

PRINCIPLE AND EVALUATION OF A REGISTRATION-BASED RANGE-DEPENDENCE COMPENSATION METHOD FOR STAP IN CASE OF ARBITRARY ANTENNA PATTERNS AND SIMULATED SNAPSHOTS

Xavier Neyt*, Fabian D. Lapierre (research fellow)†, Jacques G. Verly†

*Royal Military Academy, Department of Electrical Engineering,
Avenue de la Renaissance, 30, B-1000, Bruxelles, Belgium

†University of Liège, Department of Electrical Engineering and Computer Science,
Sart-Tilman, Building B28, B-4000 Liège, Belgium

Xavier.Neyt@elec.rma.ac.be, {F.Lapierre,Jacques.Verly}@ulg.ac.be

ABSTRACT

We propose a new method for estimating the clutter power spectrum locus in arbitrary STAP radar configurations, i.e., both monostatic and bistatic. This locus is a surface in the 3D space of spatial frequency, Doppler frequency and range. Based on the knowledge of this locus, one can perform a range dependence compensation of the snapshots to obtain an estimate of the clutter covariance matrix. The method is designed to work with omnidirectional or directional transmit and receive antenna patterns and with single-realization random snapshots. End-to-end simulations show that the method provides good results.

1. INTRODUCTION

Space-Time Adaptive Processing (STAP) can be used to mitigate ground clutter in pulse-Doppler radars. However, an estimate of the interference-plus-noise (I+N) covariance matrix at the range of interest is needed to compute the optimum filter. This matrix is typically estimated by averaging sample covariance matrices at neighboring ranges. In all configurations but the monostatic (MS) sidelooking case, clutter snapshots exhibit a range-dependent power spectral density. In these cases, estimating the I+N covariance matrix by averaging snapshot power spectral density estimates leads to severe performance degradations.

The range-dependence compensation methods proposed in [1] are based on registering the ridges of the clutter power spectrum (PS) using a mathematical theory of the clutter PS locus. These methods come in two flavors: “True Parameters (TP),” where the configuration parameters are known, and “Estimated Parameters (EP),” where they are not known and autonomously estimated from the data. However, in [1], the performances of the proposed methods were only examined using theoretical covariance matrices. Furthermore, the parameter estimation method was only intended for use with omnidirectional antennas.

In this paper, we extend these methods to simulated single-realization snapshots. The extension is straightforward for the TP method. However, this becomes more difficult for the EP method. In this case, the configuration parameter estimation is performed by fitting a mathematical model of the clutter PS locus to the snapshot PS, the model being a surface in the 3D space corresponding to spatial frequency, temporal (Doppler) frequency and range.

Section 2 describes the form of the clutter PS locus. Section 3

describes the generation of simulated random snapshots and discusses some of the properties of these snapshots. The parameter estimation method is described in Section 4. In Section 5, the performance of the estimation method is discussed, both in terms of parameter-estimation accuracy and in terms of SINR loss. Let us emphasize that the method is valid for any configuration, whether MS or bistatic (BS).

2. 2D AND 3D CLUTTER POWER SPECTRUM LOCUS

The clutter patches contributing to the ground clutter in a snapshot at a given range are evidently located on the corresponding isorange surface. The signal backscattered by each clutter patch along that isorange corresponds to a particular pair (ν_s, ν_d) of spatial and temporal (Doppler) frequencies. Each snapshot is the result of the contribution of all the clutter patches along the corresponding isorange. Hence, the energy of the PS of one snapshot is located along a continuous curve in the spatio-temporal frequency plane. That curve is the clutter PS locus. This is illustrated in Fig. 1. The

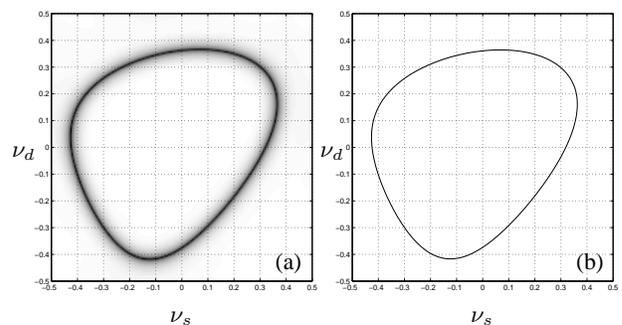


Figure 1: (a) Clutter power spectrum (PS) at a specific range and (b) corresponding 2D clutter PS locus.

analytic expression of this curve can be obtained by eliminating the location of the clutter patch between the equations of the isorange and the equations giving the spatial and temporal (Doppler) frequencies as functions of the location of the clutter patch.

The above considerations can be generalized from one particular range to a continuous variation of ranges (practically limited to a finite interval). By vertically stacking the (2D) clutter PS locus

corresponding to increasing ranges, a 3D surface is obtained. With the exception of the MS sidelooking configuration, the clutter PS locus exhibits a range dependency. This is clearly visible in Fig. 2, where the 3D clutter PS locus for two different BS configurations is plotted.

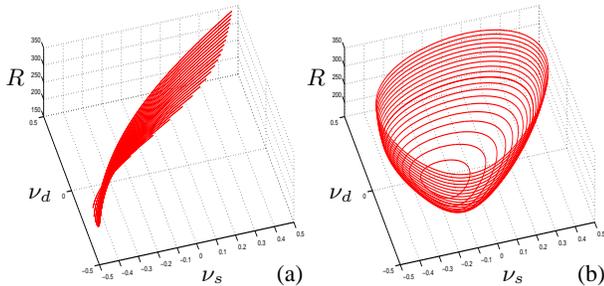


Figure 2: 3D clutter PS locus (a) for a bistatic (BS) configuration where transmitter and receiver follow each other and (b) for an arbitrary BS configuration.

3. SINGLE-REALIZATION RANDOM SNAPSHOTS

The ground clutter is modelled as described in [2]. The complex returns from individual clutter patches are assumed to have a complex circularly-symmetric gaussian distribution with zero mean and appropriate variance. The return at a particular range due to one pulse is the integral of the properly scaled and phased signals backscattered by the clutter patches located along the corresponding isorange. For each pulse, one sample is obtained at each of the N elements of the linear antenna array. Thus, a train of M (coherent) pulses results in an $N \times M$ snapshot. By stacking the 2D snapshots corresponding to successive ranges, a 3D datacube is obtained.

If the transmit (Tx) antenna has a directional pattern, the contributions to the integral come only from patches illuminated by the footprint of this antenna. Similarly, if the receive (Rx) antenna has a directional pattern, the contributions to the integral come only from patches illuminated by the footprint of this antenna. Following the same reasoning as the one in the previous section, each clutter patch can be located in the spatio-temporal frequency plane. This also means that the Tx and Rx antenna patterns produce a scaling of the intensities of the clutter PS and thus, of course, of the intensities along the clutter PS locus.

The PS of one particular realization of a random snapshot is presented in Fig. 3(a) in the case of an omnidirectional antenna. One sees that the PS energy is spread over the whole clutter PS locus. Figure 3(b) shows the PS of one particular realization of a snapshot for a directional Tx antenna with a raised-cosine diagram similar to that described in [2]. In the latter case, the PS energy is concentrated where the scatterers are illuminated by the Tx-antenna footprint. Note that Fig. 3(b) is not the simple amplitude modulation of Fig. 3(a) by the antenna pattern along the clutter PS locus. This is due to the fact that these two graphs correspond to two distinct realizations of the same random process.

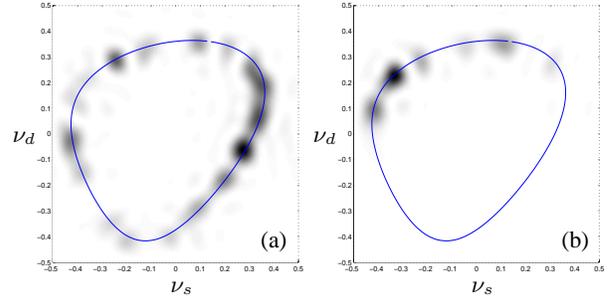


Figure 3: PS of a single-realization snapshot for an omnidirectional Rx antenna and (a) for an omnidirectional Tx antenna and (b) for a directional Tx antenna. Note that the graphs also correspond to two distinct realizations of the same random process.

4. PARAMETER-ESTIMATION METHOD

4.1. Overview of the method

As alluded to earlier, it is often useful to think of the stack of 2D PS as a continuous function of (ν_s, ν_d) and range. If the returns are essentially due to ground clutter, the 3D PS energy will concentrate along a particular surface, i.e., along the 3D clutter PS locus, examples of which are shown in Fig. 2. The parameter estimation method consists in fitting a mathematical model of the 3D clutter PS locus to the 3D PS energy. The solution of this fitting problem provides the parameters describing the radar configuration. These parameters are lumped into the parameter vector

$$\underline{\theta} = (R_{TR}, \theta, \phi, v_T, \alpha_R, \delta), \quad (1)$$

where (R_{TR}, θ, ϕ) defines the relative position of the receiver w.r.t. the transmitter in terms of spherical coordinates, v_T is the velocity of the transmitter, α_R is the relative velocity direction of the receiver w.r.t. the velocity direction of the transmitter and δ is the crab angle of the Rx antenna relative to the velocity direction of the Rx. The optimum parameter vector is obtained by minimizing some cost function. The minimization is performed by a modified simplex method. Details are provided below.

4.2. Peak extraction

The first step of the method is to locate the energy of the clutter PS. As can be seen from Fig. 3, the PS corresponding to any particular clutter snapshot is concentrated in ‘blobs’ along the clutter PS locus. Since the PS is very smooth, the location of these peaks can be found by extracting local maxima. To remove spurious low-amplitude local maxima due to noise and other artifacts, a threshold on the amplitude of the peaks is used. This is illustrated in Fig. 4.

4.3. Definition of the cost function

It is clear that the cost function ought to be related to the distance between the 3D clutter PS energy location and the 3D clutter PS locus. The chosen cost function is defined as the RMS euclidean (absolute) distance between the location of the peaks and the 3D clutter PS locus. To reduce the complexity of the method, this cost function is approximated by computing the required distances only in horizontal planes, i.e., at constant ranges.

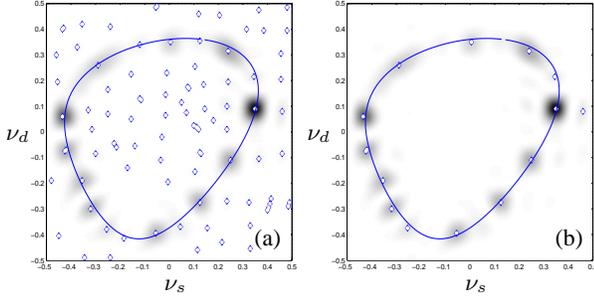


Figure 4: Extracted local maximas (white dots) in the PS of one particular clutter snapshot (a) before thresholding and (b) after thresholding.

Two strategies are used in succession to eliminate or mitigate the influence of potentially-spurious peaks on the parameter estimation to follow. First, peaks located at a distance larger than 5 times the mean distance are discarded. Second, the contribution of each peak to the cost function is weighted by the (normalized) amplitude of the peak. This is done with the rationale that low-amplitude peaks are more likely to be spurious, while large-amplitude peaks are more likely to be real.

Since the cost function is not quadratic, local minima are to be expected. These local minima correspond to natural symmetries of the system. This is illustrated in Fig. 5, which shows two dis-

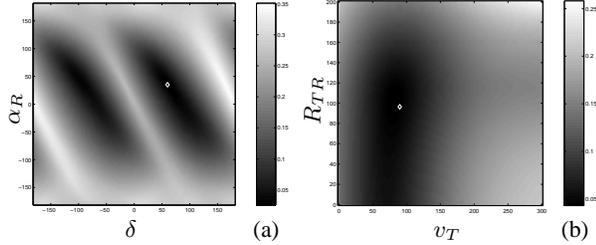


Figure 5: 2D slices in the 6D cost-function hypercube corresponding to (a) the (α_R, δ) plane and (b) the (v_T, R_{TR}) plane. The true values of the parameters are indicated by the white dots.

tinct orthogonal 2D slices in the 6D cost-function hypercube. In Fig. 5(a), two local minima are clearly visible. The corresponding true values of the parameters of interest are also shown.

4.4. Surface fitting

4.4.1. Overview

The cost function is minimized by using a variation of the simplex algorithm. This algorithm has the advantage that no analytic expression of the derivative of the cost function with respect to the parameters is needed.

It is clear that a strategy must be developed to avoid spurious local-minima solutions and to reach the global minimum. One strategy is to use several different initial values. However, without a deeper understanding of the behavior of the cost function, it is difficult to know how to select pertinent initial conditions.

The strategy we propose consists in splitting the estimation of the parameter vector in several steps. At each step, only a subset

of the parameters are estimated. The order in which the parameter subsets are estimated and the initial value used differentiate the different methods. We found that using 3 distinct (sub)methods was adequate, each working independently on the data and providing its own individual estimation of the parameter vector $\underline{\theta}$ as well as the corresponding value of the cost function. The overall method consists in picking the estimate of $\underline{\theta}$ associated with the lowest cost. Each submethod works according to the same principle and consists of 2 main phases. In Phase 1, one estimates the parameters in some manner that is adhoc but guided by the physics of the problem. In Phase 2, one reestimates all parameters using the values obtained in Phase 1 as initial values. In all cases, the estimation is done using the same algorithm, i.e., a variation of the simplex algorithm.

The first submethod works in most situations and can thus be termed “general-purpose.” The two other methods are aimed at particular pathological cases. The above discussion clearly indicates that these 3 submethods differ only by Phase 1, i.e., by the adhoc strategy used. As we shall see, each Phase 1 is tailored to a specific class of radar configurations.

One important observation lives at the heart of our strategies. As range increases, any radar configuration becomes more and more MS-like. In other words, the clutter PS locus tends towards an ellipse, which can also degenerate into a diagonal line [2, 3]. The more one approaches an MS-like configuration, the more it makes sense to assume $(R_{TR}, \theta, \phi) = (0, 0, 0)$. The more valid this assumption is, the more likely we are to correctly estimate the remaining parameters (v_T, α_R, δ) using snapshots at only long ranges. Then, we can go back and estimate the values of (R_{TR}, θ, ϕ) more precisely. To increase the sensitivity of the cost function, it even makes sense to estimate these parameters using snapshots at only short ranges. As we shall see, the above comments are at the heart of Method 1 and guide the design of Method 2 and Method 3.

We will see that a vector of initial values $\underline{\theta}_0$ can be provided to start the overall method and, thus, each of the 3 (sub)methods, unless otherwise stated. This vector could, for instance, come from the estimation of the configuration parameters during the previous coherent processing interval (CPI), with the justification that the configuration parameters do not change much from one CPI to the next. In the absence of any a priori information, this vector defaults to zero. Moreover, we assume we can constrain the parameter δ to about $\pm 10^\circ$ around its true value, since the crab angle can be measured relatively accurately. The corresponding initial value is adapted to initially satisfy this constraint.

4.4.2. Method 1

The important observation made above is the principal basis for Phase 1 of this method. The processing steps of Phase 1 are as follows:

1. Set $(R_{TR}, \theta, \phi) = (0, 0, 0)$, which corresponds to an MS configuration. Then, estimate (v_T, α_R, δ) using only long-range snapshots.
2. Continuing to use the values $(\theta, \phi) = (0, 0)$ and using the current estimates for (v_T, α_R, δ) , estimate R_{TR} using only short-range snapshots. Note that we cannot possibly estimate (θ, ϕ) in the next step below if R_{TR} remains equal to zero at the end of this step.
3. Using the current estimates for $(R_{TR}, v_T, \alpha_R, \delta)$, estimate (R_{TR}, θ, ϕ) using only short-range snapshots. If R_{TR} is

zero, the method will detect that there is no change in the value of the cost function as various values of (θ, ϕ) are tried and will abandon trying to estimate (θ, ϕ) .

As with any of the methods, we then perform Phase 2, which consists in reestimating all parameters using the values obtained at the end of Step 3 above as initial values.

Figure 6 illustrates the estimates of the 3D clutter PS locus obtained after Step 1 of Phase 1 and after Phase 2. The 3D clutter PS locus is represented from above and the +’s indicate the locations of the peaks of the PS. The characteristic shape of the locus of MS configurations is clearly visible in Fig. 6(a). If the assumption that

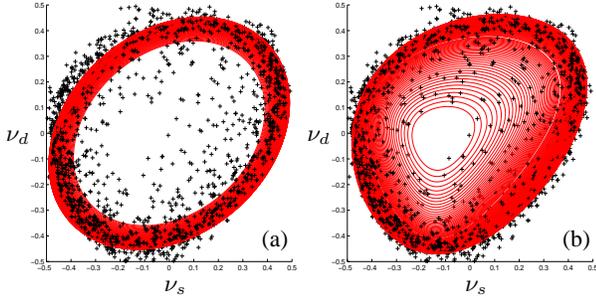


Figure 6: Results of fitting the parametric model of the 3D clutter PS locus to the extracted peaks (represented by +’s) (a) after Step 1 of Phase 1 (Method 1) and (b) after Phase 2 (Method 1).

(v_T, α_R, δ) can be estimated at long ranges is not valid, i.e., if the estimates obtained after Step 1 are too far from the true values, this method will fail. This is due to the fact that the method can get trapped in a local minimum as a result of the fact that incorrect initial values were considered. This is illustrated in Fig. 7, where

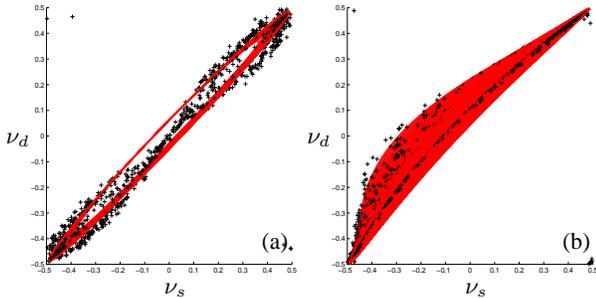


Figure 7: Results of fitting after Phase 2 (Method 1) for (a) wing-to-wing formation and (b) in-trail formation. Errors are due to getting stuck in a local minimum.

the final estimates are shown.

An explanation of the results of Fig. 7(b) follows. This figure corresponds to a scenario where Tx and Rx follow each other along the same path. In this case, the true 3D clutter PS locus is like a planar surface slightly bent out of shape (such as depicted in Fig. 2(a)). The extracted peaks, shown as +’s in Fig. 7(b), are in the general vicinity of this surface. Step 1 of Method 1 fits an ellipse to all peaks (at each range). A key observation is that the resulting ellipse essentially follows the peaks with only one of its “sides.” There are no peaks to constrain the other “side,”

which is thus free-floating. As one goes through Steps 2 and 3 of Method 1, one essentially bends the ellipse into a boomerang shape (at each range). This results in a bad estimate of the 2D clutter PS locus at each range. Both “sides” of the ellipses are so close to each other that both pick up clutter PS energy as part of the range-compensation process. This would ultimately result in a widening of the clutter notch. If we had been prescient, we would have used a diagonal line instead of an ellipse in Step 1 (at each range). Then, Steps 2 and 3 would have bent the diagonal line to match the peaks. The result would have been, not only a surface closely matching the peaks, but a good set of parameters. A similar explanation holds for Fig. 7(a).

4.4.3. Method 2

Method 2 is designed to favor parameter values for which the 3D clutter PS locus is folded onto itself, in a kind of “butterfly”- or eight-shaped surface. This corresponds to scenarios where the Tx and Rx fly wing-to-wing. This is also the scenario considered in Fig. 7(a). A major difference in strategy as compared to that of Method 1, is that, at long range, we favor an eight-shaped surface in Step 1, as opposed to an ellipse. The processing steps of Phase 1 are as follows:

1. Set $\delta = 0$ and impose a nonzero initial value for (R_{TR}, ϕ) , which favors an eight-shaped surface. Then estimate $(R_{TR}, \theta, \phi, v_T)$, using only long-range snapshots.
2. Using the current estimates for $(R_{TR}, \theta, \phi, v_T)$, estimate (v_T, α_R, δ) using all snapshots (note that v_T is reestimated).
3. Using the current estimates for all parameters, reestimate R_{TR} using only short-range snapshots. This refines the estimate of R_{TR} obtained after Step 1.

Figure 8 illustrates the 3D clutter PS locus obtained after Step 1 of Phase 1 and after Phase 2. One sees that the fit is nearly perfect after Step 1. Notice that the BS configuration considered in Fig. 8

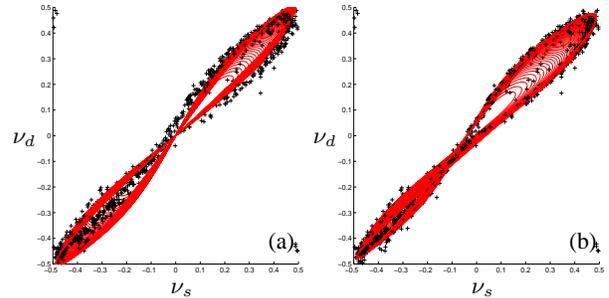


Figure 8: Result of the fitting of the parametric model of the 3D clutter PS locus (a) after Step 1 (Method 2) and (b) after Phase 2 (Method 2).

differs slightly from the one for which the method was conceived, in that the true crab angle δ is not zero, which yields a 3D clutter PS locus that has the shape of an asymmetric-eight surface.

4.4.4. Method 3

Method 3 is designed to favor parameter values for which the 3D clutter PS locus is a planar surface bent out of shape. This corresponds to Tx and Rx following each other along the same path

(this is also the scenario considered in Fig. 7(b)). A major difference in strategy, as compared to that of Method 1 and Method 2, is that, at long range, we favor a line instead of an ellipse or an eight. This is achieved by imposing $\phi = 0$ as initial value.

Phase 1 of this method thus consists of the single following step:

1. Set $(\phi, \alpha_R, \delta) = (0, 0, 0)$ and estimate (R_{TR}, θ, v_T) using all snapshots. This is done to force the estimation of meaningful initial values for (R_{TR}, θ, v_T) , before performing the global optimization. This avoids convergence to a narrow ellipse or to a narrow “eight”.

4.5. Overall method

The overall parameter estimation method consists in placing the three above (sub)methods in competition. Each Method i provides

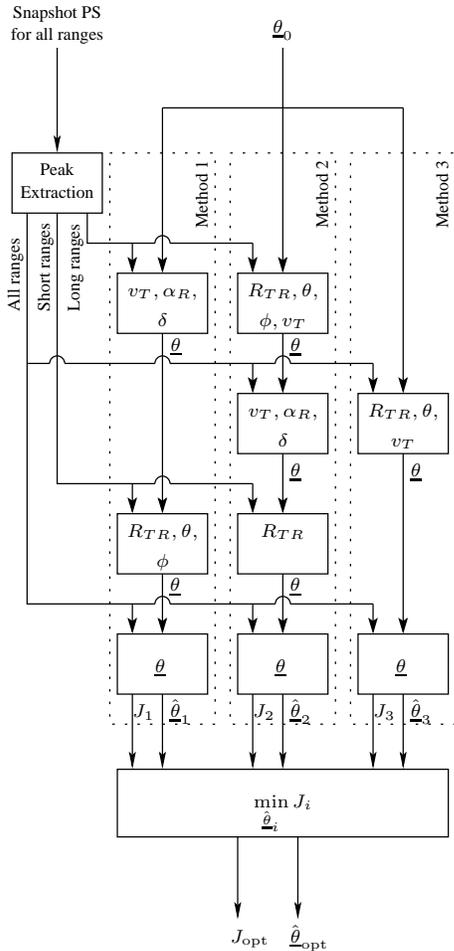


Figure 9: Block diagram of the overall method, where the 3 (sub)methods compete against each other.

an estimated parameter vector $\hat{\theta}_i$ together with the associated value J_i of the cost function. This value is the RMS distance between the extracted PS peaks and the estimated 3D clutter PS locus corresponding to $\hat{\theta}_i$. The parameter vector that is finally retained is the one that corresponds to the smallest J_i . It is indeed clear that

the corresponding $\hat{\theta}_i$ best fits the peaks extracted from the data. The block diagram of the overall parameter estimation method is depicted in Fig. 9.

5. PERFORMANCE EVALUATION

5.1. Parameter estimation

We first evaluate the performance of the proposed method w.r.t. the estimation of the parameters. The performance is essentially measured by the distance, which we denote by J_m , between the estimated 3D PS locus and the true 3D PS locus.

The estimation was performed for 10 random realizations of the STAP datacube for a particular BS configuration. The results are shown in Table 1. The successive data columns correspond to the parameters $R_{TR}, \theta, \phi, v_T, \alpha_R, \delta$, the RMS distance J between the estimated 3D PS locus and the peaks, and the value J_m of the cost function defined above. For this particular BS confi-

Method		R_{TR}	θ	ϕ	v_T	α_R	δ	J	J_m
1	mean	96.2	12.1	13.3	83.45	13.0	69.9	17.3	9.3
	stddev	1.2	0.7	1.0	0.21	1.9	0.3	1.3	1.3
2	mean	100.6	0.1	90.0	76.76	10.3	69.6	44.8	59.0
	stddev	0.9	0.1	0.0	0.81	2.8	0.4	0.5	0.8
3	mean	113.4	1.2	1.2	79.27	9.9	69.2	26.0	26.3
	stddev	0.6	0.1	0.1	0.34	0.1	0.1	0.6	0.3
Best mean		96.2	12.1	13.3	83.45	13.0	69.9	17.3	9.3
stddev		1.2	0.7	1.0	0.21	1.9	0.3	1.3	1.3
True		96.4	12.0	32.0	90.00	35.0	60.0		

Table 1: Estimated parameters for a particular BS configuration. The indicated values for J and J_m must be multiplied by 10^{-3} .

guration, Method 1 always provides the best estimates. This can be deduced from the fact that, for that method, the mean value of J_m is the lowest. Each submethod consistently provides similar results for the different STAP datacube realizations, which is reflected in the low standard-deviation value for all parameters and all submethods. The estimates are, however, less accurate for α_R and δ . This is due to the relative lack of sensitivity of the cost function w.r.t. these two parameters taken independently. This lack of sensitivity to the individual parameters α_R and δ is due to the fact that the physical quantity that influences the spatial frequency is their sum, i.e., $\alpha_R + \delta$, which is the angle between the antenna and the Tx velocity vector.

5.2. Clutter PS locus estimation

Although the method we propose is capable of estimating the BS configuration parameters, our main objective is to estimate the clutter PS locus to perform a compensation of the variation of that locus with range. Figure 10 compares the true and the estimated clutter PS loci at a specific range and in the case of omnidirectional antennas. The results are quite satisfying. Figure 11(a) shows the corresponding plots for a directive Tx antenna with a sinc-shaped antenna diagram (6 elements). The corresponding true clutter PS is depicted in Fig. 11(b). The fact that there is only a partial match between the curves of Fig. 11(a) can be explained by the fact that the clutter energy is concentrated in a relatively limited region along the true locus. It is thus impossible for the method to estimate the complete PS locus. However, a good match is achieved where energy is present. As we shall see below, this is sufficient to perform range-dependence compensation. Once again, we do not have for ultimate goal to get an accurate estimate of θ .

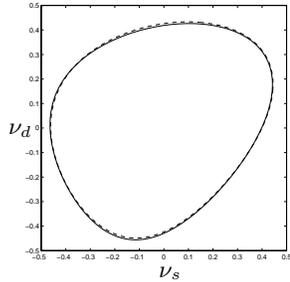


Figure 10: Comparison between the estimated (dashed) and the true (solid) clutter PS locus for omnidirectional Tx and Rx antennas.

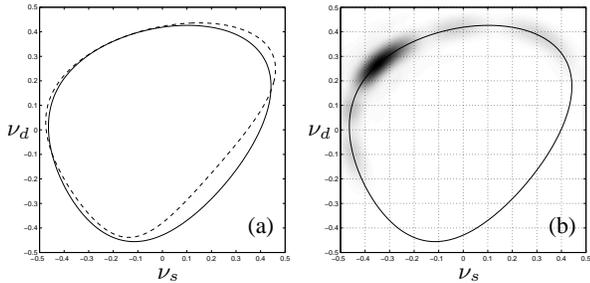


Figure 11: (a) Comparison between the estimated (dashed) and the true (solid) clutter PS locus for a sinc-shaped Tx-antenna diagram. (b) Corresponding true clutter PS and true clutter PS locus.

5.3. SINR loss

To test the real potential of the method, we used it in conjunction with the range-compensation method described in [1]. The performance of the method is compared with the performance obtained with the so-called TP method, where the configuration parameters are assumed to be known. The optimum processor, computed from the true I+N covariance matrix is used as a reference. A further comparison point is given by the straight-averaging processor, where the estimate of the I+N covariance matrix is obtained by averaging the sample covariance matrix of all snapshots.

Figure 12(a) shows cuts in the SINR loss surface for the case of omnidirectional antenna patterns. Figure 12(b) shows the corresponding graph for the case of a Tx antenna with a sinc-shaped pattern. One can see that the performances of the EP range-compensation method are undistinguishable from those of the TP method.

6. CONCLUSIONS AND FUTURE WORK

We proposed a new method for estimating the clutter PS locus given single-realization snapshots at a series of ranges. The method is shown to provide reliable estimates of the configuration parameters when the antennas have omnidirectional patterns. If directional antenna patterns are considered, the estimated parameters values differ from the true values, but the clutter PS locus is still correctly estimated in areas where the clutter PS is not zero. Fortunately, in our intended application, obtaining correct parameters values is irrelevant; all we care about is obtaining the shape of

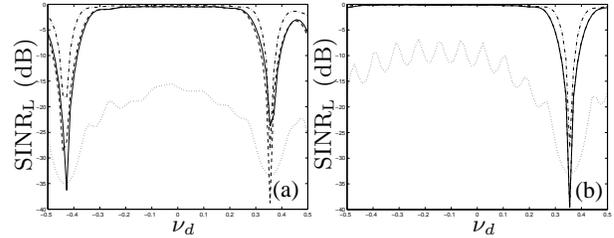


Figure 12: Cuts of SINR loss at $\nu_s = -0.2$ in the case of (a) omnidirectional antenna patterns and of (b) a sinc-shaped Tx-antenna pattern. The dash-dotted line corresponds to the optimal processor (OP), the dotted line to the straight-averaging processor (SAP), the dashed line to the range compensation performed using the true parameters (TP), and the solid line to the range compensation performed using the estimated parameters.

the clutter PS locus where energy is present. Once the clutter PS locus is obtained, it can be used to perform a range compensation of the clutter PS in order to estimate the clutter covariance matrix. The performance of the estimation method in terms of SINR losses was shown, both for omnidirectional antennas and for directional antennas.

By construction, the method is very robust to isolated scatterers such as targets. However, several issues remain to be investigated: (1) the performance in the presence of barrage jamming, (2) the performance in case of nonideal clutter (decorrelation effects, nonflat terrain, etc) and (3) the relation between the sample processor size and the accuracy of the clutter PS locus estimation.

7. REFERENCES

- [1] F. D. Lapierre and J. G. Verly, "Registration-based solutions to the range-dependence problem in STAP radars," in *Adaptive Sensor Array Processing (ASAP) Workshop*, Lincoln Laboratory, MIT, Lexington, MA, March 2003.
- [2] Richard Klemm, *Principles of space-time adaptive processing*, The Institution of Electrical Engineers (IEE), UK, 2002.
- [3] J. Ward, "Space-time adaptive processing for airborne radar," Tech. Rep. 1015, Lincoln Laboratory, MIT, Lexington, MA, Dec. 1994.