

Beyond Nyquist: Efficient Sampling of Sparse Bandlimited Signals

Joel A. Tropp, *Member, IEEE*, Jason N. Laska, *Student Member, IEEE*, Marco F. Duarte, *Student Member, IEEE*,
Justin K. Romberg, *Member, IEEE*, and Richard G. Baraniuk, *Fellow, IEEE*

Abstract—Wideband analog signals push contemporary analog-to-digital conversion systems to their performance limits. In many applications, however, sampling at the Nyquist rate is inefficient, because the signals of interest contain only a small number of significant frequencies relative to the bandlimit, although the locations of the frequencies may not be known a priori. For this type of sparse signal, other sampling strategies are possible. This paper describes a new type of data acquisition system, called a *random demodulator*, that is constructed from robust, readily available components. Let K denote the total number of frequencies in the signal, and let W denote its bandlimit in Hz. Simulations show that the random demodulator requires just $O(K \log(W/K))$ samples per second to stably reconstruct the signal. This sampling rate is *exponentially lower* than the Nyquist rate of W Hz. In contrast with Nyquist sampling, one must use nonlinear methods, such as convex programming, to recover the signal from the samples taken by the random demodulator. Finally, the paper provides a theoretical analysis of the system's performance.

Index Terms—analog-to-digital conversion, compressive sampling, sampling theory, signal recovery, sparse approximation

I. INTRODUCTION

THE Shannon sampling theorem is one of the foundations of modern signal processing. For a continuous-time signal f whose highest frequency is less than $W/2$ Hz, the theorem suggests that we sample the signal uniformly at a rate of W Hz. The values of the signal at intermediate points in time are determined completely by the *cardinal series*

$$f(t) = \sum_{n \in \mathbb{Z}} f\left(\frac{n}{W}\right) \text{sinc}(Wt - n). \quad (1)$$

In practice, one typically samples the signal at a somewhat higher rate and reconstructs with a kernel that decays faster than the sinc function [1, Ch. 4].

This well-known approach becomes impractical when the bandlimit W is large because it is challenging to build sampling hardware that operates at a sufficient rate. The demands of many modern applications already exceed the capabilities of current technology. Even though recent developments in analog-to-digital converter (ADC) technologies have increased their sampling speeds, state-of-the-art architectures are not yet adequate for emerging applications, such as ultrawideband and radar systems because of the additional requirements on power

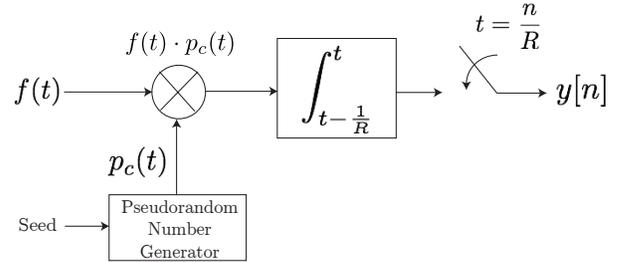


Fig. 1. Block diagram for the random demodulator.

consumption [2]. The time has come to explore alternative techniques [3].

A. The Random Demodulator

In the absence of extra information, Nyquist-rate sampling is essentially optimal for bandlimited signals [4]. Therefore, we must identify other properties that can provide additional leverage. Fortunately, in many applications, signals are also *sparse*. That is, the number of significant frequency components is often much smaller than the bandlimit allows. We can exploit this fact to design new kinds of sampling hardware.

This paper studies the performance of a new type of sampling system—called a *random demodulator*—for acquiring sparse, bandlimited signals. Figure 1 displays a block diagram for the system, and Figure 2 describes the intuition behind the design. In summary, we demodulate the signal by multiplying it with a high-rate pseudonoise sequence, which smears the tones across the entire spectrum. Then we apply a lowpass anti-aliasing filter, and we capture the signal by sampling it at a relatively low rate. As illustrated in Figure 3, the demodulation process ensures that each tone has a distinct signature within the passband of the filter. Since there are few tones present, it is possible to identify the tones and their amplitudes from the low-rate samples.

The major advantage of the random demodulator is that it bypasses the need for a high-rate ADC. Demodulation is typically much easier to implement than sampling, yet it allows us to use a low-rate ADC. As a result, the system can be constructed from robust, readily available components even while it can acquire higher-bandlimit signals than traditional sampling hardware.

We do pay a price for the slower sampling rate: It is no longer possible to express the original signal as a linear function of the samples, *à la* the cardinal series (1). Rather,

Submitted: 30 January 2009.

JAT was supported by ONR N00014-08-1-0883, DARPA/ONR N66001-06-1-2011 and N66001-08-1-2065, and NSF DMS-0503299. JNL, MFD, and RGB were supported by DARPA/ONR N66001-06-1-2011 and N66001-08-1-2065, ONR N00014-07-1-0936, AFOSR FA9550-04-1-0148, NSF CCF-0431150, and the Texas Instruments Leadership University Program. JKR was supported by NSF CCF-515632.

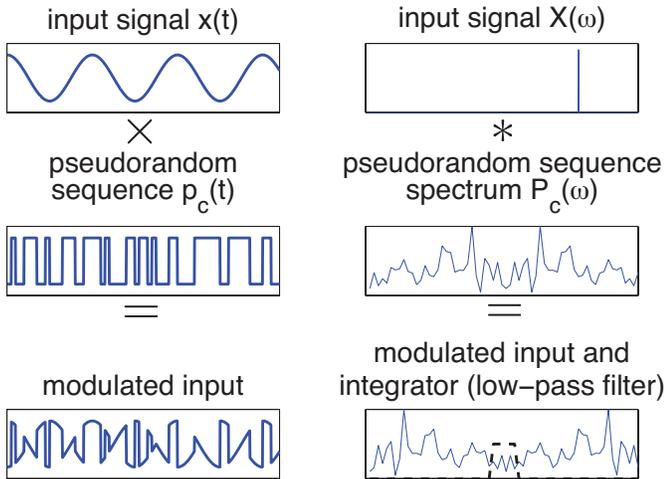


Fig. 2. The demodulation process multiplies the continuous-time input signal by a random square wave. The action of the system on a single tone is illustrated in the time domain (left) and the frequency domain (right). The dashed line indicates the frequency response of the lowpass filter. See Figure 3 for an enlargement of the filter’s passband.

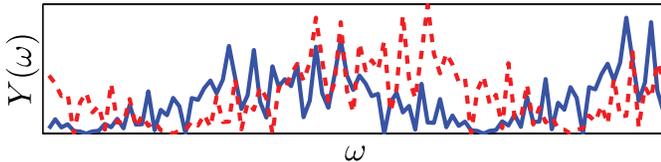


Fig. 3. The random demodulator furnishes each frequency with a unique signature that can be discerned by examining the passband of the antialiasing filter. This image enlarges the pass region of the demodulator’s output for two input tones (solid and dashed). The two signatures are nearly orthogonal when their phases are taken into account.

$f(t)$ is encoded into the measurements in a more subtle manner; the reconstruction process is highly nonlinear, and must carefully take advantage of the fact that the signal is sparse. As a result, signal recovery becomes more computationally intensive. In short, the random demodulator uses additional digital processing to reduce the burden on the analog hardware. This tradeoff seems acceptable, as advances in digital computing have outpaced those in analog-to-digital conversion.

B. Results

Our simulations provide striking evidence that the random demodulator performs. Consider a periodic signal with a bandlimit of $W/2$ Hz, and suppose that it contains K tones with random frequencies and phases. Our experiments below show that, with high probability, the system acquires enough information to reconstruct the signal after sampling at just $O(K \log(W/K))$ Hz. In words, the sampling rate is proportional to the number K of tones and the logarithm of the bandlimit W . In contrast, the usual approach requires sampling at W Hz, regardless of K . In other words, the random demodulator operates at an *exponentially* slower sampling rate! We also demonstrate that the system is effective for reconstructing simple communication signals.

Our theoretical work supports these empirical conclusions, but it results in slightly weaker bounds on the sampling rate. We have been able to prove that a sampling rate of $O(K \log W + \log^3 W)$ suffices for high-probability recovery of the random signals we studied experimentally. This analysis also suggests that there is a small startup cost when the number of tones is small, but we did not observe this phenomenon in our experiments. It remains an open problem to explain the computational results in complete detail.

The random signal model arises naturally in numerical experiments, but it does not provide an adequate description of real signals, whose frequencies and phases are typically far from random. To address this concern, we have proved that the random demodulator can acquire *all* K -tone signals—regardless of the frequencies, amplitudes, and phases—when the sampling rate is $O(K \log^6 W)$. In fact, the system does not even require the spectrum of the input signal to be sparse; the system can successfully recover any signal whose spectrum is well-approximated by K tones. Moreover, our analysis shows that the random demodulator is robust against noise and quantization errors.

This work focuses on input signals drawn from a specific mathematical model, framed in Section II. Many real signals have sparse spectral occupancy, even though they do not meet all of our formal assumptions. We propose a device, based on the classical idea of *windowing*, that allows us to approximate general signals by signals drawn from our model. Therefore, our recovery results for the idealized signal class extend to signals that we are likely to encounter in practice.

In summary, we believe that these empirical and theoretical results, taken together, provide compelling evidence that the demodulator system is a powerful alternative to Nyquist-rate sampling for sparse signals.

C. Outline

In Section II, we present a mathematical model for the class of sparse, bandlimited signals. Section III describes the intuition and architecture of the random demodulator, and it addresses the nonidealities that may affect its performance. In Section IV, we model the action of the random demodulator as a matrix. Section V describes computational algorithms for reconstructing frequency-sparse signals from the coded samples provided by the demodulator. We continue with an empirical study of the system in Section VI, and we offer some theoretical results in Section VII that partially explain the system’s performance. Section VIII discusses a windowing technique that allows the demodulator to capture nonperiodic signals. We conclude with a discussion of related work in Section IX. Finally, Appendices I, II and III contain proofs of our signal reconstruction theorems.

II. THE SIGNAL MODEL

Our analysis focuses on a class of discrete, multitone signals that have three distinguished properties:

- **Bandlimited.** The maximum frequency is bounded.
- **Periodic.** Each tone has an integral frequency in Hz.

- **Sparse.** The number of active tones is small in comparison with the bandlimit.

Our work shows that the random demodulator can recover these signals very efficiently. Indeed, the number of samples per unit time scales directly with the sparsity, but it increases only logarithmically in the bandlimit.

At first, these discrete multitone signals may appear to be simpler than the signals that arise in most applications. For example, we often encounter signals that contain non-harmonic tones or signals that contain continuous bands of active frequencies rather than discrete tones. Nevertheless, these broader signal classes can be approximated within our model by means of *windowing* techniques. We address this point in Section VIII.

A. Mathematical Model

Consider the following mathematical model for the class of discrete multitone signals. Let $W/2$ be a positive integer that exceeds the highest frequency present in the continuous-time signal f . Consider a set Ω of integer-valued frequencies

$$\Omega \subset \{0, \pm 1, \pm 2, \dots, \pm(W/2 - 1), W/2\}$$

and set of complex-valued amplitudes

$$\{a_\omega : \omega \in \Omega\}.$$

The model contains each signal of the form

$$f(t) = \sum_{\omega \in \Omega} a_\omega e^{-2\pi i \omega t} \quad \text{for } t \in [0, 1). \quad (2)$$

Typically, we write $K = |\Omega|$ for the number of active tones, and we focus on the case where K is much smaller than the bandlimit W .

To summarize, the signals of interest are *bandlimited* because they contain no frequencies above $W/2$ cycles per second; *periodic* because the frequencies are integral; and *sparse* because the number of tones $K \ll W$.

Let us emphasize several conceptual points about this model:

- We have normalized the time interval to one second for simplicity. Of course, it is possible to consider signals at another time resolution.
- We have also normalized frequencies. To consider signals whose frequencies are drawn from a set equally spaced by Δ , we would change the effective bandlimit to W/Δ .
- The model also applies to signals that are sparse and bandlimited in a single time interval. It can be extended to signals where the model (2) holds with a different set of frequencies and amplitudes in each time interval $[0, 1), [1, 2), [2, 3), \dots$. These signals are sparse not in the Fourier domain but rather in the short-time Fourier domain [5, Ch. IV].

B. Information Content of Signals

According to the sampling theorem, we can identify signals from the model (2) by sampling for one second at W Hz. Yet these signals contain only $R = O(K \log(W/K))$ bits of

information. In consequence, it is reasonable to expect that we can acquire these signals using only R digital samples.

Here is one way to establish the information bound. Stirling's approximation shows that there are about $\exp\{K \log(W/K) + O(K)\}$ ways to select K distinct integers in the range $\{1, 2, \dots, W\}$. Therefore, it takes $O(K \log(W/K))$ bits to encode the frequencies present in the signal. Each of the K amplitudes can be approximated with a fixed number of bits, so the cost of storing the frequencies dominates.

C. Examples

Many applications feature wideband signals that are sparse or locally sparse in frequency. Here are some basic examples:

- **Communications signals**, such as transmissions with a frequency hopping modulation scheme that switches a sinusoidal carrier among many frequency channels according to a predefined (often pseudorandom) sequence. Other examples include transmissions with narrowband modulation where the carrier frequency is unknown but could lie anywhere in a wide bandwidth.
- **Acoustic signals**, such as musical signals where each note consists of a dominant sinusoid with a progression of several harmonic overtones.
- **Slowly varying chirps**, as used in radar and geophysics, that slowly increase or decrease the frequency of a sinusoid over time.
- **Smooth signals** that require only a few Fourier coefficients to represent.
- **Piecewise smooth signals** that are differentiable except for a small number of step discontinuities.

III. THE RANDOM DEMODULATOR

This section describes the random demodulator system that we propose for signal acquisition. We first discuss the intuition behind the system and its design. Then we address some implementation issues and nonidealities that impact its performance.

A. Intuition

The random demodulator performs three basic actions: demodulation, lowpass filtering, and low-rate sampling. Refer back to Figure 1 for the block diagram. In this section, we offer a short explanation of why this approach allows us to acquire sparse signals.

Consider the problem of acquiring a single high-frequency tone that lies within a wide spectral band. Obviously, a low-rate sampler with an antialiasing filter is oblivious to any tone whose frequency exceeds the passband of the filter. The random demodulator deals with the problem by smearing the tone across the entire spectrum so that it leaves a signature that can be detected by a low-rate sampler.

More precisely, the random demodulator forms a (periodic) square wave that randomly alternates at or above the Nyquist rate. This random signal is a sort of periodic approximation to white noise. When we multiply a pure tone by this random

square wave, we simply translate the spectrum of the noise, as documented in Figure 2. The key point is that translates of the noise spectrum look completely different from each other, even when restricted to a narrow frequency band, as illustrated in Figure 3.

Now consider what happens when we multiply a frequency-sparse signal by the random square wave. In the frequency domain, we obtain a superposition of translates of the noise spectrum, one translate for each tone. Since the translates are so distinct, each tone has its own signature. The original signal contains few tones, so we can disentangle them by examining a small slice of the spectrum of the demodulated signal.

To that end, we perform lowpass filtering to prevent aliasing, and we sample with a low-rate ADC. This process results in coded samples that contain a complete representation of the original sparse signal. We discuss methods for decoding the samples in Section V.

B. System Design

Let us present a more formal description of the random demodulator shown in Figure 1. The first two components implement the demodulation process. The first piece is a random number generator, which produces a discrete-time sequence $\varepsilon_0, \varepsilon_1, \varepsilon_2 \dots$ of numbers that take values ± 1 with equal probability. We refer to this as the *chipping sequence*. The chipping sequence is used to create a continuous-time demodulation signal $p_c(t)$ via the formula

$$p_c(t) = \varepsilon_n, \quad t \in \left[\frac{n}{W}, \frac{n+1}{W} \right) \text{ and } n = 0, 1, \dots, W-1.$$

In words, the demodulation signal switches between the levels ± 1 randomly at the Nyquist rate of W Hz. Next, the mixer multiplies the continuous-time input $f(t)$ by the demodulation signal $p_c(t)$ to obtain a continuous-time demodulated signal

$$y(t) = f(t) \cdot p_c(t), \quad t \in [0, 1).$$

Together these two steps smear the frequency spectrum of the original signal via the convolution

$$Y(\omega) = (F * P_c)(\omega).$$

See Figure 2 for a visual.

The next two components behave the same way as a standard ADC, which performs lowpass filtering to prevent aliasing and then samples the signal. Here, the lowpass filter is simply an accumulator that sums the demodulated signal $y(t)$ for $1/R$ seconds. The filtered signal is sampled instantaneously every $1/R$ seconds to obtain a sequence $\{y_m\}$ of measurements. After each sample is taken, the accumulator is reset. In summary,

$$y_m = R \int_{m/R}^{(m+1)/R} y(t) dt, \quad m = 0, 1, \dots, R-1.$$

This approach is called *integrate-and-dump* sampling. Finally, the samples are quantized to a finite precision. (In this work, we do not model the final quantization step.)

The fundamental point here is that the rate R is much lower than the Nyquist rate W . We will see that R depends primarily on the number K of significant frequencies that participate in the signal.

C. Implementation and Nonidealities

Any reasonable system for acquiring continuous-time signals must be implementable in analog hardware. The system that we propose is built from robust, readily available components. This subsection briefly discusses some of the engineering issues.

In practice, we generate the chipping sequence with a pseudorandom number generator. It is preferable to use pseudorandom numbers for several reasons: they are easier to generate; they are easier to store; and their structure can be exploited by digital algorithms. Many types of pseudorandom generators can be fashioned from basic hardware components. For example, the Mersenne twister [6] can be implemented with shift registers. The performance of the random demodulator is unlikely to suffer from the fact that the chipping sequence is not completely random. In fact, for some applications it may suffice to fix a chipping sequence in advance.

The mixer must operate at the Nyquist rate W . Nevertheless, the chipping sequence alternates between the levels ± 1 , so the mixer only needs to reverse the polarity of the signal. It is relatively easy to perform this step using inverters and multiplexers. Most conventional mixers trade off speed for linearity, i.e., fast transitions may result in incorrect products. Since the random demodulator only needs to reverse polarity of the signal, nonlinearity is not the primary nonideality. Instead, the bottleneck for the speed of the mixer is the settling times of inverters and multiplexers, which determine the length of time it takes for the output of the mixer to reach steady state.

The sampler can be implemented with an off-the-shelf ADC. It suffers the same types of nonidealities as any ADC, including thermal noise, aperture jitter, comparator ambiguity, and so forth [7]. Since the random demodulator operates at a relatively low sampling rate, we can use high-quality ADCs, which exhibit fewer problems.

In practice, a high-fidelity integrator is not required. It suffices to perform lowpass filtering before the samples are taken. It is essential, however, that the impulse response of this filter can be characterized very accurately.

The net effect of these nonidealities is much like the addition of noise to the signal. The signal reconstruction process is very robust, so it performs well even in the presence of noise.

IV. RANDOM DEMODULATION IN MATRIX FORM

In the ideal case, the random demodulator is a linear system that maps a continuous-time signal to a discrete sequence of samples. To understand its performance, we prefer to express the system in matrix form. We can then study its properties using tools from matrix analysis and functional analysis.

A. Discrete-Time Representation of Signals

The first step is to find an appropriate discrete representation for the space of continuous-time input signals. To that end, note that each $(1/W)$ -second block of the signal is multiplied by a random sign. Then these blocks are aggregated, summed, and sampled. Therefore, part of the time-averaging performed by the accumulator commutes with the demodulation process.

In other words, we can average the input signal over blocks of duration $1/W$ without affecting subsequent steps.

Fix a time instant of the form $t_n = n/W$ for an integer n . Let x_n denote the average value of the signal f over a time interval of length $1/W$ starting at t_n . Thus,

$$\begin{aligned} x_n &= \int_{t_n}^{t_n+1/W} f(t) dt \\ &= \sum_{\omega \in \Omega} a_\omega \left[\frac{e^{-2\pi i \omega / W} - 1}{2\pi i \omega} \right] e^{-2\pi i \omega t_n} \end{aligned} \quad (3)$$

with the convention that, for the frequency $\omega = 0$, the bracketed term equals $1/W$. Since $|\omega| \leq W/2$, the bracket never equals zero. Absorbing the brackets into the amplitude coefficients, we obtain a discrete-time representation x_n of the signal $f(t)$:

$$x_n = \sum_{\omega \in \Omega} s_\omega e^{-2\pi i n \omega / W} \quad \text{for } n = 0, 1, \dots, W-1$$

where

$$s_\omega = a_\omega \left[\frac{e^{-2\pi i \omega / W} - 1}{2\pi i \omega} \right].$$

In particular, a continuous-time signal that involves only the frequencies in Ω can be viewed as a discrete-time signal comprised of the same frequencies. We refer to the complex vector \mathbf{s} as an *amplitude vector*, with the understanding that it contains the phase information as well.

The nonzero components of the length- W vector \mathbf{s} are listed in the set Ω . We may now express the discrete-time signal \mathbf{x} as a matrix-vector product. Define the $W \times W$ matrix

$$\begin{aligned} \mathbf{F} &= \frac{1}{\sqrt{W}} [e^{-2\pi i n \omega / W}]_{n,\omega} \quad \text{where} \\ & \quad n = 0, 1, \dots, W-1 \text{ and} \\ & \quad \omega = 0, \pm 1, \dots, \pm \left(\frac{W}{2} - 1 \right), \frac{W}{2}. \end{aligned}$$

The matrix \mathbf{F} is simply a permuted discrete Fourier transform (DFT) matrix. In particular, \mathbf{F} is unitary and its entries share the magnitude $W^{-1/2}$.

In summary, we can work with a discrete representation

$$\mathbf{x} = \mathbf{F}\mathbf{s}$$

of the input signal.

B. Action of the Demodulator

We view the random demodulator as a linear system acting on the discrete form \mathbf{x} of the continuous-time signal f .

First, we consider the effect of random demodulation on the discrete-time signal. Let $\varepsilon_0, \varepsilon_1, \dots, \varepsilon_{W-1}$ be the chipping sequence. The demodulation step multiplies each x_n , which is the average of f on the n th time interval, by the random sign ε_n . Therefore, demodulation corresponds to the map $\mathbf{x} \mapsto \mathbf{D}\mathbf{x}$ where

$$\mathbf{D} = \begin{bmatrix} \varepsilon_0 & & & & \\ & \varepsilon_1 & & & \\ & & \ddots & & \\ & & & \ddots & \\ & & & & \varepsilon_{W-1} \end{bmatrix}$$

is a $W \times W$ diagonal matrix.

Next, we consider the action of the accumulate-and-dump sampler. Suppose that the sampling rate is R , and assume that R divides W . Then each sample is the sum of W/R consecutive entries of the demodulated signal. Therefore, the action of the sampler can be treated as an $R \times W$ matrix \mathbf{H} whose r th row has W/R consecutive unit entries starting in column $rW/R + 1$ for each $r = 0, 1, \dots, R-1$. An example with $R = 3$ and $W = 12$ is

$$\mathbf{H} = \begin{bmatrix} 1 & 1 & 1 & 1 & & & & & & & & \\ & & & & 1 & 1 & 1 & 1 & & & & \\ & & & & & & & & & 1 & 1 & 1 & 1 \end{bmatrix}.$$

When R does not divide W , two samples may share contributions from a single element of the chipping sequence. We choose to address this situation by allowing the matrix \mathbf{H} to have fractional elements in some of its columns. An example with $R = 3$ and $W = 7$ is

$$\mathbf{H} = \begin{bmatrix} 1 & 1 & \sqrt{1/3} & & & & \\ & & \sqrt{2/3} & 1 & \sqrt{2/3} & & \\ & & & & \sqrt{1/3} & 1 & 1 \end{bmatrix}.$$

We have found that this device provides an adequate approximation to the true action of the system. For some applications, it may be necessary to exercise more care.

In summary, the matrix $\mathbf{M} = \mathbf{H}\mathbf{D}$ describes the action of the hardware system on the discrete signal \mathbf{x} . Each row of the matrix yields a separate sample of the input signal.

The matrix $\mathbf{\Phi} = \mathbf{M}\mathbf{F}$ describes the overall action of the system on the vector \mathbf{s} of amplitudes. This matrix $\mathbf{\Phi}$ has a special place in our analysis, and we refer to it as a *random demodulator matrix*.

C. Prewhitening

It is important to note that the bracket in (3) leads to a nonlinear attenuation of the amplitude coefficients. In a hardware implementation of the random demodulator, it may be advisable to apply a prewhitening filter to preserve the magnitudes of the amplitude coefficients. On the other hand, prewhitening amplifies noise in the high frequencies. We leave this issue for future work.

V. SIGNAL RECOVERY ALGORITHMS

The Shannon sampling theorem provides a simple linear method for reconstructing a bandlimited signal from its time samples. In contrast, there is no linear process for reconstructing the input signal from the output of the random demodulator because we must incorporate the highly nonlinear sparsity constraint into the reconstruction process.

Suppose that \mathbf{s} is a sparse amplitude vector, and let $\mathbf{y} = \mathbf{\Phi}\mathbf{s}$ be the vector of samples acquired by the random demodulator. Conceptually, the way to recover \mathbf{s} is to solve the mathematical program

$$\hat{\mathbf{s}} = \arg \min \|\mathbf{v}\|_0 \quad \text{subject to } \mathbf{\Phi}\mathbf{v} = \mathbf{y} \quad (4)$$

where the ℓ_0 function $\|\cdot\|_0$ counts the number of nonzero entries in a vector. In words, we seek the sparsest amplitude

vector that generates the samples we have observed. The presence of the ℓ_0 function gives the problem a combinatorial character, which may make it computationally difficult to solve completely.

Instead, we resort to one of the signal recovery algorithms from the sparse approximation or compressive sampling literature. These techniques fall in two rough classes: convex relaxation and greedy pursuit. We describe the advantages and disadvantages of each approach in the sequel.

This work concentrates on convex relaxation methods because they are more amenable to theoretical analysis. Our purpose here is not to advocate a specific algorithm but rather to argue that the random demodulator has genuine potential as a method for acquiring signals that are spectrally sparse. Additional research on algorithms will be necessary to make the technology viable.

A. Convex Relaxation

The problem (4) is difficult because of the unfavorable properties of the ℓ_0 function. A fundamental method for dealing with this challenge is to relax the ℓ_0 function to the ℓ_1 norm, which may be viewed as the convex function “closest” to ℓ_0 . Since the ℓ_1 norm is convex, it can be minimized subject to convex constraints in polynomial time [8].

Let \mathbf{s} be the unknown amplitude vector, and let $\mathbf{y} = \Phi\mathbf{s}$ be the vector of samples acquired by the random demodulator. We attempt to identify the amplitude vector \mathbf{s} by solving the convex optimization problem

$$\hat{\mathbf{s}} = \arg \min \|\mathbf{v}\|_1 \quad \text{subject to} \quad \Phi\mathbf{v} = \mathbf{y}. \quad (5)$$

In words, we search for an amplitude vector that yields the same samples and has the least ℓ_1 norm. On account of the geometry of the ℓ_1 ball, this method promotes sparsity in the estimate $\hat{\mathbf{s}}$.

The problem (5) can be recast as a second-order cone program. In our work, we use an old method, Iteratively Reweighted Least Squares (IRLS), for performing the optimization [9, 173ff]. It is known that IRLS converges linearly for certain signal recovery problems [10]. It is also possible to use interior-point methods, as proposed by Candès et al. [11].

Convex programming methods for sparse signal recovery problems are very powerful. Second-order methods, in particular, seem capable of achieving very good reconstructions of signals with wide dynamic range. On the other hand, optimization algorithms may be computationally burdensome.

B. Greedy Pursuits

To produce sparse approximate solutions to linear systems, we can also use another class of methods based on greedy pursuit. Roughly, these algorithms build up a sparse solution one step at a time by adding new components that yield the greatest immediate improvement in the approximation error. An early paper of Gilbert and Tropp analyzed the performance of an algorithm called Orthogonal Matching Pursuit for simple compressive sampling problems [12]. More recently, work of Needell and Vershynin [13], [14] and work of Needell and Tropp [15] has resulted in greedy-type algorithms whose

theoretical performance guarantees are analogous with those of convex relaxation methods.

Greedy pursuits tend to be less effective than convex optimization methods. On the other hand, they have a favorable computational profile which makes them attractive for large-scale problems.

C. Impact of Noise

In applications, signals are more likely to be *compressible* than to be sparse. (Compressible signals are not sparse but can be approximated by sparse signals.) Nonidealities in the hardware system lead to noise in the measurements. Furthermore, the samples acquired by the random demodulator are quantized. Modern convex relaxation methods and greedy pursuit methods are robust against all these departures from the ideal model. We discuss the impact of noise on signal reconstruction algorithms in Section VII-B.

In fact, the major issue with all compressive sampling methods is not the presence of noise *per se*. Rather, the process of compressing the signal’s information into a small number of samples decreases the signal-to-noise ratio. In the design of compressive sampling systems, it is essential to address this issue.

VI. EMPIRICAL RESULTS

We continue with an empirical study of the minimum measurement rate required to accurately reconstruct signals with the random demodulator. Our results are phrased in terms of the Nyquist rate W , the sampling rate R , and the sparsity level K . It is also common to combine these parameters into scale-free quantities. The *compression factor* R/W measures the improvement in the sampling rate over the Nyquist rate, and the *sampling efficiency* K/R measures the number of samples required per tone. Both these numbers range between zero and one.

Our results lead us to an empirical rule for the sampling rate necessary to recover sparse signals using the demodulator system:

$$R \approx 1.7K \log(W/K + 1). \quad (6)$$

The empirical rate matches the phase transition threshold that Donoho and Tanner calculated for compressive sensing problems with a Gaussian sampling matrix [16]. The form of this relation also echoes the form of the information bound developed in Section II-B.

A. Random Signal Model

We begin with a stochastic model for frequency-sparse discrete signals. First, let us define the *signum function*

$$\text{sgn}(re^{i\theta}) \stackrel{\text{def}}{=} \begin{cases} 0, & r = 0, \\ e^{i\theta}, & r > 0. \end{cases}$$

The model is described in the following box:

Model (A) for a random amplitude vector s		
Frequencies:	Ω	is a uniformly random set of K frequencies from $\{0, \pm 1, \dots, \pm(W/2-1), W/2\}$
Amplitudes:	$ s_\omega $ $s_\omega = 0$	is arbitrary for each $\omega \in \Omega$ for each $\omega \notin \Omega$
Phases:	$\text{sgn}(s_\omega)$	is i.i.d. uniform on the unit circle for each $\omega \in \Omega$

In our experiments, we set the amplitude of each nonzero coefficient equal to one because the success of ℓ_1 minimization does not depend on the amplitudes.

B. Experimental Design

Our experiments are intended to determine the minimum sampling rate R that is necessary to identify a K -sparse signal with a bandwidth of W Hz. In each trial, we perform the following steps:

- 1) **Input Signal.** A new amplitude vector s is drawn at random according to Model (A) in Section VI-A. The amplitudes all have magnitude one.
- 2) **Random demodulator.** A new random demodulator Φ is drawn with parameters K , R , and W . The elements of the chipping sequence are independent random variables, equally likely to be ± 1 .
- 3) **Sampling.** The vector of samples is computed using the expression $y = \Phi s$.
- 4) **Reconstruction.** An estimate \hat{s} of the amplitude vector is computed with IRLS.

For each triple (K, R, W) , we perform 500 trials. We declare the experiment a success when $s = \hat{s}$ to machine precision, and we report the smallest value of R for which the empirical failure probability is less than 1%.

C. Performance Results

We begin by evaluating the relationship between the signal bandwidth W and the sampling rate R required to achieve a high probability of successful reconstruction. Figure 4 shows the experimental results for a fixed sparsity of $K = 5$ as W increases from 128 to 2048 Hz, denoted by asterisks. The solid line shows the result

$$R = 1.69K \log(W/K + 1) + 4.51$$

of a linear regression on the experimental data. The variation about the regression line is probably due to arithmetic effects that occur when the sampling rate does not divide the bandwidth.

Next, we evaluate the relationship between the sparsity K and the sampling rate R . Figure 5 shows the experimental results for a fixed chipping rate of $W = 512$ Hz as the sparsity K increases from 1 to 64. The figure also shows that the linear regression

$$R = 1.71K \log(W/K + 1) + 1.00$$

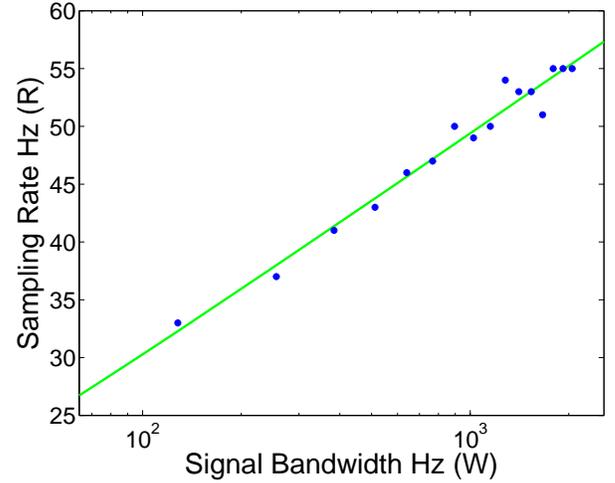


Fig. 4. Sampling rate as a function of signal bandwidth. The sparsity is fixed at $K = 5$. The asterisks denote the lowest sampling rate R that achieves successful reconstruction with probability 0.99. The solid line denotes the fit $R = 1.69K \log(W/K + 1) + 4.51$ for the experimental data.

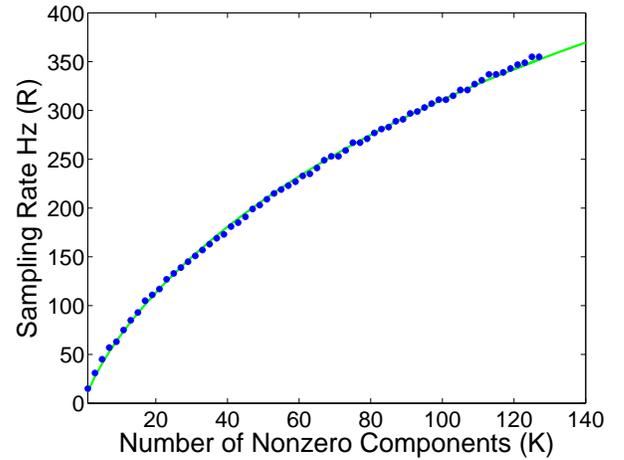


Fig. 5. Sampling rate as a function of sparsity. The bandwidth is fixed at $W = 512$ Hz. The asterisks denote the lowest sampling rate that achieves successful reconstruction with probability 0.99. The solid line denotes the fit $R = 1.71K \log(W/K + 1) + 1.00$ for the experimental data.

gives a close fit for the experimental data.

The regression lines from these two experiments suggest that successful reconstruction of signals from Model (A) occurs with high probability when the sampling rate obeys the bound

$$R \geq 1.7K \log(W/K + 1), \quad (7)$$

Thus, for a fixed sparsity, the sampling rate grows only logarithmically as the Nyquist rate increases. We also note that this bound is similar to those obtained for other measurement schemes that require fully random, dense matrices [17].

Finally, we study the threshold that denotes a change from high to low probability of successful reconstruction. This type of *phase transition* is a common phenomenon in compressive sampling. For this experiment, the chipping rate is fixed at $W = 512$ Hz, while the sparsity K and sampling R rate vary.

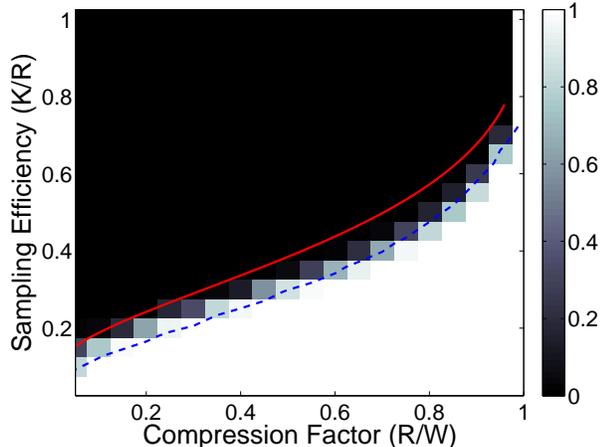


Fig. 6. Probability of success as a function of sampling efficiency and compression factor. The shade of each pixel indicates the probability of successful reconstruction for the corresponding combination of parameter values. The solid line marks the theoretical threshold for a Gaussian matrix; the dashed line traces the 99% success isocone.

We record the probability of success as the compression factor R/W and the sampling efficiency K/R vary.

The experimental results appear in Figure 6. Each pixel in this image represents the probability of success for the corresponding combination of system parameters. Lighter pixels denote higher probability. The dashed line,

$$\frac{K}{R} = \frac{0.68}{\log(W/K + 1)},$$

describes the relationship among the parameters where the probability of success drop below 99%. The numerical parameter was derived from a linear regression without intercept.

For reference, we compare the random demodulator with a benchmark system that obtains measurements of the amplitude vector \mathbf{s} by applying a matrix Φ whose entries are drawn independently from the standard Gaussian distribution. As the dimensions of the matrix grow, ℓ_1 minimization methods exhibit a sharp phase transition from success to failure. The solid line in Figure 6 marks the location of the precipice, which can be computed analytically with methods from [16].

D. Example: Analog Demodulation

We constructed these empirical performance trends using signals drawn from a synthetic model. To narrow the gap between theory and practice, we performed another experiment to demonstrate that the random demodulator can successfully recover a simple communication signal from samples taken below the Nyquist rate.

In the amplitude modulation (AM) encoding scheme, the transmitted signal $f_{AM}(t)$ takes the form

$$f_{AM}(t) = A \cos(2\pi\omega_c t) \cdot (m(t) + C), \quad (8)$$

where $m(t)$ is the original message signal, ω_c is the carrier frequency, and A, C are fixed values. When the original message signal $m(t)$ has K nonzero Fourier coefficients, then

the cosine modulated signal has only $2K + 2$ nonzero Fourier coefficients.

We consider an AM signal that encodes the message appearing in Figure 7(a). The signal was transmitted from a communications device using carrier frequency $\omega_c = 8.2$ KHz, and the received signal was sampled by an ADC at a rate of 32 KHz. Both the transmitter and receiver were isolated in a lab to produce a clean signal; however, noise is still present on the sampled data due to hardware effects and environmental conditions.

We feed the received signal into a *simulated* random demodulator, where we take the Nyquist rate $W = 32$ Hz and we attempt a variety of sampling rates R . We use IRLS to reconstruct the signal from the random samples, and we perform AM demodulation on the recovered signal to reconstruct the original message. Figures 7(b)–(d) display reconstructions for a random demodulator with sampling rates $R = 16$ KHz, 8 KHz, and 3.2 KHz, respectively. To evaluate the quality of the reconstruction, we measure the signal-to-noise ratio (SNR) between the message obtained from the received signal f_{AM} and the message obtained from the output of the random demodulator. The reconstructions achieve SNRs of 27.8 dB, 22.3 dB, and 20.9 dB, respectively.

These results demonstrate that the performance of the random demodulator degrades gracefully as the SNR increases. Note, however, that the system will not function at all unless the sampling rate is sufficiently high that we stand below the phase transition.

VII. THEORETICAL RECOVERY RESULTS

We have been able to establish several theoretical guarantees on the performance of the random demodulator system. First, we focus on the setting described in the numerical experiments, and we develop estimates for the sampling rate required to recover the synthetic sparse signals featured in Section VI-C. Qualitatively, these bounds almost match the system performance we witnessed in our experiments. The second set of results addresses the behavior of the system for a much wider class of input signals. This theory establishes that, when the sampling rate is slightly higher, the signal recovery process will succeed—even if the spectrum is not perfectly sparse and the samples collected by the random demodulator are contaminated with noise.

A. Recovery of Random Signals

First, we study when ℓ_1 minimization can reconstruct random, frequency-sparse signals that have been sampled with the random demodulator. The setting of the following theorem precisely matches our numerical experiments.

Theorem 1 (Recovery of Random Signals): Suppose that the sampling rate

$$R \geq C [K \log W + \log^3 W]$$

and that R divides W .

Let \mathbf{s} be a random amplitude vector drawn according to Model (A). Draw an $R \times W$ random demodulator matrix Φ ,

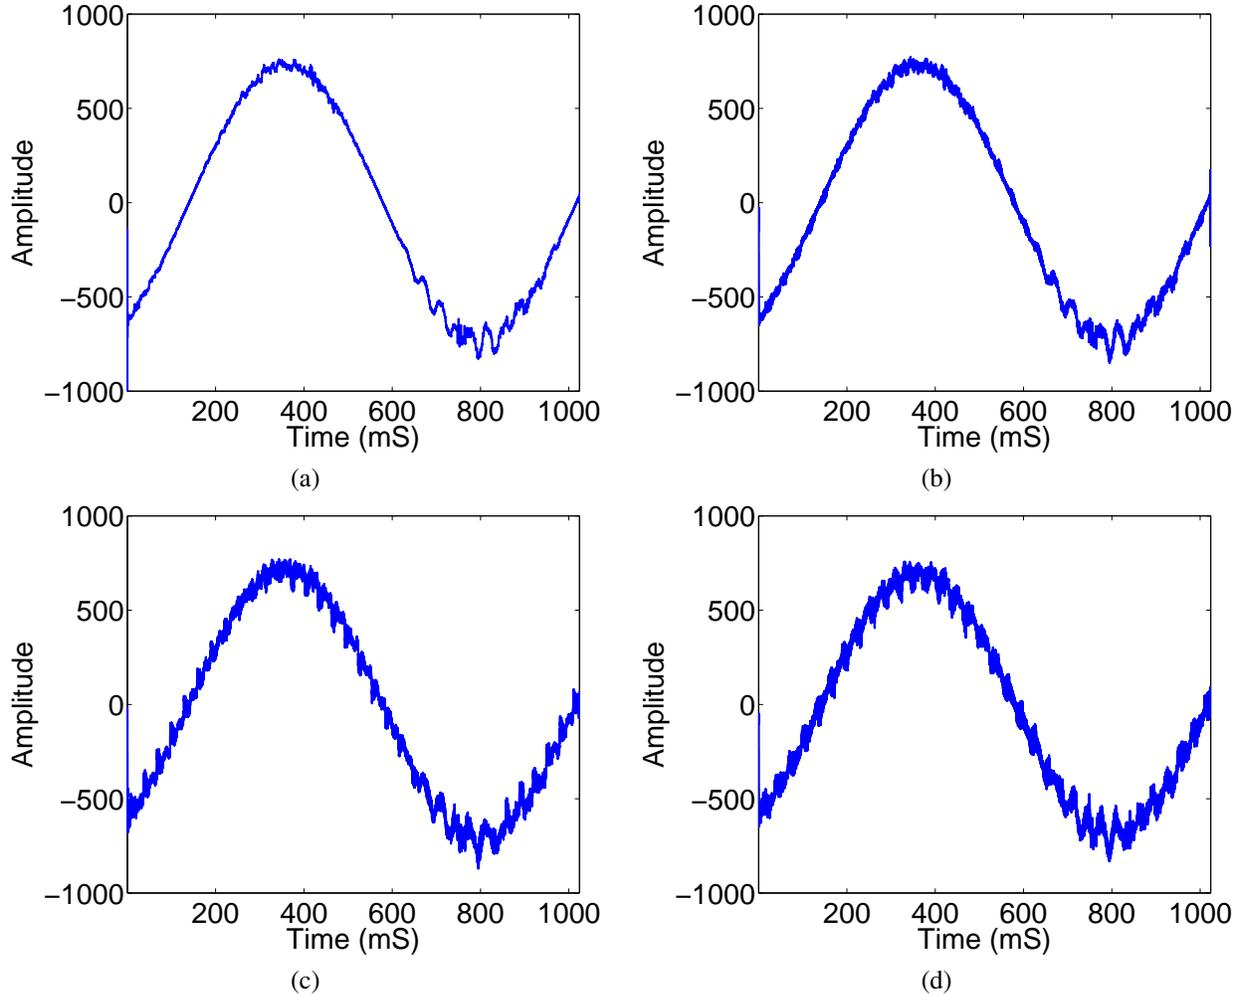


Fig. 7. Simulated acquisition and reconstruction of an AM signal. (a) Message reconstructed from signal sampled at Nyquist rate $W = 32$ KHz. (b) Message reconstructed from the output of a simulated random demodulator running at sampling rate $R = 16$ KHz (SNR = 27.8 dB). (c) Reconstruction with $R = 8$ KHz (SNR = 22.3 dB). (d) Reconstruction with $R = 3.2$ KHz (SNR = 20.9 dB).

and let $\mathbf{y} = \Phi \mathbf{s}$ be the samples collected by the random demodulator.

Then the solution $\hat{\mathbf{s}}$ to the convex program (5) equals \mathbf{s} , except with probability $O(W^{-1})$. The number C is a positive, universal constant.

We have framed the unimportant technical assumption that R divides W to simplify the lengthy argument. See Appendix II for the proof.

Our analysis demonstrates that the sampling rate R scales linearly with the sparsity level K , while it is logarithmic in the bandlimit W . In other words, the theorem supports our empirical rule (6) for the sampling rate. We believe that this result offers an attractive theoretical justification for our claims about the sampling efficiency of the random demodulator.

The theorem also suggests that there is a small startup cost. That is, a minimal number of measurements is required before the demodulator system is effective. Our experiments were not refined enough to detect whether this startup cost actually exists in practice.

B. Uniformity and Stability

Although the results described in the last section provide a satisfactory explanation of the numerical experiments, they are less compelling as a vision of signal recovery in the real world. Theorem 1 has three major shortcomings. First, it is unreasonable to assume that tones are located at random positions in the spectrum, so Model (A) is somewhat artificial. Second, typical signals are not spectrally sparse because they contain background noise and, perhaps, irrelevant low-power frequencies. Finally, the hardware system suffers from nonidealities, so it only approximates the linear transformation described by the matrix M . As a consequence, the samples are contaminated with noise.

To address these problems, we need an algorithm that provides *uniform* and *stable* signal recovery. Uniformity means that the approach works for all signals, irrespective of the frequencies and phases that participate. Stability means that the performance degrades gracefully when the signal's spectrum is not perfectly sparse and the samples are noisy. Fortunately, the compressive sampling community has developed several powerful signal recovery techniques that enjoy both these

properties.

Conceptually, the simplest approach is to modify the convex program (5) to account for the contamination in the samples [18]. Suppose that $\mathbf{s} \in \mathbb{C}^W$ is an arbitrary amplitude vector, and let $\boldsymbol{\nu} \in \mathbb{C}^R$ be an arbitrary noise vector that is known to satisfy the bound $\|\boldsymbol{\nu}\|_2 \leq \eta$. Assume that we have acquired the dirty samples

$$\mathbf{y} = \Phi \mathbf{s} + \boldsymbol{\nu}.$$

To produce an approximation of \mathbf{s} , it is natural to solve the noise-aware optimization problem

$$\hat{\mathbf{s}} = \arg \min \|\mathbf{v}\|_1 \quad \text{subject to} \quad \|\Phi \mathbf{v} - \mathbf{y}\|_2 \leq \eta. \quad (9)$$

As before, this problem can be formulated as a second-order cone program, which can be optimized using various algorithms [18].

We have established a theorem that describes how the optimization-based recovery technique performs when the samples are acquired using the random demodulator system. An analogous result holds for the CoSaMP algorithm, a greedy pursuit method with superb guarantees on its runtime [15].

Theorem 2 (Recovery of General Signals): Suppose that the sampling rate

$$R \geq CK \log^6 W$$

and that R divides W . Draw an $R \times W$ random demodulator matrix Φ . Then the following statement holds, except with probability $O(W^{-1})$.

Suppose that \mathbf{s} is an arbitrary amplitude vector and $\boldsymbol{\nu}$ is a noise vector with $\|\boldsymbol{\nu}\|_2 \leq \eta$. Let $\mathbf{y} = \Phi \mathbf{s} + \boldsymbol{\nu}$ be the noisy samples collected by the random demodulator.

Then every solution $\hat{\mathbf{s}}$ to the convex program (9) approximates the target vector \mathbf{s} :

$$\|\hat{\mathbf{s}} - \mathbf{s}\|_2 \leq C \max \left\{ \eta, \frac{1}{\sqrt{K}} \|\mathbf{s} - \mathbf{s}_K\|_1 \right\}, \quad (10)$$

where \mathbf{s}_K is a best K -sparse approximation to \mathbf{s} with respect to the ℓ_1 norm.

The proof relies on the *restricted isometry property* (RIP) of Candès–Tao [19]. To establish that the random demodulator verifies the RIP, we adapt ideas of Rudelson–Veryshnin [20]. Turn to Appendix III for the argument.

Theorem 2 is not as easy to grasp as Theorem 1, so we must spend some time to unpack its meaning. First, observe that the sampling rate has increased by several logarithmic factors. Some of these factors are probably parasitic, a consequence of the techniques used to prove the theorem. It seems plausible that, in practice, the actual requirement on the sampling rate is closer to

$$R \geq CK \log^2 W.$$

This conjecture is beyond the power of current techniques.

The earlier result, Theorem 1, suggests that we should draw a new random demodulator each time we want to acquire a signal. In contrast, Theorem 2 shows that, with high probability, a particular instance of the random demodulator acquires enough information to reconstruct any amplitude vector whatsoever. This aspect of Theorem 2 has the practical consequence that

a random chipping sequence can be chosen in advance and fixed for all time.

Theorem 2 does not place any specific requirements on the amplitude vector, nor does it model the noise contaminating the signal. Indeed, the approximation error bound (10) holds generally. That said, the strength of the error bound depends substantially on the structure of the amplitude vector. When \mathbf{s} happens to be K -sparse, then the second term in the maximum vanishes. So the convex programming method still has the ability to recover sparse signals perfectly.

When \mathbf{s} is not sparse, we may interpret the bound (10) as saying that the computed approximation is comparable with the best K -sparse approximation of the amplitude vector. When the amplitude vector has a good sparse approximation, then the recovery succeeds well. When the amplitude vector does not have a good approximation, then signal recovery may fail completely. For this class of signal, the random demodulator is not an appropriate technology.

Initially, it may seem difficult to reconcile the two norms that appear in the error bound (10). In fact, this type of mixed-norm bound is structurally optimal [21], so we cannot hope to improve the ℓ_1 norm to an ℓ_2 norm. Nevertheless, for an important class of signals, the scaled ℓ_1 norm of the tail is essentially equivalent to the ℓ_2 norm of the tail.

Let $p \in (0, 1)$. We say that a vector \mathbf{s} is *p-compressible* when its sorted components decay sufficiently fast:

$$|s|_{(k)} \leq k^{-1/p} \quad \text{for } k = 1, 2, 3, \dots \quad (11)$$

Harmonic analysis shows that many natural signal classes are compressible [22]. Moreover, the windowing scheme of Section VIII results in compressible signals.

The critical fact is that compressible signals are well approximated by sparse signals. Indeed, it is straightforward to check that

$$\begin{aligned} \|\mathbf{s} - \mathbf{s}_K\|_1 &\leq C_p K^{1-1/p} \\ \|\mathbf{s} - \mathbf{s}_K\|_2 &\leq D_p K^{1/2-1/p}. \end{aligned}$$

where C_p and D_p are constants that depend only on the level p of compressibility.

For a p -compressible signal, the error bound (10) reads

$$\|\hat{\mathbf{s}} - \mathbf{s}\|_2 \leq C \max \left\{ \eta, K^{1/2-1/p} \right\}.$$

We see that the right-hand side is comparable with the ℓ_2 norm of the tail of the signal. This quantitative conclusion reinforces the intuition that the random demodulator is efficient when the amplitude vector is well approximated by a sparse vector.

C. Extension to Bounded Orthobases

We have shown that the random demodulator is effective at acquiring signals that are spectrally sparse or compressible. In fact, the system can acquire a much more general class of signals. The proofs of Theorem 1 and Theorem 2 indicate that the crucial feature of the sparsity basis \mathbf{F} is its *incoherence* with the Dirac basis. That is, we exploit the fact that the entries of the matrix \mathbf{F} have small magnitude. This insight allow us to extend the recovery theory to other sparsity bases with the same property. We avoid a detailed exposition. See [23] for more discussion of this type of result.

VIII. PREWINDING

We have shown that the random demodulator can acquire periodic multitone signals. The effectiveness of the system for a general signal f depends on how closely we can approximate f on the time interval $[0, 1)$ by a periodic multitone signal. In this section, we argue that many common types of nonperiodic signals can be approximated well using *prewindowing* techniques. In particular, this approach allows us to capture nonharmonic sinusoids and multiband signals with small total bandwidth. We offer a brief overview of the ideas here, deferring a detailed analysis for a later publication.

A. Nonharmonic Tones

First, there is the question of capturing and reconstructing *nonharmonic* sinusoids, signals of the form

$$f(t) = a_{\omega'} e^{-2\pi i \omega' t / W}, \quad \omega' \notin \mathbb{Z}.$$

It is notoriously hard to approximate f on $[0, 1)$ using harmonic sinusoids, because the coefficients $a_{\omega'}$ in (2) are essentially samples of a sinc function. Indeed, the signal f_K , defined as the best approximation of f using K harmonic frequencies, satisfies only

$$\|f - f_K\|_{L_2[0,1]} \lesssim K^{-1/2}, \quad (12)$$

a painfully slow rate of decay.

There is a classical way to remedy this problem. Instead of acquiring f directly, we acquire a smoothly windowed version of f . To fix ideas, suppose that ψ is a window function that vanishes outside the interval $[0, 1)$. Instead of measuring f , we measure $g = \psi \cdot f$. The goal is now to reconstruct this *windowed* signal g .

As before, our ability to reconstruct the windowed signal depends on how well we can approximate it using a periodic multitone signal. In fact, the approximation rate depends on the smoothness of the window. When the continuous-time Fourier transform $\hat{\psi}$ decays like ω^{-r} for some $r \geq 1$, we have that g_K , the best approximation of g using K harmonic frequencies, satisfies

$$\|g - g_K\|_{L_2[0,1]} \lesssim K^{-r+1/2},$$

which decays to zero much faster than the error bound (12). Thus, to achieve a tolerance of ε , we need only about $\varepsilon^{-1/(r-1/2)}$ terms.

In summary, windowing the input signal allows us to closely approximate a nonharmonic sinusoid by a periodic multitone signal; g will be compressible as in (11). If there are multiple nonharmonic sinusoids present, then the number of harmonic tones required to approximate the signal to a certain tolerance scales linearly with the number of nonharmonic sinusoids.

B. Multiband Signals

Windowing also enables us to treat signals that occupy a small band in frequency. Suppose that f is a signal whose continuous-time Fourier transform \hat{f} vanishes outside an interval of length B . The Fourier transform of the windowed signal $g = \psi \cdot f$ can be broken into two pieces. One part is nonzero only on the support of \hat{f} ; the other part decays like

ω^{-r} away from this interval. As a result, we can approximate g to a tolerance of ε using a multitone signal with about $B + \varepsilon^{-1/(r-1/2)}$ terms.

For a multiband signal, the same argument applies to each band. Therefore, the number of tones required for the approximation scales linearly with the total bandwidth.

C. A Vista on Windows

To reconstruct the original signal f over a long period of time, we must window the signal with overlapping shifts of the window ψ . This window needs to have additional properties for this scheme to work. Suppose that the set of half-integer shifts of ψ form a partition of unity, i.e.,

$$\sum_k \psi(t - k/2) = 1, \quad t \in \mathbb{R}.$$

After measuring and reconstructing $g_k(t) = \psi(t - k/2) \cdot f(t)$ on each subinterval, we simply add the reconstructions together to obtain a complete reconstruction of f .

This prewindowing strategy relies on our ability to measure the windowed signal $\psi \cdot f$. To accomplish this, we require a slightly more complicated system architecture. To perform the windowing, we use an amplifier with a time-varying gain. Since subsequent windows overlap, we will also need to measure g_k and g_{k+1} simultaneously, which creates the need for two random demodulator channels running in parallel. In principle, none of these modifications diminishes the practicality of the system.

IX. DISCUSSION

This section develops some additional themes that arise from our work. First, we discuss the technological impact of the random demodulator system. Afterward, we present a comparison with nonuniform sampling, which is a competing technique for acquiring spectrally sparse signals. We conclude with an overview of other related research.

A. Technological Impact

A significant benefit of the random demodulator is that we can control the SNR performance of the reconstruction by optimizing the sampling rate of the back-end ADC, as demonstrated in Section VI-D. When input signals are spectrally sparse, the random demodulator system can outperform a standard ADC by a substantial margin.

To quantify the SNR performance of an ADC, we consider a standard metric, the effective number of bits (ENOB), which is roughly the actual number of quantization levels possible at a given SNR. The ENOB is calculated as

$$\text{ENOB} = (\text{SNR} - 1.76)/6.02, \quad (13)$$

where the SNR is expressed in dB. We estimate the ENOB of a modern ADC using the information in Walden's 1999 survey [7].

When the input signal has sparsity K and bandlimit W , we can acquire the signal using a random demodulator. We can estimate the rate R required of the demodulator using the relation (7). As noted, this rate is typically much lower than

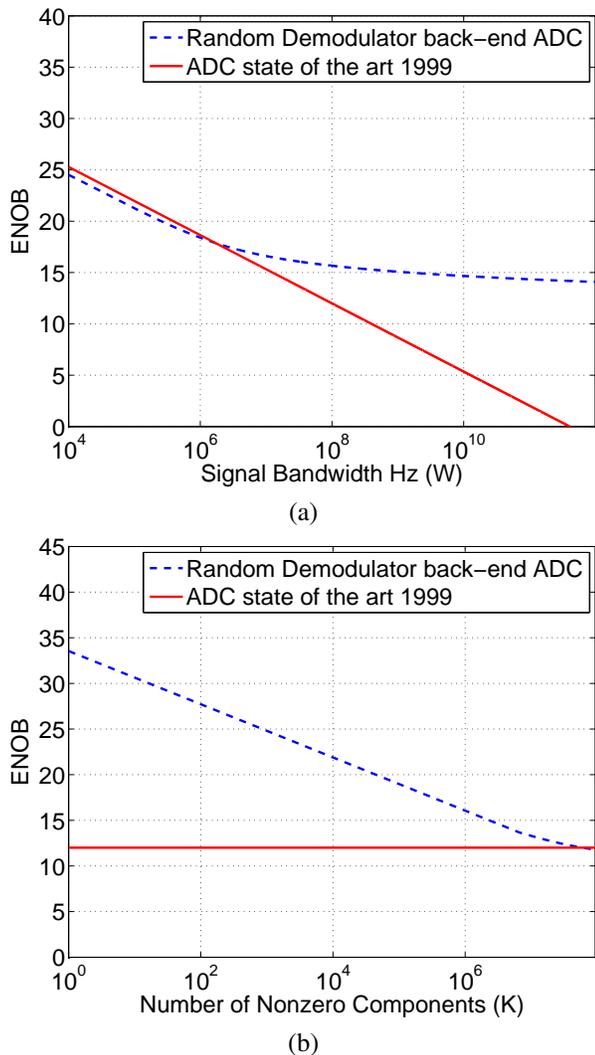


Fig. 8. *ENOB for random demodulator versus a standard ADC.* The solid lines represent state-of-the-art ADC performance in 1999. The dashed lines represent the random demodulator performance using an ADC running at the sampling rate R suggested by our empirical work. (a) Performance as a function of the bandlimit W with fixed sparsity $K = 10^6$. (b) Performance as a function of the sparsity K with fixed bandlimit $W = 10^8$. Note that the performance of a standard ADC does not depend on K .

the Nyquist rate. Since low-rate ADCs exhibit much better SNR performance than high-rate ADCs, the demodulator can produce a higher ENOB.

Figure 8 charts how much the random demodulator improves on state-of-the-art ADCs when the input signals are assumed to be sparse. The first panel, Figure 8(a), compares the random demodulator with a standard ADC for signals with the sparsity $K = 10^6$ as the Nyquist rate W varies up to 10^{12} Hz. In this setting, the clear advantages of the random demodulator can be seen in the slow decay of the curve. The second panel, Figure 8(b), estimates the ENOB as a function of the sparsity K when the bandlimit is fixed at $W = 10^8$ Hz. Again, a distinct advantage is visible. But we stress that this improvement is contingent on the sparsity of the signal.

Although this discussion suggests that the random demod-

ulator is a clear favorite, we have neglected some important factors. The estimated performance of the demodulator does not take into account SNR degradations due to the chain of hardware components upstream of the sampler. These degradations are caused by factors such as the nonlinearity of the multiplier and jitter of the pseudorandom modulation signal. Thus, it is critical to choose high-quality components in the design. Even under this constraint, the system has a more attractive feasibility and cost profile than a high-rate ADC.

It is worth noting an additional benefit of the random demodulator. The samples collected by the system automatically compress a sparse signal into a minimal amount of data. As a result, the storage and communication requirements of the signal acquisition system are also reduced. Furthermore, no extra processing is required after sampling to compress the data, which decreases the complexity of the required hardware and its power consumption.

B. Comparison with Nonuniform Sampling

Other research [11], [24] has shown that frequency-sparse, bandlimited signals can be acquired by sampling nonuniformly in time at an average rate comparable with (15). This type of nonuniform sampling can be implemented by changing the clock input to a standard ADC. Although the random demodulator system involves more components, it has several advantages over nonuniform sampling.

First, nonuniform samplers are extremely sensitive to *timing jitter*. Consider the problem of acquiring a signal with high-frequency components by means of nonuniform sampling. Since the signal values change rapidly, a small error in the sampling time can result in an erroneous sample value. The random demodulator, on the other hand, benefits from the integrator, which effectively lowers the bandwidth of the input into the ADC. Moreover, the random demodulator uses a uniform clock, which is more stable to generate.

Second, the SNR in the measurements from a random demodulator is much higher than the SNR in the measurements from a nonuniform sampler. Suppose that we are acquiring a single sinusoid with unit amplitude. On average, each sample has magnitude $2^{-1/2}$, so if we take W samples at the Nyquist rate, the total energy in the samples is $W/2$. If we take R nonuniform samples the total energy will be on average $R/2$. In contrast, if we take R samples with the random demodulator, each sample has an approximate magnitude of $\sqrt{W/R}$, so the total energy in the samples is about W . In consequence, signal recovery using samples from the random demodulator is more robust against additive noise.

C. Related Work

The classical approach to recovering bandlimited signals from time samples is, of course, the well-known method associated with the names of Shannon, Nyquist, Whittaker, Kotelnikov, and others. In the 1960s, Landau demonstrated that *stable* reconstruction of a bandlimited signal demands a sampling rate no less than the Nyquist rate [4]. Landau also considered multiband signals, those bandlimited signals whose spectrum is supported on a collection of frequency intervals.

Roughly speaking, he proved that a sequence of time samples stably determines a multiband signal if and only if the average sampling rate exceeds the measure of the occupied part of the spectrum [25].

In the 1990s, researchers began to consider practical sampling schemes for acquiring multiband signals. The earliest effort is probably a paper of Feng and Bresler, who assumed that the band locations were known in advance [26]. See also the subsequent work [27]. Afterward, Mishali and Eldar showed that related methods can be used to acquire multiband signals *without* knowledge of the band locations [28]. Eldar also considered the case of general shift-invariant spaces [29].

During the 1990s and the early years of the present decade, a separate strand of work appeared in the literature on theoretical computer science. Researchers developed computationally efficient algorithms for approximating the Fourier transform of a signal, given a small number of random time samples from a structured grid. The earliest work in this area was due to Kushilevitz and Mansour [30], while the method was perfected by Gilbert et al. [31], [32]. See [24] for a tutorial on these ideas. These schemes provide an alternative approach to sub-Nyquist ADCs [33], [34].

Vetterli and collaborators have developed an alternative framework for sub-Nyquist sampling and reconstruction, called “finite rate of innovation” sampling, that passes an analog signal f having K degrees of freedom per second through a linear time-invariant filter and then samples at a rate above $2K$ Hz. Reconstruction is then performed by rooting a high-order polynomial [35], [36], [37], [38]. While this sampling rate is less than the $O(K \log(W/K))$ required by the random demodulator, a proof of the numerical stability of this method has proved elusive.

The random demodulator was introduced in two earlier papers [39], [40], which offer preliminary work on the system architecture and experimental studies of the system’s behavior. The current paper can be seen as an expansion of these articles, because we offer detailed information on performance trends as a function of the signal bandlimit, the sparsity, and the number of samples. We have also developed novel theoretical foundations that support these empirical results.

Very recently, it has been shown that a bank of random demodulators can be used for blind acquisition of multiband signals [41]. This approach processes the samples in a rather different fashion from the methods described in the present work.

Finally, we must mention that the current work can be placed naturally within the framework of compressive sampling. This theory, initiated in the papers [11], [42], has shown that random measurements are an efficient and practical method for acquiring compressible signals. For the most part, compressive sampling has concentrated on finite-length, discrete-time signals. One of the innovations in this work is to transport the continuous-time signal acquisition problem into a setting where established techniques apply. Indeed, our empirical and theoretical estimates for the sampling rate of the random demodulator echo established results from the compressive sampling literature regarding the number of measurements needed to acquire sparse signals.

APPENDIX I THE RANDOM DEMODULATOR MATRIX

This appendix collects some facts about the random demodulator matrix that we use to prove the recovery bounds.

A. Notation

Let us begin with some notation. First, we abbreviate $\llbracket W \rrbracket = \{1, 2, \dots, W\}$. We write $*$ for the complex conjugate transpose of a scalar, vector, or matrix. The symbol $\|\cdot\|$ denotes the spectral norm of a matrix, while $\|\cdot\|_F$ indicates the Frobenius norm. We write $\|\cdot\|_{\max}$ for the maximum absolute entry of a matrix. Other norms will be introduced as necessary. The probability of an event is expressed as $\mathbb{P}\{\cdot\}$, and we use \mathbb{E} for the expectation operator. For a random variable Z , we define its L_p norm

$$\mathbb{E}_p Z = (\mathbb{E} |Z|^p)^{1/p}.$$

Finally, we remind the reader of the analyst’s convention that roman letters c , C , etc. denote universal constants that may change at every appearance.

B. Background

This section contains a potpourri of basic results that will be helpful to us.

We begin with a simple technique for bounding the moments of a maximum. Consider an arbitrary set $\{Z_1, \dots, Z_N\}$ of random variables. It holds that

$$\mathbb{E}_p \max_j Z_j \leq N^{1/p} \max_j \mathbb{E}_p Z_j. \quad (14)$$

To check this claim, simply note that

$$\begin{aligned} \left[\mathbb{E} \max_j |Z_j|^p \right]^{1/p} &\leq \left[\mathbb{E} \sum_j |Z_j|^p \right]^{1/p} \\ &= \left[\sum_j \mathbb{E} |Z_j|^p \right]^{1/p} \leq [N \cdot \max_j \mathbb{E} |Z_j|^p]^{1/p}. \end{aligned}$$

In many cases, this inequality yields essentially sharp results for the appropriate choice of p .

The simplest probabilistic object is the *Rademacher* random variable, which takes the two values ± 1 with equal probability. A sequence of independent Rademacher variables is referred to as a *Rademacher sequence*. A *Rademacher series* in a Banach space X is a sum of the form

$$\sum_{j=1}^{\infty} \xi_j \mathbf{x}_j$$

where $\{\mathbf{x}_j\}$ is a sequence of coefficients in X and $\{\xi_j\}$ is an (independent) Rademacher sequence.

For Rademacher series with scalar coefficients, the most important result is the inequality of Khintchine. The following sharp version is due to Haagerup [43].

Proposition 3 (Khintchine): Let $p \geq 2$. For every sequence $\{a_j\}$ of complex scalars,

$$\mathbb{E}_p \left| \sum_j \xi_j a_j \right| \leq C_p \left[\sum_j |a_j|^2 \right]^{1/2},$$

where the optimal constant

$$C_p = \left[\frac{p!}{2^{p/2}(p/2)!} \right]^{1/p} \leq 2^{1/2p} e^{-0.5} \sqrt{p}.$$

This inequality is typically established only for real scalars, but the real case implies that the complex case holds with the same constant.

Rademacher series appear as a basic tool for studying sums of independent random variables in a Banach space, as illustrated in the following proposition [44, Lem. 6.3].

Proposition 4 (Symmetrization): Let $\{Z_j\}$ be a finite sequence of independent, zero-mean random variables taking values in a Banach space X . Then

$$\mathbb{E}_p \left\| \sum_j Z_j \right\|_X \leq 2 \mathbb{E}_p \left\| \sum_j \xi_j Z_j \right\|_X,$$

where $\{\xi_j\}$ is a Rademacher sequence independent of $\{Z_j\}$.

In words, the moments of the sum are controlled by the moments of the associated Rademacher series. The advantage of this approach is that we can condition on the choice of $\{Z_j\}$ and apply sophisticated methods to estimate the moments of the Rademacher series.

Finally, we need some facts about *symmetrized* random variables. Suppose that Z is a zero-mean random variable Z that takes values in a Banach space X . We may define the symmetrized variable $Y = Z - Z'$, where Z' is an independent copy of Z . The tail of the symmetrized variable Y is closely related to the tail of Z . Indeed,

$$\mathbb{P} \{ \|Z\|_X > 2 \mathbb{E} \|Z\|_X + u \} \leq \mathbb{P} \{ \|Y\|_X > u \}. \quad (15)$$

This relation follows from [44, Eqn. (6.2)] and the fact that $\text{Med } Y \leq 2 \mathbb{E} Y$ for every nonnegative random variable Y .

C. The Random Demodulator Matrix

We continue with a review of the structure of the random demodulator matrix, which is the central object of our affection. Throughout the appendices, we assume that the sampling rate R divides the bandlimit W . That is,

$$W/R \in \mathbb{Z}.$$

Recall that the $R \times W$ random demodulator matrix Φ is defined via the product

$$\Phi = \mathbf{H} \mathbf{D} \mathbf{F}.$$

We index the R rows of the matrix with the Roman letter r , while we index the W columns with the Greek letters α, ω . It is convenient to summarize the properties of the three factors.

The matrix \mathbf{F} is a permutation of the $W \times W$ discrete Fourier transform matrix. In particular, it is unitary, and each of its entries has magnitude $W^{-1/2}$.

The matrix \mathbf{D} is a random $W \times W$ diagonal matrix. The entries of \mathbf{D} are the elements $\varepsilon_1, \dots, \varepsilon_W$ of the chipping sequence, which we assume to be a Rademacher sequence. Since the nonzero entries of \mathbf{D} are ± 1 , it follows that the matrix is unitary.

The matrix \mathbf{H} is an $R \times W$ matrix with 0–1 entries. Each of its rows contains a block of W/R contiguous ones, and the rows are orthogonal. These facts imply that

$$\|\mathbf{H}\| = \sqrt{W/R}. \quad (16)$$

To keep track of the locations of the ones, the following notation is useful. We write

$$j \sim r \quad \text{when} \quad (r-1)W/R < j \leq rW/R.$$

Thus, each value of r is associated with W/R values of j . The entry $h_{rj} = 1$ if and only if $j \sim r$.

The spectral norm of Φ is determined completely by its structure.

$$\|\Phi\| = \|\mathbf{H} \mathbf{D} \mathbf{F}\| = \|\mathbf{H}\| = \sqrt{W/R} \quad (17)$$

because \mathbf{F} and \mathbf{D} are unitary.

Next, we present notations for the entries and columns of the random demodulator. For an index $\omega \in \llbracket W \rrbracket$, we write φ_ω for the ω th column of Φ . Expanding the matrix product, we can express

$$\varphi_\omega = \sum_{j=1}^W \varepsilon_j f_{j\omega} \mathbf{h}_j \quad (18)$$

where $f_{j\omega}$ is the (j, ω) entry of \mathbf{F} and \mathbf{h}_j is the j th column of \mathbf{H} . Similarly, we write $\varphi_{r\omega}$ for the (r, ω) entry of the matrix Φ . The entry can also be expanded as a sum

$$\varphi_{r\omega} = \sum_{j \sim r} \varepsilon_j f_{j\omega}. \quad (19)$$

Finally, given a set Ω of column indices, we define the column submatrix Φ_Ω of Φ whose columns are listed in Ω .

D. A Componentwise Bound

The first key property of the random demodulator matrix is that its entries are small with high probability. We apply this bound repeatedly in the subsequent arguments.

Lemma 5: Let Φ be an $R \times W$ random demodulator matrix. When $2 \leq p \leq 4 \log W$, we have

$$\mathbb{E}_p \|\Phi\|_{\max} \leq \sqrt{\frac{6 \log W}{R}}.$$

Furthermore,

$$\mathbb{P} \left\{ \|\Phi\|_{\max} > \sqrt{\frac{10 \log W}{R}} \right\} \leq W^{-1}.$$

Proof: As noted in (19), each entry of Φ is a Rademacher series:

$$\varphi_{r\omega} = \sum_{j \sim r} \varepsilon_j f_{j\omega}.$$

Observe that

$$\sum_{j \sim r} |f_{j\omega}|^2 = \frac{W}{R} \cdot \frac{1}{W} = \frac{1}{R}.$$

because the entries of \mathbf{F} all share the magnitude $W^{-1/2}$. Khintchine's inequality, Proposition 3, provides a bound on the moments of the Rademacher series:

$$\mathbb{E}_p |\varphi_{r\omega}| \leq \frac{C_p}{\sqrt{R}}$$

where $C_p \leq 2^{0.25} e^{-0.5} \sqrt{p}$.

We now compute the moments of the maximum entry. Let

$$M = \|\Phi\|_{\max} = \max_{r,\omega} |\varphi_{r\omega}|.$$

Select $q = \max\{p, 4 \log W\}$, and invoke Hölder's inequality.

$$\mathbb{E}_p M \leq \mathbb{E}_q \max_{r,\omega} |\varphi_{r\omega}|.$$

Inequality (14) yields

$$\mathbb{E}_p M \leq (RW)^{1/q} \max_{r,\omega} \mathbb{E}_q |\varphi_{r\omega}|.$$

Apply the moment bound for the individual entries to reach

$$\mathbb{E}_p M \leq \frac{C_q (RW)^{1/q}}{\sqrt{R}}.$$

Since $R \leq W$ and $q \geq 4 \log W$, we have $(RW)^{1/q} \leq e^{0.5}$. Simplify the constant to discover that

$$\mathbb{E}_p M \leq 2^{1.25} \sqrt{\frac{\log W}{R}}.$$

A numerical bound yields the first conclusion.

To obtain a probability bound, we apply Markov's inequality, which is the standard device. Indeed,

$$\mathbb{P}\{M > u\} = \mathbb{P}\{M^q > u^q\} \leq \left[\frac{\mathbb{E}_q M}{u} \right]^q.$$

Choose $u = e^{0.25} \mathbb{E}_q M$ to obtain

$$\mathbb{P}\left\{M > 2^{1.25} e^{0.25} \sqrt{\frac{\log W}{R}}\right\} \leq e^{-\log W} = W^{-1}.$$

Another numerical bound completes the demonstration. ■

E. The Gram Matrix

Next, we develop some information about the inner products between columns of the random demodulator. Let α and ω be column indices in $[[W]]$. Using the expression (18) for the columns, we find that

$$\langle \varphi_\alpha, \varphi_\omega \rangle = \sum_{j,k=1}^W \varepsilon_j \varepsilon_k \eta_{jk} f_{j\alpha}^* f_{k\omega}$$

where we have abbreviated $\eta_{jk} = \langle \mathbf{h}_k, \mathbf{h}_j \rangle$. Since $\eta_{jj} = 1$, the sum of the diagonal terms is

$$\sum_j f_{j\alpha}^* f_{j\omega} = \begin{cases} 1, & \alpha = \omega \\ 0, & \alpha \neq \omega \end{cases}$$

because the columns of \mathbf{F} are orthonormal. Therefore,

$$\langle \varphi_\alpha, \varphi_\omega \rangle = \delta_{\alpha\omega} + \sum_{j \neq k} \varepsilon_j \varepsilon_k \eta_{jk} f_{j\alpha}^* f_{k\omega}$$

where $\delta_{\alpha\omega}$ is the Kronecker delta.

The Gram matrix $\Phi^* \Phi$ tabulates these inner products. As a result, we can express the latter identity as

$$\Phi^* \Phi = \mathbf{I} + \mathbf{X}$$

where

$$x_{\alpha\omega} = \sum_{j \neq k} \varepsilon_j \varepsilon_k \eta_{jk} f_{j\alpha}^* f_{k\omega}.$$

It is clear that $\mathbb{E} \mathbf{X} = \mathbf{0}$, so that

$$\mathbb{E} \Phi^* \Phi = \mathbf{I}. \quad (20)$$

We interpret this relation to mean that, on average, the columns of Φ form an orthonormal system. This ideal situation cannot obtain since Φ has more columns than rows. The matrix \mathbf{X} measures the discrepancy between reality and this impossible dream. Indeed, most of argument can be viewed as a collection

of norm bounds on \mathbf{X} that quantify different aspects of this deviation.

Our central result provides a uniform bound on the entries of \mathbf{X} that holds with high probability. In the sequel, we exploit this fact to obtain estimates for several other quantities of interest.

Lemma 6: Suppose that $R \geq 2 \log W$. Then

$$\mathbb{P}\left\{\|\mathbf{X}\|_{\max} > C \sqrt{\frac{\log W}{R}}\right\} \leq W^{-1}.$$

Proof: The object of our attention is

$$M = \|\mathbf{X}\|_{\max} = \max_{\alpha,\omega} \left| \sum_{j \neq k} \varepsilon_j \varepsilon_k \eta_{jk} f_{j\alpha} f_{k\omega}^* \right|.$$

A random variable of this form is called a *second-order Rademacher chaos*, and there are sophisticated methods available for studying its distribution. We use these ideas to develop a bound on $\mathbb{E}_p M$ for $p = 2 \log W$. Afterward, we apply Markov's inequality to obtain the tail bound.

For technical reasons, we must rewrite the chaos before we start making estimates. First, note that we can express the chaos in a more symmetric fashion:

$$M = \max_{\alpha,\omega} \left| \sum_{j \neq k} \varepsilon_j \varepsilon_k \cdot \eta_{jk} \cdot \frac{1}{2} (f_{j\alpha} f_{k\omega}^* + f_{k\alpha} f_{j\omega}^*) \right|.$$

It is also simpler to study the real and imaginary parts of the chaos separately. To that end, define

$$a_{jk}^{\alpha\omega} = \begin{cases} \eta_{jk} \cdot \frac{1}{2} \operatorname{Re}(f_{j\alpha} f_{k\omega}^* + f_{k\alpha} f_{j\omega}^*), & j \neq k \\ 0, & j = k \end{cases}$$

$$b_{jk}^{\alpha\omega} = \begin{cases} \eta_{jk} \cdot \frac{1}{2} \operatorname{Im}(f_{j\alpha} f_{k\omega}^* + f_{k\alpha} f_{j\omega}^*), & j \neq k \\ 0, & j = k \end{cases}$$

With this notation,

$$M \leq \max_{\alpha,\omega} \left| \sum_{j,k} \varepsilon_j \varepsilon_k a_{jk}^{\alpha\omega} \right| + \max_{\alpha,\omega} \left| \sum_{j,k} \varepsilon_j \varepsilon_k b_{jk}^{\alpha\omega} \right|$$

$$\stackrel{\text{def}}{=} M_{\text{Re}} + M_{\text{Im}}.$$

We focus on the real part since the imaginary part receives an identical treatment.

The next step toward obtaining a tail bound is to decouple the chaos. Define

$$Y = \max_{\alpha,\omega} \left| \sum_{j,k} \varepsilon'_j \varepsilon_k a_{jk}^{\alpha\omega} \right|$$

where $\{\varepsilon'_j\}$ is an independent copy of $\{\varepsilon_j\}$. Standard decoupling results [45, Prop. 1.9] imply that

$$\mathbb{E}_p M_{\text{Re}} \leq 4 \mathbb{E}_p Y.$$

So it suffices to study the decoupled variable Y .

To approach this problem, we first bound the moments of each random variable that participates in the maximum. Fix a pair (α, ω) of frequencies. We omit the superscript α and ω to simplify notations. Define the auxiliary random variable

$$Z = Z_{\alpha\omega} = \left| \sum_{j,k} \varepsilon'_j \varepsilon_k a_{jk} \right|$$

Construct the matrix $\mathbf{A} = [a_{jk}]$. As we will see, the variation of Z is controlled by spectral properties of the matrix \mathbf{A} .

For that purpose, we compute the Frobenius norm and spectral norm of \mathbf{A} . Each of its entries satisfies

$$|a_{jk}| \leq \frac{1}{\sqrt{W}} \eta_{jk} = \frac{1}{\sqrt{W}} (\mathbf{H}^* \mathbf{H})_{jk}$$

owing to the fact that the entries of \mathbf{F} are uniformly bounded by $W^{-1/2}$. The structure of \mathbf{H} implies that $\mathbf{H}^* \mathbf{H}$ is a symmetric, 0–1 matrix with exactly W/R nonzero entries in each of its W rows. Therefore,

$$\|\mathbf{A}\|_{\text{F}} \leq \frac{1}{\sqrt{W}} \|\mathbf{H}^* \mathbf{H}\|_{\text{F}} = \frac{1}{\sqrt{W}} \sqrt{W \cdot \frac{W}{R}} = \frac{1}{\sqrt{R}}.$$

The spectral norm of \mathbf{A} is bounded by the spectral norm of its entrywise absolute value $\text{abs}(\mathbf{A})$. Applying the fact (16) that $\|\mathbf{H}\| = \sqrt{W/R}$, we obtain

$$\|\mathbf{A}\| \leq \|\text{abs}(\mathbf{A})\| \leq \frac{1}{\sqrt{W}} \|\mathbf{H}^* \mathbf{H}\| = \frac{1}{\sqrt{W}} \|\mathbf{H}\|^2 = \frac{1}{R}.$$

Let us emphasize that the norm bounds are uniform over all pairs (α, ω) of frequencies. Moreover, the matrix $\mathbf{B} = [b_{jk}]$ satisfies identical bounds.

To continue, we estimate the mean of the variable Z . This calculation is a simple consequence of Hölder's inequality:

$$\begin{aligned} \mathbb{E} Z &\leq (\mathbb{E} Z^2)^{1/2} = \left[\mathbb{E} \left| \sum_j \varepsilon'_j \left(\sum_k \varepsilon_k a_{jk} \right) \right|^2 \right]^{1/2} \\ &= \left[\mathbb{E} \sum_j \left| \sum_k \varepsilon_k a_{jk} \right|^2 \right]^{1/2} = \left[\sum_{j,k} |a_{jk}|^2 \right]^{1/2} \\ &= \|\mathbf{A}\|_{\text{F}} = \frac{1}{\sqrt{R}}. \end{aligned}$$

Chaos random variables, such as Z , concentrate sharply about their mean. Deviations of Z from its mean are controlled by two separate variances. The probability of large deviation is determined by

$$U = \sup_{\|\mathbf{u}\|_2=1} \left| \sum_{j,k} u_j u_k a_{jk} \right|,$$

while the probability of a moderate deviation depends on

$$V = \mathbb{E} \sup_{\|\mathbf{u}\|_2=1} \left| \sum_{j,k} u_j \varepsilon_k a_{jk} \right|.$$

To compute the first variance U , observe that

$$U = \|\mathbf{A}\| \leq \frac{1}{R}.$$

The second variance is not much harder. Using Jensen's inequality, we find

$$\begin{aligned} V &= \mathbb{E} \left[\sum_j \left| \sum_k \varepsilon_k a_{jk} \right|^2 \right]^{1/2} \\ &\leq \left[\mathbb{E} \sum_j \left| \sum_k \varepsilon_k a_{jk} \right|^2 \right]^{1/2} = \frac{1}{\sqrt{R}}. \end{aligned}$$

We are prepared to appeal to the following theorem which bounds the moments of a chaos variable [46, Cor. 2].

Theorem 7 (Moments of chaos): Let Z be a decoupled, symmetric, second-order chaos. Then

$$\mathbb{E}_p |Z - \mathbb{E} Z| \leq K [\sqrt{p} V + p U]$$

for each $p \geq 1$.

Substituting the values of the expectation and variances, we reach

$$\mathbb{E}_p \left[Z - \frac{1}{\sqrt{R}} \right] \leq K \left[\sqrt{\frac{p}{R}} + \frac{p}{R} \right].$$

The content of this estimate is clearer, perhaps, if we split the bound at $p = R$.

$$\mathbb{E}_p \left[Z - \frac{1}{\sqrt{R}} \right] \leq \begin{cases} 2K\sqrt{p/R}, & p \leq R \\ 2Kp/R, & p > R. \end{cases}$$

In words, the variable Z exhibits subgaussian behavior in the moderate deviation regime, but its tail is subexponential.

We are now prepared to bound the moments of the maximum of the ensemble $\{Z_{\alpha\omega}\}$. When $p = 2 \log W \leq R$, inequality (14) yields

$$\begin{aligned} \mathbb{E}_p \max_{\alpha, \omega} \left[Z_{\alpha\omega} - \frac{1}{\sqrt{R}} \right] &\leq W^{2/p} \max_{\alpha, \omega} \mathbb{E}_p \left[Z_{\alpha\omega} - \frac{1}{\sqrt{R}} \right] \\ &\leq e \cdot 2K \sqrt{\frac{p}{R}}, \end{aligned}$$

Recalling the definitions of Y and $Z_{\alpha\omega}$, we reach

$$\mathbb{E}_p Y = \mathbb{E}_p \max_{\alpha, \omega} Z_{\alpha\omega} \leq C \sqrt{\frac{\log W}{R}}.$$

In words, we have obtained a moment bound for the decoupled version of the real chaos M_{Re} .

To complete the proof, remember that $\mathbb{E}_p M_{\text{Re}} \leq 4 \mathbb{E}_p Y$. Therefore,

$$\mathbb{E}_p M_{\text{Re}} \leq 4C \sqrt{\frac{\log W}{R}}.$$

An identical argument establishes that

$$\mathbb{E}_p M_{\text{Im}} \leq 4C \sqrt{\frac{\log W}{R}}.$$

Since $M \leq M_{\text{Re}} + M_{\text{Im}}$, it follows inexorably that

$$\mathbb{E}_p M \leq 8C \sqrt{\frac{\log W}{R}}.$$

Finally, we invoke Markov's inequality to obtain the tail bound

$$\mathbb{P} \left\{ M \geq e^{0.5} \cdot 8C \sqrt{\frac{\log W}{R}} \right\} \leq e^{-0.5p} = W^{-1}.$$

This endeth the lesson. ■

F. Column Norms and Coherence

Lemma 6 has two important and immediate consequences. First, it implies that the columns of Φ essentially have unit norm.

Theorem 8 (Column Norms): Suppose the sampling rate

$$R \geq C \delta^{-2} \log W.$$

An $R \times W$ random demodulator Φ satisfies

$$\mathbb{P} \left\{ \max_{\omega} \left| \|\varphi_{\omega}\|_2^2 - 1 \right| \geq \delta \right\} \leq W^{-1}.$$

Lemma 6 also shows that the *coherence* of the random demodulator is small. The coherence, which is denoted by μ , bounds the maximum inner product between distinct columns

of Φ , and it has emerged as a fundamental tool for establish the success of compressive sampling recovery algorithms. Rigorously,

$$\mu = \max_{\alpha \neq \omega} |\langle \varphi_\alpha, \varphi_\omega \rangle|.$$

We have the following probabilistic bound.

Theorem 9 (Coherence): Suppose the sampling rate

$$R \geq 2 \log W.$$

An $R \times W$ random demodulator Φ satisfies

$$\mathbb{P} \left\{ \mu \geq C \sqrt{\frac{\log W}{R}} \right\} \leq W^{-1}.$$

For a general $R \times W$ matrix with unit-norm columns, we can verify [47, Thm. 2.3] that its coherence

$$\mu \geq \frac{1}{\sqrt{R}} \left[1 - \frac{R}{W} \right] \approx \frac{1}{\sqrt{R}}.$$

Since the columns of the random demodulator are essentially normalized, we conclude that its coherence is nearly as small as possible.

APPENDIX II

RECOVERY FOR THE RANDOM PHASE MODEL

In this appendix, we establish Theorem 1, which shows that ℓ_1 minimization is likely to recover a random signal drawn from Model (A). Appendix III develops results for general signals.

The performance of ℓ_1 minimization for random signals depends on several subtle properties of the demodulator matrix. We must discuss these ideas in some detail. In the next two sections, we use the bounds from Appendix I to check that the required conditions are in force. Afterward, we proceed with the demonstration of Theorem 1.

A. Cumulative Coherence

The coherence measures only the inner product between a single pair of columns. To develop accurate results on the performance of ℓ_1 minimization, we need to understand the total correlation between a fixed column and a collection of distinct columns.

Let Φ be an $R \times W$ matrix, and let Ω be a subset of $\llbracket W \rrbracket$. The *local 2-cumulative coherence* of the set Ω with respect to the matrix Φ is defined as

$$\mu_2(\Omega) = \max_{\alpha \notin \Omega} \left[\sum_{\omega \in \Omega} |\langle \varphi_\alpha, \varphi_\omega \rangle|^2 \right]^{1/2}.$$

The coherence μ provides an immediate bound on the cumulative coherence:

$$\mu_2(\Omega) \leq \mu \sqrt{|\Omega|}.$$

Unfortunately, this bound is completely inadequate.

To develop a better estimate, we instate some additional notation. Consider the matrix norm $\|\cdot\|_{1 \rightarrow 2}$, which returns the maximum ℓ_2 norm of a column. Define the *hollow Gram matrix*

$$\mathbf{G} = \Phi^* \Phi - \text{diag}(\Phi^* \Phi),$$

which tabulates the inner products between *distinct* columns of Φ . Let \mathbf{R}_Ω be the $W \times W$ orthogonal projector onto the coordinates listed in Ω . Elaborating on these definitions, we see that

$$\mu_2(\Omega) = \|\mathbf{R}_\Omega \mathbf{G} (\mathbf{I} - \mathbf{R}_\Omega)\|_{1 \rightarrow 2} \leq \|\mathbf{R}_\Omega \mathbf{G}\|_{1 \rightarrow 2}. \quad (21)$$

In particular, the cumulative coherence $\mu_2(\Omega)$ is dominated by the maximum column norm of \mathbf{G} .

When Ω is chosen at random, it turns out that we can use this upper bound to improve our estimate of the cumulative coherence $\mu_2(\Omega)$. To incorporate this effect, we need the following result, which is a consequence of Theorem 3.2 of [48] combined with a standard poissonization argument (e.g., see [49, Lem. 14]).

Proposition 10: Fix a $W \times W$ matrix \mathbf{G} . Let \mathbf{R} be an orthogonal projector onto K coordinates, chosen randomly from $\llbracket W \rrbracket$. For $p = 2 \log W$,

$$\mathbb{E}_p \|\mathbf{R} \mathbf{G}\|_{1 \rightarrow 2} \leq 8 \sqrt{\log W} \|\mathbf{G}\|_{\max} + 2 \sqrt{\frac{K}{W}} \|\mathbf{G}\|_{1 \rightarrow 2}.$$

With these facts at hand, we can establish the following bound.

Theorem 11 (Cumulative Coherence): Suppose the sampling rate

$$R \geq C [K \log W + \log^3 W].$$

Draw an $R \times W$ random demodulator Φ . Let Ω be a random set of K coordinates in $\llbracket W \rrbracket$. Then

$$\mathbb{P} \left\{ \mu_2(\Omega) \geq \frac{1}{\sqrt{16 \log W}} \right\} \leq 3W^{-1}.$$

Proof: Under our hypothesis on the sampling rate, Theorem 9 demonstrates that, except with probability W^{-1} ,

$$\|\mathbf{G}\|_{\max} \leq \frac{c}{\log W},$$

where we can make c as small as we like by increasing the constant in the sampling rate. Similarly, Theorem 8 ensures that

$$\max_{\omega} \|\varphi_\omega\|_2 \leq 2,$$

except with probability W^{-1} . We condition on the event F that these two bounds hold. We have $\mathbb{P}\{F^c\} \leq 2W^{-1}$.

On account of the latter inequality and the fact (17) that $\|\Phi\| = \sqrt{W/R}$, we obtain

$$\begin{aligned} \|\mathbf{G}\|_{1 \rightarrow 2} &\leq \|\Phi^* \Phi\|_{1 \rightarrow 2} = \max_{\omega} \|\Phi^* \Phi \mathbf{e}_\omega\|_2 \\ &= \max_{\omega} \|\Phi^* \varphi_\omega\|_2 \leq \|\Phi\| \cdot \max_{\omega} \|\varphi_\omega\|_2 \leq 2 \sqrt{W/R}. \end{aligned}$$

Let \mathbf{R}_Ω be the (random) projector onto the K coordinates listed in Ω . For $p = 2 \log W$, relation (21) and Proposition 10 yield

$$\begin{aligned} \mathbb{E}_p [\mu_2(\Omega) | F] &\leq \mathbb{E}_p [\|\mathbf{R}_\Omega \mathbf{G}\|_{1 \rightarrow 2} | F] \\ &\leq 8 \sqrt{\log W} \cdot \frac{c}{\log W} + 2 \sqrt{\frac{K}{W}} \cdot 2 \sqrt{\frac{W}{R}} \\ &\leq \frac{c'}{\sqrt{\log W}} \end{aligned}$$

under our assumption on the sampling rate. Finally, we apply Markov's inequality to obtain

$$\mathbb{P} \left\{ \mu_2(\Omega) > \frac{c'e^{0.5}}{\sqrt{\log W}} \mid F \right\} \leq e^{-0.5p} = W^{-1}.$$

By selecting the constant in the sampling rate sufficiently large, we can ensure that c' is sufficiently small that

$$\mathbb{P} \left\{ \mu_2(\Omega) > \frac{1}{\sqrt{16 \log W}} \mid F \right\} \leq W^{-1}.$$

We reach the conclusion

$$\mathbb{P} \left\{ \mu_2(\Omega) > \frac{1}{\sqrt{16 \log W}} \right\} \leq W^{-1} + \mathbb{P} \{F^c\} \leq 3W^{-1}$$

when we remove the conditioning. \blacksquare

B. Conditioning of a Random Submatrix

We also require information about the conditioning of a random set of columns drawn from the random demodulator matrix Φ . To obtain this intelligence, we refer to the following result, which is a consequence of [48, Thm. 1] and a poissonization argument.

Proposition 12: Let \mathbf{A} be a $W \times W$ Hermitian matrix, split into its diagonal and off-diagonal parts: $\mathbf{A} = \mathbf{E} + \mathbf{G}$. Draw an orthogonal projector \mathbf{R} onto K coordinates, chosen randomly from $\llbracket W \rrbracket$. For $p = 2 \log W$,

$$\begin{aligned} \mathbb{E}_p \|\mathbf{R}\mathbf{A}\mathbf{R}\| &\leq C \left[\log W \|\mathbf{A}\|_{\max} \right. \\ &\quad \left. + \sqrt{\frac{K \log W}{W}} \|\mathbf{A}\|_{1 \rightarrow 2} + \frac{K}{W} \|\mathbf{A}\| \right] + \|\mathbf{E}\|. \end{aligned}$$

Suppose that Ω is a subset of $\llbracket W \rrbracket$. Recall that Φ_Ω is the column submatrix of Φ indexed by Ω . We have the following result.

Theorem 13 (Conditioning of a Random Submatrix): Suppose the sampling rate

$$R \geq C [K \log W + \log^3 W].$$

Draw an $R \times W$ random demodulator, and let Ω be a random subset of $\llbracket W \rrbracket$ with cardinality K . Then

$$\mathbb{P} \{ \|\Phi_\Omega^* \Phi_\Omega - \mathbf{I}\| \geq 0.5 \} \leq 3W^{-1}.$$

Proof: Define the quantity of interest

$$Q = \|\Phi_\Omega^* \Phi_\Omega - \mathbf{I}\|.$$

Let \mathbf{R}_Ω be the random orthogonal projector onto the coordinates listed in Ω , and observe that

$$Q = \|\mathbf{R}_\Omega(\Phi^* \Phi - \mathbf{I})\mathbf{R}_\Omega\|.$$

Define the matrix $\mathbf{A} = \Phi^* \Phi - \mathbf{I}$. Let us perform some background investigations so that we are fully prepared to apply Proposition 12.

We split the matrix

$$\mathbf{A} = \mathbf{E} + \mathbf{G},$$

where $\mathbf{E} = \text{diag}(\Phi^* \Phi) - \mathbf{I}$ and $\mathbf{G} = \Phi^* \Phi - \text{diag}(\Phi^* \Phi)$ is the hollow Gram matrix. Under our hypothesis on the sampling rate, Theorem 8 provides that

$$\|\mathbf{E}\| = \max_\omega \left| \|\varphi_\omega\|_2^2 - 1 \right| \leq 0.15,$$

except with probability W^{-1} . It also follows that

$$\|\text{diag}(\Phi^* \Phi)\| = \max_\omega \|\varphi_\omega\|_2^2 \leq 1.15.$$

We can bound $\|\mathbf{G}\|_{1 \rightarrow 2}$ by repeating our earlier calculation and introducing the identity (17) that $\|\Phi\| = \sqrt{W/R}$. Thus,

$$\|\mathbf{G}\|_{1 \rightarrow 2} \leq \|\Phi\| \max_\omega \|\varphi_\omega\|_2 \leq \sqrt{\frac{1.15W}{R}}.$$

Since $W/R \geq 1$, it also holds that

$$\|\mathbf{G}\| = \|\Phi^* \Phi - \text{diag}(\Phi^* \Phi)\| \leq \|\Phi\|^2 + 1.15 \leq \frac{2.15W}{R}.$$

Meanwhile, Theorem 9 ensures that

$$\|\mathbf{G}\|_{\max} \leq \frac{c}{\log W},$$

except with probability W^{-1} . As before, we can make c as small as we like by increasing the constant in the sampling rate. We condition on the event F that these estimates are in force. So $\mathbb{P} \{F^c\} \leq 2W^{-1}$.

Now, invoke Proposition 12 to obtain

$$\begin{aligned} \mathbb{E}_p[Q \mid F] &\leq C \left[\log W \cdot \frac{c}{\log W} \right. \\ &\quad \left. + \sqrt{\frac{K \log W}{W}} \cdot \sqrt{\frac{1.15W}{R}} + \frac{K}{W} \cdot \frac{1.15W}{R} \right] + 0.15 \end{aligned}$$

for $p = 2 \log W$. Simplifying this bound, we reach

$$\mathbb{E}_p[Q \mid F] \leq C \left[c + \sqrt{\frac{C'K \log W}{R}} \right] + 0.15.$$

By choosing the constant in the sampling rate R sufficiently large, we can guarantee that

$$\mathbb{E}_p[Q \mid F] \leq 0.3.$$

Markov's inequality now provides

$$\mathbb{P} \{ Q \geq 0.3e^{0.5} \mid F \} \leq e^{-0.5p} = W^{-1}.$$

Note that $0.3e^{0.5} < 0.5$ to conclude that

$$\mathbb{P} \{ Q \geq 0.5 \} \leq W^{-1} + \mathbb{P} \{F^c\} \leq 3W^{-1}.$$

This is advertised conclusion. \blacksquare

C. Recovery of Random Signals

The literature contains powerful results on the performance of ℓ_1 minimization for recovery of random signals that are sparse with respect to an incoherent dictionary. We have finally acquired the keys we need to start this machinery. Our major result for random signals, Theorem 1, is a consequence of the following result, which is an adaptation of [50, Thm. 14].

Proposition 14: Let Φ be an $R \times W$ matrix, and let Ω be a subset of $\llbracket W \rrbracket$ for which

- $\mu_2(\Omega) \leq (16 \log W)^{-1/2}$, and
- $\|(\Phi_\Omega^* \Phi_\Omega)^{-1}\| \leq 2$.

Suppose that \mathbf{s} is a vector that satisfies

- $\text{supp}(\mathbf{s}) \subset \Omega$, and
- $\text{sgn}(s_\omega)$ are i.i.d. uniform on the complex unit circle.

Then \mathbf{s} is the unique solution to the optimization problem

$$\min \|\mathbf{v}\|_1 \quad \text{subject to} \quad \Phi \mathbf{v} = \Phi \mathbf{s} \quad (22)$$

except with probability $2W^{-1}$.

Our main result, Theorem 1, is a corollary of this result.

Corollary 15: Suppose that the sampling rate satisfies

$$R \geq C [K \log W + \log^3 W].$$

Draw an $R \times W$ random demodulator Φ . Let \mathbf{s} be a random amplitude vector drawn according to Model (A). Then the solution $\hat{\mathbf{s}}$ to the convex program (22) equals \mathbf{s} except with probability $8W^{-1}$.

Proof: Since the signal \mathbf{s} is drawn from Model (A), its support Ω is a uniformly random set of K components from $[W]$. Theorem 11 ensures that

$$\mu_2(\Omega) \leq \frac{1}{\sqrt{16 \log W}},$$

except with probability $3W^{-1}$. Likewise, Theorem 13 guarantees

$$\|\Phi_\Omega^* \Phi_\Omega - \mathbf{I}\| \leq 0.5,$$

except with probability $3W^{-1}$. This bound implies that all the eigenvalues of $\Phi_\Omega^* \Phi_\Omega$ lie in the range $[0.5, 1.5]$. In particular,

$$\|(\Phi_\Omega^* \Phi_\Omega)^{-1}\| \leq 2.$$

Meanwhile, Model (A) provides that the phases of the nonzero entries of \mathbf{s} are uniformly distributed on the complex unit circle. We invoke Proposition 14 to see that the optimization problem (22) recovers the amplitude vector \mathbf{s} from the observations $\Phi \mathbf{s}$, except with probability $2W^{-1}$. The total failure probability is $8W^{-1}$. ■

APPENDIX III STABLE RECOVERY

In this appendix, we establish much stronger recovery results under a slightly stronger assumption on the sampling rate. This work is based on the *restricted isometry property* (RIP), which enjoys the privilege of a detailed theory.

A. Background

The RIP is a formalization of the statement that the sampling matrix preserves the norm of all sparse vectors up to a small constant factor. Geometrically, this property ensures that the sampling matrix embeds the set of sparse vectors in a high-dimensional space into the lower-dimensional space of samples. Consequently, it becomes possible to recover sparse vectors from a small number of samples. Moreover, as we will see, this property also allows us to approximate compressible vectors, given relatively few samples.

Let us shift to rigorous discussion. We say that an $R \times W$ matrix Φ has the RIP of order N with restricted isometry

constant $\delta_N \in (0, 1)$ when the following inequalities are in force:

$$\left| \frac{\|\Phi \mathbf{x}\|_2^2 - \|\mathbf{x}\|_2^2}{\|\mathbf{x}\|_2^2} \right| \leq \delta_N \quad \text{whenever} \quad \|\mathbf{x}\|_0 \leq N. \quad (23)$$

Recall that the function $\|\cdot\|_0$ counts the number of nonzero entries in a vector. In words, the definition states that the sampling operator produces a small relative change in the squared ℓ_2 norm of an N -sparse vector. The RIP was introduced by Candès and Tao in an important paper [19].

For our purposes, it is more convenient to rephrase the RIP. Observe that the inequalities (23) hold if and only if

$$\left| \frac{\mathbf{x}^* (\Phi^* \Phi - \mathbf{I}) \mathbf{x}}{\mathbf{x}^* \mathbf{x}} \right| \leq \delta_N \quad \text{whenever} \quad \|\mathbf{x}\|_0 \leq N.$$

The extreme value of this Rayleigh quotient is clearly the largest magnitude eigenvalue of any $N \times N$ principal submatrix of $\Phi^* \Phi - \mathbf{I}$.

Let us construct a norm that packages up this quantity. Define

$$\|\mathbf{A}\| = \sup_{|\Omega| \leq N} \|\mathbf{A}|_{\Omega \times \Omega}\|.$$

In words, the triple-bar norm returns the least upper bound on the spectral norm of any $N \times N$ principal submatrix of \mathbf{A} . Therefore, the matrix Φ has the RIP of order N with constant δ_N if and only if

$$\|\Phi^* \Phi - \mathbf{I}\| \leq \delta_N.$$

Referring back to (20), we may write this relation in the more suggestive form

$$\|\Phi^* \Phi - \mathbb{E} \Phi^* \Phi\| \leq \delta_N.$$

In the next section, we strive to achieve this estimate.

B. RIP for Random Demodulator

Our major result is that the random demodulator matrix has the restricted isometry property when the sampling rate is chosen appropriately.

Theorem 16 (RIP for Random Demodulator): Fix $\delta > 0$. Suppose that the sampling rate

$$R \geq C \delta^{-2} \cdot N \log^6(W).$$

Then an $R \times W$ random demodulator matrix Φ has the restricted isometry property of order N with constant $\delta_N \leq \delta$, except with probability $O(W^{-1})$.

There are three basic reasons that the random demodulator matrix satisfies the RIP. First, its gram matrix averages to the identity. Second, its rows are independent. Third, its entries are uniformly bounded. The proof shows how these three factors interact to produce the conclusion.

In the sequel, it is convenient to write $\mathbf{z} \otimes \mathbf{z}$ for the rank-one matrix $\mathbf{z} \mathbf{z}^*$. We also number constants for clarity.

Most of the difficulty of argument is hidden in a lemma due to Rudelson and Vershynin [20, Lem. 3.6].

Lemma 17 (Rudelson–Vershynin): Suppose that $\{\mathbf{z}_r\}$ is a sequence of R vectors in \mathbb{C}^W where $R \leq W$, and assume

that each vector satisfies the bound $\|z_r\|_\infty \leq B$. Let $\{\xi_r\}$ be an independent Rademacher series. Then

$$\mathbb{E} \left\| \sum_r \xi_r z_r \otimes z_r \right\| \leq \beta \left\| \sum_r z_r \otimes z_r \right\|^{1/2},$$

where

$$\beta \leq C_1 B \sqrt{N} \log^2(W).$$

The next lemma employs this bound to show that, in expectation, the random demodulator has the RIP.

Lemma 18: Fix $\delta \in (0, 1)$. Suppose that the sampling rate

$$R \geq C_2 \delta^{-2} \cdot N \log^5 W.$$

Let Φ be an $R \times W$ random demodulator. Then

$$\mathbb{E} \|\Phi^* \Phi - \mathbf{I}\| \leq \delta.$$

Proof: Let z_r^* denote the r th row of Φ . Observe that the rows are mutually *independent* random vectors, although their distributions differ. Recall that the entries of each row are uniformly bounded, owing to Lemma 5. We may now express the Gram matrix of Φ as

$$\Phi^* \Phi = \sum_{r=1}^R z_r \otimes z_r.$$

We are interested in bounding the quantity

$$\begin{aligned} E &= \mathbb{E} \|\Phi^* \Phi - \mathbf{I}\| = \mathbb{E} \left\| \sum_r z_r \otimes z_r - \mathbf{I} \right\| \\ &= \mathbb{E} \left\| \sum_r (z_r \otimes z_r - \mathbb{E} z_r' \otimes z_r') \right\|, \end{aligned}$$

where $\{z_r'\}$ is an independent copy of $\{z_r\}$.

The symmetrization result, Proposition 4, yields

$$E \leq 2 \mathbb{E} \left\| \sum_r \xi_r z_r \otimes z_r \right\|,$$

where $\{\xi_r\}$ is a Rademacher series, independent of everything else. Conditional on the choice of $\{z_r\}$, the Rudelson–Vershynin lemma results in

$$E \leq 2 \mathbb{E} \left(\beta \cdot \left\| \sum_r z_r \otimes z_r \right\|^{1/2} \right),$$

where

$$\beta \leq C_1 B \sqrt{N} \log^2 W$$

and $B = \max_{k,\omega} |\varphi_{k\omega}|$ is a random variable.

According to Lemma 5,

$$(\mathbb{E} \beta^2)^{1/2} \leq \sqrt{6} C_1 \sqrt{\frac{N \log^5(W)}{R}}.$$

We may now apply the Cauchy–Schwarz inequality to our bound on E to reach

$$E \leq 2\sqrt{6} C_1 \sqrt{\frac{N \log^5 W}{R}} \left(\mathbb{E} \left\| \sum_r z_r \otimes z_r \right\| \right)^{1/2}.$$

Add and subtract an identity matrix inside the norm, and invoke the triangle inequality. Note that $\|\mathbf{I}\| = 1$, and identify a copy of E on the right-hand side:

$$\begin{aligned} E &\leq \sqrt{\frac{24C_1^2 N \log^5 W}{R}} \left(\mathbb{E} \left\| \sum_r z_r \otimes z_r - \mathbf{I} \right\| + \|\mathbf{I}\| \right)^{1/2} \\ &= \sqrt{\frac{24C_1^2 N \log^5 W}{R}} (E + 1)^{1/2}. \end{aligned}$$

Solutions to the relation $E \leq \gamma(E + 1)^{1/2}$ obey $E \leq 2\gamma$ whenever $\gamma \leq 1$.

We conclude that

$$E \leq \sqrt{\frac{C_2 N \log^5 W}{R}}$$

whenever the fraction is smaller than one. To guarantee that $E \leq \delta$, then, it suffices that

$$R \geq C_2 \delta^{-2} \cdot N \log^5 W.$$

This point completes the argument. \blacksquare

To finish proving the theorem, it remains to develop a large deviation bound. Our method is the same as that of Rudelson and Vershynin: We invoke a concentration inequality for sums of independent random variables in a Banach space. See [20, Thm. 3.8], which follows from [44, Thm. 6.17].

Proposition 19: Let Y_1, \dots, Y_R be independent, symmetric random variables in a Banach space X , and assume each random variable satisfies the bound $\|Y_r\|_X \leq B$ almost surely. Let $Y = \left\| \sum_r Y_r \right\|_X$. Then

$$\mathbb{P}\{Y > C_3 [u \mathbb{E} Y + tB]\} \leq e^{-u^2} + e^{-t}$$

for all $u, t \geq 1$.

In words, the norm of a sum of independent, symmetric, bounded random variables has a tail bound consisting of two parts. The norm exhibits subgaussian decay with “standard deviation” comparable to its mean, and it exhibits subexponential decay controlled by the size of the summands.

Proof: [Theorem 16] We seek a tail bound for the random variable

$$Z = \left\| \sum_r z_r \otimes z_r - \mathbf{I} \right\| = \left\| \sum_r (z_r \otimes z_r - \mathbb{E} z_r' \otimes z_r') \right\|.$$

As before, z_r^* is the r th row of Φ and z_r' is an independent copy of z_r .

To begin, we express the constant in the stated sampling rate as $C = C_2 \cdot c^{-2}$, where c will be adjusted at the end of the argument. Therefore, the sampling rate may be written as

$$R \geq C_2 \left(\frac{c\delta}{\sqrt{\log W}} \right)^{-2} N \log^5 W.$$

Lemma 18 implies that

$$\mathbb{E} Z \leq \frac{c\delta}{\sqrt{\log W}} \leq c\delta.$$

As described in the background section, we can produce a tail bound for Z by studying the symmetrized random variable

$$Y = \left\| \sum_r (z_r \otimes z_r - z_r' \otimes z_r') \right\|,$$

For future reference, note that

$$\mathbb{E} Y \leq 2 \mathbb{E} Z \leq \frac{2c\delta}{\sqrt{\log W}},$$

where the first inequality follows from the triangle inequality and identical distribution.

Next, observe that Y is the norm of a sum of independent, symmetric random variables in a Banach space. The tail bound of Proposition 19 requires the summands to be bounded in

norm. To that end, we demand that the vectors \mathbf{z}_r and \mathbf{z}'_r are uniformly bounded. Define the event

$$F = \left\{ \max_r \left\{ \|\mathbf{z}_r\|_\infty^2, \|\mathbf{z}'_r\|_\infty^2 \right\} \leq \frac{10 \log W}{R} \right\}.$$

In other terms, F is the event that two independent random demodulator matrices both have bounded entries. Therefore,

$$\mathbb{P}\{F^c\} \leq 2W^{-1}$$

after two applications of Lemma 5.

Under the event F , we have a hard bound on the norm of each summand in Y .

$$\begin{aligned} B &= \max_r \|\mathbf{z}_r \otimes \mathbf{z}_r - \mathbf{z}'_r \otimes \mathbf{z}'_r\| \\ &\leq \max_r (\|\mathbf{z}_r \otimes \mathbf{z}_r\| + \|\mathbf{z}'_r \otimes \mathbf{z}'_r\|). \end{aligned}$$

For each r , we may compute

$$\begin{aligned} \|\mathbf{z}_r \otimes \mathbf{z}_r\| &= \sup_{|\Omega| \leq N} \left\| (\mathbf{z}_r \otimes \mathbf{z}_r)|_{\Omega \times \Omega} \right\| \\ &= \sup_{|\Omega| \leq N} \|\mathbf{z}_r|_\Omega\|_2^2 \leq N \|\mathbf{z}_r\|_\infty^2 \leq \frac{10N \log W}{R}, \end{aligned}$$

where the last inequality follows from F . An identical estimate holds for the terms involving \mathbf{z}'_r . Therefore,

$$B \leq \frac{20N \log W}{R}.$$

Our hypothesized lower bound for the sampling rate R gives

$$B \leq 20N \log W \cdot \frac{(c\delta/\sqrt{\log W})^2}{C_2 N \log^5 W} \leq \frac{c\delta}{\log W},$$

where the second inequality certainly holds if c is a sufficiently small constant.

We now apply Proposition 19 to the symmetric variable Y , conditional on the event F . The bound reads

$$\mathbb{P}\{Y > C_3(u\mathbb{E}Y + tB) \mid F\} \leq e^{-u^2} + e^{-t}.$$

Select $u = \sqrt{\log W}$ and $t = \log W$, and recall the bounds on $\mathbb{E}Y$ and on B to obtain

$$\mathbb{P}\{Y > C_3(2c\delta + c\delta) \mid F\} \leq 2W^{-1}.$$

In summary,

$$\begin{aligned} \mathbb{P}\{Y > 3C_3c\delta\} &\leq \mathbb{P}\{Y > 3C_3c\delta \mid F\} + \mathbb{P}\{F^c\} \\ &\leq 4W^{-1}. \end{aligned}$$

The tail bound for Y yields a tail bound for Z via relation (15):

$$\mathbb{P}\{Z > 2\mathbb{E}Z + u\} \leq 2\mathbb{P}\{Y > u\}.$$

As noted, $\mathbb{E}Z \leq c\delta \leq C_3c\delta$. Therefore,

$$\begin{aligned} \mathbb{P}\{Z > 5C_3c\delta\} &\leq \mathbb{P}\{Z > 2\mathbb{E}Z + 3C_3c\delta\} \\ &\leq 2\mathbb{P}\{Y > 3C_3c\delta\} \\ &\leq 8W^{-1}. \end{aligned}$$

To complete the proof, we select $c \leq (5C_3)^{-1}$. ■

C. Signal Recovery under the RIP

When the sampling matrix has the restricted isometry property, the samples contain enough information to approximate general signals extremely well. Candès, Romberg, and Tao have shown that signal recovery can be performed by solving a convex optimization problem. This approach produces exceptional results in theory and in practice [18]

Proposition 20: Suppose that Φ is an $R \times W$ matrix that verifies the RIP of order $2K$ with restricted isometry constant $\delta_{2K} \leq c$. Then the following statement holds. Consider an arbitrary vector \mathbf{s} in \mathbb{C}^W , and suppose we collect noisy samples

$$\mathbf{y} = \Phi \mathbf{s} + \boldsymbol{\nu}$$

where $\|\boldsymbol{\nu}\|_2 \leq \eta$. Every solution $\hat{\mathbf{s}}$ to the optimization problem

$$\min \|\mathbf{v}\|_1 \quad \text{subject to} \quad \|\Phi \mathbf{v} - \mathbf{y}\|_2 \leq \eta, \quad (24)$$

approximates the target signal:

$$\|\hat{\mathbf{s}} - \mathbf{s}\|_2 \leq C \max \left\{ \eta, \frac{1}{\sqrt{K}} \|\mathbf{s} - \mathbf{s}_K\|_1 \right\},$$

where \mathbf{s}_K is a best K -sparse approximation to \mathbf{s} with respect to the ℓ_1 norm.

An equivalent guarantee holds when the approximation $\hat{\mathbf{s}}$ is computed using the CoSaMP algorithm [15, Thm. A].

Combining Proposition 20 with Theorem 16, we obtain our major result, Theorem 2.

Corollary 21: Suppose that the sampling rate

$$R \geq CK \log^6 W.$$

An $R \times W$ random demodulator matrix verifies the RIP of order $2K$ with constant $\delta_{2K} \leq c$, except with probability $O(W^{-1})$. Thus, the conclusions of Proposition 20 are in force.

REFERENCES

- [1] A. V. Oppenheim, R. W. Schaffer, and J. R. Buck, *Discrete-Time Signal Processing*, Prentice Hall, 2nd edition, 1999.
- [2] B. Le, T. W. Rondeau, J. H. Reed, and C. W. Bostian, "Analog-to-digital converters: A review of the past, present, and future," *IEEE Signal Proc. Mag.*, vol. 22, no. 6, pp. 69–77, Nov. 2005.
- [3] D. Healy, "Analog-to-information," 2005, BAA #05-35. Available from <http://www.darpa.mil/mto/solicitations/baa05-35/s/index.html>.
- [4] H. Landau, "Sampling, data transmission, and the Nyquist rate," *Proc. IEEE*, vol. 65, no. 10, pp. 1701–1706, 1967.
- [5] S. Mallat, *A Wavelet Tour of Signal Processing*, Academic Press, London, 2nd edition, 1999.
- [6] M. Matsumoto and T. Nishimura, "Mersenne Twister: A 623-dimensionally equidistributed uniform pseudorandom number generator," *ACM Trans. Modeling and Computer Simulation*, vol. 8, no. 1, pp. 3–30, Jan. 1998.
- [7] Robert H. Walden, "Analog-to-digital converter survey and analysis," *IEEE Journal on Selected Areas in Communications*, vol. 17, no. 4, pp. 539–550, April 1999.
- [8] S. S. Chen, D. L. Donoho, and M. A. Saunders, "Atomic decomposition by Basis Pursuit," *SIAM J. Sci. Comput.*, vol. 20, no. 1, pp. 33–61, 1999.
- [9] Å. Björck, *Numerical Methods for Least Squares Problems*, SIAM, Philadelphia, 1996.
- [10] I. Daubechies, R. DeVore, M. Fornasier, and S. Güntürk, "Iteratively reweighted least squares minimization for sparse recovery," *Comm. Pure Appl. Math.*, 2009, To appear.
- [11] E. Candès, J. Romberg, and T. Tao, "Robust uncertainty principles: Exact signal reconstruction from highly incomplete Fourier information," *IEEE Trans. Inform. Theory*, vol. 52, no. 2, pp. 489–509, Feb. 2006.

- [12] J. A. Tropp and A. C. Gilbert, "Signal recovery from random measurements via orthogonal matching pursuit," *IEEE Trans. Inform. Theory*, vol. 53, no. 12, pp. 4655–4666, 2007.
- [13] D. Needell and R. Vershynin, "Uniform uncertainty principle and signal recovery via regularized orthogonal matching pursuit," *Foundations of Comput. Math.*, 2009, To appear.
- [14] D. Needell and R. Vershynin, "Signal recovery from incomplete and inaccurate measurements via regularized orthogonal matching pursuit," Submitted for publication, October 2007.
- [15] D. Needell and J. A. Tropp, "CoSaMP: Iterative signal recovery from incomplete and inaccurate samples," *Appl. Comp. Harmonic Anal.*, 2009, To appear.
- [16] D. L. Donoho, "High-dimensional centrally-symmetric polytopes with neighborliness proportional to dimension," *Disc. Comput. Geometry*, vol. 35, no. 4, pp. 617–652, 2006.
- [17] R. Baraniuk, M. Davenport, R. DeVore, and M. Wakin, "A simple proof of the restricted isometry property for random matrices," *Constructive Approx.*, vol. 28, no. 3, pp. 253–263, Dec. 2008.
- [18] E. J. Candès, J. Romberg, and T. Tao, "Stable signal recovery from incomplete and inaccurate measurements," *Comm. Pure Appl. Math.*, vol. 59, pp. 1207–1223, 2006.
- [19] E. J. Candès and T. Tao, "Near-optimal signal recovery from random projections: Universal encoding strategies?," *IEEE Trans. Inform. Theory*, vol. 52, no. 12, pp. 5406–5425, Dec. 2006.
- [20] M. Rudelson and R. Vershynin, "Sparse reconstruction by convex relaxation: Fourier and Gaussian measurements," in *Annual Conference on Information Sciences and Systems (CISS)*, Princeton, NJ, Mar. 2006.
- [21] A. Cohen, W. Dahmen, and R. DeVore, "Compressed sensing and best k -term approximation," *J. Amer. Math. Soc.*, vol. 22, no. 1, pp. 211–231, Jan. 2009.
- [22] D. L. Donoho, M. Vetterli, R. A. DeVore, and I. Daubechies, "Data compression and harmonic analysis," *IEEE Trans. Inform. Theory*, vol. 44, no. 6, pp. 2433–2452, October 1998.
- [23] E. Candès and J. Romberg, "Sparsity and incoherence in compressive sampling," *Inverse Problems*, vol. 23, no. 3, pp. 969–985, 2007.
- [24] 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.
- [25] H. Landau, "Necessary density conditions for sampling and interpolation of certain entire functions," *Acta Math.*, vol. 117, pp. 37–52, 1967.
- [26] Ping Feng and Y. Bresler, "Spectrum-blind minimum-rate sampling and reconstruction of multiband signals," *IEEE Int. Conf. Acoustics, Speech and Signal Proc. (ICASSP)*, vol. 3, pp. 1688–1691 vol. 3, May 1996.
- [27] R. Venkataramani and Y. Bresler, "Perfect reconstruction formulas and bounds on aliasing error in sub-Nyquist nonuniform sampling of multiband signals," *IEEE Trans. Inform. Theory*, vol. 46, no. 6, pp. 2173–2183, Sep 2000.
- [28] M. Mishali and Y. Eldar, "Blind multi-band signal reconstruction: Compressed sensing for analog signals," *IEEE Trans. Signal Proc.*, 2007, Submitted for publication.
- [29] Y. Eldar, "Compressed sensing for analog signals," *IEEE Trans. Signal Proc.*, 2008, Submitted for publication.
- [30] E. Kushilevitz and Y. Mansour, "Learning decision trees using the Fourier spectrum," *SIAM J. Comput.*, vol. 22, no. 6, pp. 1331–1348, 1993.
- [31] A. C. Gilbert, S. Guha, P. Indyk, S. Muthukrishnan, and M. J. Strauss, "Near-optimal sparse Fourier representations via sampling," in *Proc. ACM Symposium on Theory of Computing*, Montreal, Canada, May 2002.
- [32] 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.
- [33] J. N. Laska, S. Kirolos, Y. Massoud, R. G. Baraniuk, A. C. Gilbert, M. Iwen, and M. J. Strauss, "Random sampling for analog-to-information conversion of wideband signals," in *Proc. IEEE Dallas Circuits and Systems Workshop (DCAS)*, Dallas, TX, Oct. 2006.
- [34] T. Ragheb, S. Kirolos, J. Laska, A. Gilbert, M. Strauss, R. Baraniuk, and Y. Massoud, "Implementation models for analog-to-information conversion via random sampling," in *Midwest Symp. Circuits and Systems (MWSCAS)*, Aug. 2007, pp. 325–328.
- [35] M. Vetterli, P. Marziliano, and T. Blu, "Sampling signals with finite rate of innovation," *IEEE Trans. Signal Proc.*, vol. 50, no. 6, pp. 1417–1428, Jun 2002.
- [36] Pier Luigi Dragotti, Martin Vetterli, and Thierry Blu, "Sampling Moments and Reconstructing Signals of Finite Rate of Innovation: Shannon Meets Strang-Fix," *IEEE Trans. Signal Proc.*, vol. 55, no. 5, Part 1, pp. 1741–1757, 2007.
- [37] T. Blu, P.-L. Dragotti, M. Vetterli, P. Marziliano, and L. Coulot, "Sparse sampling of signal innovations," *IEEE Signal Proc. Mag.*, vol. 25, no. 2, pp. 31–40, March 2008.
- [38] J. Berent and P.L. Dragotti, "Perfect reconstruction schemes for sampling piecewise sinusoidal signals," *IEEE Int. Conf. Acoustics, Speech and Signal Proc. (ICASSP)*, vol. 3, pp. III–III, May 2006.
- [39] S. Kirolos, J. N. Laska, M. B. Wakin, M. F. Duarte, D. Baron, T. Ragheb, Y. Massoud, and R. G. Baraniuk, "Analog to information conversion via random demodulation," in *IEEE Dallas Circuits and Systems Workshop (DCAS)*, Dallas, TX, Oct. 2006.
- [40] J. N. Laska, S. Kirolos, M. F. Duarte, T. Ragheb, R. G. Baraniuk, and Y. Massoud, "Theory and implementation of an analog-to-information conversion using random demodulation," in *IEEE Int. Symposium on Circuits and Systems (ISCAS)*, New Orleans, LA, May 2007.
- [41] M. Mishali, Y. Eldar, and J. A. Tropp, "Efficient sampling of sparse wideband analog signals," in *IEEE Conv. Electrical and Electronic Engineers in Israel (IEEEI)*, Eilat, Dec. 2008.
- [42] D. L. Donoho, "Compressed sensing," *IEEE Trans. Inform. Theory*, vol. 52, no. 4, pp. 1289–1306, Apr. 2006.
- [43] U. Haagerup, "The best constants in the Khintchine inequality," *Studia Math.*, vol. 70, pp. 231–283, 1981.
- [44] M. Ledoux and M. Talagrand, *Probability in Banach Spaces: Isoperimetry and Processes*, Springer, 1991.
- [45] J. Bourgain and L. Tzafriri, "Invertibility of "large" submatrices with applications to the geometry of Banach spaces and harmonic analysis," *Israel J. Math.*, vol. 57, no. 2, pp. 137–224, 1987.
- [46] R. Adamczak, "Logarithmic Sobolev inequalities and concentration of measure for convex functions and polynomial chaoses," *Bull. Pol. Acad. Sci. Math.*, vol. 53, no. 2, pp. 221–238, 2005.
- [47] T. Strohmer and R. W. Heath, "Grassmannian frames with applications to coding and communication," *Appl. Comp. Harmonic Anal.*, vol. 14, no. 3, pp. 257–275, May 2003.
- [48] J. A. Tropp, "Norms of random submatrices and sparse approximation," *C. R. Acad. Sci. Paris Ser. I Math.*, vol. 346, pp. 1271–1274, 2008.
- [49] J. A. Tropp, "On the linear independence of spikes and sines," *J. Fourier Anal. Appl.*, vol. 14, no. 838–858, 2008.
- [50] J. A. Tropp, "On the conditioning of random subdictionaries," *Appl. Comp. Harmonic Anal.*, vol. 25, pp. 1–24, 2008.