

Auxiliary Particle Implementation of the Probability Hypothesis Density Filter

Nick Whiteley Sumeetpal Singh Simon Godsill
Signal Processing Laboratory, Department of Engineering, University of Cambridge
Trumpington Street, Cambridge, UK

Abstract

Optimal Bayesian multi-target filtering is, in general, computationally impractical due to the high dimensionality of the multi-target state. Recently Mahler, [9], introduced a filter which propagates the first moment of the multi-target posterior distribution, which he called the Probability Hypothesis Density (PHD) filter. While this reduces the dimensionality of the problem, the PHD filter still involves intractable integrals in many cases of interest. Several authors have proposed Sequential Monte Carlo (SMC) implementations of the PHD filter. However, these implementations are the equivalent of the Bootstrap Particle Filter, and the latter is well known to be inefficient. Drawing on ideas from the Auxiliary Particle Filter of Pitt and Shephard [10], we present a SMC implementation of the PHD filter which employs auxiliary variables to enhance its efficiency. Numerical examples are also presented.

1. Introduction

Multi-target filtering is a dynamic state estimation problem in which both the number of hidden targets and the locations of the targets are unknown. Additionally, the targets appear and terminate at random times. The modelling of multi-target dynamics in this manner naturally incorporates track initiation and termination, a procedure that has mostly been performed separately in traditional tracking algorithms.

As in the single-target case, optimal multi-target filtering involves the propagation of the posterior distribution through Bayes' law. Exact optimal multi-target filtering is impossible in many cases of interest due to the presence of intractable integrals in the filtering recursion. The application of numerical methods (Monte Carlo or otherwise) to approximate the optimal filter is extremely computationally intensive due to the high dimensionality of the multi-target state.

Consider the state space of a single target, $E \subset \mathbb{R}^d$. Each point in this space may specify, for example, the position and velocity of the target. Multi-target filtering involves computation of a distribution on the number of targets and each of their locations in E . The multi-target posterior is therefore a probability distribution on the disjoint

union $\biguplus_{k \geq 0} E^k$.

A more tractable alternative to the optimal multi-target filter is the Probability Hypothesis Density filter of Mahler, [9], which propagates the first moment of the multi-target posterior, known as the intensity function or PHD.

While multi-target tracking is a mature field, the Point Process formalism behind the derivation of the PHD filter is new to the area. The PHD filter has recently been the focus of much interest due to its favourable performance in multi-target tracking as compared with traditional approaches. In this paper we present a new SMC implementation of the PHD filter which is much more efficient than the algorithms proposed in the literature thus far [13], [11], [15]. We refer the reader to [1] for other approaches to multi-target tracking not based on the Point Process formalism.

2. The PHD Filter

The intensity function, $\alpha : E \rightarrow \mathbb{R}_+$, of the multi-target posterior is useful because it yields the expected number of targets in any region of the state space:

$$\mathbb{E}[N(A)] = \int_A \alpha(x) dx, \quad A \in \mathcal{B}(E).$$

where $N(A)$ is the number of targets in the set A (for theoretical details, see [3]). Peaks in the intensity function can be used to estimate target locations and the total mass of the intensity function provides an estimate of the total number of targets. A filtering scheme which propagates only this intensity function, as opposed to the full posterior, is attractive as the dimensionality of the problem is effectively reduced to the dimensionality of E .

The PHD filter consists of a prediction and update operation which propagates the intensity function of the multi-state posterior recursively in time [9]:

$$\alpha_n(x_n) = \int_E f(x_n|x_{n-1})p_S(x_{n-1})\hat{\alpha}_{n-1}(x_{n-1})dx_{n-1} + \gamma(x_n), \quad (1)$$

$$\hat{\alpha}_n(x_n) = \left[1 - p_D(x_n) + \sum_{m=1}^{M_n} \frac{\psi_{n,m}(x_n)}{\langle \psi_{n,m}, \alpha_n \rangle + \kappa(z_{n,m})} \right] \alpha_n(x_n), \quad (2)$$

where

$$\psi_{n,m}(x) = p_D(x)g(z_{n,m}|x),$$

$$\langle \psi_{n,m}, \alpha \rangle = \int_E \psi_{n,m}(x)\alpha(x)dx.$$

In this notation, α_n and $\hat{\alpha}_n$ are respectively the predicted and updated intensities on the state space E , $p_S(x)$ is the probability that a target at x survives, $f(x_n|x_{n-1})$ is the transition density of a single target (it is assumed that all targets follow the same transition model), γ is the birth intensity, $p_D(x)$ is the probability that a target at x is detected, $g(z_{n,m}|x)$ is the likelihood for the m th observation at time n , M_n is the total number of observations at time n and κ is the clutter intensity.

Whilst the PHD filter reduces the dimensionality of the problem, the PHD recursion still involves intractable integrals in many cases of interest, the exception being the ‘linear-Gaussian’ case, [12]. In general Monte Carlo methods are called for.

3. Sequential Monte Carlo

Sequential Monte Carlo methods have become a standard tool for computation in non-linear optimal filtering problems, [5], and in this context have been termed *particle filters*. These algorithms recursively propagate a set of weighted random samples, termed *particles*, which are used to estimate integrals of interest. A typical SMC algorithm recursively proposes and weights samples, with occasional resampling from the discrete distribution defined by the particle weights.

In order for a SMC scheme to be efficient, it is important to ensure that the variance of the weights is minimised. Weights with high variance will rapidly degenerate over time, concentrating all mass on a single particle, yielding poor estimates. Under these circumstances, resampling must be performed frequently, which further increases the variance of estimates, locally in time. Therefore an important factor in the practical efficiency of SMC methods is the mechanism by which particles are proposed. If degeneracy of the weights is to be avoided, this mechanism should take into account information from the observations and drive particles into regions of high probability.

The Auxiliary Particle Filter (APF) of Pitt and Shephard, [10], involves the selection of particles for propagation by drawing particle indices (the auxiliary random variables) from a discrete distribution defined in terms of an additional set of particle weights. These weights are defined to reflect

each particle’s ability to explain the observation at the next time step. Intuitively, this scheme can work well because it selects particles which are likely to end up in regions of high probability.

4. Particle PHD Filter

A particle implementation of the PHD filter in its full generality was proposed in [13], around the same time as two other independent works [11] and [15]. In [11], only the special case without clutter was considered. On the other hand, [15] describes an implementation for the special case with neither birth nor spawning. The theme common to these approaches is the propagation of a particle approximation to the intensity function through the PHD recursion (1) and (2).

One iteration of existing particle PHD filters is outlined as follows. Samples are drawn from a proposal distribution, conditionally upon the previous particle set, and weighted in accordance with the prediction operation. Supplementary particles dedicated to the birth term are then added. This yields an approximation to the predicted intensity α_n , which is itself used in its entirety to approximate the integrals of the form $\langle \psi_{n,m}, \alpha_n \rangle$ in the denominator of (2). The particles are then re-weighted according to the update operator and resampling is performed.

In this framework it is not obvious how to choose the proposal distribution in order to minimise the variance of the weights. In practice the prior distribution is chosen as the proposal, which is sub-optimal, and this is the analogue of the ‘bootstrap’ particle filter, which is well known to be inefficient.

Convergence results establishing the theoretical validity of the particle PHD filter have been obtained in [14], [7] and [2].

5. Auxiliary Particle PHD Filter

5.1 Outline of the Approach

The basic idea of our proposed auxiliary particle PHD filter, drawing on ideas from [10], is to redefine the sampling problem on a higher dimensional space by introducing auxiliary random variables which index particles and observations. In this framework the construction of proposals tailored to observations is straightforward. This is in contrast to existing particle PHD filters, for which it is not obvious how to construct such proposals.

In order to ease exposition, suppose that, at time $n - 1$ we have available a particle set $\{x_{n-1}^{(i)}, w_{n-1}^{(i)}\}_{i=1:N}$ which approximates $\hat{\alpha}_{n-1}$, in the sense that for some test function φ :

$$\sum_{i=1}^N \varphi(x_{n-1}^{(i)})w_{n-1}^{(i)} \xrightarrow[N \rightarrow \infty]{a.s.} \int \varphi(x)\hat{\alpha}_{n-1}(x)dx. \quad (3)$$

As a notational device, we will introduce an additional particle, $(x_{n-1}^{(N+1)}, w_{n-1}^{(N+1)})$, with the conventions that:

$$\begin{aligned} w_{n-1}^{(N+1)} &= \int \gamma(x) dx, \\ f(x|x_{n-1}^{(N+1)}) &= \frac{\gamma(x)}{\int \gamma(x) dx}, \\ p_S(x_{n-1}^{(N+1)}) &= 1. \end{aligned}$$

The value of $x_{n-1}^{(N+1)}$ can be assigned to any point in the state space as it is irrelevant. This ‘source’ particle does *not* enter into estimation of an integral of the form (3) at time $n-1$, but will merely act symbolically as a source of intensity at time n , unifying notation when expressing the birth term in the prediction operation (1). We augment the particle set with this $N+1$ th particle to yield $\{x_{n-1}^{(i)}, w_{n-1}^{(i)}\}_{i=1:N+1}$.

Furthermore, we introduce a null M_n+1 th observation index to unify notation when dealing with the missed-detection term in (2), and define $\mathcal{M}_n = \{1, 2, \dots, M_n+1\}$.

Let $\eta_n : E \times \{1, 2, \dots, N+1\} \times \mathcal{M}_n \rightarrow \mathbb{R}_+$, be defined as follows:

$$\eta_n(x, j, m) = G_{n,m}(x) f(x|x_{n-1}^{(j)}) p_S(x_{n-1}^{(j)}) w_{n-1}^{(j)},$$

where for $m \in \{1, 2, \dots, M_n\}$,

$$G_{n,m}(x) = \frac{\psi_{n,m}(x)}{S_{n,m}^N + \kappa(z_{n,m})},$$

and for $m = M_n+1$,

$$G_{n,m}(x) = 1 - p_D(x).$$

$S_{n,m}^N$ is a consistent estimate of $\langle \psi_{n,m}, \alpha_n \rangle$, discussed in further detail below. Then from (2), we have, for some test function φ ,

$$\begin{aligned} & \sum_{j=1}^{N+1} \sum_{m=1}^{M_n+1} \int \varphi(x) \eta_n(x, j, m) dx \\ &= \sum_{j=1}^{N+1} \sum_{m=1}^{M_n+1} \int \varphi(x) G_{n,m}(x) f(x|x_{n-1}^{(j)}) p_S(x_{n-1}^{(j)}) w_{n-1}^{(j)} dx \\ & \quad \xrightarrow[N \rightarrow \infty]{a.s.} \int \varphi(x) \hat{\alpha}_n(x) dx. \end{aligned} \quad (4)$$

Thus, asymptotically in the number of particles at the previous time step, the marginal of η_n is the intensity function of interest, $\hat{\alpha}_n$. The approach of the auxiliary particle PHD filter is to approximate the integral in (4) using importance sampling. The advantage of this approach, as further described below, is that it naturally accommodates an efficient proposal mechanism.

The method consists of drawing samples from a proposal distribution $q_n(x, j, m)$ defined on $E \times \{1, 2, \dots, N+1\} \times \mathcal{M}_n$, whose support includes that of $\eta_n(x, j, m)$, and weighting them accordingly, having calculated each $S_{n,m}^N$ as a local Monte Carlo integral.

As recommended in [13], the number of particles should be adjusted at each time step to reflect the estimated number of targets. Such adaptation can be straightforwardly accommodated in the auxiliary particle scheme presented here. However, for ease of presentation and without loss of generality we describe the algorithm for fixed number, N , particles.

5.2 Choice of Proposal Distribution

We can design a proposal distribution by first noting that we are able to factorise $\eta_n(x, j, m) = \eta_n(x|j, m) \eta_n(j|m) \eta_n(m)$, where for $m \in \{1, 2, \dots, M_n\}$,

$$\begin{aligned} \eta_n(x|j, m) &= \frac{\psi_{n,m}(x) f(x|x_{n-1}^{(j)}) p_S(x_{n-1}^{(j)}) w_{n-1}^{(j)}}{\langle \psi_{n,m}, f(\cdot|x_{n-1}^{(j)}) p_S(x_{n-1}^{(j)}) w_{n-1}^{(j)} \rangle}, \\ \eta_n(j|m) &= \frac{\langle \psi_{n,m}, f(\cdot|x_{n-1}^{(j)}) p_S(x_{n-1}^{(j)}) w_{n-1}^{(j)} \rangle}{\sum_{j=1}^{N+1} \langle \psi_{n,m}, f(\cdot|x_{n-1}^{(j)}) p_S(x_{n-1}^{(j)}) w_{n-1}^{(j)} \rangle}, \\ \eta_n(m) &= \frac{\sum_{j=1}^{N+1} \langle \psi_{n,m}, f(\cdot|x_{n-1}^{(j)}) p_S(x_{n-1}^{(j)}) w_{n-1}^{(j)} \rangle}{S_{n,m}^N + \kappa(z_{n,m})}. \end{aligned} \quad (5)$$

Similarly, for $m = M_n+1$,

$$\begin{aligned} \eta_n(x|j, m) &= \frac{[1 - p_D(x)] f(x|x_{n-1}^{(j)}) p_S(x_{n-1}^{(j)}) w_{n-1}^{(j)}}{\langle [1 - p_D(\cdot)], f(\cdot|x_{n-1}^{(j)}) p_S(x_{n-1}^{(j)}) w_{n-1}^{(j)} \rangle}, \\ \eta_n(j|m) &= \frac{\langle [1 - p_D(\cdot)], f(\cdot|x_{n-1}^{(j)}) p_S(x_{n-1}^{(j)}) w_{n-1}^{(j)} \rangle}{\sum_{j=1}^{N+1} \langle [1 - p_D(\cdot)], f(\cdot|x_{n-1}^{(j)}) p_S(x_{n-1}^{(j)}) w_{n-1}^{(j)} \rangle}, \\ \eta_n(m) &= \sum_{j=1}^{N+1} \langle [1 - p_D(\cdot)], f(\cdot|x_{n-1}^{(j)}) p_S(x_{n-1}^{(j)}) w_{n-1}^{(j)} \rangle. \end{aligned} \quad (6)$$

Note $\eta_n(m)$ is not normalised and this is not the only possible decomposition.

A natural choice of proposal distribution factorises in the same manner, i.e. $q_n(x, j, m) = q_n(x|j, m) q_n(j|m) q_n(m)$, with each term being a normalised approximation to the corresponding term in (5) and (6). The act of sampling from $q_n(m)$ essentially allocates the number of particles which are going to be used to explore each observation, and sampling from $q_n(j|m)$ selects which particles to propagate for each observation.

Intuitively, this scheme can be efficient for two reasons. Firstly, through careful choice of $q_n(m)$, it allows us to concentrate effort on those observations which are likely to originate from true targets. Secondly, through careful choice of $q_n(j|m)$ and $q_n(x|j, m)$, this approach allows the automatic selection and proposal of particles which are best suited to each of these observations, as in the standard auxiliary particle filter [10].

A method which was proposed in [10] can be used here to define the discrete distributions $q_n(m)$ and $q_n(j|m)$. The idea is to approximate (5) and (6) by approximating each of

the integrals therein using either the mode of $f(x|x_{n-1}^{(j)})$ or its mean. For $m \in \{1, 2, \dots, M_n\}$,

$$q_n(m) \propto \frac{\sum_{j=1}^{N+1} \psi_{n,m}(\hat{x}_n^{(j)}) p_S(x_{n-1}^{(j)}) w_{n-1}^{(j)}}{\sum_{j=1}^{N+1} \psi_{n,m}(\hat{x}_n^{(j)}) p_S(x_{n-1}^{(j)}) w_{n-1}^{(j)} + \kappa(z_m)},$$

$$q_n(j|m) = \frac{\psi_{n,m}(\hat{x}_n^{(j)}) p_S(x_{n-1}^{(j)}) w_{n-1}^{(j)}}{\sum_{j=1}^{N+1} \psi_{n,m}(\hat{x}_n^{(j)}) p_S(x_{n-1}^{(j)}) w_{n-1}^{(j)}}.$$

and for $m = M_n + 1$,

$$q_n(m) \propto \frac{\sum_{j=1}^{N+1} [1 - p_d(\hat{x}_n^{(j)})] p_S(x_{n-1}^{(j)}) w_{n-1}^{(j)}}{\sum_{j=1}^{N+1} p_S(x_{n-1}^{(j)}) w_{n-1}^{(j)}},$$

$$q_n(j|m) = \frac{[1 - p_d(\hat{x}_n^{(j)})] p_S(x_{n-1}^{(j)}) w_{n-1}^{(j)}}{\sum_{j=1}^{N+1} [1 - p_d(\hat{x}_n^{(j)})] p_S(x_{n-1}^{(j)}) w_{n-1}^{(j)}},$$

where,

$$\sum_{m=1}^{M_n+1} q_n(m) = 1,$$

and

$$\hat{x}_n^{(i)} = \int x f(x|x_{n-1}^{(i)}) dx \quad \text{or} \quad \underset{x}{\operatorname{argmax}} f(x|x_{n-1}^{(i)}).$$

Variants of the proposed scheme involve assigning a fixed or minimum number of particles to each observation, but we do not discuss these approaches further here. For $j = N+1$, the above scheme will only be of use when $\gamma(x)$ is unimodal and localised. Various methods, as in the standard particle filter, can be used to construct an efficient $q_n(x|j, m)$, see [6] for some examples.

As in existing particle PHD filters, in order to compute the particle weights, it is necessary to obtain Monte Carlo approximations to the integrals of the form $\langle \psi_{n,m}, \alpha_n \rangle$. Let $\{x_n^{(i)}, j_n^{(i)}, m_n^{(i)}\}_{i=1:N}$ be N samples from $q_n(x, j, m)$. One possibility, applicable when targets are well separated, is to compute each integral using the N_m particles assigned to the corresponding observation. In this case, we have:

$$S_{n,m}^{N_m} = \frac{1}{N_m} \sum_{i \in \Upsilon_{n,m}} \left[\psi_{n,m}(x_n^{(i)}) \times \frac{f(x_n^{(i)} | x_{n-1}^{(j_n^{(i)})}) p_S(x_{n-1}^{(j_n^{(i)})}) w_{n-1}^{(j_n^{(i)})}}{q_n(x_n^{(i)} | j_n^{(i)}, m) q_n(j_n^{(i)} | m)} \right],$$

where $\Upsilon_{n,u} = \{i \in \{1, 2, \dots, N\} : m_n^{(i)} = u\}$ and $N_m = \#\Upsilon_{n,m}$. Note that under the auxiliary particle PHD filtering scheme, in evaluating the importance weights we need not compute $S_{n,m}^{N_m}$ for any m such that $N_m = 0$.

The algorithm for the auxiliary particle PHD filter is given below, with k_0 being the expected initial number of targets. We advocate the use of low variance sampling methods when drawing from the discrete distributions $q_n(j|m)$ and $q_n(m)$, for example as in the residual and stratified resampling schemes, see [4] for comparisons.

Algorithm 1 Auxiliary SMC PHD Filter

```

1:  $n = 0$ 
2: for  $i = 1$  to  $N$  do
3:    $x_1^{(i)} \sim q_0(x)$ 
4:    $w_1^{(i)} = \frac{1}{N} k_0$ 
5: end for
6:  $n \leftarrow n + 1$ 
7: for  $i = 1$  to  $N$  do
8:    $m_n^{(i)} \sim q_n(m)$ 
9:    $j_n^{(i)} \sim q_n(j | m_n^{(i)})$ 
10:   $X_n^{(i)} \sim q_n(x | j_n^{(i)}, m_n^{(i)})$ 
11: end for
12: for  $m = 1$  to  $M_n$  do
13:   if  $N_m > 0$  then
14:     $S_{n,m}^{N_m} = \frac{1}{N_m} \sum_{i \in \Upsilon_{n,m}} \left[ \psi_{n,m}(x_n^{(i)}) \right. \\ \left. \times \frac{f(x_n^{(i)} | x_{n-1}^{(j_n^{(i)})}) p_S(x_{n-1}^{(j_n^{(i)})}) w_{n-1}^{(j_n^{(i)})}}{q_n(x_n^{(i)} | j_n^{(i)}, m) q_n(j_n^{(i)} | m)} \right]$ 
15:   end if
16: end for
17: for  $i = 1$  to  $N$  do
18:    $w_n^{(i)} = \frac{1}{N} \frac{\eta_n(x_n^{(i)}, j_n^{(i)}, m_n^{(i)})}{q_n(x_n^{(i)} | j_n^{(i)}, m_n^{(i)})}$ 
19: end for

```

6. Results

We present simulation results to demonstrate the improvement in efficiency over the bootstrap particle PHD filter which is possible under the proposed scheme.

Consider a constant velocity tracking model for a vehicle whose position is specified in two dimensions, restricted to the window $[0, 100] \times [0, 100]$. The state of a single target is specified by a 4 dimensional vector $x_n = [x_{n,1} \ x_{n,2}; x_{n,3} \ x_{n,4}]^T$; $[x_{n,1} \ x_{n,3}]^T$ specifies position and $[x_{n,2} \ x_{n,4}]^T$ specifies velocity. The target dynamics are defined by:

$$x_n = \begin{bmatrix} 1 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 1 \end{bmatrix} x_{n-1} + \begin{bmatrix} v_1 \\ v_2 \\ v_3 \\ v_4 \end{bmatrix}$$

where v_1, v_3 and v_2, v_4 are i.i.d. zero mean Gaussian with variance $\sigma_{x,1}^2 = 0.09$ and v_2, v_4 are i.i.d. zero mean Gaussian with variance $\sigma_{x,2}^2 = 0.0025$. Probability of survival is set $p_S = 0.98$ and the birth intensity is defined as $\gamma = 0.2\mathcal{N}(\cdot; x_b, Q_b)$, where:

$$x_b = \begin{bmatrix} 30 \\ 0 \\ 30 \\ 0 \end{bmatrix}, \quad Q_b = \begin{bmatrix} 9 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 9 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

The position of the target is observed in additive, isotropic Gaussian noise with variance $\sigma_z^2 = 0.04$. The clutter intensity is set $\kappa = 0.001$ uniform on $[0, 100] \times [0, 100]$, corre-

sponding to an average number of 10 clutter points per scan. Without loss of generality, in this example we set $p_D = 1$.

While the structure of this model is simple, the low observation noise is a challenge for SMC algorithms. The localisation of the likelihood means that blind proposals have little chance of putting particles in regions of high weight.

The auxiliary particle PHD scheme described above was run with 3000 particles at each iteration. The proposal distribution was designed as specified in section (5.2), with $q_n(x|j, m) = f(x|x_{n-1}^{(j)})$. Therefore results illustrate the improvement in performance due to auxiliary variable techniques alone, and further improvements in performance can be expected through more structured design of $q_n(x|j, m)$. Results show comparison with the particle PHD scheme of [13], proposing from the prior with a total of 3000 particles including 1000 particles allocated for births. Both algorithms were initialised with $k_0 = 2$, sampling from $\mathcal{N}(\cdot; x_b, Q_b)$.




Figure 1. Observations.

Figure (1) shows a typical observation record generated from the model for the ground truth data displayed in figure (2). As noted in section (2), the peaks in the intensity function can be used to obtain state estimates. Typically this requires heuristic clustering of the particles, followed by estimation within each cluster. For the auxiliary particle PHD filter, there is a natural method which can be employed to this end without the increased computational cost of employing a clustering algorithm. This consists of simply computing estimates from particles clustered by which observation they are assigned to, and can be expected to work well when the clutter intensity is not too high. MMSE estimates computed using this method for a single run of the auxiliary particle algorithm are shown in figure (2).

Figures (4) and (5) show histograms of estimates of the number of targets from the two algorithms, averaged over 200 runs, each with a different observation set. The ground truth is shown in figure (3). Note that estimates of the number of targets are not affected by heuristic clustering as they are made on the basis of the total mass of the particle set.




Figure 2. Ground truth target positions (solid line) and state estimates (crosses) for auxiliary particle PHD filter.




Figure 3. Ground truth number of targets.




Figure 4. Auxiliary particle PHD filter. Estimated number of targets averaged over 200 observation realisations.

The estimates from the bootstrap algorithm show more uncertainty, frequently under estimating the number of targets. This is due to the bootstrap algorithm loosing track of targets and failing to identify the birth of a target at $n = 11$.

Figure (6) shows the effective sample size (ESS) of the normalised particle sets, calculated at each iteration and




Figure 5. Bootstrap particle PHD filter. Estimated number of targets averaged over 200 observation realisations.




Figure 6. Effective sample size for auxiliary and bootstrap algorithms averaged over 200 observation realisations.

then averaged across the 200 runs. The ESS, introduced in [8], provides a measure of the degeneracy of the weights, and has a maximum possible value of 100%. The decrease in ESS at $n = 11$ for the auxiliary particle algorithm is caused by the birth of a target. Birth particles, i.e. those with $j = N + 1$, are proposed from the prior and cannot benefit from auxiliary particle index effects from the previous time step. Performance at this time step is therefore the same as for the bootstrap algorithm. A more constructive choice of $q_n(x|j, m)$ would lead to improvements under these circumstances. At all other time steps, the performance of the auxiliary algorithm is significantly better than that of the bootstrap algorithm.

7. Conclusions

We have introduced an auxiliary particle implementation of the PHD filter. The proposed scheme involves auxiliary random variables which index particles and observations. The resulting algorithm samples on a higher dimensional space than previous particle implementations of the PHD filter and naturally permits more efficient proposals. Forthcoming work will give a more detailed analysis of the proposed particle algorithm and its variants.

8. Acknowledgements

The research of the second author was partially funded by Boeing.

References

- [1] Y. Bar-Shalom and X. R. Li, *Multitarget-Multisensor Tracking: Principles and Techniques*, YBS Publishing, 1995.
- [2] D. E. Clark and J. Bell, "Convergence results for the particle PHD filter", *IEEE Transactions on Signal Processing*, 54(7), July 2006, pp. 2652–2661.
- [3] D. J. Daley and D. Vere-Jones, *An Introduction to the Theory of Point Processes*, volume I: Elementary Theory and Methods of *Probability and Its Applications*, Springer, New York, second edition, 2003.
- [4] R. Douc, O. Cappé, and E. Moulines, "Comparison of resampling schemes for particle filtering", In *Proceedings of the IEEE 4th International Symposium on Image and Signal Analysis*, pp. 64–69, 2005.
- [5] A. Doucet, N. de Freitas, and N. Gordon, editors, *Sequential Monte Carlo Methods in Practice*, Statistics for Engineering and Information Science, Springer Verlag, New York, 2001.
- [6] A. Doucet, S. Godsill, and C. Andrieu, "On sequential Monte Carlo sampling methods for Bayesian filtering", *Statistics and Computing*, 10, 2000, pp. 197–208.
- [7] A. M. Johansen, S. Singh, A. Doucet, and B. Vo, "Convergence of the SMC implementation of the PHD filter", *Methodology and Computing in Applied Probability*, 8(2), June 2006, pp. 265–291.
- [8] A. Kong, J. S. Liu, and W. H. Wong, "Sequential imputations and Bayesian missing data problems", *Journal of the American Statistical Association*, 89(425), March 1994, pp. 278–288.
- [9] R. P. S. Mahler, "Multitarget bayes filtering via first-order multitarget moments", *IEEE Transactions on Aerospace and Electronic Systems*, October 2003, pp. 1152–1178.
- [10] M. K. Pitt and N. Shephard, "Filtering via simulation: Auxiliary particle filters", *Journal of the American Statistical Association*, 94(446), 1999, pp. 590–599.
- [11] H. Siddenblath, "Multi-target particle filtering for the probability hypothesis density", In *Proceedings of the International Conference on Information Fusion, Cairns, Australia*, pp. 800–806, 2003.
- [12] B. Vo and W.-K. Ma, "The Gaussian mixture probability hypothesis density filter", *IEEE Transactions on Signal Processing*, 54(11), November 2006, pp. 4091–4104.
- [13] B. Vo, S. Singh, and A. Doucet, "Random finite sets and sequential monte carlo methods in multi-target tracking", In *Proceedings of the International Conference on Information Fusion, Cairns, Australia*, pp. 792–799, 2003.
- [14] B. Vo, S. Singh, and A. Doucet, "Sequential Monte Carlo methods for multitarget filtering with random finite sets", *IEEE Transactions on Aerospace and Electronic Systems*, 41(4), October 2005, pp. 1224–1245.
- [15] T. Zajic and R. P. S. Mahler, "Particle-systems implementation of the PHD multitarget tracking filter", In *Proceedings of SPIE*, pp. 291–299, 2003.