Efficient Particle-Based Track-Before-Detect in Rayleigh Noise

Mark G. Rutten,[a,b] **Neil J. Gordon**[a] **and Simon Maskell**[c]

[a]Intelligence Surveillance and Reconnaissance Division, Defence Science and Technology Organisation
PO Box 1500, Edinburgh, SA 5111, Australia
[b]Cooperative Research Centre for Sensor Signal and Information Processing, Department of Electrical and Electronic Engineering, The University of Melbourne, Parkville, Victoria 3010, Australia
[c]QinetiQ Ltd., Malvern Technology Centre, St Andrews Rd, Malvern, Worcestershire, UK

Abstract – *Standard target state estimation schemes typically use detections as their source of measurements, which are produced by thresholding the output of a sensor's signal processing stage. This work exploits a track-before-detect (TBD) technique, which simultaneously detects and tracks a target without needing to threshold the sensor data. By removing the need for thresholding, TBD can potentially detect and track targets with a much lower signal-to-noise ratio (SNR) than conventional systems.*

The signal processing that is modelled in this work is designed to match that which might be found in a sensor such as radar. In these systems, the data used by a tracking filter is the magnitude of a complex spectrum. This gives rise to a signal in Rayleigh distributed noise. This paper presents a particle-based TBD filter, operating on the output of this signal processing stage, which estimates the target state incorporating target existence, position, velocity and target signal strength. The main contributions of this work are the development of an efficient method of calculating the probability of target existence and the derivation of a TBD filter which operates in Rayleigh noise. It is shown through simulation that a target can be detected and tracked with an SNR as low as 3dB.

Keywords: Track-before-detect (TBD), particle filter, target existence, Rayleigh, Ricean.

1 Introduction

A typical tracking filter will form detections by thresholding the output of the sensor's signal processing stage and then filtering the resulting detections. The detection threshold is generally chosen as a compromise between the potential number of false alarms and the number of missed detections, based on known noise statistics and signal models. This implies that targets with a low signal-to-noise ratio (SNR) may not be reliably detected if the signal power is below the required threshold. Track-before-detect (TBD) techniques eliminate the need for a detection threshold, simultaneously detecting and tracking targets with much lower signal-to-noise ratios than conventional trackers.

Particle filter implementations of track-before-detect were introduced by Salmond et. al. [1, 2], giving several advantages over previous approaches, including a reduction in computation complexity. Ristic [3] has extended the work of Salmond by providing a detailed explanation of an implementation, giving results of the detection performance of a particle filtering TBD algorithm and deriving the Cramér-Rao lower bound on the estimation error for the technique.

The work of Boers and Driessen [4] develops a TBD algorithm similar in structure to Ristic but extended to consider multiple targets. The algorithm presented in this paper is based heavily on that of Ristic, with the main differences being that the likelihood function is modified to more accurately match radar signal processing and a more efficient method for calculating the probability of target existence is developed.

In this work the simulated data used as the input to the tracker is a uniform grid, with intensities defined for each bin. The intensities are modelled as the magnitude of a complex spectrum, which corresponds to the output of signal processing stages in operational sensors such as radar. The envelope of complex Gaussian noise is Rayleigh distributed and the magnitude of a signal in complex Gaussian noise is Ricean distributed [5]. These densities are properly incorporated into the likelihood function of the TBD filter.

This paper introduces a TBD algorithm for processing data with an efficient method for calculating the probability of target existence. The algorithm simultaneously estimates target parameters including existence, position, velocity and signal intensity. Section 2 sets the problem mathematically, with a summary of the models used for target motion and the assumed signal processing model. Section 3 derives the filter, incorporating target existence, while Section 4 details the implementation of this filter using sequential Monte-Carlo techniques. Section 5 gives details of the simulation used to test the filter and shows results gained from repeated simulation. The final section summarises the key contributions of this paper and offers further opportunities for research.

2 Target and Sensor Models

2.1 Target Model

The target state is represented by a five-dimensional state vector ψ_k with the elements x_k, $\dot{x}_k$, y_k, $\dot{y}_k$ and A_k. These correspond to position and velocity in the x and y directions and the intensity of the target, which could be related to the target's sensor cross-section. The time evolution of the target state is modelled as a linear Gaussian process

$$\psi_{k+1} = F\psi_k + u_k, \tag{1}$$

where F is the process transition matrix and u_k is a Gaussian noise process with zero mean and covariance Q.

Along with target position, this work attempts to detect the *existence* of a target in the data, in analogy with [6, 7]. The random variable $E_k \in \{e, \bar{e}\}$, denotes the existence or non-existence of the target. This process is modelled as a Hidden Markov Model (HMM), where the transition relations can be defined as

- $p_{ee} = P(E_k = e | E_{k-1} = e)$, target stays alive
- $p_{\bar{e}e} = P(E_k = \bar{e} | E_{k-1} = e)$, target dies
- $p_{e\bar{e}} = P(E_k = e | E_{k-1} = \bar{e})$, target is born
- $p_{\bar{e}\bar{e}} = P(E_k = \bar{e} | E_{k-1} = \bar{e})$, target stays dead,

with the Markovian constraints that

$$p_{ee} + p_{\bar{e}e} = 1, \text{ and} \quad (2)$$

$$p_{e\bar{e}} + p_{\bar{e}\bar{e}} = 1. \quad (3)$$

Hence the Markov transition matrix is completely specified by defining the target birth probability, $p_{e\bar{e}}$, and the target death probability, $p_{\bar{e}e}$.

Equation (1) describes the evolution of the target state when the target is assumed to exist at both times k and $k+1$, hence (1) describes the target state density $p(\psi_k|\psi_{k-1}, E_k = e, E_{k-1} = e)$. However, if the target does not exist at time $k-1$, there is no valid state to evolve into time k and the prior density, $p(\psi_k|E_k = e, E_{k-1} = \bar{e})$, is required. The prior is chosen to be uniform in the state space over a volume related to the surveillance region.

2.2 Signal Processing Model

The measurement model used in this work is as close as possible to that which might be used in an operational sensor, such as radar, without describing details such as beam-forming or the receiving process. It is assumed that a target signal can be represented by a complex sinusoid in two dimensions. This signal can be written in terms of the target return amplitude, A_k, and the location coordinates, x_k and y_k

$$w_k(\psi_k) = A_k \exp\left(j2\pi\left[x_k \frac{l}{L} + y_k \frac{p}{P} + \phi_k\right]\right), \quad (4)$$

where ϕ_k is some arbitrary phase. The indices $l \in \{0 \ldots L-1\}$, $p \in \{0 \ldots P-1\}$ correspond to the signal indices, which could be for example time (related to range) or receiver index (related to azimuth).

The resulting signal is windowed [8] in each dimension to reduce peak side-lobe levels, before a three-dimensional discrete Fourier transform (DFT) is applied. This gives data which is indexed over the location coordinates x and y. In order to remove the influence of the signal's unknown phase component, ϕ_k, the *magnitude* of the spectrum in each DFT bin forms the set of measurements for a single time index, k. Rather than use the entire resulting spectrum, which may have a large number of samples in either dimension (for example this simulation uses 128 bins for each), a small section around the location of the target is extracted to form the set of measurements used by the filtering process. In this case a contiguous subset of bins is chosen with sizes $X \leq L$ and $Y \leq P$.

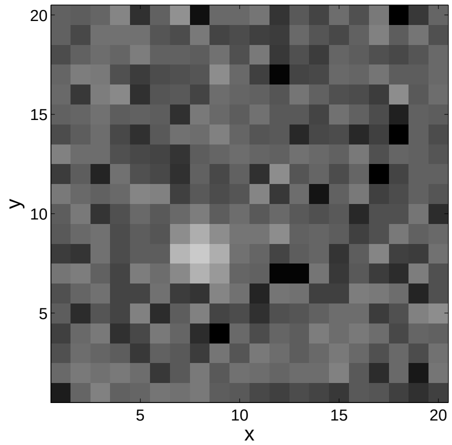

Fig. 1: Simulated sensor data.

Figure 1 shows data simulated using this method. Each axis contains 20 bins from the 128 total bins used in each dimension for this simulation. A single 20 dB peak is displayed in the data at location $(8,8)$. It is apparent that the peak is smeared along both axes, due to the broadening of the peak by the DFT windowing functions.

Following the radar model described above, the measurement equation is described in terms of each DFT bin, and takes the form

$$z_k^{(x,y)} = \begin{cases} h(W^{(x,y)}(\psi_k), \omega_k^{(x,y)}), & \text{target present } (E_k = e) \\ h(0, \omega_k^{(x,y)}), & \text{otherwise } (E_k = \bar{e}), \end{cases} \quad (5)$$

where $\omega_k^{(x,y)}$ is a complex, zero mean Gaussian noise process independently and identically distributed in the real and imaginary axes, with variance σ_ω^2. The function $W^{(x,y)}(\psi_k)$ represents the signal $w_k(\psi_k)$ after the windowing and DFT process, indexed by bin in x and y coordinates. Equation 5 results in two corresponding likelihood functions. The signal likelihood has a Ricean distribution [5]

$$P(z_k^{(x,y)}|\psi_k, E_k = e) = \frac{z_k^{(x,y)}}{\alpha^2} I_0\left(\frac{z_k^{(x,y)} W^{(x,y)}(\psi_k)}{\alpha^2}\right) \times \exp\left(-\frac{[z_k^{(x,y)}]^2 + [W^{(x,y)}(\psi_k)]^2}{2\alpha^2}\right), \quad (6)$$

where $\alpha^2 = \sigma_\omega^2/2$ is related to the noise variance, $I_0(\cdot)$ is the modified Bessel function of order zero and $z_k^{(x,y)}$ must be non-negative. The noise likelihood has a Rayleigh distribution [5]

$$P(z_k^{(x,y)}|E_k = \bar{e}) = \frac{z_k^{(x,y)}}{\alpha^2} \exp\left(-\frac{[z_k^{(x,y)}]^2}{2\alpha^2}\right), \quad (7)$$

where $z_k^{(x,y)}$ must again be non-negative. The noise in each bin is assumed to be independent and so the complete likelihood function is then a product over all of the contributions from each bin

$$P(z_k|\psi_k, E_k = e) = \prod_{x=1}^{X}\prod_{y=1}^{Y} P(z_k^{(x,y)}|\psi_k, E_k) \quad (8)$$

$$P(z_k|E_k = \bar{e}) = \prod_{x=1}^{X}\prod_{y=1}^{Y} P(z_k^{(x,y)}|E_k). \quad (9)$$

The particle weighting function (see Section 4) is written in terms of a likelihood ratio

$$\ell(\cdot) = \frac{P(\cdot)}{P(z_k^{(x,y)}|E_k = \bar{e})}, \quad (10)$$

where the most important case is $\ell(z_k^{(x,y)}|\psi_k, E_k)$, which using (6) and (7) gives

$$\ell(z_k^{(x,y)}|\psi_k, E_k) = \exp\left(-\frac{[W^{(x,y)}(\psi_k)]^2}{2\sigma_\omega^2}\right) \times I_0\left(\frac{z_k^{(x,y)} W^{(x,y)}(\psi_k)}{\sigma_\omega^2}\right). \quad (11)$$

The complete likelihood ratio for a single time k is used in the filter derivation of the following section and can be defined in terms of (9), giving

$$L(\cdot) = \frac{P(\cdot)}{P(z_k|E_k = \bar{e})}. \quad (12)$$

3 Filter Derivation

The desired state estimate at time k, $\hat{\psi}_k$, will be a function of the joint probability density function, $p(\psi_k, E_k|z_{1:k})$, which describes the state of the target, including its existence or non-existence. This density can be expanded as

$$p(\psi_k, E_k|z_{1:k}) = p(\psi_k|E_k, z_{1:k})P(E_k|z_{1:k}). \quad (13)$$

The two factors on the right-hand side of (13) are thus required for calculation of the joint density. Since E_k is Boolean, the following relation must hold

$$P(E_k = \bar{e}|z_{1:k}) = 1 - P(E_k = e|z_{1:k}), \quad (14)$$

and hence there is only need to calculate one of these quantities, which in this case is chosen to be $P(E_k = e|z_{1:k})$, the probability that the target exists. Furthermore, the density $p(\psi_k|E_k, z_{1:k})$ is not physical if the target does not exist, so the only quantities of interest are the posterior target density, $p(\psi_k|E_k = e, z_{1:k})$, and the probability of existence, $P(E_k = e|z_{1:k})$. For this reason, and for notational brevity, E_k will be assumed to refer to $E_k = e$ unless explicitly specified in the remainder of this paper.

3.1 Joint State Densities

The density of ψ_k required by (13) can be expanded over the target's existence at time $k-1$

$$p(\psi_k|E_k, z_{1:k}) = \sum_{E_{k-1}\in\{e,\bar{e}\}} p(\psi_k|E_k, E_{k-1}, z_{1:k}) \times P(E_{k-1}|E_k, z_{1:k}). \quad (15)$$

Equation (15) shows that the desired density, $p(\psi_k|E_k, z_{1:k})$, can be separated into a mixture of two calculable densities. The first of these, $p(\psi_k|E_k, E_{k-1} = e, z_{1:k})$, called the *continuing density* in the sequel, describes the case where the target exists at time $k-1$ and continues to exist at time k. The second of these, $p(\psi_k|E_k, E_{k-1} = \bar{e}, z_{1:k})$, called the *birth density*, describes the case where the target does not exist at time $k-1$, but comes into existence between times $k-1$ and k.

The densities, $p(\psi_k|E_k, E_{k-1}, z_{1:k})$, in (15) can be evaluated using Bayes' rule

$$p(\psi_k|E_k, E_{k-1}, z_{1:k}) = \frac{P(z_k|\psi_k, E_k, E_{k-1}, z_{1:k-1})p(\psi_k|E_k, E_{k-1}, z_{1:k-1})}{P(z_k|E_k, E_{k-1}, z_{1:k-1})} \quad (16)$$

$$= \frac{P(z_k|\psi_k, E_k)p(\psi_k|E_k, E_{k-1}, z_{1:k-1})}{P(z_k|E_k, E_{k-1}, z_{1:k-1})}. \quad (17)$$

Since the likelihood that the target is not present, $P(z_k|E_k = \bar{e})$, is independent of the target state, the density can be written in terms of a likelihood ratio

$$p(\psi_k|E_k, E_{k-1}, z_{1:k}) = \frac{L(z_k|\psi_k, E_k)p(\psi_k|E_k, E_{k-1}, z_{1:k-1})}{L(z_k|E_k, E_{k-1}, z_{1:k-1})} \quad (18)$$

$$\propto L(z_k|\psi_k, E_k)p(\psi_k|E_k, E_{k-1}, z_{1:k-1}). \quad (19)$$

In the case where the target existed at time $k-1$, $E_{k-1} = e$, the predicted density can be calculated in terms of the target dynamic model, that is

$$p(\psi_k|E_k, E_{k-1} = e, z_{1:k-1}) = \int p(\psi_k|\psi_{k-1}, E_k, E_{k-1} = e) \times p(\psi_{k-1}|E_{k-1} = e, z_{1:k-1})d\psi_{k-1}. \quad (20)$$

If the target did not exist at time $k-1$ then

$$p(\psi_k|E_k, E_{k-1} = \bar{e}, z_{1:k-1}) = p(\psi_k|E_k, E_{k-1} = \bar{e}), \quad (21)$$

which represents the prior density of a target which has started to exist between times $k-1$ and k.

3.2 Mixing Terms

The mixing terms required by (15), can be rearranged using Bayes' rule

$$P(E_{k-1}|E_k, z_{1:k}) = \frac{P(E_k, z_k|E_{k-1}, z_{1:k-1})}{P(z_k, E_k|z_{1:k-1})}P(E_{k-1}|z_{1:k-1}) \quad (22)$$

$$= \frac{L(z_k|E_k, E_{k-1}, z_{1:k-1})P(E_k|E_{k-1})}{L(z_k, E_k|z_{1:k-1})} \times P(E_{k-1}|z_{1:k-1}). \quad (23)$$

The first term in the numerator of (23) is identical to the normalising term of (18) and expands to

$$L(z_k|E_k, E_{k-1}, z_{1:k-1}) = \int L(z_k|\psi_k, E_k)p(\psi_k|E_k, E_{k-1}, z_{1:k-1})d\psi_k, \quad (24)$$

while the second term represents the existence transition relation. The third term in the numerator of (23) is the posterior probability of existence at time $k-1$. Since the denominator of (23) is independent of E_{k-1}, it can be considered to be a normalising constant, which can be simply calculated using

$$L(z_k, E_k|z_{1:k-1}) = \sum_{E_{k-1}=\{e,\bar{e}\}} L(z_k|E_k, E_{k-1}, z_{1:k-1}) \times P(E_k|E_{k-1}, z_{1:k-1})P(E_{k-1}|z_{1:k-1}). \quad (25)$$

3.3 Probability of Existence

Using Bayes' rule, the probability of existence required by (13) is given by

$$P(E_k|z_{1:k}) = \frac{P(z_k|E_k, z_{1:k-1})P(E_k|z_{1:k-1})}{P(z_k|z_{1:k-1})} \quad (26)$$

$$= \frac{L(z_k|E_k, z_{1:k-1})P(E_k|z_{1:k-1})}{L(z_k|z_{1:k-1})}, \quad (27)$$

where, as before, $L(\cdot)$ refers to the likelihood ratio as given in (12). The first term in the numerator of (27) can be expanded as follows

$$L(z_k|E_k, z_{1:k-1}) = \sum_{E_{k-1}=\{e,\bar{e}\}} L(z_k, E_{k-1}|E_k, z_{1:k-1}) \quad (28)$$

$$= \sum_{E_{k-1}=\{e,\bar{e}\}} P(E_{k-1}|E_k, z_{1:k})L(z_k|E_k, E_{k-1}, z_{1:k-1}). \quad (29)$$

The mixing terms, $P(E_{k-1}|E_k, z_{1:k})$, are exactly those required by the conditional target state in (15) and given by (23). The likelihood in (29) was derived in the calculation of the mixing terms and is given by (24). The second term in the numerator of (27) is the predicted probability of existence and is given by

$$P(E_k|z_{1:k-1}) = p_{ee}P(E_{k-1} = e|z_{1:k-1}) + p_{e\bar{e}}\left[1 - P(E_{k-1} = e|z_{1:k-1})\right], \quad (30)$$

where p_{ee} and $p_{e\bar{e}}$ are defined by the Markovian existence model (Section 2.1) and $P(E_{k-1} = e|z_{1:k-1})$ is the posterior probability of existence at time $k-1$.

The denominator of (27) can be expanded as follows

$$L(z_k|z_{1:k-1}) = \sum_{E_k \in \{e,\bar{e}\}} L(z_k, E_k|z_{1:k-1}) \quad (31)$$

$$= L(z_k|E_k = e, z_{1:k-1})P(E_k = e|z_{1:k-1}) + L(z_k|E_k = \bar{e}, z_{1:k-1})P(E_k = \bar{e}|z_{1:k-1}). \quad (32)$$

The calculation of the first likelihood ratio has been derived in (29) and by definition the second likelihood ratio $L(z_k|E_k = \bar{e}, z_{1:k-1}) = L(z_k|E_k = \bar{e}) = 1$. Obviously the predicted probability that the target does not exist is related to the predicted probability of target existence by $P(E_k = \bar{e}|z_{1:k-1}) = 1 - P(E_k = e|z_{1:k-1})$. Equation (32) then simplifies as

$$L(z_k|z_{1:k-1}) = L(z_k|E_k, z_{1:k-1})P(E_k|z_{1:k-1}) + \left[1 - P(E_k|z_{1:k-1})\right]. \quad (33)$$

Hence the probability of existence is given by

$$P(E_k|z_{1:k}) = \frac{L(z_k|E_k, z_{1:k-1})P(E_k|z_{1:k-1})}{L(z_k|E_k, z_{1:k-1})P(E_k|z_{1:k-1}) + \left[1 - P(E_k|z_{1:k-1})\right]}, \quad (34)$$

where $L(z_k|E_k, z_{1:k-1})$ is calculated using (29) and $P(E_k|z_{1:k-1})$ is calculated using (30).

3.4 Summary of the Derivation

As can be seen from the derivations, each of the desired terms can be calculated as functions of

- the prior probability of existence, $P(E_{k-1}|z_{1:k-1})$,
- the Markov transition terms, p_{ee} and $p_{e\bar{e}}$,
- the likelihood ratio $L(z_k|\psi_k, E_k)$,
- a prior state density assuming that the target existed at time $k-1$, $p(\psi_{k-1}|E_{k-1} = e, z_{1:k-1})$,
- a transition density assuming that the target continued to exist through times $k-1$ and k, $p(\psi_k|\psi_{k-1}, E_k, E_{k-1} = e)$, and
- a prior state density assuming that the target started to exist between times $k-1$ and k, $p(\psi_k|E_k, E_{k-1} = \bar{e})$.

All of the above are quantities that have been defined by the process model and the measurement model presented in previous sections.

4 Particle Filter Implementation

The basic algorithm is based upon sequential Monte-Carlo methods (particle filtering) [9, 10], with a novel modification to detect target existence. The algorithm incorporates four dependent processes that estimate the target state and the probability of target existence. The continuing and birth densities (15) are estimated using separate particle filters. The third process mixes the two approximate densities together to form the complete posterior state density, while a fourth process calculates the probability of target existence.

4.1 Calculation of the Continuing Density

Following from (18), the continuing density, $p(\psi_k|E_k, E_{k-1} = e, z_{1:k})$, is calculated using a standard SIR particle filtering algorithm (Algorithm 4 of [10]), without a resampling step. In this algorithm resampling is performed after the mixing stage. Assuming that the prior density, $p(\psi_{k-1}|E_{k-1} = e, z_{1:k-1})$, is represented by the set of particles $i \in \{1 \ldots N_c\}$, with position ψ^i_{k-1} and weights $w^i_{k-1} = {}^1/_{N_c}$, the algorithm proceeds as

1. The system dynamics form the importance sampling density, and so the process equation is used to propose the particle positions at time k

$$\psi_k^{(c)i} = F\psi_{k-1}^i + \mu_k^i, \quad \forall i \in \{1 \dots N_s\}, \tag{35}$$

where μ_k^i is a random sample drawn from the process noise distribution.

2. The new particle weights are calculated using the likelihood ratio (12). Thus the set of unnormalised weights at time k become

$$\tilde{w}_k^{(c)i} = \frac{1}{N_c} L(z_k | \psi_k^{(c)i}, E_k), \tag{36}$$

for all particles $i \in \{1 \dots N_c\}$.

3. The weights can then be normalised

$$w_k^{(c)i} = \frac{\tilde{w}_k^{(c)i}}{\sum_{j=1}^{N_s} \tilde{w}_k^{(c)j}}, \quad \forall i \in \{1 \dots N_c\}. \tag{37}$$

4. The set of particles $\psi_k^{(c)i}$ with corresponding weights $w_k^{(c)i}$ thus form an approximation to the posterior density $p(\psi_k | E_k, E_{k-1} = e, z_{1:k})$.

Since the application of this algorithm is for low SNR targets, using the system dynamics as the importance density is adequate. If the algorithm was applied to higher SNR targets, then an importance density based on the data would give better performance.

4.2 Calculation of the Birth Density

Unlike the calculation of the continuing density, the calculation of the birth density, $p(\psi_k | E_k, E_{k-1} = \bar{e}, z_{1:k})$, does not have the systems dynamics as prior information, since the target does not exist at time $k-1$. It follows that in this case the importance density cannot be base on the system dynamics. From [10], the particle weights can be defined as

$$w_k^{(b)i} \propto \frac{p(\psi_k^{(b)i} | E_k = e, E_{0:k-1} = \bar{e}, z_{1:k})}{q(\psi_k^{(b)i} | E_k = e, E_{0:k-1} = \bar{e}, z_{1:k})}, \tag{38}$$

for $i \in \{1 \dots N_b\}$, where $q(\cdot)$ is the importance density and $E_{0:k-1} = \bar{e}$ signifies that $E_t = \bar{e}$ for all $t \in \{0 \dots k-1\}$. The weight calculation does not account for $\psi_{0:k-1}$, since the non-existence of the target prior to k implies that the target state is undefined. In analogy with (19), the density in the numerator of (38) factorises as

$$p(\psi_k | E_k = e, E_{0:k-1} = \bar{e}, z_{1:k}) \propto L(z_k | \psi_k, E_k) p(\psi_k | E_k = e, E_{k-1} = \bar{e}), \tag{39}$$

while the importance density is chosen as

$$q(\psi_k | E_k, E_{0:k-1} = \bar{e}, z_{1:k}) = q(\psi_k | E_k, E_{k-1} = \bar{e}, z_k). \tag{40}$$

Thus the weight calculation becomes

$$w_k^{(b)i} \propto \frac{L(z_k | \psi_k^{(b)i}, E_k) p(\psi_k^{(b)i} | E_k = e, E_{k-1} = \bar{e})}{q(\psi_k^{(b)i} | E_k, E_{k-1} = \bar{e}, z_k)}. \tag{41}$$

Choosing an importance density equal to the prior described in Section 2.1 would imply that the particle weights become normalised likelihoods. However, since the importance density depends on z_k, the data at time k can give some hints as to a plausible target state. In this case peaks in the data are used to initialise the target state, under the assumption that even quiet targets will instantaneously disturb the underlying noise. Bins in the data which have an intensity that exceed a certain threshold are chosen as the initialising bins. Uniform samples are drawn from within the area enclosed by each initialising bin and the target intensity is similarly sampled from the measurement noise, with the intensity of the initialising bin as the mean.

The importance density can then be written as

$$q(\psi_k | E_k, E_{k-1} = \bar{e}, z_k) = \frac{1}{N_P V} \sum_{(x,y) \in N_P} \mathcal{N}(A_k; z_k^{(x,y)}, \sigma_\omega^2), \tag{42}$$

where N_P is the set of of bins in the data which have intensity that exceeds the threshold, $z_k^{(x,y)}$ is the measured intensity in bin (x,y) and V is the area of a single DFT bin. Note that since the measurements do provide information about $\dot{x}$ and $\dot{y}$, they are sampled directly from the prior.

Hence the algorithm for computing the birth density is

1. Draw N_b samples, $\psi_k^{(b)i}$, from the importance density,

$$\psi_k^{(b)i} \sim q(\psi_k | E_k, E_{k-1} = \bar{e}, z_k), \quad i \in \{1 \dots N_b\}. \tag{43}$$

2. The unnormalised particle weights are then calculated using the likelihood ratio

$$\tilde{w}_k^{(b)i} = \frac{L(z_k | \psi_k^{(b)i}, E_k)}{N_b V_T q(\psi_k^{(b)i} | E_k, E_{k-1} = \bar{e}, z_k)}, \tag{44}$$

for all particles $i \in \{1 \dots N_b\}$, where V_T is the total volume encapsulated by the target state prior.

3. The weights can then be normalised

$$w_k^{(b)i} = \frac{\tilde{w}_k^{(b)i}}{\sum_{j=1}^{N_s} \tilde{w}_k^{(b)j}}, \quad \forall i \in \{1 \dots N_b\}. \tag{45}$$

4. The set of particles, $\psi_k^{(b)i}$, and their weights, $w_k^{(b)i}$, then approximate the density $p(\psi_k | E_k, E_{k-1} = \bar{e}, z_{1:k})$.

4.3 Calculation of the Mixing Parameters

In order to determine the complete posterior distribution, the mixing parameters (23) must be calculated. For both the continuing and birth densities, the mixing parameter requires calculation of the integral

$$\int L(z_k | \psi_k, E_k) p(\psi_k | E_k, E_{k-1}, z_{1:k-1}) d\psi_k, \tag{46}$$

where $E_k = e$, but $E_{k-1} \in \{e, \bar{e}\}$. Notice that from the derivation of the densities in Section 3, the normalising term in (18) is

$$L(z_k|E_k, E_{k-1}, z_{1:k-1}) = \int L(z_k|\psi_k, E_k)p(\psi_k|E_k, E_{k-1}, z_{1:k-1})d\psi_k, \quad (47)$$

which is the required quantity. This quantity is implicitly evaluated when normalising the densities, which occurs through normalisation of the particle weights. Hence the integral required for the calculation of the mixing parameters is simply approximated by the sum of the unnormalised weights. This is not a new observation, for example see [11]. The unnormalised mixing term for the continuing density is thus

$$\tilde{M}_c = p_{ee}P(E_{k-1} = e|z_{1:k-1})\sum_{i=1}^{N_c}\tilde{w}_k^{(c)i}. \quad (48)$$

Similarly, the mixing term for the birth density is

$$\tilde{M}_b = p_{e\bar{e}}P(E_{k-1} = \bar{e}|z_{1:k-1})\sum_{i=1}^{N_b}\tilde{w}_k^{(b)i}. \quad (49)$$

These can then be normalised, giving

$$M_i = \frac{\tilde{M}_i}{\tilde{M}_c + \tilde{M}_b}, \quad i \in \{c, b\}, \quad (50)$$

where the M_i denote $M_c = P(E_{k-1} = e|E_k, z_{1:k})$ and $M_b = P(E_{k-1} = \bar{e}|E_k, z_{1:k})$.

4.4 Forming the Posterior Target State Density

The two sets of particle/weight pairs, $\{(\psi_k^{(c)i}, w_k^{(c)i})|i = 1\dots N_c\}$ and $\{(\psi_k^{(b)i}, w_k^{(b)i})|i = 1\dots N_b\}$, representing the continuing and birth densities, can be joined into a single set of particles representing the complete posterior (15). This super-set of particles can then be resampled, leaving N_c particles describing the posterior target state density. Using the mixing parameters calculated in the previous section, the algorithm is straightforward

1. Multiply the weight of each particle $i \in \{1\dots N_c\}$ in the continuing density by the corresponding mixing parameter
$$\hat{w}_k^{(c)i} = M_c w_k^{(c)i}. \quad (51)$$
2. Similarly for the birth density
$$\hat{w}_k^{(b)i} = M_b w_k^{(b)i}. \quad (52)$$
3. Form the set of particles representing the complete posterior from the re-weighted set of samples
$$\left\{(\psi_k^{(j)i}, \hat{w}_k^{(j)i})|i \in \{1\dots N_j\}, j \in \{c, b\}\right\} \quad (53)$$
4. Resample the full set of $N_c + N_b$ particles, reducing their number down to N_c, resulting in the particles ψ_k^i, $i \in \{1\dots N_c\}$ with uniform weights.

4.5 Calculating the Probability of Existence

In order to evaluate the probability of existence, the two terms given by (29) and (30) must be calculated. As discussed in Section 4.3, the likelihood required by (29) is given by the sum of unnormalised particle weights and the mixing parameters are the same as those required for the state densities. Hence

$$L(z_k|E_k, z_{1:k-1}) \approx M_c\sum_{i=1}^{N_c}\tilde{w}_k^{(c)i} + M_b\sum_{i=1}^{N_b}\tilde{w}_k^{(b)i}. \quad (54)$$

The predicted probability of existence is straightforward to calculate from (30). Writing $P_k^E = P(E_k|z_{1:k})$,

$$P(E_k|z_{1:k-1}) = p_{ee}P_{k-1}^E + p_{e\bar{e}}[1 - P_{k-1}^E]. \quad (55)$$

Using (54) and (55) the probability of existence, P_k^E, can then be calculated by substitution into (34).

5 Simulation and Results

This section aims to show the efficacy of the track-before-detect algorithm in detecting and tracking a single target through simulated data.

The simulation generates random synthetic target trajectories according to the dynamics given in (1). The process noise covariance matrix used for the target process, Q, is

$$Q = \begin{bmatrix} \frac{1}{2}Q_s & 0_{2\times2} & 0_{2\times1} \\ 0_{2\times2} & \frac{1}{2}Q_s & 0_{2\times1} \\ 0_{2\times1} & 0_{2\times1} & \frac{1}{10}T \end{bmatrix}, \quad (56)$$

where the sub-matrix Q_s is given by

$$Q_s = \begin{bmatrix} T^3/3 & T^2/2 \\ T^2/2 & T \end{bmatrix}, \quad (57)$$

$0_{a\times b}$ is a $a \times b$ matrix of zeros and $T = 1$ is the period of time between measurements. The birth probability, $p_{e\bar{e}}$, and death probability, $p_{\bar{e}e}$ are both set at 0.1.

The 20×20 grid of data given to the tracker is a subset of the full 128×128 measurement data set. Setting the noise variance, σ_ϖ^2, at unity allows the amplitude of the complex sinusoids (4) to control the signal-to-noise ratio. Hanning DFT windows are used in both the x and y directions, although this is not a limitation of the algorithm. As long as the windowing function used to simulate the data matches (or is closely approximated by) that used to generate each particle's expected signal, any function (including the flat-top window) could be used.

In each simulation the target state is initialised at

$$\psi_0 = \begin{bmatrix} 8 & 0.2 & 8 & 0.2 & 10^{A_{\mathrm{dB}}/20} \end{bmatrix}^T, \quad (58)$$

where A_{dB} is the simulated signal amplitude in decibels. The target position is assumed to remain within the measured region, the velocity in each dimension is assumed to be constrained to $|\dot{x}_k|, |\dot{y}_k| < 1$ and the signal intensity constrained to $0 < A_k < 10$. 2000 particles were used for both the birth and continuing densities, for a total of 4000 particles.

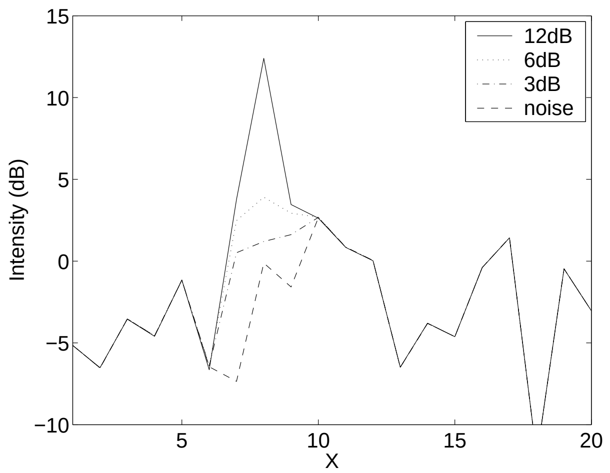

Fig. 2: Simulated data with signals of different strengths.

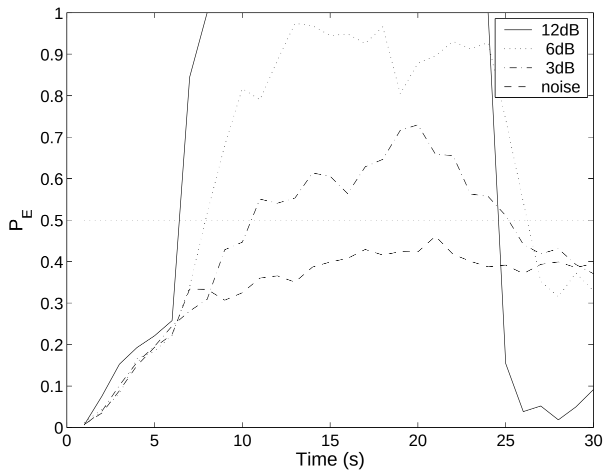

Fig. 3: Average probability of existence.

Simulated data containing a target at several different signal strengths are compared in Figure 2. This figure shows a slice through each set of data along the x axis at $y = 8$ (c.f. Figure 1), showing the intensities, in dB, in each bin. For the purposes of comparison, this figure shows each target signal simulated in the same set of noise. The figure shows that, for this snapshot of noise and target intensities, the 12 dB SNR target is relatively easy to pick out from the noise, whereas the 3 dB SNR target is much more difficult to discern. This illustrates the difficulty that a standard detection scheme would have in detecting these low SNR targets.

Figure 3 shows the probability of existence calculate by the TBD algorithm, averaged over ten simulations, for SNR set at 12 dB, 6 dB, 3 dB and for noise only. The noise only example is included to show the resulting probability of existence when there is no target present. In these examples the target appears after 7 seconds and then dies at the 25 second mark. The target is assumed to be present if the probability of existence is greater than 0.5, which is plotted as a horizontal line in Figure 3.

It can be seen from Figure 3 that, on average, the algorithm performs very well at 12 dB SNR, with no track initiation delay and no track termination delay. For a 6 dB SNR target, the detection confidence is also quite high, but there is an average delay of one second on both track initiation and track termination. The detection of the target at 3 dB SNR is much less confident, but passes the threshold with approximately 4 seconds track initiation delay and 1 second track termination delay. The asymptotic value of P_E, which can be seen prominently in the noise only curve, is related to the probabilities of target birth and death in the existence transition matrix.

6 Conclusions and Further Work

This paper has developed a novel track-before-detect algorithm based on the implementation of an efficient particle filter. The data expected by the filter is closely related to that which would be produced by operational sensors. Results produced from the algorithm using a total of 4000 particles have shown that this technique can reliably detect and track targets with a per-bin signal-to-noise ratio of 6dB. Targets with a 3dB SNR can be detected and tracked, although with a much lower confidence.

There are several continuing areas of research for improving the algorithm developed in this paper. There is no need to have the same number of particles in the birth and continuing densities. In the interests of improving the algorithm efficiency, the effect of changing the ratio of the number of particles used to estimate each density could be explored. The extension of the current single target algorithm to multi-target situations is an on-going area of research.

Acknowledgements

Mark Rutten would like to acknowledge his useful discussions with Professors Bill Moran (Department of EEE, The University of Melbourne) and Rob Evans (CSSIP and the Department of EEE, The University of Melbourne). Simon Maskell's contribution was sponsored in part by the UK Ministry of Defence. All authors gratefully acknowledge that the Anglo-Australian Memorandum Of Understanding on Research, AAMOUR, facilitated the collaborative work documented in this paper.

References

[1] D. J. Salmond and H. Birch, "A particle filter for track-before-detect," in *Proc. American Control Conf.*, Arlington, VA, USA, June 2001, pp. 3755–3760.

[2] M. Rollason and D. Salmond, "A particle filter for track-before-detect of a target with unknown amplitude," in *IEE Int. Seminar on Target Tracking: Algorithms and Applications*, Enschede, Netherlands, Oct. 2001, pp. 14/1–4.

[3] B. Ristic, "Detection and tracking of stealthy targets," in *Beyond the Kalman Filter: Particle Filters for Tracking Applications*. Artech House, 2004, ch. 11.

[4] Y. Boers and H. Driessen, "Particle filter based track before detect algorithms," in *Proc. SPIE Conf. on Signal and Data Processing of Small Targets*, vol. 5204, San Diego, CA, USA, Aug. 2003, pp. 20–30.

[5] M. I. Skolnik, *Introduction to Radar Systems*, 3rd ed. McGraw-Hill, 2001.

[6] S. B. Colegrove, A. W. Davis, and J. K. Ayliffe, "Track initiation and nearest neighbours incorporated into probabilistic data association," *Journal of Electrical and Electronics Engineers, Australia*, vol. 6, no. 3, pp. 191–198, Sept. 1986.

[7] D. Mušicki, R. Evans, and S. Stankovic, "Integrated probabilistic data association," *IEEE Trans. Automat. Contr.*, vol. 39, no. 6, pp. 1237–1240, June 1994.

[8] F. J. Harris, "On the use of windows for harmonic analysis with the discrete Fourier transform," *Proc. IEEE*, vol. 66, no. 1, pp. 51–83, Jan. 1978.

[9] N. J. Gordon, D. J. Salmond, and A. F. M. Smith, "Novel approach to nonlinear/non-Gaussian Bayesian state estimation," *IEE Proc.-F: Radar and Signal Processing*, vol. 140, no. 2, pp. 107–113, Apr. 1993.

[10] M. S. Arulampalam, S. Maskell, N. Gordon, and T. Clapp, "A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking," *IEEE Trans. Signal Processing*, vol. 50, no. 2, pp. 174–188, Feb. 2002.

[11] Y. Boers and H. Driessen, "A particle-filter-based detection scheme," *IEEE Signal Processing Lett.*, vol. 10, no. 10, pp. 300–302, Oct. 2003.