Detection Performance of Compressively Sampled Radar Signals

Bruce Pollock and Nathan A. Goodman
Department of Electrical and Computer Engineering
The University of Arizona
Tucson, Arizona
brpolloc@email.arizona.edu; goodman@ece.arizona.edu

Abstract— Compressed sensing and sparse reconstruction techniques have been applied to radar signal acquisition, reconstruction of the range-Doppler map, and radar imaging. Most prior work in this area focuses on imaging and reconstruction while little attention has been paid to the effects of these techniques on detection performance. In this paper, we study the detection performance of signals acquired in an undersampled manner via random projections. We compare detection performance for signals acquired via traditional sampling of the matched filter, for correlation detection that operates directly on compressed measurements, and for the matched filter applied to the signal reconstructed via basis pursuit denoising.

I. Introduction

In order to overcome the reduction in signal strength due to propagation loss and small target radar cross section (RCS), and also to accurately measure the target's delay and Doppler shift, radars often use frequency- or phase-modulated pulses to achieve high time-bandwidth products. Acquisition of such signals is usually done with matched filtering and analog-to-digital (A/D) conversion at the bandwidth of the signal, which can require costly high-rate A/D converters (ADC) or a compromise between A/D sampling rate and number of bits (stretch processing of linear frequency modulated (LFM) signals is an exception – the sample rate can be kept low, but the pulsewidth must be long compared to the range swath). On the other hand, compressed sensing (CS) [1] techniques can be used to acquire signals at a sample rate that is below the standard Nyquist rate, potentially relieving the hardware burden and replacing it with intelligent signal processing and algorithm design.

Because data acquired via CS methods are undersampled with respect to the Nyquist rate, signal reconstruction is ill-conditioned and some form of regularization must be applied when reconstructing the full signal. In compressed sensing, this regularization takes the form of a sparsity constraint, which means that the signal can be represented by some set of basis signals with only a few non-zero basis coefficients. For example, it may be known that a signal is *sparse* in the frequency domain, meaning that the signal always consists of only a few frequency components. While the exact frequencies and their amplitudes may be unknown, the full signal can still be reconstructed under certain requirements on a) how many measurements are taken; and b) the structure of those measurements. The requirement on the structure of CS measurements is that the measurement *kernels* must be incoherent [2] with the sparse representation basis. It is now well known that random measurement kernels meet this incoherence requirement with high probability [3], and measurements taken with random kernels are generally called random projections.

Thus, while certain signals can be reconstructed despite being measured in an undersampled manner, the method of collecting these samples cannot be just a typical receiver operating at a lower rate. For example, if we were to simply undersample the output of the radar matched filter, it is likely that any peak resulting from the correlation of the matched filter with a reflected signal would be missed altogether. Furthermore, samples of the resulting range sidelobes would be insufficient, especially in the presence of noise, to determine the presence and/or range of the target. Instead, the receiver must be re-designed such that the time interval contributing to a particular sample actually spans multiple resolution cells. This measurement technique, which has a non-local kernel in the time domain, ensures that no matter the delay of the target, at least a few adjacent samples capture some of the signal's energy.

To begin thinking about taking radar measurements in a new way, consider first a function $r(t)$ that represents the signal to be captured. We can consider any measurement taken by the radar to be in the form

$$x = \int r(t)\,k(t)\,dt$$

where $k(t)$ is the measurement kernel. For example, a sample of this signal obtained at time $t = \tau$ by a conventional ADC occurs when the measurement kernel $k(t) = \delta(t - \tau)$ where $\delta(t)$ is the Dirac delta function. A radar receiver that samples after an analog implementation of a matched filter captures the

978-1-4244-8902-2/11/$26.00 ©2011 IEEE

measurement

$$x = \int r(t)\, p^*(t-\tau)\, dt$$

where $p(t)$ is the transmitted pulse. In other words, the radar receiver correlates the received waveform with a replica of the transmitted pulse. A full array of fast-time measurements is then obtained by correlating the received waveform against many replicas of the transmitted pulse at different delays, i.e.,

$$x_m = \int r(t)\, k_m(t)\, dt = \int r(t)\, p^*(t-\tau_m)\, dt\,.$$

Of course, this array of correlations is usually implemented as a matched filter followed by an ADC operating at the waveform bandwidth.

If we desire to reduce the rate of the ADC, then the measurement kernels can no longer be matched to the transmitted pulse. At the output of the matched filter, the pulse is compressed such that it's time support is approximately the reciprocal of the pulse bandwidth. Because the filter output is now localized in time, we cannot sample it with a conventional ADC (which takes temporally localized samples) unless we intend to sample every possible delay. In CS terminology, we cannot undersample because the measurement kernel is *not* incoherent [3] with the basis in which the signal can be sparsely represented – in fact, a matched filter implements fully coherent measurements. To undersample, we must instead find a new receiver architecture that implements measurements that are incoherent with the radar pulse.

The random demodulator [5] and the modulated wideband converter [6] implement random measurement kernels, which are known to form an incoherent measurement basis. The random demodulator, shown in Fig. 1, multiplies the analog waveform with a pseudo-random binary sequence. The resulting product is then integrated over a time interval equal to the sample period, and the result of the integration is stored to produce a measurement.

In this paper, we demonstrate the ability to reconstruct high-SNR signals acquired via a generalized version of the structure in Fig. 1. (We use various random signals for the measurement kernel, allow the integration time to be longer than the sample period, and the integration is approximated with a lowpass filter.) However, the ultimate objective of a radar system is to detect and locate targets. Although signals can be reconstructed in certain circumstances, the reconstruction methods are iterative and non-linear, and it is not clear how they will affect the detection statistics. Furthermore, signals acquired via undersampled methods do not achieve the full SNR gain provided by the matched filter, so we expect at least some detection loss that varies with the amount of compression. In this paper, we begin to study the detection performance of signals acquired via compressed sensing techniques. We show histograms of detection statistics in both the target-absent and target-present cases. We also compare detection performance when operating directly on the compressed measurements to performance of the matched filter applied to a signal reconstructed via basis pursuit denoising (BPDN) [7].

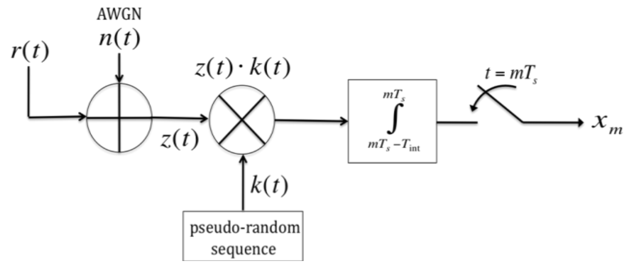

Figure 1. Block diagram of random demodulator.

In Section II, we briefly describe the signal model being used and how knowledge of the radar pulse can be used to form a sparse basis. In Section III, we present an example of a signal reconstructed via BPDN, and in Section IV, we show various detection performance analyses. We make our conclusions in Section V.

II. Signal Model and Random Demodulation

A. Signal Model

We consider a simple single-pulse model. Let the complex baseband radar waveform be denoted by $p(t)$. If a point target with reflection coefficient α is present at delay τ, then the (baseband equivalent) reflected signal that arrives at the receiving antenna will be $\alpha p(t - \tau)$. For multiple targets, we can express the received signal at the antenna as

$$r(t) = \sum_{i=1}^{N_t} \alpha_i\, p(t-\tau_i) \tag{1}$$

where N_t is the number of targets. The receiving antenna, low-noise amplifier, transmission lines, and other components all add noise to the signal. We model this noise as complex additive white Gaussian noise (AWGN) $n(t)$ (see Fig. 1). Thus, the noise-corrupted signal that enters the sampling structure of the receiver is

$$z(t) = \sum_{i=1}^{N_t} \alpha_i\, p(t-\tau_i) + n(t)\,. \tag{2}$$

For typical radar waveforms that use some form of phase or frequency modulation, the noise-free signal in (1) spans both a broad time interval (the pulse width) and a reasonably wide bandwidth. Thus, prior to compression the signal is not sparse in either the time or the frequency domain, but instead is sparse in the pulse basis. That is, assuming that the number of targets is small, (1) clearly shows that the received signal can be represented as a linear combination of just a few scaled and time-shifted radar pulses. Hence, a sparse representation basis is the set of pulse waveforms spanning the set of possible target delays.

B. CS Fundamentals

Compressed sensing deals with signals that are sparse. If a basis exists in which the vector representation of the signal contains mostly zeros, then with respect to some coherency conditions, a series of projections can be designed so that the vector is mapped into a new, more compact representation and is completely recoverable. Let a vector representation of some signal be

$$\mathrm{r} = \sum_i \alpha_i \mathrm{p}_i = \mathrm{P}\boldsymbol{\alpha}, \quad (3)$$

where p_i denotes a column vector of the matrix P and α_i is an element of the length-N vector $\boldsymbol{\alpha}$. For $p \leq 1$, if $\|\boldsymbol{\alpha}\|_p$ is small relative to N then a suitable sensing matrix K can be designed to form a more compact representation of $\boldsymbol{\alpha}$. In this paper, P is a basis of continuous-time, square-integrable pulses $p(t - \tau_i)$, represented by their discrete samples $p_i[n] = p(nT_s - \tau_i)$, where T_s satisfies the Nyquist criterion for the pulse bandwidth.

A measurement operator can be designed that performs the necessary projections. Denote the measurement operator as K, then

$$\mathrm{A} = \mathrm{KP} \quad (4)$$

where A is equivalent to the sparse basis projected into the measurement space. The measurement operator K will act upon r to form the length-M measurement vector x,

$$\mathrm{x} = \mathrm{Kr}\,. \quad (5)$$

Let $V = \|\boldsymbol{\alpha}\|_0$. As long as K and P obey the *Restricted Isometry Property* (RIP) [3], (that is, the rows of K cannot be represented sparsely with the columns of P and vice-versa) and M is on the order of $V\log(N/V)$, then r can be recovered from x. The recovery process involves some form of regularization in conjunction with a linear matrix solver (either Basis Pursuit or some kind of Gradient Method) [1-4].

C. *Random Demodulator*

The random demodulator multiplies the input signal by a pseudo-random sequence, and then integrates to produce a measurement sample. The combination of multiplying by a pseudo-random sequence and then integrating is equivalent to performing an inner product between the signal and pseudo-random sequence over a particular time interval. Letting the integration period be T_{int} and the time between samples be T_s, the mth measurement is

$$x_m = x(mT_s) = \int_{mT_s - T_{\mathrm{int}}}^{mT_s} z(t)\,k(t)\,dt \quad (6)$$

where $k(t)$ is the pseudo-random measurement kernel. Note that if $T_s > 1/2B$ where B is the radar pulse (baseband) bandwidth, then the received signal is undersampled. Moreover, we could allow $T_{int} > T_s$, which would cause adjacent projections to overlap. Stacking the samples into a measurement vector produces the vector $\mathrm{x} = [x_1\ x_2\ \ldots\ x_M]$. Our analysis involves performing correlation-based detection directly on the compressed measurement vector x and also on the Nyquist signal representation reconstructed via BPDN.

III. Reconstruction Example

Consider an LFM pulse with pulsewidth equal to 0.2 µs and (bandpass) bandwidth of 200 MHz for a time-bandwidth product of 40 [8]. Figure 2 shows a reconstructed pulse with a delay of 0.375 µs. The SNR, defined here as the ratio of squared pulse amplitude to E[$|n(t)|^2$], in Fig. 2 is 20 dB. The undersampling factor/measurement compression ratio is 6. Thus, instead of sampling at a rate of 200 MHz, the receiver sampling rate is 33.3 MHz. The integration time per sample is four times longer than the sampling interval, or 0.12 µs. Fig. 2 shows (the real part of) the noisy waveform, the noise-free waveform, and the waveform reconstructed from the compressed measurements using BPDN. The compressed measurements themselves are shown in Fig. 3. Parameters of the sparse reconstruction have not been optimized here, yet the reconstructed signal is quite good and seems to have much of the noise removed. The waveform reconstruction shown in Fig. 2 demonstrates the potential for acquiring radar signals at rates lower than the waveform bandwidth.

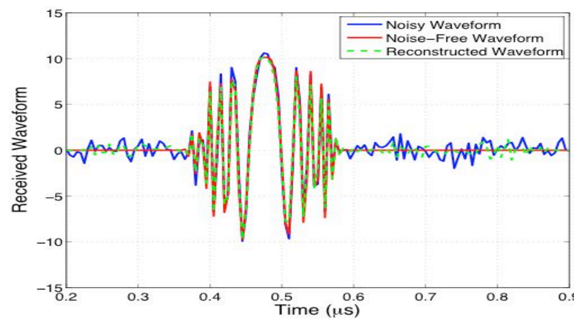

Figure 2. LFM waveform reconstruction example.

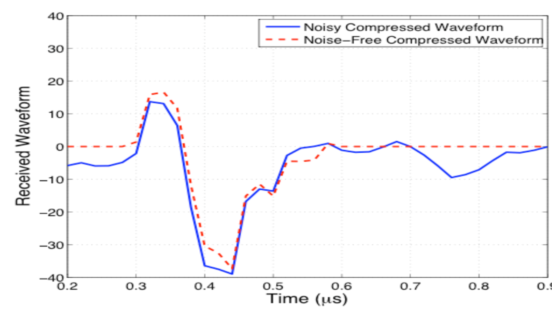

Figure 3. LFM waveform compression example (output of the random demodulator – the samples form the measurement vector x).

If the reconstructed signal were passed through the filter matched to the LFM waveform, it appears that the resulting output would have a peak in the proper location corresponding to the target range. However, it's obviously more difficult to obtain a good signal reconstruction in low-SNR environments where radars must operate in order to maximize detection range. In low-SNR scenarios, the signal is often weaker than the receiver noise and cannot be seen or detected until after the SNR gain is realized by compressing the radar pulse. An important question, therefore, is whether sparse reconstruction methods can perform well enough in low-SNR scenarios to realize pulse compression gain. Moreover, since sparse reconstruction methods are non-linear, it is not clear what the distribution of the detection statistics will be, or even if they can be predicted at all. In the next section, we show sample results that begin to explore the behavior of CS and sparse reconstruction with respect to detection performance in realistic environments.

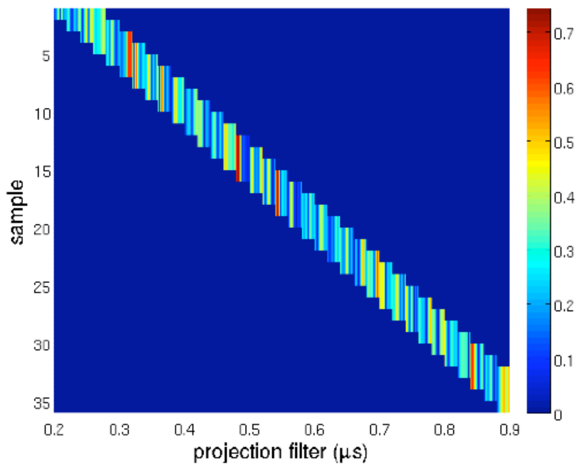

Figure 4. Projection/measurement kernel (K).

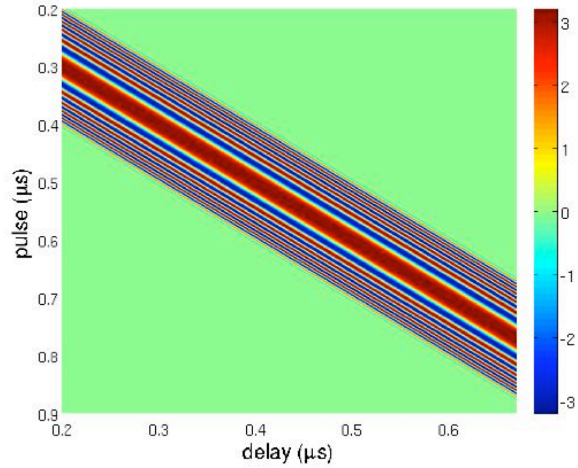

Figure 5. Representation basis matrix P for sparse vector α.

IV. Detection Performance Analysis

In this section, we begin exploring the detection performance of a radar system based on fast-time compressive sampling via random projections. The projections in (6) can be expressed in matrix form via a matrix projection kernel K. For convenience, we simulate the received signal in discrete-time form where the initial representation is at or above the Nyquist rate (to model the "analog" signal in our simulations, we sampled at 10x the Nyquist rate). The projection kernel K is made up of zero-mean Gaussian distributed samples passed through a low-pass FIR filter matched to the bandwidth of signal of interest $r(t)$. Each row of the measurement matrix is then a repeat of this random kernel, with all but the time-shifted interval of integration for that measurement set to zero.

Fig. 4 depicts the structure of the measurement matrix. Let the vector of discrete-time values prior to compression be z. The compressed measurements can then be defined as $x = Kz$. Each row of K implements a projection onto a different time interval of the received signal. If the projections overlap, then some of the entries will be repeated in successive rows. Breaking z into signal and noise components, we have $x = K(r + n) = Kr + Kn$.

Now consider detection of r in the case where r is known (for example, detection of a target at a specific range) except for its amplitude and phase. We wish to compare detection performance for three different approaches. First, our baseline approach is to perform optimum detection of r directly on the uncompressed data z. For a signal that is known except for phase embedded in AWGN, the optimum detector is the magnitude of the matched filter or correlation output; therefore, the detection statistic is

$$\zeta = \left| r^{H} z \right| . \tag{7}$$

where $(\cdot)^{H}$ denotes the conjugate transpose operation. Performance of this detector is well known [9].

For the second approach, we consider detection of the signal from the compressed measurements. After compression, the desired signal has the distorted form $r_c = Kr$, and in addition, the noise is no longer white. The covariance of the noise is now $C = P_n K K^{H}$ where P_n is the noise power before compression. The optimum detector for this case is

$$\zeta = \left| \left(C^{-1} r_c \right)^{H} x \right| = \left| \left(\left(K K^{H} \right)^{-1} r_c \right)^{H} x \right| . \tag{8}$$

The C^{-1} term acts as a pre-whitening filter, de-correlating the noise in x.

Finally, for the third approach, we first perform the sparse reconstruction of the original signal. Let P represent the pulse basis of the received noise-free waveform. This is made up of uniformly spaced, time-delayed versions of the LFM waveform presented in Fig. 2. The image of the basis matrix in Fig. 5 shows the delay for each pulse in P. We then define a vector of possible reflection coefficients $\boldsymbol{\alpha} = [\alpha_1,...,\alpha_N]^{T}$. This allows us to use BPDN [10] to find the reconstruction $\hat{z}$ of the signal z. The algorithm solves for $\boldsymbol{\alpha}$ according to

$$\hat{\boldsymbol{\alpha}} = \arg\min \|\boldsymbol{\alpha}\|_1 \ \text{subject to} \ \|A\boldsymbol{\alpha} - \mathbf{x}\|_2 \le \sigma \tag{9}$$

where the regularization condition is the l_1 norm of $\boldsymbol{\alpha}$ (the requirement that $\boldsymbol{\alpha}$ is sparse). The MxN matrix A is equal to KP. Once $\hat{\boldsymbol{\alpha}}$ is found, $\hat{z}$ can be formed from a simple matrix-vector multiply onto the pulse basis ($\hat{z} = P\hat{\boldsymbol{\alpha}}$). Since noise is white before compression, the noise formed from the direction of the largest gain in the measurement matrix will dominate in the post-compression noise covariance matrix C. The noise power in this direction is equivalent to the largest eigenvalue of C; therefore, for σ we use the square root of the largest eigenvalue of C.

The processing that produces $\hat{z}$ is non-linear, so it is difficult to know the distribution of the noise contained in $\hat{z}$; in fact, there is no separable noise component like there is in (2). Thus, without further information at this point, we attempt to apply the correlation detector directly to the reconstructed signal to produce the detection statistic.

$$\zeta = \left| r^{H} \hat{z} \right| . \tag{10}$$

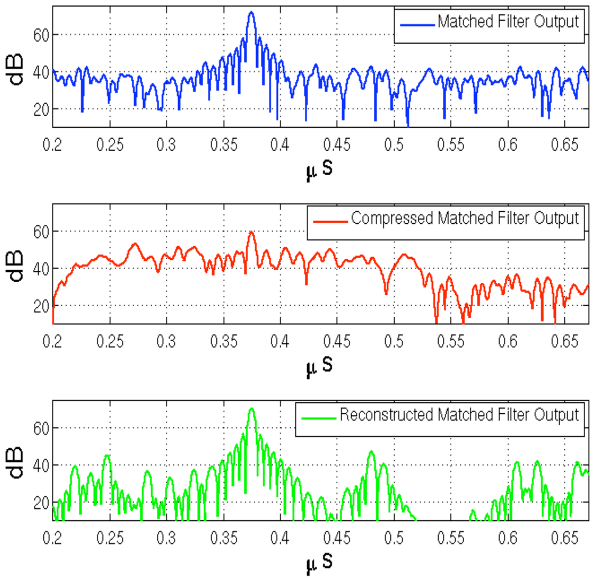

Figure 6. Matched Filter Results.

Fig. 6 shows each statistic at the output of the matched filter for the waveform shown in Fig. 2. The 16-dB pulse compression gain (equivalent to a time-bandwidth product of 40) added to the 20-dB input SNR is quite clear in the 1st graph. This gives a total gain of 36 dB over the noise floor. The second graph shows the result of passing the compressed measurements through a filter matched to $r_c = Kr$. The compression was set to 4:1 relative to Nyquist, and the resulting 12-dB loss relative to the uncompressed signal output is evident in the peak. Also the range sidelobe structure has become uniform across the time-of-arrival of the signal. The reconstructed waveform (the third graph) is quite interesting as it shows complete signal recovery in the matched filter, seemingly outperforming compression alone. However, the price paid for recovering the output signal peak seems to be a non-stationarity in the noise power over different time delays.

Fig. 7 shows performance curves for the three detection approaches applied to a single detection scenario. The target was placed at a delay of 0.375 μs (same as in Fig. 2), and the SNR was set to -10 dB (6 dB at the output of the matched filter) in order to provide small overlap between the two distributions of the test statistic in the uncompressed data case (thus, the signal strength in this scenario is right on the boundary for reasonable detection performance). The sample vector x has length $M = 36$, and the sparse vector $\boldsymbol{\alpha}$ is length $N = 943$. The length of $\boldsymbol{\alpha}$ comes from the number of basis vectors needed to represent a signal over a 0.5μs window of time delays. In the noiseless case, the number of required measurements is on the order of 7 (one non-zero element in $\boldsymbol{\alpha}$ out of 943 for $\log(N = 943) \sim 7$) [11]. We use $M = 36$ samples, which provides 4x undersampling relative to the Nyquist rate.

We see the strong performance degradation that results from compressed acquisition of the waveform. Fig. 8 shows the histogram of the noise-only detection statistic and the

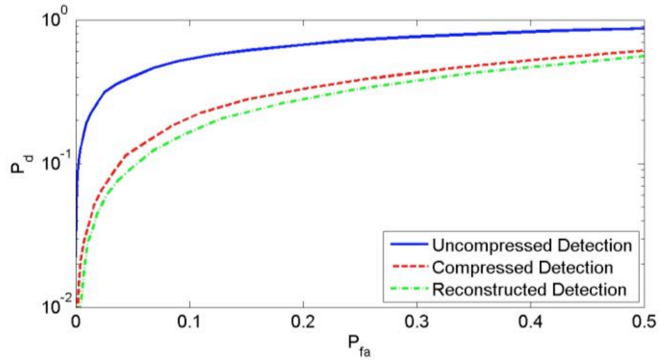

Figure 7. ROC's for compressed & uncompressed data.

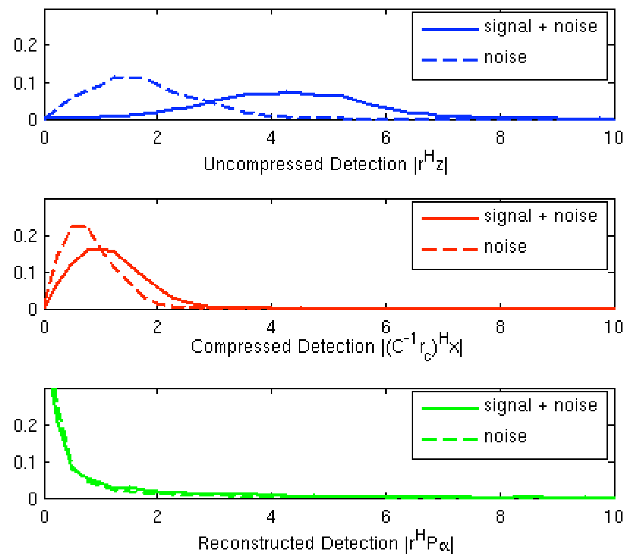

Figure 8. Histograms for compressed & uncompressed data.

histogram of the signal-plus-noise detection statistic for each detection technique. The loss in performance of the compressed detector is due to the loss in SNR from a reduction in samples, on the order of the undersample rate.

The loss in performance of the reconstructed detection is less straightforward. Certainly, some degradation is expected since compression does not follow the matched filter, and the SNR loss that occurs due to compression cannot be undone, but still the performance degradation is surprising. Figure 8 implies that the reconstructed signal has very little signal component. The nature of the random measurement basis is good for incoherency because no matter where in space the sparse signal exists, at least a small portion of it will be mapped into the compact sensing space. However, because white noise is spherical, random kernels map a large amount of noise power onto the compact sensing space, but without the corresponding gain in signal power provided by matched filtering. Thus, there is a tradeoff between SNR and sampling rate when implementing compressive sensing via random projections in the RF domain.

V. CONCLUSIONS

We have presented several results concerning the detection of radar signals acquired via a compressive receiver employing random measurement kernels. In the high-SNR regime, signal reconstruction may be possible such that radar imaging or other functions can be performed. In the low-SNR regime, radar detection performance will depend greatly on the structure of the radar waveform as well as the design of the measurement basis. In future research we will investigate the potential for optimizing measurement kernels to perform detection at a given compression ratio with minimal performance loss compared to full-rate sampling.

ACKNOWLEDGMENT

The authors acknowledge support from the Defense Advanced Research Projects Agency via grant #N66001-10-1-4079.

REFERENCES

[1] S. Kirolos, J. Laska, M. Walkin, M. Duarte, D. Baron, T. Ragheb, Y. Massoud, and R. Baraniuk, "Analog-to-information conversion via random demodulation," *IEEE Workshop on Design, Applications, Integration and Software*, pp. 71-74, Dallas/CAS, 2006.

[2] D. Donoho, "Compressed sensing," *IEEE Trans. On Information Theory*, 52(2), pp. 489-509, February 2006.

[3] E.J. Candes and M.B. Wakin, "An introduction to compressive sensing," *IEEE Signal Processing Magazine*, pp. 20-30, March 2008.

[4] E.J. Candes, T. Tao, "Near-optimal signal recovery from random projections: universal encoding strategies?," *IEEE Trans. On Information Theory*, 52(12), pp. 5406-5425, December 2006.

[5] J.A. Tropp, J.N. Laska, M.F. Duarte, J.K. Romberg, and R.G. Baraniuk, "Beyond Nyquist: efficient sampling of sparse bandlimited signals," *IEEE Trans. On Information Theory*, 56(1), pp. 520-544, January 2010.

[6] M. Mishali, Y.C. Eldar, "From theory to practice: sub-Nyquist sampling of sparse wideband analog signals," *IEEE J. Sel. Topics in Sig. Proc.*, vol. 4, no. 2, pp. 375 – 391, April 2010.

[7] S.S. Chen, D.L. Donoho, and M.A. Saunders, "Atomic decomposition by basis pursuit," *SIAM J. Scientific Computing*, vol. 20, no. 1, pp. 33 – 61, 1998.

[8] N. Levanon and Eli Mozeson, *Radar Signals*. New Jersey: John Wiley & Sons, 2004.

[9] Steven M. Kay, *Fundamentals of Statistical Signal Processing Volume II Detection Theory*. New Jersey: Prentice-Hall, 1998.

[10] E.Van Den Berg and M.P. Friedlander, "Probing the Pareto frontier for basis pursuit solutions," *SIAM J. Scientific Computing*, vol. 31, no. 2, pp. 890-912, November 2008.

[11] R. Baraniuk and Philippe Steeghs, "Compressive radar imaging," in *Proc. 2007 IEEE Radar Conference*, pp. 128 – 133, April 2007.