DYNAMIC RANGE AND COMPRESSIVE SENSING ACQUISITION RECEIVERS

John R. Treichler — Applied Signal Technology, Inc., Sunnyvale, California

Mark A. Davenport — Stanford University, Stanford, California

Jason N. Laska, Richard G. Baraniuk — Rice University, Houston, Texas

ABSTRACT

Compressive sensing (CS) exploits the sparsity present in many signal environments to reduce the number of measurements needed for digital acquisition and processing. We have previously introduced the concept and feasibility of using CS techniques to build a wideband signal acquisition systems. This paper extends that work to examine such a receiver's performance as a function of several key design parameters. In particular we show that that the system noise figure is predictably degraded as the subsampling implicit in CS is made more aggressive. Conversely we show that the dynamic range of a CS-based system can be substantially improved as the subsampling factor grows. The ability to control these aspects of performance provides an engineer new degrees of freedom in the design of wideband acquisition systems. A specific practical example, a multi-collector emitter geolocation system, is included to illustrate that point.

1. INTRODUCTION

Compressive sensing (CS) [1–3] exploits the sparsity present in many common signals to reduce the number of measurements needed for digital acquisition. With this reduction comes, in theory, commensurate reductions in the size, weight, power consumption, and/or monetary cost of both the signal sensors and any associated communication links. A previous DASP paper [4] examined the use of CS techniques to build a wideband acquisition receiver that would operate 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. In [4] we showed that such a receiver was feasible, at least in theory, and often desirable, but that the subsampling associated with compressive sensing had the negative effect of increasing the noise figure of the receiver. In this paper we discuss the other side of the coin — the positive effect that CS can have on the overall dynamic range (DR) of the acquisition system. We examine this effect theoretically, and then discuss, with a practical example, the new types of tradeoffs that use of CS permits a systems designer.

This paper is organized as follows. Section 2.1 restates the practical design problem laid out in [4] and reviews a set of requirements that a receiver should meet to be highly attractive for practical use. Section 2.2 reviews the relevant CS theory and the results from [4] and [5] that describe the performance of such a receiver in the presence of white noise. Rather than repeat the analysis recently presented in [5], Section 3 reviews the formulation of the claim that CS can substantially improve a system's DR performance and outlines the proof presented in [5]. Section 4 examines the new engineering tradeoffs that CS makes available to a designer and follows a particular example to illustrate the point. Section 5 collects various recommendations for additional study and investigation.

2. REVIEW OF TECHNICAL OBJECTIVES AND PAST RESULTS

2.1. Background and problem statement

Our objective in this paper is to continue the exploration of CS with the intent of using it in practical radio signal receiving systems. We began this in [4] by examining how it might be applied to meet a specific set of requirements. We review that example again briefly here, since it remains the reference point for this paper as well.

The particular application we 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. Many types of acquisition receivers have been designed, built, and sold over the years, but we chose in [4] a set of putative requirements for such a receiver to ease comparisons and analysis. The reader is invited to repeat the comparison for other parameter choices.

The attributes that characterize an acquisition receiver typically fall into two categories: technical specifications — such as instantaneous bandwidth — and various "costs" — such as size, weight, and power consumption (SWAP) and monetary cost. In [4], [5], and this paper we will address only the few most important technical specifications:

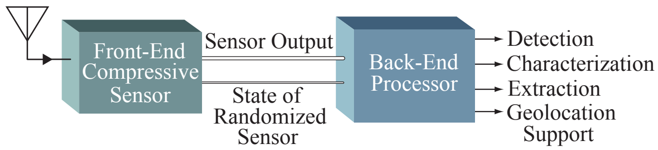

Fig. 1. The processing asymmetry assumed in a CS wideband acquisition receiver. The low size, weight, power and cost of the compressive sensor usually implies the need for substantial computation at the "backend" of the system.

- *Instantaneous bandwidth* — the RF range over which signals will be accepted by the receiver and handled with full fidelity.
- *Instantaneous dynamic range* — the ratio of the maximum to minimum signal power level for which received signals can be handled with full fidelity.
- *SNR degradation* — usually 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.
- *Maximum signal bandwidth* — the maximum combined bandwidth of the constituent signals in the acquisition bandwidth of the receiver.
- *Datalink bit rate* — the transmission rate required to carry the sampled output stream to the central processing facility.

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, using 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. If and when it can be applied, CS offers the promise to (*i*) increase the instantaneous input bandwidth, (*ii*) lower all of the cost attributes at the receiver, and (*iii*) move the computationally intensive portions of the acquisition process away from the sensor and toward a central processing facility. The "processing asymmetry" induced by CS, as identified in point (*iii*), is illustrated in Figure 1.

For the purposes of the comparisons to be made in this paper, we have assumed in Table 1 a set of requirements for an acquisition system that are rather audacious and would at the least stress conventional implementations at the present time. To meet the bandwidth and DR 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 settings without the need for scanning at the receiver, then they would have broad application.

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

Attribute		Specification
Instantaneous bandwidth	$B/2$	500 MHz
Instantaneous dynamic range	DR	96 dB
SNR degradation/noise figure	NF	12 dB
Maximum signal bandwidth	$W/2$	200 kHz
Required data link bit rate	r	150 Mb/s

In order to apply CS, we must make two last, but important, assumptions:

1. *Signal sparsity* — In order to meet the first-order assumption of all CS techniques, in this paper we assume that the input signal is sparse. To be concrete, in Table 1 we assume that the sum of the bandwidths of all signals present in the full acquisition band is no more than 200 kHz. Note that this is significantly smaller than the instantaneous bandwidth of 500 MHz. Thus we are assuming that the RF input to the receiver is significantly sparse in the frequency domain (the instantaneous bandwidth is only $1/2500$ occupied). Although 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 theoretically and in simulation, for both the case where the input is noise-free, so that 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 objective is to minimize all receiver and data link costs, i.e., the SWAP and monetary cost of the receiver and the bandwidth required for transmission. We assume that once data is acquired and transmitted, we are prepared to invest heavily in a (possibly centralized) system that can do as much processing as needed to detect, characterize, and/or recover the signal of interest. In other words, we assume that there is no cost to processing the receiver output, while there is high cost to the receiver acquisition and data forwarding processes. This separation of computation is illustrated in Figure 1.

2.2. Key result regarding noise folding in CS

The bulk of the CS literature focuses on acquisition and recovery in the face of *measurement noise* [6–10]. Moreover, most of this literature also focuses on the setting where the noise is *bounded*. In [4] and [5] we examined the effect of measurement noise as well as any *signal noise* that may be present in the signal itself. Specifically we examined the harmful impact of white additive noise at the receiver's input. The careful analysis presented in [5] reveals, under a reasonable set of circumstances, a surprisingly simple result. Theorem 4.3 from [5] states that if the noiseless input is sparse, if the additive noise is white, and if the CS measurement process satisfies the "restricted isometry property" (RIP), then the recovered SNR (RSNR) is related to the "in-band" SNR (IBSNR) of the received signal by

$$\rho\frac{1-\delta}{1+\delta} \leq \frac{\mathrm{IBSNR}}{\mathrm{RSNR}} \leq \rho\frac{1+\delta}{1-\delta}. \tag{1}$$

Here, ρ is the decimation rate, or the "subsampling ratio," and $\delta \in (0, 1)$ is a constant determined by the CS measurement process. It can be shown that ρ must be less than a critical value, denoted ρ_{CS}, which depends on $\frac{B}{W}$, the degree of sparsity of the input signal.

Further simplification of (1) yields the main result. Specifically, if we measure the ratio in dB, then we have that

$$\frac{\mathrm{IBSNR}}{\mathrm{RSNR}} \approx 10\log_{10}(\rho).$$

Thus, every time we double the subsampling factor ρ (a one octave increase), up to ρ_{CS}, the SNR loss increases by 3 dB. In other words, *for the acquisition of a sparse signal in white noise, the RSNR of the recovered signal decreases by 3 dB for every octave increase in the amount of subsampling*.

The 3dB/octave SNR degradation represents an important tradeoff in the design of CS receivers. It yields the engineering design rule for CS receivers of $NF \approx 10\log_{10}(\rho)$, where NF is the noise figure as defined in Section 2. This result implies that for a fixed signal bandwidth $W/2$ there is a practical limit to the instantaneous bandwidth $B/2$ for which we can obtain a desired RSNR. In Section 3.4 we match this theoretical result against the results of multiple simulations.

Although the noise folding behavior of CS systems imposes a very real cost, it does not necessarily preclude its use in practice, one example of which is discussed in Section 4. The dramatic sampling rate reduction enabled by CS can lead, in some cases, to significant improvements in the DR of the system. This issue is examined in the next section.

3. DYNAMIC RANGE OF A CS ACQUISITION RECEIVER

3.1. General strategy

A fundamental attribute of CS is that it enables a significantly lower sampling rate for sparse signals than would otherwise be required for full Nyquist-band sampling. This, in turn, enables the use of slower, but higher-resolution, ADCs. By exploiting this fact, a CS acquisition system should be able to provide a significantly larger DR than a conventional Nyquist-rate acquisition system within the same instantaneous bandwidth. Our strategy for demonstrating this falls into two parts.

- We first examine the literature on ADC device technology to confirm that lower sampling rates permit the use of devices with higher intrinsic DR.
- We then prove, by reference to [5], that in a properly designed CS receiver, the ADC's quantizer is the only component that limits the system's DR. Hence any improvement in the DR of the underlying ADC will result in a commensurate improvement in the DR of the CS receiver.

With these in hand, we can produce practical design rules that characterize the DR of a CS receiver.

3.2. Conversion speed versus dynamic range for ADCs

Assuming that we can prove that the CS process itself does not degrade the DR of a signal acquisition system (beyond perhaps a signal-dependent bias), the DR performance of the overall receiver depends on the low-rate ADC used to obtain the CS measurements. Rather than review the lengthy literature on the design and implementation of ADCs, we refer the reader to [11], an excellent tutorial on the topic. This paper examines the factors that affect ADC performance, predicts that performance and, the item that is important for us, compares those predictions with a large amount of empirical data, one key presentation of which appears here as Figure 2. We draw two points from [11]:

- Walden [12] predicted that the performance of an ADC (measured in several ways — effective number of bits (ENOB), DR, and SNR) should degrade at a rate of 1 bit per octave of sampling rate, over a broad range of sampling rates.
- Performance evaluation of the "best of breed" ADC converters has shown that Walden's rule is not matched precisely but, as a general trend, it is true. Specifically, there is a broad range of conversion rates (between roughly 10 kHz and 1 MHz) in which each factor-of-two reduction in sampling rate increases the DR by 1.3 bits (about 8 dB), and another range (roughly 100 MHz and above) in which each factor-of-two reduction increases the DR by about 0.9 bits (about 5.5 dB).

While it is clear that the exact value of the improvement that can be attained will depend on the exact speed and the exact ADC design, we proceed forward with the assumption that the CS-enabled sampling rate reduction can increase the system DR, by roughly one bit (and therefore roughly 6 dB) for every factor of two that CS permits the ADC sampling rate to be reduced.

(a)

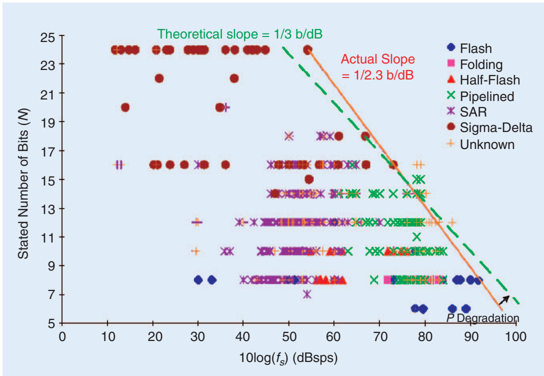

(b)

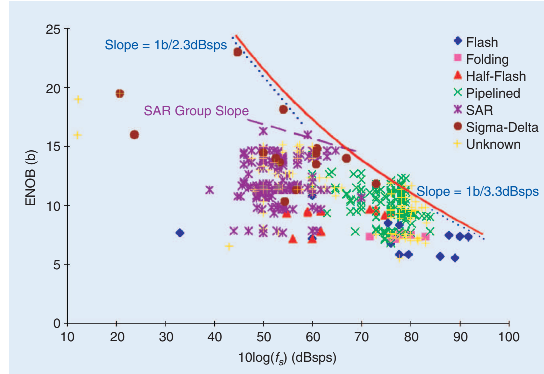

Fig. 2. Several studies, including the one illustrated here [11], have shown a clear empirical relationship between the quantization rate achieved by a practical ADC and the precision with it can make its measurements. (a) Stated number of bits vs. sample rate. (b) Effective number of bits (ENOB) vs. sample rate. Figure courtesy of [11].

3.3. The dynamic range of a compressive receiver

It remains to be shown that a properly designed CS-based receiver does not intrinsically degrade the system's DR performance. Since our recent paper [5] has laid out the argument in detail and provided the relevant proofs, here we only sketch the method and associated results.

- We first provide first a careful definition of DR, one that brings practical intuition but has the needed mathematical structure.
- We then observe that the ADC's quantizer is the key limitation to the receiver's DR. We extend this by assuming that the other components of the receiver are designed well enough that the quantizer is the *only* factor controlling the DR of the system.
- We then prove that the CS process does not affect the DR of the system, other than by a small bias, which can be positive or negative, that depends on the nature of the input signal (for example, its peak-to-average ratio).

3.4. Simulation results

Simulation work on the performance of CS receivers has been reported upon in both [4] and [5]. The key results are reproduced here in Figure 3. The two sets of curves illustrate two different aspects of SNR performance as a function of the subsampling ratio ρ. The set of curves on the left shows the effect of subsampling-induced noise folding. Just as predicted in [4] and more carefully in [5], the output SNR of a CS receiver in the presence of white additive input noise is bounded above by a term that degrades by 3 dB for every increase in the CS-enabled subsampling ratio by a factor of 2. Note also that in Figure 3 the precise value of the output SNR depends as well on a number of other factors including the input SNR, the design of the CS receiver, and the performance of the estimation algorithms located at the "backend" processor.

The set of curves on the right, conversely, shows the improvement that can be expected in DR as the subsamping ratio is increased. As outlined in the sections above, and more carefully analyzed in [5] these curves capture two effects: (*i*) the simulation-supported theoretical result that the DR performance of a properly designed CS receiver will, up to a point, be affected only by the performance of the ADC's quantizer, and (*ii*) the empirical improvement in DR, as captured in [11], of the ADCs themselves as the sampling rate is decreased. Note that in Figure 3(b) the rate of DR improvement with ρ is substantially larger than the rate of noise floor degradation shown in Figure 3(a). We observe that this is not a theoretical effect but rather one that results from practical issues associated with the ADC implementation. Advances in ADC technology might change this relationship, but, at long as it lasts, it turns out that using even more aggressive CS-based sampling, up to the limit imposed by ρ_{CS}, can produce more DR improvement than it costs in noise figure.

4. EXAMINATION OF CS FOR A SPECIFIC APPLICATION

The simulation work presented in [5] and Section 3 provides an initial validation of the engineering design relationships established in the preceding sections. To see how they might be used in practice, we return to the example set of system requirements from Table 1. In Table 2 we repeat these requirements and add two columns, the first being the specifications that might be attained by a classical wideband digital acquisition system, and the second being those that we think are attainable by a CS-based system.

For the conventional receiver we assume the use of a modern 8-bit flash ADC (e.g., the National ADC08D1500), which is capable of sampling at a rate of above 1 GHz and is advertised to provide roughly 7.3 effective bits of precision. For the purpose of comparison we will assume that the ADC is

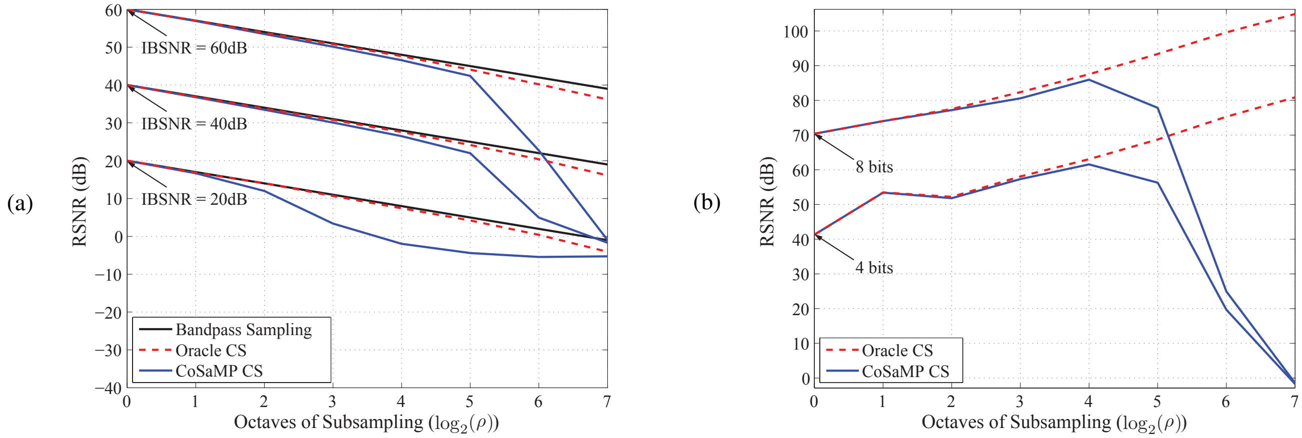

Fig. 3. Impact of the noise and quantization on the recovered SNR (RSNR) as a function of the subsampling ratio ρ and a simple sparse input consisting of a single unmodulated sinusoid. (a) shows the "3 dB/octave" noise figure increase induced by CS, while (b) shows that the DR of the receiver improves with sampling rate reduction by about 6 dB per octave.

the only source of signal degradation other than additive input noise, and therefore the system noise figure is small. Under these conditions, this classical system can be expected to be able to monitor the entire 500 MHz of instantaneous bandwidth, but with a limited dynamic range and a very large data link requirement.

By applying the rules laid out in [5] and Section 2.1, we find ρ_{CS}, the maximum subsampling factor in a CS system, is about 160. This implies that the sampling rate can be reduced from 1 GHz to 6.25 MHz. This sampling rate reduction has three key impacts:

- The noise figure goes up approximately 22 dB, thus reducing the recovered SNR of all received signals by that amount.
- However, a dynamic range improvement equivalent to an additional 9–10 bits of quantization accuracy might be achieved thanks to the lower sampling rate. In this case, if we assume the use of an 8-bit convertor for the conventional receiver, then the compressive sensing receiver should be able to achieve 17 bits or more, leading to a system dynamic range of greater than 100 dB.
- The data link bandwidth is reduced substantially. In this case the sampling rate can be reduced by a factor of 160, but the number of bits captured by the slower ADC will be greater (say 17 instead of 8). Thus the required datalink bandwidth is lowered by a factor of approximately 75, still a very substantial reduction.

Comparing these results with the objectives laid out in Table 2 shows the remarkable result that a CS-based acquisition system can theoretically meet the very stringent and rarely attained instantaneous bandwidth and dynamic range requirements, but at the cost of a worse SNR.

In aggregate, these results imply that CS introduces new tradeoffs in the design of signal acquisition systems. While a poorer noise figure reduces the sensitivity of a receiver, at the "systems level" that might be acceptable in trade for what one gets for it—much wider instantaneous bandwidth, improved dynamic range, and reduction of virtually all elements of the "cost vector" at the sensor end of the system, where it usually matters the most.

An example of how this tradeoff can be exploited is illustrated in Figure 4. Figure 4(a) portrays a traditional three-sensor arrangement for performing radio emitter geolocation. It is well known that the location accuracy of such a system is determined by, among other things, the SNRs of the signals arriving at the three sensors. It is common for these sensors to be located some distance away from the emitters and to be quite expensive and complex. Consider now the scenario shown in Figure 4(b), where one of the three sensors is brought down to a much lower altitude and implemented using the CS techniques discussed in this paper. It can be shown that in many practical cases the reduction in altitude and associated improvement in SNR at the sensor more than compensates for the CS-induced elevation of the noise figure. In fact, geolocation accuracy can actually be improved while simultaneously reducing SWAP and cost of the receiver, increasing the dynamic range, and increasing the instantaneous acquisition bandwidth enormously! All of this assumes, of course, that the input to the receiver satisfies the sparsity conditions required to employ CS. Certainly not all acquisition systems operate in such environments, but some important ones do.

Table 2. A comparison of the theoretical performance of two technical approaches to building a wideband signal acquisition receiver: Conventional high-speed digitization vs. exploiting signal sparsity through CS.

Attribute		Specification	Conventional	Compressive
Instantaneous bandwidth	$B/2$	500 MHz	500 MHz	500 MHz
Instantaneous dynamic range	DR	96 dB	44 dB	103 dB
SNR degradation/noise figure	NF	12 dB	3 dB	25 dB
Maximum signal bandwidth	$W/2$	200 kHz	500 MHz	200 kHz
Data link bit rate	r	150 Mb/s	8 Gb/s	107 Mb/s

5. CONCLUSIONS, IMPLICATIONS AND RECOMMENDATIONS FOR FUTURE WORK

This paper has examined how the positive impact that CS has on the dynamic range of a system provides an exciting new degree of freedom in the design of high-performance signal acquisition systems. Specifically, the results reported in this paper and two of its antecedents [4], [5] can be captured succinctly as follows:

- From [4] and [5] we have established that designing an RF receiver based on CS techniques is indeed feasible, and that it should reduce the size, weight, power consumption, and monetary cost of the receiver, but at the costs of increasing the receiver's noise figure and hence degrading the recovered SNR, and increasing the amount of computation required at the downstream processing center.

- We have established that the amount of subsampling by the CS receiver affects the noise figure and hence recovered SNR in a predicable way in the presence of white additive noise.

- We have also established that since CS permits the use of lower-rate, but higher performance, ADCs, the introduction of CS can actually substantially improve the dynamic range of a receiver system.

In aggregate, these results mean that CS introduces new tradeoffs in the design of signal acquisition systems. While a worse noise figure reduces the sensitivity of a receiver, at the "systems level" that might be acceptable in trade for what one gets in return — much wider instantaneous bandwidth, improved dynamic range and reduction of virtually all elements of the "cost vector" at the sensor end of the system, where it usually matters the most.

Thus we conclude that further investigation in this area is an area that will produce both theoretical and practical fruit. There are three areas in which we recommend immediate emphasis: (*i*) verification that CS receivers can actually be physically implemented with performance we have theoretically predicted, (*ii*) more work on practical and efficient processing center algorithms for signal reconstruction, or, equivalently, parameter estimation (e.g., emitter location) from the incoming compressed measurements, and (*iii*) closer examination of the effect of compressive sensors on signals with high peak-to-average ratios, which is yet another area in which CS-based systems might prove to have important advantages over conventionally designed systems [5]. Successes in the first two areas will make CS an important tool in the toolbox of radio system designers, while success in the third will only make the approach more attractive.

6. ACKNOWLEDGEMENTS

This work was supported by the grants NSF DMS-1004718, NSF CCF-0431150, CCF-0728867, CCF-0926127, CNS-0435425, and CNS-0520280, DARPA/ONR N66001-08-1-2065, N66001-11-1-4090, ONR N00014-07-1-0936, N00014-08-1-1067, N00014-08-1-1112, and N00014-08-1-1066, AFOSR FA9550-07-1-0301 and FA9550-09-1-0432, ARO MURI W911NF-07-1-0185 and W911NF-09-1-0383, and the Texas Instruments Leadership University Program.

The authors further wish to acknowledge the contributions of DARPA's Dr. Dennis Healy to the field of compressive sensing. Conversions with Dennis provided key insight into the tradeoffs between signal quality and dynamic range that are the subject of this paper.

7. REFERENCES

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

[2] E. Candès, "Compressive sampling," in *Proc. Int. Congress of Math.*, Madrid, Spain, Aug. 2006.

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

[4] J. Treichler, M. Davenport, and R. Baraniuk, "Application of compressive sensing to the design of wideband signal acquisition receivers," in *Proc. Defense Apps. of Signal Processing (DASP)*, Lihue, Hawaii, Sept. 2009.

[5] M. Davenport, J. Laska, J. Treichler, and R. Baraniuk, "The pros and cons of compressive sensing for wideband signal acquisition: Noise folding vs. dynamic range," Preprint, 2011.

[6] 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.

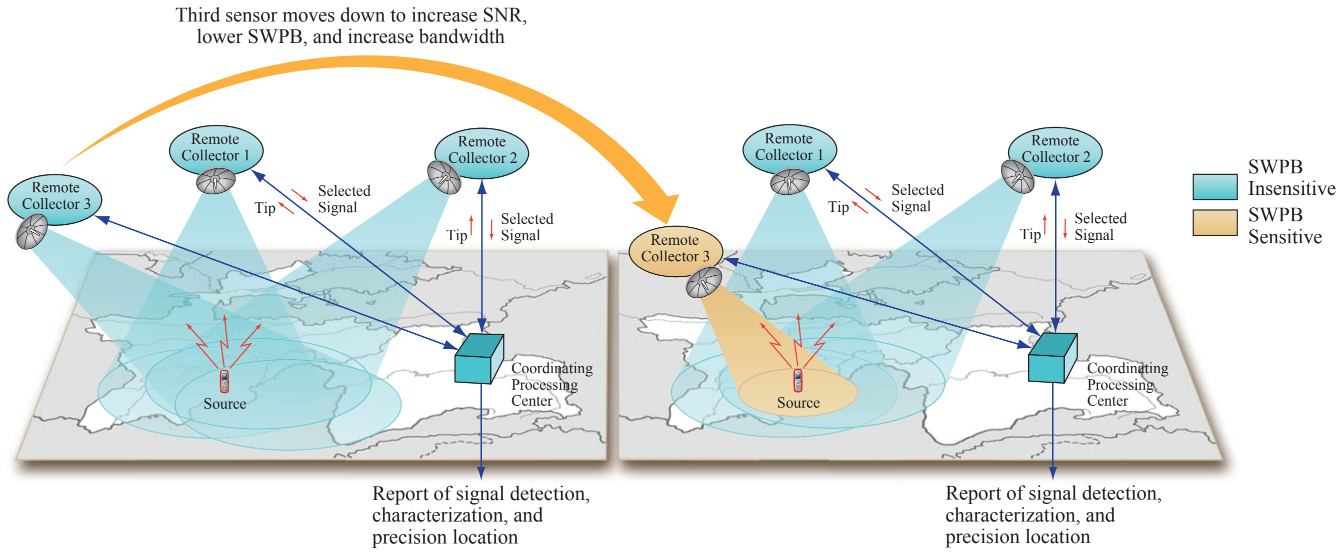

Fig. 4. The practical application of performing cross-platform emitter geolocation. (a) illustrates a classical technique in which three receiving platforms relay signals to a common processing point. The performance of such a system is strongly dependent on the SNRs seen at the collectors. (b) illustrates how a compressive sensor might be used to improve over all performance of the system. Even with the elevated noise figure that a CS receiver might have, its low SWAP might permit it to operate very close to the emitter of interest and actually improve the overall performance of the system in several ways.

[7] D. Needell and J. Tropp, "CoSaMP: Iterative signal recovery from incomplete and inaccurate samples," *Appl. Comput. Harmon. Anal.*, vol. 26, no. 3, pp. 301–321, 2009.

[8] T. Blumensath and M. Davies, "Iterative hard thresholding for compressive sensing," *Appl. Comput. Harmon. Anal.*, vol. 27, no. 3, pp. 265–274, 2009.

[9] A. Cohen, W. Dahmen, and R. DeVore, "Instance optimal decoding by thresholding in compressed sensing," in *Int. Conf. Harmonic Analysis and Partial Differential Equations*, Madrid, Spain, Jun. 2008.

[10] E. Candès and T. Tao, "The Dantzig selector: Statistical estimation when p is much larger than n," *Ann. Stat.*, vol. 35, no. 6, pp. 2313–2351, 2007.

[11] B. Le, T. W. Rondeau, J. H. Reed, and C. W. Bostian, "Analog-to-digital converters," *IEEE Sig. Proc. Mag.*, Nov. 2005.

[12] R. Walden, "Analog-to-digital converter survey and analysis," *IEEE J. Selected Areas Comm.*, vol. 17, no. 4, pp. 539–550, 1999.