On multi-Bernoulli approximations to the Bayes multi-target filter

B.-T. Vo
WATRI
University of Western Australia
Crawley, WA, Australia.
vob@watri.org.au

B.-N. Vo
EEE Department
University of Melbourne
Parkville, VIC, Australia.
bv@ee.unimelb.edu.au

A. Cantoni
WATRI
University of Western Australia
Crawley, WA, Australia.
cantoni@watri.org.au

Abstract - ***Mahler recently proposed the Multi-target Multi-Bernoulli (MeMBer) recursion as a tractable approximation to the Bayes multi-target recursion, and outlined a Gaussian mixture solution under linear Gaussian assumptions. These proposals are speculative in the sense that, to date, no implementations have been reported.***

In this paper, it is shown analytically that the MeMBer recursion has a significant bias in cardinality that results in a high number of false tracks. A novel approximation that alleviates the bias problem is proposed. In addition, a sequential Monte Carlo implementation (for generic models) and a Gaussian mixture implementation (for linear Gaussian models) are given. Comparisons with Mahler's original MeMBer filter via simulations show significant reduction of false tracks.

Keywords: Tracking, estimation, random sets, point processes, finite set statistics, multi-Bernoulli.

1 Introduction

Multi-target filtering, involves the joint estimation of the number of targets and their individual states from a sequence of observations in the presence of detection uncertainty, association uncertainty and clutter. Mahler's finite set statistics (FISST) is an elegant Bayesian formulation of multi-target filtering based on random finite set (RFS) theory, which has generated substantial interest in recent years due to the development of the Probability Hypothesis Density (PHD) filter [1] and the Cardinalized PHD (CPHD) filter [2]. The PHD/CPHD filters are moment approximations of the Bayes multi-target filter, which operate on the single-object state space and avoid the combinatorial problem that arises from data association. Sequential Monte Carlo (SMC) implementations with provable convergence [3], and closed form solutions [4, 5] to the PHD/CPHD recursions have opened the doors to numerous novel extensions and applications, see for example [6, 7] and the references therein.

In addition to the PHD/CPHD filters, Mahler recently proposed the **M**ulti-targ**e**t **M**ulti-**Ber**noulli (MeMBer) recursion as a tractable approximation to the Bayes multi-target recursion under low clutter density scenarios [6]. Unlike the PHD/CPHD recursions which propagate moments and cardinality distributions, the MeMBer recursion propagates (approximately) the multi-target posterior density. A Gaussian mixture solution to the MeMBer recursion under linear Gaussian assumptions was also proposed in [6]. However, no implementations have been reported to-date.

In this paper, we show analytically that the MeMBer filter significantly over-estimates the number of targets, which would result in a high number of false tracks. This bias is caused by a multi-Bernoulli approximation in the MeMBer update step. We propose a novel multi-Bernoulli approximation, called the cardinality-balanced MeMBer update, that alleviates the bias problem (under the same signal setting). A generic sequential Monte Carlo (SMC) implementation applicable to nonlinear dynamic and measurement models with state-dependent probability of detection is given. Unlike the SMC-PHD filters, the SMC-MeMBer approach does not require clustering to extract estimates of target states. A Gaussian mixture implementation for linear Gaussian models is also given. We compare the proposed filter with Mahler's original MeMBer filter via simulations, and demonstrate significant reduction of false tracks whilst maintaining similar state estimation accuracy.

The paper is organised as follows. The necessary background on RFSs and multi-target filtering is given in Section 2. In Section 3 we present a review of Mahler's MeMBer recursion, followed by a derivation of the cardinality bias, and descriptions of the cardinality-balanced MeMBer update. A generic sequential Monte Carlo implementation for non-linear

multi-target models (with simulation results) is described in Section 4. A Gaussian mixture implementation for linear multi-target models (with simulation results) is given in Section 5. Closing remarks are given in Section 6.

2 Background

This section introduces RFSs, multi-target system models and the multi-target Bayes filter.

2.1 Random Finite Sets

Intuitively, a *random finite set* (RFS), is a random (spatial) point pattern, e.g. measurements on a radar screen or positions of rain drops on a spatial area. What distinguishes an RFS from a random vector is that: the number of points is random; the points themselves are random and unordered. In essence, an RFS is simply a finite-set-valued random variable. At the fundamental level, like any other random variable, the randomness of an RFS is captured by its probability distribution (or probability density if one exists).

Another fundamental descriptor of an RFS, which has direct relevance to this paper, is the probability generating functional (PGFl). Let $\mathcal{F}(\mathcal{X})$ denote the space of finite subsets of $\mathcal{X}$. Suppose X is an RFS on $\mathcal{X}$, i.e. X is a random variable taking values in $\mathcal{F}(\mathcal{X})$. Following [8, 9], the *probability generating functional* (PGFl) $G[\cdot]$ of X is defined by

$$G[h] \equiv \mathbb{E}[h^X], \tag{1}$$

where $\mathbb{E}$ denotes the expectation operator, h is any (reasonable) real-valued function on $\mathcal{X}$ such that $0 \leq h(x) \leq 1$, and

$$h^X = \prod_{x \in X} h(x),$$

with $h^{\emptyset} = 1$ by convention. The probability generating function $G(\cdot)$ of the cardinality distribution of X can be obtained by substituting $h(x) = y$ into the PGFl.

2.2 Poisson RFSs

An RFS X on $\mathcal{X}$ is said to be *Poisson* with *intensity function* v (defined on $\mathcal{X}$) if its cardinality distribution, $\Pr(|X| = n)$, is Poisson with mean $\bar{N} = \int v(x)dx$, and for any finite cardinality, the elements x of X are independently and identically distributed (i.i.d) according to the probability density $v(\cdot)/\bar{N}$ [8, 9]. A Poisson RFS is completely characterized by its intensity function, also known in the tracking literature as the *Probability Hypothesis Density* (PHD). A Poisson RFS with intensity function v has PGFl

$$G[h] = e^{\langle v, h-1 \rangle}$$

where $\langle v, h \rangle = \int v(x)h(x)dx$. The probability density[1] of a Poisson RFS can also be explicitly expressed in terms of v as follows

$$\pi(X) = e^{-\bar{N}} \prod_{x \in X} v(x) = e^{-\bar{N}} v^X.$$

2.3 Multi-Bernoulli RFS

A *Bernoulli* RFS X on $\mathcal{X}$ has probability $1 - r$ of being empty, and probability r of being a singleton whose (only) element is distributed according to a probability density p (defined on $\mathcal{X}$). The cardinality distribution of a Bernoulli RFS is a Bernoulli distribution with parameter r. The PGFl of a Bernoulli RFS is given by

$$G[h] = 1 - r + r \langle p, h \rangle . \tag{2}$$

A *multi-Bernoulli* RFS X on $\mathcal{X}$ is a union of a fixed number of independent Bernoulli RFSs $X^{(i)}$ with existence probability $r^{(i)} \in (0, 1)$ and probability density $p^{(i)}$ (defined on $\mathcal{X}$), $i = 1, ..., M$, i.e.

$$X = \bigcup_{i=1}^{M} X^{(i)} \tag{3}$$

The PGFl of a multi-Bernoulli RFS is

$$G[h] = \prod_{i=1}^{M} \left(1 - r^{(i)} + r^{(i)} \langle p^{(i)}, h \rangle\right). \tag{4}$$

A multi-Bernoulli RFS is thus completely described by the multi-Bernoulli parameter set $\{(r^{(i)}, p^{(i)})\}_{i=1}^{M}$. Moreover, the probability density π is given by $\pi(\emptyset) = \prod_{j=1}^{M} \left(1 - r^{(j)}\right)$ and

$$\pi(\{x_1, ..., x_n\}) = n! \sum_{1 \leq i_1 < \cdots < i_n \leq M} \prod_{j=1}^{n} \frac{\left(1 - r^{(j)}\right) r^{(i_j)} p^{(i_j)}(x_j)}{1 - r^{(i_j)}}. \tag{5}$$

Also, the mean of the cardinality distribution of a multi-Bernoulli RFS is given by $\bar{N} = \mathbb{E}[|X|] = \sum_{i=1}^{M} r^{(i)}$. Throughout this paper, we abbreviate probability density of the form (5) by $\pi = \{(r^{(i)}, p^{(i)})\}_{i=1}^{M}$. We also refer to a PGFl of the form (4), or a multi-target density of the form (5), as a multi-Bernoulli.

2.4 Multi-target System Model

Suppose that at time k, there are $N(k)$ target states $x_{k,1}, \ldots, x_{k,N(k)}$, each taking values in a state space $\mathcal{X} \subseteq \mathbb{R}^{n_x}$, and $M(k)$ measurements $z_{k,1}, \ldots, z_{k,M(k)}$ each taking values in an observation space $\mathcal{Z} \subseteq \mathbb{R}^{n_z}$. In the random finite set approach, the finite sets of targets and observations, at time k, [10, 1, 6] are treated as the *multi-target state* and *multi-target observation*, respectively

$$X_k = \{x_{k,1}, \ldots, x_{k,N(k)}\} \in \mathcal{F}(\mathcal{X}),$$
$$Z_k = \{z_{k,1}, \ldots, z_{k,M(k)}\} \in \mathcal{F}(\mathcal{Z}).$$

Using RFS theory, we can construct stochastic models for the time evolution of the multi-target state and the multi-target observations as follows.

[1] For simplicity, in this paper, we shall not distinguish FISST set derivative and probability density. While the FISST set derivative of an RFS is not a probability density [6], it is equivalent to a probability density (see [3]).

Given a multi-target state X_{k-1} at time $k-1$, each $x_{k-1} \in X_{k-1}$ either continues to exist at time k with probability $p_{S,k}(x_{k-1})$ and move to a new state x_k with probability density $f_{k|k-1}(x_k|x_{k-1})$, or dies with probability $1 - p_{S,k}(x_{k-1})$. Thus, for a given target with state $x_{k-1} \in X_{k-1}$ at time $k-1$, its behaviour at the next time step is modeled by the Bernoulli RFS

$$S_{k|k-1}(x_{k-1})$$

with $r = p_{S,k}(x_{k-1})$ and $p(\cdot) = f_{k|k-1}(\cdot|x_{k-1})$. This transition is commonly known in point process theory as a Markov shift [8]. The RFS of the multi-target state X_k at time k is given by the union

$$X_k = \left[\bigcup_{x_{k-1} \in X_{k-1}} S_{k|k-1}(x_{k-1}) \right] \cup \Gamma_k, \qquad (6)$$

where Γ_k denotes the multi-Bernoulli RFS of spontaneous births. The RFS multi-target transition equation (6) incorporates target motion, birth and death. Assuming that the RFSs constituting the union in (6) are mutually independent, X_k is a multi-Bernoulli RFS conditional on X_{k-1}.

A given target $x_k \in X_k$, at time k, is either detected with probability $p_{D,k}(x_k)$ and generates an observation z_k with likelihood $g_k(z_k|x_k)$, or missed with probability $1 - p_{D,k}(x_k)$, i.e. each state $x_k \in X_k$ generates a Bernoulli RFS

$$\Theta_k(x_k)$$

with $r = p_{D,k}(x_k)$ and $p(\cdot) = g_k(\cdot|x_k)$. In addition, the sensor also receives a set of false alarms or clutter which can be modeled as a Poisson RFS K_k with intensity function $\kappa_k(\cdot)$. Thus, at time k, the multi-target measurement Z_k generated by a multi-target state X_k is formed by the union

$$Z_k = \left[\bigcup_{x \in X_k} \Theta_k(x) \right] \cup K_k \qquad (7)$$

The RFS multi-target measurement equation (7) accounts for detection uncertainty and clutter. It is assumed that the RFSs constituting the union in (7) are independent of one another. Note that conditional on X_k, the target generated measurements in (7) form a multi-Bernoulli RFS.

2.5 Multi-target Bayes Recursion

The multi-target filtering problem can be posed as a Bayesian filtering problem with state space $\mathcal{F}(\mathcal{X})$ and observation space $\mathcal{F}(\mathcal{Z})$. Let $\pi_k(\cdot|Z_{1:k})$ denote the *multi-target posterior density* at time k. Then, the *multi-target Bayes recursion* propagates $\pi_k(\cdot|Z_{1:k})$ in time [10, 1, 3] according to

$$\pi_{k|k-1}(X_k|Z_{1:k-1}) = \int f_{k|k-1}(X_k|X)\pi_{k-1}(X|Z_{1:k-1})\delta X, \qquad (8)$$

$$\pi_k(X_k|Z_{1:k}) = \frac{g_k(Z_k|X_k)\pi_{k|k-1}(X_k|Z_{1:k-1})}{\int g_k(Z_k|X)\pi_{k|k-1}(X|Z_{1:k-1})\delta X}, \qquad (9)$$

where the integrals in the above recursion are FISST set integrals (see [10, 1, 6]), $f_{k|k-1}(\cdot|\cdot)$ is the *multi-target transition density*[2] and $g_k(\cdot|\cdot)$ is the *multi-target likelihood*[2]. The multi-target transition encapsulates the underlying models of target motions, births and deaths described in (6), while the multi-target likelihood encapsulates the underlying models of detections and false alarms described in (7). Explicit expressions for $f_{k|k-1}(X_k|X_{k-1})$ and $g_k(Z_k|X_k)$ can be derived from (6) and (7) using FISST, see for example [10, 1, 6].

3 MeMBer Approximations

The MeMBer recursion, proposed by Mahler in [6], is an approximation to the full multi-target Bayes recursion (8-9) using multi-Bernoulli RFS. Intuitively, the MeMBer recursion propagates the multi-target posterior probability density in time by propagating a finite but time-varying number of hypothesized tracks, each characterized by the probability of existence and the probability density of the current hypothesized state.

In this section, we briefly summarize Mahler's MeMBer recursion and derive the cardinality bias that occurs in the update step. Moreover, we propose a MeMBer update step that is unbiased in cardinality, herein referred to as the cardinality-balanced MeMBer update.

3.1 The Original MeMBer Recursion

The premise of the MeMBer recursion is that the multi-target RFS at each time step is approximated by a multi-Bernoulli RFS, based on the following modelling assumptions:

- Each target evolves and generates measurements independently,
- Target births follow a multi-Bernoulli RFS independent of target survivals,
- Clutter follows a Poisson RFS, not too dense, and is independent of target-generated measurements.

Proposition 1 (MeMBer Prediction [6]). *Suppose that, at time $k-1$, the posterior multi-target density is a multi-Bernoulli of the form*

$$\pi_{k-1} = \{(r_{k-1}^{(i)}, p_{k-1}^{(i)})\}_{i=1}^{M_{k-1}}.$$

Then, the predicted multi-target density is also a multi-Bernoulli and is given by

$$\pi_{k|k-1} = \{(r_{P,k|k-1}^{(i)}, p_{P,k|k-1}^{(i)})\}_{i=1}^{M_{k-1}} \cup \{(r_{\Gamma,k}^{(i)}, p_{\Gamma,k}^{(i)})\}_{i=1}^{M_{\Gamma,k}}, \qquad (10)$$

[2] The same notation is used for multi-target and single-target densities. There is no danger of confusion since in the single-target case the arguments are vectors whereas in the multi-target case the arguments are finite sets.

where

$$r_{P,k|k-1}^{(i)} = r_{k-1}^{(i)} \langle p_{k-1}^{(i)}, p_{S,k} \rangle, \tag{11}$$

$$p_{P,k|k-1}^{(i)}(x) = \frac{\langle f_{k|k-1}(x|\cdot), p_{k-1}^{(i)} p_{S,k} \rangle}{\langle p_{k-1}^{(i)}, p_{S,k} \rangle}, \tag{12}$$

$f_{k|k-1}(\cdot|\zeta)$ = *single target transition density at time k, given previous state ζ,*

$p_{S,k}(\zeta)$ = *probability of target existence at time k, given previous state ζ,*

$\{(r_{\Gamma,k}^{(i)}, p_{\Gamma,k}^{(i)})\}_{i=1}^{M_{\Gamma,k}}$ = *parameters of the multi-Bernoulli RFS of births at time k.*

In essence, the multi-Bernoulli parameter set for the predicted multi-target density $\pi_{k|k-1}$ is formed by the union of the multi-Bernoulli parameter sets for the surviving targets and births. The total number of predicted hypothesized tracks is $M_{k|k-1} = M_{k-1} + M_{\Gamma,k}$.

Proposition 2 (MeMBer Update [6]). *Suppose that, at time k, the predicted multi-target density is a multi-Bernoulli of the form*

$$\pi_{k|k-1} = \{(r_{k|k-1}^{(i)}, p_{k|k-1}^{(i)})\}_{i=1}^{M_{k|k-1}}.$$

Then, the posterior multi-target density can be approximated by a multi-Bernoulli as follows

$$\pi_k \approx \{(r_{L,k}^{(i)}, p_{L,k}^{(i)})\}_{i=1}^{M_{k|k-1}} \cup \{(r_{U,k}(z), p_{U,k}(\cdot; z))\}_{z \in Z_k}, \tag{13}$$

where

$$r_{L,k}^{(i)} = r_{k|k-1}^{(i)} \frac{1 - \langle p_{k|k-1}^{(i)}, p_{D,k} \rangle}{1 - r_{k|k-1}^{(i)} \langle p_{k|k-1}^{(i)}, p_{D,k} \rangle}, \tag{14}$$

$$p_{L,k}^{(i)}(x) = p_{k|k-1}^{(i)}(x) \frac{1 - p_{D,k}(x)}{1 - \langle p_{k|k-1}^{(i)}, p_{D,k} \rangle}, \tag{15}$$

$$r_{U,k}(z) = \frac{\langle \tilde{v}_{k|k-1}, \psi_{k,z} \rangle}{\kappa_k(z) + \langle \tilde{v}_{k|k-1}, \psi_{k,z} \rangle}, \tag{16}$$

$$p_{U,k}(x; z) = \frac{\tilde{v}_{k|k-1}(x) \psi_{k,z}(x)}{\langle \tilde{v}_{k|k-1}, \psi_{k,z} \rangle}, \tag{17}$$

$$\tilde{v}_{k|k-1}(x) = \sum_{i=1}^{M_{k|k-1}} \frac{r_{k|k-1}^{(i)} p_{k|k-1}^{(i)}(x)}{1 - r_{k|k-1}^{(i)} \langle p_{k|k-1}^{(i)}, p_{D,k} \rangle}, \tag{18}$$

$\psi_{k,z}(x) = g_k(z|x) p_{D,k}(x)$,

Z_k = *measurement set at time k,*

$g_k(\cdot|x)$ = *single target measurement likelihood at time k, given current state x,*

$p_{D,k}(x)$ = *probability of target detection at time k given current state x,*

$\kappa_k(\cdot)$ = *intensity of Poisson clutter at time k.*

It is implicitly assumed that $p_{D,k}$ and $r_{k|k-1}^{(i)}, i = 1, ..., M_{k|k-1}$ cannot all be equal to 1. In essence, the multi-Bernoulli parameter set for the updated multi-target density π_k is formed by the union of the multi-Bernoulli parameter sets for the legacy (missed-detection) tracks and measurement-corrected tracks. The total number of posterior hypothesized tracks is $M_k = M_{k|k-1} + |Z_k|$.

3.2 Cardinality Bias

It follows from Proposition 1 that the mean cardinality of the predicted multi-target state is

$$\bar{N}_{k|k-1} = \sum_{i=1}^{M_{k-1}} r_{P,k|k-1}^{(i)} + \sum_{i=1}^{M_{\Gamma,k}} r_{\Gamma,k}^{(i)}.$$

Proposition 2 implies that the mean posterior cardinality, $\bar{N}_k$, is approximated by

$$\tilde{N}_k = \sum_{i=1}^{M_{k|k-1}} r_{L,k}^{(i)} + \sum_{z \in Z_k} r_{U,k}(z). \tag{19}$$

In the following we derive the mean posterior cardinality. As shown in [6], if clutter is not too dense, the PGFl corresponding to the updated multi-target density π_k is given by

$$G_k[h] = \prod_{i=1}^{M_{k|k-1}} G_{L,k}^{(i)}[h] \prod_{z \in Z_k} G_{U,k}[z; h] \tag{20}$$

where

$$G_{L,k}^{(i)}[h] = \frac{1 - r_{k|k-1}^{(i)} + r_{k|k-1}^{(i)} \langle p_{k|k-1}^{(i)}, h q_{D,k} \rangle}{1 - r_{k|k-1}^{(i)} + r_{k|k-1}^{(i)} \langle p_{k|k-1}^{(i)}, q_{D,k} \rangle}, \tag{21}$$

$$G_{U,k}[h; z] = \frac{\kappa_k(z) + \sum_{i=1}^{M_{k|k-1}} G_{U,k}^{(i)}[h; z]}{\kappa_k(z) + \langle \tilde{v}_{k|k-1}, \psi_{k,z} \rangle}, \tag{22}$$

$$G_{U,k}^{(i)}[h; z] = \frac{r_{k|k-1}^{(i)} \langle p_{k|k-1}^{(i)}, h \psi_{k,z} \rangle}{1 - r_{k|k-1}^{(i)} + r_{k|k-1}^{(i)} \langle p_{k|k-1}^{(i)}, h q_{D,k} \rangle}, \tag{23}$$

$$q_{D,k} = 1 - p_{D,k}.$$

The product, $\prod_{i=1}^{M_{k|k-1}} G_{L,k}^{(i)}[h]$, corresponds to the set of legacy tracks and is a multi-Bernoulli, since (21) can be rewritten in Bernoulli form

$$G_{L,k}^{(i)}[h] = 1 - r_{L,k}^{(i)} + r_{L,k}^{(i)} \langle p_{L,k}^{(i)}, h \rangle.$$

with $r_{L,k}^{(i)}$ and $p_{L,k}^{(i)}$ given as in Proposition 2. Hence the mean cardinality of the legacy tracks is

$$\bar{N}_{L,k} = \sum_{i=1}^{M_{k|k-1}} r_{L,k}^{(i)}. \tag{24}$$

The product, $\prod_{z \in Z_k} G_{U,k}[h; z]$, corresponds to the set of measurement-updated tracks, which is not a multi-Bernoulli RFS. Nonetheless, its mean cardinality $\bar{N}_{U,k}$ can be determined by: substituting $h(x) = y$ into the PGFl (22); then differentiating at $y = 1$; and summing over all $z \in Z_k$, yielding

$$\bar{N}_{U,k} = \sum_{z \in Z_k} \frac{\langle v_{k|k-1}^{(1)}, \psi_{k,z} \rangle}{\kappa_k(z) + \langle \tilde{v}_{k|k-1}, \psi_{k,z} \rangle}, \tag{25}$$

where

$$v_{k|k-1}^{(1)}(x) = \sum_{i=1}^{M_{k|k-1}} \frac{r_{k|k-1}^{(i)} (1 - r_{k|k-1}^{(i)}) p_{k|k-1}^{(i)}(x)}{(1 - r_{k|k-1}^{(i)} \langle p_{k|k-1}^{(i)}, p_{D,k} \rangle)^2}. \tag{26}$$

Proposition 3. *Under the premises of Proposition 2, the posterior mean number of targets at time k is*

$$\bar{N}_k = \bar{N}_{L,k} + \bar{N}_{U,k}, \tag{27}$$

where $\bar{N}_{L,k}$ and $\bar{N}_{U,k}$ are given by (24) and (25) respectively.

Corollary to Proposition 3. *Under the premises of Proposition 2, the posterior cardinality bias at time k is*

$$\tilde{N}_k - \bar{N}_k = \sum_{z\in Z_k} \frac{\langle \tilde{v}_{k|k-1} - v^{(1)}_{k|k-1}, \psi_{k,z}\rangle}{\kappa_k(z) + \langle \tilde{v}_{k|k-1}, \psi_{k,z}\rangle}. \tag{28}$$

Since

$$\tilde{v}_{k|k-1}(x) - v^{(1)}_{k|k-1}(x) = \sum_{i=1}^{M_{k|k-1}} \frac{(r^{(i)}_{k|k-1})^2(1-\langle p^{(i)}_{k|k-1}, p_{D,k}\rangle)p^{(i)}_{k|k-1}(x)}{(1-r^{(i)}_{k|k-1}\langle p^{(i)}_{k|k-1}, p_{D,k}\rangle)^2},$$

each term of the sum in (28) is non-negative and equals zero only when $p_{D,k} = 1$. Hence the bias $\tilde{N}_k - \bar{N}_k$ is non-negative and is zero only when $p_{D,k} = 1$. Experimental results, see Sections 4 and 5, demonstrated significant cardinality bias.

3.3 Cardinality Balancing

Observe from (22) that $G_{U,k}[h;z]$ is not in Bernoulli form, and it is necessary to find a Bernoulli approximation in order to express $G_k[\cdot]$ as a multi-Bernoulli. Mahler achieved this by setting $h = 1$ in the denominator of $G^{(i)}_{U,k}[h;z]$, i.e.

$$G^{(i)}_{U,k}[h;z] \approx \frac{r^{(i)}_{k|k-1}\langle p^{(i)}_{k|k-1}, h\psi_{k,z}\rangle}{1-r^{(i)}_{k|k-1}\langle p^{(i)}_{k|k-1}, p_{D,k}\rangle}, \tag{29}$$

which, after substituting into (22), yields the following Bernoulli approximation

$$G_{U,k}[h;z] \approx 1 - r_{U,k}(z) + r_{U,k}(z)\langle p_{U,k}(\cdot;z), h\rangle, \tag{30}$$

and subsequently the update step in Proposition 2. This particular approximation leads to a bias in the cardinality of the measurement-updated tracks and consequently the posterior cardinality bias (28).

A simple way of correcting the bias is to set $r_{U,k}(z)$ in (16) to $\bar{N}_{U,k}$ the mean cardinality of the measurement-updated tracks, given in (25). It is also possible to derive other cardinality-unbiased approximate updates by first setting $r_{U,k}(z)$ to $\bar{N}_{U,k}$ and then linearizing the PGFl $G_{U,k}[\cdot;z]$ to obtain the probability density $p_{U,k}(\cdot;z)$ (note that linearizing $G_{U,k}[\cdot;z]$ alone does not necessarily result in a Bernoulli PGFl). Using this strategy with linearization at $h = 0$ leads to the following update.

Proposition 4 (Cardinality-Balanced MeMBer Update). *Under the premises of Proposition 2, the posterior multi-target density can be approximated with unbiased cardinality by a multi-Bernoulli as follows*

$$\pi_k \approx \{(r^{(i)}_{L,k}, p^{(i)}_{L,k})\}_{i=1}^{M_{k|k-1}} \cup \{(r^*_{U,k}(z), p^*_{U,k}(\cdot;z))\}_{z\in Z_k}, \tag{31}$$

where, $r^{(i)}_{L,k}$, $p^{(i)}_{L,k}$ are given in (14-15),

$$r^*_{U,k}(z) = \frac{\langle v^{(1)}_{k|k-1}, \psi_{k,z}\rangle}{\kappa_k(z) + \langle \tilde{v}_{k|k-1}, \psi_{k,z}\rangle}, \tag{32}$$

$$p^*_{U,k}(x;z) = \frac{\tilde{v}^*_{k|k-1}(x)\psi_{k,z}(x)}{\langle \tilde{v}^*_{k|k-1}, \psi_{k,z}\rangle}, \tag{33}$$

$$\tilde{v}^*_{k|k-1}(x) = \sum_{i=1}^{M_{k|k-1}} \frac{r^{(i)}_{k|k-1}p^{(i)}_{k|k-1}(x)}{1-r^{(i)}_{k|k-1}}, \tag{34}$$

$v^{(1)}_{k|k-1}$ and $\tilde{v}_{k|k-1}$ are given in (26) and (18) respectively, $\psi_{k,z}$ and $\kappa_k(z)$ are defined in Proposition 2.

4 Sequential Monte Carlo Implementation

In this section, we present a generic SMC implementation of the cardinality-balanced MeMBer recursion which can accommodate non-linear dynamic and measurement models, as well as non-constant probabilities of survival and detection. This technique directly extends to Mahler's original MeMBer recursion.

4.1 Prediction

Suppose that at time $k-1$ the (multi-Bernoulli) posterior multi-target density $\pi_{k-1} = \{(r^{(i)}_{k-1}, p^{(i)}_{k-1})\}_{i=1}^{M_{k-1}}$ is given and each $p^{(i)}_{k-1}$, $i = 1, ..., M_{k-1}$, is comprised of a set of weighted samples $\{w^{(i,j)}_{k-1}, x^{(i.j)}_{k-1}\}_{j=1}^{L^{(i)}_{k-1}}$, i.e.

$$p^{(i)}_{k-1}(x) = \sum_{j=1}^{L^{(i)}_{k-1}} w^{(i,j)}_{k-1}\delta_{x^{(i.j)}_{k-1}}(x).$$

Then, given importance (or proposal) densities $q^{(i)}_k(\cdot|x_{k-1}, Z_k)$ and $b^{(i)}_k(\cdot|Z_k)$ such that

$$\begin{aligned} support(p^{(i)}_k) &\subseteq support(q^{(i)}_k) \\ support(p^{(i)}_{\Gamma,k}) &\subseteq support(b^{(i)}_k) \end{aligned}$$

the predicted (multi-Bernoulli) multi-target density $\pi_{k|k-1} = \{(r^{(i)}_{P,k|k-1}, p^{(i)}_{P,k|k-1})\}_{i=1}^{M_{k-1}} \cup \{(r^{(i)}_{\Gamma,k}, p^{(i)}_{\Gamma,k})\}_{i=1}^{M_{\Gamma,k}}$ can be computed as follows

$$r^{(i)}_{P,k|k-1} = r^{(i)}_{k-1}\sum_{j=1}^{L^{(i)}_{k-1}} w^{(i,j)}_{k-1}p_{S,k}(x^{(i.j)}_{k-1}), \tag{35}$$

$$p^{(i)}_{P,k|k-1}(x) = \sum_{j=1}^{L^{(i)}_{k-1}} \tilde{w}^{(i,j)}_{P,k|k-1}\delta_{x^{(i.j)}_{P,k|k-1}}(x), \tag{36}$$

$$r^{(i)}_{\Gamma,k} = \text{parameter given by birth model}, \tag{37}$$

$$p^{(i)}_{\Gamma,k}(x) = \sum_{j=1}^{L^{(i)}_{\Gamma,k}} \tilde{w}^{(i,j)}_{\Gamma,k}\delta_{x^{(i.j)}_{\Gamma,k}}(x), \tag{38}$$

where

$$
\begin{aligned}
x_{P,k|k-1}^{(i.j)} &\sim q_k^{(i)}(\cdot|x_{k-1}^{(i.j)}, Z_k),\ j = 1, ..., L_{k-1}^{(i)} \\
w_{P,k|k-1}^{(i,j)} &= \frac{w_{k-1}^{(i,j)} f_{k|k-1}(x_{P,k|k-1}^{(i.j)}|x_{k-1}^{(i.j)}) p_{S,k}(x_{k-1}^{(i.j)})}{q_k^{(i)}(x_{P,k|k-1}^{(i.j)}|x_{k-1}^{(i.j)}, Z_k)}, \\
\tilde{w}_{P,k|k-1}^{(i,j)} &= w_{P,k|k-1}^{(i,j)} / \sum_{j=1}^{L_{k-1}^{(i)}} w_{P,k|k-1}^{(i,j)} \\
x_{\Gamma,k}^{(i.j)} &\sim b_k^{(i)}(\cdot|Z_k)\ j = 1, ..., L_{\Gamma,k}^{(i)} \\
w_{\Gamma,k}^{(i,j)} &= \frac{p_{\Gamma,k}(x_{\Gamma,k}^{(i.j)})}{b_k^{(i)}(x_{\Gamma,k}^{(i.j)}|Z_k)}, \\
\tilde{w}_{\Gamma,k}^{(i,j)} &= w_{\Gamma,k}^{(i,j)} / \sum_{j=1}^{L_{\Gamma,k}^{(i)}} w_{\Gamma,k}^{(i,j)}.
\end{aligned}
$$

4.2 Update

Suppose that at time k the predicted (multi-Bernoulli) multi-target density $\pi_{k|k-1} = \{(r_{k|k-1}^{(i)}, p_{k|k-1}^{(i)})\}_{i=1}^{M_{k|k-1}}$ is given and each $p_{k|k-1}^{(i)}$, $i = 1, ..., M_{k|k-1}$, is comprised of a set of weighted samples $\{w_{k|k-1}^{(i,j)}, x_{k|k-1}^{(i.j)}\}_{j=1}^{L_{k-1}^{(i)}}$, i.e.

$$
p_{k|k-1}^{(i)} = \sum_{j=1}^{L_{k|k-1}^{(i)}} w_{k|k-1}^{(i,j)} \delta_{x_{k|k-1}^{(i.j)}}(x).
$$

Then, (the multi-Bernoulli approximation of) the updated multi-target density $\pi_k = \{(r_{L,k}^{(i)}, p_{L,k}^{(i)})\}_{i=1}^{M_{k|k-1}} \cup \{(r_{U,k}(z), p_{U,k}(\cdot; z))\}_{z \in Z_k}$, can be computed as follows

$$
r_{L,k}^{(i)} = r_{k|k-1}^{(i)} \frac{1 - \varrho_{L,k}^{(i)}}{1 - r_{k|k-1}^{(i)} \varrho_{L,k}^{(i)}}, \tag{39}
$$

$$
p_{L,k}^{(i)}(x) = \sum_{j=1}^{L_{k|k-1}^{(i)}} \tilde{w}_{L,k}^{(i,j)} \delta_{x_{k|k-1}^{(i.j)}}(x), \tag{40}
$$

$$
r_{U,k}^{*}(z) = \frac{\sum_{i=1}^{M_{k|k-1}} \frac{r_{k|k-1}^{(i)}(1 - r_{k|k-1}^{(i)}) \varrho_{U,k}^{(i)}(z)}{(1 - r_{k|k-1}^{(i)} \varrho_{L,k}^{(i)})^2}}{\kappa_k(z) + \sum_{i=1}^{M_{k|k-1}} \frac{r_{k|k-1}^{(i)} \varrho_{U,k}^{(i)}(z)}{1 - r_{k|k-1}^{(i)} \varrho_{L,k}^{(i)}}}, \tag{41}
$$

$$
p_{U,k}^{*}(x; z) = \sum_{i=1}^{M_{k|k-1}} \sum_{j=1}^{L_{k|k-1}^{(i)}} \tilde{w}_{U,k}^{*(i,j)}(z) \delta_{x_{k|k-1}^{(i.j)}}(x), \tag{42}
$$

where

$$
\begin{aligned}
\varrho_{L,k}^{(i)} &= \sum_{j=1}^{L_{k|k-1}^{(i)}} w_{k|k-1}^{(i,j)} p_{D,k}(x_{k|k-1}^{(i.j)}) \\
\tilde{w}_{L,k}^{(i,j)} &= w_{L,k}^{(i,j)} / \sum_{j=1}^{L_{k|k-1}^{(i)}} w_{L,k}^{(i,j)} \\
w_{L,k}^{(i,j)} &= w_{k|k-1}^{(i,j)} (1 - p_{D,k}(x_{k|k-1}^{(i.j)})) \\
\varrho_{U,k}^{(i)}(z) &= \sum_{j=1}^{L_{k|k-1}^{(i)}} w_{k|k-1}^{(i,j)} \psi_{k,z}(x_{k|k-1}^{(i.j)}) \\
\tilde{w}_{U,k}^{*(i,j)}(z) &= w_{U,k}^{*(i,j)}(z) / \sum_{i=1}^{M_{k|k-1}} \sum_{j=1}^{L_{k|k-1}^{(i)}} w_{U,k}^{*(i,j)}(z) \\
w_{U,k}^{*(i,j)}(z) &= w_{k|k-1}^{(i,j)} \frac{r_{k|k-1}^{(i)}}{1 - r_{k|k-1}^{(i)}} \psi_{k,z}(x_{k|k-1}^{(i.j)}).
\end{aligned}
$$

4.3 Implementation Issues

Notice that the number of particles required to represent the predicted density increases due to the birth of targets in the prediction, and that the number of particles required to represent the posterior density increases due to involvement of the pseudo-PHD $\tilde{v}_{k|k-1}^{*}$ in the update. Moreover, since the MeMBer is an inherent multi-target filter, similar to the SMC-PHD filter [3] it is desirable to allocate the number of particles representing the posterior density at each time step to be proportional to the expected number of targets present. For example, this can be achieved by allocating the number of particles for each hypothesized track to be proportional to its probability of existence, i.e. if a maximum of $L_{\max}$ particles is allocated per hypothesized track, then during the prediction we sample $L_{\Gamma,k}^{(i)} = r_{\Gamma,k}^{(i)} L_{\max}$ particles per birth term, and following the update we resample $L_k^{(i)} = r_k^{(i)} L_{\max}$ for each measurement-corrected track. In this case, it is also important to impose a minimum of $L_{\min}$ number of particles per track to ensure sufficient particle diversity.

4.4 Demonstration

In this example, we demonstrate the performance of the SMC cardinality-balanced MeMBer recursion for bearings and range tracking of a time varying number of targets in clutter. A nearly constant turn model having varying turn rate together with bearing and range measurements is considered. The observation region is the half disc of radius $2000m$. A total of 10 targets is considered in which there are various births and deaths throughout the simulation. The target state variable $x_k = [\ \tilde{x}_k^T, \omega_k\ ]^T$ comprises the planar position and velocity $\tilde{x}_k^T = [\ p_{x,k}, \dot{p}_{x,k}, p_{y,k}, \dot{p}_{y,k}\ ]^T$ as well as the turn rate ω_k. The state transition model is

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

where

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

$w_{k-1} \sim \mathcal{N}(\cdot; 0, \sigma_w^2 I)$, and $u_{k-1} \sim \mathcal{N}(\cdot; 0, \sigma_u^2 I)$ with $\Delta = 1s$, $\sigma_w = 15m/s^2$, and $\sigma_u = \pi/180 rad/s$. The sensor observation is a noisy bearing and range vector given by

$$
z_k = \begin{bmatrix} \arctan(p_{x,k}/p_{y,k}) \\ \sqrt{p_{x,k}^2 + p_{y,k}^2} \end{bmatrix} + \varepsilon_k
$$

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 birth process is multi Bernoulli with density $\pi_\Gamma = \{(r_\Gamma, p_\Gamma^{(i)})\}_{i=1}^4$ where $r_\Gamma = 0.1$, $p_\Gamma^{(i)}(x) = \mathcal{N}(x; m_\gamma^{(i)}, P_\gamma)$, $m_\gamma^{(1)} = [\ -1500, 0, 250, 0\ 0\ ]^T$,

$m_\gamma^{(2)} = [\ -250, 0, 1000, 0\ \ 0\ ]^T$, $m_\gamma^{(3)} = [\ 250, 0, 750, 0\ \ 0\ ]^T$, $m_\gamma^{(4)} = [\ 1000, 0, 1500, 0\ \ 0\ ]^T$, $P_\gamma = \mathrm{diag}([\ 50,\ 50,\ 50,\ 50,\ 6(\pi/180)\ ]^T)^2$. The probability of target survival and detection are $p_{S,k}(x) = 0.99$ and $p_{D,k}(x) = 0.98\mathcal{N}([p_{x,k}, p_{y,k}]^T; 0, 6000^2 I_2)/\mathcal{N}(0; 0, 6000^2 I_2)$ respectively. Clutter is Poisson with intensity $\lambda_c = 4.8 \times 10^{-4}\ (radm)^{-1}$ over the region $[-\pi/2, \pi/2]rad \times [0, 2000]m$ (giving an average of 3 clutter points per scan).

At each time step, a maximum and minimum of $L_{\max} = 1000$ and $L_{\min} = 300$ particles per hypothesized track is imposed, and resampling is performed for each hypothesized track so that the number of particles representing each track is proportional to its probability of existence. Additionally, at each time step track pruning is performed with a weight threshold of $P = 10^{-3}$ and a maximum of $T_{max} = 100$ tracks. State estimates are calculated by first using an MAP estimate on the cardinality distribution for the number of targets, and then using an EAP estimate on each of the corresponding number of hypothesized tracks with the highest probabilities of existence for the states.

Fig. 1 shows the x and y components of the measurements, true trajectories and filter output against time. It can be seen that the SMC cardinality-balanced MeMBer filter is able to identify all target births and track the non-linear motion well, as well as the target deaths at $k = 66$ and $k = 80$. Notice also that the filter has no trouble handling the three targets which cross paths at time $k = 50$, and another two which cross paths at $k = 80$.

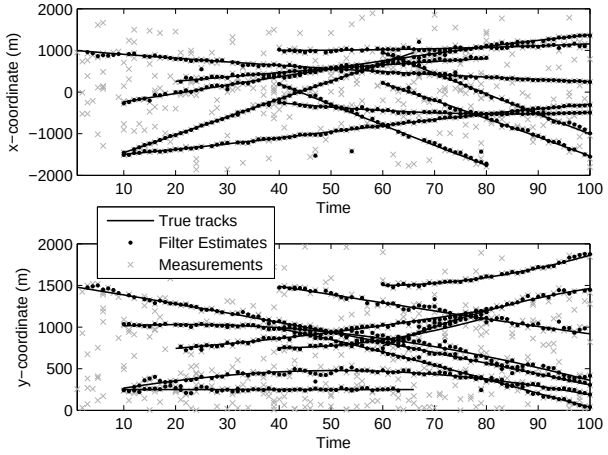

Figure 1: Measurements, true tracks, and filter estimates.

We now verify the performance of the SMC cardinality-balanced MeMBer filter by simulating over 1000 Monte Carlo (MC) runs. For each run, the same target tracks are used but the measurement data is randomly generated. Fig. 2 shows the MC average of the mean and standard deviation of the estimated cardinality distribution against time. Fig. 2 also compares these results with the output from the original MeMBer filter. It can be seen that cardinality-balanced filter provides unbiased estimates of the number of targets, whereas the original filter has a significant positive bias in this respect. Moreover, these observations confirm the theoretical developments presented in Propositions 3 and 4.

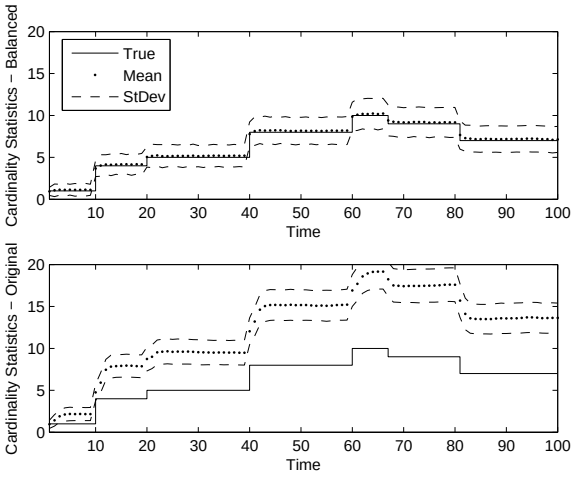

Figure 2: Cardinality statistics for cardinality-balanced filter (top), and original filter (bottom)

To evaluate the overall performance of the cardinality-balanced and original MeMBer filters, Fig. 3 shows the MC average of the Wasserstein distance (WD) [11] between the estimated and true multi-target state against time. These results suggest that the cardinality-balanced filter outperforms the original filter, overall in terms of the accuracy of state and cardinality estimates jointly. Nonetheless, it is worth reporting the following empirical observations regarding their performance. In this example at least, the bias in measurement-corrected terms of the original MeMBer recursion subsequently causes it to overestimate the existence probabilities for the legacy tracks (the tracks corresponding to missed detections). Furthermore, the estimated existence probabilities for these legacy tracks are usually similar in magnitude to those of the measurement-corrected tracks, and hence for each target present the filter tends to choose as candidates for state estimates both the corresponding legacy and measurement-corrected tracks. The net result is that the original MeMBer filter takes its two best guesses for each target present, namely both the predicted track and its best measurement-corrected track. Such behaviour becomes significant in low signal-to-noise ratio (SNR) conditions, such as the relatively lower probability of detection in this example, since even with missed detections the original MeMBer filter's multi-target state estimate still contains the predicted state (and thus may appear to have good localization performance albeit at the expense of significant cardinality error). On the other hand, the cardinality-balanced MeMBer filter having the correct estimate for the number of targets, only takes one guess for each target present, or less when it underestimates the number of targets present. Consequently, the cardinality-balanced MeMBer filter's multi-target state estimate

occasionally drops target estimates even though it retains their tracks in its propagation of the posterior (and thus may be penalized on localization performance).

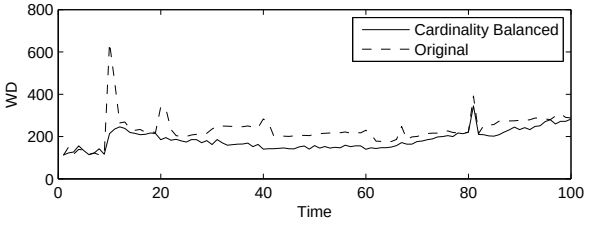

Figure 3: Wasserstein distance (WD) for the cardinality-balanced and original filters against time

5 Closed Form Linear Gaussian Implementation

In this section, we present a closed form solution to the MeMBer recursion for the class of linear Gaussian multi-target models which consists of standard linear Gaussian assumptions for the transition and observation models of individual targets, as well as certain assumptions on birth, death and detection:

- Each target follows a linear Gaussian dynamical model i.e.

$$f_{k|k-1}(x|\zeta) = \mathcal{N}(x; F_{k-1}\zeta, Q_{k-1}), \quad (43)$$
$$g_k(z|x) = \mathcal{N}(z; H_k x, R_k), \quad (44)$$

  where $\mathcal{N}(\cdot; m, P)$ denotes a Gaussian density with mean m and covariance P, F_{k-1} is the state transition matrix, Q_{k-1} is the process noise covariance, H_k is the observation matrix, and R_k is the observation noise covariance.

- The survival and detection probabilities are state independent, i.e.

$$p_{S,k}(x) = p_{S,k}, \quad (45)$$
$$p_{D,k}(x) = p_{D,k}. \quad (46)$$

- The birth model is a multi-Bernoulli with parameter set $\{(r^{(i)}_{\Gamma,k}, p^{(i)}_{\Gamma,k})\}_{i=1}^{M_{\Gamma,k}}$ where $p^{(i)}_{\Gamma,k}$, $= 1..., M_{\Gamma,k}$ are Gaussian mixtures of the form

$$p^{(i)}_{\Gamma,k}(x) = \sum_{j=1}^{J^{(i)}_{\Gamma,k}} w^{(i,j)}_{\Gamma,k} \mathcal{N}(x; m^{(i,j)}_{\Gamma,k}, P^{(i,j)}_{\Gamma,k}), \quad (47)$$

  with $w^{(i,j)}_{\Gamma,k}$, $m^{(i,j)}_{\Gamma,k}$, $P^{(i,j)}_{\Gamma,k}$ denoting the weights, means and covariances of the jth component.

For the linear Gaussian multi-target model, the following two subsections present a closed form solution to the MeMBer recursion by showing how the posterior density is analytically propagated in time.

5.1 Prediction

Suppose that at time $k-1$ the (multi-Bernoulli) posterior multi-target density $\pi_{k-1} = \{(r^{(i)}_{k-1}, p^{(i)}_{k-1})\}_{i=1}^{M_{k-1}}$ is given and each probability density $p^{(i)}_{k-1}$, $i = 1, ..., M_{k-1}$, is comprised of Gaussian mixtures of the form

$$p^{(i)}_{k-1}(x) = \sum_{j=1}^{J^{(i)}_{k-1}} w^{(i,j)}_{k-1} \mathcal{N}(x; m^{(i.j)}_{k-1}, P^{(i.j)}_{k-1}).$$

Then, the predicted (multi-Bernoulli) multi-target density $\pi_{k|k-1} = \{(r^{(i)}_{P,k|k-1}, p^{(i)}_{P,k|k-1})\}_{i=1}^{M_{k-1}} \cup \{(r^{(i)}_{\Gamma,k}, p^{(i)}_{\Gamma,k})\}_{i=1}^{M_{\Gamma,k}}$ can be computed as follows: $(r^{(i)}_{\Gamma,k}, p^{(i)}_{\Gamma,k})$, $i, ..., M_{\Gamma,k}$ are given by the birth model (47), while

$$r^{(i)}_{P,k|k-1} = r^{(i)}_{k-1} p_{S,k}, \quad (48)$$

$$p^{(i)}_{P,k|k-1}(x) = \sum_{j=1}^{J^{(i)}_{k-1}} w^{(i,j)}_{k-1} \mathcal{N}(x; m^{(i,j)}_{P,k|k-1}, P^{(i,j)}_{P,k|k-1}), \quad (49)$$

where

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

5.2 Update

Suppose that at time k the predicted (multi-Bernoulli) multi-target density $\pi_{k|k-1} = \{(r^{(i)}_{k|k-1}, p^{(i)}_{k|k-1})\}_{i=1}^{M_{k|k-1}}$ is given and each probability density $p^{(i)}_{k|k-1}$, $i = 1, ..., M_{k-1}$, is comprised of Gaussian mixtures of the form

$$p^{(i)}_{k|k-1}(x) = \sum_{j=1}^{J^{(i)}_{k-1}} w^{(i,j)}_{k-1} \mathcal{N}(x; m^{(i.j)}_{k-1}, P^{(i.j)}_{k-1}).$$

Then, (the multi-Bernoulli approximation of) the updated density $\pi_k = \{(r^{(i)}_{L,k}, p^{(i)}_{L,k})\}_{i=1}^{M_{k|k-1}} \cup \{(r_{U,k}(z), p_{U,k}(\cdot; z))\}_{z \in Z_k}$ can be computed as follows

$$r^{(i)}_{L,k} = r^{(i)}_{k|k-1} \frac{1 - p_{D,k}}{1 - r^{(i)}_{k|k-1} p_{D,k}}, \quad (50)$$

$$p^{(i)}_{L,k}(x) = p^{(i)}_{k|k-1}(x), \quad (51)$$

$$r^*_{U,k}(z) = \frac{\sum_{i=1}^{M_{k|k-1}} \frac{r^{(i)}_{k|k-1}(1 - r^{(i)}_{k|k-1}) \varrho^{(i)}_{U,k}(z)}{(1 - r^{(i)}_{k|k-1} p_{D,k})^2}}{\kappa_k(z) + \sum_{i=1}^{M_{k|k-1}} \frac{r^{(i)}_{k|k-1} \varrho^{(i)}_{U,k}(z)}{1 - r^{(i)}_{k|k-1} p_{D,k}}}, \quad (52)$$

$$p^*_{U,k}(x; z) = \frac{\sum_{i=1}^{M_{k|k-1}} \sum_{j=1}^{J^{(i)}_{k-1}} w^{(i,j)}_{U,k}(z) \mathcal{N}(x; m^{(i.j)}_{U,k}, P^{(i.j)}_{U,k})}{\sum_{i=1}^{M_{k|k-1}} \sum_{j=1}^{J^{(i)}_{k-1}} w^{(i,j)}_{U,k}(z)}, \quad (53)$$

where

$$\varrho_{U,k}^{(i)}(z) = p_{D,k} \sum_{j=1}^{J_{k-1}^{(i)}} w_{k-1}^{(i,j)} q_k^{(i,j)}(z)$$
$$q_k^{(i,j)}(z) = \mathcal{N}(z; H_k m_{k|k-1}^{(i,j)}, H_k P_{k|k-1}^{(i,j)} H_k^T + R_k),$$
$$w_{U,k}^{(i,j)}(z) = \frac{r_{k|k-1}^{(i)}}{1 - r_{k|k-1}^{(i)}} p_{D,k} w_{k-1}^{(i,j)} q_k^{(i,j)}(z),$$
$$m_{U,k}^{(i,j)}(z) = m_{k|k-1}^{(i,j)} + K_{U,k}^{(i,j)}(z - H_k m_{k|k-1}^{(i,j)}),$$
$$P_{U,k}^{(i,j)} = [I - K_{U,k}^{(i,j)} H_k] P_{k|k-1}^{(i,j)},$$
$$K_{U,k}^{(i,j)} = P_{k|k-1}^{(i,j)} H_k^T \left[H_k P_{k|k-1}^{(i,j)} H_k^T + R_k \right]^{-1}.$$

5.3 Implementation Issues

Notice that the number of Gaussians required to represent the predicted density increases due to the birth of targets in the prediction, and that the number of Gaussians required to represent the posterior density increases due to involvement of the pseudo-PHD $\tilde{v}_{k|k-1}^*$ in the update. Analogous to the Gaussian mixture PHD and CPHD filters [4, 5] it is necessary to perform pruning and merging of Gaussian components at each time step in order to mitigate the unbounded growth of mixture components. Specifically, for each hypothesized track this involves eliminating components with weights below a threshold T, merging components within a distance U of each other, and imposing a maximum of J_{max} components, see [4, 5] for the exact meaning of these parameters.

5.4 Demonstration

In this example, we demonstrate the performance of the linear Gaussian cardinality-balanced MeMBer recursion for a time varying number of targets observed in clutter over a two dimensional region. The surveillance region is the square $[-1000, 1000]m \times [-1000, 1000]m$. A 10 target scenario is considered here with births and deaths at various times and locations. The target state variable is a vector of planar position and velocity $x_k = [\ p_{x,k}, p_{y,k}, \dot{p}_{x,k}, \dot{p}_{y,k}\ ]^T$. The single-target transition model is linear Gaussian specified by

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

where I_n and 0_n denote the $n \times n$ identity and zero matrices, $\Delta = 1s$ is the sampling period, and $\sigma_\nu = 5(m/s^2)$ is the standard deviation of the process noise. The probability of target survival $p_{S,k} = 0.99$. The birth process is multi Bernoulli with density $\pi_\Gamma = \{(r_\Gamma, p_\Gamma^{(i)})\}_{i=1}^4$ where $r_\Gamma = 0.03$, $p_\Gamma^{(i)}(x) = \mathcal{N}(x; m_\Gamma^{(i)}, P_\Gamma)$, $m_\Gamma^{(1)} = [\ 0, 0, 0, 0\ ]^T$, $m_\Gamma^{(2)} = [\ 400, 0, -600, 0\ ]^T$, $m_\Gamma^{(3)} = [\ -800, 0, -200, 0\ ]^T$, $m_\Gamma^{(4)} = [\ -200, 0, -800, 0\ ]^T$, $P_\Gamma = \mathrm{diag}([\ 10,\ 10,\ 10,\ 10\ ]^T)^2$. The probability of target detection is $p_{D,k} = 0.98$. The single-target measurement model is also linear Gaussian with

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

where $\sigma_\varepsilon = 10m$, is the standard deviation of the measurement noise. Clutter is Poisson with intensity $\kappa_k(z) = \lambda_c V u(z)$, where $u(\cdot)$ is a uniform probability density over the surveillance region, $V = 4 \times 10^6 m^2$ is the 'volume' of the surveillance region, and $\lambda_c = 2.5 \times 10^{-6} m^{-2}$ is the clutter intensity (giving an average of 10 clutter returns per scan).

At each time step, Gaussian pruning and merging is performed for each hypothesized track using a weight threshold of $T = 10^{-5}$, a merging threshold of $U = 4m$, and a maximum of $J_{max} = 100$ components. Additionally, at each time step track pruning is performed with a weight threshold of $P = 10^{-3}$ and a maximum of $T_{max} = 100$ tracks. State estimates are calculated by first using an MAP estimate on the cardinality distribution for number of targets, and then using an MAP estimate on each of the corresponding number of hypothesized tracks with the highest probabilities of existence for the states.

Fig. 4 plots the x and y components of the measurements, true trajectories and filter estimates against time. It can be seen that the linear Gaussian MeMBer filter correctly identifies all target births and tracks them accordingly, as well as the two target deaths at time $k = 70$. Here, the filter has no difficulty handling the three targets which cross paths at time $k = 40$, and another two which cross paths at time $k = 60$.

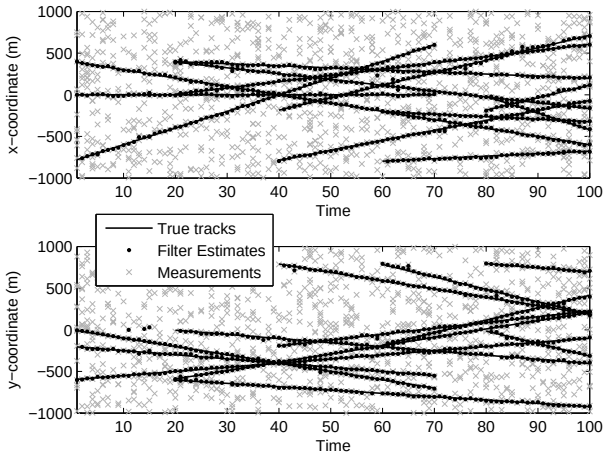

Figure 4: Measurements, true tracks, and filter estimates.

Similar to the SMC example, we verify the performance of the linear Gaussian cardinality-balanced MeMBer filter over 1000 Monte Carlo (MC) runs using the one set of target tracks but with randomly generated measurement data. Fig. 5 shows the MC average of the mean and standard deviation of the estimated cardinality distribution against time, and compares these results with the output from the original MeMBer filter. Again, these results confirm that the cardinality-balanced filter provides unbiased estimates

of the number of targets, whereas the original filter provides significantly biased ones.

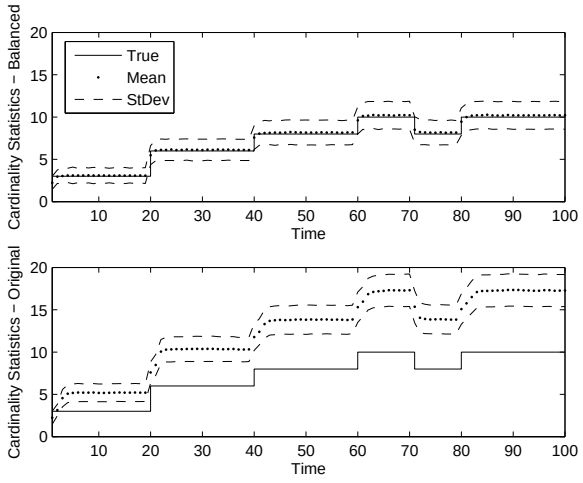

Figure 5: Cardinality statistics for cardinality-balanced filter (top), original filter (bottom)

Similarly, to evaluate the overall performance, Fig. 6 shows the MC average of the Wasserstein distance (WD) [11] between the estimated and true multi-target state against time. Again, these results suggest that the cardinality-balanced filter outperforms the original filter, overall in terms of the accuracy of state and cardinality estimates jointly. In regards to false tracks and accuracy of state estimates, similar empirical observations reported in the SMC examples also apply to this linear Gaussian example. However, the effect of this behaviour is less pronounced in this example since the probability of detection is relatively high.

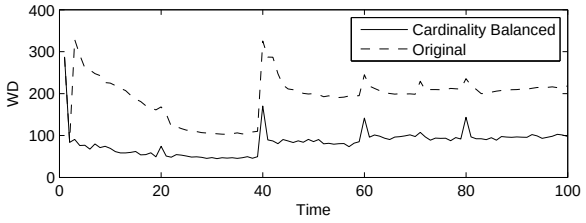

Figure 6: Wasserstein distance (WD) for the cardinality-balanced and original filters against time

6 Conclusion

It has been shown, analytically and experimentally, in this paper that Mahler's MeMBer filter has a bias in the number of targets. Moreover, this bias vanishes only for unity probability of detection. The cause of the bias has been identified in the update step of the filter and a new update has been proposed, which eliminates the posterior cardinality bias. Experiments for linear and nonlinear scenarios confirm that the proposed remedy is unbiased whilst maintaining similar state estimation accuracy as the original MeMBer filter.

Similar to the closed form implementations of the PHD/CPHD filters [4, 5], extensions of the closed form implementation of the MeMBer filter to mildly nonlinear Gaussian models with constant probability of detection can be achieved via linearization and unscented transformation.

References

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

[2] ——, "PHD filters of higher order in target number (to appear)," *IEEE Trans. AES*, vol. 43, no. 3, July 2007.

[3] B.-N. Vo, S. Singh, and A. Doucet, "Sequential Monte Carlo methods for multi-target filtering with random finite sets," in *IEEE Trans. AES*, vol. 41, no. 4, pp. 1224–1245, 2005.

[4] B.-N. Vo and W. Ma, "The Gaussian mixture probability hypothesis density filter," *IEEE Trans. SP.*, vol. 54, no. 11, pp. 4091–4104, Nov. 2006.

[5] B.-T. Vo, B.-N. Vo, and A. Cantoni, "Analytic implementations of the cardinalized probability hypothesis density filter," *IEEE Trans. SP.*, vol. 55, no. 7, pp. 3553–3567, July. 2007.

[6] R. Mahler, *Statistical Multisource-Multitarget Information Fusion.* Artech House, 2007.

[7] ——, "A survey of PHD filter and CPHD filter implementations," *Signal Processing, Sensor Fusion, and Target Recognition XV, SPIE Defense & Security Symposium*, April 2007.

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

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

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

[11] J. Hoffman and R. Mahler, "Multitarget miss distance via optimal assignment," *IEEE Trans. Sys., Man, and Cybernetics-Part A*, vol. 34, no. 3, pp. 327–336, 2004.