PARTICLE PHD FORWARD FILTER-BACKWARD SIMULATOR FOR TARGETS IN CLOSE PROXIMITY

*Ramona Georgescu, Peter Willett**

University of Connecticut
Department of Electrical and Computer Engineering
Storrs, CT, USA

Lennart Svensson

Chalmers University of Technology
Department of Signals and Systems
Gothenburg, Sweden

ABSTRACT

In this work, we introduce the particle PHD forward filter - backward simulator (PHD-FFBSi) capable of dealing with uncertainties in the labeling of tracks that appear when tracking two targets in close proximity with measurements that do not discriminate between them. The Forward Filter Backward Simulator is a smoothing technique based on rejection sampling for the calculation of the probabilities of association between targets and tracks. The forward filter is a particle implementation of the Probability Hypothesis Density (PHD) filter that presents advantages over an SIR filter. Difficulties that arise due to the presence of target birth and death processes are addressed through modifications to the fast FFBSi. Simulations show the new particle filter of asymptotically linear complexity in the number of particles calculates correct target label probabilities at varying levels of measurement noise.

Index Terms— PHD, particle filter, FFBSi, smoothing, closely spaced targets

1. INTRODUCTION

In theory, particle filters provide optimal solutions to any target tracking problem. However, in practice they have significant problems in scenarios with targets that have been closely spaced for some time and later dissolve. Therein, uncertainties arise as to which target is which, what we refer to as target labeling uncertainties. Based on the information obtained from, e.g., a radar sensor, it is typically not possible to resolve uncertainties in labeling once they have appeared. Still, due to degeneracy, which is an inherent part of any particle filter (with finite number of particles), the filter will soon claim it knows the correct labeling with probability 1 [4]. Clearly, this does not reflect the true posterior distribution and moreover, it may have significant repercussions in many applications.

Consequently, recent years have brought significant interest in developing particle filters able to maintain accurate target label probabilities. A good formulation of the track labeling problem and summary of existing approaches can be found in [1]. In the following, we go briefly over the publications directly relevant to the current manuscript.

In [3], Blom and Bloem propose a new particle filter able to provide an estimated track swap probability for the case of two closely spaced linear Gaussian targets. At its core lies a unique decomposition of the joint conditional density as a mixture of a permutation invariant density and a permutation strictly variant density.

In [7], Garcia-Fernandez et al. introduce a particle filter able to maintain the multimodality of the posterior pdf after the targets have moved in close proximity and thus, to extract information about target labels. The drawback is that complexity grows as $O(N^2)$, where N is the number of particles, due to the association of a probability vector to each particle.

In [12], a closed form smoothing solution for the Probability Hypothesis Density (PHD) filter under linear Gaussian multi-target assumptions has been derived.

In [8], we introduced a particle filter of asymptotically linear complexity that consists of a sampling-importance-resampling (SIR) forward filter and a backward smoother. Target labels from the forward filter are ignored and target identity probabilities are calculated based only on the particle trajectories generated by backwards rejection sampling.

The $O(N^2)$ complexity of [3], [7] is prohibitive. [12] does not perform labeling and [8] assumes the forward filter is a particle filter, which is not always desirable. In this work, the PHD serves as the forward filter coupled with the FFBSi, bringing access to more difficult scenarios at the cost of relying on a potentially unreliable state extraction step.

2. FORWARD FILTERING-BACKWARD SIMULATION

The PHD-FFBSi algorithm starts with running the forward filter, here a particle implementation of the PHD filter, followed by backward sampling to smooth the forward PHD surface, achieved by the fast FFBSi recursion. Note that, in order to deal with the additional challenges of target birth and death, the fast FFBSi requires modifications that will be described in Section 2.2.2. Afterwards, target states are extracted from

*The work of R. Georgescu and P. Willett was supported by the Office of Naval Research under contract N00014-10-10412.

the smoothed PHD surface as they are a prerequisite for the computation of target label probabilities, the final step of the algorithm. Next, we go over the PHD-FFBSi in detail.

2.1. Particle PHD filter

In our simulations, we have used the particle implementation of the PHD filter in [6]. Here, we briefly sketch its steps.

1) Prediction: Assuming there are N particles at time $k-1$, their propagated states and corresponding weights are:

$$\xi_{k|k-1}^{(i)} = F\xi_{k-1}^{(i)} + v(k) \tag{1}$$

$$w_{k|k-1}^{(i)} = w_{k-1}^{(i)} P_s \tag{2}$$

for $i = 1, \dots, N$, where P_s is the target survival probability.

2) Target Birth: Each measurement z is treated as a potential new target and represented by N_z particles with locations sampled from $\mathcal{N}(z, \sigma_z^2)$. Birth particles have equal weights:

$$w_{k|k-1}^{(i)} = \frac{b_{k|k-1}(S)}{N_{new}} \tag{3}$$

where $b_{k|k-1}(S)$ is the total target birth weight for the surveillance region and N_{new}=$N_z|Z(k)|$ is the total number of particles proposed for the "investigation" of newborn targets.

3) Update: Particle weights are updated by:

$$w_k^{(i)} = \Big[1 - P_d + \sum_{z\in Z} \frac{P_d f(z|\xi^{(i)})}{\lambda_k c_k(z) + C(z)} + \Big] w_{k|k-1}^{(i)} \tag{4}$$

where $C(z) = \sum_{j=1}^{N+N_{new}} P_d f(z|\xi^{(j)}) w_{k|k-1}^{(j)}$.

4) Resampling: Particles $\{w_k^{(i)}, \xi^{(i)}\}_{i=1}^{i=N+N_{new}}$ are resampled proportionally to their weight in order to preserve the total weight of the PHD surface $\mathcal{T}_{k|k}$ (see Eq. (5)). After resampling, each particle is given equal weight. Further description and complete notation are in [11] and [6].

2.2. Backward Simulator

The fast FFBSi algorithm proposed in [5] is the backbone of our smoothing of the PHD filter. The original FFBSi formulation has $O(N^2)$ complexity. Instead, by using a rejection sampling, the fast FFBSi algorithm used here has asymptotically linear complexity in the number of particles.

2.2.1. Fast FFBSi

Input: Sequences of weighted particles $\{x_k^i, w_k^i\}_{i=1}^N$ describing forward filtering distributions $\pi(x_k|z_{1:k})$, $k = 1, \cdots, T$.

1. Initialize the backward trajectories:
   $\{I(j)\}_{j=1}^M \sim \{w_T^i\}_{i=1}^N$, $\tilde{x}_T^j = x_T^{I(j)}$, $j = 1, \cdots, M$
2. for $k = T-1 : 1$
3. $\quad L = \{1, \cdots, M\}$
4. $\quad$ while $L \neq \emptyset$
5. $\qquad n = |L|$
6. $\qquad \delta = \emptyset$
7. $\qquad$ Sample independently $\{C(q)\}_{q=1}^n \sim \{w_k^i\}_{i=1}^N$
8. $\qquad$ Sample independently $\{U(q)\}_{q=1}^n \sim \mathcal{U}([0,1])$
9. $\qquad$ for $q = 1 : n$
10. $\qquad\quad$ if $U(q) \leq \frac{f(\tilde{x}_{k+1}^{L(q)}|x_k^{C(q)})}{\rho}$
11. $\qquad\qquad I(L(q)) = C(q)$
12. $\qquad\qquad \delta = \delta \cup \{L(q)\}$
13. $\qquad\quad$ end
14. $\qquad$ end
15. $\qquad L = L \backslash \delta$
16. $\quad$ end
17. $\quad$ Append the samples to the backward trajectories.
   $\tilde{x}_k^j = x_k^{I(j)}$, $\tilde{x}_{k:T}^j = \{\tilde{x}_k^j, \tilde{x}_{k+1:T}^j\}$, $j = 1, \cdots, N$
18. end

Output: Collection of backward trajectories $\{\tilde{x}_{1:T}^j\}_{j=1}^M$ describing the joint smoothing distribution $\pi(x_{1:T}|z_{1:T})$.

The core of the above algorithm relies on updating L, the index list of samples at time k that still need assignment (smoothing particles) based on C, the index list of candidate samples at time k (filter particles) by testing whether the forward filter particle with (random) index $C(q)$ should be accepted as the smoothing particle with index $L(q)$ [10].

2.2.2. Birth and death of targets

Target birth is automatically dealt with by the algorithm in Section 2.2.1. If a target is born at time k, there are particles representing it on the forward PHD surface from time k onward as described in Section 2.1. In the backward pass, the fast FFBSi tests if sampled particles from the forward PHD surface at k (some of which would be representing the new target) can be accepted as particles on the smoothed PHD surface at k by checking whether the sampled particles from the forward PHD surface at k could give rise to the sampled particles from the smoothed PHD surface at $k+1$ in accordance with the target motion model.

On the other hand, target death requires an additional step to be introduced into the fast FFBSi (between lines 3 and 4). If a target dies at time k, there will not be particles representing it on the forward PHD surface from time k onward and hence, in the backward pass, sampled particles from the forward PHD surface at k representing this target would not pass the test as there would be no sampled particles from the smoothed PHD surface at $k+1$ that could be predicted from them using the target motion model. Therefore, in order to account for target deaths, at each scan k, $N(1-P_s)$ sampled particles from the forward PHD surface at k are directly accepted as particles on the smoothed PHD surface at k.

2.3. Peak Extraction

To estimate target states, the smoothed PHD surface is approximated by a Gaussian Mixture using the Expectation-

Maximization algorithm [2]. At each scan, the EM algorithm fits $M = |Z(k)| + [\mathcal{T}_{k-1|k-1}]$ Gaussian modes, where

$$\mathcal{T}_{k|k} = \sum_{j=1}^{N+N_{new}} w_k^{(j)} \quad (5)$$

The heaviest $\mathcal{T}_{k|k}$ modes are used in the computation of target label probabilities. Equations for weight, location and covariance of the resulting Gaussian modes are given in [6].

2.4. Computation of target label probabilities

Most tracking algorithms define target labels at the start of the scenario with the objective to track these labeled targets over time. We instead define the target labels at scenario end time. For the scans in the interval $[k_{start}, k_{end}]$ when there are two targets in the surveillance area, there are two possibilities of target to track labeling. Let $Pr(x_1, x_2; k)$ be the probability that target x_1 at time k, $\forall k \in [k_{start}, k_{end}]$, is target 1 at time k_{end} and thus target x_2 at time k is target 2 at time k_{end}. Clearly, for two targets, $Pr(x_2, x_1; k) = 1 - Pr(x_1, x_2; k)$.

Based on the backward trajectories created by the FF-BSi (see line 17 in Section 2.2.1), we count how many particles are assigned/close to which target at k_{end} and stay assigned/close to the same target at time k to form the probabilities of which target is which track. Note that $Pr(x_1, x_2; k)$ is not necessarily decreasing because trajectories can change assignment back and forth between targets from time k to k_{end}.

3. RESULTS

In our simulation, we considered the same scenario as in our previous work [8] involving two targets that come in close proximity of each other and subsequently diverge. State vectors are scalars evolving according to a Gaussian random walk model $x_k^j = x_{k-1}^j + v_{k-1}^j$, where the process noise is $v_{k-1}^j \sim \mathcal{N}(0, \sigma_v^2)$ for $j = 1, 2$. Both targets are detected at all times and there are no false alarms. Target detections are modeled as $z_k^j = x_k^j + u_k^j$, with measurement noise $u_k^j \sim \mathcal{N}(0, \sigma_u^2)$ but we do not know which detection is associated to which target. Three cases of measurement noise strength were investigated: low ($\sigma_u = 0.01$), medium ($\sigma_u = 0.1$) and high ($\sigma_u = 0.25$).

Additional difficulty is introduced compared to the scenarios in [8] by adding birth and death of targets. Here, target 1 appears at $k = 1$ and dies at time $k_{end} = 78$ and target 2 appears at $k_{start} = 23$ and is present throughout until $T = 100$. The same number of particles, i.e. $M = N = 2000$ are used for both the forward filter and the smoother.

Fig. 1 displays a typical Monte Carlo run in the case of medium measurement noise. Fig. 1a plots the measurements available to the PHD filter, Fig. 1b shows the target location estimates after the forward pass, i.e. the locations of the two heaviest peaks extracted from the forward PHD surface and Fig. 1c shows the target location estimates ater the backward pass, i.e. the locations of the two heaviest peaks extracted from the smoothed PHD surface. Expectedly, due to smoothing, the target location estimates in Fig. 1c are closer to the true trajectories than the target location estimates in Fig. 1b.

Fig. 2 shows the target label probabilities averaged over 100 Monte Carlo runs at scans in which both targets are present, for the same levels of measurement noise as above. In the case of low measurement noise, there is minimal overlap between the particle clouds when the two targets move in close proximity of each other and this is reflected in the values of $Pr(x_1, x_2; k)$, i.e. the probability that target x_1 at scan $k_{end} = 78$ is target 1 at scan k and target x_2 at scan k_{end} is target 2 at scan k, stays far from the permutation variant point of 0.5. Specifically, after the backward pass, we are about 75% certain that target x_1 at scan $k_{end} = 78$ is target 1 at $k_{start} = 23$ and that target x_2 at k_{end} is target 2 at k_{start}.

In the case of a medium measurement noise strength, after the backward pass we are still able to resolve the targets with $Pr(x_1, x_2; k_{start} = 23) \approx 55\%$. Note that $Pr(x_1, x_2; k)$ decreases dramatically between scans $k = 55$ to $k = 45$, the time in which the targets move in close proximity.

In extremely high measurement noise ($\sigma_u = 0.25$)[1], target label probabilities are slightly indicative of a switch as $Pr(x_1, x_2; k = 23) \approx 45\%$. This value is accurate and due to significant overlap between the particle clouds following the two targets induced by the very large process noise.

4. CONCLUSION

4.1. Relationship to our prior work

In our initial work on extracting target label probabilities in scenarios with targets that move in close proximity [8], an SIR filter was used as the forward filter resulting in the SIR-FFBSi. Next, we give the reasons for investigating the PHD as the forward filter and the differences with respect to the SIR. Our intention is not to compare performance but to point out that, when combined with the fast FFBSi, the PHD filter and the SIR lead to different approaches.

The PHD filter is better suited for use in multiple targets scenarios because the dimension of the PHD surface is the dimension of a single target state while SIR particle filters don't scale well in the number of targets. Moreover, the PHD filter is more versatile than the SIR when handling target birth/death, missed detections and false alarms. Note that in [8], a variable number of targets was not considered.

The PHD takes an unlabeled approach to tracking which on one hand, prevents us from outputing MMOSPA estimates [9] as we did for the SIR-FFBSi; instead, we output target states through peak extraction which could become less reliable in complex scenarios. On the other hand, due to the PHD being an unlabeled tracker, we avoid the step of removing the

[1]Note that the intertarget distance while targets are in close proximity is $d = 0.5$, i.e. $d = 2\sigma_u$ in high measurement noise!

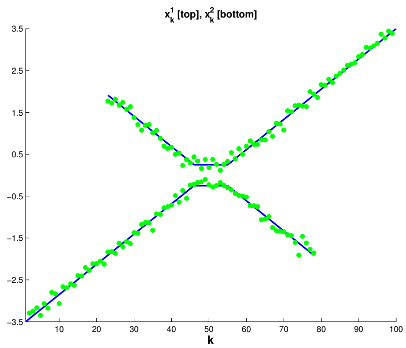

(a) Measurements

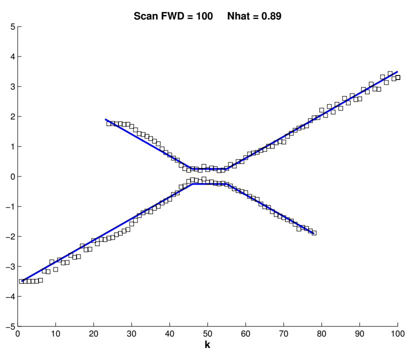

(b) After forward pass

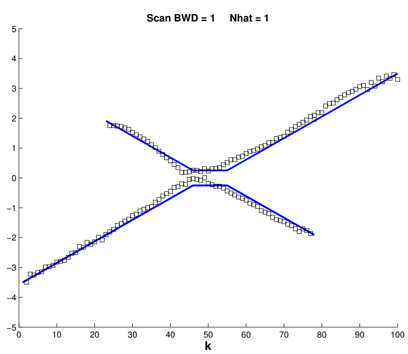

(c) After backward pass

Fig. 1. Target location estimates in a typical Monte Carlo run.

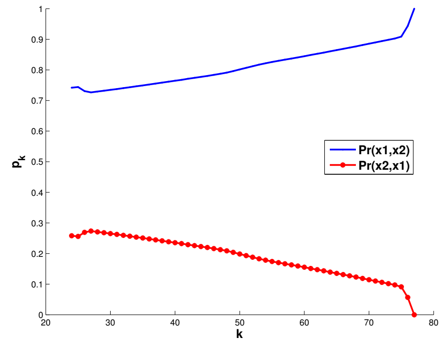

(a) Low

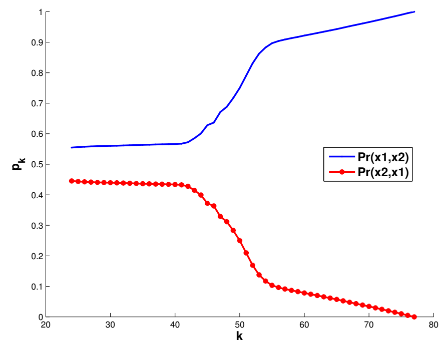

(b) Medium

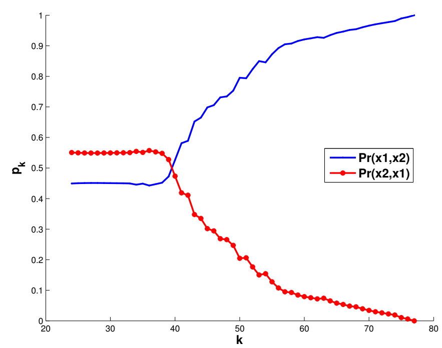

(c) High

Fig. 2. Target label probabilities vs. measurement noise.

target labels generated by the forward filter (in our case, the SIR) prior to smoothing, a step that required the creation of $N_T!$ copies for each particle in the forward filter, where N_T is the number of targets [8].

Issues such as the presence of 3+ targets, missed detections and false alarms and using a Gaussian Mixture implementation of PHD filter will be addressed in future work.

4.2. Relationship to other prior work

The emerging area of track labeling has been very active in recent years but there is quite a bit more work to be done [1].

The PHD-FFBSi algorithm presented here is a novel particle filter able to maintain accurate target label probabilities in challenging scenarios in which targets move in close proximity and later diverge. Our approach is based on smoothing (hence, it is different from e.g. the decomposed particle filtering proposed in [3]) and has approximately linear complexity in the number of particles as opposed to the squared complexity of the particle filter in [7]. Additionally, while the work in [12] improves past estimates, it does so without connecting them over time; instead, the PHD-FFBSi has the advantage of computing target trajectories. While the current study fits within recent approaches to joint multitarget tracking and labeling (see [1] and references therein), it stands out due to its flexibility in the choice of forward filter, low (linear) complexity, and promise of automated track management based on the particle trajectories output during smoothing.

5. REFERENCES

[1] E. Aoki, A. Bagchi, P. K. Mandal, and Y. Boers, "A theoretical analysis of Bayes-optimal multi-target tracking and labelling," Department of Applied Mathematics, University of Twente, Enschede, Netherlands, Memorandum 1953, 2011.

[2] J. Bilmes, "A gentle tutorial on the EM algorithm and its application to parameter estimation for Gaussian Mixture and Hidden Markov Models," University of Berkeley, Tech. Rep. ICSI-TR-97-02, 1998.

[3] H. A. P. Blom and E. A. Bloem, "Decomposed particle filtering and track swap estimation in tracking two closely spaced targets," in *Proc. of 14th International Conf. on Information Fusion*, Chicago, IL, 2011.

[4] Y. Boers and H. Driessen, "The mixed labeling problem in multi target particle filtering," in *Proc. of 10th International Conf. on Information Fusion*, Quebec City, Canada, 2007.

[5] R. Douc, A. Garivier, E. Moulines, and J. Olsson, "Sequential Monte Carlo smoothing for general state space hidden Markov models," *Submitted to Annals of Applied Probability*, 2010.

[6] O. Erdinc, P. Willett, and Y. Bar-Shalom, "Probability hypothesis density filter for multitarget multisensor tracking," in *Proc. of 8th International Conf. on Information Fusion*, Philadelphia, PA, 2005.

[7] A. F. Garcia-Fernandez, M. R. Morelande, and J. Grajal, "Particle filter for extracting target label information when targets move in close proximity," in *Proc. of 14th International Conf. on Information Fusion*, Chicago, IL, 2011.

[8] R. Georgescu, P. Willett, L. Svensson, and M. Morelande, "Two linear complexity particle filters capable of maintaining target label probabilities for targets in close proximity," in *Proc. of 15th International Conf. on Information Fusion*, Singapore, 2012.

[9] M. Guerriero, L. Svensson, D. Svensson, and P. Willett, "Shooting two birds with two bullets: how to find Minimum Mean OSPA estimates," in *Proc. of 13th International Conf. on Information Fusion*, Edinburgh, Scotland, 2010.

[10] F. Lindsten, "Rao-Blackwellised particle methods for inference and identification," Master's thesis, Linkoping University, Linkoping, Sweden, 2011.

[11] B.-N. Vo, S. Singh, and A. Doucet, "Sequential Monte Carlo implementation of the PHD filter for multi-target tracking," in *Proc. of 6th International Conf. on Information Fusion*, Cairns, Australia, 2003.

[12] B.-N. Vo, B.-T. Vo, and R. P. S. Mahler, "Closed-form solutions to forward-backward smoothing," *IEEE Trans. Signal Process.*, vol. 60, no. 1, pp. 2–17, 2012.