APPLICATION OF COMPRESSIVE SENSING TO THE DESIGN OF WIDEBAND SIGNAL ACQUISITION RECEIVERS

John Treichler

Applied Signal Technology, Inc.
Sunnyvale, California

Mark Davenport, Richard Baraniuk

Rice University
Houston, Texas

ABSTRACT

Compressive sensing (CS) exploits the sparsity present in many signals to reduce the number of measurements needed for digital acquisition. With this reduction would come, in theory, commensurate reductions in the size, weight, power consumption, and/or monetary cost of both signal sensors and any associated communication links. This paper examines the use of CS in environments where the input signal takes the form of a sparse combination of narrowband signals of unknown frequencies that appear anywhere in a broad spectral band. We formulate the problem statement for such a receiver and establish a reasonable set of requirements that a receiver should meet to be practically useful. The performance of a CS receiver for this application is then evaluated in two ways: using the applicable (and still evolving) CS theory and using a set of computer simulations carefully constructed to compare the CS receiver against the performance expected from a conventional implementation. This sets the stage for work in a sequel that will use these results to produce comparisons of the size, weight, and power consumption of a CS receiver against an exemplar of a conventional design.

1. INTRODUCTION

Compressive sensing (CS) [1–3] exploits the sparsity present in many signals to reduce the number of measurements needed for acqusition. It has been shown theoretically that, under the right set of circumstances, CS can dramatically reduce the number of measurements needed to detect, characterize, and/or extract signals, and therefore can reduce by the same factor the storage and/or transmission rate needed to handle the signal at the sensing point. Conversely, signals with much larger bandwidths could be accepted by existing acquisition systems. If these reductions were found to proportionally reduce the size, weight, and power consumption (SWAP) and cost of operational signal acquisition systems, then the practical impact could be transformative.

This paper examines the potential practicality of CS to build signal acquisition receivers for the specific, but important, case where the receiver's input signal takes the form of a sparse combination of narrowband signals of unknown frequencies that can appear anywhere in a broad spectral band. Our objective here is to examine the specific application of the acquisition receiver, thus providing the opportunity to test the robustness of CS techniques to imperfect match with its underlying theoretical assumptions.

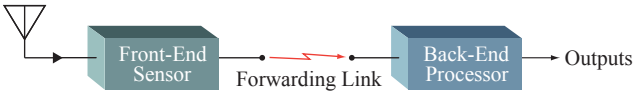

Fig. 1. A wideband signal acquisition receiver.

Section 2 formulates the problem statement for the signal receiver and establishes a set of requirements that a receiver should meet to be highly attractive for practical use. Section 3 applies the existing CS theory to determine how well a CS-based receiver should work in this case, while Section 4 presents the interim results of a set of simulation-based evaluations designed to measure the expected performance of a CS-based receiver in comparison with theory and with a conventional receiver design. These results are discussed in Section 5, and recommendations for additional study and investigation appear in Section 6.

2. PUTATIVE REQUIREMENTS

Our objective in this paper is to explore the attributes and capabilities of CS by examining how it might be applied to meet a specific set of requirements. The particular application addressed is a wideband radio frequency (RF) signal acquisition receiver, a device commonly used in both commercial and military systems to monitor a wide band of radio frequencies for the purposes of *(i)* detecting the presence of signals, *(ii)* characterizing them, and, where appropriate, *(iii)* extracting a specific signal from the several that might be present within that band. A high-level system diagram incorporating such a receiver is shown in Figure 1. While many types of acquisition receivers have been designed, built, and sold over the years, we will choose here a set of putative requirements for such a receiver so that comparisons and analysis can be done. The reader is free to, and in fact invited to, choose

Table 1. A putative set of specifications for an advanced RF signal acquisition receiver.

Instantaneous bandwidth	B	500 MHz
Instantaneous dynamic range	D	96 dB
SNR degradation/noise figure	NF	12 dB
Maximum signal bandwidth	W	200 kHz

other operational parameters and repeat the comparison.

The attributes that characterize an acquisition receiver typically fall into two categories: technical specifications — such as instantaneous bandwidth — and various "costs"— such as SWAP and monetary cost. In this paper we will address only the few most important technical specifications:

- *Instantaneous bandwidth* — the RF range over which signals will be accepted by the receiver and handled with their full fidelity;
- *Instantaneous dynamic range* — the ratio of the maximum to minimum signal power level with which received signals can be handled with full fidelity;
- *SNR degradation* — sometimes termed "noise figure", a measure of the tendency of the receiver to lower the input signal-to-noise ratio (SNR) of a received signal, usually measured in dB. The root cause of this degradation has historically depended on the technology used to build the receiver.
- *Maximum signal bandwidth* — the maximum combined bandwidth of the constituent signals in the acquisition bandwidth of the receiver

These requirements must be met subject to many constraints, including, at least, SWAP and monetary cost. There are also typically system-level constraints, such as the bandwidth available for communicating what the receiver has discovered to other assets or a central processing facility.

Historically RF signal acquisition receivers were first built using purely analog technology, then, more recently, with analog technology conditioning the signal environment sufficiently to employ a high-rate analog-to-digital converter (ADC) followed by digital processing, storage, and/or transmission. When and if it can be applied, CS offers the promise to *(i)* increase the instantaneous input bandwidth, *(ii)* lower all of the cost attributes, and *(iii)* move the computationally intensive portions of the acquisition process away from the sensor and toward a central processing facility.

For the purposes of the comparisons to be made in this paper, we will assume a set of requirements, listed in Table 1, for an acquisition system that are quite audacious and would stress, at the least, conventional implementations. To meet the bandwidth and dynamic range requirements, conventional designs would typically be forced to use techniques based on scanning narrowband receivers across the band. If CS-based systems can be shown to work in such cases without the need for scanning at the receiver, then they would have broad application.

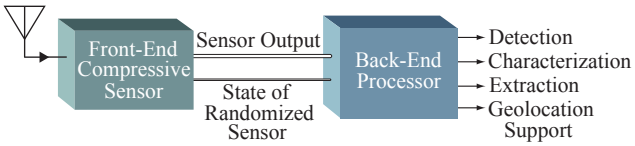

Fig. 2. The "processing asymmetry" assumed in a CS signal acquisition receiver.

We make two last, but important, assumptions:

1. *Signal sparsity* — In order to meet the first-order assumption of all CS techniques that the input signal be sparse in some way, we assume that the sum of the bandwidths of all signals present in the full acquisition band, which we denote W, is no more than 200 kHz. Note that this is significantly smaller than the instantaneous bandwidth B of 500 MHz. Thus we are assuming that the RF input to the receiver is quite sparse in the frequency domain (the instantaneous bandwidth is only $1/2500$ occupied in this case). While inputs with this level of spectral sparsity are not common, they exist often enough to make a solution useful if it can be found. To test the impact of the sparsity assumption for this application, we will evaluate the performance, both theoretical and in simulation, for both the case where the input is noise-free, in which case the input signal is truly sparse, and in the more practical case where the input is contaminated with additive white noise.

2. *Processing asymmetry* — Our presumption is that there is no cost to the computation done to process the output of the compressive sensor. Our objective is to minimize all costs, including the transmission bandwidth between the sensor and the processor, and, for the purposes of this paper, we are prepared to do as much processing as needed to detect, characterize, and/or recover the signal of interest. This assumed separation of functions is illustrated in Figure 2.

3. COMPRESSIVE SENSING AND ITS APPLICATION TO WIDEBAND RADIO RECEIVERS

3.1. CS theory

In laying out the putative requirements in Section 2, we assumed that the real-valued, continuous-time signal of interest, which we will denote by $x(t)$, has instantaneous bandwidth B and maximum signal bandwidth W. Thus, each segment of

$x(t)$ of duration $T = 1$ seconds has the Fourier representation

$$x(t) = \mathbf{\Psi}(\boldsymbol{\alpha}) = \sum_{k=0}^{2B-1} \alpha_k \psi_k(t), \tag{1}$$

where $\psi_k(t) = e^{j2\pi kt/(2B)}$ are the Fourier basis functions and where $\boldsymbol{\alpha} = [\alpha_0, \alpha_1, \ldots, \alpha_{2B-1}]$ is a complex-valued vector of length $2B$ that has $2W$ nonzeros corresponding to the active frequencies of the constituent signals in the acquisition bandwidth B.[1] For the moment, we assume that $\alpha_k = 0$ exactly outside of these $2W$ nonzeros, in which case we say that the vector $\boldsymbol{\alpha}$ is $2W$-*sparse*. Finally, note that while we focus on Fourier-sparse signals in this paper, the CS theory is general and extends to signals sparse in arbitrary bases.

The Shannon-Nyquist sampling theorem tells us that $2B$ uniform samples of $x(t)$ per $T = 1$ second contain all of the information in $x(t)$. Our goal in CS is to do better: to acquire $x(t)$ via a set of $2B/Q$ measurements with $Q \geq 1$ as large as possible. The *subsampling factor* Q strongly affects the various costs (i.e., SWAP and monetary cost) described in Section 2. Observe that if the locations of the $2W$ non-zeros of $\boldsymbol{\alpha}$ are known *a priori*, then by filtering and decimation, we could drive Q as large as $Q_f = B/W$. Our aim is to show that CS-based techniques will enable us to acquire $x(t)$ via a set of $2B/Q$ fixed, nonadaptive, linear measurements that require no *a priori* knowledge of the locations of the non-zeros of $\boldsymbol{\alpha}$, with Q nearly as large as Q_f.

Towards this end, we take the measurements

$$\boldsymbol{y} = \mathbf{\Phi}(x(t)) + \boldsymbol{e}, \tag{2}$$

where $\mathbf{\Phi}$ is a linear *measurement operator* that maps continuous-time functions defined on the time interval $t \in [0, 1]$ to a length $2B/Q$ vector $\boldsymbol{y}$ of measurements, and where $\boldsymbol{e}$ is a length $2B/Q$ vector that represents *measurement noise* generated by the acquisition hardware. The central theoretical question in CS is how to design the measurement operator $\mathbf{\Phi}$ to ensure that we will be able to recover the signal $x(t)$ from the measurements $\boldsymbol{y}$. While there are many approaches to solving this problem, the most common method is to split this question into two parts: *(i)* What properties of $\mathbf{\Phi}$ will ensure that there exists *some* algorithm that can recover $x(t)$ from $\boldsymbol{y}$? and *(ii)* What algorithms can perform this recovery in an efficient manner?

The answer to the first question is rather intuitive. While alternative properties have been studied, the majority of the work in CS assumes that the measurement operator $\mathbf{\Phi}$ satisfies the so-called *restricted isometry property* (RIP) [4]. In our setting, where $x(t)$ has a $2W$-sparse representation with respect to the basis $\mathbf{\Psi}$, the RIP requires that there exists a constant $\delta \in (0, 1)$ such that

$$\sqrt{1-\delta} \leq \frac{\|\mathbf{\Phi}(\mathbf{\Psi}(\boldsymbol{\alpha})) - \mathbf{\Phi}(\mathbf{\Psi}(\boldsymbol{\beta}))\|_2}{\|\boldsymbol{\alpha} - \boldsymbol{\beta}\|_2} \leq \sqrt{1+\delta}, \tag{3}$$

holds for all $2W$-sparse $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$. In words, $\mathbf{\Phi\Psi}$ preserves the Euclidean distance between vectors that are $2W$-sparse. Intuitively, the RIP means that for any particular set of measurements $\boldsymbol{y}$, there exists at most one possible $2W$-sparse signal consistent with these measurements, since if there were two distinct $2W$-sparse signals that mapped to the same set of measurements, then this would violate the lower inequality of (3). In principle, this means that it should be possible (in the noise-free setting at least) to exactly recover any $2W$-sparse signal. Furthermore, if a small amount of measurement noise is added to the measurements $\boldsymbol{y}$ as in (2), then the RIP provides a guarantee that any $2W$-sparse signal consistent with the perturbed measurements will be close to the original signal, and so the RIP ensures that the system has a degree of stability and robustness to measurement noise.

We now consider how to design an operator $\mathbf{\Phi}$ satisfying the RIP. An important result from the CS theory is that for any given basis $\mathbf{\Psi}$ (not just the Fourier basis we focus on here), if we draw a random matrix $\boldsymbol{R}$ of size $2B/Q \times 2B$ whose entries r_{ij} are independent realizations from a Gaussian, Rademacher (± 1-valued), or more generally, any bounded, zero-mean distribution, then with overwhelmingly high probability $\mathbf{\Phi} = \boldsymbol{R}\mathbf{\Psi}^{-1}$ will satisfy (3) for $2W$-sparse signals provided that

$$Q \leq Q_c = \kappa_0 \frac{Q_f}{\ln Q_f} \tag{4}$$

where $\kappa_0 < 1$ is a constant that depends on B and the probability with which (3) holds [5]. From this we conclude that CS-based measurement operators $\mathbf{\Phi}$ pay a small penalty, quantified by $\kappa_0/\ln(Q_f)$, for not exploiting any *a priori* knowledge of the locations of the nonzero frequencies.

In general, the theoretical analysis is somewhat lacking in terms of the precise value of κ_0. For instance, an asymptotic analysis of the results in [5] would suggest that $\kappa_0 \approx 1/50$ would be sufficient, but this value seems to be far more conservative than what is required in practice. Thus, if a specific value for κ_0 is required, one must determine this value experimentally. This is typically accomplished via Monte Carlo simulations that identify how many measurements are sufficient to ensure exact recovery in the noise-free setting on at least, say, 99% of trials. As an example, it is shown in [6] that $Q <\approx 0.6Q_f/\ln(Q_f)$ is sufficient to enable exact recovery in the noise-free setting. Note that this is not the same as demonstrating that $Q <\approx 0.6Q_f/\ln(Q_f)$ is sufficient to ensure that $\mathbf{\Phi}$ satisfies the RIP, but it is highly suggestive that the true value of κ_0 is much greater than the conservative estiamtes provided by the theory. We will observe this phenomenon for ourselves in Section 4.

Since the random matrix approach is somewhat impractical to build in hardware, several hardware architectures have

[1] We set the segment-length to $T = 1$ second for notational convenience; for snippets of other lengths one can merely replace B and W below with BT and WT, *mutatis mutandis*.

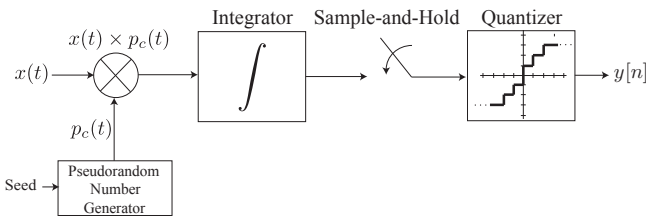

Fig. 3. Random demodulator for obtaining compressive measurements of analog signals.

been implemented and/or proposed that enable compressive samples to be acquired in practical settings. Examples include the random demodulator [6], random filtering [7], and random convolution [8, 9].

We briefly describe the random demodulator as an example of such a system; see Figure 3 for a block diagram. The four key components are a pseudo-random ± 1 "chipping sequence" $p_c(t)$ operating at the Nyquist rate $2B$ or higher, a low pass filter, represented by an ideal integrator with reset, and a low-rate ADC consisting of a sample-and-hold circuit and a quantizer. The input analog signal $x(t)$ is modulated by the chipping sequence and integrated. The output of the integrator is sampled and quantized, and the integrator is reset after each sample. The output measurements from the sample-and-hold circuit are then quantized. The subsampling factor Q_c determines the sampling rate C at which the sample-and-hold circuit and quantizer operate. Specifically, we must have $C \geq 2B/Q_c$.

Mathematically, the linear measurement operator $\boldsymbol{\Phi}$ of the random demodulator can be decomposed into the product $\boldsymbol{\Phi} = \boldsymbol{R}\boldsymbol{\Psi}^{-1}$, where $\boldsymbol{R}$ is a matrix of size of size $2B/Q \times 2B$. The resulting $\boldsymbol{R}$ matrix, while randomized, typically contains some degree of structure. For example, the random convolution architecture [8, 9] endows $\boldsymbol{R}$ with a Toeplitz structure. While theoretical analysis of structured random matrices remains a topic of active study in the CS community, there do exist theoretical guarantees for some architectures [6, 8]. The amount of subsampling possible (Q_c) with these constructions is generally consistent with fully random measurements as given in (4), although the denominator is sometimes raised to a small power (e.g., 2 or 4) for provability.

3.2. CS recovery algorithms

We now address the question of how to recover the signal $x(t)$ from the measurements $\boldsymbol{y}$. The original CS theory proposed ℓ_1-minimization as a recovery technique when dealing with noise-free measurements [1, 2]. Noisy measurements as in (2) can be easily handled using similar techniques provided that the noise $\boldsymbol{e}$ is bounded, meaning that $\|\boldsymbol{e}\|_2 \leq \epsilon$. In this case, assuming that $\boldsymbol{\Phi} = \boldsymbol{R}\boldsymbol{\Psi}^{-1}$ satisfies the RIP, in which case we can write $\boldsymbol{y} = \boldsymbol{R}\boldsymbol{\Psi}^{-1}(\boldsymbol{\Psi}(\boldsymbol{\alpha})) = \boldsymbol{R}\boldsymbol{\alpha}$, the convex program

$$\widehat{\boldsymbol{\alpha}} = \operatorname{argmin}_{\boldsymbol{\alpha}} \|\boldsymbol{\alpha}\|_1 \text{ s.t. } \|\boldsymbol{R}\boldsymbol{\alpha} - \boldsymbol{y}\|_2 \leq \epsilon \tag{5}$$

can recover any sparse signal $\boldsymbol{\alpha}$. Thus by setting $\widehat{x}(t) = \boldsymbol{\Psi}(\widehat{\boldsymbol{\alpha}})$, we can recover $x(t)$. More specifically, this is made precise in [10] which establishes that for $2W$-sparse signals, the recovery error can be bounded by

$$\|\widehat{x}(t) - x(t)\|_2 \leq \kappa_1 \epsilon, \tag{6}$$

where $\kappa_1 \geq 2$ is a constant that depends on the subsampling factor Q. This constant is similarly difficult to determine theoretically, but in practice it should be very close to 2 provided that $Q < Q_c$. Thus, measurement noise has a controlled impact on the amount of noise in the reconstruction. A similar guarantee can be obtained for approximately sparse, or *compressible*, signals[2].

While convex optimization techniques like (5) are powerful methods for CS signal recovery, there also exist a variety of alternative algorithms that are commonly used in practice and for which performance guarantees comparable to that of (6) can be established. In particular, iterative algorithms such as CoSaMP, iterative hard thresholding (IHT), and various other thresholding algorithms are known to satisfy similar guarantees to (6) [11–13]. Most of these algorithms are built on similar techniques and can be easily understood by breaking the recovery problem into two separate sub-problems: identifying the locations of the nonzero coefficients of $\boldsymbol{\alpha}$ and estimating the values of the nonzero coefficients of $\boldsymbol{\alpha}$. The former problem is clearly somewhat challenging, but once solved, the latter is relatively straightforward and can be solved using standard techniques like least squares. In particular, again using the factorization of $\boldsymbol{\Phi} = \boldsymbol{R}\boldsymbol{\Psi}^{-1}$, suppose that we have identified the indices of the nonzero coefficients of $\boldsymbol{\alpha}$ (its *support*), denoted by the index set J, and let $\boldsymbol{R}_J$ denote the submatrix of $\boldsymbol{R}$ that contains only the columns corresponding to the index set J. Then an optimal recovery strategy is to solve the problem:

$$\widehat{\boldsymbol{\alpha}} = \operatorname*{argmin}_{\boldsymbol{\alpha}} \|\boldsymbol{y} - \boldsymbol{R}_J\boldsymbol{\alpha}\|_2, \tag{7}$$

which is a standard least-squares problem that can be solved via the pseudo-inverse of $\boldsymbol{R}_J$, denoted $\boldsymbol{R}_J^\dagger$:

$$\widehat{\boldsymbol{\alpha}} = \boldsymbol{R}_J^\dagger \boldsymbol{y} = (\boldsymbol{R}_J^T \boldsymbol{R}_J)^{-1} \boldsymbol{R}_J \boldsymbol{y}. \tag{8}$$

Note that in the noise-free case, if $\boldsymbol{\alpha}$ is $2W$-sparse and the support estimate J is correct , then $\boldsymbol{y} = \boldsymbol{R}_J\boldsymbol{\alpha}$, and so plugging this into (8) yields $\widehat{\boldsymbol{\alpha}} = \boldsymbol{\alpha}$. Thus, the central challenge

[2]By compressible, we mean signals that are well approximated by a sparse signal. The guarantee on the recovery error for compressible signals is similar to (6) but includes an addititional additive component that quantifies the error incurred by approximating $x(t)$ with a sparse signal. Therefore, if a signal is very close to being sparse, then this error is negligible, while if a signal is not sparse at all, then this error can be quite large.

in recovery is to identify the correct locations of the nonzeros. CoSaMP and many related algorithms solve this problem by iteratively identifying likely nonzeros, estimating their values, and then improving the estimate of which coefficients are nonzero.

3.3. Noise folding and CS measurements

Most of the CS literature focuses on ensuring the stability of the CS acquisition and recovery process in the face of *measurement noise* $\boldsymbol{e}$ in (2) [10–13]. In this section, we take a careful look at the effect of *signal noise* that is present in the signal before acquisition. Specifically, we now assume that the Fourier representation of $x(t)$ consists of a $2W$-sparse signal corrupted by additive Gaussian noise $\boldsymbol{n}$. We will assume that the noise $\boldsymbol{n}$ is zero-mean and spectrally white with covariance matrix $\boldsymbol{\Sigma_n} = \sigma^2 \boldsymbol{I}_{2B}$ and added across the entire $2B$-dimensional Fourier spectrum $\boldsymbol{\alpha}$ and not just the $2W$ non-zeros of $\boldsymbol{\alpha}$. Thus, we acquire the measurements

$$\boldsymbol{y} = \boldsymbol{\Phi}(\boldsymbol{\Psi}(\boldsymbol{\alpha} + \boldsymbol{n})) = \boldsymbol{R}\boldsymbol{\alpha} + \boldsymbol{R}\boldsymbol{n}. \tag{9}$$

The noise situation is subtly different from (2), because the noise in the measurements $\boldsymbol{y}$ is now scaled by the matrix $\boldsymbol{R}$. Our chief interest here is to understand how $\boldsymbol{R}$ impacts the signal noise and how it manifests itself in the final recovered signal.

We first observe that $\boldsymbol{Rn}$ is zero mean with covariance matrix $\boldsymbol{\Sigma_{Rn}} = \sigma^2 \boldsymbol{R}\boldsymbol{R}^T$. To further simplify this expression, we make some relatively intuitive assumptions concerning $\boldsymbol{R}$: *(i)* each column of $\boldsymbol{R}$ has norm (energy) approximately equal to 1,[3] *(ii)* the rows of $\boldsymbol{R}$ have approximately equal norm, and *(iii)* the rows of $\boldsymbol{R}$ are orthogonal. The first assumption simply means that the columns of $\boldsymbol{R}$ are weighted equally, which means that the noise n_k on each coefficient α_k will contribute roughly an equal amount to the total noise. If the locations of the nonzeros were known *a priori*, then we could do better by assigning more weight to the columns corresponding to $\alpha_k \neq 0$ and less weight to columns corresponding to $\alpha_k = 0$ (which are purely noise). However, in the absence of this knowledge, if we want the impact of the noise to be independent of the location of the nonzeros, then there is no alternative but to weight the columns equally. Similarly, the second assumption requires that each measurement should have equal weight. Finally, the third assumption is not strictly necessary, but it seems reasonable that if we wish to take as few measurements as possible, then each measurement should provide as much new information about the signal as possible. Note that these assumptions hold for randomly generated $\boldsymbol{R}$ matrices, as well the $\boldsymbol{R}$ matrices corresponding to more practical architectures such as the random demodulator [6].

By combining *(i)* and *(ii)*, we can infer that each row of $\boldsymbol{R}$ should have norm of approximately $\sqrt{2B/(2B/Q)} = \sqrt{Q}$.

[3] Note that this property is implied by the RIP by considering signals with only one nonzero coefficient.

Combining this with *(iii)*, we obtain the approximation

$$\boldsymbol{R}\boldsymbol{R}^T \approx Q\boldsymbol{I}_{2B/Q}. \tag{10}$$

Thus, the covariance matrix of $\boldsymbol{Rn}$ is approximately

$$\boldsymbol{\Sigma_{Rn}} \approx \boldsymbol{R}\boldsymbol{R}^T(\sigma^2 \boldsymbol{I}_{2B}) = \sigma^2 Q \boldsymbol{I}_{2B/Q}. \tag{11}$$

From (11) we see that the noise $\boldsymbol{Rn}$ is spectrally white but amplified by the subsampling factor Q. This makes sense intuitively, since we are projecting all of the noise in the $2B$-dimensional input signal ($\boldsymbol{\alpha}$ or equivalently $x(t)$) down into the $2B/Q$-dimensional measurements $\boldsymbol{y}$, and all of the noise power must be preserved. In the literature, this effect is known as *noise folding*.

Noise folding has a significant impact on the amount of noise present in CS measurements. Specifically, when acquiring a $2W$-sparse signal, as we double the subsampling factor Q (a one octave increase), the signal-to-noise ratio (SNR) of the measurements $\boldsymbol{y}$ decreases by 3dB. In words, *for the acquisition of a noisy signal of fixed sparsity, the SNR of the CS measurements decreases by 3 dB for every octave increase in the subsampling factor*.

We note that alternative signal acquisition techniques like *bandpass sampling* (sampling a narrowband signal uniformly at a sub-Nyquist rate to preserve the values but not the locations of its large Fourier coefficients) are affected by an identical 3dB/octave SNR degradation. However, in practice bandpass sampling suffers from the limitation that it is impossible to determine the original component center frequencies, and furthermore, if there are multiple narrowband signals present then it causes irreversible aliasing, in which case the components can overlap and will be impossible to separate. In contrast to bandpass sampling, however, CS acquisition preserves sufficient information to enable the recovery of both the values and the locations of the large Fourier coefficients.

3.4. Noise folding and CS reconstruction

We now examine the impact of noise folding on CS signal reconstruction. Rather than directly analyzing a particular reconstruction algorithm, we will instead consider the performance of an *oracle*-based recovery algorithm that has perfect knowledge of the true location of the $2W$ nonzeros of $\boldsymbol{\alpha}$, which we have denoted by J. While an oracle is clearly impractical, it characterizes the best that we can hope to achieve using any practical algorithm. And, in fact, we find that practical algorithms like CoSaMP typically perform almost as well as the oracle-based recovery algorithm.

Given J, the optimal estimate of $\boldsymbol{\alpha}$ is given by (8), resulting in

$$\widehat{\boldsymbol{\alpha}} = \boldsymbol{\alpha} + \boldsymbol{R}_J^{\dagger}\boldsymbol{R}\boldsymbol{n}. \tag{12}$$

Thus, we recover the signal $\boldsymbol{\alpha}$ plus the additive noise now scaled by the matrix $\boldsymbol{R}_J^{\dagger}\boldsymbol{R}$. To understand the properties of

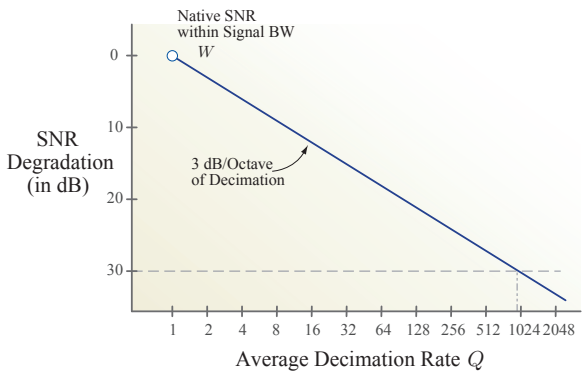

Fig. 4. Signal-to-noise ratio (SNR) degredation as a function of the subsampling factor Q. The output SNR is lower than the input SNR by 3dB per octave of decimation.

this noise, we examine its covariance matrix

$$\begin{aligned}\boldsymbol{\Sigma}_{\boldsymbol{R}_J^\dagger \boldsymbol{R}\boldsymbol{n}} &= \left(\boldsymbol{R}_J^\dagger \boldsymbol{R}\right)\left(\boldsymbol{R}_J^\dagger \boldsymbol{R}\right)^T \\ &= \boldsymbol{R}_J^\dagger \boldsymbol{R}\boldsymbol{R}^T (\boldsymbol{R}_J^\dagger)^T. \end{aligned} \tag{13}$$

If we assume that Φ has the RIP, then for any J of size K and for any $\boldsymbol{v} \in \mathbb{R}^K$, $\|\boldsymbol{R}_J \boldsymbol{v}\|_2^2 = \boldsymbol{v}^T \boldsymbol{R}_J^T \boldsymbol{R}_J \boldsymbol{v} \approx \|\boldsymbol{v}\|_2^2$. Thus

$$\boldsymbol{R}_J^T \boldsymbol{R}_J \approx \boldsymbol{I}_{2W}. \tag{14}$$

Hence, we arrive at

$$\begin{aligned}\boldsymbol{\Sigma}_{\boldsymbol{R}_J^\dagger \boldsymbol{R}\boldsymbol{n}} &\approx \sigma^2 Q \boldsymbol{R}_J^\dagger (\boldsymbol{R}_J^\dagger)^T \\ &= \sigma^2 Q \left(\boldsymbol{R}_J^T \boldsymbol{R}_J\right)^{-1} \boldsymbol{R}_J^T \boldsymbol{R}_J \left(\boldsymbol{R}_J^T \boldsymbol{R}_J\right)^{-1} \\ &= \sigma^2 Q \left(\boldsymbol{R}_J^T \boldsymbol{R}_J\right)^{-1} \\ &\approx \sigma^2 Q \boldsymbol{I}_{2W}. \end{aligned} \tag{15}$$

As with the case of the compressive measurements $\boldsymbol{y}$, the covariance matrix (15) is amplified by the factor Q. Thus, the oracle-reconstructed signal suffers from the same 3dB/octave SNR degradation as the compressive measurements.

The 3dB/octave SNR degradation represents a potentially important tradeoff in the design of CS receivers. Figure 4 illustrates the impact of the theoretically predicted growth of the noise floor, and hence noise figure, as a function of Q. This yields the the engineering design rule for CS receivers of $NF \approx 10\log_{10}(Q)$, where NF is the noise figure as defined in Section 2. We see that for a signal of fixed signal bandwidth W and input SNR, there is a practical limit to the instantaneous bandwidth B within which we can effectively acquire the signal. In Section 4 we match this theoretical result against the results of multiple simulations.

3.5. Dynamic range and CS

We now consider the impact of the subsampling factor Q on the dynamic range of a CS acquisition system. By dynamic range we mean the ratio of the maximum to minimum signal level with which received signals can be handled with full fidelity. Recall that a conventional signal receiver comprises linear input signal amplification and demodulation stages followed by a sample-and-hold circuit and a (nonlinear) quantizer, and so its dynamic range is determined by the number of bits available to the quantizer. Similarly, a CS acquisition system comprises a linear measurement operator Φ followed by a sample-and-hold circuit and then a quantizer (e.g., see Figure 3). Hence, its dynamic range will also be determined by the number of quantizer bits. We can thus identify another potential advantage of CS-based receivers: since they operate at a lower sampling rate than conventional receivers, we can potentially re-allocate some of the SWAP saved from lowering the sampling rate towards increasing the number of bits and hence the system dynamic range.

4. TESTING STRATEGY

Hardware devices are currently under construction that will permit laboratory testing of CS-based acquisition receivers [6]. In the interim, however, we have conducted a set of computer simulations to validate the theoretical performance results of the information extraction algorithms that have been developed to date. We present here the strategy associated with this simulation-based testing and an example of the results obtained so far.

The use of CS implies that the input to the receiver must be reasonably close to the definition of "sparse." To match that model, our simulation work assumed that the receiver input consists of no more than P signals, each of bandwidth no greater than W and with a total bandwidth no greater than W. We also assume the presence of white additive noise across the input band. We assume that the P input components do not overlap in the Fourier domain but otherwise might appear anywhere within the instantaneous bandwidth of the receiver B.

To make it possible to measure SNR and, in particular, the SNR degradation induced by the receiver, we have assumed a particular structure for the input signal components. Each is assumed to model a single voice signal, with uniform energy density between 300 and 3400 Hz, but with a notch at a specific frequency (2500 Hz in this case). The use of such a signal, as shown in Figure 5, permits the use of two techniques for measuring the output SNR – the classical method of computing the summed squared difference between the input and output signals and a technique called the *noise power ratio* (NPR) by which the ouput SNR is measured by comparing the power density in the notch to the power density in the rest of the signal. As [14] and others report, the technique, historically used in the telephone industry and in testing A/D

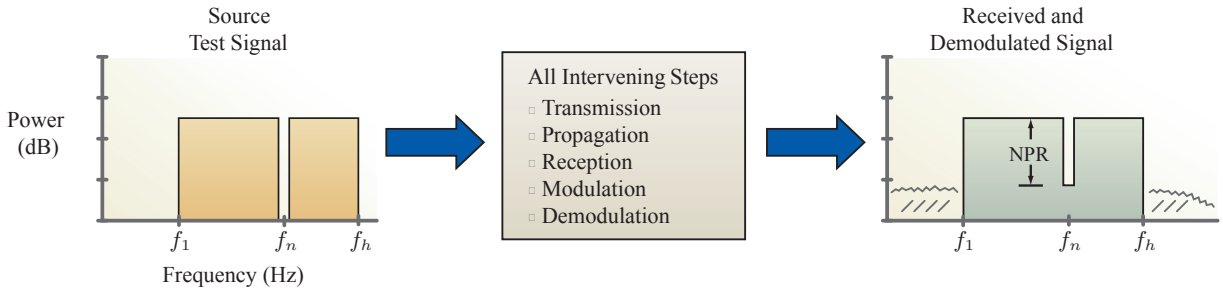

Fig. 5. The use of noise power ratio (NPR) to measure per-channel output signal-to-noise ratio (SNR).

convertors, has the advantages that it is easily computed and that it is affected by additive noise, system nonlinearities, and receiver imperfections in predictable ways. We use the classic method in the simulation-based results presented in this paper, but are using both in ongoing investigations.

The simulation testing suite assembled for this effort permits from 1 to P of these "voice-like" signals to be modulated, if desired, translated up to arbitrary frequencies, summed, corrupted with additive white noise, and then applied to the CS receiver. The quality of the recovered individual components is then determined using the NPR or classical SNR measurement technique.

Figure 6 shows the results of one set of receiver performance evaluations. In this case the transmitted signal is modeled in the simplest possible way — as a single 3.1 kHz-wide unmodulated voice signal single-side-band-upconverted to a frequency within a 1 MHz input bandwidth of the receiver. In this case performance is measured as a function of the subsamping factor Q, which was identified in Section 3 to be the key parameter affecting the cost and performance of practical receiver. The testing shown in Figure 6 was conducted at three input SNRs — 60, 40, and 20 dB — where SNR in this case is measured as the ratio of the signal power to that of the noise within the 3.1 kHz bandwidth occupied by the signal. The output SNR, measured classically within the 3.1 kHz signal bandwidth, was evaluated three ways:

- Bandpass sampling — This is not a recommended practical technique, but it does serve as a benchmark since it is "filterless" like CS. It is important to note that this method "folds" the input spectrum so that signal frequencies can no longer be unambiguously determined at the receiver.
- Idealized CS-based signal extraction using an "oracle" that knows the spectral support of the signal component — While not practical, again, the oracle provides a way to determine what portion of any observed received quality degradation is totally unavoidable within the CS framework and what portion is due to the recovery algorithm's inability to determine the spectral support.
- Practical CS-based signal extraction using CoSaMP to determine the spectral support of the input signal.

We can make the following observations from the experimental results depicted in Figure 6:

- The output SNR of both the bandpass sampled signal and the oracle-aided CS technique is degraded at a rate of 3 dB for each octave increase in Q, exactly as predicted by theory.
- The oracle-aided CS SNR closely follows the subsampled SNR until Q begins to approach the theoretical CS limit:
$$Q_c = \kappa_0 \frac{Q_f}{\ln Q_f} = \kappa_0 \frac{B/W}{\ln(B/W)}.$$
Note that for these experiments, $Q_f = (2 \cdot 10^6)/(3.1 \cdot 10^3) \approx 645$, and thus $\log_2(Q_f/\ln(Q_f)) \approx 6.6$. In Figure 6 we observe that we do not begin to observe a dramatic difference between the performance of oracle-aided CS and CoSaMP until $\log_2(Q) > 7$. While in general we should not expect to always perform this well (this example would suggest that performance does not begin to decay until κ_0 becomes less than 1), this does demonstrate that the practical impact of κ_0 on the maximum subsampling factor Q_c is likely to be minimal.
- The SNR performance of the CoSaMP algorithm generally tracks the others, but performs progressively more poorly for high subsampling factors. Moreover, its performance collapses as the theoretical limit is reached and as the input SNR falls below a critical level.
- In the regimes where the SNR performance of CoSaMP is significantly worse than that of oracle-aided CS, we observe that oracle-aided CS continues to match the SNR of the bandpass sampled signal. This indicates that in these regimes, CoSaMP is unable to identify the correct locations of the nonzero Fourier coefficients, since if it could it would match the oracle-aided CS

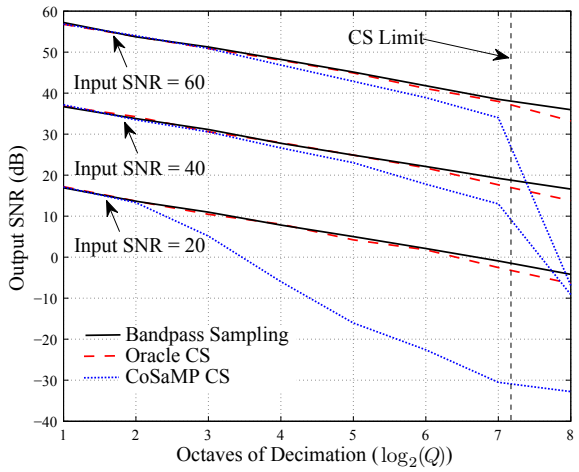

Fig. 6. One example of simulation results. Receiver ouput SNR as a function of receiver subsampling ratio Q for an environment consisting of a single unmodulated voice channel in the presence of additive white noise.

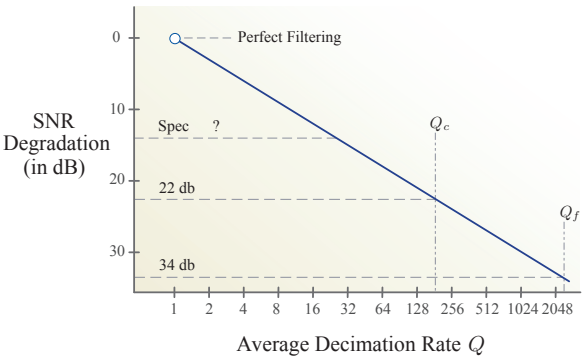

Fig. 7. Theoretical performance of CS-based receiver in a specific example.

performance, i.e., support estimation is the harder part of CS recovery (as opposed to coefficient estimation). Thus, if any side information concerning the likely locations of these nonzeros were available, one could expect that exploiting this information would have a significant impact on the SNR performance.

Many more simulation-based experiments are underway, exploring a variety of input signal mixtures and characteristics. Of particular note so far is the result that, as theoretically expected, CS-based signal extraction can successfully recover signal components without the aliasing that causes self-interference when using subsampling techniques based on simple unfiltered decimation.

5. USING THE DESIGN RULES TO EVALUATE A CS RECEIVER

The design implications of the relationships developed in Section 3 and validated, at least to first order, in Section 4 can be illustrated in Figure 7 using the system performance objectives listed in Table 1. Applying the appropriate equations from Section 3 and using the rule of thumb $\kappa_0 \approx 0.5$, we find that:

$$\begin{aligned} Q_f &= \text{subsampling factor achievable with bandpass sampling} \\ &= B/W = 2500, \\ Q_c &= \text{maximum subsampling factor for a CS-based system} \\ &= \kappa_0 \frac{Q_f}{\ln Q_f} \approx 160 \end{aligned}$$

$$\begin{aligned} C &= \text{measurement rate of a CS-based system} \\ &= 2B/Q_c \approx 6250 \text{ kHz}, \\ NF &= \text{decimation-induced noise magnification} \\ &\approx 10 \log_{10} Q_c \approx 22 \text{ dB}. \end{aligned}$$

Thus, given the set of objectives in Table 1, we find that the sampling can be reduced by a factor of 160 (and hopefully the entire "cost vector" with it) at the price of increasing the noise floor, and hence reducing the signal SNR, by 22 dB. This should be compared with bandpass sampling schemes that can reduce the sampling rate even more (by a factor of 2500) but at the price of *(i)* high computational cost if predecimation filtering is performed or *(ii)* irretrievable aliasing of the signal components if it is not.

Not addressed in this evaluation is the dynamic range achieved by the system. As discussed in Section 3, the process of CS does not theoretically degrade or improve the dynamic range of the receiver, since the power ratios among the various input components remains unchanged. There is strong evidence, however, that the sample rate reduction brought with CS techniques can have the practical effect of increasing the system's achievable dynamic range, and, further, that specific design choices can produce a CS-based receiver that excels over conventional designs. Work on this topic is underway, and the potential for improvement depends strongly on the specific circuit components available. The rationale for our expectation of potential improvement can be seen in Figure 8. It shows a compressive sensor which employs the randomized demodulator discussed in Section 3. Note that an attribute of this design is that both the sample-and-hold circuit and the quantizer at the output of the demodulator operate at rate C, which is much slower than the rate required with conventional receivers which must digitize at rates of $2B$ or greater. This rate reduction is expected to permit the use of circuit components with much greater amplitude sampling precision, and therefore capable of preserving the dynamic range of the signal with much more care.

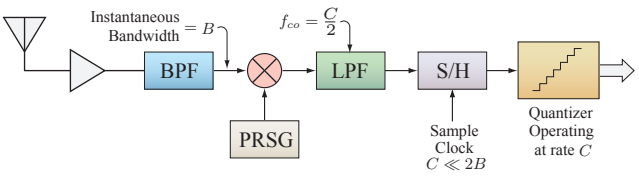

Fig. 8. The potential for CS-based dynamic range extension.

6. CONCLUSIONS AND THE WAY FORWARD

The results reported in this paper to this point can be captured as follows:

- The application of the existing CS theory to the problem of designing a high-performance RF signal acquisition receiver indicates that the approach should work, and moreover that it should be able to reduce the system SWAP and monetary cost considerably but at the cost of increasing the noise figure of the system.
- The amount of decimation (unfiltered or CS) affects SNR in a predictable way in the presence of additive white noise.
- The limited signal simulation results available so far indicate that a properly designed CS receiver can approach, or even meet, the theoretically predicted performance if the input SNR is high enough.
- Dynamic range is theoretically unaffected but, as a practical matter, might be significantly affected in the favor of certain CS receiver designs because of practical A/D implementation considerations.

These results allow the design conundrum to be clearly stated: the move to high levels of decimation, which should yield big reductions in SWAP and monetary cost, relay bandwidth, and perhaps dynamic range improvement, is achieved at the expense of "input noise amplification", which will limit the system's signal sensitivity (aka the "minimum discernible signal" (MDS)) for the acquisition system. The bottom line is that CS-based techniques are likely to be very useful in selected but important situations, such as ones with a few relatively strong narrowband signals with unknown characteristics dispersed over a very large frequency band. An example of this is illustrated in Figure 9, where a compressive acquisition receiver is used "at the point of the spear" in a tactical signal geolocation system. In this application the CS receiver samples the wideband environment and sends the resulting data to a central, non-power-constrained processing and coordination point. With this data the central processor detects the presence of emitters of interest, determines their center frequencies and bandwidths, and tasks the other signal collectors that were previously unaware of the signal. By combining the signals collected by the "tipped" collectors with the

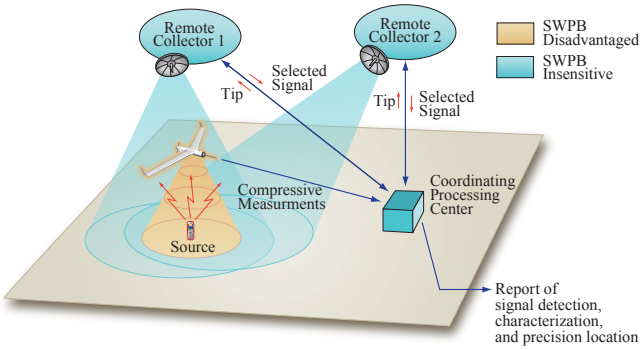

Fig. 9. Using a compact CS receiver as the primary sensor in a cross-platform geolocation system.

signals recovered from the CS receiver's data, the locations of the emitters can be determined, and with better precision than with many other collector configurations. These results, and estimates of expected performance, will appear in a sequel.

7. ACKNOWLEDGEMENTS

This paper is dedicated to Dennis Healy, whose vision and leadership have transformed compressive sensing from a theory to a practical reality through the analog-to-information research program at DARPA.

For many insightful discussions, thanks to Glenn Collins, Jeffrey Harp, and Jared Sorenen at AST, and Jason Laska, Stephen Schnelle, J. P. Slavinsky, and Andrew Waters at Rice.

This work was supported by the grants NSF CCF-0431150, CCF-0728867, CNS-0435425, and CNS-0520280, DARPA / ONR N66001-08-1-2065, ONR N00014-07-1-0936, N00014-08-1-1067, N00014-08-1-1112, and N00014-08-1-1066, AFOSR FA9550-07-1-0301, ARO MURI W311NF-07-1-0185, and the Texas Instruments Leadership University Program.

8. REFERENCES

[1] E. Candès, "Compressive sampling," in *Proc. Int. Congress Math.*, Madrid, Spain, 2006, vol. 3, pp. 1433–1452.

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

[3] R. G. Baraniuk, "Lecture notes on compressive sensing," *IEEE Signal Proc. Magazine*, vol. 24, pp. 118–124, July 2007.

[4] E. Candès and T. Tao, "Decoding by linear programming," *IEEE Trans. Inform. Theory*, vol. 51, no. 12, pp. 4203–4215, 2005.

[5] R. Baraniuk, M. Davenport, R. DeVore, and M. Wakin, "A simple proof of the restricted isometry property for random matrices," *Const. Approx.*, vol. 28, no. 3, pp. 253–263, 2008.

[6] J. Tropp, J. Laska, M. Duarte, J. Romberg, and R. Baraniuk, "Beyond Nyquist: Efficient sampling of sparse bandlimited signals," *to appear in IEEE Trans. Inform. Theory*, 2009.

[7] J. A. Tropp, M. Wakin, M. F. Duarte, D. Baron, and R. G. Baraniuk, "Random filters for compressive sampling and reconstruction," in *IEEE Int. Conf. Acoustics, Speech, Signal Proc. (ICASSP)*, May 2006.

[8] W. Bajwa, J. Haupt, G. Raz, S. Wright, and R. Nowak, "Toeplitz-structured compressed sensing matrices," in *IEEE Workshop on Statistical Signal Proc. (SSP)*, August 2007.

[9] J. Romberg, "Compressive sensing by random convolution," 2008, Preprint.

[10] E. Candès, J. Romberg, and T. Tao, "Stable signal recovery from incomplete and inaccurate measurements," *Comm. Pure Appl. Math.*, vol. 59, no. 8, pp. 1207–1223, 2006.

[11] D. Needell and J. Tropp, "CoSaMP: Iterative signal recovery from incomplete and inaccurate samples," *Appl. Comp. Harmonic Anal.*, vol. 26, no. 3, pp. 291–432, 2008.

[12] T. Blumensath and M. Davies, "Iterative hard thresholding for compressive sensing," *to appear in Appl. Comp. Harmonic Anal.*, 2008.

[13] A. Cohen, W. Dahmen, and R. DeVore, "Instance optimal decoding by thresholding in compressed sensing," 2008, Preprint.

[14] M. J. Tant, *The White Noise Book*, White Cresent Press Limited, 1974.