Compressive CFAR Radar Processing

Laura Anitori, Wim van Rossum
and Matern Otten
TNO, The Netherlands
Email: laura.anitori@tno.nl

Arian Maleki
Columbia Univeristy
New York City, USA

Richard Baraniuk
Rice University
Houston, USA

Abstract—In this paper we investigate the performance of a combined Compressive Sensing (CS) Constant False Alarm Rate (CFAR) radar processor under different interference scenarios using both the Cell Averaging (CA) and Order Statistic (OS) CFAR detectors. Using the properties of the Complex Approximate Message Passing (CAMP) algorithm, we demonstrate that the behavior of the CFAR processor is independent of the combination with the non-linear recovery and therefore its performance can be predicted using standard radar tools. We also compare the performance of the CS CFAR processor to that of an ℓ_1-norm detector using an experimental data set.

I. Introduction

Compressive Sensing (CS) is a novel data acquisition scheme that enables reconstruction of sparse signals from highly undersampled measurements. In many radar applications, such as air traffic control, obstacle avoidance, and wide area surveillance, it is reasonable to assume that the scene is sparse, since the number of targets is much smaller than the number of resolution cells in the illuminated area. Examples of CS applied to radar can be found in [1]–[5].

However, while classical radar architectures use well-established processing algorithms and detection schemes, such as Matched Filtering (MF) and Constant False Alarm Rate (CFAR) detectors, the reconstruction of the target scene from the CS measurements involves the use of highly nonlinear algorithms such as ℓ_1-norm minimization. These algorithms have a number of parameters that must be tuned properly in order to achieve good performance. The optimal value of the parameters depend on both the underlying noise power and the number of non-zero coefficients. Hence, in a practical scenario, where neither the disturbance variance nor the number of targets are known a priori, it is not clear how to tune these parameters to achieve the desired performance.

In most operational radars, to deal with the uncertainties about the background and the interference scenario, CFAR processors are widely used for adaptive target detection. Several CFAR schemes have been designed to attain good performance in the presence of different types of clutter and target scenarios [6]–[8]. The modeling and prediction of False Alarm Probability (FAP) is essential for the design of CFAR schemes. This in turn requires some level of knowledge of the underlying noise (or clutter) distribution that is input to the detector. Designing CFAR schemes seem to be out of reach for CS radar systems, due to the so far unknown relations between FAP/noise statistics and the parameters involved in the ℓ_1-norm reconstruction.

In [5], [9] we show that, using the properties of the Complex Approximate Message Passing (CAMP) [10], CFAR processing can be combined with ℓ_1-minimization to obtain fully adaptive detection schemes. In this paper, we further investigate the performance of the joint CS CFAR detector in combination with both the Cell Averaging (CA) and the Order Statistic (OS) CFAR processors under different interference scenarios using a set of CS radar measurements.

II. Complex Approximate Message Passing (CAMP)

In CS, we are concerned with the problem of recovering a k-sparse signal $\mathbf{x_0} \in \mathbb{C}^N$ from an undersampled set of linear measurements $\mathbf{y} \in \mathbb{C}^n$ of the form

$$\mathbf{y} = \mathbf{A}\mathbf{x_0} + \mathbf{n}, \tag{1}$$

where $\mathbf{A} \in \mathbb{C}^{n \times N}$ is the sensing matrix, and $\mathbf{n}$ is complex white Gaussian noise with variance σ_{in}^2. Let $n < N$ and define $\delta = n/N$ and $\rho = k/n$.

Since the number of measurements n is smaller than the number of signal samples N, the problem of recovering $\mathbf{x_0}$ is ill-posed. However, under certain conditions on $\mathbf{A}$, n, and k the following convex optimization problem, known in the literature as the LASSO [11] or Basis Pursuit Denoising (BPDN) [12], recovers a close approximation of $\mathbf{x_0}$ [13], [14]:

$$\widehat{\mathbf{x}} = \min_{\mathbf{x}} \frac{1}{2}\|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2^2 + \lambda \|\mathbf{x}\|_1, \tag{2}$$

where λ is a regularization parameter that controls the trade off between the sparsity of the solution and the ℓ_2-norm of the residual. Finding the "optimal" value of λ is a major practical problem when dealing with CS reconstruction algorithms. In particular, for radar applications the relations between the parameter λ and the detection and false alarm rates are unknown.

The Complex Approximate Message Passing (CAMP) is an iterative algorithm for solving (2) for signals in the complex domain.[1] Interestingly, the CAMP algorithm has a number properties that enable us to solve both the problem of optimal tuning and adaptive target detection. These properties are summarized in P1–P3 [10], [15], [16]:

P1: Under an appropriate tuning of the regularization parameter used in CAMP and the parameter λ in (2), CAMP solves LASSO exactly. See Section 3.4 in [10].

---

[1] A detailed description of the algorithm and its properties can be found in [5], [9], [10].

P2: At every iteration, $\tilde{\mathbf{x}}^t$ can be considered as $\mathbf{x_0}+\mathbf{w}^t$, where the distribution of $\mathbf{w}^t$ converges to complex Gaussian with zero mean and variance σ_t^2. See Section 3.4 in [10].

P3: The performance of CAMP can be predicted theoretically by the so-called state evolution equation. See Section 3.1 in [10].

An important relation derived from the analytical framework used in CAMP is that the variance of the total noise σ_t^2 present in the signal $\tilde{\mathbf{x}}$ at each iteration t is expressed as a linear combination of the input noise variance and the MSE of the solution: $\mathrm{MSE}_t = \frac{\|\hat{\mathbf{x}}^t - \mathbf{x_0}\|^2}{N}$.[2] In CAMP an estimate $\hat{\sigma}_t^2$ of the noise variance is computed at each iteration by means of median filtering.

Also, using the signal-plus-noise model described in P2, the problem of tuning the regularization parameter in CAMP, which we refer to as τ, can be easily solved. Amongst all τ, the optimal threshold τ_o in CAMP is the one that achieves the minimum MSE or, equivalently, the minimum σ_∞^2. For the practical case of unknown signal and noise statistics, we can use the *Adaptive CAMP* algorithm described in [9] to obtain a good estimate $\hat{\tau}_o$ of the optimal threshold multiplier τ_o. The optimum estimated threshold $\hat{\tau}_o$ is the one that minimizes the estimated CAMP output noise variance. This choice, in turn, also maximizes the recovery SNR of CAMP.

III. CS TARGET DETECTION USING CAMP

In radar, the detection problem is to determine the presence or absence of a target in a given range/Doppler bin when the received signal is corrupted by noise and clutter. In practice, both the noise and clutter power are unknown a priori, and therefore an adaptive detection scheme must be designed. Also, it is desirable that the detector has the CFAR property. We consider here two different CS CAMP based architectures, whose block diagrams are shown in Figure 1.

In the first system, the CS reconstruction is considered as the detector itself. This implies that in CAMP we should set the threshold, say τ_α, such that the desired FAP α is achieved. It is shown in [5] that for complex signals, if $\mathbf{x_0} = 0$, then setting the CAMP threshold $\tau_\alpha = \sqrt{-\ln\alpha}$ results in a FAP equal to α. We will refer to this detection strategy as Architecture 1; its block diagram is shown in Figure 1(a).

However, theoretical and empirical results [9] show that better performance can be achieved in terms of detection (P_d) and false alarm probability (P_{fa} or FAP) if the CS recovery is followed by a second detector. This means that, just as in conventional radar processing, we can use the recovery stage (a Matched Filter (MF) in classical architectures) to maximize the recovery SNR (i.e., maximize detection for a given FAP), and later use the detector to obtain the desired FAP. In this

---

[2]Specifically, for the case of Gaussian sensing matrices, at the fixed point solution ($t \to \infty$), the relation $\sigma_\infty^2 = \sigma_{in}^2 + \frac{1}{\delta}\mathrm{MSE}_\infty$ holds; see [10], [15] for a more detailed analysis on the (C)AMP input/output relations. From the previous equation, it is clear that the CAMP total output noise power is the sum of the effective system noise plus noise introduced by the recovery itself. Consequently, for a given input SNR, minimizing MSE also minimizes the output noise variance and therefore maximizes the reconstruction (or recovery) SNR.

(a) Architecture 1
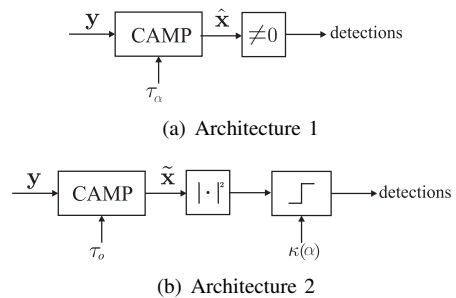
(b) Architecture 2

Fig. 1. Detection schemes based on CAMP. Note that in Architecture 2 the output of CAMP is the noisy version of the estimated signal $\tilde{\mathbf{x}}$.

case the CAMP threshold is selected to achieve the minimum MSE at the output of CAMP, i.e., $\tau = \tau_o$, and τ_o is estimated using Adaptive CAMP. We refer to this scheme as Architecture 2. In this architecture, the input to the detector is the signal $\tilde{\mathbf{x}}$. According to P2, this signal can be modeled as the sum of the true observable $\mathbf{x_0}$ plus Gaussian noise.

Thanks to the statistical properties of CAMP summarized in P1–P3, the CAMP thresholds can now be set, adaptively in Architecture 2 and as a function of the FAP for Architecture 1.

Ideally, if the noise statistics were homogeneous, stationary and known, the detector threshold in Architecture 2 could be set once and remain fixed. This represents the ideal case of a fixed threshold (FT) detector. In practice, however, these conditions are never satisfied and CFAR processors are employed to adaptively estimate the detector threshold $\kappa(\alpha)$ when the noise statistics are not known in advance. In CFAR schemes the cell under test (CUT) is tested for the presence of a target against a threshold that is derived based on an estimated clutter plus noise power. The cells surrounding the CUT (CFAR window) are used to derive an estimate of the background and they are assumed to be target free. The great advantage of CFAR schemes is that they are able to maintain a constant false alarm rate via adaptation of the threshold to a changing environment. It is known that for the case of homogeneous Gaussian background, the optimum CFAR processor is the well-known Cell Average CFAR (CA-CFAR) detector [6]. However, in situations in which the clutter changes rapidly or in the presence of interfering targets in the CFAR window, or when the clutter and noise distribution are not Gaussian, the CA-CFAR detector performance degrades severely. For this reason many alternative CFAR schemes have been developed in the past, such as the Order Statistic (OS) CFAR detector [7], [8]. In OS-CFAR processing, the power received from the cells in the CFAR window are rearranged in increasing order and the kth ordered cell (order statistic) is used as an estimate of the environment. OS-CFAR processing has the advantage of being robust against interfering targets in the CFAR window and clutter power transitions, while preserving reasonably good performance in homogenous background.

IV. Experimental Data

In the this section, we compare the performance of the proposed detection schemes under different interference scenarios using a set of experimental CS radar measurements.

A. Experimental Set-up

In our experiments, we consider the case of a one dimensional radar operating in the range domain. We use as targets five stationary corner reflectors with different Radar Cross Sections (RCS). For each transmitted waveform 300 measurements (with the same set-up) were performed.

The measurements were carried out at Fraunhofer FHR, in Germany, using the LabRadOr experimental radar system described in [5]. We used a stepped frequency (SF) waveform and the TX signal consists of a number of discrete frequencies f_m. In the Nyquist case (that represents unambiguous mapping of ranges to phases over the whole bandwidth) we transmit $N = 200$ frequencies over a bandwidth of 800 MHz. The achievable range resolution is therefore $\delta_R = 18.75$ cm. Each frequency is transmitted during 0.512 μs, corresponding to a bandwidth of $B_f = 1.95$ MHz, and sequential frequencies are separated by $\Delta f = 4$ MHz, resulting in an unambiguous range of $\Delta R = 37.5$ m.

In the CS case, the number of TX frequencies is reduced from N to n ($n < N$). The subset of transmitted frequencies is chosen uniformly at random within the total transmitted bandwidth, with the constraints that we always use the first and last frequencies in the bandwidth (to span the same total bandwidth to preserve range resolution), and we also force at least two of the transmitted frequencies to be separated by the nominal frequency separation Δf, to guarantee that the unambiguous range is preserved.

After reception and demodulation each range bin maps to n phases proportional to the n transmitted frequencies, and the n samples y_m, $m = 1, \cdots, n$, of the compressed measurement vector $\mathbf{y}$ are given by

$$y_m = \frac{1}{\sqrt{n}} \sum_{i=1}^{N} e^{-j4\pi f_m r_i/c} x_{0,i}, \qquad (3)$$

where $r_i = r_0 + i \Delta R/N$, and $i = 1, \ldots, N$ is the range bin index.

B. Results

In this section, we use ROC curves to analyze the performance of the the two CAMP based detection schemes for both interfering and non-interfering target scenarios, which we obtain by changing the CFAR window size. For Architecture 2, we combine the CAMP recovery with both the CA and OS CFAR processors.

Figure 2 exhibits the signals reconstructed by using the two CAMP based architectures introduced in Section III in addition to the MF, which represent the reference case. We use $\delta = 0.5$ for the CS measurements and $N = 200$ measurements for the MF. There are five corner reflectors (T1–T5) at ranges from 20m to 36m. For Architecture 1, τ_α was set using $\alpha = 10^{-4}$.
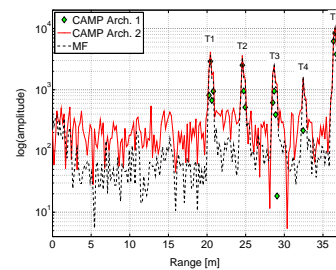

Fig. 2. Reconstructed range profile using CAMP Architectures 1 and 2, and the MF. For the MF, $N = 200$ (i.e., no subsampling); for all other schemes $n = 100$ and $\delta = 0.5$. The y-axis is in log scale and arbitrary units [au].

Notice that the signal from Architecture 2 is a noisy version of the estimated sparse signal before soft thresholding is applied, whereas the signal estimated from Architecture 1 is the sparse signal, where each non-zero coefficient represents a detection.

In the interest of space, we report only the ROC curves for target T3. For the other targets, the behavior of the detectors is the same, although the actual values of P_d are different due to the different SNRs of both the desired target and the interferers. For estimating P_d, we used the detection at the location of the highest target peak. For Architecture 2 we use both the CA and OS CFAR processors, preceded by a Square Law (SL) detector. For the CFAR processors, we use 4 guard cells and 3 different CFAR windows of length 20, 40 and 90 respectively. For the OS-CFAR, the selected order statistic is chosen as $k_{OS} = 0.6\%$.

Note that for all detector cases (adaptive and non-adaptive), the CAMP reconstruction threshold τ_o of Architecture 2 is always adaptive, whereas in Architecture 1 the threshold τ_α is non-adaptive and fixed. Furthermore, the performance of the two architectures are upper bounded by the performance of Architecture 2 that uses an ideal (non-adaptive) fixed threshold (FT) detector instead of a CFAR one.

Figure 3(a) shows the ROC curve for T3 with a CFAR window of length $M = 20$. For this choice of M, none of the other targets fall in the CFAR window of T3, and therefore the CA-CFAR processor performs better than the OS one. Furthermore, we can see that it also outperforms Architecture 1, where the noise variance is estimated inside the CAMP algorithm using the median estimator. Therefore, Architecture 1 is similar to an OS processor that uses the entire range as the CFAR window and $k_{OS} = 0.5$. Clearly, in this case the CA-CFAR performs better than both the OS and Architecture 1, since it excludes the other targets from the (local) estimation of the noise level, therefore resulting in an unbiased estimate. For this window size, CA is the best choice since there are no noise/clutter power transition, and the targets are never in the reference window of one another.

Figures 3(b) and 3(c) show the results for the same data set but for CFAR windows of sizes 40 and 90, which result in 2 and 3 interferers in the CFAR window of the target of interest. In both cases we observe that, in accordance with conventional CFAR processing, Architecture 2 with OS-CFAR outperforms
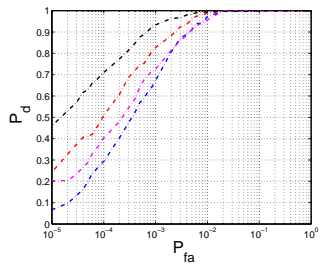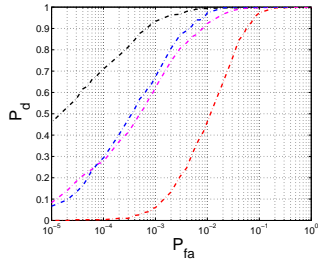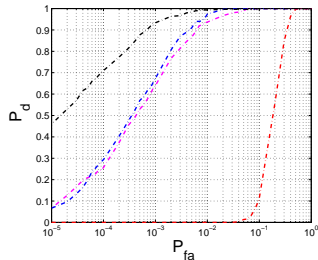
(a) CFAR window length $M = 20$

(b) CFAR window length $M = 40$

(c) CFAR window length $M = 90$

Fig. 3. ROC curves for T3 using Architecture 1 (blue), and Architecture 2 in combination with FT detector (black), CA-CFAR (red), and OS-CFAR (magenta) processors. For the OS-CFAR processor, $k_{OS} = 0.6M$. $\delta = 0.5$.

Architecture 2 with CA-CFAR but performs very similarly to Architecture 1. Furthermore, the performance of the CA-CFAR processor degrades as the number of interfering targets in the reference window increases. Note that the ROC curve of Architecture 1 (and also Architecture 2 with FT detector) is unchanged for different CFAR window sizes. In fact, for Architecture 1 we do not use a CFAR processor and the CAMP reconstruction is independent of the locations of the targets, as it uses the whole range response. It is clear that, in cases where there might be multiple interfering targets either an OS-CFAR processor should be used after Architecture 2 or otherwise the theoretically suboptimum Architecture 1 can represent a simple, effective alternative to CFAR processing. However, the disadvantage of Architecture 1 is that it lacks the local adaptivity provided by CFAR processing. Clearly, there is a trade off between the number of range bins used for the noise power estimation and the bias in the estimate that can be caused by including in the reference window interfering targets and /or noise and clutter power transitions.

V. CONCLUSIONS

In this paper we compare the results of different CS based radar detection architectures. From the experimental results we conclude that:

- the combination of CS with standard CFAR processing does not alter the behavior of the CFAR processor compared to the case when this is used in combination with a standard MF;
- in the presence of interfering targets in the CFAR window, as expected, OS is better than CA-CFAR processing;
- although the performance of Architectures 1 and Architecture 2 plus OS-CFAR are similar, Architecture 2 seems to be preferable as it leaves the user the freedom to chose the most appropriate processing parameters and it allows to perform a local adaptation of the threshold. With the combined architecture, the CFAR loss can be controlled by changing both the type of CFAR processor and the window length.

ACKNOWLEDGMENT

REFERENCES

[1] R. G. Baraniuk and T. P. H. Steeghs. Compressive radar imaging. In *Proc. IEEE Radar Conf.*, 2007.

[2] M. A. Herman and T. Strohmer. High-resolution radar via compressed sensing. *IEEE Trans. Signal Process.*, 57(6):2275–2284, 2009.

[3] J. H. G. Ender. On compressive sensing applied to radar. *Elsevier J. Signal Process.*, 90(5):1402–1414, May 2010.

[4] L. C. Potter, E. Ertin, J. T. Parker, and M. Cetin. Sparsity and compressed sensing in radar imaging. *Proc. IEEE*, 98(6):1006–1020, Jun. 2010.

[5] L. Anitori, A. Maleki, M. Otten, R. G. Baraniuk, and P. Hoogeboom. Design and analysis of compressive sensing radar detectors. Feb. 2013.

[6] P. P. Gandhi and S.A. Kassam. Analysis of CFAR processors in homogeneous background. *IEEE Trans. Aerosp. Electron. Syst.*, 24(4):427–445, Jul. 1988.

[7] H. Rohling. Radar CFAR thresholding in clutter and multiple target situations. *IEEE Trans. Aerosp. Electron. Syst.*, 19(4):608–621, Jul. 1983.

[8] S. Blake. Os-cfar theory for multiple targets and nonuniform clutter. *IEEE Trans. Aerosp. Electron. Syst.*, 24(6):785–790, Nov. 1988.

[9] L. Anitori, A. Maleki, W. van Rossum, R. Baraniuk, and M. Otten. Compressive CFAR radar detection. In *Proc. IEEE Radar Conf.*, 2012.

[10] A. Maleki, L. Anitori, Y. Zai, and R. G. Baraniuk. Asymptotic analysis of complex LASSO via complex approximate message passing (CAMP). submitted to *IEEE Trans. Inf. Theory*, 2011.

[11] Robert Tibshirani. Regression shrinkage and selection via the LASSO. *J. Roy. Stat. Soc., Series B*, 58(1):pp. 267–288, 1996.

[12] S.S. Chen, D.L. Donoho, and M.A. Saunders. Atomic decomposition by basis pursuit. *SIAM J. on Sci. Computing*, 20:33–61, 1998.

[13] E. Candès and T. Tao. Near optimal signal recovery from random projections: Universal encoding strategies? *IEEE Trans. Inf. Theory*, 52(12):5406–5425, Dec. 2006.

[14] D. L. Donoho. Compressed sensing. *IEEE Trans. Inf. Theory*, 52(4):1289 – 1306, Apr. 2006.

[15] D. L. Donoho, A. Maleki, and A. Montanari. Message passing algorithms for compressed sensing. *Proc. Natl. Acad. Sci.*, 106(45):18914–18919, 2009.

[16] D. L. Donoho, A. Maleki, and A. Montanari. The noise-sensitivity phase transition in compressed sensing. *IEEE Trans. Inf. Theory*, 57(10):6920–6941, Oct. 2011.

[17] D. P. Meyer and H. A. Mayer. *Radar target detection: handbook of theory and practice*. Academic Press Inc., New York, NY, 1973.