Hindawi Publishing Corporation
EURASIP Journal on Advances in Signal Processing
Volume 2008, Article ID 326259, 9 pages
doi:10.1155/2008/326259

Research Article

Track-before-Detect Using Swerling 0, 1, and 3 Target Models for Small Manoeuvring Maritime Targets

Michael McDonald and Bhashyam Balaji

Surveillance Radar Group, Defence R&D Canada, 3701 Carling Avenue, Ottawa, ON, Canada K1A 0Z4

Correspondence should be addressed to Michael McDonald, mike.mcdonald@drdc-rddc.gc.ca

Received 13 April 2007; Revised 25 September 2007; Accepted 26 December 2007

Recommended by Lawrence Stone

Real-radar data containing a small manoeuvring boat in sea clutter is processed using a finite difference (FD) implementation of continuous-discrete filtering with a four-dimensional constant velocity model. Measurement data is modelled assuming a Rayleigh sea clutter model with embedded Swerling 0, 1, or 3 target signal models. The results are examined to obtain a qualitative understanding of the effects of using the different target models. The Swerling 0 model is observed to exhibit a heightened sensitivity to changes in measured signal strength and provides enhanced detection of the maritime target examined at the cost of more peaked or multimodal posterior density in comparison with Swerling 1 and 3 targets.

1. INTRODUCTION

Traditional approaches to detecting and tracking maritime targets typically rely on a multiphase strategy in which target detections are fed to a Kalman-filter-based tracker. This is commonly combined with a simple linear track-before-detect (TkBD) scheme to provide noncoherent integration prior to the final detection step.

For highly manoeuvrable targets, the above approach provides poor performance due to the failure of the target dynamics to meet the constant velocity requirements. In addition, the methodology also fails to utilize all available target signal information as a "hard" detection decision must be made prior to tracking via a Kalman filter with a linear measurement function. The loss of information due to pretracking detection is particularly detrimental for targets possessing very low signal-to-interference ratios (SIRs).

Two approaches which are commonly used for solving nonlinear continuous-discrete problems with nonanalytical measurement functions are particle filters (PFs) and grid-based methods (see, e.g., [1, 2]). In this study, we focus on a grid-based approach and in particular a finite difference (FD) algorithm. Unlike simple TkBD methods, the FD approach utilizes the complete measurement set from each scan to evolve the complete state conditional probability density. For radar surveillance, the "measurement" at each time step corresponds to the filtered amplitude radar echoes or returns from all range-azimuth bins within a scanned sector.

The solution of the TkBD problem based on the continuous-discrete filtering and continuous-continuous filtering (based on the Duncan-Mortensen-Zakai equation) has previously been studied in the context of SAR [3], and IR [4] images, or ground moving target indicator (GMTI) radar measurements [5]. However, these studies utilized simulated targets and/or simulated clutter.

2. THEORY

2.1. Continuous-discrete filtering

In continuous-discrete filtering theory (see, e.g., [2]), the state model is given by the Itô stochastic differential equation of the form

$$d\mathbf{x}(t) = f(\mathbf{x}(t), t)dt + e(\mathbf{x}(t), t)d\mathbf{v}(t). \tag{1}$$

Here, $\mathbf{x}(t)$ is an $\mathbb{R}^n$-valued process, $f(\mathbf{x}(t), t) \in \mathbb{R}^n$, $e(\mathbf{x}(t), t) \in \mathbb{R}^{n\times p}$, and $\mathbf{v}(t)$ is an $\mathbb{R}^p$-valued Brownian process

with covariance $Q(t)$. The forward diffusion operator, $\mathcal{L}$, of the state process generated by (1) is given by

$$\mathcal{L}(\cdot) = -\sum_{i=1}^{n} \frac{\partial(\cdot f_i)}{\partial x_i} + \frac{1}{2}\sum_{i,j=1}^{n} \frac{\partial^2\left[\cdot (eQe^T)_{ij}\right]}{\partial x_i \partial x_j}. \tag{2}$$

The general continuous-discrete filtering problem considers the following signal and measurement processes:

$$\begin{aligned} d\mathbf{x}(t) &= f(\mathbf{x}(t), t)dt + e(\mathbf{x}(t), t)d\mathbf{v}(t), \\ \mathbf{y}(t_k) &= h(\mathbf{x}(t_k), t_k, \mathbf{w}(t_k)). \end{aligned} \tag{3}$$

Here, $\mathbf{y}$ is an $\mathbb{R}^m$-valued process, $h(\mathbf{x}(t), t, \mathbf{w}(t)) \in \mathbb{R}^m$, and $\mathbf{w}(t)$ is an $\mathbb{R}^q$-valued Brownian process.

The continuous-discrete filtering problem is solved as follows. Let the initial distribution be $\sigma_0(x)$ and let the measurements be collected at time instants $t_1, t_2, \ldots, t_k, \ldots$. We use the notation $Y(\tau) = \{y(t_l) : t_0 < t_l \le \tau\}$. Then, at observation at time t_k, the conditional density is given by

$$p(t_k, x|Y(t_k)) = \frac{p(y(t_k)|x)\, p(t_k, x|Y(t_{k-1}))}{\int p(y(t_k)|\xi)\, p(t_k, \xi|Y(t_{k-1}))\{d^n\xi\}}, \tag{4}$$

and $p(t_k, x|Y(t_{k-1}))$ is given by the solution of the Fokker-Planck-Kolmogorov forward equation (FPKfe)

$$\frac{\partial}{\partial t} p(t, x|Y(t_{k-1})) = \mathcal{L}(p(t, x|Y(t_{k-1}))), \quad t_{k-1} \le t < t_k, \tag{5}$$

with initial condition $p(t_{k-1}, x|Y(t_{k-1}))$. Often, the signal and measurement model is described by the following system:

$$\begin{aligned} d\mathbf{x}(t) &= f(\mathbf{x}(t), t)dt + e(\mathbf{x}(t), t)d\mathbf{v}(t), \\ \mathbf{y}(t_k) &= h(\mathbf{x}(t_k), t_k)dt + \mathbf{w}(t_k), \quad k = 1, 2, \ldots, \end{aligned} \tag{6}$$

where $\mathbf{y}(t) \in \mathbb{R}^{m\times 1}$, $h \in \mathbb{R}^{m\times 1}$, and the noise process is described by $\mathbf{w}(t)$.

2.2. Signal and measurement models

The state model we consider is the constant velocity (CV) model on the plane so that the resulting state model is four-dimensional. If

$$\left[\ \mathbf{x}_1(t)\ \ \mathbf{x}_2(t)\ \ \mathbf{x}_3(t)\ \ \mathbf{x}_4(t)\ \right] = \left[\ x(t)\ \ v_x(t)\ \ y(t)\ \ v_y(t)\ \right], \tag{7}$$

then the model is

$$\begin{bmatrix} d\mathbf{x}_1(t) \\ d\mathbf{x}_2(t) \\ d\mathbf{x}_3(t) \\ d\mathbf{x}_4(t) \end{bmatrix} = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} \mathbf{x}_1(t) \\ \mathbf{x}_2(t) \\ \mathbf{x}_3(t) \\ \mathbf{x}_4(t) \end{bmatrix} + \begin{bmatrix} 0 & 0 \\ 1 & 0 \\ 0 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} \sigma_2 d\mathbf{v}_2(t) \\ \sigma_4 d\mathbf{v}_4(t) \end{bmatrix}. \tag{8}$$

The FPKfe for the CV model (where $Q(t) = 1$) is

$$\frac{\partial u}{\partial t}(t, x) = \left(\frac{\sigma_2^2}{2}\frac{\partial^2}{\partial x_2^2} + \frac{\sigma_4^2}{2}\frac{\partial^2}{\sigma x_4^2} - x_2\frac{\partial}{\partial x_1} - x_4\frac{\partial}{\partial x_3}\right)u(t, x). \tag{9}$$

The measurement model is specified by $p(y(t_k)|x)$, where a measurement $y(t_k)$ is the return amplitude on a grid.

Another possible state model is the integrated Ornstein-Uhlenbeck model (see, e.g., [6]). This has the nice property that the variance of the velocity is bounded, which is especially relevant if no measurements are available over a long period of time. In the data that was processed in this study, the measurements were available at short-time intervals, and the CV model was found to be adequate.

2.3. Multiplicative operator splitting

The solution of the continuous-discrete filtering problem thus requires the solution of a PDE of the following form:

$$\frac{\partial u}{\partial t}(t, x) = \sum_{i=1}^{s} \mathcal{L}_i u(t, x). \tag{10}$$

For the CV model we have

$$\begin{aligned} \mathcal{L}_1 &= \frac{\sigma_2^2}{2}\frac{\partial^2}{\partial x_2^2}, \qquad \mathcal{L}_2 = \frac{\sigma_4^2}{2}\frac{\partial^2}{\partial x_4^2}, \\ \mathcal{L}_3 &= -x_2\frac{\partial}{\partial x_1}, \qquad \mathcal{L}_4 = -x_4\frac{\partial}{\partial x_3}. \end{aligned} \tag{11}$$

In the forward Euler explicit scheme (see, e.g., [7]), (10) is numerically solved using the following approximation:

$$\frac{u(t + \Delta t, x) - u(t, x)}{\Delta t} = \sum_{i=1}^{s} \mathcal{L}_i u(t, x), \tag{12}$$

so that

$$\begin{aligned} u(t + \Delta t, x) &= \left(1 + \Delta t \sum_{i=1}^{s} \mathcal{L}_i\right) u(t, x) \\ &\approx \prod_{i=1}^{s} (1 + \Delta t \mathcal{L}_i) u(t, x) + O(\Delta t)^2. \end{aligned} \tag{13}$$

The advantage of the multiplicative form is that it reduces the s-dimensional problem into s one-dimensional problems thus resulting in memory and computational savings.

Note that if the time step is too large, the explicit scheme is unstable. This problem is evaded by splitting up the time interval between measurements into N_T time steps prior to applying the forward Euler scheme, that is,

$$(1 + \Delta t \mathcal{L}_i) \approx \left(1 + \frac{\Delta t}{N_T}\mathcal{L}_i\right)^{N_T}. \tag{14}$$

Furthermore, stability of the discretization of the convection operators requires that "upwind differencing" be used for the first order derivative operator [7]. This also ensures that the probability remains positive.

Figure 1: Speedboat used as target in trials.

Table 1: Radar and aircraft operating parameters.

Radar resolution	<1 meter
Radar scan rate	36°/s
Sector width	100°
Pulse repetition frequency (PRF)	500 Hz
Aircraft altitude	1000 feet

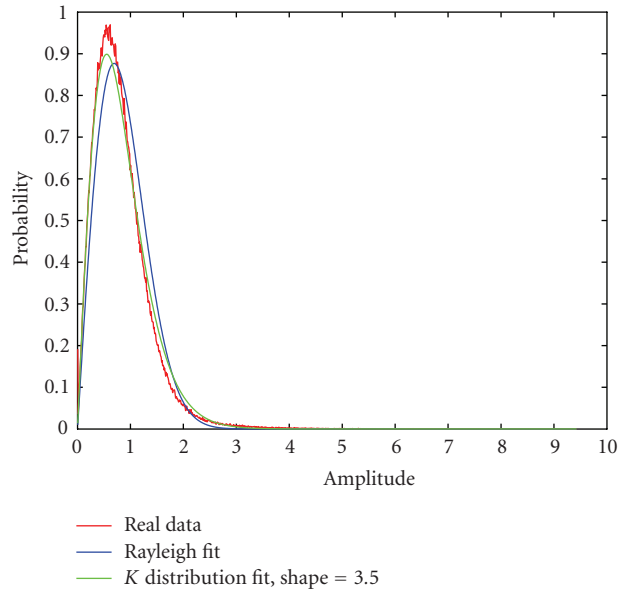

Figure 2: Histogram of real-data distribution and plots of Rayleigh and K distribution probability distribution functions with variance and/or shape parameter matched to real data.

Often, the backward Euler (or Laasonen) implicit scheme is used and the following approximation is made:

$$\frac{u(t+\Delta t,x)-u(t,x)}{\Delta t} = \Delta t \sum_{i=1}^{s} \mathcal{L}_i u(t+\Delta t,x), \tag{15}$$

or

$$u(t+\Delta t,x) = \prod_{i=1}^{s} (1-\Delta t\mathcal{L}_i)^{-1} u(t,x). \tag{16}$$

Since the matrices are tridiagonal, the inverses may be computed eficiently using the Thomas tridiagonal method [8]. This implicit scheme has the advantage of being stable even for large time steps. However, it is not as accurate as the explicit scheme. Since the time-step restrictions on the explicit scheme is not severe in this application, we used the explicit scheme.

Finally, the Dirichlet boundary conditions are applied which ensures that the probability vanishes at the boundary.

3. DATA DESCRIPTION

The data used in this study was collected near the mouth of Halifax harbour in Nova Scotia, Canada using the DRDC Ottawa X-band Wideband Experimental Airborne Radar (XWEAR). A small, highly manoeuvrable speedboat, see Figure 1, was fielded. The relevant radar and aircraft operating parameters are given in Table 1.

After pulse compression, the radar returns were normalized using a cell averaging (CA) approach so as to remove large scale fluctuations in underlying power levels (see, e.g., [9]). In particular, the estimated mean power of the cell under test (CUT) was calculated using a total of 256 range bins equally distributed on both sides of the CUT with a guard band of 30 cells to prevent self-nulling of the target signal via target contamination of the mean. The choice of 256 as a background sample dimension was arrived at through trial and error experimentation against a range of smaller and larger background sizes. For this data set, the chosen values of background and guard band size produced the best results.

Figures 2 and 3 present a histogram of the measured clutter returns for the data set and the corresponding Rayleigh probability distribution functions (pdfs) matched to the variance of the real data. It is evident that the Rayleigh distribution is not a particularly good match to the real data, which exhibits a much longer tail corresponding to a greater probability of high-amplitude outliers. This tail is commonly observed in high-resolution sea clutter measurements and is caused by the presence of sea spikes. Past studies have shown that the K distribution often provides a better fit to the real data [10, 11]. The corresponding K distribution is also shown in Figures 2 and 3. The measured shape parameter of 3.5 was calculated from the real data using the $z \log z$ method [12].

Figures 4 and 5 present the measured boat velocity and change of velocity, respectively, at each time step as measured using onboard GPS. The bearing and change in bearing is also plotted. It can be seen that the boat was manoeuvring strongly. It should also be noted that for approximately the first 10 scans the boat was moving very slowly after which time it rapidly accelerated to a velocity of greater then 10 m/s (20 knots). A final observation is made on the SIR of the boat. During the early portion of the data set, the signature of the boat is much less visible against the clutter background (not

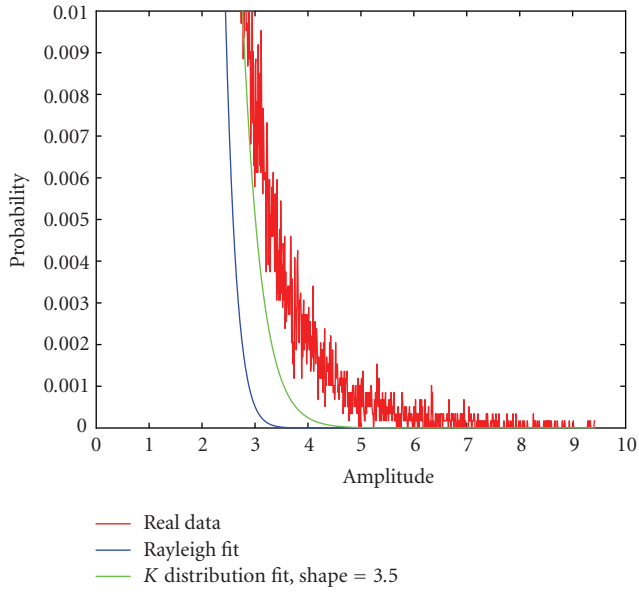

Figure 3: Close up of Figure 2 on high-amplitude tail of distribution.

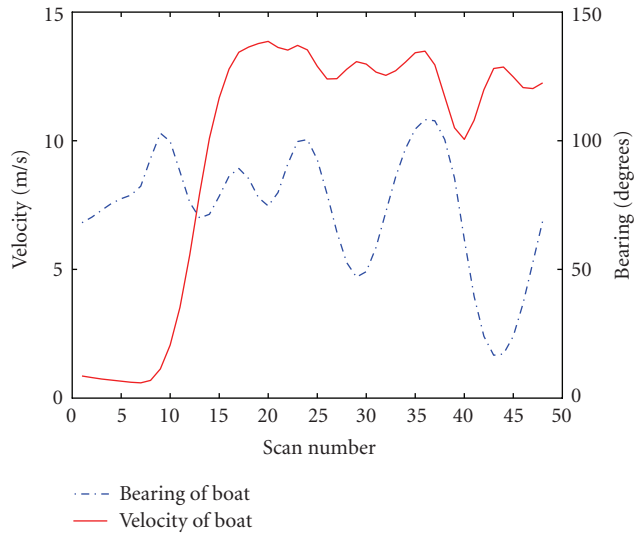

Figure 4: Velocity and bearing of boat for each scan during trial collection period.

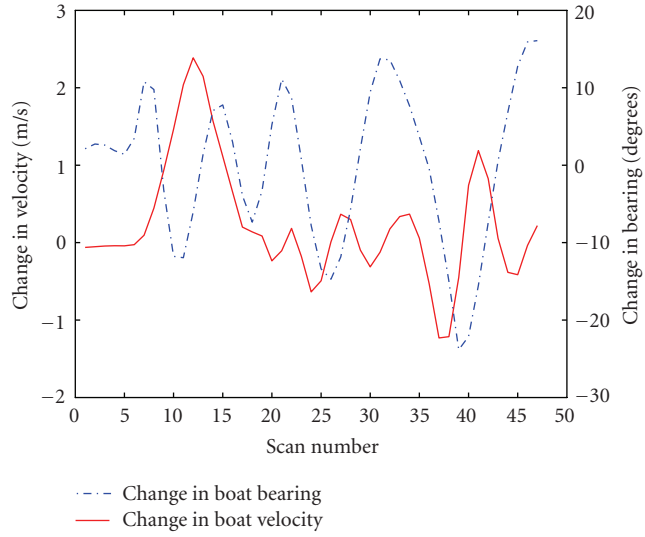

Figure 5: Change of velocity and bearing of boat between scans of trial collection period.

shown). After approximately 15 scans, the relative strength of the target signal is seen to increase greatly with respect to the clutter background, probably due to a combination of changes in incidence and viewing angle.

4. IMPLEMENTATION OF THE MEASUREMENT CORRECTION

The measurement correction corresponds to the application of $p(y(t_k)|x)$ in (4). Implementation of the measurement correction is practically difficult due to the departure of real-life target and clutter characteristics from the analytically tractable stochastic models that must be used. In this paper, we will examine the application of a variety of commonly used target pdfs against the real data. To model the radar clutter we use the well-known Rayleigh pdf. As discussed above this choice tends to underrepresent the proportion of high-amplitude returns observed in the real data. While the K distribution appears to offer a slightly better fit, it does not permit a closed form expression for the signal-plus-noise pdf and requires the use of a numerical integration. This imposes a significant computational load and has not been implemented in this study. It will be examined in future studies. The implications of choosing the Rayleigh model are discussed further below.

A common choice for a target model is the Swerling 0, or constant amplitude, target model. Unfortunately, highly manoeuvrable small cross-section targets tend to undergo very large cross-section fluctuations between scan-to-scan and even between pulse-to-pulse measurements. In recognition of this limitation, Rutten et al. [13] suggested the application of the fading Swerling target models in Rayleigh clutter. In this paper, we implement models for Swerling 1 and Swerling 3 targets in Rayleigh clutter and compare it with Swerling 0 target results. While Rutten et al. indicated that the application of Swerling 1 and Swerling 3 models would necessitate the use of a numerical integral, this is not strictly true for the case where the measurement samples are statistically independent. For this case, simple closed form expressions are available for the Swerling 0 and Swerling 3 pdfs [14]. In general, the Swerling 1 target model is applicable to a complex target comprised of numerous independent scatterers of similar cross-section, while Swerling 3 is representative of a target comprised of one large scatterer and numerous smaller cross-section scatterers. The corresponding pdf for each target type is shown in Figure 6.

It is easily shown that calculating $p(y(t_k)|x)$ is equivalent to the calculation of the product of the likelihood ratios for

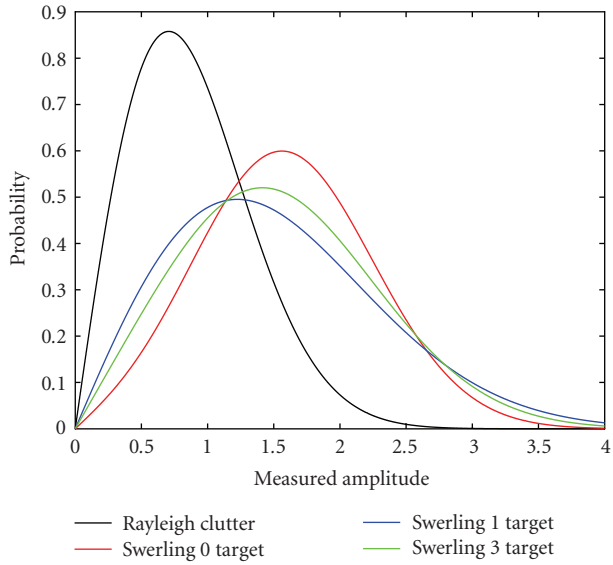

FIGURE 6: Sample probability distribution function curves for pure Rayleigh clutter and Swerling 0, 1, and 3 targets in Rayleigh clutter. Average target signal power is 3 dB above clutter power.

each measurement point within the target-spread function to within a common normalization factor [15]. The following sections detail the likelihood function corresponding to each of the clutter-plus-target models utilized in this study.

For all models the spread function of the target, h_n^i, is assumed to be Gaussian and is given by

$$h_n^i = I\exp\left[-\frac{(x^i - x^n)^2 + (y^i - y^n)^2}{2\sigma_s^2}\right], \tag{17}$$

where x^i and y^i are the x- and y-locations of the current measurement and x^n and y^n are the x- and y-locations of the current state point. I is the postulated amplitude of the target and σ_s^2 is the variance of the Gaussian spread function. In regions where no component of the target spread function exists, the likelihood ratio reduces to a value of 1. It should again be emphasized that $p(y(t_k)|x)$ corresponds to the product of all the likelihood ratios formed from all the measurement locations (x^i and y^i) within the spread function, that is,

$$l(y(t_k)|x^n(t_k)) = \prod_i \frac{p(y^i(t_k)|x^n(t_k), H_1)}{p(y^i(t_k)|x^n(t_k), H_0)}, \tag{18}$$

where H_1 and H_0 denote the target "present" and "not present" cases, respectively.

4.1. Rayleigh distribution model with swerling 0 target

The pdf for a Swerling 0 target in Rayleigh clutter is given by

$$p(y^i(t_k)|x^n(t_k), H_1) = \frac{2y^i(t_k)}{P}\exp\left(-\frac{y^i(t_k)^2 + (h_n^i)^2}{P}\right) I_0\left(\frac{2h_n^i y^i(t_k)}{P}\right), \tag{19}$$

(see, e.g., [13]) where I_0 is a modified Bessel function of the first kind and P is the average Rayleigh clutter power determined by calculating the mean power across the measurement frame. The pdf for Rayleigh clutter without a target is given by

$$p(y^i(t_k)|x^n(t_k), H_0) = \frac{2y^i(t_k)}{P}\exp\left(-\frac{y^i(t_k)^2}{P}\right). \tag{20}$$

The corresponding likelihood ratio for the Rayleigh case is therefore given by

$$l(y^i(t_k)|x^n(t_k)) = \exp\left(-\frac{(h_n^i)^2}{P}\right) I_0\left(\frac{2h_n^i y^i(t_k)}{P}\right). \tag{21}$$

4.2. Rayleigh distribution model with swerling 1 target

In this case,

$$p(y^i(t_k)|x^n(t_k), H_1) = \left(1 + \frac{(h_n^i)^2}{P}\right)\exp\left(-\frac{y^i(t_k)^2}{P + (h_n^i)^2}\right), \tag{22}$$

where the intensity used in calculation now corresponds to the average target intensity [14]. The corresponding likelihood ratio is formed using (22) and (20).

4.3. Rayleigh distribution model with swerling 3 target

In this instance,

$$p(y^i(t_k)|x^n(t_k), H_1) = \frac{1}{\left(1+(h_n^i)^2/2P\right)}\left[1+\frac{y^i(t_k)^2/P}{\left(1+2P/(h_n^i)^2\right)}\right] \times \exp\left[-\frac{y^i(t_k)^2/P}{\left(1+(h_n^i)^2/2P\right)}\right], \tag{23}$$

where the intensity used in calculation now corresponds to the average target intensity [14]. The corresponding likelihood ratio is formed using (23) and (20).

4.4. Comments

To calculate the spread function, the target intensity, I, is required. This is typically not known a priori although a reasonable estimate may be formed through knowledge of desired target types. Methodologies for choosing an optimum intensity value are not examined in this paper, rather, an optimum intensity value is determined by trial and error. For the purposes of this study, the optimum intensity is considered to be that which achieves target detection on the maximum number of scans. The concept of a "detection" is discussed in further detail in Section 5 below. In all cases, the target intensity is held constant across all scans.

Even presuming the optimum intensity value has been accurately chosen, significant performance degradation can

still occur due to the mismatch between the real and postulated target-plus-clutter models. Significant problems arise due to the enhanced high-amplitude tail of the real clutter that was observed in Figure 2. Since the Rayleigh distribution does not "anticipate" this increased prevalence of high-amplitude clutter spikes, its produces a larger likelihood ratio than is warranted.

The result is a filtered state probability distribution, which fails to maintain a "lock" on the actual target location. Instead, the locations of the state-distribution peaks fluctuate rapidly from scan to scan, the latest peak location corresponding to the most recent clutter spike. The problem is compounded for very large spread functions, that is, large number of measurement points. In this study, the observed 3 dB width of the target spread function contains over 600 separate measurement points. Ideally, more measurements will result in greater integration gain, but in practise, the mismatch between the real and postulated target-plus-clutter models introduces a small error to each calculated likelihood ratio. These small errors translate into huge errors when the overall likelihood ratio is formed from the product of all individual likelihood ratios. The resulting overall likelihood ratios can be many orders of magnitude larger than warranted. Care was taken during the processing in this analysis to restrict the extent of the spread function in order to prevent this exponential growth of overall likelihood ratio error. In addition, a simple ad hoc approach was adopted in this study whereby the value of the overall likelihood function was limited to a maximum upper value so as to suppress the effect of the largest clutter spikes. This approach was seen to provide a significant improvement in performance. All three target models were observed to exhibit similar performance gains when likelihood limiting was employed.

It is instructive to further examine the properties of the Swerling target models with respect to one another. Figure 6 presents the pdf curves for Swerling 0, 1, and 3 targets in Rayleigh clutter. As anticipated, the Swerling 1 and 3 distributions exhibit a larger variance than the Swerling 0 target model, with the Swerling 1 being the broadest. The effect of increased variance on the calculated likelihood ratio is illustrated in Figures 7 and 8.

From Figure 7, it is observed that likelihoods ratios of the Swerling 0 model are most sensitive to the average target power assumption (corresponding to choosing intensity, I, in the spread function calculation detailed above) while the Swerling 1 model is least sensitive. In a practical implementation, the choice of a correct target power is not a trivial task and the apparent desensitization of this parameter could prove highly advantageous. The flipside of the relationship is illustrated in Figure 8 where the likelihood ratio is plotted against received power for a fixed underlying target power. It is readily observed that the Swerling 0 target model provides the most aggressive promotion and demotion of the posterior density. Namely, when a weak signal power is observed, the posterior density is more heavily suppressed (i.e., the likelihood ratio is less than one); while for strong signals, the posterior is most strongly promoted (i.e., likelihood ratio greater then one). While the difference in individual likelihood ratios between Swerling models may appear small, it can become very large when the product of a large number of likelihoods ratios is computed to determine the overall likelihood ratio.

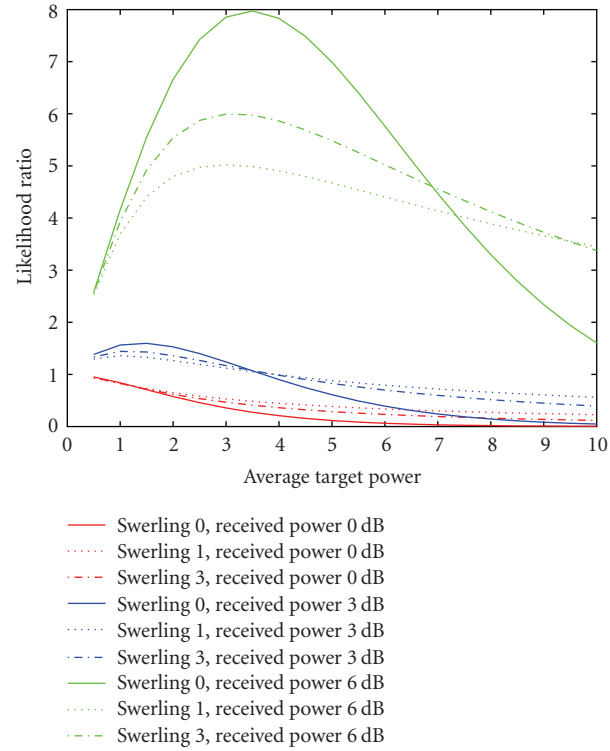

Figure 7: Variation of calculated likelihood ratio against postulated average target power for received measurement powers of 0 dB, 3 dB and 6 dB above clutter power. Curves are shown for Swerling 0, 1, and 3 target models.

Neither of the above behaviors is unexpected but must be fully appreciated when comparing results via different target models. In practical terms, one would expect to find that the posterior distributions associated with Swerling 0 targets would be more strongly peaked than those with the Swerling 1 and 3 targets. The broader the variance of the underlying target model, the flatter the expected posterior distribution. The temporal behavior of the posterior evolution is also likely to differ, the low variance Swerling 0 model will likely show a greater sensitivity to anomalous clutter peaks (with an associated peak in the posterior) but the signature of transient events will more rapidly decay with time due to the greater subsequent suppression.

5. RESULTS

The data was processed across a 1.5 km by 1.25 km region with 100 grid steps in the x- and y-directions (i.e., 10^4 state points). For a CV model, two additional dimensions, corresponding to the x- and y-velocities, are also required. A relatively coarse velocity spacing corresponding to ± 40 m/s spread across 10 grid steps along the velocity dimensions was used. The CV state grid thus comprises 10^6 state points.

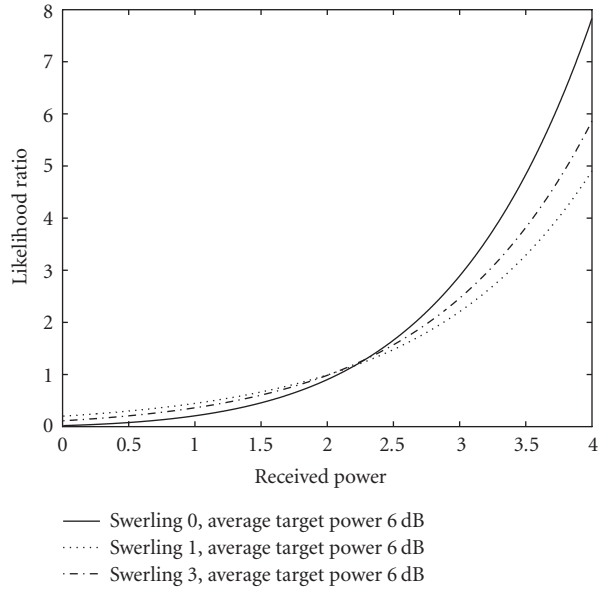

Figure 8: Variation of calculated likelihood ratios against received measurement power for postulated average target power 6 dB above clutter power.

A "detection" was determined as follows. The state density was summed across the velocity dimensions. The location of the maximum value of the collapsed state density was identified and the probabilities were then zeroed for all state points falling within ±5 grid points. This zeroing acts as a crude multiple detection pruning technique. The process can be repeated to extract progressively smaller maximum state localities. In the following text, detections will be referred to as a first maximum detection, second maximum detection, etc. to identify the order of extraction.

It should be noted that while this approach bears some passing resemblance to the concept of specifying a given PFA (i.e., allowing one mode to be chosen would crudely correspond to a PFA of 1/# of independent state locations, allowing three modes would correspond to 3/# of independent state locations, etc.), the approach is not strictly CFAR. In addition, this concept of modal "detection" differs from the more commonly utilized approach of applying a threshold to the measured target amplitude in a number of subtle but significant ways.

The first important difference is that the output of the TkBD processing is a posterior measure of the probability that the target is present at a given location within the field of view rather then some function of the current measured radar return. The posterior measure reflects the entire history of measurements up to that point in time, hence it is more representative of a track than of an individual detection; and the applicability of detection performance measures such as probability of detection (PD) and probability of false alarm rate (PFA) become much less clearly defined. Characterization of tracking performance is a much more difficult problem than detection performance and one is typically forced to resort to a broad range of measures such as false track rate, false track length, association changes, missed object history, omitted tracks, track establishment delay, and so on to quantitatively capture the results. In most cases, there is not one definitive combination of specifications that characterize ideal performance, rather the tracker designer must choose the mixture that best suits their application. This sort of analysis is beyond the scope of this study and is in fact not possible here due to the limited size of the data set and the inability to calculate statistically meaningful values.

As will be discussed later, it is envisioned by the authors that a practical implementation of TkBD will require a follow-on tracker stage to refine the tracklet input from the TkBD and identify the real target tracks. Under this scenario, the overall performance is an intimately coupled function of the TkBD and follow-on tracker design. Further investigations of these aspects are reserved as a topic for future study. The focus of the following discussion is to highlight the impact of utilizing different targets models and, in particular, understand the effect of this choice on the posterior and the distribution of modes within it. General observations of how posterior distributions and modal 'detections' are affected by target choice are presented but definitive statements on the superiority of one model with respect to another cannot be provided for the reasons discussed above.

Figure 9 compares results obtained using a Swerling 0, 1, and 3 target model. To generate this plot, the top three maximum detection localities have been identified on each scan. Only the detections that clearly coincide with the actual target location (as determined by GPS truth and raw signal intensity plots) are retained and plotted.

It is evident from Figure 9 that the Swerling 0 target provides the greatest detection performance as it successfully detects the target on almost all scans. The Swerling 1 and 3 models, which produce virtually identical results to each, suffer from large drop-out regions, particularly near the beginning of the data set, in which the target is not detected. At first glance, this result seems somewhat surprising as the Swerling 0 target model permits the smallest variation in target signal and would be expected to be least tolerant of changes in target strength across the data set. Figure 10 sheds further light on this discrepancy. In Figure 10, only the first maximum detection from each scan is extracted and only those detections corresponding to the actual target are plotted in the figure. The result is a strong degradation of Swerling 0 target model performance with respect to the Swerling 1 and 3 models in comparison with those observed in Figure 9. It is difficult to be definitive on the precise mechanisms at work due to the complicated environmental conditions, but some of the differences are likely explained by the enhanced promotion/demotion characteristics of the Swerling 0 target model discussed above. The ability of the Swerling 0 target model to provide second or third maximum detections of the target reflects its greater sensitivity to small changes in received signal strength. This effect should be evident as a strongly peaked posterior distribution. The difference between the target models is likely to be most pronounced during the earlier scans when the target signal was observed to be much weaker and less stable, and in fact,

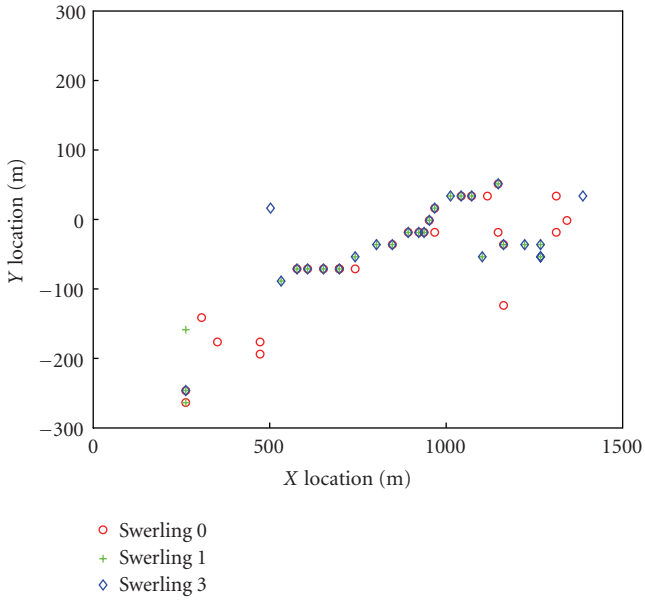

Figure 9: True target detections for Swerling 0, 1, and 3 target model when top three maximum localities are considered.

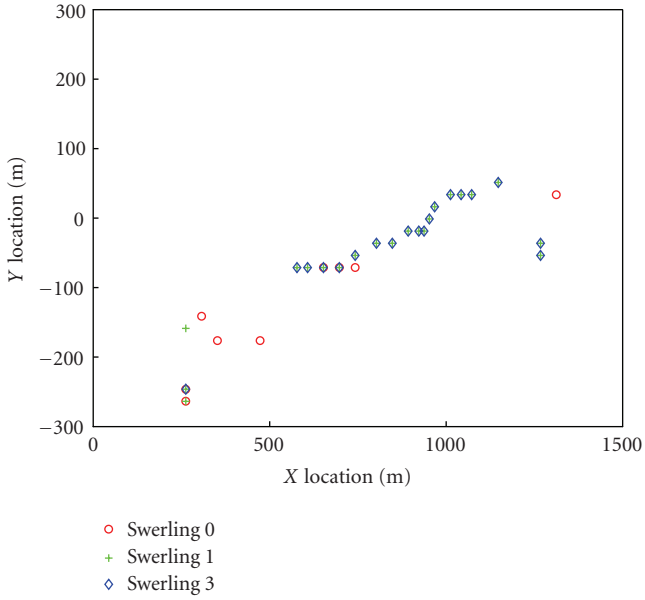

Figure 10: True target detections for Swerling 0, 1, and 3 target model when top three maximum localities are considered.

it is in this time frame that the Swerling 1 and 3 targets suffer from a detection drop-out. However, the enhanced sensitivity also means that the approach is more sensitive to clutter spikes, which leads to the increased prevalence of second or third maximum detections of the actual target in later scans.

Figures 11 and 12 present sample contour plots of the posterior density for one scan derived using the Swerling 0 and Swerling 1 target models, respectively. The Swerling 3 results are very similar to the Swerling 1 results and are

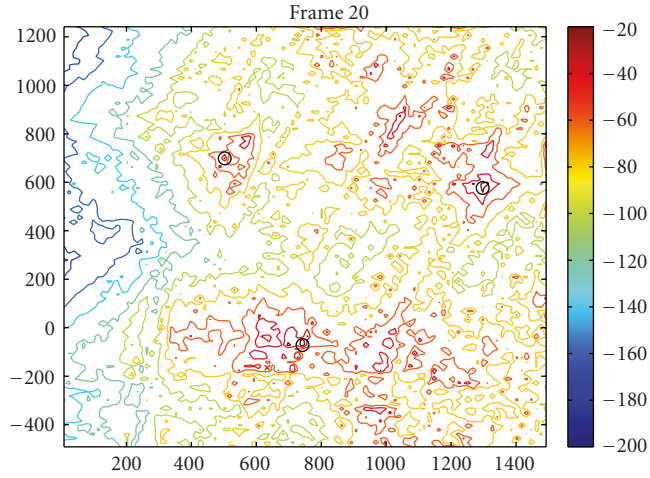

Figure 11: Contour plot of posterior density function for Swerling 0 model on selected scan.

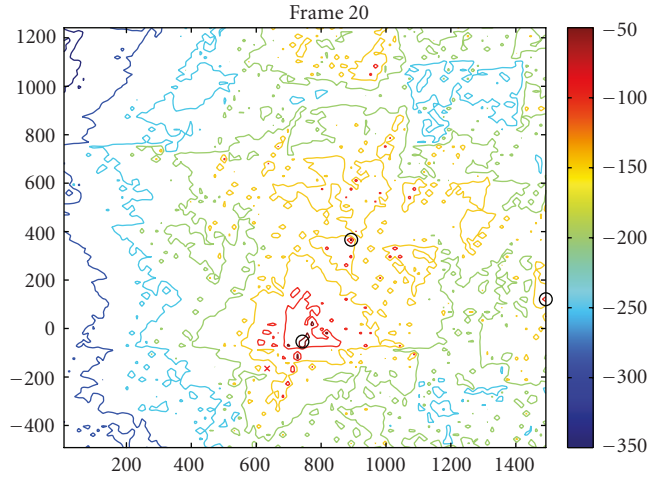

Figure 12: Contour plot of posterior density function for Swerling 1 model on selected scan.

omitted for conciseness. The enhanced multimodal character of the Swerling 0 results with respect to the Swerling 1 results is readily apparent and supports the discussion above.

6. CONCLUSION

The results of the TkBD nonlinear filtering using Swerling 0, 1, and 3 target models provide some insights into the applicability of the models for the detection of small manoeuvring targets in high-resolution sea clutter. The Swerling 0 model is observed to exhibit a heightened sensitivity to changes in measured signal strength, at least for the current data set. This provides enhanced detection of the maritime target but at the cost of more strongly peaked or multimodal posterior density. None of the Swerling models tested provides universally superior detection performance. The choice of Swerling model will likely need to be considered in conjunction with the design for any post-TkBD tracking that might be applied. The Swerling 0 model appears to be

most effective when several posterior peaks are identified as potential targets or tracklets, where it is recognized that many will represent false targets. However, this approach requires that the post-TkBD tracking algorithms have the capability to reliably promote the tracklets to firm track status or terminate them. Conversely, the use of the Swerling 1 or 3 target models may allow for a simplified detection and post-TkBD algorithm but at the cost of detection sensitivity.

REFERENCES

[1] B. Ristic, S. Arulampalam, and N. Gordon, Beyond the Kalman Filter: Particle Filters for Tracking Applications , Artech House, Norwood, Mass, USA, 2004.

[2] A. H. Jazwinski, Stochastic Processes and Filtering Theory , Dover Publications, New York, NY, USA, 2007.

[3] Z. S. Haddad and S. R. Simanca, "Filtering image records using wavelets and the Zakai equation," IEEE Transactions on Pattern Analysis and Machine Intelligence , vol. 17, no. 11, pp. 1069–1078, 1995.

[4] S. Kligys, B. Rozovsky, and A. Tartakovsky, "Detection algorithms and track before detect architecture based on nonlinear filtering for infrared search and track systems," *Tech. Rep. CAMS-98.9.1* , Center for Applied Mathematical Sciences, University of Southern California, Los Angeles, Calif, USA, 1998.

[5] K. Kastella and C. Kreucher, "Multiple model nonlinear filtering for low signal ground target applications," IEEE Transactions on Aerospace and Electronic Systems , vol. 41, no. 2, pp. 549–564, 2005.

[6] l. D. Stone, C. A. Barlow, and T. L. Corwin, Bayesian Multiple Target Tracking , Artech House, Norwood, Mass, USA, 1999.

[7] W. F. Ames, Numerical Methods for Partial Differential Equations , Academic Press, New York, NY, USA, 2nd edition, 1977.

[8] G. H. Golub and C. F. V. Loan, Matrix Computations , John Hopkins University Press, Baltimore, Md, USA, 2nd edition, 1989.

[9] B. C. Armstrong and H. D. Griffiths, "CFAR detection of fluctuating targets in spatially correlated K-distributed clutter," IEE Proceedings—Radar, Sonar and Navigation , vol. 138, no. 2, pp. 139–152, 1991.

[10] D. Walker, "Experimentally motivated model for low grazing angle radar Doppler spectra of the sea surface," IEE Proceedings—Radar, Sonar and Navigation , vol. 147, no. 3, pp. 114–120, 2000.

[11] K. D. Ward, C. J. Baker, and S. Watts, "Maritime surveillance radar—part 1: radar scattering from the ocean surface," IEE Proceedings—Radar, Sonar and Navigation , vol. 137, no. 2, pp. 51–62, 1990.

[12] D. Blacknell and R. J. A. Tough, "Parameter estimation for the K-distribution based on [z log(z)]," IEE Proceedings—Radar, Sonar and Navigation , vol. 148, no. 6, pp. 309–312, 2001.

[13] M. G. Rutten, N. J. Gordon, and S. Maskell, "Recursive track-before-detect with target amplitude fluctuations," IEE Proceedings—Radar, Sonar and Navigation , vol. 152, no. 5, pp. 345–352, 2005.

[14] D. P. Meyer and H. A. Mayer, Radar Target Detection , Academic Press, New York, NY, USA, 1973.

[15] D. J. Salmond and H. Birch, "A particle filter for track-before-detect," in *Proceedings of the American Control Conference (ACC '01)*, vol. 5, pp. 3755–3760, Arlington, Va, USA, June 2001.