Optimal Detection and Tracking of Randomly Moving Targets *

Marcelo G. S. Bruno and José M. F. Moura
Depart. Elect. and Comp. Eng., Carnegie Mellon University, Pittsburgh, PA, 15213
ph: (412)268-6341; fax: (412)268-3890; email:moura@ece.cmu.edu

Abstract: *This paper considers optimal joint detection and tracking of targets that move randomly on finite grids. Detection decisions and target position estimates are based on raw radar measurements corrupted by noise which may have heavy tails statistics.*

1 Introduction

The problem we consider is to detect and track moving targets using noisy radar measurements. Assuming that the sensor device has a finite resolution, a target is, at each instant n, either absent or centered in a given position on a finite discrete grid that represents the different resolution cells of the sensor. The recorded sensor measurements may correspond to true returns from a target that is actuallly present and/or false returns representing spurious reflectors plus measurement noise. Real targets move randomly along the finite discrete sensor grid according to a known stochastic model.

The traditional approach to this problem involves the separation of the detection and tracking tasks [1]. By contrast, we propose to integrate the detection problem into the same framework in which the tracking problem is solved.

We apply nonlinear stochastic filtering [4] to design the optimal joint tracker/detector. We use a Bayesian estimation strategy to compute recursively the posterior probability mass function (pmf) of the target's centroid position at the nth sensor scan conditioned on all previous scans. This framework allows us to consider both Gaussian and non-Gaussian background noise and results in better performance when compared to more conventional algorithms such as the spatial matched filter and the linearized Kalman-Bucy filter.

*This work was supported by ONR grant no. N0014-97-0040. The first author was partially supported by CNPq-Brazil

2 Joint Detection/Tracking

We consider the simplest problem of tracking a single target with 1D motion. We restrict possible targets to the class of rigid bodies with translational motion and time-invariant signatures. In the absence of rotation, the tracking problem reduces therefore to tracking the target's centroid, whose position at each sensor scan n we denote by z_n.

Due to the finite resolution of the sensor, the range of values that z_n assumes is discretized by a lattice where each site represents one of the resolution cells of the sensor. We assume a finite lattice $\mathcal{L} = \{l: l = 1, \ldots, L\}$, with z_n restricted to taking values in $\mathcal{L}$. Once the centroid crosses the boundary of the lattice, the target is declared to have become absent. The target's motion is described by a finite state machine or discrete Markov chain (MC) with state space $\mathcal{L}$.

Given a set of observations (sensor returns) from instant 0 up to instant n, our problem is to determine at instant n whether the target is absent or present (detection) and, if it is present, to follow its transitions in the lattice (tracking).

To incorporate the target detection problem into the same framework in which the target tracking problem is solved, we introduce an additional zero state that represents the absence of target. The scalar random variable z_n is defined then on the modified discrete lattice $\mathcal{L} = \{0, 1, \ldots, L\}$ such that $z_n = i$, $1 \leq i \leq L$, means that the target's centroid occupies cell i at instant n and $z_n = 0$ means that the target is absent at instant n. The transition probabilities are stored in a $(L+1)$ x $(L+1)$ transition matrix $\mathbf{P}_T$ whose general element (k, j) is

$$P_T(k,j) = \text{Prob}(z_n = k \mid z_{n-1} = j) \qquad (1)$$

where $0 \leq k, j \leq L$.

Our model for the nth sensor scan $\mathbf{y}_n$, assuming real returns, is given by

$$\mathbf{y}_n = \sum_{k=-m_1}^{m_2} a_k \mathbf{e}_{z_n+k} + \mathbf{v}_n \qquad (2)$$

0-7803-5030-8/98/$10.00 ©1998 IEEE

where $\mathbf{v}_n$ is the observation noise and $\mathbf{e}_r$, $r \in \mathcal{Z}$, is a L x 1 vector defined as follows: if $r = i$, $i = 1$, 2, ..., L, $\mathbf{e}_r$ has all entries zero except for the ith position whose entry is 1. Otherwise, if $r > L$, $\mathbf{e}_r$ is an identically zero vector. The coefficients a_k in (2) define the shape of the target.

The optimal statistical solution for the detection/tracking problem in this context follows a Bayesian estimation strategy. First, we define $\mathbf{\Pi}^f[n]$ such that, for $0 \leq i \leq L$,

$$\Pi_i^f[n] = P(z_n = i \mid \mathbf{y}_0^n) \tag{3}$$

where vector

$$\mathbf{y}_0^n = [\mathbf{y}_0\ \mathbf{y}_1\ \cdots\ \mathbf{y}_n] \tag{4}$$

collects all the observations from instant 0 up to instant n. Vector $\mathbf{\Pi}^f[n]$ is the vector of the posterior probabilities of the $(L+1)$ states at instant n given the observations $\mathbf{y}_0^n$. Likewise, we define $\mathbf{\Pi}^p[n]$ such that

$$\Pi_i^p[n] = P(z_n = i \mid \mathbf{y}_0^{n-1}) \quad 0 \leq i \leq L\ . \tag{5}$$

The solution to the optimal joint detection and tracking problem is now divided in two steps:

Prediction Step

Predict the next state of the target from the previous one using the Markov chain's transition matrix. By the Theorem of Total Probability and using the fact that, conditioned on $\mathbf{z}_{n-1}$, $\mathbf{z}_n$ is independent from $\mathbf{y}_0^{n-1}$, we get

$$P(z_n \mid \mathbf{y}_0^{n-1}) = \sum_{z_{n-1}} P(z_n \mid z_{n-1}) P(z_{n-1} \mid \mathbf{y}_0^{n-1}) \tag{6}$$

In matrix notation, equation (6) is written as

$$\mathbf{\Pi}^p[n] = \mathbf{P}_T \mathbf{\Pi}^f[n-1]\ . \tag{7}$$

Filtering Step

Correct the prediction with the new information given by the observations. From Bayes' Law and using the fact that, conditioned on $\mathbf{z}_n$, $\mathbf{y}_n$ is independent from $\mathbf{y}_0^{n-1}$, we can write

$$P(z_n \mid \mathbf{y}_0^n) = C_n p(\mathbf{y}_n \mid z_n) P(z_n \mid \mathbf{y}_0^{n-1})\ . \tag{8}$$

In matrix notation, equation (8) is rewritten as the pointwise vector multiplication

$$\mathbf{\Pi}^f[n] = C_n \mathbf{S}_n \odot \mathbf{\Pi}^p[n] \tag{9}$$

where

$$S_n[i] = p(\mathbf{y}_n \mid z_n = i) \quad 0 \leq i \leq L\ . \tag{10}$$

The constant C_n is a normalization constant such that

$$\sum_{i=0}^{L} \Pi_i^f[n] = 1\ . \tag{11}$$

We now consider detection and tracking.

Detection

The probability of a target being absent at instant n conditioned on the observations is given by $\Pi_0^f[n]$. Representing by H_0 the hypothesis that the target is absent and by H_1, the hypothesis of target present, the minimum probability of error detector follows the decision strategy

$$\Pi_0^f[n] \underset{H_1}{\overset{H_0}{\gtrless}} 1 - \Pi_0^f[n]\ . \tag{12}$$

Tracking

We introduce now the conditional probability

$$\begin{aligned} Q_i^f[n] &= P(\mathbf{z}_n = i \mid \text{target is present}, \mathbf{y}_0^n) \\ &= \frac{\Pi_i^f[n]}{1 - \Pi_0^f[n]}\ . \end{aligned} \tag{13}$$

A maximum a posteriori (MAP) tracking strategy estimates the actual position of the target's centroid assuming that it is present as

$$\hat{z}_{\text{map}}[n] = \arg \max_{1 \leq i \leq L} Q_i^f[n]\ . \tag{14}$$

Remark: Tracking Gates

The main source of computational burden in this algorithm comes from the fact that the raw algorithm computes the posterior pmf of the position of the target's centroid over the entire surveillance region. A reasonable alternative is to define for the target a particular tracking gate which is a smaller subset of the larger surveillance region. The tracking algorithm now looks for each target only within the corresponding gate. The tracking gates are defined following an acquisition stage during which the tracker looks for targets in the entire surveillance space.

3 Simulation Examples

In the following simulations, we assume the simplest possible MC model for the target's centroid. Assuming that the velocity of the target is decomposed into a mean deterministic component plus a turbulent random velocity, the 1D target motion is described by

$$z_n = z_{n-1} + d + \varepsilon_n \tag{15}$$

where z_n is a defined on a discrete grid $(1, 2, \ldots L)$, d is a positive integer representing the constant (positive) mean drift of the target and ε_n is a fluctuation of the velocity around the mean drift. The fluctuation is modeled as a random walk that takes values $\{-1, 0, 1\}$ with probabilities r, $1 - q - r$, and q respectively. Both r and q were set at 0.40.

At a given sensor scan, we assume that at most one target can be present. However, whenever a present target becomes absent, a new target can reappear randomly at any cell of the grid. The probability of a new target entering the grid was set at 0.30. A present target is assumed to be pointwise with a constant unit signature, i.e., the coefficientes a_k in (2) are such that $a_0 = 1$ and $a_k = 0$, $-m_1 \leq k \leq m_2$, $k \neq 0$.

3.1 Detection and tracking in a Gaussian background

Assuming the measurement noise vector $\mathbf{v}_n$ in (2) to be Gaussian with zero mean and covariance $\mathbf{R}$, the probability density function (pdf) of the observations conditioned on the position of the target is given simply by $N(\mathbf{e}_{z_n}, \mathbf{R})$ where $N(.)$ is the Gauss function with mean $\mathbf{e}_{z_n}$ and covariance $\mathbf{R}$.

From the point of view of detection, the scheme proposed in this paper corresponds to a minimum probability of error detector in a Bayesian sense.If instead of setting the threshold in (12) to 1, we vary the threshold over a wide range, the detector can be changed from a minimum probability of error test to a Neyman-Pearson scheme.

For different values of threshold, the detector operates at different (fixed) values of probability of false alarm. Figure 1 shows experimental (Monte Carlo) ROC curves showing different combinations of probabilities of detection and false alarm assuming white Gaussian noise ($\mathbf{R} = \sigma_w^2 \mathbf{I}$) with SNR per scan equal to 10, 6 and 3 dB respectively. Each point in the curves was obtained from an average on 20000 sensor scans with each scan corresponding to 80 cell returns.

Even under relatively low levels of SNR, the ROC curves tend to a "step-like" shape, i.e, for very low levels of false alarm, we observe correspondingly much higher levels of detection. The shape of the ROC curves highlights the very good performance of the detector, even in adverse noisy environments.

As far as tracking is concerned, the simplest tracking scheme, assuming no knowledge of the target's motion, is the memoryless maximum likelihood (ML) estimator based on one single sensor scan. Under Gaussian noise, the ML estimator reduces to the spatial matched filter, which simply slides the target's template along the noisy sensor image until it finds the best match.

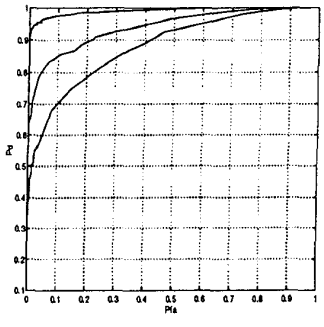

Figure 1: Monte Carlo ROC curves, white Gaussian noise, SNR= 10, 6 and 3 dB

If we have just a pointwise target, i.e., $a_0 = \alpha$ and $a_k = 0$ otherwise, the spatial matched filter reduces to a peak detector, given by

$$z_{ML} = \arg \max_{1 \leq i \leq L} y_n[i] \quad . \tag{16}$$

One possible improvement on the raw peak detector is to consider its output as a noisy observation of a state variable to which an approximate linearized Kalman-Bucy filter (KBF) is associated. In the linearized KBF, the position of the target is the rounded output of a linear filter driven by white Gaussian noise. The variance of the driving noise is set to match approximately the Markov chain model. The variance of the observations error is obtained from a Monte Carlo estimate of the peak detector error.

Figure 2 compares the tracking results for one simulated trajectory with white Gaussian clutter using: 1) the raw peak detector (dashed line), 2) the peak detector corrected by a linearized KBF (dashdotted line) and 3) the proposed PDF tracker (marked by '+'). The solid line is the real trajectory and the SNR per scan is 10 dB. The total number of cells is 300. As expected, the performance of the peak detector is very poor. The linearized KBF improves tracking due to inertia in the prediction step that reduces the magnitude of peak detector errors. However, the linearized KBF still performs far worse than the nonlinear PDF tracker.

3.2 Detection and tracking in non-Gaussian clutter

In this section, we consider a simulated experiment where the recorded measurements at instant n cor-

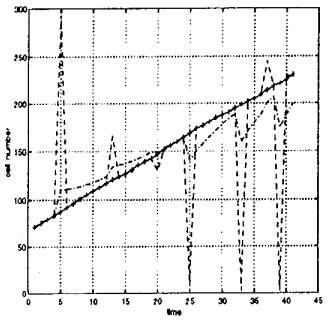

Figure 2: Tracking results for peak detector, linearized KBF and PDF tracker, SNR= 10 dB

respond to a sampling of the returned complex envelope and are given by the even-sized vector

$$\mathbf{y}_n = [y^1_{c_n}\, y^1_{s_n}\, \ldots\, y^L_{c_n}\, y^L_{s_n}] \tag{17}$$

where L is the number of resolution cells and

$$\begin{aligned} y^k_{c_n} &= s^k_{c_n} + v^k_{c_n} \\ y^k_{s_n} &= s^k_{s_n} + v^k_{s_n} \quad 1 \le k \le L \end{aligned} \tag{18}$$

represents some averaging of the samples of the low-pass quadrature returns over a time interval Δt corresponding to the scan of one single resolution cell.

We arbitarily assume that the background noise vector

$$\mathbf{v}_n = [v^1_{c_n}\, v^1_{s_n}\, \ldots\, v^L_{c_n}\, v^L_{s_n}] \tag{19}$$

is such that the sequence of random variables

$$e_k = \sqrt{(v^k_{c_n})^2 + (v^k_{s_n})^2} \quad 1 \le k \le L \tag{20}$$

is identically distributed with a probability density function different from a Rayleigh distribution. We are interested in analyzing how the tracker performs against a background noise whose envelope has "longer tails" than a Rayleigh pdf. We are particularly interested in two statistical models for e_k, namely

Weibull-PDF

$$p_E(e) = a\, c\, e^{c-1} \exp(-a\, e^c) \;\; e \ge 0 \tag{21}$$

where c is a shape parmeter and a is related to the average power σ^2 of the quadrature components by

$$2\sigma^2 = a^{-2/c}\, \Gamma(1 + \frac{2}{c}) \;. \tag{22}$$

K-PDF

$$p_E(e) = \frac{b^{\nu+1}\, e^{\nu}}{2^{\nu-1}\Gamma(\nu)}\, K_{\nu-1}(b\, e) \;\; e \ge 0 \tag{23}$$

where ν is a shape parameter, $\Gamma(.)$ is the Eulerian function, $K_{\nu-1}(.)$ is a modified Bessel function of the second kind and b is related to σ^2 by

$$b^2 = \frac{2\,\nu}{\sigma^2} \;. \tag{24}$$

Simulation of the background clutter

The background noise vector is generated using the spherically invariant random vector (SIRV) technique introduced originally in the context of radar clutter by Conte and Longo [2].

Specifically, let $\mathbf{v} = [v_{c_1}\, v_{s_1}\, \ldots\, v_{c_N} v_{s_N}]$ be a 2N x 1 vector representing the quadrature components of clutter returns. All components of $\mathbf{v}$ are assumed to have zero mean with common variance σ^2. We also assume that the quadrature components v_{c_i} are v_{s_i} , $i = 1, \ldots N$ are orthogonal. Vector $\mathbf{v}$ is said to be a white SIRV if it has a joint multivariate pdf of the form [2]

$$p_{\mathbf{V}}(\mathbf{v}) = (2\pi)^{-N}\, \sigma_w^{-2N}\, h_{2N}(-\frac{\mathbf{v}^T\mathbf{v}}{2\sigma_w^2}) \;. \tag{25}$$

The class of functions $h_{2N}(y)$ in (25) is defined by the expression

$$h_{2N}(y) = \int_0^\infty s^{-2N} \exp(-\frac{y}{2s^2}) p_S(s) ds \tag{26}$$

where $p_S(s)$ is a function greater or equal than zero for all $s \ge 0$ to which we refer as the characteristic pdf of the SIRV. The variance σ_w^2 in (25) is such that $\sigma^2 = E[s^2]\, \sigma_w^2$ and $E[s^2]$ is defined as

$$E[s^2] = \int_0^\infty s^2 p_S(s) ds \;. \tag{27}$$

An analytic expression for h_{2N} can be obtained from the previously defined envelope pdf of $\mathbf{v}$, $p_E(e)$, using the equation [2]

$$h_{2N}(y) = (-2)^{N-1} \frac{d^{N-1}}{dy^{N-1}} \left[\sigma_w y^{-1/2}\, p_E(\sigma_w y^{1/2})\right] \;. \tag{28}$$

It is clear from the previous equation that for a given pdf envelope $p_E(e)$ to correspond to a valid SIRV, $h_2(y)$ and its higher order derivatives must be alternatively positive decreasing and negative increasing, starting with $h_2(y)$ positive decreasing. This requirement ensures that $h_{2N}(y) \ge 0$ for all $N = 1, 2, \ldots$.

Equations (25) and (26) suggest a physical interpretation of a white SIRV as a generalized Gaussian random vector obtained by the product of a Gaussian random vector $\mathbf{w}$ with zero mean and covariance $\sigma_w^2\mathbf{I}$, by a nonnegative random variable s

which is statistically independent from w and has a pdf $p_S(s)$ that coincides with the characteristic pdf of the SIRV.

In general, the modulating rv s is such that $E[s^2] = 1$ and, hence, the power of the white SIRV componentes σ^2 coincides with the power σ_w^2 of the premodulated white Gaussian noise components.The characteristic pdf for given desired SIRV statistics (corresponding to an envelope pdf which is admissible according to the monothonicity criteria) may be normally obtained [2] from the solution of the integral equation in (26) for $N = 1$.

The appropriate $p_S(s)$, already normalized for unit mean square ($E[s^2] = 1$), for the K-distributed envelope is given by the generalized chi-pdf [2]

$$p_S(s) = \frac{2\nu^\nu}{\Gamma(\nu)} s^{2\nu-1} \exp(-\nu s^2) \quad s \geq 0 \qquad (29)$$

where ν is the shape parameter of the K-envelope as in (23). The corresponding h_{2N} is tabulated in [2] as

$$h_{2N}(y) = \frac{b^{2N}}{\Gamma(\nu)} \frac{(b\sqrt{y})^{\nu-N}}{2^{\nu-1}} K_{N-\nu}(b\sqrt{y}) \ . \qquad (30)$$

In order to generate a $2L$ x 1 noise vector with K-distributed envelope pdf, we generate windows of white Gaussian noise of size $2N$ (with $N << L$) and modulate each windows with a random number independently drawn from $p_S(s)$ in (31). Clearly, if we use a sufficiently large number of white Gaussian windows and subsequently compute $e_k = \sqrt{v_{c_k}^2 + v_{s_k}^2}$ over all windows, the actual histogram of e_k should then closely match the desired K distribution. The match is illustrated in figure 3, where the histogram of the envelope was computed over 8000 generated complex quadrature clutter samples and compared to the K-pdf in (23), with $\nu = 1.5$. We normalize the power of the quadrature components to 1 by setting the parameter b in (23) to $\sqrt{2\nu}$.

For a Weibull envelope as in (21), SIRV admissibility holds for the parameter range $0 \leq b \leq 2$. The simulation, however, is more complex since there is no closed form solution for the integral equation in (26). We resort then to an alternative technique discussed in [3] which is based on the SIRV representation theorem in hyperspherical coordinates. If a $2N$ x 1 white SIRV is represented in hyperspherical coordinates, the statistics of the hyperphases are invariant to the choice of the SIRV. Hence, any realization of hyperphases obtained from a white Gaussian random vector with zero mean and identity covariance can be used to generate a white SIRV with the same mean and covariance. The desired SIRV realization and the white Gaussian noise realization will differ only in the hyperradius coordinate, which is related to the SIRV metrics by the equation [3]

$$p_R(r) = \frac{r^{2N-1}}{2^{N-1}\Gamma(N)} h_{2N}(r^2) \ \ r \geq 0 \ . \qquad (31)$$

In order to illustrate the match between the desired Weibull envelope and the envelope statistics in the computer generated clutter, again we generate independent white noise windows of size $2N$ and modulate each window by the corresponding R/R_w where the modulating coefficients R are random numbers drawn independently for each window from the appropriate $p_R(r)$ and R_w is the norm of the white noise window. Figure 4 shows the very good match between the histogram of the envelope and the Weibull pdf in (21) for 8000 generated samples.

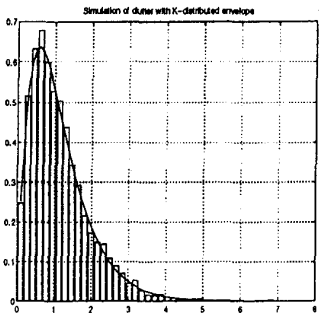

Figure 3: Observed and desired envelope statistics for K clutter

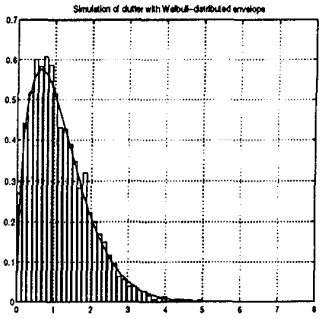

Figure 4: Observed and desired envelope statistics for Weibull clutter

Detection and Tracking Examples

In the following examples we assume the simplest possible model for the target, i.e., a deterministic, time-invariant return of the form

$$[s_{c_n}^k \ s_{s_n}^k] = [0\,0] \quad k \neq z_n$$

$$[s^k_{c_n}\ s^k_{s_n}] \quad = \quad [1\,0] \quad k = z_n \qquad (32)$$

where z_n is the position of the target in the discrete lattice $\mathcal{L}$ at instant n.

In order to generate the measurement vector $\mathbf{v}_n$ for each sensor scan, we simply extract L/N statistically independent SIRV windows from the 8000 samples we used to draw the histograms in figure 5. The window size was set to $N = 2$ (or 4 complex returns).

We present initially an example corresponding to tracking/detection of a single target against a white K-Noise background. The parameters of the noise are $\nu = 1.5$ and $b = 1.7312$. The sensor is assumed to have 64 resolution cells, corresponding to 128 returns at each time scan. The simulation includes 60 time scans. As before, there is only one target present at each time scan but, once a target disappears from the sensor range, another target can appear randomly at any resolution cell with a 30 per cent probability. The average deterministic drift of a present target corresponds to 4 cells and the fluctuation probability of one cell around the deterministic displacement is 0.40.

Figure 5 illustrates the performance of the tracker over 60 time scans with SNR = 6 dB. The symbol '+' represents the correct position, while the symbol 'o' represents the output of the tracker. Notice that even under this extreme noise conditions, the overall performance of the tracker is fairly good. Early misses when a new target appears are corrected as more data becomes available to the tracker.

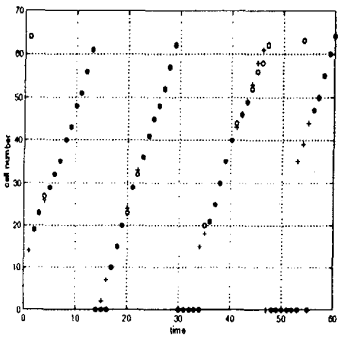

Figure 5: Single target tracking in white K-noise, SNR=6 dB

In the sequel, we present an example where a single target is tracked against white Weibull clutter with parameters $a = 0.6777$ and $c = 1.5$. The real trajectory is now plotted as a solid line when the target is present. The signal-to-noise ratio for a given frame is 10 dB. Notice that, except for a miss at instant 17, we achieve perfect detection/tracking with very high immunity to false alarms.

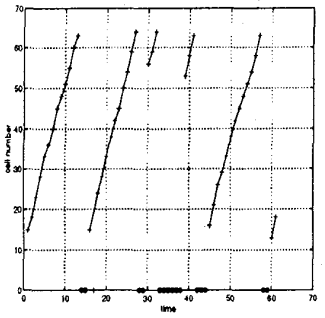

Figure 6: Single target tracking in white Weibull noise, SNR=10 dB

4 Conclusion

In this paper, we present an optimal pdf joint detector/tracker for a single target moving randomly on a finite discrete grid. Simulation examples assuming both white Gaussian background noise and non-Gaussian clutter with K and Weibull envelopes indicate good detection/tracking performance even in adverse noisy scenarios.The algorithm outperforms conventional trackers such as the peak detector or the linearized KBF. Examples with multiple targets and spatially correlated noise are presented in future papers.

References

[1] Roy E. Bethel and George J. Paras. A PDF Multitarget Tracker. *IEEE Trans. on Aerospace and Electronic Systems* vol.30, n.2, pp 386-403, April 1994.

[2] E. Conte, M. Longo and M.Lops. Modelling and simulation of non-rayleigh radar clutter. *IEEE Proceedings-F* Vol. 138, No 2, April 1991, pp 121-130.

[3] M. Rangaswamy, D. Weiner and A. Öztürk. Computer Generation of Correlated Non-Gaussian Radar Clutter. *IEEE Trans. on Aerospace and Electronic Systems* Vol 31, No. 1, January 1995, pp 106-116.

[4] José M. F. Moura and Carlos A. C. Belo. Threshold Extension by Nonlinear Techniques. *Underwater Acoustic Data Processing*, pp 433-452, 1989.