Improved SMC implementation of the PHD filter

Branko Ristic
ISR Division
DSTO
Melbourne
Australia
branko.ristic@dsto.defence.gov.au

Daniel Clark
EECE EPS
Heriot-Watt University
Edinburgh
United Kingdom
D.E.Clark@hw.ac.uk

Ba-Ngu Vo
School of EECE
The University of Western Australia
Crawley, WA 6009,
Australia
bnvo@ee.uwa.edu.au

Abstract – *The paper makes two contributions. First, a new formulation of the PHD filter which distinguishes between persistent and newborn objects is presented. This formulation results in an efficient sequential Monte Carlo (SMC) implementation of the PHD filter, where the placement of newborn object particles is determined by the measurements. The second contribution is a novel method for the state and error estimation from an SMC implementation of the PHD filter. Instead of clustering the particles in an ad-hoc manner after the update step (which is the current approach), we perform state estimation and, if required, particle clustering, within the update step in an exact and principled manner. Numerical simulations indicate a significant improvement in the estimation accuracy of the proposed SMC-PHD filter.*

Keywords: Tracking, PHD filter, multi-object estimation, sequential Monte Carlo, particle filter.

1 Introduction

Recently Mahler [7] proposed a systematic generalisation of the single-object recursive Bayesian estimator to the multi-object case. In this formulation, the moving objects can appear and disappear over time and the goal of the multi-object Bayes filter is to estimate sequentially both the number of objects and their individual states. Estimation is based on (uncertain) *prior knowledge* (in the form of the measurement model and object dynamic model) and *measurements* corrupted by noise and detection uncertainty.

The Bayes filter propagates the posterior probability density function (pdf) through a two-step procedure: the prediction and update. In the multi-object case, the multi-object posterior pdf is formulated by Mahler using the finite set statistics (FISST), a set of practical mathematical tools from point process theory. The propagation of this multi-object posterior, however, is computationally very intensive due to the high dimensionality of the multi-object state. If the state space of a single object is $\mathcal{X}$, the multi-object posterior pdf is defined on $\mathcal{F}(\mathcal{X})$, the space of finite subsets of $\mathcal{X}$. To overcome the high dimensionality of the multi-object Bayes filter, Mahler introduced the Probability Hypothesis Density (PHD) filter [8], which propagates the first moment of the multi-object posterior known as the intensity function or PHD (instead of propagating the full multi-object posterior). The PHD is defined on a single-object state-space $\mathcal{X}$. The PHD filter has subsequently attracted a great deal of interest, with applications in passive radar [16], sonar [2], computer vision [6,11], SLAM [10], traffic monitoring [1], cell tracking [5], etc. A generalisation of the PHD filter, referred to as the Cardinalised PHD filter, has been proposed by Mahler [9] to improve the estimate of the number of objects. The CPHD filter has also attracted interest with applications to GMTI tracking [17], tracking in aerial videos [12], etc. An alternative derivation of PHD filters (avoiding FISST and process point theory) is presented in [4].

PHD filters are implemented either via the sequential Monte Carlo (SMC) method [20, 23] or using finite Gaussian mixtures (GM) [19, 22]. The GM method is attractive because it provides a closed form algebraic solution to PHD filtering equation, with the state estimate (and it covariance) easily accomplished. However, the GM method is based on somewhat restrictive assumptions that single-object transitional densities and likelihood functions are Gaussian, and that the probability of survival and the probability of detection are constant [7, 19]. The SMC method imposes no such restrictions and should therefore provide a more general framework for PHD filtering, but in reality is affected by different kind of problems [7].

The first problem with the SMC implementation of the PHD filter is related to drawing the samples (particles) from the birth density. If the birth density is non-restrictive, that is the new objects can appear anywhere in the state space $\mathcal{X}$, the birth density has to cover the entire $\mathcal{X}$. For an SMC-PHD filter to work properly in this case, a massive number of "new-born" particles is typically required. The second drawback of the SMC implementation is related to the state estimate and its covariance: (a) the instantaneous estimate of the number of objects (i,e, the cardinality estimate) obtained from the particle weights is very unstable for the PHD filter; (b) the point estimates and their covariances require the particles to be clustered beforehand. The particle clustering is particularly problematic: clustering techniques are either ad-hoc (e.g. k-means) or require the underlying density to be Gaussian (e.g. EM algorithm). This clearly defeats the purpose of the SMC method, that is the propagation of possibly non-Gaussian single-object posterior densities by random samples. Moreover, data clustering algorithms require the number of clusters to be specified as an input parameter. If for this purpose we use the instantaneous cardinality estimate obtained from the particle weights, the clustering output will be further degraded.

In this paper we propose remedies for both of the outlined problems. The presentation will focus on the PHD filter, whereas the Cardinalised PHD results will appear in the future [13]. We point out that the proposed techniques are complementary with recent attempts to improve the efficiency of the SMC-PHD filter by preselecting particles for propagation (the so-called auxiliary particle PHD filter) presented in [23].

2 Background

Suppose that at time k there are n_k object states $\mathbf{x}_{k,1},\ldots,\mathbf{x}_{k,n_k}$, each taking values in a state space $\mathcal{X}\subseteq\mathbb{R}^{n_x}$, and m_k measurements (detections) $\mathbf{z}_{k,1},\ldots,\mathbf{z}_{k,m_k}$, each taking values in the observation space $\mathcal{Z}\subseteq\mathbb{R}^{n_z}$. A multi-object state and a multi-object observation are then represented by the finite sets:

$$\mathbf{X}_k = \{\mathbf{x}_{k,1},\ldots,\mathbf{x}_{k,n_k}\}\in\mathcal{F}(\mathcal{X}), \quad (1)$$

$$\mathbf{Z}_k = \{\mathbf{z}_{k,1},\ldots,\mathbf{z}_{k,m_k}\}\in\mathcal{F}(\mathcal{Z}), \quad (2)$$

respectively. Here $\mathcal{F}(\mathcal{X})$ and $\mathcal{F}(\mathcal{Z})$ are the finite subsets of $\mathcal{X}$ and $\mathcal{Z}$, respectively. At each time step some objects may disappear (die), others may survive and transition into a new state, and new objects may appear. Due to the imperfections in the detector, some of the surviving and newborn objects may not be detected, whereas the observation set $\mathbf{Z}_k$ may include false detections (or clutter). The evolution of the objects and the origin of measurements are unknown. Uncertainty in both multi-object state and multi-object measurement is naturally modelled by random finite sets.

The objective of the recursive multi-object Bayesian estimator [7] is to determine at each time step k the posterior probability density of multi-object state $f_{k|k}(\mathbf{X}_k|\mathbf{Z}_{1:k})$, where $\mathbf{Z}_{1:k}=(\mathbf{Z}_1,\ldots,\mathbf{Z}_k)$ denotes the accumulated observation sets up to time k. The multi-object posterior can be computed sequentially via the prediction and the update steps. Suppose that $f_{k-1|k-1}(\mathbf{X}_{k-1}|\mathbf{Z}_{1:k-1})$ is known and that a new set of measurements $\mathbf{Z}_k$ corresponding to time k has been received. Then the predicted and updated multi-object posterior densities are calculated as follows [7]:

$$f_{k|k-1}(\mathbf{X}_k|\mathbf{Z}_{1:k-1}) = \int \Pi_{k|k-1}(\mathbf{X}_k|\mathbf{X}_{k-1})\, f_{k-1|k-1}(\mathbf{X}_{k-1}|\mathbf{Z}_{1:k-1})\delta\mathbf{X}_{k-1} \quad (3)$$

$$f_{k|k}(\mathbf{X}_k|\mathbf{Z}_{1:k}) = \frac{\vartheta_k(\mathbf{Z}_k|\mathbf{X}_k)f_{k|k-1}(\mathbf{X}_k|\mathbf{Z}_{1:k-1})}{\int\vartheta_k(\mathbf{Z}_k|\mathbf{X})f_{k|k-1}(\mathbf{X}|\mathbf{Z}_{1:k-1})\delta\mathbf{X}}, \quad (4)$$

where $\Pi_{k|k-1}(\mathbf{X}_k|\mathbf{X}_{k-1})$ is a multi-object transition density and $\vartheta_k(\mathbf{Z}_k|\mathbf{X}_k)$ is a multi-object likelihood. We should note that this recursion is a non-trivial generalization, since the transition density needs to consider the uncertainty in target number, which can change over time due to targets entering and leaving the state space, and the multi-object likelihood needs to consider detection uncertainty and false alarms. It is also clear that the integrals in the recursion (3)-(4) are non-standard set integrals.

Since $f_{k|k}(\mathbf{X}_k|\mathbf{Z}_{1:k})$ is defined over $\mathcal{F}(\mathcal{X})$, practical implementation of the multi-object Bayes nonlinear filter is challenging and usually limited to a small number of objects [14,18]. In order to overcome this limitation, Mahler [8] proposed to propagate only the first-order statistical moment of $f_{k|k}(\mathbf{X}|\mathbf{Z}_{1:k})$, the intensity function or PHD $D_{k|k}(\mathbf{x}|\mathbf{Z}_{1:k})=\int\delta_{\mathbf{X}}(\mathbf{x})f_{k|k}(\mathbf{X}|\mathbf{Z}_{1:k})\delta\mathbf{X}$, defined over $\mathcal{X}$. The PHD has the property that

$$\int_{\mathcal{X}} D_{k|k}(\mathbf{x}|\mathbf{Z}_{1:k})d\mathbf{x} = \nu_{k|k}\in\mathbb{R} \quad (5)$$

is the expected number of objects in the state space $\mathcal{X}$.

Using abbreviation $D_{k|k}(\mathbf{x}|\mathbf{Z}_{1:k}) \overset{\text{abbr}}{=} D_{k|k}(\mathbf{x})$, the prediction equation of the PHD filter is given by[1] [8]:

$$D_{k|k-1}(\mathbf{x}) = \gamma_{k|k-1}(\mathbf{x}) + \left\langle\, p_S\, D_{k-1|k-1},\ \pi_{k|k-1}(\mathbf{x}|\cdot)\,\right\rangle \quad (6)$$

where

- $\gamma_{k|k-1}(\mathbf{x})$ is the PHD of object births between time k and $k+1$;
- $p_S(\mathbf{x}') \overset{\text{abbr}}{=} p_{S,k|k-1}(\mathbf{x}')$ is the probability that a target with state $\mathbf{x}'$ at time $k-1$ will survive until time k;
- $\pi_{k|k-1}(\mathbf{x}|\mathbf{x}')$ is the single-object transition density from time $k-1$ to k;
- $\langle g,f\rangle=\int f(\mathbf{x})\,g(\mathbf{x})\,d\mathbf{x}$.

[1] We do not consider object spawning in this paper.

The first term on the RHS of (6) refers to the newborn objects, while the second represents the persistent objects. Upon receiving the measurement set $\mathbf{Z}_k$ at time k, the update step of the PHD filter is computed according to:

$$D_{k|k}(\mathbf{x}) = [1 - p_D(\mathbf{x})]\, D_{k|k-1}(\mathbf{x}) + \sum_{\mathbf{z}\in\mathbf{Z}_k} \frac{p_D(\mathbf{x}) g_k(\mathbf{z}|\mathbf{x})\, D_{k|k-1}(\mathbf{x})}{\kappa_k(\mathbf{z}) + \langle p_D\; g_k(\mathbf{z}|\cdot), D_{k|k-1}\rangle} \quad (7)$$

where

- $p_D(\mathbf{x}) \overset{\text{abbr}}{=} p_{D,k}(\mathbf{x})$ is the probability that an observation will be collected at time k from a target with state $\mathbf{x}$;
- $g_k(\mathbf{z}|\mathbf{x})$ is the single-object measurement likelihood at time k;
- $\kappa_k(\mathbf{z})$ is the PHD of clutter at time k.

The SMC implementation of the PHD filter [7, p16.5.2] is based on the random sample approximation of the PHD:

$$D_{k|k}(\mathbf{x}) \approx \sum_{n=1}^{N} w^n_{k|k}\, \delta_{\mathbf{x}^n_{k|k}}(\mathbf{x}) \quad (8)$$

where $\delta_{\mathbf{y}}(\mathbf{x})$ is the Dirac delta function and $\{(w^n_{k|k}, \mathbf{x}^n_{k|k})\}_{n=1}^N$ is the weighted particle set, consisting of random samples $\mathbf{x}^n_{k|k}$ and their weights $w^n_{k|k}$. According to (5), for large N, the sum of importance weights approximates the expected number of objects, i.e.

$$\sum_{n=1}^{N} w^n_{k|k} \approx \nu_{k|k}. \quad (9)$$

The prediction step forms two types of particles, corresponding to persisting and newborn objects. The particles for newborn objects are drawn from density $\gamma_{k|k-1}(\mathbf{x})/\int \gamma_{k|k-1}(\mathbf{x}) d\mathbf{x}$. The number of newborn particles corresponds to the nearest integer to $\int \gamma_{k|k-1}(\mathbf{x}) d\mathbf{x}$.

3 Birth density driven by measurements

Observe that in the above formulation of the PHD filter, new objects are "born" in the prediction step (6). The intensity function of newborn objects $\gamma_{k|k-1}(\mathbf{x})$ is independent of measurements, and in the general case, where the objects can appear anywhere in the state space, it has to cover the entire $\mathcal{X}$. This is significant for the SMC implementation of the PHD filter, because the newborn object particles need to cover the entire state-space with reasonable density for the SMC-PHD filter to work properly. Clearly this is inefficient and wasteful.

Conventional target tracking algorithms, on the other hand, initialise new tracks based on measurements. Typically, the measurements at each scan are first associated with predicted tracks, and the remaining un-associated measurements are used to initialise new "tentative" or "preliminary" tracks (unreported to the operator).

We propose to use the same approach in the PHD filter, that is to place the newborn particles in the region of the state-space $\mathbf{x} \in \mathcal{X}$ for which the likelihood $g_k(\mathbf{z}|\mathbf{x})$ will have high values. This is not a new idea, it has been proposed in [20]. However, we show in this paper that if the birth density is made dependent on measurements, *the PHD equations must be applied in a different form* (which has not been done in [20]), otherwise the PHD filter cardinality estimates will be biased.

3.1 New formulation of the PHD filter

We start from (6) and (7), where the state vector consists of the usual kinematic/feature component (position, velocity, amplitude, etc) denoted by $\mathbf{y}$ and a mark or a label β which distinguishes a newborn target from the persistent object, i.e. $\mathbf{x} = (\mathbf{y}, \beta)$ where

$$\beta = \begin{cases} 0 & \text{for a persistent object} \\ 1 & \text{for a newborn object} \end{cases} \quad (10)$$

and $\mathbf{y} \in \mathcal{Y}$. The birth PHD is then:

$$\gamma_{k|k-1}(\mathbf{x}) = \gamma_{k|k-1}(\mathbf{y}, \beta) = \begin{cases} \gamma_{k|k-1}(\mathbf{y}), & \beta = 1 \\ 0, & \beta = 0. \end{cases} \quad (11)$$

Note a slight abuse of notation in using the same symbol $\gamma_{k|k-1}$ for both functions of $\mathbf{x}$ and $\mathbf{y}$. Similar abuse will be used throughout this section, but the meaning should be clear from the context.

A newborn object becomes a persisting object at the next time, but a persisting object cannot become a newborn object. Thus the mark β can only change from 1 to 0 but not vice-versa. The transition model is then

$$\begin{aligned} \pi_{k|k-1}(\mathbf{x}|\mathbf{x}') &= \pi_{k|k-1}(\mathbf{y}, \beta|\mathbf{y}', \beta') \\ &= \pi_{k|k-1}(\mathbf{y}|\beta, \mathbf{y}', \beta')\pi_{k|k-1}(\beta|\mathbf{y}', \beta') \\ &= \pi_{k|k-1}(\mathbf{y}|\beta, \mathbf{y}', \beta')\pi_{k|k-1}(\beta|\beta') \\ &= \pi_{k|k-1}(\mathbf{y}|\beta, \mathbf{y}')\pi_{k|k-1}(\beta|\beta') \\ &= \pi_{k|k-1}(\mathbf{y}|\mathbf{y}')\pi_{k|k-1}(\beta|\beta') \end{aligned} \quad (12)$$

with

$$\pi_{k|k-1}(\beta|\beta') = \begin{cases} 0, & \beta = 1 \\ 1, & \beta = 0. \end{cases} \quad (13)$$

The probability of survival does not depend on β and hence

$$p_S(\mathbf{x}) = p_S(\mathbf{y}, \beta) = p_S(\mathbf{y}) \quad (14)$$

The PHD filter prediction equation (6) for augmented state vector is given by:

$$D_{k|k-1}(\mathbf{y},\beta) = \gamma_{k|k-1}(\mathbf{y},\beta)+ \sum_{\beta'=0}^{1} \int D_{k-1|k-1}(\mathbf{y}',\beta')p_S(\mathbf{y}',\beta')\pi_{k|k-1}(\mathbf{y},\beta|\mathbf{y}',\beta')d\mathbf{y}'. \tag{15}$$

Upon substitution of expressions (11)-(14) into (15) we obtain the new form of the PHD filter prediction:

$$D_{k|k-1}(\mathbf{y},\beta) = \begin{cases} \gamma_{k|k-1}(\mathbf{y}), & \beta=1 \\ \Big\langle D_{k-1|k-1}(\cdot,1)+D_{k-1|k-1}(\cdot,0), \\ \qquad p_S\pi_{k|k-1}(\mathbf{y}|\cdot)\Big\rangle & \beta=0 \end{cases} \tag{16}$$

Now we carry out similar manipulations for the update step. First, assume that new targets are always detected, so we can write:

$$p_D(\mathbf{x}) = p_D(\mathbf{y},\beta) = \begin{cases} 1, & \beta=1 \\ p_D(\mathbf{y}), & \beta=0 \end{cases} \tag{17}$$

The measurement does not depend on the mark, hence

$$g_k(\mathbf{z}|\mathbf{x}) = g_k(\mathbf{z}|\mathbf{y},\beta) = g_k(\mathbf{z}|\mathbf{y}) \tag{18}$$

The PHD update equation (7) for augmented state vector is given by:

$$D_{k|k}(\mathbf{y},\beta) = [1-p_D(\mathbf{y},\beta)]D_{k|k-1}(\mathbf{y},\beta)+ \sum_{\mathbf{z}\in\mathbf{Z}_k} \frac{p_D(\mathbf{y},\beta)g_k(\mathbf{z}|\mathbf{y},\beta)D_{k|k-1}(\mathbf{y},\beta)}{\kappa_k(\mathbf{z})+\sum_{\beta=0}^{1}\big\langle p_D(\cdot,\beta)g_k(\mathbf{z}|\cdot,\beta),D_{k|k-1}(\cdot,\beta)\big\rangle} \tag{19}$$

After some manipulations and using (17) and (18), it follows that the update step for persisting objects ($\beta = 0$), is given by:

$$D_{k|k}(\mathbf{y},0) = [1-p_D(\mathbf{y})]D_{k|k-1}(\mathbf{y},0)+ \sum_{\mathbf{z}\in\mathbf{Z}_k} \frac{p_D(\mathbf{y})g_k(\mathbf{z}|\mathbf{y})D_{k|k-1}(\mathbf{y},0)}{\kappa_k(\mathbf{z})+\big\langle g_k(\mathbf{z}|\cdot),\gamma_{k|k-1}\big\rangle+\big\langle p_D g_k(\mathbf{z}|\cdot),D_{k|k-1}(\cdot,0)\big\rangle} \tag{20}$$

while the update step for newborn objects ($\beta = 1$) is given by:

$$D_{k|k}(\mathbf{y},1) = \sum_{\mathbf{z}\in\mathbf{Z}_k} \frac{g_k(\mathbf{z}|\mathbf{y})\,\gamma_{k|k-1}(\mathbf{y})}{\kappa_k(\mathbf{z})+\big\langle g_k(\mathbf{z}|\cdot)\,\gamma_{k|k-1}\big\rangle+\big\langle p_D g_k(\mathbf{z}|\cdot),D_{k|k-1}(\cdot,0)\big\rangle} \tag{21}$$

In summary, the update step is performed differently for newborn and persistent objects. The two types of objects are added and predicted together in the prediction step (16).

3.2 Implementation with an example

The pseudo-code of the proposed SMC-PHD filter is given in Algorithm 1. The newborn object particles are formed using measurement set $\mathbf{Z}_k$ before application of the update step (20)-(21). For each $\mathbf{z}\in\mathbf{Z}_k$ we generate ρ newborn particles $\mathbf{y}^n_{k|k-1,b}$ such that $\mathbf{z}$ can be considered as a random sample from $g_k(\mathbf{z}|\mathbf{y}^n_{k|k-1,b})\gamma_{k|k-1}(\mathbf{y}^n_{k|k-1,b})$. The weights of newborn particles $\mathbf{y}^n_{k|k-1,b}$ before application of the update step are uniform

$$w^n_{k|k-1,b} = \frac{\nu^b_{k|k-1}}{\rho\, m_k} \tag{22}$$

where $m_k = |\mathbf{Z}_k|$ and $\nu^b_{k|k-1} = \sum_{j=1}^{m_k}\sum_{i=1}^{\rho} g_k(\mathbf{z}_{k,j}|\mathbf{y}^{i+(j-1)m_k}_{k|k-1,b})\gamma_{k|k-1}(\mathbf{y}^{i+(j-1)m_k}_{k|k-1,b})$ is a parameter corresponding to the expected number of newborn objects between $k-1$ and k.

Before the application of the update step (20)-(21), the PHD corresponding to persistent objects, $D_{k|k-1}(\mathbf{y},0)$, is approximated by a weighted particle set $\{(w^n_{k|k-1,p},\mathbf{y}^n_{k|k-1,p})\}_{n=1}^{N_{k-1}}$, see lines 7, 8 in Alg. 1. Similarly, the PHD of newborn objects, $D_{k|k-1}(\mathbf{y},1) = \gamma_{k|k-1}(\mathbf{y})$ is approximated by a weighted particle set $\{(w^n_{k|k-1,b},\mathbf{y}^n_{k|k-1,b})\}_{n=1}^{\rho\, m_k}$, see lines 25, 26 in Alg. 1. The weights for persistent and newborn objects are updated according to (20):

$$w^n_{k|k,p} = [1-p_D(\mathbf{y}^n_{k|k-1,p})]w^n_{k|k-1,p}+ \sum_{\mathbf{z}\in\mathbf{Z}_k}\frac{p_D(\mathbf{y}^n_{k|k-1,p})g_k(\mathbf{z}|\mathbf{y}_{k|k-1,p})w^n_{k|k-1,p}}{\mathcal{L}(\mathbf{z})} \tag{23}$$

and (21)

$$w^n_{k|k,b} = \sum_{\mathbf{z}\in\mathbf{Z}_k}\frac{w^n_{k|k-1,b}}{\mathcal{L}(\mathbf{z})} \tag{24}$$

respectively, where

$$\mathcal{L}(\mathbf{z}) = \kappa_k(\mathbf{z})+\sum_{n=1}^{\rho m_k} w^n_{k|k-1,b} + \sum_{n=1}^{N_{k-1}} p_D(\mathbf{y}^n_{k|k-1,p})g_k(\mathbf{z}|\mathbf{y}^n_{k|k-1,p})w^n_{k|k-1,p} \tag{25}$$

The expected number of persistent objects at time k can be approximated then by $\nu^p_k = \sum_{n=1}^{N_{k-1}} w^n_{k|k,p}$, with $w^n_{k|k,p}$ given by (23). The two particle sets (persistent and newborn) are then resampled *separately*, based on their corresponding updated weights in (23) and (24). After resampling, the updated PHD function for persistent objects $D_{k|k}(\mathbf{y},0)$ is approximated by a particle set: $\{(w^n_{k,p},\mathbf{y}^n_{k,p}); n=1,\dots,[\eta\nu^p_k]\}$, where $[\cdot]$ defines the nearest integer function, η is a parameter which defines the number of particles per persisting object, see lines 19-21 in Alg. 1. Similarly, the PHD function for newborn objects $D_{k|k}(\mathbf{y}, 0)$ is approximated by $\{(w^n_{k,b}, \mathbf{y}^n_{k,b}); n = 1, \ldots, \rho m_k\}$, with uniform weights $w^n_{k,b}$, see line 31 in Alg.1.

The PHD filter at time k reports only the PHD of persistent objects (this is similar to tentative tracks not being reported to the operator at the time they were initialised). The two particle sets (corresponding to newborn and persistent objects) are merged only in the prediction step, see line 5 in Alg. 1. This follows from (16), case $\beta = 0$, where two intensity functions are added, before being predicted to the future time. Addition of PHD functions corresponds to the union of weighted particle sets.

Next we illustrate the performance of the proposed SMC-PHD filter with a numerical example. The scenario shown in Fig.1 consist of up to 10 objects all moving towards the centre of the observation area. The objects are observed with a sensor located at

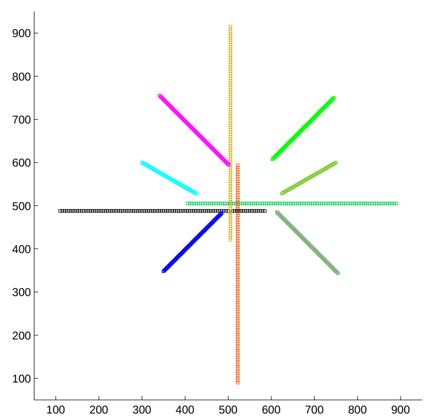

Figure 1: Scenario with up to 10 moving objects

$(-100, 100)$m, which measures range and bearings, with measurement noise standard deviation of 1 degree and 3 meters, respectively. The probability of detection is constant, $p_D(\mathbf{y}) = 0.95$. The clutter (false detections) are generated as a Poisson random finite set with intensity function $\kappa(\mathbf{z}) = \lambda\, c(\mathbf{z})$, where the average number of false detections is $\lambda = 10$, and $c(\mathbf{z})$ is uniform in angle between 0 and $\pi/2$ and uniform in range between 0 and 1600m. The SMC-PHD filter was implemented with the following parameters: $\gamma_{k|k-1}(\mathbf{y}) = const$ (objects are equally likely to appear anywhere in the state space), with $\nu^b_{k|k-1} = 1.0$. The number of newborn object particles per measurement is $\rho = 5$, the number of particles per persisting object is $\eta = 100$. Probability of survival $p_S(\mathbf{y}) = 0.99$ and transition $\pi_{k|k-1}(\mathbf{y}|\mathbf{y}')$ is based on the nearly constant velocity model with a small amount of process noise. Fig.2.(a) shows the true cardinality (blue solid line) and the estimated cardinality of persistent objects $\nu^p_k = \sum_{n=1}^{N_{k-1}} w^n_{k|k,p}$ averaged over 100 Monte Carlo runs (red solid line), with plus/minus one-standard deviation indicated by red dashed lines. The proposed SMC-PHD filter produced unbiased car-

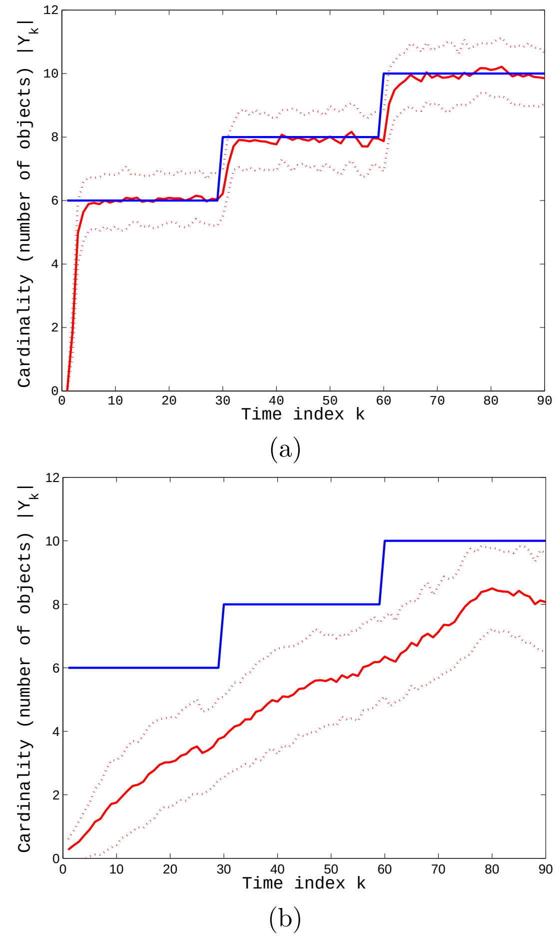

Figure 2: Cardinality over time: the true value is in blue; red is the estimate (200 Monte Carlo runs); (a) the proposed SMC-PHD estimate ($\rho = 5$); (b) the SMC-PHD estimate when new object particles are placed according to $\gamma_{k|k-1}(\mathbf{y})$ ($\rho = 5$)

dinality estimates, with the standard deviation of the cardinality estimate very small (approximately 1) and slightly increasing with cardinality. We point out that if the placement of newborn object particles is driven by measurements, but the PHD filter update is not carried out as described here, the resulting cardinality estimate would be biased, with the value of the bias proportional to parameter $\nu^b_{k|k-1}$. Fig.2.(b) shows the cardinality estimate of the SMC-PHD filter with exactly the same parameters, where new-object particles are placed according to $\gamma_{k|k-1}(\mathbf{y})$ (in this case uniform across the state-space) and using the standard PHD filter equations (6) and (7). Only for a very large number of particles per measurement ρ, this SMC implementation can achieve similar performance as the one we propose.

4 State and error estimation

So far we have been concerned with the cardinality estimate $\hat{n}_k = [\nu^p_k]$ from the output of the SMC-PHD filter. The second problem we address in this paper is how to compute the single-object state estimates, $\hat{\mathbf{y}}_{k,1}, \ldots, \hat{\mathbf{y}}_{k,\hat{n}_k}$, and their covariances, from the SMC-PHD filter. To illustrate the difficulty, Fig.3.(a) shows the particles corresponding to persistent objects (position only) of the proposed SMC-PHD filter at $k = 75$, for the scenario described above. There are ten objects, four are in close proximity to each other. In order to compute the estimates and their covariances, we need first to cluster the particles into $\hat{n}_k$ groups. Note that the clusters of particles corresponding to single objects in general approximate non-Gaussian probability density functions. This means that clustering based on the EM algorithm, for example (which assumes a model, typically Gaussian, of the underlying pdfs) is inappropriate in general. Other clustering techniques are ad-hoc and their use only spoils the elegance of the Bayesian multi-object nonlinear filtering framework. In this section we show that grouping of particles corresponding to individual objects can be done within the update step of the proposed SMC-PHD filter, thus eliminating the need for ad-hoc post-update-step particle clustering.

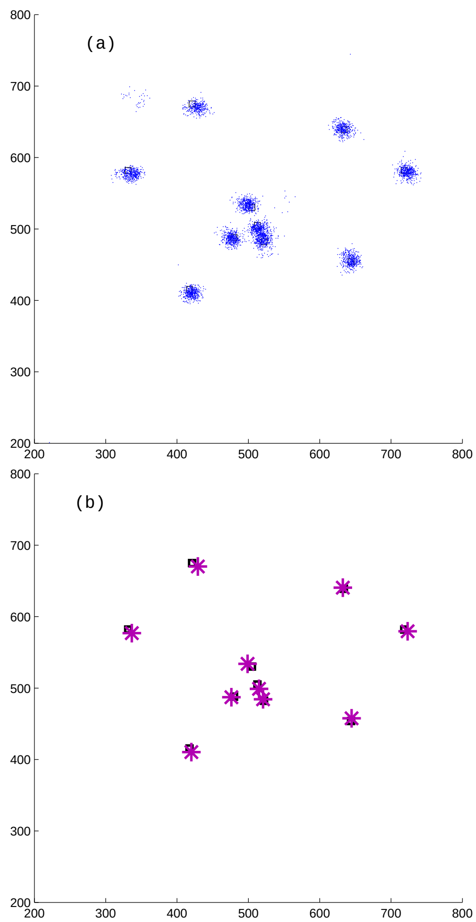

Figure 3: (a) SMC-PHD particles (position only) corresponding to ten persistent objects at $k = 75$; (b) point estimates (denoted by $*$) and true object locations (denoted by $\square$)

4.1 Grouping of particles

A close inspection of (23) reveals the that the update step can be interpreted as a creation of $m_k + 1$ replicas of the particle set $\{\mathbf{y}^n_{k|k-1}; n = 1, \ldots, N_{k-1}\}$, because the summation of $m_k + 1$ terms in (23) is identical to the union of replicated particle sets. Let us denote the weights corresponding to particle set replica $j = 0, 1, \ldots, m_k$ by $\{w^{n,j}_{k|k,p}, n = 1, \ldots, N_{k-1}\}$. Index $j > 0$ corresponds to the index of a measurement at time k, see (2): $\mathbf{Z}_k = \{\mathbf{z}_{k,j}; j = 1, \ldots, m_k\}$. The weights $w^{n,j}_{k|k,p}$, according to (23), can be computed as:

$$w^{n,j}_{k|k,p} = \begin{cases} [1 - p_D(\mathbf{y}^n_{k|k-1,p})] w^n_{k|k-1,p}, & j = 0 \\ \frac{p_D(\mathbf{y}^n_{k|k-1,p}) g_k(\mathbf{z}_{k,j}|\mathbf{y}^n_{k|k-1,p}) w^n_{k|k-1,p}}{\mathcal{L}(\mathbf{z}_{k,j})}, & j = 1, \ldots, m_k \end{cases} \tag{26}$$

with $\mathcal{L}(\mathbf{z}_{k,j})$ specified in (25). Eq.(26) is implemented in lines 10-12 in Alg.1. The sum $W^j_{k,p} = \sum_{n=1}^{N_{k-1}} w^{n,j}_{k|k,p}$ can be interpreted as the total weight assigned to the particle set replica $j = 1, \ldots, m_k$. According to (23)-(25), note that $0 \leq W^j_{k,p} < 1$. Moreover, if $\mathbf{z}_{k,j}$ is a measurement which results in non-zero likelihood for at least some particles (i.e. likely to have originated from an object), $W^j_{k,p}$ will have high value, and if not, it will tend to zero.

For $j = 1, \ldots, m_k$, the state estimate $\hat{\mathbf{y}}_{k,j}$ and its covariance matrix $\mathbf{P}_{k,j}$ are computed as:

$$\hat{\mathbf{y}}_{k,j} = \sum_{n=1}^{N_{k-1}} w^{n,j}_{k|k,p} \mathbf{y}^n_{k|k-1,p} \tag{27}$$

$$\mathbf{P}_{k,j} = \sum_{n=1}^{N_{k-1}} w^{n,j}_{k|k,p} \left(\mathbf{y}^n_{k|k-1,p} - \hat{\mathbf{y}}_{k,j}\right) \left(\mathbf{y}^n_{k|k-1,p} - \hat{\mathbf{y}}_{k,j}\right)^\intercal \tag{28}$$

Note that the SMC-PHD filter assigns the total weight $W^j_{k,p}$ to each pair $(\hat{\mathbf{y}}_{k,j}, \mathbf{P}_{k,j})$. Only if the value of $W^j_{k,p}$, $j = 1, \ldots, m_k$ is above a certain threshold value, the state estimate $\hat{\mathbf{y}}_{k,j}$ and its covariance $\mathbf{P}_{k,j}$ will be reported (lines 14-18 in Alg.1.). Fig.3.(b) illustrates the state estimates obtained using the described method. The weights $W^j_{k,p}$ for all ten estimated objects were above 0.8 in this example.

If we want to find not only the state estimate and its covariance, but also the particle approximation of the posterior pdf for a particular object, we then need to resample the particles $\{\mathbf{y}^n_{k|k-1,p}; n = 1, \ldots, N_{k-1}\}$ using the weights $\{w^{n,j}_{k|k,p}, n = 1, \ldots, N_{k-1}\}$ computed for measurement $\mathbf{z}_{k,j}$. This particle approximation of the posterior pdf is important in the context of nonlinear/non-Gaussian estimation. Note that the output of the proposed SMC-PHD filter can effectively be presented in the same form as the output of the MeMBer filter [21]: for each object it provides the weight $W^j_{k,p}$, which corresponds to the probability of existence, and the particle approximation of its posterior pdf.

Remark. Only objects that have been detected in scan k have a chance to be reported according to the proposed scheme. This is in agreement with the lack of "memory" or inertia of the PHD filter, a characteristic discussed in [7, Sec.16.6], [3].

4.2 Numerical comparisons

The performance of the proposed method for state estimation is compared to the post-update step clustering method via k-means algorithm. In the latter, in order to improve its performance, the k-means algorithm was repeated 5 times (in each run) using different initialisation. The performance of the two methods is measured using the optimal sub-pattern assignment (OSPA) distance [15] between the estimated and the true multi-target states, $\hat{\mathbf{Y}}_k$ and $\mathbf{Y}_k$. We use the OSPA metric because it jointly captures the differences in both the cardinality and the individual elements between two finite sets, in a mathematically consistent yet intuitively meaningful way. The parameters of OSPA metric are the *order* p (which determines the sensitivity to outliers) and the *cut-off* c (the relative weighting of the penalties assigned to cardinality and localization errors). We adopt $p = 2$ and $c = 150$m in this comparison.

Fig.4 shows the OSPA metric for the two state estimation methods of SMC-PHD filters, averaged over 100 Monte Carlo runs. The scenario and the parameters were the same as in Sec.3.2. The results indicate a dramatic improvement in the accuracy of the proposed method. This improvement is mainly due to the reduction in the localisation error, while the cardinality errors are similar. In addition, the proposed method is much faster to run.

5 Summary

The paper presented two important improvements in the sequential Monte Carlo implementation of the PHD filter. The first improvement enabled use of a significantly reduced number of particles for newborn objects, without any loss in accuracy. The second improvement dealt with reporting the state estimates, their covariances or individual object posterior probability density functions. The resulting SMC-PHD filter is significantly more efficient and accurate than the existing SMC-PHD filters. We emphasize that both proposed techniques are complementary to the auxiliary particle PHD filter of [23]. Future work will focus on improving the SMC implementation of the Cardinalised PHD filter using similar ideas.

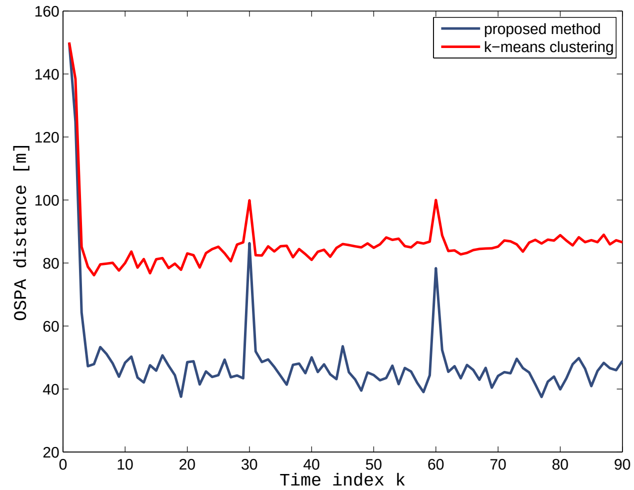

Figure 4: OSPA distance (averaged over 100 Monte Carlo runs) for two methods of state estimation:

References

[1] G. Battistelli, L. Chisci, S. Morrocchi, F. Papi, A. Benavoli, A. Di Lallo, A. Farina, and A. Graziano. Traffic intensity estimation via PHD filtering. In *Proc. 5th European Radar Conf.*, pages 340–343, Amsterdam, The Netherlands, Oct. 2008.

[2] D. Clark, I. T. Ruiz, Y. Petillot, and J. Bell. Particle PHD filter multiple target tracking in sonar image. *IEEE Trans. Aerospace & Electronic Systems*, 43(1):409–416, 2007.

[3] O. Erdinc, P. Willett, and Y. Bar-Shalom. Probability hypothesis density filter for multitarget multisensor tracking. In *Proc. Int. Conf. Information Fusion*, pages 146–153, Philadelphia, PA, USA, July 2005.

[4] O. Erdinc, P. Willett, and Y. Bar-Shalom. The bin-occupancy filter and its connection to the PHD filters. *IEEE Trans. Signal Processing*, 57(11):4232–4246, 2009.

[5] R. R. Juang, A. Levchenko, and P. Burlina. Tracking cell motion using GM-PHD. In *Int. Symp. Biomedical Imaging*, pages 1154–1157, June/July 2009.

[6] E. Maggio, M. Taj, and A. Cavallaro. Efficient multitarget visual tracking using random finite sets. *IEEE Trans. Circuits & Systems for Video Technology*, 18(8):1016–1027, 2008.

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

[8] R. P. S. Mahler. Multi-target bayes filtering via first-order multi-target moments. *IEEE Trans. Aerospace & Electronic Systems*, 39(4):1152–1178, 2003.

[9] R. P. S. Mahler. PHD filters of higher order in target number. *IEEE Trans. Aerospace & Electronic Systems*, 43(4):1523–1543, 2007.

[10] J. Mullane, Ba-Ngu Vo, M. D. Adams, and W. S. Wijesoma. A random set formulation for Bayesian SLAM. In *IEEE/RSJ Int. Confer. Intelligent Robots and Systems*, Nice, France, Sept. 2008.

Algorithm 1 Pseudo-code of the SMC-PHD filter (processing steps at time k)

Input:
  a) Particle set (persistent objects): $\{(w^n_{k-1,p}, \mathbf{y}^n_{k-1,p})\}_{n=1}^{L_{k-1}}$
  b) Particle set (newborn objects): $\{(w^n_{k-1,b}, \mathbf{y}^n_{k-1,b})\}_{n=1}^{S_{k-1}}$
  c) Measurement set at k: $\mathbf{Z} = \{\mathbf{z}_{k,1}, \ldots, \mathbf{z}_{k,m_k}\}$;
$\{(w^n_{k-1}, \mathbf{y}^n_{k-1})\}_{n=1}^{N_{k-1}} = \{(w^n_{k-1,p}, \mathbf{y}^n_{k-1,p})\}_{n=1}^{L_{k-1}} \bigcup \{(w^n_{k-1,b}, \mathbf{y}^n_{k-1,b})\}_{n=1}^{S_{k-1}}$ ▷ Union of input particle sets
for $n = 1, \ldots, N_{k-1}$ **do** ▷ Random sample approximation of $D_{k|k-1}(\mathbf{y}, 0)$
  Draw $\mathbf{y}_{k|k-1,p} \sim \pi_{k|k-1}(\mathbf{y}|\mathbf{y}^n_{k-1})$
  Compute weight $w^n_{k|k-1,p} = p_S(\mathbf{y}^n_{k-1}) w^n_{k-1}$
  Compute weight $w^n_{k|k,p}$ according to (23)
  for $j = 1, \ldots, m_k$ **do**
    Compute weights $w^{n,j}_{k|k,p}$ according to (26)
  end for
end for
for $j = 1, \ldots, m_k$ **do** ▷ State end error estimation
  if $W^j_{k,p} = \left(\sum_{n=1}^{N_{k-1}} w^{n,j}_{k|k,p}\right) > \tau$ **then**, ▷ Reporting threshold $\tau \in (0, 1)$
    Compute $\hat{\mathbf{y}}_{k,j}, \mathbf{P}_{k,j}$ according to (27),(28); report $\left(W^j_{k,p}, \hat{\mathbf{y}}_{k,j}, \mathbf{P}_{k,j}\right)$
  end if
end for
$\nu^p_k = \sum_{n=1}^{N_{k-1}} w^n_{k|k,p}$; $\hat{n}_k = [\nu^p_k]$ ▷ Estimated cardinality; $[\cdot]$ is "nearest integer"
$L_k = \hat{n}_k \cdot \eta$ ▷ parameter η: number of particles per persistent object
Resample L_k times from $\{(w^n_{k|k,p}/\nu^p_k, \mathbf{y}^n_{k|k-1,p})\}_{n=1}^{N_{k-1}}$ to obtain $\{(w^n_{k,p}, \mathbf{y}^n_{k,p})\}_{n=1}^{L_k}$, with $w^n_{k,p} = \nu^p_k / L_p$.
for $j = 1, \ldots, m_k$ **do** ▷ Random sample approximation of $D_{k|k-1}(\mathbf{y}, 1)$
  for $\ell = 1, \ldots, \rho$ **do** ▷ parameter ρ: number of particles per newborn object
    $n = \ell + (j-1) m_k$
    Draw $\mathbf{y}^n_{k|k-1,b} \sim b(\mathbf{y}|\mathbf{z}_{k,j})$
    Compute weight $w^n_{k|k-1,b} = \nu^b_{k|k-1} / (\rho\, m_k)$ ▷ parameter $\nu^b_{k|k-1}$: expected num. newborn objects
    Compute weight $w^n_{k|k,b}$ according to (24)
  end for
end for
$S_k = \rho \cdot m_k$
Resample S_k times from $\{(w^n_{k|k,b} / \sum_{n=1}^{S_k} w^n_{k|k,b}, \mathbf{y}^n_{k|k-1,b})\}_{n=1}^{S_k}$ to obtain $\{(w^n_{k,b}, \mathbf{y}^n_{k,b})\}_{n=1}^{S_k}$.

[11] Nam Trung Pham, Weimin Huang, and S. H.; Ong. Tracking multiple objects using Probability Hypothesis Density filter and color measurements. In *Proc. IEEE Int. Conf. Multimedia and Expo*, pages 1511 – 1514, July 2007.

[12] E. Pollard, A. Plyer, B. Pannetier, F. Champagnat, and G. Le Besnerais. GM-PHD filters for multi-object tracking in uncalibrated aerial videos. *Proc. 12th Int. Conf. Information Fusion*, pages 1171–1178, July 2009.

[13] B. Ristic, D. Clark, B.-N. Vo, and B.-T. Vo. Improved SMC implementation of PHD and CPHD filters. In preparation.

[14] B. Ristic and B.-N. Vo. Sensor management for SMC implementation of multi-object filtering with random finite sets. In *Proc. COGIS'09*, Paris, France, Nov. 2009. SEE & IET.

[15] D. Schuhmacher, B.-T. Vo, and B.-N. Vo. A consistent metric for performance evaluation of multi-object filters. *IEEE Trans. Signal Processing*, 56(8):3447–3457, Aug. 2008.

[16] M. Tobias and A.D. Lanterman. Probability hypothesis density-based multitarget tracking with bistatic range and Doppler observations. *IEE Proc.-Radar Sonar Navig*, 152(3):195–205, 2005.

[17] M. Ulmke, O. Erdinc, and Peter Willett. GMTI tracking via the Gaussian mixture cardinalized probability hypothesis density filter. *IEEE Trans. Aerospace & Electronic Systems*, 2009. To appear.

[18] M. Vihola. Rao-Blackwellised particle filtering in random set multitarget tracking. *IEEE Trans. Aerospace & Electronic Systems*, 43(2):689–705, 2007.

[19] B.-N. Vo and W.-K. Ma. The Gaussian mixture probability hypothesis density filter. *IEEE Trans. Signal Processing*, 54(11):4091–4104, Nov. 2006.

[20] B. N. Vo, S. Singh, and A. Doucet. Sequential Monte Carlo methods for multi-target filtering with random finite sets. *IEEE Trans. Aerospace & Electronic Systems*, 41(4):1224–1245, Oct. 2005.

[21] B.-T. Vo, B.N. Vo, and A. Cantoni. The cardinality balanced multi-target multi-Bernoulli filter and its implementations. *IEEE Trans. Signal Processing*, 57(2):409–423, 2009.

[22] Ba-Tuong Vo, Ba-Ngu Vo, and A. Cantoni. Analytic implementations of the cardinalized probability hypothesis density filter. *IEEE Trans. Signal Processing*, 55(7):3553–3567, 2007.

[23] N. P. Whiteley, S. S. Singh, and S. J. Godsill. Auxiliary particle implementation of the probability hypothesis density filter. *IEEE Trans. on Aerospace & Electronic Systems*, 2009. To appear.