Bayesian Filtering with Random Finite Set Observations

Ba-Tuong Vo
WATRI
The University of Western Australia
Crawley WA 6009, Australia
Email: vob@watri.org.au

Ba-Ngu Vo
EEE Department
University of Melbourne
Parkville VIC 3010, Australia
Email: bv@ee.unimelb.edu.au

Antonio Cantoni
WATRI
The University of Western Australia
Crawley WA 6009, Australia
Email: cantoni@watri.org.au

***Abstract*— This paper presents a Bayes recursion for tracking a target that generates multiple measurements with state dependent sensor field of view and clutter. Our Bayesian formulation is mathematically well-founded due to our use of a mathematically consistent likelihood function derived from random finite set theory. A particle implementation of the proposed filter is given. Under linear Gaussian assumptions, an exact closed form solution to the proposed recursion is derived, and efficient implementations are given.**

I. Introduction

The objective of target tracking is to estimate the state of the target from measurement sets collected by the sensor at each time step. This is a challenging problem since the target can generate multiple measurements which are not always detected by the sensor, and the sensor receives a set of spurious measurements (clutter) not generated by the target. Existing techniques for handling this problem rest on the simplifying assumptions that the target generates at most one measurement and that the sensor field of view is constant. Such assumptions are not realistic, for example, in extended object tracking or tracking in the presence of electronic counter measures, which have increasingly becoming important due to high resolution capabilities of modern sensors. Nonetheless, these assumptions have formed the basis of a plethora of works e.g. the probabilistic data association (PDA) filter [1], the multiple hypothesis tracker (MHT) [2], [3] and their variants. However, such techniques are not easily adapted to accommodate multiple target generated-measurements and state dependent field of view. Moreover, it is not clear how such techniques are mathematically consistent with the Bayesian paradigm.

This paper presents a mathematically well-founded Bayesian approach to tracking a target that can generate multiple measurements, in the presence of detection uncertainty and clutter. In our formulation, the collection of observations at any time is treated as a set-valued observation which encapsulates the underlying models of multiple target-generated measurements, state dependent sensor field of view, and clutter [4], [5]. Since the observation space is now the space of finite sets, the usual Euclidean notion of a density is not suitable. Random finite set (RFS) or point process theory provides an elegant and rigorous framework to derive mathematically consistent densities needed in the Bayes recursion [4], [5].

The key contributions of this paper are

- A mathematically consistent Bayes recursion together with a particle implementation that accommodates multiple target-generated measurements, state dependent field of view and clutter using RFS theory.
- A closed form solution to the proposed recursion for the special class of linear Gaussian single target models.

Assuming no clutter and that the target generates exactly one measurement, the proposed recursion reduces to the usual Bayes recursion and the particle filter implementation reduces to the standard particle filter. Under additional linear Gaussian assumptions, our closed form recursion reduces to the celebrated Kalman filter. In the case of a linear Gaussian model with at most one target-generated measurement, constant field of view, and uniform clutter, the proposed closed form recursion reduces to the Gaussian mixture filter given in [6]. Moreover, if at each time step, gating is performed and the Gaussian mixture posterior density is collapsed to a single Gaussian component, then the proposed recursion reduces to the PDA filter [1].

II. Background

A. The Bayes Recursion

In the classical Bayes filter the hidden state x_k is assumed to follow a first order Markov process on the state space $\mathcal{X} \subseteq \mathbb{R}^{n_x}$ according to a transition density $f_{k|k-1}(x_k|x_{k-1})$. The observation $z_k \in \mathcal{Z} \subseteq \mathbb{R}^{n_z}$ is assumed conditionally independent given the states x_k and is characterized by a likelihood $g_k(z_k|x_k)$, which is the probability density that, at time k, the target with state x_k produces a measurement z_k. Under these assumptions, the classical Bayes recursion propagates the posterior density p_k in time according to

$$p_{k|k-1}(x_k|z_{1:k-1}) = \int f_{k|k-1}(x_k|x)p_{k-1}(x|z_{1:k-1})dx, \quad (1)$$

$$p_k(x_k|z_{1:k}) = \frac{g_k(z_k|x_k)p_{k|k-1}(x_k|z_{1:k-1})}{\int g_k(z_k|x)p_{k|k-1}(x|z_{1:k-1})dx}. \quad (2)$$

where $z_{1:k} = [z_1, ..., z_k]$. All inference on the target state at time k is derived from the posterior density p_k at time k.

The Bayes recursion (1)-(2) is formulated for single-target single-measurement systems. In practice due to multi-path reflections, electronic counter measures, etc. the target may generate multiple measurements, in addition to spurious measurements not generated by the target. Note that the order of appearance of these measurements has no physical significance. Hence, the sensor effectively receives an unordered set of measurements denoted by Z_k, and the observation space is now the space of finite subsets of $\mathcal{Z}$, denoted by $\mathcal{F}(\mathcal{Z})$. Consequently, the Bayes update (2) is not directly applicable.

To accommodate set-valued measurements, we require a mathematically consistent generalization of the likelihood $g_k(z_k|x_k)$ to the set-valued case. In other words, we need a mathematically well-founded notion of the probability density of the set Z_k given x_k. However, the notion of such densities is not straightforward because the space $\mathcal{F}(\mathcal{Z})$ does not inherit the usual Euclidean notions of volume and integration on $\mathcal{Z}$. RFS theory or point process theory provides rigorous notions of volume and integration on $\mathcal{F}(\mathcal{Z})$ needed to define a mathematically consistent likelihood [7].

B. Random Finite Sets

Intuitively, an RFS is simply a finite-set-valued random variable (for further details, see e.g. [7]–[9]). Let $(\Omega, \sigma(\Omega), P)$ be a probability space, where Ω is the sample space, $\sigma(\Omega)$ is a σ-algebra on Ω, and P a probability measure on $\sigma(\Omega)$. A *random finite set* Z on $\mathcal{Z}$ is defined as a measurable mapping

$$Z : \Omega \to \mathcal{F}(\mathcal{Z}). \tag{3}$$

with respect to the Borels sets of $\mathcal{F}(\mathcal{Z})$, generated by the Mathéron topology [10]. The *probability distribution* of the RFS Z is given in terms of the probability measure P by

$$Pr(Z \in \mathcal{T}) = P(\{\omega \in \Omega : Z(\omega) \in \mathcal{T}\}) \tag{4}$$

where $\mathcal{T}$ is any Borel subset $\mathcal{F}(\mathcal{Z})$. The probability distribution of the RFS Z can be equivalently characterized by a discrete probability distribution and a family of joint probability distributions. The discrete distribution characterizes the cardinality (the number of elements) of the RFS, whilst for a given cardinality, an appropriate distribution characterizes the joint distribution of the elements of the RFS [7]–[9].

The *probability density* p_Z of Z is given by the Radon-Nikodým derivative of its probability distribution with respect to an appropriate dominating measure μ. The conventional choice of dominating measure is the unnormalized distribution of a Poisson point process [11]

$$\mu(\mathcal{T}) = \textstyle\sum_{r=0}^{\infty} \lambda^r(\chi^{-1}(\mathcal{T}) \cap \mathcal{Z}^r)/r! \tag{5}$$

where λ^r is the rth product (unitless) Lebesque measure, and χ is a mapping of vectors to sets defined by $\chi([z_1, ..., z_r]^T) = \{z_1, ..., z_r\}$. The integral of a measurable function $f : \mathcal{F}(\mathcal{Z}) \to [0, \infty)$ with respect to μ is defined as follows

$$\begin{aligned}&\textstyle\int_{\mathcal{T}} f(Z)\mu(dZ) = \\ &\quad \textstyle\sum_{r=0}^{\infty} \frac{1}{r!} \int_{\chi^{-1}(\mathcal{T})\cap E^r} f(\{z_1, ..., z_r\})\lambda^r(dz_1...dz_r).\end{aligned} \tag{6}$$

The 1st-order moment of a random finite set Z on $\mathcal{Z}$, also called the *intensity function*, is a non-negative function v_Σ on $\mathcal{Z}$ with the property that for any closed subset $S \subseteq \mathcal{Z}$

$$\mathbb{E}\left[|Z \cap S|\right] = \textstyle\int_S v_Z(x)dx$$

where $|Z|$ denotes the cardinality of Z. In other words, for a given point x, the intensity $v_Z(x)$ is the density of expected number of targets per unit volume at x.

An important class of RFSs are the *Poisson* RFSs, which are completely characterized by their intensity functions. The cardinality of a Poisson RFS Z is Poisson distributed with mean $N_Z = \int v_Z(x)dx$, and for a given cardinality the elements of Z are each independent and identically distributed (i.i.d) with probability density v_Z/N_Z.

For simplicity in notation, we shall use the same symbol for an RFS and its realizations hereon.

III. The RFS Single Target Bayes Recursion

As previously argued, in the presence of detection uncertainty and clutter, the measurement is set-valued. In this section, we describe a RFS measurement model and derive the corresponding likelihood function.

A. RFS Measurement Model

The collection of measurements obtained at time k is represented as a finite subset Z_k of the original observation space $\mathcal{Z} \subseteq \mathbb{R}^{n_z}$. More concisely, if $M(k)$ observations $z_{k,1}, \ldots, z_{k,M(k)} \in \mathcal{Z}$ are received at time k, then

$$Z_k = \{z_{k,1}, \ldots, z_{k,M(k)}\} \in \mathcal{F}(\mathcal{Z}). \tag{7}$$

Suppose at time k that the target is in state x_k. The measurement process is given by the RFS measurement equation

$$Z_k = \Theta_k(x_k) \cup E_k(x_k) \cup W_k, \tag{8}$$

where $\Theta_k(x_k)$ is the RFS of the primary target measurement, $E_k(x_k)$ is the RFS of extraneous target measurements, and W_k is the RFS of clutter. For example $\Theta_k(x_k)$ may represent a single direct path measurement, $E_k(x_k)$ may represent measurements generated by multi-path effects or counter measures, and W_k may represent state independent spurious measurements. It is assumed that $\Theta_k(x_k)$, $E_k(x_k)$ and W_k are independent RFSs.

We model $\Theta_k(x_k)$ as a binary RFS

$$\Theta_k(x_k) = \begin{cases} \varnothing & \text{with probability } 1 - p_{D,k}(x_k) \\ \{z_k^*\} & \text{with probability density } p_{D,k}(x_k)\, g_k(z_k^*|x_k) \end{cases}$$

where $p_{D,k}(x_k)$ is the probability of detection for the primary measurement given x_k, and $g_k(z_k^*|x_k)$ is the primary measurement likelihood given x_k. Hence, the probability of not obtaining the primary measurement is $1 - p_{D,k}(x_k)$, and conversely, given that there is a primary measurement the probability density of obtaining the primary measurement z_k^* given x_k is $g_k(z_k^*|x_k)$.

We model $E_k(x_k)$ and W_k as Poisson RFSs with intensities $v_{W,k}(\cdot)$ and $v_{E,k}(\cdot|x_k)$ respectively. For convenience we group these RFSs together as

$$K_k(x_k) = E_k(x_k) \cup W_k. \tag{9}$$

Since $K_k(x_k)$ is a union of statistically independent Poisson RFSs, it is also a Poisson RFS with intensity

$$v_{K,k}(z_k|x_k) = v_{W,k}(z_k) + v_{E,k}(z_k|x_k). \tag{10}$$

The cardinality distribution $\rho_{K,k}(\cdot|x_k)$ of $K_k(x_k)$ is Poisson with mean $\int v_{K,k}(z_k|x_k)\,dz_k$. Hence, if the target is in state x_k at time k, the probability of $K_k(x_k)$ having exactly n_k measurements is $\rho_{K,k}(n_k|x_k)$, whilst each measurement z_k is independent and identically distributed according to

$$c_k(z_k|x_k) = v_{K,k}(z_k|x_k) \,/ \int v_{K,k}(z_k|x_k)\,dz_k. \quad (11)$$

Proposition 1 *Suppose that measurements follow the RFS model in (8). Then, the probability density that the state x_k at time k produces the measurement set Z_k is given by*

$$\begin{aligned}\eta_k(Z_k|x_k) = {} & [1-p_{D,k}(x_k)]\rho_{K,k}(|Z_k|\,|x_k)\,|Z_k|!\prod_{z_k\in Z_k} c_k(z_k|x_k) \\ & + p_{D,k}(x_k)\cdot\rho_{K,k}(|Z_k|-1|x_k)\cdot(|Z_k|-1)! \\ & \times \sum_{z_k^*\in Z_k} g_k(z_k^*|x_k)\prod_{z_k\neq z_k^*} c_k(z_k|x_k) \end{aligned} \quad (12)$$

in the sense that $\eta_k(\cdot|x_k)$ is the Radon-Nikodým derivative of the probability distribution of Z_k given x_k with respect to the dominating measure (6).

The likelihood (12) has $|Z_k|+1$ terms each of which admits an intuitive interpretation. The first term relates to a missed primary measurement detection, whilst each of the remaining $|Z_k|$ terms relates to a primary measurement detection. To explain the first term, notice that when there is a missed primary measurement detection, $Z_k = K_k(x_k)$. Hence, the likelihood of Z_k comprises: $1 - p_{D,k}(x_k)$, the probability of a missed primary measurement detection; $\rho_{K,k}(|Z_k|\,|x_k)$, the probability that $K_k(x_k)$ has exactly $|Z_k|$ measurements; $\prod_{z_k\in Z_k} c_k(z_k|x_k)$, the joint density of the measurements; and a factorial term to account for all possible permutations of Z_k. To explain each of the $|Z_k|$ remaining terms, notice that when there is a primary measurement detection, $\Theta_k(x_k) = \{z_k^*\}$ and $K_k(x_k) = Z_k\backslash\{z_k^*\}$. Hence, the likelihood of Z_k comprises: $p_{D,k}(x_k)$, the probability of a primary measurement detection; $\rho_{K,k}(|Z_k|-1|x_k)$, the probability that $K_k(x_k)$ has exactly $|Z_k|-1$ measurements; $g_k(z_k^*|x_k)\prod_{z_k\neq z_k^*} c_k(z_k|x_k)$, the joint density of the measurements and a factorial term to account for all possible permutations of $K_k(x_k)$.

B. RFS Single Target Bayes Recursion

The Bayes recursion (1)-(2) can be generalized to accommodate multiple target generated measurements, detection uncertainty and clutter, by replacing the standard likelihood $g_k(z_k|x_k)$ with the RFS measurement likelihood (12). Hence, the posterior density $p_k(\cdot|Z_{1:k})$ can be propagated as follows

$$p_{k|k-1}(x_k|Z_{1:k-1}) = \int f_{k|k-1}(x_k|x)p_{k-1}(x|Z_{1:k-1})dx, \quad (13)$$

$$p_k(x_k|Z_{1:k}) = \frac{\eta_k(Z_k|x_k)p_{k|k-1}(x_k|Z_{1:k-1})}{\int \eta_k(Z_k|x)p_{k|k-1}(x|Z_{1:k-1})dx}, \quad (14)$$

where $Z_{1:k} = [Z_1, ..., Z_k]$. In general, no analytic solution exists for this recursion. In Section IV we present a generic sequential Monte Carlo implementation whilst in Section V a closed form solution to this recursion is derived under linear Gaussian assumptions.

Remark: If there is always a primary target generated measurement, no extraneous target generated measurements and no clutter, then $\eta_k(\{z_k\}|x_k) = g_k(z_k|x_k)$ and the recursion (13)-(14) reduces to the classical Bayes recursion (1)-(2).

Remark: Multiple sensors can be handled with a straightforward extension. Suppose that there are S mutually independent sensors and that each sensor is modelled by a likelihood $\eta_k^{(s)}(\cdot|\cdot)$ at time k where $s = 1, \ldots, S$. If each sensor receives a measurement set $Z_k^{(s)}$ at time k, where $s = 1, \ldots, S$, then the combined likelihood accounting for all sensors is $\eta_k(Z_k^{(1)}, \ldots, Z_k^{(S)}|x_k) = \prod_{s=1}^{S} \eta_k^{(s)}(Z_k^{(s)}|x_k)$.

IV. Sequential Monte Carlo Implementation

In this section, we describe a generic sequential Monte Carlo (SMC) (see also [12], [13]) implementation of the RFS single target Bayes recursion (13)-(14) and demonstrate the proposed filter on a non-linear tracking example.

Suppose at time $k-1$ that the posterior density $p_{k-1}(\cdot)$ is represented by set of weighted particles $\{w_{k-1}^{(i)}, x_{k-1}^{(i)}\}_{i=1}^{N}$, i.e.

$$p_{k-1}(x_{k-1}|Z_{1:k-1}) \approx \textstyle\sum_{i=1}^{N} w_{k-1}^{(i)}\delta_{x_{k-1}^{(i)}}(x_{k-1}). \quad (15)$$

Then, for a given proposal density $q_k(\cdot|x_{k-1}^{(i)}, Z_k)$ satisfying $support(p_k) \subseteq support(q_k)$, the particle filter approximates the posterior density $p_k(\cdot)$ by a new set of weighted particles $\{w_k^{(i)}, x_k^{(i)}\}_{i=1}^{N}$, i.e.

$$p_k(x_k|Z_{1:k}) \approx \textstyle\sum_{i=1}^{N} \tilde{w}_k^{(i)}\delta_{x_k^{(i)}}(x_k) \quad (16)$$

where

$$\begin{aligned} x_k^{(i)} &\sim q_k(\cdot|x_{k-1}^{(i)}, Z_k), & (17)\\ \tilde{w}_k^{(i)} &= w_k^{(i)} / \textstyle\sum_{i=1}^{N} w_k^{(i)}, & (18)\\ w_k^{(i)} &= w_{k-1}^{(i)}\frac{\eta_k(Z_k|x_k^{(i)})f_{k|k-1}(x_k^{(i)}|x_{k-1}^{(i)})}{q_k(x_k^{(i)}|x_{k-1}^{(i)}, Z_k)}. & (19)\end{aligned}$$

The recursion is initialized by generating a set of weighted particles $\{w_0^{(i)}, x_0^{(i)}\}_{i=1}^{N}$ representing p_0. Equations (17)-(19) then provide a recursion for computing the set of weighted particles representing p_k from those representing p_{k-1} when a new measurement arrives.

A resampling step is usually performed after each update to minimize particle degeneracy and after resampling, an optional Markov Chain Monte Carlo (MCMC) step can be used to increase particle diversity (see [12], [14] for further details).

A. Non-Linear Example

In this section, a non-linear scenario is used to demonstrate the performance of the particle implementation of the proposed filter. In particular, a nearly constant turn model with varying turn rate [15] together with bearing and range measurements is considered. The observation region is the half disc of radius $2000m$. The state variable $x_k = [\ \tilde{x}_k^T, \omega_k\ ]^T$ comprises the planar position and velocity $\tilde{x}_k^T = [\ p_{x,k}, \dot{p}_{x,k}, p_{y,k}, \dot{p}_{y,k}\ ]$ as well as the turn rate ω_k. The state transition model is

$$\begin{aligned}\tilde{x}_k &= F(\omega_{k-1})\tilde{x}_{k-1} + Gw_{k-1},\\ \omega_k &= \omega_{k-1} + \Delta u_{k-1},\end{aligned}$$

where

$$F(\omega) = \begin{bmatrix} 1 & \frac{\sin\omega\Delta}{\omega} & 0 & -\frac{1-\cos\omega\Delta}{\omega} \\ 0 & \cos\omega\Delta & 0 & -\sin\omega\Delta \\ 0 & \frac{1-\cos\omega\Delta}{\omega} & 1 & \frac{\sin\omega\Delta}{\omega} \\ 0 & \sin\omega\Delta & 0 & \cos\omega\Delta \end{bmatrix}, \; G = \begin{bmatrix} \frac{\Delta^2}{2} & 0 \\ \Delta & 0 \\ 0 & \frac{\Delta^2}{2} \\ 0 & \Delta \end{bmatrix},$$

$w_{k-1} \sim \mathcal{N}(\cdot; 0, \sigma_w^2 I)$, and $u_{k-1} \sim \mathcal{N}(\cdot; 0, \sigma_u^2 I)$ with $\Delta = 1s$, $\sigma_w = 5m/s^2$, and $\sigma_u = \pi/180 rad/s$. The observation region is the half disc $[-\pi/2, \pi/2]rad \times [0, 2000]m$. The primary target measurement is a noisy bearing and range vector

$$z_k = \left[\text{atan}(p_{x,k}/p_{y,k}), \sqrt{p_{x,k}^2 + p_{y,k}^2} \right]^T + \varepsilon_k,$$

where $\varepsilon_k \sim \mathcal{N}(\cdot; 0, R_k)$, with $R_k = \text{diag}([\ \sigma_\theta^2, \sigma_r^2\ ]^T)$, $\sigma_\theta = 2(\pi/180)rad$, and $\sigma_r = 10m$. The sensor field of view is modelled by

$$p_{D,k}(x) = \mathcal{N}([p_{x,k}, p_{y,k}]^T; 0, 2000 I_2).$$

where I_n denotes an $n \times n$ identity matrix. Extraneous measurements are modelled as a Poisson RFS with intensity

$$v_{E,k}(z|x) = \lambda_k^{(1)} \mathcal{N}(z; \left[\text{atan}(p_{x,k}/p_{y,k}), 2\sqrt{p_{x,k}^2 + p_{y,k}^2} \right]^T, D_k),$$

where $\lambda_k^{(1)} = 3$, $D_k = \sigma_\iota^2 I_2$ and $\sigma_\iota = 10m$. Clutter is modelled as a Poisson RFS with intensity

$$v_{W,k}(z) = \lambda_k^{(0)} u(z)$$

where $u(\cdot)$ is the uniform probability density over the observation region, $\lambda_k^{(0)} = \lambda_{c,k} V$ is the expected number of clutter returns with $V = 2000\pi radm$ is the 'volume' of observation region and $\lambda_{c,k} = 3.2 \times 10^{-3}\ (radm)^{-1}$ (giving an average of $\lambda_k^{(0)} = 20$ clutter returns per scan).

The transition density is used as the proposal, and resampling is performed at every step. A total of $N = 1000$ weighted particles is used to represent the posterior density at each time step. Figure 1 show the tracks, measurements and filter estimates for x and y coordinates versus time for each approximation on a typical sample run. This figure suggests that the proposed filter satisfactorily tracks the target in the presence of multiple target generated measurements, clutter and state dependent field of view.

For performance comparison purposes, we consider the non-linear analogue of the Gaussian mixture filter in [6]. Our reason for choosing this filter is that it subsumes many popular traditional techniques for tracking in clutter including the PDA. A typical sample run of this filter on the same set of data is also superimposed on Figure 1, which suggests that the traditional approach is prone to track loss. This is further reinforced in Figure 2, which shows the root mean square error (RMSE) versus clutter rate for both the proposed filter and the traditional filter. The RMSE for each clutter rate is obtained from 1000 Monte Carlo (MC) runs on the same target trajectory but with independently generated measurements for each trial. Figure 2 suggests that across a wide range of clutter conditions, the proposed RFS single-target Bayes filter performs well over traditional methods. The former correctly identifies the track, whereas the latter consistently looses track.

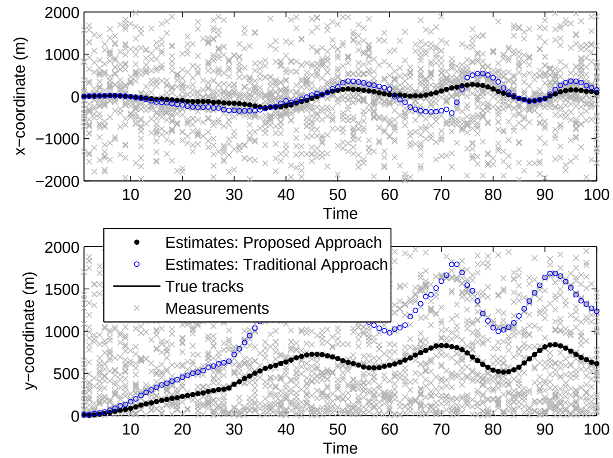

Fig. 1. Particle RFS single target Bayes filter estimates and true target tracks in x and y coordinates versus time.

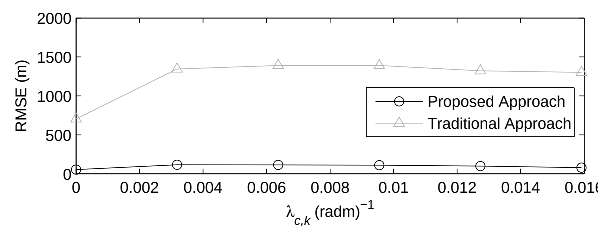

Fig. 2. RMSE from 1000 MC runs for varying $\lambda_{c,k}$.

V. Closed Form Solution for Linear Gaussian Models

In this section, we derive a closed form solution to the proposed filter for the class of linear Gaussian single target models. In addition to linear Gaussian transition and likelihood

$$f_{k|k-1}(x|\zeta) = \mathcal{N}(x; F_{k-1}\zeta, Q_{k-1}) \quad (20)$$
$$g_k(z|x) = \mathcal{N}(z; H_k x, R_k), \quad (21)$$

the linear Gaussian single target model assumes a constant sensor field of view, i.e. $p_{D,k}(x) = p_{D,k}$ and linear Gaussian intensity of extraneous target measurements i.e.

$$v_{E,k}(z|x) = \lambda_k^{(1)} \mathcal{N}(z; B_k x + b_k, D_k). \quad (22)$$

Proposition 2 *Suppose at time $k-1$ that the posterior density $p_{k-1}(\cdot)$ is a Gaussian mixture of the form*

$$p_{k-1}(x) = \sum_{j=1}^{J_{k-1}} w_{k-1}^{(j)} \mathcal{N}(x; m_{k-1}^{(j)}, P_{k-1}^{(j)}). \quad (23)$$

Then, the predicted density $p_{k|k-1}(\cdot)$ is also a Gaussian mixture and is given by

$$p_{k|k-1}(x) = \sum_{i=1}^{J_{k-1}} w_{k-1}^{(i)} \mathcal{N}(x; m_{k|k-1}^{(i)}, P_{k|k-1}^{(i)}) \quad (24)$$

where $m^{(i)}_{k|k-1} = F_{k-1}m^{(i)}_{k-1}$, $P^{(i)}_{k|k-1} = Q_{k-1} + F_{k-1}P^{(i)}_{k-1}F^T_{k-1}$.

For the closed form update equation, it is convenient to define two intermediate operators $\Phi_{k,z}$ and $\Psi_{k,z}$ on $\mathcal{X}$ by

$$(\Phi_{k,z}\phi)(x) = \Xi_z[g_k,\phi](x), \quad (25)$$

$$(\Psi_{k,z}\phi)(x) = \Xi_z[c_k,\phi](x), \quad (26)$$

where $(\Xi_z[s,\phi])(x) = s(z|x)\phi(x)$. If

$$s(z|x) = \bar{s}(z) + w_s\mathcal{N}(z; H_s x + b_s, P_s) \quad (27)$$

$$\phi(x) = \sum\nolimits_{u=1}^{U} w^{(u)}_\phi \mathcal{N}(x; m^{(u)}_\phi, P^{(u)}_\phi), \quad (28)$$

then, $(\Xi_z[s,\phi])(\cdot)$ is a Gaussian mixture and is given by

$$(\Xi_z[s,\phi])(x) = \bar{s}(z)\phi(x) + \sum_{u=1}^{U} w^{(u)}_\Xi(z)\mathcal{N}(x; m^{(u)}_\Xi(z), P^{(u)}_\Xi), \quad (29)$$

where

$$w^{(u)}_\Xi(z) = w_s w^{(u)}_\phi q^{(u)}_\Xi(z), \quad (30)$$

$$q^{(u)}_\Xi(z) = \mathcal{N}(z; \eta^{(u)}_\Xi + b_s, S^{(u)}_\Xi), \quad (31)$$

$$\eta^{(u)}_\Xi = H_s m^{(u)}_\phi, \quad (32)$$

$$S^{(u)}_\Xi = P_s + H_s P^{(u)}_\phi H^T_s, \quad (33)$$

$$m^{(u)}_\Xi(z) = m^{(u)}_\phi + K^{(u)}_\Xi(z - \eta^{(u)}_\Xi - b_s), \quad (34)$$

$$P^{(u)}_\Xi = (I - K^{(u)}_\Xi H_s)P^{(u)}_\phi, \quad (35)$$

$$K^{(u)}_\Xi = P^{(u)}_\phi H^T_s (S^{(u)}_\Xi)^{-1}. \quad (36)$$

Proposition 3 *Suppose at time k that the predicted density $p_{k|k-1}(\cdot)$ is a Gaussian mixture of the form*

$$p_{k|k-1}(x) = \sum_{j=1}^{J_{k|k-1}} w^{(j)}_{k|k-1}\mathcal{N}(x; m^{(j)}_{k|k-1}, P^{(j)}_{k|k-1}). \quad (37)$$

Then, $\eta_k(Z_k|x)p_{k|k-1}(x)$ is also a Gaussian mixture and is given by

$$\eta_k(Z_k|x)p_{k|k-1}(x) = \bar{d}_k(x) + \sum_{z^*\in Z_k} d_k(x; z^*) \quad (38)$$

where

$$\begin{aligned}\bar{d}_k(x) &= \rho_{K,k}(|Z_k|)\cdot|Z_k|!\cdot(1-p_{D,k}) \\ &\times \left(\left[\prod\nolimits_{z\in Z_k}\Psi_{k,z}\right]p_{k|k-1}\right)(x), \quad (39)\\ d_k(x;z^*) &= \rho_{K,k}(|Z_k|-1)\cdot(|Z_k|-1)!\cdot p_{D,k} \\ &\times \left(\left[\prod\nolimits_{z\neq z^*}\Psi_{k,z}\right]\left[\Phi_{k,z^*}p_{k|k-1}\right]\right)(x), \quad (40)\end{aligned}$$

and by convention a product of operators denotes a composition, i.e. $\prod_{i=1}^{N(k)}\Psi_{k,z_{k,i}} = \Psi_{k,z_{k,1}}\circ\Psi_{k,z_{k,2}}\circ\cdots\circ\Psi_{k,z_{k,N(k)}}$.

Remark: The Gaussian mixture (38) can also be written in as $\eta_k(Z_k|x)p_{k|k-1}(x) = \sum_{j=1}^{J_k} w^{(j)}_k\mathcal{N}(x; m^{(j)}_k, P^{(j)}_k)$ Consequently, the posterior density is given by $p_k(x) = \sum_{j=1}^{J_k}\tilde{w}^{(j)}_k\mathcal{N}(x; m^{(j)}_k, P^{(j)}_k)$ where $\tilde{w}^{(j)}_k = w^{(j)}_k/\sum_{j=1}^{J_k}w^{(j)}_k$ and $\sum_{j=1}^{J_k}w^{(j)}_k$ is the normalizing constant.

It follows by induction from Propositions 2 and 3 that if the initial density p_0 is a Gaussian mixture, then all subsequent predicted $p_{k|k-1}$ and posterior densities p_k are also Gaussian mixtures. Proposition 2 provides closed form expressions for computing the weights, means and covariances of $p_{k|k-1}$, whilst Proposition 3 provides closed form expressions for computing the weights, means and covariances of p_k when a new set of measurements arrives.

If the posterior at time $k-1$ has J_{k-1} components, then the posterior at time k has $J_{k-1}\left[2^{|Z_k|} + |Z_k|2^{|Z_k|-1}\right] = \mathcal{O}(J_{k-1}\cdot 2^{|Z_k|})$ components. To reduce this complexity, we only retain measurements that fall within a standard elliptical validation region around each mixture component of $p_{k|k-1}$; we truncate $(\Phi_{k,z}\phi)(\cdot)$ or $(\Phi_{k,z}\phi)(\cdot)$ to the $\hat{J}_{\max}$ terms with highest weights (whilst ensuring the sum of the weights before and after truncation must the same), we also discard components of p_k with negligible weights and merge components that are close together.

A. Linear Gaussian Example

The following linear Gaussian single-target model is used. The target state is a vector of position and velocity $x_k = [\, p_{x,k}, p_{y,k}, \dot{p}_{x,k}, \dot{p}_{y,k} \,]^T$ and follows a linear Gaussian transition model (20) with

$$F_k = \begin{bmatrix} I_2 & \Delta I_2 \\ 0_2 & I_2 \end{bmatrix}, \qquad Q_k = \sigma^2_\nu \begin{bmatrix} \frac{\Delta^4}{4}I_2 & \frac{\Delta^3}{2}I_2 \\ \frac{\Delta^3}{2}I_2 & \Delta^2 I_2 \end{bmatrix},$$

where I_n and 0_n denote the $n\times n$ identity and zero matrices respectively, $\Delta = 1s$ is the sampling period, and $\sigma_\nu = 5(m/s^2)$ is the standard deviation of the process noise. The primary measurement likelihood is linear Gaussian (21) with

$$H_k = \begin{bmatrix} I_2 & 0_2 \end{bmatrix}, \qquad R_k = \sigma^2_\varepsilon I_2,$$

where $\sigma_\varepsilon = 10m$ is the standard deviation of the measurement noise. The observation region is the square $\mathcal{Z} = [-1000, 1000]\times[-1000, 1000]$ (units are in m). The corresponding probability of detection is fixed at $p_{D,k} = 0.98$. Extraneous target measurements are modelled as a Poisson RFS with linear Gaussian intensity

$$v_{E,k}(z|x) = \lambda^{(1)}_k\mathcal{N}(z; B_k x, D_k),$$

where $\lambda^{(1)}_k = 3$, $B_k = \begin{bmatrix} 2I_2 & 0_2 \end{bmatrix}$, $D_k = \sigma^2_\iota I_2$ and $\sigma_\iota = 10m$. Clutter is modelled as a Poisson RFS with intensity

$$v_{W,k}(z) = \lambda^{(0)}_k u(z)$$

where $u(\cdot)$ is the uniform probability density over $\mathcal{Z}$, $\lambda^{(0)}_k = \lambda_{c,k}V$, $\lambda_{c.k} = 1.25\times 10^{-5}m^{-2}$ is the average clutter intensity and $V = 4\times 10^6 m^2$ is the 'volume' of $\mathcal{Z}$ (giving an average of $\lambda^{(0)}_k = 50$ clutter returns per scan).

In this scenario, the target follows a curved path with varying velocity. The filter is initialized with the true initial location. Figure 3 illustrates a typical sample run showing the tracks, measurements and filter estimates for x and y coordinates versus time. This figure suggests that our proposed filter correctly identifies the track, and does not suffer any track losses in the presence of multiple target generated measurements and clutter.

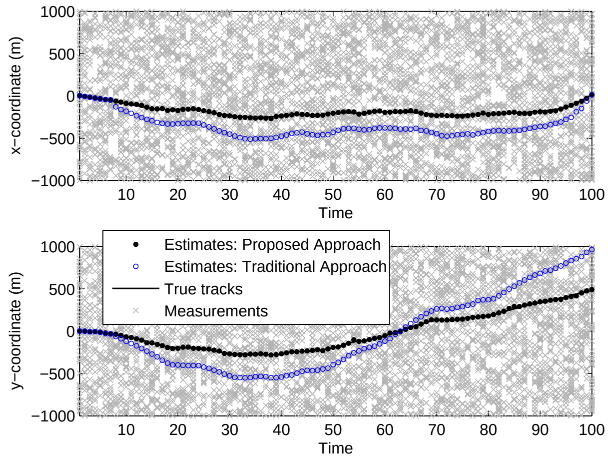

Fig. 3. Linear Gaussian RFS single-target Bayes filter estimates and true target tracks in x and y coordinates versus time.

Similar to the non-linear example, we compare with the Gaussian mixture filter in [6]. Figure 3 has superimposed a typical sample run for the same data. Again, it can be seen that traditional approaches tend to loose the track and erroneously follow the pattern of the extraneous target measurements. This observation is supported by the results of 1000 Monte Carlo (MC) runs performed on the same target trajectory but with independently generated measurements for each trial. The MC runs are performed for our proposed filter, and for the filter in [6]. In Figure 4, the RMSE is shown versus the clutter rate suggesting that across a wide range of conditions, the proposed RFS single-target Bayes filter performs well over traditional methods. The former correctly identifies the true tracks, whereas the latter consistently loses the true track.

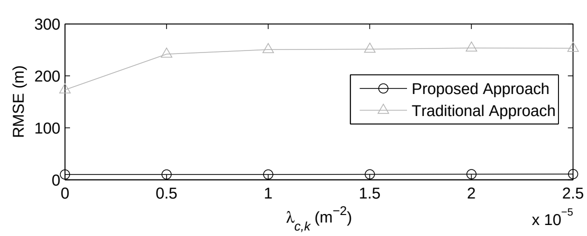

Fig. 4. RMSE values from 1000 MC runs for varying $\lambda_{c,k}$

VI. Conclusion

This paper has presented a Bayes recursion that formally accommodates multiple target-generated measurements, detection uncertainty and clutter. The proposed Bayes recursion (referred to as the RFS single-target Bayes recursion) has been derived from the random finite set or point process framework using standard measure theoretic probability. A particle implementation has been given and a closed form solution has been derived for linear Gaussian single-target models. The closed formed solution can be easily extended to nonlinear models via linearization or unscented transformation. Simulations have suggested that the proposed filter performs well compared to traditional techniques in the presence of multiple target generated measurements, clutter and detection uncertainty.

References

[1] Y. Bar-Shalom and T. E. Fortmann, *Tracking and Data Association*. Academic Press, San Diego, 1988.

[2] D. Reid, “An algorithm for tracking multiple targets,” *IEEE Trans. AC*, vol. AC-24, no. 6, pp. 843–854, 1979.

[3] S. Blackman, *Multiple Target Tracking with Radar Applications*. Artech House, Norwood, 1986.

[4] I. Goodman, R. Mahler, and H. Nguyen, *Mathematics of Data Fusion*. Kluwer Academic Publishers, 1997.

[5] R. Mahler, “Multi-target Bayes filtering via first-order multi-target moments,” *IEEE Trans. AES*, vol. 39, no. 4, pp. 1152–1178, 2003.

[6] D. Salmond, “Mixture reduction algorithms for target tracking in clutter,” in *Proc. SPIE Signal and Data Processing of Small Targets*, vol. 1305, pp. 434–445, 1990.

[7] D. Daley and D. Vere-Jones, *An introduction to the theory of point processes*. Springer-Verlag, 1988.

[8] D. Stoyan, D. Kendall, and J. Mecke, *Stochastic Geometry and its Applications*. John Wiley & Sons, 1995.

[9] J. Moller and R. Waagepetersen, *Statistical Inference and Simulation for Spatial Point Processes*. Chapman & Hall CRC, 2004.

[10] G. Mathéron, *Random sets and integral geometry*. J. Wiley, 1975.

[11] M. N. van Lieshout, *Markov Point Processes and their Applications*. Imperial College Press, 2000.

[12] A.Doucet, N. de Freitas, and N. Gordon, *Sequential Monte Carlo Methods in Practice*. Springer-Verlag, 2001.

[13] S. Arulampalam, S. Maskell, N. J. Gordon, and T. Clapp, “A tutorial on particle filters for on-line non-linear/non-Gaussian Bayesian tracking,” *IEEE Trans. SP*, vol. 50, no. 2, pp. 174–188, 2002.

[14] R. Gilks and C. Berzuini, “Following a moving target: Monte Carlo inference for dynamic Bayesian models,” in *Journal of the Royal Statistical Society,*, vol. 63, no. 1, pp. 127–146, 2001.

[15] X.-R. Li and V. Jilkov, “Survey of maneuvering target tracking,” *IEEE Trans. AES*, vol. 39, no. 4, pp. 1333–1364, 2003.