf{x}}$.

initialize: $\widehat{\mathbf{x}}_0 = 0$, $\mathbf{r} = \mathbf{y}$, $i = 0$

while halting criterion false **do**

$i \leftarrow i + 1$

$\{\widehat{\omega}_k, \widehat{a}_k\}_{k=1}^K \leftarrow \mathbb{T}'(\widehat{\mathbf{x}}_{i-1} + \Phi^T(\mathbf{y} - \Phi\widehat{\mathbf{x}}_{i-1}), K)$ {obtain parameter estimates}

$\widehat{\mathbf{x}}_i \leftarrow \sum_{k=1}^K \widehat{a}_k \mathbf{e}(\widehat{\omega}_k)$ {form signal estimate}

end while

return $\widehat{\mathbf{x}} \leftarrow \widehat{\mathbf{x}}_i$, $\{\widehat{\omega}_k, \widehat{a}_k\}_{k=1}^K$

resolution of neighboring frequencies that it can distinguish. This limitation is inherited from the frequency and coefficient estimation methods used by SIHT, which are based on the periodogram.

Fortunately, we can leverage the line spectrum estimation methods described in Section III-C; recall that these methods return a set of dominant K frequencies for the input signal, with K being a controllable parameter. Since these methods do not rely on redundant frames, we do not need to leverage the features of SIHT that control the effect of coherence. We simply employ the structured sparse approximation algorithm $\mathbb{T}'(\mathbf{x}, K)$ from Section III-C in IHT, resulting in Algorithm 2. While analytical results for this new algorithm have proven difficult to establish, we show experimentally below that its performance matches or exceeds that of SIHT via periodogram while exhibiting a much simpler implementation.

V. Experimental Results

In this section, we report experimental results for the performance of the two SIHT recovery algorithms as compared to standard CS recovery using the IHT algorithm. We probe the robustness of the algorithms to varying amounts of measurement noise and varying frequency oversampling factors c. We also test the algorithms on a real-world communications signal. A Matlab toolbox containing implementations of the SCS recovery algorithms, together with scripts that generate all figures in this paper, is available at http://dsp.rice.edu/scs.

Our first experiment compares the performance of standard IHT using the orthonormal DFT basis against that of the SIHT algorithms (Algorithms 1 and 2 from Section IV). Our experiments use signals of length $N = 1024$ samples containing $K = 20$ complex-valued sinusoids. For varying M, we executed

100 independent trials using random measurement matrices Φ of size $M \times N$ with i.i.d. Gaussian entries and signals $\mathbf{x} = \sum_{k=1}^{K} \mathbf{e}(\omega_k)$, where each pair of frequencies $\omega_i,\ \omega_j,\ 1 \leq i, j \leq K,\ i \neq j$ are spaced by at least $10\pi/1024$ radians/sample. For each CS matrix / sparse signal pair we obtain the measurements $\mathbf{y} = \Phi\mathbf{x}$ and calculate estimates of the signal $\widehat{\mathbf{x}}$ using IHT with the orthonormal DFT basis, SIHT via periodogram with frequency oversampling factor $c = 10$ and maximum allowed coherence $\mu = 0.1$ (Algorithm 1), and SIHT via Root MUSIC (Algorithm 2). We use a window size $W = N/10$ in Root MUSIC to estimate the autocorrelation matrix $\mathbf{R_{xx}}$. We study the performance of these three algorithms in three different regimes: (*i*) the *average* case, in which the frequencies are selected randomly to machine precision; (*ii*) the *best* case, in which the frequencies are randomly selected and rounded to the closest integral frequency, resulting in zero spectral leakage; and (*iii*) the *worst* case, in which each frequency is half-way in between two consecutive integral frequencies, resulting in maximal spectral leakage. The results are summarized in Fig. 3 and show first that the average performance of standard IHT is very close to its worst-case performance, and second that both SIHT algorithms perform significantly better on the same signals. We also note that the SIHT algorithms work well in the average case even though the resulting signals do not match exactly match the sparse-in-DFT-frame assumption. Thus, the proposed algorithms are robust to mistmatch in the values for the frequencies in the signal model (1). We use this experimental setup in the rest of this section, but we restrict ourselves to the average case regime.

Our second experiment tests the robustness of the SIHT algorithms to additive noise in the measurements. We set the experiment parameters to $N = 1024$, $K = 20$ and $M = 300$, and we add i.i.d. Gaussian noise of variance σ to each measurements. For each value of σ, we perform 1000 independent noise trials; in each trial, we generate the matrices Φ and signals $\mathbf{x}$ randomly as in the previous experiment. Fig. 4 shows the average norm of the recovery error as a function of the noise variance σ; the linear relationship indicates stability to additive noise to the measurements, confirming the guarantee given in (9).

Our third experiment studies the impact of the frequency oversampling factor c on the performance of the SIHT algorithms. We use the same matrix and signal setup as in the previous experiment: we execute 10000 independent trials for each value of c. The results, shown in Fig. 5, indicate a linear dependence between the granularity of the DFT frame Δ and the norm of the recovery error. This sheds light on the tradeoff between the computational complexity and the performance of the recovery algorithm, as well as between the oversampling factor M/K (dependent on $\log c$) and the recovery performance. These results also experimentally confirm Lemma 1.

Our fourth experiment tests the capacity of standard CS and SCS recovery algorithms to resolve

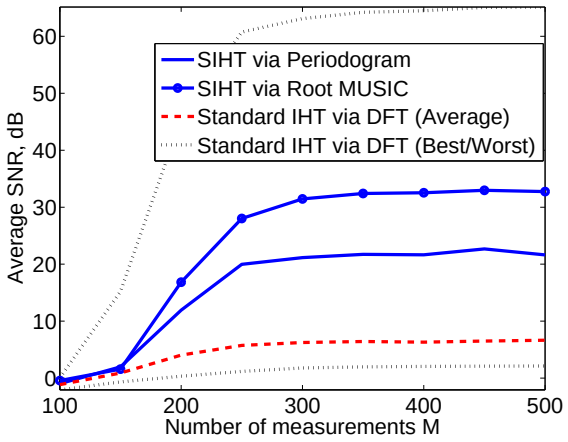

Fig. 3. *Performance of CS signal recovery via IHT using the orthonormal DFT basis, SIHT via periodogram, and SIHT via Root MUSIC. We use signals of length $N = 1024$ containing $K = 20$ complex-valued sinusoids. The dotted lines indicate the performance of IHT via the orthonormal DFT basis for the best case (when the frequencies of the sinusoids are integral) and the worst case (when each frequency is half way in between two consecutive integral frequencies). The performance of IHT for arbitrary frequencies is close to its worst-case performance, while both SIHT algorithms perform significantly better for arbitrary frequencies. All quantities are averaged over 100 independent trials.*

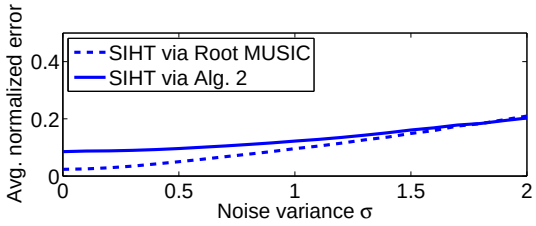

Fig. 4. *Performance of SIHT via periodogram and SIHT via Root MUSIC for SCS recovery from noisy measurements. We use signals of length $N = 1024$ containing $K = 20$ complex-valued sinusoids and take $M = 300$ measurements. We add noise of varying variances σ and calculate the average normalized error magnitude over 1000 independent trials. The linear relationship between the noise variance and the recovery error indicates the robustness of the recovery algorithm to noise.*

neighboring frequencies in frequency-sparse signals. For this experiment, the signal consists of 2 real-valued sinusoids (i.e., $K = 4$) of length $N = 1024$ with frequencies that are separated by a value δ_ω varying between $0.1 - 5$ cycles/sample ($2\pi/100N - 10\pi/N$ rad/sample); we obtain $M = 100$ measurements of the signal. We measure the performance of standard IHT via DFT, SIHT via periodogram

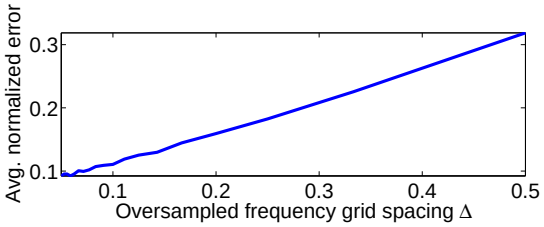

Fig. 5. *Performance of SIHT via periodogram for SCS recovery under varying grid spacing resolutions* $\Delta = 2\pi/cN$. *We use signals of length* $N = 1024$ *containing* $K = 20$ *complex-valued sinusoids and take* $M = 300$ *measurements. We average the recovery error over 10000 independent trials. There is a linear dependence between the granularity of the DFT frame and the norm of the recovery error.*

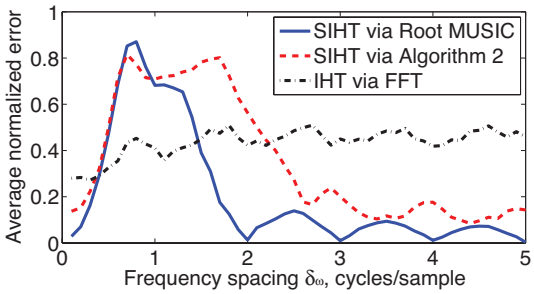

Fig. 6. *Performance of IHT and SIHT algorithms for frequency-sparse signals with components at closely spaced frequencies. We use signals of length* $N = 1024$ *containing 2 real sinusoids (*$K = 4$*) with frequencies separated by* δ_ω*, and measure the signal recovery performance of IHT and SIHT via periodogram and Root MUSIC from* $M = 100$ *measurements as a function of* δ_ω*. The results verify the limitations of periodogram-based methods and the markedly improved performance of line spectrum estimation methods used by the different versions of SIHT. Additionally, we see that standard IHT outperforms the SIHT algorithms only when the signal of interest is not contained in the class of frequency-sparse signals with incoherent components.*

with frequency oversampling factor $c = 10$ and maximum allowed coherence $\mu = 0.1$, and SIHT via Root MUSIC, all as a function of the frequency spacing δ_ω. For this experiment, we modify the window size parameter of the Root MUSIC algorithm to $W = N/3$ to improve its estimation accuracy at the cost of higher computational complexity. For each value of δ_ω, we execute 100 independent trials as detailed in previous experiments. The results, shown in Fig. 6, verify the limitation of periodogram-based methods as well as the improved resolution performance afforded by line spectrum estimation methods like Root MUSIC. Standard IHT only outperforms the SIHT algorithms when the signal does not belong in the class of frequency-sparse signals with incoherent components (that is, very small frequency spacing δ_ω).

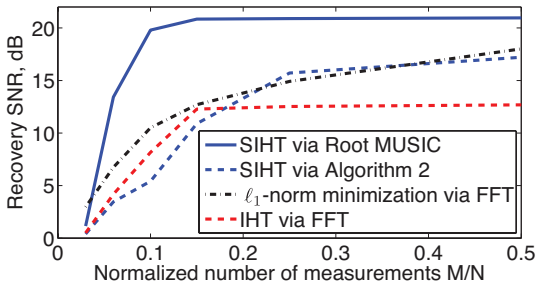

Fig. 7. *Performance of ℓ_1-norm minimization, IHT via orthonormal DFT, and SIHT via Root MUSIC algorithms on a real-world AM communications signal of length $N = 32768$ for a varying number of measurements M. SIHT significantly outperforms its standard counterparts.*

Our last experiment tests the performance of standard CS and SCS recovery algorithms on a real-world signal. We use the amplitude modulated (AM) signal from [7, Figure 7], measured in a lab, and simulate its acquisition using the random demodulator [7]. The signal has length $N = 32768$ samples. We recover the signal in half-overlapping blocks of length $N' = 1024$ using IHT and SIHT via Root MUSIC and compare their recovery performance as a function of the number of measurements M against that of the ℓ_1-norm minimization algorithm used in [7]. We set the target signal sparsity to $K = N/100$ in the IHT and SIHT algorithms. The recovered AM signals are then demodulated, and the recovery performance is measured in terms of the distortion against the message obtained by demodulating the signal at Nyquist rate. We average the performance over 20 trials for the random demodulator chipping sequence. The results in Fig. 7 shows that SIHT consistently outperforms its standard CS counterparts.

To summarize, our experiments have shown that SCS achieves significantly improved signal recovery performance for the overwhelming majority of frequency-sparse signals when compared with standard CS recovery algorithms and sparsity bases. The two SCS recovery algorithms we introduced and tested inherit some attractive properties from their standard counterparts, including robustness to model mismatch and measurement noise.

VI. Related work

The most relevant prior work in CS of frequency-sparse signals employs a block-sparsity model, where the DFT coefficients are grouped into local bins [35]. Such a model is motivated by the locality of the frequency spectra of real-world signals and is enforced via a two-stage recovery algorithm. The first stage aims to find the locations of the bins containing the nonzero DFT coefficients of the signal

using standard CS; the second stage then recovers the content in those bins by obtaining enough regular samples for each of the occupied bins. When the bins are correctly identified, the recovery of the signal content succeeds independently of the frequencies of the components in (1).

A recent paper [36] independently studied the poor performance of recovery of frequency-sparse signals when the DFT basis is used. The paper provides a generic framework for sparsity basis mismatch in which an inaccurate sparsity basis is used for CS recovery, and determines a bound for the approximation error as a function of the basis mismatch. The paper shows that in the noiseless setting, CS via the DFT basis provides lower accuracy that linear prediction methods on subsampled sinusoids. However, such linear prediction methods are very sensitive to noise, and thus are not suitable for use in the CS recovery approach provided in Section IV.

There also exists a related body of related work on compressive acquisition of signals governed by a small number of continuous-valued parameters. In finite rate of innovation (FROI) sampling [37], certain classes of signals governed by a small number of parameters can be acquired by uniformly sampling them using a specially designed kernel; the samples are then processed to obtain an annihilating filter, which in turn is used to estimate the values of the parameters. The application of FROI to frequency-sparse signals results in the linear prediction method used in [36], where the arguments of the complex roots of an annihilating filter reveal the frequencies ω_k of the signal components in (1). In fact, line spectral estimation algorithms have been proposed to extend FROI to noisy sampling settings [38, 39], albeit without performance guarantees.

Sparse approximation algorithms for frames characterized by continuously varying parameters have also been considered [40]. Here, one designs a frame composed of vectors corresponding to a sampling of the parameter space; this sampled frame can be used with a modified greedy algorithm to obtain an initial estimate of the parameter value, followed by a refinement via gradient descent. To date the analysis of such approaches has been limited to the convergence rate of the sparse approximation error $\|\mathbf{y} - \Phi\widehat{\mathbf{x}}\|_2$, which is not exactly relevant to CS applications where we seek low error in the sparse representation (i.e. $\|\mathbf{x} - \widehat{\mathbf{x}}\|_2$) instead of simple approximation fidelity.

VII. Conclusions

In this paper we have developed a new framework for CS recovery of frequency-sparse signals, which we have dubbed spectral compressive sensing (SCS). The framework uses a redundant frame of sinusoids corresponding to an oversampled frequency grid together with a coherence-inhibiting structured signal

model that prevents the usual loss of performance due to the frame coherence. We have provided both performance guarantees for SCS signal recovery and a bound on the number of random measurements needed to enable these guarantees. We have also presented adaptations of standard line spectrum estimation methods to achieve recovery of combinations of sinusoids with arbitrarily close frequencies while achieving low computational complexity. As Fig. 3 indicates, SCS recovery significantly outperform CS recovery based on the orthonormal DFT basis (up to 25dB in the figure).

Our SCS framework can be extended to other signal recovery settings where each component of the signal's sparse representation is governed by a small set of parameters. While such classes of signals are well suited for manifold models when only one component is present in the signal [22], they fall short for linear combinations of a varying number of components. Following SCS, we can construct a frame whose elements correspond to a uniformly obtained sampling of the parameter space. However, when the manifold model is very smooth, the resulting frame will be highly coherent, limiting the performance of sparse approximation algorithms. By posing coherence-inhibiting models such as that of (7), we can enable accurate recovery of sparse signals under the redundant frame. Similarly to [40], we can evolve from the parameter values corresponding to the chosen frame vectors to more accurate estimates of these parameters through the use of a least-squares fit of a functional form or through the use of gradient descent when the manifold has a tractable analytical form. Immediate applications of this formulation include sparsity-based localization [15–18], radar imaging [19–21], and sparse time-frequency representations [8].

Further work includes integrating our frequency inhibition approach into more powerful iterative [24] and ℓ_1-norm minimization recovery algorithms as well as better characterizing the performance of the line spectrum estimators used in the algorithms of Section IV. The performance of these algorithms (accuracy, robustness, and resolution) might be different when they are applied to signal estimates obtained from compressive measurements. We are also interested in extensions to other CS recovery algorithms to be used in conjunction with parameterized frame models. SCS can also be applied to signal ensembles; when a microphone or antenna array is used and the emitter is static, the dominant frequencies are the same for each of the sensors, following the common sparse supports joint sparsity model [41]. For mobile emitters, the changes in the frequency values can be modeled according to the Doppler effect, which increases the number of parameters for the signal ensemble observed from two (for emitter position) to four (for emitter position and velocity).

APPENDIX

Proof of Lemma 1. We start with a K-term approximation in the frame $\Psi(c)$:

$$\mathbf{x}' = \sum_{k=1}^{K} a'_k \mathbf{e}(\omega'_k),$$

where $a'_k = a_k \operatorname{sinc}\frac{(\omega_k - \omega'_k)N}{2\pi}$ and $\omega'_k = \Delta \operatorname{round}(\omega_k/\Delta)$. We then have

$$\begin{aligned}
\|\mathbf{x}_k - \mathbf{x}_k\|_2 &\leq \|\mathbf{x} - \mathbf{x}'\|_2 = \left\| \sum_{k=1}^{K} a_k \mathbf{e}(\omega_k) - \sum_{k=1}^{K} a'_k \mathbf{e}(\omega'_k) \right\|_2, \\
&= \left\| \sum_{k=1}^{K} \left(a_k \mathbf{e}(\omega_k) - a_k \operatorname{sinc}\left(\frac{(\omega_k - \omega'_k)N}{2\pi} \right) \mathbf{e}(\omega'_k) \right) \right\|_2, \\
&\leq \sum_{k=1}^{K} |a_k| \left\| \mathbf{e}(\omega_k) - \operatorname{sinc}\left(\frac{(\omega_k - \omega'_k)N}{2\pi} \right) \mathbf{e}(\omega'_k) \right\|_2 = \sum_{k=1}^{K} |a_k| \sqrt{1 - \operatorname{sinc}^2 \frac{(\omega_k - \omega'_k)N}{2\pi}}, \\
&\leq \sum_{k=1}^{K} |a_k| \sqrt{1 - \operatorname{sinc}^2 \frac{\Delta N}{4\pi}} = \sqrt{1 - \operatorname{sinc}^2 \frac{1}{2c}} \sum_{k=1}^{K} |a_k| = \sqrt{1 - \operatorname{sinc}^2 \frac{1}{2c}} \|\mathbf{a}\|_1,
\end{aligned}$$

proving the lemma. □

ACKNOWLEDGEMENTS

Thanks to Volkan Cevher, Yuejie Chi, Mark Davenport, Chinmay Hegde, and Cédric Vonesch for helpful discussions. This paper is dedicated to the memory of Dennis M. Healy; his insightful discussions with us inspired much of this project.

REFERENCES

[1] E. J. Candès, "Compressive sampling," in *Proc. International Congress of Mathematicians*, vol. 3, Madrid, Spain, 2006, pp. 1433–1452.

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

[3] R. G. Baraniuk, "Compressive sensing," *IEEE Signal Processing Mag.*, vol. 24, no. 4, pp. 118–120, 124, July 2007.

[4] A. C. Gilbert, S. Guha, P. Indyk, S. Muthukrishnan, and M. Strauss, "Near-optimal sparse Fourier representations via sampling," in *ACM Symposium on Theory of Computing (STOC)*. New York, NY, USA: ACM, 2002, pp. 152–161.

[5] A. C. Gilbert, S. Muthukrishnan, and M. J. Strauss, "Improved time bounds for near-optimal sparse Fourier representations," in *Proc. Wavelets XI at SPIE Optics and Photonics*, San Diego, CA, Aug. 2005.

[6] A. C. Gilbert, M. J. Strauss, and J. A. Tropp, "A tutorial on fast Fourier sampling," *IEEE Signal Proc. Mag.*, pp. 57–66, Mar. 2008.

[7] J. Tropp, J. N. Laska, M. F. Duarte, J. K. Romberg, and R. G. Baraniuk, "Beyond Nyquist: Efficient sampling of bandlimited signals," *IEEE Trans. Info. Theory*, vol. 56, no. 1, pp. 1–26, Jan. 2010.

[8] S. Mallat, *A wavelet tour of signal processing*. Academic Press, 1999.

[9] J. A. Tropp, "Greed is good: Algorithmic results for sparse approximation," *IEEE Trans. Inform. Theory*, vol. 50, no. 10, pp. 2231–2242, Oct. 2004.

[10] H. Rauhut, K. Schnass, and P. Vandergheynst, "Compressed sensing and redundant dictionaries," *IEEE Trans. Info. Theory*, vol. 54, no. 5, pp. 2210–2219, May 2008.

[11] S. M. Kay, *Fundamentals of statistical signal processing: Estimation theory*. Englewood Cliffs, NJ: Prentice-Hall, 1998.

[12] ——, *Modern spectral estimation: Theory and application*. Englewood Cliffs, NJ: Prentice Hall, 1988.

[13] P. Stoica and R. L. Moses, *Introduction to spectral analysis*. Upper Saddle River, NJ: Prentice Hall, 1997.

[14] R. G. Baraniuk, V. Cevher, M. F. Duarte, and C. Hegde, "Model-based compressive sensing," *IEEE Trans. Info. Theory*, vol. 56, no. 4, pp. 1982–2001, Apr. 2010.

[15] I. F. Gorodnitsky and B. D. Rao, "Sparse signal reconstruction from limited data using FOCUSS: A re-weighted minimum norm algorithm," *IEEE Trans. Signal Processing*, vol. 45, no. 3, pp. 600–616, 1997.

[16] D. Malioutov, M. Cetin, and A. S. Willsky, "A sparse signal reconstruction perspective for source localization with sensor arrays," *IEEE Trans. Signal Processing*, vol. 53, no. 8, pp. 3010–3022, 2005.

[17] V. Cevher, M. F. Duarte, and R. G. Baraniuk, "Localization via spatial sparsity," in *European Signal Processing Conference (EUSIPCO)*, Lausanne, Switzerland, 2008.

[18] V. Cevher, A. C. Gurbuz, J. H. McClellan, and R. Chellappa, "Compressive wireless arrays for bearing estimation," in *Proc. IEEE Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP)*, Las Vegas, NV, Apr. 2008, pp. 2497–2500.

[19] R. G. Baraniuk and P. Steeghs, "Compressive radar imaging," in *IEEE Radar Conf.*, Boston, MA, Apr. 2007, pp. 128–133.

[20] K. R. Varshney, M. Cetin, J. W. Fisher, and A. S. Willsky, "Sparse representation in structured dictionaries with application to synthetic aperture radar," *IEEE Trans. Signal Processing*, vol. 56, no. 8, pp. 3548–3561, Aug. 2008.

[21] M. Herman and T. Strohmer, "High resolution radar via compressive sensing," *IEEE Trans. Signal Processing*, vol. 57, no. 6, pp. 2275–2284, June 2009.

[22] R. G. Baraniuk and M. B. Wakin, "Random projections of smooth manifolds," *Foundations of Computational Mathematics*, vol. 9, no. 1, pp. 51–77, Feb. 2009.

[23] S. Chen, D. Donoho, and M. Saunders, "Atomic decomposition by basis pursuit," *SIAM J. on Sci. Comp.*, vol. 20, no. 1, pp. 33–61, 1998.

[24] D. Needell and J. Tropp, "CoSaMP: Iterative signal recovery from incomplete and inaccurate samples," *Applied and Computational Harmonic Analysis*, vol. 26, no. 3, pp. 301–321, May 2009.

[25] M. Figueiredo and R. Nowak, "An EM algorithm for wavelet-based image restoration," *IEEE Trans. Image Processing*, vol. 12, no. 8, pp. 906–916, 2003.

[26] I. Daubechies, M. Defrise, and C. De Mol, "An iterative thresholding algorithm for linear inverse problems with a sparsity constraint," *Comm. Pure Appl. Math.*, vol. 57, pp. 1413–1457, 2004.

[27] E. J. Candès and J. K. Romberg, "Signal recovery from random projections," in *Proc. Computational Imaging III at SPIE Electronic Imaging*, vol. 5674. San Jose, CA: SPIE, Jan. 2005, pp. 76–86.

[28] T. Blumensath and M. E. Davies, "Iterative hard thresholding for compressed sensing," *Applied and Computational Harmonic Analysis*, vol. 27, no. 3, pp. 265–274, Nov. 2009.

[29] M. A. Davenport and M. B. Wakin, "Analysis of orthogonal matching pursuit using the restricted isometry property," *IEEE Trans. Info. Theory*, 2010, to appear.

[30] T. Blumensath and M. E. Davies, "Sampling theorems for signals from the union of finite-dimensional linear subspaces," *IEEE Trans. Info. Theory*, vol. 55, no. 4, pp. 1872–1882, Apr. 2009.

[31] Y. M. Lu and M. N. Do, "Sampling signals from a union of subspaces," *IEEE Signal Processing Mag.*, vol. 25, no. 2, pp. 41–47, Mar. 2008.

[32] D. J. Thomson, "Spectrum estimation and harmonic analysis," *Proceedings of the IEEE*, vol. 70, no. 9, pp. 1055–1094, Sept. 1982.

[33] C. Hegde, M. F. Duarte, and V. Cevher, "Compressive sensing recovery of spike trains using a structured sparsity model," in *Workshop on Signal Processing with Adaptive Sparse Structured Representations (SPARS)*, Saint Malo, France, Apr. 2009.

[34] G. L. Nemhauser and L. A. Wolsey, *Integer and combinatorial optimization*. Wiley-Interscience, 1999.

[35] Y. Eldar and M. Mishali, "Blind multi-band signal reconstruction: Compressed sensing for analog signals," *IEEE Trans. Signal Processing*, vol. 57, no. 3, pp. 993–1009, Mar. 2009.

[36] Y. Chi, L. Scharf, A. Pezeshki, and R. Calderbank, "The sensitivity to basis mismatch of compressed sensing in spectrum analysis and beamforming," in *Workshop on Defense Applications of Signal Processing (DASP)*, Lihue, HI, Oct. 2009.

[37] M. Vetterli, P. Marziliano, and T. Blu, "Sampling signals with finite rate of innovation," *IEEE Trans. Signal Processing*, vol. 50, no. 6, pp. 1417–1428, June 2002.

[38] I. Maravić and M. Vetterli, "Sampling and reconstruction of signals with finite innovation in the presence of noise," *IEEE Trans. Signal Processing*, vol. 53, no. 8, pp. 2788–2805, Aug. 2005.

[39] A. Ridolfi, I. Maravić, J. Kusuma, and M. Vetterli, "Sampling signals with finite rate of innnovation: the noisy case," École Polytechnique Fédérale de Laussane, Lausanne, Switzerland, Tech. Rep. LCAV-REPORT-2002-001, Dec. 2002.

[40] L. Jacques and C. De Vleeschouwer, "A geometrical study of matching pursuit parametrization," *IEEE Trans. Signal Processing*, vol. 56, no. 7, pp. 2835–2848, July 2008.

[41] D. Baron, M. F. Duarte, M. B. Wakin, S. Sarvotham, and R. G. Baraniuk, "Distributed compressive sensing," Rice University, Houston, TX, ECE Department Tech. Report TREE-0612, Nov. 2006.