

Cardinality-Consensus-based PHD Filtering for Distributed Multitarget Tracking

Tiancheng Li, Franz Hlawatsch, *Fellow, IEEE*, and Petar M. Djurić, *Fellow, IEEE*

Abstract—We present a distributed probability hypothesis density (PHD) filter for multitarget tracking in decentralized sensor networks with severely constrained communication. The proposed “cardinality consensus” (CC) scheme uses communication only to estimate the number of targets (or, the cardinality of the target set) in a distributed way. The CC scheme allows for different implementations—e.g., using Gaussian mixtures or particles—of the local PHD filters. Although the CC scheme requires only a small amount of communication and of fusion computation, our simulation results demonstrate large performance gains compared with noncooperative local PHD filters.

Index Terms—Distributed multitarget tracking, cardinality consensus, probability hypothesis density filter, PHD filter.

I. INTRODUCTION

THE probability hypothesis density (PHD) filter is a popular method for multitarget tracking [1]–[3]. In decentralized sensor networks [4]–[6], a distributed PHD (D-PHD) filter can be used where each sensor runs a local PHD (L-PHD) filter and exchanges information with neighboring sensors. Most existing D-PHD filters, such as [7]–[11], rely on geometric averaging fusion, which is also known as generalized covariance intersection (GCI) or Chernoff fusion [12]–[14]. Other D-PHD filters [15], [16] use arithmetic averaging fusion, which was originally proposed for centralized multisensor PHD filtering [17]. These D-PHD filters disseminate Gaussian mixture (GM) parameters or a large number of particles and weights. However, distributed algorithms with a small amount of intersensor communication and of fusion computation are often desirable. Therefore, a D-PHD filter with only a small amount of disseminated and fused data is of great interest.

In this letter, we propose a D-PHD filter that disseminates and fuses merely a single positive parameter per sensor, namely, an estimate of the number of targets, i.e., of the cardinality of the set of target states. The local cardinality estimates are fused via a distributed protocol for “reaching consensus” on an average value of the local estimates. Accordingly, we refer to our cooperation scheme as “cardinality consensus” (CC). Compared with noncooperative L-PHD filters, the proposed CC-based D-PHD filter achieves a significant perfor-

mance improvement with very little communication and fusion computation. Moreover, in contrast to existing D-PHD filters, such as [7]–[11], [15], [16], the L-PHD filters are allowed to use different implementations, such as GM implementations [3] and sequential Monte Carlo (SMC) implementations [2].

This letter is organized as follows. The system model and the basic L-PHD filter are described in Section II. The proposed CC scheme is presented in Section III. Simulation results are reported in Section IV.

II. SYSTEM MODEL AND LOCAL PHD FILTERS

We consider N_k targets with random states $\mathbf{x}_k^{(n)} \in \mathbb{R}^d$, $n = 1, 2, \dots, N_k$. The collection of target states at time k is modeled by a random finite set (RFS) $\mathcal{X}_k = \{\mathbf{x}_k^{(1)}, \mathbf{x}_k^{(2)}, \dots, \mathbf{x}_k^{(N_k)}\}$ with random cardinality $N_k = |\mathcal{X}_k|$ [1]. There are S sensors indexed by $s \in \{1, 2, \dots, S\}$. At time k , each sensor s observes an RFS of $M_{s,k}$ measurements $\mathbf{z}_{s,k}^{(m)}$, i.e., $\mathcal{Z}_{s,k} = \{\mathbf{z}_{s,k}^{(1)}, \dots, \mathbf{z}_{s,k}^{(M_{s,k})}\}$. The sensor network is connected. We denote by $\mathcal{S}_s \subseteq \{1, 2, \dots, S\} \setminus \{s\}$ the set of “neighbor” sensors that communicate with sensor s .

Each sensor runs an L-PHD filter [1]–[3], which propagates across time k the local posterior PHD $D_k(\mathbf{x}|\mathcal{Z}_{s,1:k}) \geq 0$. Here, $\mathbf{x} \in \mathbb{R}^d$ and $\mathcal{Z}_{s,1:k} \triangleq (\mathcal{Z}_{s,1}, \dots, \mathcal{Z}_{s,k})$. The significance of $D_k(\mathbf{x}|\mathcal{Z}_{s,1:k})$ is based on the property that its integral over a region $\mathcal{R} \subseteq \mathbb{R}^d$ gives the local posterior expectation of the number of targets whose states are in \mathcal{R} , i.e., $\int_{\mathcal{R}} D_k(\mathbf{x}|\mathcal{Z}_{s,1:k}) d\mathbf{x} = \mathbb{E}[|\mathcal{X}_k \cap \mathcal{R}| | \mathcal{Z}_{s,1:k}]$ [18]. In particular, for $\mathcal{R} = \mathbb{R}^d$, since $|\mathcal{X}_k \cap \mathbb{R}^d| = |\mathcal{X}_k| = N_k$, we have

$$\int_{\mathbb{R}^d} D_k(\mathbf{x}|\mathcal{Z}_{s,1:k}) d\mathbf{x} = \mathbb{E}[N_k | \mathcal{Z}_{s,1:k}] = \sum_{n=0}^{\infty} n p(n | \mathcal{Z}_{s,1:k}), \quad (1)$$

where $p(n | \mathcal{Z}_{s,1:k}) = \Pr[N_k = n | \mathcal{Z}_{s,1:k}]$ is the local posterior probability mass function of N_k at sensor s . According to (1), $\int_{\mathbb{R}^d} D_k(\mathbf{x}|\mathcal{Z}_{s,1:k}) d\mathbf{x}$ is the minimum mean square error (MMSE) estimate of N_k from $\mathcal{Z}_{s,1:k}$, defined as $\hat{N}_{s,k}^{\text{MMSE}} = \mathbb{E}[N_k | \mathcal{Z}_{s,1:k}]$ [19]. That is,

$$\hat{N}_{s,k}^{\text{MMSE}} = \int_{\mathbb{R}^d} D_k(\mathbf{x}|\mathcal{Z}_{s,1:k}) d\mathbf{x}. \quad (2)$$

III. CARDINALITY CONSENSUS

Let $\hat{N}_{s,k}$ denote the local cardinality estimate produced by the L-PHD filter of sensor s at time k . In the proposed CC-based D-PHD filtering scheme, the objective of intersensor information exchange is merely to estimate the cardinality N_k by calculating an approximation of the average of all the $\hat{N}_{s,k}$ s,

T. Li is with the Institute of Systems Engineering, Xi’an Jiaotong University, Xi’an, China and with the BISITE group, School of Science, University of Salamanca, 37007 Salamanca, Spain, e-mail: t.c.li@usal.es. F. Hlawatsch is with the Institute of Telecommunications, TU Wien, 1040 Vienna, Austria, e-mail: franz.hlawatsch@tuwien.ac.at. P. M. Djurić is with the Department of Electrical and Computer Engineering, Stony Brook University, Stony Brook 11794, NY, USA, e-mail: petar.djuric@stonybrook.edu. This work was partially supported by the Marie Skłodowska-Curie Individual Fellowship under grant 709267, by the Austrian Science Fund (FWF) under grant P27370-N30, and by the NSF under grant CCF-1618999.

$$\bar{N}_k \triangleq \frac{1}{S} \sum_{s=1}^S \hat{N}_{s,k}. \quad (3)$$

Each sensor s obtains a “quasi-consensual” cardinality estimate $\hat{N}_{s,k}$ that approximates \bar{N}_k , and which is then used in the L-PHD filter to improve the filtering performance.

A. The Cardinality Consensus Algorithm

The CC algorithm relies on the following approximative expansion of the local posterior PHD at each sensor s :

$$D_k(\mathbf{x}|\mathcal{Z}_{s,1:k}) \approx \sum_{j=1}^{J_{s,k}} w_{s,k}^{(j)} \psi_{s,k}^{(j)}(\mathbf{x}), \quad (4)$$

with $J_{s,k}$ weights $w_{s,k}^{(j)} \geq 0$ and functions $\psi_{s,k}^{(j)}(\mathbf{x})$ satisfying $\int_{\mathbb{R}^d} \psi_{s,k}^{(j)}(\mathbf{x}) d\mathbf{x} = 1$. A PHD expansion of the form (4) has been previously used in two PHD filter implementations, here referred to as GM-PHD and SMC-PHD filter. In the GM-PHD filter [3], $\psi_{s,k}^{(j)}(\mathbf{x}) = \mathcal{N}(\mathbf{x}; \mathbf{m}_{s,k}^{(j)}, \mathbf{P}_{s,k}^{(j)})$, which denotes a Gaussian probability density function with mean vector $\mathbf{m}_{s,k}^{(j)}$ and covariance matrix $\mathbf{P}_{s,k}^{(j)}$. In the SMC-PHD filter [2], $\psi_{s,k}^{(j)}(\mathbf{x}) = \delta(\mathbf{x} - \mathbf{x}_{s,k}^{(j)})$, where $\delta(\mathbf{x})$ is the Dirac delta function and $\mathbf{x}_{s,k}^{(j)}$ is a particle. Our CC scheme now consists of the following three steps, which are performed at each time k .

Step 1—Local cardinality estimation: Each sensor s forms a local estimate $\hat{N}_{s,k}$ of N_k . Inserting (4) into (2) yields $\hat{N}_{s,k}^{\text{MMSE}} \approx \sum_{j=1}^{J_{s,k}} w_{s,k}^{(j)}$. Thus, we use

$$\hat{N}_{s,k} \triangleq \sum_{j=1}^{J_{s,k}} w_{s,k}^{(j)}. \quad (5)$$

Step 2—Fusion of local cardinality estimates: Next, the local cardinality estimates $\hat{N}_{s,k}$ are fused to obtain at each sensor s a quasi-consensual approximation $\hat{\hat{N}}_{s,k}$ of the networkwide cardinality average \bar{N}_k . To this end, each sensor s participates in a distributed dissemination/fusion scheme, based on communication with its neighbors $r \in \mathcal{S}_s$. This scheme may make use of average consensus [5], [20], [21], gossip [15], [22], [23], flooding [24], or diffusion [25], [26] (see Section III-B).

Step 3—Weight scaling: At each sensor s , the local weights $w_{s,k}^{(j)}$ are scaled as

$$\tilde{w}_{s,k}^{(j)} = \frac{\hat{\hat{N}}_{s,k}}{\hat{N}_{s,k}} w_{s,k}^{(j)}, \quad j = 1, \dots, J_{s,k}. \quad (6)$$

Here, the scaling factor is chosen as $\hat{\hat{N}}_{s,k}/\hat{N}_{s,k}$ because then the new weights $\tilde{w}_{s,k}^{(j)}$ are consistent with the quasi-consensual cardinality estimate $\hat{\hat{N}}_{s,k}$ in the sense that $\sum_{j=1}^{J_{s,k}} \tilde{w}_{s,k}^{(j)} = \hat{\hat{N}}_{s,k}$, in analogy to (5). Indeed, we have $\sum_{j=1}^{J_{s,k}} \tilde{w}_{s,k}^{(j)} = (\hat{\hat{N}}_{s,k}/\hat{N}_{s,k}) \sum_{j=1}^{J_{s,k}} w_{s,k}^{(j)} = (\hat{\hat{N}}_{s,k}/\hat{N}_{s,k}) \hat{N}_{s,k} = \hat{\hat{N}}_{s,k}$, where (5) was used. At the next time step $k+1$, the scaled weights $\tilde{w}_{s,k}^{(j)}$ are used by the L-PHD filter instead of the original weights $w_{s,k}^{(j)}$. We note that (6) corresponds to an analogous scaling of the local posterior PHD $D_k(\mathbf{x}|\mathcal{Z}_{s,1:k})$: substituting $\tilde{w}_{s,k}^{(j)}$ for $w_{s,k}^{(j)}$ in (4), one obtains

$$\sum_{j=1}^{J_{s,k}} \tilde{w}_{s,k}^{(j)} \psi_{s,k}^{(j)}(\mathbf{x}) \approx \frac{\hat{\hat{N}}_{s,k}}{\hat{N}_{s,k}} D_k(\mathbf{x}|\mathcal{Z}_{s,1:k}),$$

and it is easily verified that the integral of $(\hat{\hat{N}}_{s,k}/\hat{N}_{s,k}) \times D_k(\mathbf{x}|\mathcal{Z}_{s,1:k})$ is approximately equal to $\hat{\hat{N}}_{s,k}$.

An important advantage of the CC scheme presented above is that the concrete nature of the expansion (4)—i.e., the type of the functions $\psi_{s,k}^{(j)}(\mathbf{x})$ and the order $J_{s,k}$ —may be different across the L-PHD filters used at the individual sensors.

For a comparison of the communication requirements per dissemination/fusion iteration, consider a 4-D state vector composed of 2-D position and velocity. In the proposed CC-based D-PHD filter, each sensor communicates only one real value (its local cardinality estimate). In a GM-based D-PHD filter, each sensor communicates the GM parameters, which amount to $15 \cdot J_{s,k}$ real values (for each Gaussian, there are four mean values, 10 values describing the covariance matrix, and one weight), and in an SMC-based D-PHD filter, each sensor communicates $5 \cdot J_{s,k}$ real values (the state and weight of each particle amount to a total of five real values). Clearly, the communication requirements of the CC-based D-PHD filter are lower by orders of magnitude.

B. Fusion Schemes

We will briefly review several alternative fusion schemes that can be used in Step 2. First, the average consensus [5], [20], [21] is an iterative distributed algorithm where the iterated cardinality estimate at sensor s is updated as

$$\hat{N}_{s,k}^{[t]} = \sum_{r \in \{s\} \cup \mathcal{S}_s} \omega_{s,r} \hat{N}_{r,k}^{[t-1]}, \quad t = 1, 2, \dots \quad (7)$$

A popular choice of the weights $\omega_{s,r}$ is given by what is known as Metropolis weights [20]. The recursion (7) is initialized as $\hat{N}_{s,k}^{[0]} = \hat{N}_{s,k}$. If the network is connected, as we assumed in Section II, then $\lim_{t \rightarrow \infty} \hat{N}_{s,k}^{[t]} = \bar{N}_k$. After a finite number t_F of iterations, $\hat{N}_{s,k}^{[t_F]}$ provides only an approximation of \bar{N}_k , and the $\hat{N}_{s,k}^{[t_F]}$ s at different sensors s will be (slightly) different.

Alternatively, one can use a flooding scheme [24] where at iteration $t \in \{1, 2, \dots, t_F\}$, each sensor forwards all the newly received cardinality estimates to its neighbors. Let $\mathcal{S}_s(t_F)$ denote the set of sensors that are at most t_F hops away from sensor s , including sensor s . Then, after the final iteration t_F , each sensor received the local cardinality estimates $\{\hat{N}_{r,k}\}_{r \in \mathcal{S}_s(t_F)}$, and it computes their average, i.e.,

$$\hat{N}_{s,k}^{[t_F]} = \frac{1}{|\mathcal{S}_s(t_F)|} \sum_{r \in \mathcal{S}_s(t_F)} \hat{N}_{r,k}.$$

If $t_F \geq D$, where D is the diameter of the network, then $\hat{N}_{s,k}^{[t_F]} = \bar{N}_k$. If $t_F < D$, then $\hat{N}_{s,k}^{[t_F]}$ is only an approximation of \bar{N}_k .

Instead of the arithmetic average \bar{N}_k in (3), one may consider using the geometric average, i.e., $\bar{N}_k^G \triangleq \left(\prod_{s=1}^S \hat{N}_{s,k} \right)^{1/S}$. Writing $\bar{N}_k^G = \exp(L_k)$, with $L_k \triangleq \ln \bar{N}_k^G = \frac{1}{S} \sum_{s=1}^S \ln \hat{N}_{s,k}$, we see that an approximation of L_k can be obtained by applying the average consensus scheme or the flooding scheme to the quantities $\ln \hat{N}_{s,k}$ instead of $\hat{N}_{s,k}$, $s = 1, \dots, S$.

The convergence of the average consensus and flooding schemes is discussed in [20], [21] and in [24], respectively.

C. Accuracy of Cardinality Estimation

For a simple analysis of the accuracy of the quasi-consensual cardinality estimates provided by the CC, $\hat{N}_{s,k}$, we make the following idealizing assumptions: (A1) The local cardinality estimates $\hat{N}_{s,k}$ (see (5)) at different sensors s are conditionally statistically independent given the target state sequence $\mathcal{X}_{1:k} \triangleq (\mathcal{X}_1, \dots, \mathcal{X}_k)$. (A2) The $\hat{N}_{s,k}$ s are conditionally unbiased, i.e., $E[\hat{N}_{s,k} | \mathcal{X}_{1:k}] = N_k$, and they have equal conditional variances, i.e., $\text{var}[\hat{N}_{s,k} | \mathcal{X}_{1:k}] =: V_k$. (A3) The quasi-consensual estimates $\hat{N}_{s,k}$ are exactly equal to the average \bar{N}_k in (3). One can then show the following results on the conditional statistics of $\hat{N}_{s,k}$ given $\mathcal{X}_{1:k}$:

- $\hat{N}_{s,k}$ is conditionally unbiased, i.e., $E[\hat{N}_{s,k} | \mathcal{X}_{1:k}] = N_k$, and its conditional variance is $\text{var}[\hat{N}_{s,k} | \mathcal{X}_{1:k}] = V_k/S$, i.e., smaller by the factor S —the number of sensors—than that of the local estimate $\hat{N}_{s,k}$.
- The conditional mean-square error (MSE) of $\hat{N}_{s,k}$, $\epsilon_{\mathcal{X}_{1:k}} \triangleq E[(\hat{N}_{s,k} - N_k)^2 | \mathcal{X}_{1:k}]$, is given by $\epsilon_{\mathcal{X}_{1:k}} = V_k/S$, which again decreases linearly with S .
- The conditional probability that $\hat{N}_{s,k}$ deviates from the true cardinality N_k by more than a given positive number ν is upper bounded as $\Pr[|\hat{N}_{s,k} - N_k| \geq \nu | \mathcal{X}_{1:k}] \leq V_k/(S\nu^2)$. This is smaller by the factor S than an analogous bound involving the local cardinality estimate $\hat{N}_{s,k}$ (which reads as $\Pr[|\hat{N}_{s,k} - N_k| \geq \nu | \mathcal{X}_{1:k}] \leq V_k/\nu^2$).

We can also analyze the unconditional statistics of $\hat{N}_{s,k}$. Note that as we no longer condition on $\mathcal{X}_{1:k}$, $N_k = |\mathcal{X}_k|$ is a random variable. We obtain that the CC estimator is also unconditionally unbiased, i.e., $E[\hat{N}_{s,k}] = E[N_k]$, and its variance is $\text{var}[\hat{N}_{s,k}] = \text{var}[N_k] + V_k/S$. Further, the unconditional MSE $\epsilon \triangleq E[(\hat{N}_{s,k} - N_k)^2]$ is V_k/S , and thus equal to $\epsilon_{\mathcal{X}_{1:k}}$. Finally, $\Pr[|\hat{N}_{s,k} - E[N_k]| \geq \nu] \leq (\text{var}[N_k] + V_k/S)/\nu^2$.

IV. SIMULATION RESULTS

We simulated eight targets in the region of interest (ROI) $[-1000\text{m}, 1000\text{m}] \times [-1000\text{m}, 1000\text{m}]$. Figure 1 shows the sensor network and the target trajectories. The target states consist of position and velocity, i.e., $\mathbf{x}_k = [x_k \ \dot{x}_k \ y_k \ \dot{y}_k]^T$. The target survival probability is 0.98. The states of the surviving targets evolve independently according to a nearly constant velocity model, i.e., $\mathbf{x}_k = \mathbf{F}\mathbf{x}_{k-1} + \mathbf{G}\mathbf{u}_k$, where $\mathbf{F} \in \mathbb{R}^{4 \times 4}$ and $\mathbf{G} \in \mathbb{R}^{4 \times 2}$ are as given in [27, Eq. (14)] with sampling period $\Delta = 1\text{s}$ and \mathbf{u}_k is an independent and identically distributed, zero-mean, Gaussian, 2-D system process with component standard deviation 5m/s^2 . Target birth is modeled by a Poisson RFS with intensity $\gamma_k(\mathbf{x}_k) = 0.05 \cdot \mathcal{N}(\mathbf{x}_k; \mathbf{m}_1, \mathbf{Q}) + 0.05 \cdot \mathcal{N}(\mathbf{x}_k; \mathbf{m}_2, \mathbf{Q})$, where $\mathbf{m}_1 = [250\text{m} \ 0\text{m/s} \ 250\text{m} \ 0\text{m/s}]^T$, $\mathbf{m}_2 = [-250\text{m} \ 0\text{m/s} \ -250\text{m} \ 0\text{m/s}]^T$, and $\mathbf{Q} = \text{diag}\{100\text{m}^2, 25\text{m}^2/\text{s}^2, 100\text{m}^2, 25\text{m}^2/\text{s}^2\}$.

There are $S = 20$ sensors, which are placed and connected as shown in Fig. 1. The field of view of each sensor is a disc of radius 3000m centered at the sensor position; this disc always covers the entire ROI. The target-originated mea-




Fig. 1. ROI, sensor network, and target trajectories. Blue squares and red circles indicate the positions of the sensors using UGM-PHD and SMC-PHD filters, respectively, black dashed lines indicate the communication links between neighboring sensors, and colored lines with dots indicate the target trajectories (with starting and ending times noted).

surements conform to the nonlinear range-bearing model, i.e., $\mathbf{z}_{s,k} = [\rho(\mathbf{x}_k) \ \theta(\mathbf{x}_k)]^T + \mathbf{v}_{s,k}$, with $\rho(\mathbf{x}_k) \triangleq ((x_k - x^{(s)})^2 + (y_k - y^{(s)})^2)^{1/2}$ and $\theta(\mathbf{x}_k) \triangleq \tan^{-1}(\frac{x_k - x^{(s)}}{y_k - y^{(s)}})$. Here, $x^{(s)}$ and $y^{(s)}$ are the coordinates of sensor s , and the components of the measurement noise $\mathbf{v}_{s,k}$ are independent, zero-mean, and Gaussian with standard deviations $\sigma_\rho = 10\text{m}$ and $\sigma_\theta = (\pi/90)\text{rad}$. These target-originated measurements are produced with target state dependent detection probability $p_D(\mathbf{x}_k) = 0.95 \cdot \mathcal{N}(\boldsymbol{\mu}_D(\mathbf{x}_k); \mathbf{0}, \sigma_D^2 \mathbf{I}_2) / \mathcal{N}(\mathbf{0}; \mathbf{0}, \sigma_D^2 \mathbf{I}_2)$, where $\boldsymbol{\mu}_D(\mathbf{x}_k) \triangleq [x_k - x^{(s)} \ |y_k - y^{(s)}|]^T$ and $\sigma_D = 6000\text{m}$. There are also clutter measurements, which are uniformly distributed over the sensor's field of view with an average rate of 10 points per time step, corresponding to clutter intensity $\kappa_k = 10/(2\pi \cdot 3000) = 5.31 \cdot 10^{-4}$.

The L-PHD filters are unscented GM-PHD (briefly UGM-PHD) filters [3] at ten sensors and SMC-PHD filters [2] at the other ten sensors, as shown in Fig. 1. The UGM-PHD filters use at most 100 Gaussian components; they prune components with weights below 10^{-5} and merge components with Mahalanobis distance below 4. In the SMC-PHD filters, the number of particles is adapted via resampling to be $200 \cdot \hat{N}_{s,k}$ if $\hat{N}_{s,k} \geq 0.5$ and 100 otherwise. The output of the UGM-PHD filters is obtained by extracting the Gaussian components with weights larger than 0.5, while the output of the SMC-PHD filters is obtained as described in [28], using the quasi-consensual cardinality estimate $\hat{N}_{s,k}$.

We performed 100 simulation runs with 80 time steps each, using the target trajectories from Fig. 1 and different measurement noise and initial particles. We compare the following fusion schemes (cf. Section III-B): average consensus computing the arithmetic or geometric average and flooding computing the arithmetic average; these are labeled A-CC, G-CC, and F-CC, respectively. We also consider genie-aided L-PHD filters that know N_k and thus scale the particle weights such that their sum equals N_k (i.e., the scaling factor $\hat{N}_{s,k}/N_k$ in (6) is replaced by $N_k/\hat{N}_{s,k}$). The accuracy of distributed cardinality estimation is measured by the root MSE (RMSE), $\text{RMSE}_{s,k}$

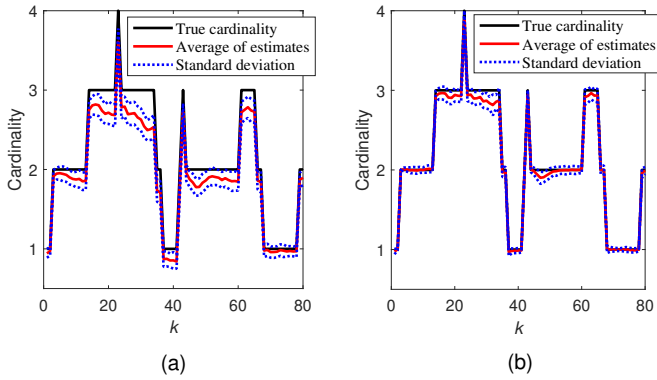


Fig. 2. Comparison of the accuracy of noncooperative and cooperative cardinality estimation versus time: Average (over sensors and simulation runs) and standard deviation of (a) the local cardinality estimates $\hat{N}_{s,k}$ and (b) the cooperative F-CC cardinality estimates $\hat{N}_{s,k}^{[t_F=5]}$.

$\triangleq \left(\frac{1}{100} \sum_{l=1}^{100} (\hat{N}_{s,k}^{(l)} - N_k)^2 \right)^{1/2}$, where $\hat{N}_{s,k}^{(l)}$ denotes the quasi-consensual cardinality estimate obtained in simulation run l . Furthermore, the target detection and position estimation performance of the D-PHD filters is measured by the optimal subpattern assignment (OSPA) error with cut-off parameter $c=1000\text{m}$ and order $\rho=2$ [29].

Figure 2(a) shows the noncooperative, local cardinality estimates $\hat{N}_{s,k}$, averaged over all the sensors s and all the 100 simulation runs, along with their standard deviation, versus time k . Figure 2(b) shows the average and standard deviation of the cooperative results of F-CC obtained after five flooding iterations, i.e., $\hat{N}_{s,k}^{[t_F=5]}$. Note that as the network diameter is $D=5$, $\hat{N}_{s,k}^{[t_F=5]}$ equals the arithmetic average \bar{N}_k . It is seen that the cooperative results are significantly better than the purely local results in that their average is closer to the true cardinality and their standard deviation is smaller.

Figure 3 shows the cardinality RMSE and the position OSPA error obtained with A-CC, G-CC, and F-CC versus the fusion iteration index $t \in \{0, 1, \dots, 5\}$. The figure shows averages taken over all the simulation runs, times, and sensors. Note that $t=0$ corresponds to no cooperation. For all the three fusion schemes, both the cardinality RMSE and the OSPA error are seen to decrease significantly with t . In fact, there is a large reduction already after a single iteration ($t=1$).

More specifically, for F-CC after five iterations ($t=5$), the cardinality RMSE and the OSPA error are reduced to, respectively, 25.1% and 53.0% of the values obtained without cooperation. For A-CC, these percentages are 36.7% and 57.7%, respectively, and for G-CC, they are 48.7% and 70.9%, respectively. Thus, the cardinality RMSE and the OSPA error of F-CC are lower than those of the other schemes (except the genie-aided filter), whereas those of G-CC are much higher. This means that the arithmetic cardinality average leads to a significantly better performance than the geometric cardinality average. We also observed that when the detection probability $p_D(\mathbf{x}_k)$ is reduced, the OSPA performance obtained with the geometric average is further degraded whereas that obtained with the arithmetic average is not changed significantly. The OSPA error of the F-CC-based filter for $t=5$ is seen to be very close to that of the genie-aided filter; this means that

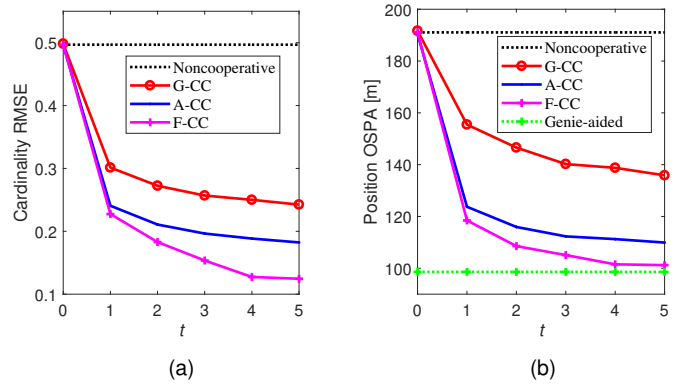


Fig. 3. Accuracy of distributed cardinality estimation and of the resulting D-PHD filter versus the number of dissemination/fusion iterations: (a) Average cardinality RMSE, (b) average position OSPA.

when each L-PHD filter uses the exact arithmetic cardinality average \bar{N}_k , the OSPA performance is almost as good as if each L-PHD filter used the true cardinality N_k . Finally, the cardinality RMSE of F-CC for $t=5$ (25.1% of the noncooperative value) is quite close to the square root of the cardinality MSE predicted by our strongly idealizing theoretical analysis in Section III-C, which is 22.4% of the noncooperative value. (Indeed, the theoretical RMSE with and without cooperation are obtained from Section III-C as $\sqrt{V_k/S}$ and $\sqrt{V_k}$, respectively, and thus their ratio is $\sqrt{1/S} \approx 0.224$.)

The number of real values broadcast by one sensor at one time step during five dissemination/fusion iterations, averaged over all the sensors, was measured as 19.0 for F-CC and 5.0 for A-CC and G-CC. Thus, the communication cost of F-CC is much higher than the communication costs of A-CC and G-CC (which are equal); this is the price paid for the faster convergence of F-CC (cf. [24]). We note that the communication cost of F-CC, A-CC, and G-CC does not depend on the number of targets.

V. CONCLUSION

We proposed a “cardinality consensus” (CC) scheme for distributed PHD filtering in which the local PHD filters at the various sensors adjust their weights based on a quasi-consensual estimate of the cardinality of the target set. This estimate was obtained by averaging the local cardinality estimates in a distributed way, using an average consensus or a flooding protocol. The CC scheme achieves a large improvement in target detection/tracking performance with only a small amount of intersensor communication and of fusion computation, and it allows for different implementations of the local PHD filters. Our simulations also showed that the flooding protocol outperforms the average consensus protocol at the cost of higher communication requirements, and that the arithmetic cardinality average outperforms the geometric cardinality average. We finally note that the CC scheme can also be combined with the fusion strategies underlying state-of-the-art distributed PHD filters, and will then typically improve the performance of these filters.

REFERENCES

- [1] R. P. S. Mahler, "Multitarget Bayes filtering via first-order multitarget moments," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 39, no. 4, pp. 1152–1178, Oct. 2003.
- [2] 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.
- [3] B. N. Vo and W. K. Ma, "The Gaussian mixture probability hypothesis density filter," *IEEE Trans. Signal Process.*, vol. 54, no. 11, pp. 4091–4104, Nov. 2006.
- [4] F. Zhao and L. Guibas, *Wireless Sensor Networks: An Information Processing Approach*. San Francisco, CA, USA: Morgan Kaufmann, 2004.
- [5] R. Olfati-Saber, J. A. Fax, and R. M. Murray, "Consensus and cooperation in networked multi-agent systems," *Proc. IEEE*, vol. 95, no. 1, pp. 215–233, Jan. 2007.
- [6] O. Hlinka, F. Hlawatsch, and P. M. Djurić, "Distributed particle filtering in agent networks: A survey, classification, and comparison," *IEEE Signal Process. Mag.*, vol. 30, no. 1, pp. 61–81, Jan. 2013.
- [7] G. Battistelli, L. Chisci, C. Fantacci, A. Farina, and A. Graziano, "Consensus CPHD filter for distributed multitarget tracking," *IEEE J. Sel. Topics Signal Process.*, vol. 7, no. 3, pp. 508–520, Jun. 2013.
- [8] M. Üney, D. E. Clark, and S. J. Julier, "Distributed fusion of PHD filters via exponential mixture densities," *IEEE J. Sel. Topics Signal Process.*, vol. 7, no. 3, pp. 521–531, Jun. 2013.
- [9] G. Battistelli, L. Chisci, and C. Fantacci, "Parallel consensus on likelihoods and priors for networked nonlinear filtering," *IEEE Signal Process. Lett.*, vol. 21, no. 7, pp. 787–791, Jul. 2014.
- [10] M. Gunay, U. Orguner, and M. Demirekler, "Chernoff fusion of Gaussian mixtures based on sigma-point approximation," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 52, no. 6, pp. 2732–2746, Dec. 2016.
- [11] G. Li, W. Yi, M. Jiang, and L. Kong, "Distributed fusion with PHD filter for multi-target tracking in asynchronous radar system," in *Proc. RadarConf 2017*, Seattle, WA, USA, May 2017, pp. 1434–1439.
- [12] R. P. S. Mahler, "Toward a theoretical foundation for distributed fusion," in *Distributed Data Fusion for Network-Centric Operations*, D. Hall, C.-Y. Chong, J. Llinas, and M. Liggins, Eds. Boca Raton, FL, USA: CRC Press, 2012, pp. 199–224.
- [13] D. Clark, S. Julier, R. P. S. Mahler, and B. Ristic, "Robust multi-object sensor fusion with unknown correlations," in *Proc. SSPD 2010*, London, UK, Sep. 2010.
- [14] T. Bailey, S. Julier, and G. Agamennoni, "On conservative fusion of information with unknown non-Gaussian dependence," in *Proc. FUSION 2012*, Singapore, Jul. 2012, pp. 1876–1883.
- [15] J. Y. Yu, M. Coates, and M. Rabbat, "Distributed multi-sensor CPHD filter using pairwise gossiping," in *Proc. IEEE ICASSP 2016*, Shanghai, China, Mar. 2016, pp. 3176–3180.
- [16] T. Li, J. M. Corchado, and S. Sun, "On generalized covariance intersection for distributed PHD filtering and a simple but better alternative," in *Proc. FUSION 2017*, Xi'an, China, Jul. 2017.
- [17] R. L. Streit, "Multisensor multitarget intensity filter," in *Proc. FUSION 2008*, Cologne, Germany, Jun. 2008.
- [18] R. P. S. Mahler, *Statistical Multisource-Multitarget Information Fusion*. Norwood, MA, USA: Artech House, 2007.
- [19] S. M. Kay, *Fundamentals of Statistical Signal Processing: Estimation Theory*. Upper Saddle River, NJ, USA: Prentice-Hall, 1993.
- [20] L. Xiao and S. Boyd, "Fast linear iterations for distributed averaging," *Syst. Contr. Lett.*, vol. 53, no. 1, pp. 65–78, 2004.
- [21] L. Xiao, S. Boyd, and S. Lall, "A scheme for robust distributed sensor fusion based on average consensus," in *Proc. IPSN 2005*, Los Angeles, CA, USA, Apr. 2005, pp. 63–70.
- [22] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah, "Randomized gossip algorithms," *IEEE Trans. Inf. Theory*, vol. 52, no. 6, pp. 2508–2530, Jun. 2006.
- [23] A. G. Dimakis, S. Kar, J. M. F. Moura, M. G. Rabbat, and A. Scaglione, "Gossip algorithms for distributed signal processing," *Proc. IEEE*, vol. 98, no. 11, pp. 1847–1864, Nov. 2010.
- [24] T. Li, J. Corchado, and J. Prieto, "Convergence of distributed flooding and its application for distributed Bayesian filtering," *IEEE Trans. Signal Inform. Process. Netw.*, vol. 3, no. 3, pp. 580–592, Sep. 2017.
- [25] F. S. Cattivelli and A. H. Sayed, "Diffusion LMS strategies for distributed estimation," *IEEE Trans. Signal Process.*, vol. 58, no. 3, pp. 1035–1048, Mar. 2010.
- [26] K. Dedecius and P. M. Djurić, "Sequential estimation and diffusion of information over networks: A Bayesian approach with exponential family of distributions," *IEEE Trans. Signal Process.*, vol. 65, no. 7, pp. 1795–1809, Apr. 2017.
- [27] X. R. Li and V. P. Jilkov, "Survey of maneuvering target tracking. Part I. Dynamic models," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 39, no. 4, pp. 1333–1364, Oct. 2003.
- [28] T. Li, J. M. Corchado, S. Sun, and H. Fan, "Multi-EAP: Extended EAP for multi-estimate extraction for SMC-PHD filter," *Chin. J. Aeronaut.*, vol. 30, no. 1, pp. 368–379, 2017.
- [29] D. Schuhmacher, B. T. Vo, and B. N. Vo, "A consistent metric for performance evaluation of multi-object filters," *IEEE Trans. Signal Process.*, vol. 56, no. 8, pp. 3447–3457, Aug. 2008.