A LINEAR COMPLEXITY PARTICLE APPROACH TO THE EXACT MULTI-SENSOR PHD

Paolo Braca[⋆†] *Stefano Marano*[⋆] *Vincenzo Matta*[⋆] *Peter Willett*[†∗]

[⋆†]Research Department, NATO STO-CMRE, braca@cmre.nato.int
[⋆]DIEII, University of Salerno, {marano, vmatta}@unisa.it
[†] ECE Department, University of Connecticut, willett@engr.uconn.edu

ABSTRACT

Recently it has been shown that the Multi-Sensor Probability Hypothesis Density (MS-PHD) has some optimality properties in the regime of large number of sensors [1, 2], achieving the same performance of the Bayes multi-sensor/multi-target posterior in the Random Finite Set (RFS) framework [3]. However, when the number of sensors N is relatively large, the traditional PHD filter loses its computational efficiency, the complexity being exponential in N.

On the other hand, the complexity of the full Bayes posterior is only linear in N, and this paper suggests an idea for its computation using Sequential Monte Carlo (SMC) methods. The MS-PHD is then evaluated, and numerical examples show that it is possible to deal with a scenario where the number of sensors is very large while targets, appearing and disappearing, evolve in time.

Index Terms— Random finite sets, RFS, probability hypothesis density, PHD, multiple sensors.

1. RELATED WORK AND CONTRIBUTIONS

The objective of target tracking is to estimate the states of targets from measurement sets collected by some sensors at each time step. This is a challenging problem since the target can generate multiple measurements which are not always detected by sensors, and the sensors receive a set of spurious measurements (clutter) not generated by the target.

A number of effective idea sets to track multiple targets are of interest within the tracking and information fusion communities, amongst which we will mention algorithms such as Multiple Hypotheses Tracking (MHT) [4, 5] and Joint Probabilistic Data Association (JPDA) [4]; and such tools as RFS [6, 7], and Point Process (PP) theories [8].

At a conceptual level, estimation of a multitarget state Θ, from a collection of measurements Z, is fully accomplished once that the multi-target posterior density $f(\Theta|Z)$ is made available. Unfortunately, the evaluation thereof is usually computationally expensive, essentially due to the possible high dimensionality of the *multi-target* state-space (T one-dimensional targets means working in $\mathbb{R}^T$).

For this reason, a lot of effort in the multi-target field has been devoted to the search of *computationally efficient* estimation algorithms. In this class a major role is currently played by the so-called Probability Hypothesis Density [3], which is the first-moment of the multi-target state, and, as such, it is a more compact (and less informative) descriptor than the full Bayes posterior $f(\Theta|Z)$. The PHD function $D(\theta|Z)$ represents the expected density of targets at a given point θ, so that its integral gives the average number of targets, while its peaks can be used to estimate the target locations. One appeal of the PHD is that it lives in a space which has the (lower) dimensionality of the *single-target* state-space. In the single sensor case, this key property has been used to derive a very efficient algorithm for PHD filtering [11]. Unfortunately, in the multi-sensor case, generalizing this approach implies a growth in complexity which is exponential in the number of sensors N, see [12]. It would involve enumerating and updating the PHD under all combinations of measurement association events, for example that both (of two measurements at two sensors) measurements are false, that both come from different targets, that one (or the other) is false and the other true, and that both come from the same target. Attempts have been made to find alternative solutions, amongst which we mention: the iterated-corrector [12], which is only an approximation of the MS-PHD, where a distortion is reintroduced continually after a small number of updates; the partitioning approach proposed in [10], which is effective in the simplest case where many sensors exhibit non-overlapping Field of Views (FOVs).

A fundamental point, which has been perhaps overlooked, is that the complexity in evaluating the posterior $f(\Theta|Z)$ grows instead only *linearly* with N, due to the conditional independence among sensors. In this paper, we adopt this viewpoint and propose a particle representation of the posterior $f(\Theta|Z)$, generalizing the SMC approaches of [11] and [13] , which both refer to the single-sensor case.

We thus show that multi-sensor/multi-target tracking is in fact possible, though computationally demanding, provided that the number of interacting targets is not too large, even with a very large number of sensors.

∗P. Willett was supported by the U.S. Office of Naval Research under Grant N00014-10-10412.

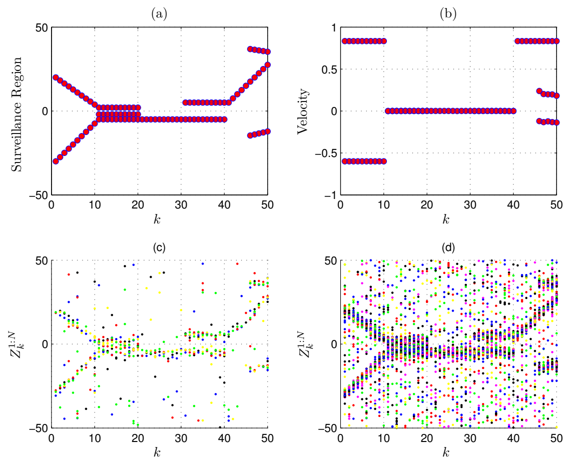

Fig. 1. Panel (a) target locations. Panel (b) target veloc ties. Panel $(c) - (d)$ data observed by $N = 5$ and $N = 5$ respectively, all sensors overlain.

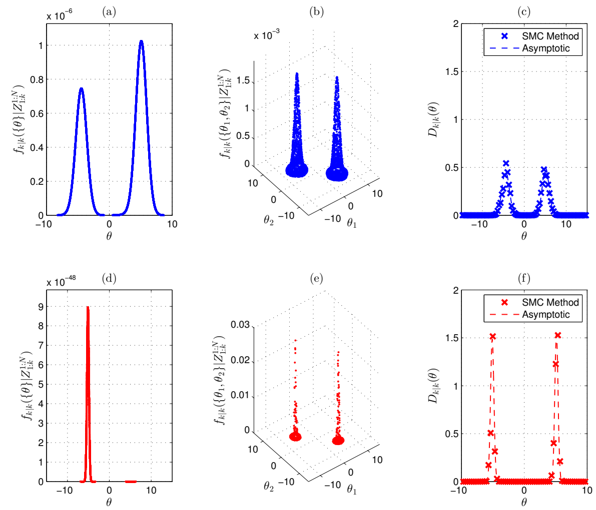

Fig. 2. Posterior distributions and MS-PHD at time scan $k = 35$ when two targets are present in locations $\{-5, 5\}$. Panels $(a)-(b)-(c)$ and $(d)-(e)-(f)$ refer to $N = 5$ and $N = 50$, respectively. In $(a)-(d)$ the posterior is computed in the case of a single target, while in $(b)-(e)$ in the case of two targets. Panels $(c)-(f)$ show the asymptotic approximation of the MS-PHD (4) and SMC approximation via Algorithm 1.

Once the posterior has been computed, a number of estimation algorithms can be in principle devised.

Among the many solutions, we here opt for using the MS-PHD as our compact descriptor to estimation purposes. This is suboptimal, and perhaps unnecessary, but it is made for two reasons. First, recent results showed that the MS-PHD and the full Bayes posterior present the same information about the targets' states as N gets large, namely, they become asymptotically efficient [1, 2]. Second, a feasible MS-PHD will at last allow it to compete directly in data fusion scenarios where the inference information stream is rich. We believe that this should motivate the search of a hybrid approach to find computationally efficient solutions, which should try to combine the more compact description provided by the PHD w.r.t. the posterior, without losing the advantages of sensors' independence.

2. PROBLEM FORMULATION AND DEFINITIONS

A multi-target system can be defined by the collection of individual targets' state (multi-target state) and the sensors' measurements (multi-target measurement). As the multi-target state and multi-target measurement evolve in time, the number of individual targets and measurements may change, i.e. the dimensions of the multi-target state and multi-target measurement also evolve in time. Moreover, there is no ordering for the elements of the multi-target state and measurement.

The multi-target state and multi-target measurement for the sensor $n \in \{1, \ldots, N\}$ at time k are naturally represented as finite subsets Θ_k and Z_k^n respectively [3, 11, 13].

The multi-target dynamics and observation can be described as follows. Given a realization Θ_{k-1} at time $k-1$, the multi-target state at time k is modeled by

$$\Theta_k = S_k(\Theta_{k-1}) \cup N_k(\Theta_{k-1}), \tag{1}$$

where $S_k(\Theta_{k-1})$ denotes the RFS of targets that have survived at time k, and $N_k(\Theta_{k-1})$ is the RFS of new targets. All aspects of multi-target motion such as the time-varying number of targets, individual target motion, target birth, spawning and target interactions are taken into account.

Similarly, given a realization Θ_k at time k, the multi-target observation at sensor n is modeled by

$$Z_k^n = Y_k^n(\Theta_k) \cup C_k^n, \tag{2}$$

where $Y_k^n(\Theta_k)$ represents the target-originated measurements that have been detected while C_k^n denotes the RFS of clutter, and then considering all sensor characteristics such as measurement noise, sensor field of view (i.e., state-dependent probability of detection) and false alarms. It is assumed that the sensors' data Z_k^n are independent conditioned on the target state Θ_k. The multitarget dynamic model (1)-(2) is formally ruled by the multi-target Markov transition distribution $f_{k|k-1}(\Theta_k | \Theta_{k-1})$, and the joint likelihood among the sensors $\prod_{n=1}^{N} g(Z_k^n | \Theta_k)$, see further details in [2, 11].

The multi-sensor/multi-target problem concerns the estimation of Θ_k at k given $Z_{1:k}^{1:N} \stackrel{def}{=} (Z_1^{1:N}, \ldots, Z_k^{1:N})$,

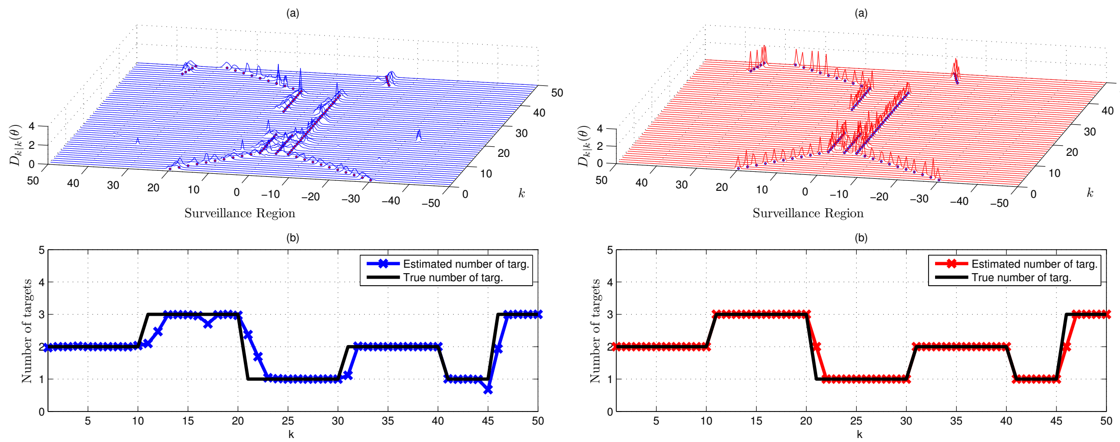

Fig. 3. Left side, panel (a) shows the entire evolution of the MS-PHD filter for a number of sensors $N = 5$. Right side, panel (a) shows the the MS-PHD for $N = 50$. Left and right side panel (b) show the true and the expected number of targets for $N = 5$ and $N = 50$, respectively.

where $Z_k^{1:N} \stackrel{def}{=} (Z_k^1, \ldots, Z_k^N)$. The object of interest in Bayesian estimation is the posterior probability density $f_{k|k}\left(\Theta_k | Z_{1:k}^{1:N}\right)$.

3. PARTICLE MULTI-TARGET MULTI-SENSOR FILTER

The single-target particle filter can be generalized to the multi-target case using as particles the *finite sets* [11, 13], i.e. the particles themselves can be of varying dimensions. Assume that at time $k-1$, a set of weighted particles $\left\{w_{k-1}^{(i)}, \Theta_{k-1}^{(i)}\right\}_{i=1}^{P}$ representing the multi-target posterior is available

$$f_{k-1|k-1}\left(\Theta_{k-1} \left| Z_{1:k-1}^{1:N}\right.\right) \approx \sum_{i=1}^{P} w_{k-1}^{(i)} \delta_{\Theta_{k-1}}\left(\Theta_{k-1}^{(i)}\right), \quad (3)$$

where $\delta_A(\cdot)$ is a delta function concentrated at A. The particle filter proceeds to approximate the multi-target posterior at time k by a new set of weighted particles $\left\{w_k^{(i)}, \Theta_k^{(i)}\right\}_{i=1}^{P}$ following Algorithm 1. The number of sensors N has an impact in terms of complexity, which is basically related to the computation of the likelihood $\prod_{n=1}^{N} g\left(Z_k^n \left| \widetilde{\Theta}_k^{(i)}\right.\right)$. Then the algorithm has a linear complexity with respect to N. On the other hand, the number of interacting targets should not be too large, in order to avoid the curse of dimensionality in the particle representation of the posterior.

At time k it will be possible to evaluate numerically the MS-PHD indicated by $D_{k|k}(\theta)$, where θ has the same dimensionality of the generic element in the set Θ_k.

From now on we shall focus on the case that the target state is composed by a position variable θ and a velocity variable $\dot{\theta}$, and the sensors collect measurements of the targets' positions.

Algorithm 1 Particle Multi-Target Multi-Sensor Filter

At time $k \geq 1$

- **Sampling Step**

- For $i = 1, \ldots, P$, sample $\widetilde{\Theta}_k^{(i)} \sim q\left(\cdot \left| \Theta_{k-1}^{(i)}, Z_k^{1:N}\right.\right)$, set $\tilde{w}_k^{(i)} = \dfrac{\prod_{n=1}^{N} g\left(Z_k^n \middle| \widetilde{\Theta}_k^{(i)}\right) f_{k|k-1}\left(\widetilde{\Theta}_k^{(i)} \middle| \Theta_{k-1}^{(i)}\right)}{q\left(\widetilde{\Theta}_k^{(i)} \middle| \Theta_{k-1}^{(i)}, Z_k^{1:N}\right)} w_{k-1}^{(i)}$.

- Normalise weights: $\sum_{i=1}^{P} \tilde{w}_k^{(i)} = 1$.

- **Resampling Step**

- Resample $\left\{\tilde{w}_k^{(i)}, \widetilde{\Theta}_k^{(i)}\right\}_{i=1}^{P}$ to get $\left\{w_k^{(i)}, \Theta_k^{(i)}\right\}_{i=1}^{P}$.

3.1. Approximation in the Regime of Large Number of Sensors

For N large enough, the posterior density is expected to be multimodal, exhibiting different peaks approximately located around the true targets' positions, with a spread decreasing as the number of sensors increases. Consequently, the MS-PHD has a similar shape, i.e., it is concentrated around the true targets' positions.

The above behavior is indeed predicted by the theoretical asymptotic results provided in [1, 2]. To elaborate, let us consider a one-dimensional scenario and introduce the vector of *ordered* targets' positions, namely, $[\xi_{1,k}, \xi_{2,k}, \ldots, \xi_{T_k,k}]$, and the Clairvoyant Maximum Likelihood (C-ML) estimator that knows in advance the correct number of targets T_k, say $\left[\widehat{\xi}_{1,k}, \widehat{\xi}_{2,k}, \dots, \widehat{\xi}_{T_k,k}\right]$. The term $\widehat{\xi}_{i,k}$ converges to the true target position $\xi_{i,k}$ as N increases, and the error $\widehat{\xi}_{i,k} - \xi_{i,k}$ is approximately a Gaussian random variable with a variance $\sigma^2_{i,k}/N$, where $\sigma^2_{i,k}$ is the i^{th} diagonal element of the inverse of the single-sensor FIM $\boldsymbol{I}_k$ at step k.

The large-sample shape of the MS-PHD related to the targets' positions is that of a Gaussian mixture, whose components are centered in the C-ML estimator points and have variances ruled by the Fisher information terms [1, 2]

$$D_{k|k}(\theta) \approx \sum_{i=1}^{T_k} \sqrt{\frac{N}{2\pi\sigma^2_{i,k}}} \exp\left\{ -N \frac{\left(\theta - \widehat{\xi}_{i,k}\right)^2}{2\,\sigma^2_{i,k}} \right\}. \quad (4)$$

Note that integrating (4) over the surveillance region we obtain the true number of targets T_k.

4. COMPUTER EXPERIMENTS

A target state is composed of position θ and velocity $\dot{\theta}$ variables. The targets evolution is modeled according to a nearly constant velocity. The birth/death process is modeled as Poisson with a constant rate, new targets are spread uniformly in the state space. The test scenario lasts 50 time steps and involves 6 targets (which appears and disappears), targets positions and velocities are depicted in the uppermost subplots of Figure 1 $(a) - (b)$. Two cases are analyzed with $N = 5, 50$ sensors which provide measurements with an independent Gaussian noise on the variable θ and false alarms uniformly spread inside the surveillance region $[-50, 50]$. The sensors' contacts are superimposed in the lowermost subplots of Figure 1 $(c) - (d)$. The contacts are given by the measurements of a targets' positions, if detected, and clutter.

In Figure 2 the posterior distributions and the MS-PHD of the targets' positions, evaluated using Algorithm 1, are represented at time scan $k = 35$. Panels $(a) - (b) - (c)$ and $(d) - (e) - (f)$ refer to $N = 5$ and $N = 50$, respectively. There are two targets, $T_k = 2$, and their true positions are in $\{-5, 5\}$. In $(a) - (d)$ the argument of the posterior is a singleton $\{\theta\}$ while in $(b) - (e)$ has two elements $\{\theta_1, \theta_2\}$. The posterior density has a multimodal shape and in particular has a significant mass around all the label permutation of targets' positions, i.e. around the points $(\theta_1, \theta_2) = (-5, 5)$ and $(\theta_1, \theta_2) = (5, -5)$. A very good agreement with the asymptotic approximation of the MS-PHD (4) is achieved, see panels $(c) - (e)$. As it is expected the shape of the MS-PHD and the posterior is peakier when the number of sensors is larger, or in other words when N is larger the estimation performance is better.

The left and right uppermost plots of Figure 3 show the entire evolution of the MS-PHD filter for $N = 5$ and $N = 50$, respectively. The lowermost plot of Figure 3 shows the expected number of targets (given the observation), computed integrating the MS-PHD over the surveillance region, this exhibits quite good performance for the estimation of the number of targets.

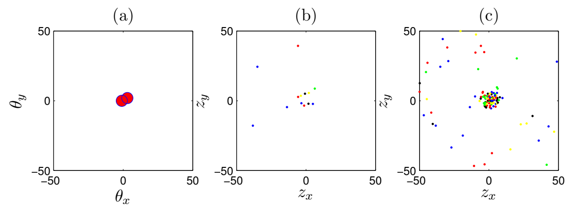

Fig. 4. Panel (a) two targets in locations $\{(-1, 0), (3, 2)\}$. Panel $(b) - (c)$ data observed by $N = 5$ and $N = 50$, respectively, all sensors overlain.

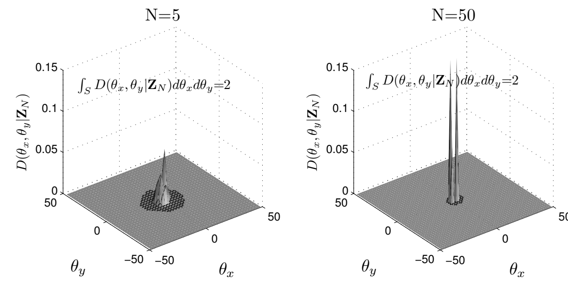

Fig. 5. MS-PHD in a 2D scenario for a number of sensors $N = 5, 50$.

The extension to a 2D scenario, with target locations (θ_x, θ_y), is given in Figure 4 and Figure 5. The MS-PHD is peakier when the number of sensors is larger, see Figure 5, indeed for $N = 5$ it is possible to localize the targets in the region close to the origin but not to distinguish one from the other, while for $N = 50$ the targets' locations are estimated with a good accuracy.

5. CONCLUSION

Sequential Monte Carlo methods are used to compute the full Bayes posterior of the targets' states for tracking purposes. This may be the first practical exact MS-PHD, and allows us finally to see the idea at work. In addition, it is shown that the the theoretical limit performance predicted by [1, 2] is effectively met.

6. REFERENCES

[1] P. Braca, S. Marano, V. Matta, and P. Willett, "Multitarget-multisensor ML and PHD: Some asymptotics," in *Proc. of Intern. Conf. on Information Fusion (FUSION)*, 2012.

[2] ——, "Asymptotic efficiency of the PHD in multitarget/multisensor estimation," *IEEE J. Sel. Topics Signal Process.* (submitted).

[3] R. Mahler, "Multitarget Bayes filtering via first-order multitarget moments," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 39, no. 4, pp. 1152–1178, Jan. 2003.

[4] Y. Bar-Shalom, P. Willett, and X. Tian, *Tracking and Data Fusion: A Handbook of Algorithms*. Storrs, CT: YBS Publishing, 2011.

[5] S. Blackman and R. Popoli, *Design and Analysis of Modern Tracking Systems*. Artech House, 1999.

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

[7] B.-T. Vo, B.-N. Vo, and A. Cantoni, "Bayesian filtering with random finite set observations," *IEEE Trans. Signal Process.*, vol. 56, no. 4, pp. 1313–1326, Apr. 2008.

[8] D. Daley and D. Vere-Jones, *An Introduction to the Theory of Point Processes, Volume I: Elementary Theory and Methods*. Springer, 2003.

[9] R. Mahler, "Approximate multisensor CPHD and PHD filters," in *Proc. of Intern. Conf. on Information Fusion (FUSION)*, Jul. 2010, pp. 1–8.

[10] E. Delande, E. Duflos, D. Heurguier, and P. Vanheeghe, "Multi-sensor PHD: Construction and implementation by space partitioning," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process. (ICASSP)*, May 2011, pp. 3632–3635.

[11] B.-N. Vo, S. Singh, and A. Doucet, "Sequential Monte Carlo methods for multitarget filtering with random finite sets," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 41, no. 4, pp. 1224 – 1245, Oct. 2005.

[12] R. Mahler, "The multisensor PHD filter: I. General solution via multitarget calculus," in *Proc. SPIE*, vol. 7336, May, 2009.

[13] C. Kreucher, K. Kastella, and A. O. Hero, "Multitarget Tracking using the Joint Multitarget Probability Density," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 41, pp. 1396–1413, Oct. 2005.