Distributed peer-to-peer multitarget tracking with association-based track fusion

G. Battistelli, L. Chisci, C. Fantacci, N. Forti
DINFO, Università di Firenze, Florence, Italy
{giorgio.battistelli, luigi.chisci, claudio.fantacci, nicola.forti}@unifi.it

A. Farina, A. Graziano
Selex ES, Rome, Italy
{alfonso.farina, antonio.graziano}@selex-es.com

***Abstract*—The paper addresses distributed multitarget tracking over a peer-to-peer network of sensors with limited target observability. It is shown how to extend the Cheap Joint Probabilistic Data Association (CJPDA) filter to this setting by devising suitable distributed, consensus-based, procedures for track initiation and track fusion. The resulting, association-based, Distributed-CJPDA (D-CJPDA) tracker is then compared to a previously proposed random set approach for the same problem based on the Gaussian Mixture Cardinalized Probability Hypothesis Density (CPHD) filter.**

I. Introduction

The problem of detecting and tracking moving objects within a certain area of interest, known as multitarget tracking [1]-[7], is of paramount importance in many defence, civilian and security applications. The widespread use of sensor networks for surveillance and monitoring purposes has recently attracted growing interest towards distributed multitarget tracking [4, chapter 8], [8]-[12]. Major attention has been focused on centralized or hierarchical architectures [8], [9], [10] wherein one or multiple fusion nodes receive data from lower-level sensor nodes, combine them and deliver results to the users. In many circumstances, however, it is preferable to adopt the *peer-to-peer (P2P)* paradigm according to which all nodes of the network are at the same level, acting at the same time as sensors, fusion nodes and users. Each node (peer) aims to detect and estimate the states of all targets in the surveillance region by processing local sensor measurements, exchanging data with the neighbors, and fusing local information with information from the neighbors. On one hand, P2P architectures provide significant advantages in terms of scalability, fault tolerance, energy efficiency and prolonged network lifetime but, on the other hand, they pose challenges on target detection and tracking in presence of nodes with limited observability. In fact, a node with limited observability has not sufficient information to initiate and update a target track and must, therefore, rely on cooperation with the neighboring nodes in order to properly accomplish these tasks. A first contribution on distributed P2P multitarget tracking, based on the random set approach, can be found in [11]. Specifically, [11] introduced a distributed P2P random-set tracker, named *Consensus Gaussian Mixture - Cardinalized Probability Hypothesis Density* (CGM-CPHD) filter. Recall that the CPHD filter propagates in time multitarget statistics in terms of both the discrete probability distribution of the number of targets (cardinality distribution) and the continuous probability distribution of such targets over the state space (location distribution), by exploiting a random set multitarget multisensor model which embeds target birth and death as well as missed detections and false alarms. In particular, the GM-CPHD filter represents the location distribution as a Gaussian mixture. In the CGM-CPHD, each node runs a local GM-CPHD filter to update its cardinality and location distributions with the local measurements and then performs consensus (repeated fusion) [13] with the neighbors in order to spread information across the network.

The aim of the present paper is to tackle the distributed P2P multitarget tracking problem from a different, association-based, perspective. In this respect, a new distributed P2P tracker based on *Joint Probabilistic Data Association* (JPDA) is developed and will be referred to hereafter as *Distributed - Cheap JPDA* (D-CJPDA). D-CJPDA employs a CJPDA filter [14] in each peer node for updating the local set of tracks with the measurements available "in situ" and exploits cooperation among neighboring peers in two ways, i.e., (1) to initiate new tracks and (2) to associate and fuse tracks of different peers. More specifically, a distributed procedure is devised for initiating target tracks whenever the individual peers are characterised by limited target observability. The idea is to spatially discretize the surveillance region into a finite number of cells and to associate, to each cell, a target birth probability. Then, such a discrete birth probability distribution is updated in each peer with the locally unassociated measurements and undergoes a consensus procedure with the neighboring peers; target tracks are finally initiated from the peaks of the target birth distribution resulting from consensus. A further key ingredient of D-CJPDA is a distributed association-based track fusion procedure by which each node performs pairwise track-to-track associations with the neighbors followed by fusion of the associated tracks.

The paper is organised as follows. Section II introduces the P2P distributed multitarget tracking problem, the P2P network model and some notation. Section III presents the new, association-based, D-CJPDA distributed P2P multitarget tracker. Section III briefly reviews the random set CGM-CPHD tracker presented in [11]. Section IV provides a comparative performance evaluation of the two trackers. Concluding remarks and perspectives for future work are in section V.

II. Problem setting and notation

This work considers multitarget tracking over a network of heterogeneous and non co-located nodes with sensing, processing and communication capabilities. It is assumed that the network operates in a P2P fashion, i.e. the nodes (peers) have no hierarchical organisation acting at the same time as sensing devices, information processors and information users. More specifically, each peer aims to get situation awareness (i.e., as accurate as possible knowledge at each time of the number, positions and velocities of targets within the surveillance area) by processing the local measurements and the information from the neighbors.

Mathematically, the network will be characterized by a directed communication graph $(\mathcal{N}, \mathcal{A})$ where $\mathcal{N}$ is the set of nodes and $\mathcal{A} \subseteq \mathcal{N} \times \mathcal{N}$ the set of arcs, representing links (connections). In particular, (j, i) belongs to $\mathcal{A}$ whenever node i can receive from node j. For each node $i \in \mathcal{N}$, $\mathcal{N}_i = \{j \in \mathcal{N} : j \neq i \text{ and } (j,i) \in \mathcal{A}\}$ will denote the set of in-neighbors of node i.

For notational simplicity, time dependence of quantities will be omitted throughout the paper. Hence, it will be implicitly assumed that all quantities pertain to the same sampling instant t_k while quantities relative to the subsequent instant t_{k+1} will have a $+$ superscript, e.g. x and x^+ will denote x_k and, respectively, x_{k+1}. The notation $\mathcal{G}(\cdot; \hat{x}, P)$ will denote the Gaussian *probability density function* (PDF) with mean $\hat{x}$ and covariance P.

III. Distributed CJPDA filter

This approach runs a local CJPDA tracker in each node of the network and exploits cooperation with the neighbors for track management. The overall *Distributed-CJPDA* (D-CJPDA) tracker is schematized in the block-diagram of Fig. 1. At each sampling interval, the local tracker of sensor node i inputs the set of measurements Y_i and manages (i.e. initiates, updates, terminates) target tracks as described in the sequel. Let us assume that, at the beginning of the sampling interval, the set of tracks (either preliminary or confirmed) $\mathcal{T}_i$ be available at node i and that each track $j \in \mathcal{T}_i$ be characterized by the filtered state estimate $\hat{x}_{ij}$ and relative covariance P_{ij}.

Track prediction - Exploiting the selected target motion model, each track $(\hat{x}_{ij}, P_{ij})$ is one-step-ahead predicted to get $\left(\hat{x}_{ij}^+, P_{ij}^+\right)$.

Measurement-to-track association - An association gate is set around the predicted measurement of each track $\left(\hat{x}_{ij}^+, P_{ij}^+\right)$ and the subset $Y_{ij} \subseteq Y_i$ of measurements falling within the gate is associated to track j.

Distributed track initiation - A sensible approach to track initiation is to update a prior birth intensity defined over the surveillance area by means of the unassociated observations (which, presumably, provide useful information on new targets) and then initialize new tracks originating from the peaks of such a birth intensity. For each node i, let

$$\overline{Y}_i = Y_i \setminus \bigcup_{j \in \mathcal{T}_i} Y_{ij} \tag{1}$$

denote the set of unassociated measurements. For computational simplicity, a spatially discretized approach [15]-[18] is adopted. To this end, let us assume that the surveillance area Z be partitioned into disjoint cells (bins) as $Z = \cup_{m \in \mathcal{M}} Z_m$ with $\mathcal{M} = \{1, \ldots, M\}$. Let b_{im} denote the birth probability in cell $m \in \mathcal{M}$ computed by node i on the basis of measurements $\overline{Y}_i$ and b_m the prior birth probability in cell m. Then, following the update (correction) step of the bin-occupancy filter [18], such probabilities are updated as follows

$$\begin{aligned} b'_{im} &= (1 - P_{d,m})\, b_m + \sum_{y \in \overline{Y}_i} \frac{P_{d,m}\, \ell_i(z_m, y)\; b_m}{c_m + \sum\limits_{m \in \mathcal{M}} P_{d,m}\, \ell_i(z_m, y)\, b_m} \\ b_{im} &= \frac{b'_{im}}{\left(\sum_{m \in \mathcal{M}} b'_{im}\right)} \end{aligned} \tag{2}$$

where: $z_m = [\xi_m, \eta_m]^\top$ is the position of the center of Z_m; $P_{d,m}$ is the detection probability in cell m; c_m is the clutter probability in cell m; the likelihood $\ell_i(\cdot, \cdot)$ is defined as follows

$$\ell_i(z, y) = \exp\left[-\frac{1}{2}\,(y - h_i(z))^T\, R_i^{-1}\,(y - h_i(z))\right], \tag{3}$$

$h_i(\cdot)$ and R_i being the measurement function and, respectively, measurement noise covariance of position sensor i. Recalling that, in the limit for the cell dimension going to zero, the bin-occupancy filter tends to the PHD filter [18], the quantities b_{im} can be interpreted as a discrete-space location density of newborn targets. Notice that the density is normalized to 1 since, in the proposed algorithm, only the location information is used for track initiation.

Due to the possible presence of sensors with incomplete observability, however, track initiation needs cooperation among nodes. In this respect, the idea is to carry out consensus among the neighbors so as to spread birth probability information across the network. In fact, consensus has emerged as a convenient tool for distributed computation (e.g. averaging, minimization, maximization) over networks [19], [20] and has been widely used in distributed parameter/state estimation algorithms (see [13] and the references therein). In its basic form, a consensus algorithm updates, in each network node, the local information by averaging it with the information coming from the neighbors, i.e., by computing a *regional average*. The same operation is then repeated with the objective that all the network nodes reach an agreement about the information of interest. Convergence results of consensus algorithms, depending on the consensus weights, can be found in [19], [20].

In the present context, the information is represented by the discrete-space location densities b_{im}, $i \in \mathcal{N}$. As discussed in [11], [12], [21], a theoretically-sound way for fusing location densities consists of computing their normalized geometric

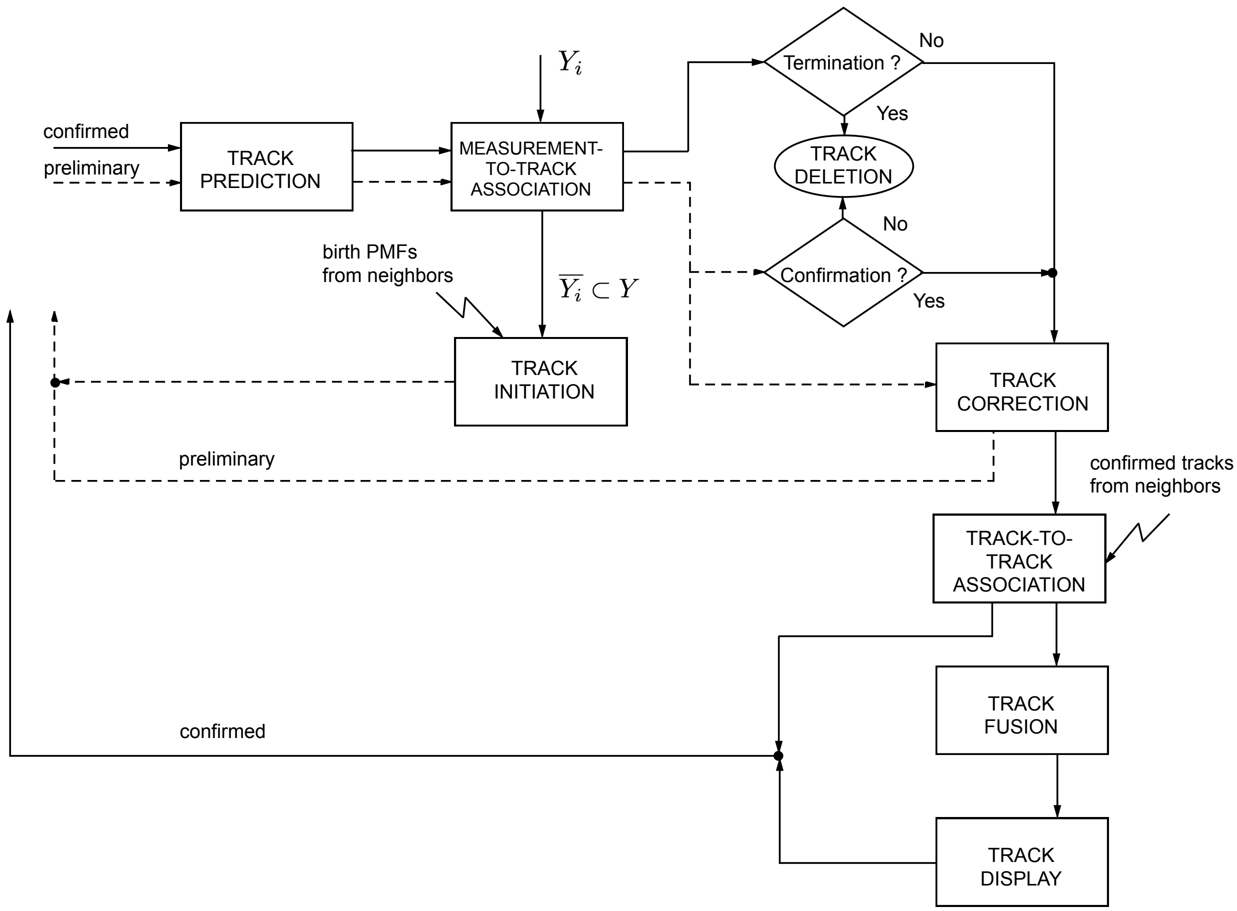

Fig. 1. Block-diagram of the distributed CJPDAF tracker

mean. Accordingly, in the proposed approach, each node i performs a given number of consensus iterations of the form

$$b_{im} \leftarrow b_{im}^{\omega_{ii}} \prod_{k \in \mathcal{N}_i} b_{km}^{\omega_{ik}}, \quad b_{im} \leftarrow b_{im} \left(\sum_{l \in \mathcal{M}} b_{il} \right)^{-1} \tag{4}$$

where the consensus weights satisfy

$$\omega_{ik} \geq 0, \quad \omega_{ii} + \sum_{k \in \mathcal{N}_i} \omega_{ik} = 1 \tag{5}$$

Then, node i initializes preliminary tracks in the local maxima (peaks) of the resulting birth location density $\{b_{im}\}_{m \in \mathcal{M}}$.

The use of a PHD-like representation to aid initiation of tracks using a JPDA-like algorithm has also been studied recently in [22].

Track correction via CJPDA filtering - Exploiting the associated measurements Y_{ij}, each (preliminary or confirmed) track $\left(\hat{x}_{ij}^{+}, P_{ij}^{+}\right)$ is updated by means of a local JPDA filter [3], [23], [24] to get $(\hat{x}_{ij}, P_{ij})$. For computational simplicity, the CJPDA algorithm [14], in which the association probabilities are computed in an approximate way, is actually used.

Track-to-track association - Only the subsets of confirmed tracks $\overline{\mathcal{T}}_i \subseteq \mathcal{T}_i$ undergo track-to-track association among neighboring nodes. To avoid the combinatorial complexity of n-dimensional ($n > 2$) assignment, pairwise associations are considered. Let $k \in \mathcal{N}_i$ be a neighbor of node i, then association between $\overline{\mathcal{T}}_i$ and $\overline{\mathcal{T}}_k$ is carried out by solving the following $2-D$ assignment problem

$$\begin{array}{ll} \min\limits_{a_{jl} \in \{0,1\}} & \sum\limits_{j \in \overline{\mathcal{T}}_i \cup \{0\}} \sum\limits_{l \in \overline{\mathcal{T}}_k \cup \{0\}} a_{jl} \; c_{jl} \\ \text{subject to} & \left\{ \begin{array}{l} \sum\limits_{j \in \overline{\mathcal{T}}_i \cup \{0\}} a_{jl} = 1, \quad \forall l \in \overline{\mathcal{T}}_k \\ \sum\limits_{l \in \overline{\mathcal{T}}_k \cup \{0\}} a_{jl} = 1, \quad \forall j \in \overline{\mathcal{T}}_i \end{array} \right. \end{array} \tag{6}$$

where: the association costs c_{jl} are defined as in [5, p. 631]; a virtual track 0 has been joined to both $\overline{\mathcal{T}}_i$ and $\overline{\mathcal{T}}_k$ in order to allow a track j of node i not to be associated to any track l of node k and viceversa. The optimisation problem (6) can be solved in polynomial time by means of the Hungarian algorithm [25], [26].

Track fusion - Associated tracks need to be fused. This can be done in many different ways, either in one shot at the end of all pairwise track-to-track associations between neighbors or sequentially, i.e. neighbor by neighbor. A sequential approach, detailed in Table I, is adopted in this paper. Track-to-track association between node i and a neighbor $k \in \mathcal{N}_i$ is performed and then the associated tracks are fused via *covariance intersection* [27]. The procedure is repeated with the fused tracks of node i and another neighbor, until all neighbors have been processed. Whenever a track of node i remains unassociated, it is discarded from the subsequent associations-fusions. Finally, only the subset $\widetilde{\mathcal{T}}_i \subseteq \overline{\mathcal{T}}_i$ of tracks associated and fused with all neighbors are actually displayed, whereas the remaining tracks are kept for the subsequent sampling intervals.

While track-to-track association and fusion are performed sequentially, the pairwise combination weights are chosen so that the resulting fused tracks are obtained via covariance intersection from the original local tracks. In particular, it can be seen that the algorithm of Table I results in the following substitutions

$$
\begin{aligned}
P'_{ij} &= \left(\omega_{ii} P_{ij}^{-1} + \sum_{k \in \mathcal{N}_i} \omega_{ik} P_{k\,l_{ik}(j)}^{-1} \right)^{-1} \\
\hat{x}'_{ij} &= P'_{ij} \left(\omega_{ii} P_{ij}^{-1} \hat{x}_{ij} + \sum_{k \in \mathcal{N}_i} \omega_{ik} P_{k\,l_{ik}(j)}^{-1} \hat{x}_{k\,l_{ik}(j)} \right) \\
P_{ij} &= P'_{ij}\,, \quad \hat{x}_{ij} = \hat{x}'_{ij}
\end{aligned} \tag{7}
$$

for any $j \in \widetilde{\mathcal{T}}_i$, where $l_{ik}(j)$ indicates the index of the track of node k associated with track j of node i.

TABLE I
SEQUENTIAL TRACK FUSION

```
N̄_i = N_i;  P^0_ij = P_ij;  x̂^0_ij = x̂_ij;  T̃_i = T̄_i;
P_ij ⟵ ω_ii^{-1} P_ij,  ∀j ∈ T̄_i;
while N̄_i ≠ ∅ do
    choose k ∈ N̄_i
    perform track-to-track association between T̃_i and T̄_k
    for any track j ∈ T̃_i
        find l such that a_jl = 1
        if l > 0 then
            P'_ij = [P_ij^{-1} + ω_ik P_kl^{-1}]^{-1};
            x̂'_ij = P'_ij [P_ij^{-1} x̂_ij + ω_ik P_kl^{-1} x̂_kl];
            P_ij = P'_ij;  x̂_ij = x̂'_ij;
        if l = 0 then
            T̃_i = T̃_i \ {j};
            x̂_ij = x̂^0_ij;
            P_ij = P^0_ij;
    end for
    N̄_i = N̄_i \ {k};
end while
```

Track confirmation and termination - Although also the confirmation of preliminary tracks and the termination of confirmed tracks might benefit from the adoption of distributed cooperative strategies, the currently implemented solution relies on simple local methods. Specifically, M/N logic [1, pp. 203-221] is exploited for track confirmation, while termination of track j in node i is performed whenever at least one of the following conditions occurs: 1) track j has been associated to no measurement, i.e. Y_{ij} has been empty, for a given number of consecutive sampling intervals; 2) the track covariance trace $tr\, P_{ij}$ exceeds a given threshold; 3) track j has not been associated to all neighbors, i.e. $j \in \overline{\mathcal{T}}_i \backslash \widetilde{\mathcal{T}}_i$, for a given number of consecutive sampling intervals.

IV. CONSENSUS CPHD FILTER

CGM-CPHD propagates, in each node i of the network, the target cardinality PMF (Probability Mass Function) $\{p_i(n)\}_{n=0}^{n_{max}}$ and the PHD function

$$
d_i(x) = \overline{n}_i \; s_i(x) \tag{8}
$$

where $\overline{n}_i = \sum_n np(n)$ is the expected number of targets and the PDF $s_i(\cdot)$ is the location density represented as a Gaussian Mixture, i.e.

$$
s_i(x) \;=\; \sum_{j=1}^{M_i} \alpha_{ij}\, \mathcal{G}\left(x; \hat{x}_{ij}, P_{ij}\right), \quad \alpha_{ij} > 0, \;\; \sum_{j=1}^{M_i} \alpha_{ij} = 1 \tag{9}
$$

At each sampling interval, a GM-CPHD filter locally updates $\{p_i(n)\}_{n=0}^{n_{max}}$ and $\{\alpha_{ij}, \hat{x}_{ij}, P_{ij}\}_{j=1}^{M_i}$, exploiting the measurements Y_i and a random set evolution model incorporating target birth and death as well as clutter generation. Then, consensus takes place in each node i involving the subnetwork $\mathcal{N}_i$. Specifically, each node i receives the cardinality PMF $\{p_k(n)\}_{n=0}^{n_{max}}$ and the GM parameters $\{\alpha_{kj}, \hat{x}_{kj}, P_{kj}\}_{j=1}^{M_k}$ of the location density from the neighbors $k \in \mathcal{N}_i$ and then fuses such information [28], repeating the procedure for a given number L of consensus iterations. The theoretical foundations and algorithmic details can be found in [11].

V. SIMULATION EXPERIMENTS

This section develops a performance comparison between two distributed P2P multi target trackers: CGM-CPHD proposed in [11] and briefly reviewed in section III and D-CJPDA proposed in this paper (see section II). To this end, a 2-dimensional (planar) multitarget tracking scenario is considered over a surveillance area of $50 \times 50 [km^2]$, wherein the sensor network of Fig. 2 is deployed. The scenario involves 5 targets as depicted in Fig. 3.

The target state is denoted by $x = \left[\xi,\ \dot{\xi},\ \eta,\ \dot{\eta}\right]^T$ where (ξ, η) and $(\dot{\xi}, \dot{\eta})$ represent the target Cartesian position and, respectively, velocity components. The motion of targets is modeled by the filters according to the nearly-constant velocity model:

$$
x^+ = \begin{bmatrix} 1 & T_s & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & T_s \\ 0 & 0 & 0 & 1 \end{bmatrix} x + w \tag{10}
$$

where w is a zero-mean white noise process with covariance

$$
Q = \sigma_w^2 \begin{bmatrix} \frac{1}{4}T_s^4 & \frac{1}{2}T_s^3 & 0 & 0 \\ \frac{1}{2}T_s^3 & T_s^2 & 0 & 0 \\ 0 & 0 & \frac{1}{4}T_s^4 & \frac{1}{2}T_s^3 \\ 0 & 0 & \frac{1}{2}T_s^3 & T_s^2 \end{bmatrix}, \tag{11}
$$

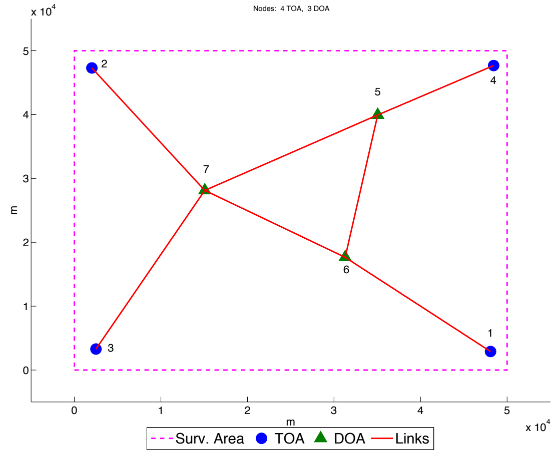

Fig. 2. Network with 7 sensors: 4 TOA and 3 DOA.

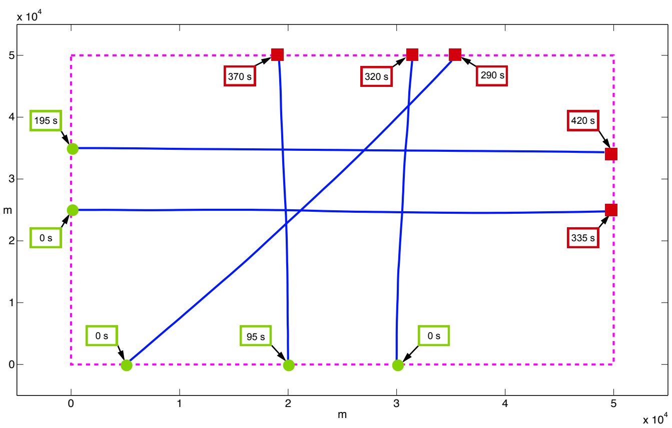

Fig. 3. Target trajectories considered in the simulation experiment. The start/end point for each trajectory is denoted, respectively, by a bullet/square.

$\sigma_w = 2[m/s^2]$ and the sampling interval is $T_s = 5[s]$.

As it can be seen from Fig. 2, the sensor network considered in the simulation consists of 4 *range-only* (TOA = Time Of Arrival) and 3 *bearing-only* (DOA = Direction Of Arrival) sensors characterized by the following measurement functions:

$$h_i(x) = \begin{cases} \angle[(\xi - \xi_i) + j\,(\eta - \eta_i)], & \text{DOA} \\ \sqrt{(\xi - \xi_i)^2 + (\eta - \eta_i)^2}, & \text{TOA} \end{cases} \quad (12)$$

where (ξ^i, η^i) represents the known position of sensor i. The standard deviation of DOA and TOA measurement noises are respectively $\sigma_{DOA} = 1[°]$ and $\sigma_{TOA} = 100[m]$. Because of the non linearity of the aforementioned sensors, the *Unscented Kalman Filter* (UKF) [29], [30] is exploited in each sensor to update means and covariances, of either the Gaussian components in the CGM-CPHD or the target tracks in the D-CJPDA .

Clutter is modeled as a Poisson Process with parameter $\lambda_c = 5$ and uniform spatial distribution over the surveillance area; the probability of target detection is $P_d = 0.9$.

In the considered scenario, targets travel through the surveillance area with no prior information for target birth locations. The following parameters (see [11] for their definitions) have been selected for CGM-CPHD: number of Gaussian components of the birth intensity $N_b = 40$ and weight of such components $\alpha = 1.5\ 10^{-3}$; target survival probability $P_s = 0.99$; maximum number of Gaussian components $N_{max} = 25$; merging threshold $\gamma_m = 9$; truncation threshold $\gamma_t = 10^{-4}$; extraction threshold $\gamma_e = 0.5$

Conversely, for D-CJPDA a 50×50 grid of 1×1 $[km^2]$ cells has been adopted for target initiation. Further the following parameters have been chosen: $M = 4$ and $N = 7$ for the M/N track confirmation logic; validation gate threshold $\gamma_v = 8$; gate threshold for the track-to-track association $\gamma_f = 15$. The consensus weights used in (4) and (7) have been chosen as uniform, i.e. $\omega_{ik} = (|\mathcal{N}_i| + 1)^{-1}$ for any $k \in \mathcal{N}_i \cup \{i\}$. Multitarget tracking performance is evaluated in terms of estimated number of targets as well as the OSPA (Optimal SubPattern Analysis) metric [31].

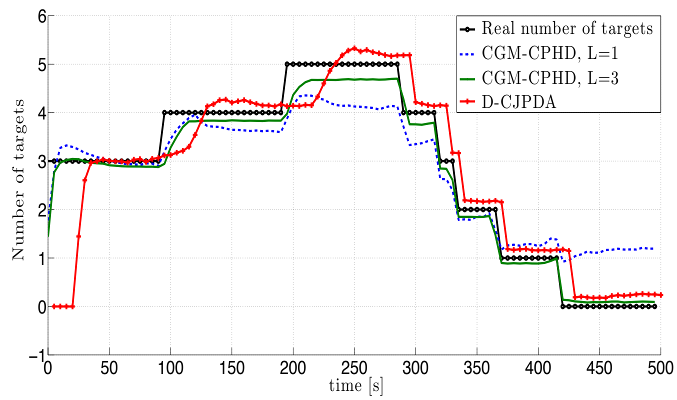

Fig. 4. Cardinality statistics for D-CJPDA, CGM-CPHD with $L = 1$ and CGM-CPHD with $L = 3$ consensus steps.

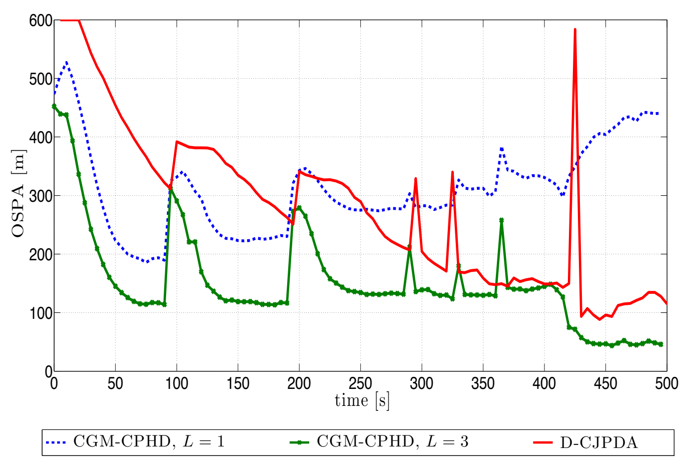

Fig. 5. Performance comparison, using OSPA, between D-CJPDA and CGM-CPHD respectively with $L = 1$ and $L = 3$ consensus steps.

The reported metrics are averaged over all nodes and $N_{mc} = 200$ Monte Carlo trials for the same target trajectories but

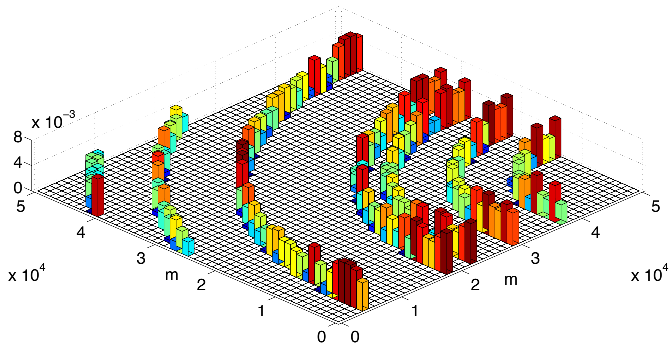

Fig. 6. Birth probabilities at the initial time instant in node 1 (TOA), before carrying out consensus among the neighbors.

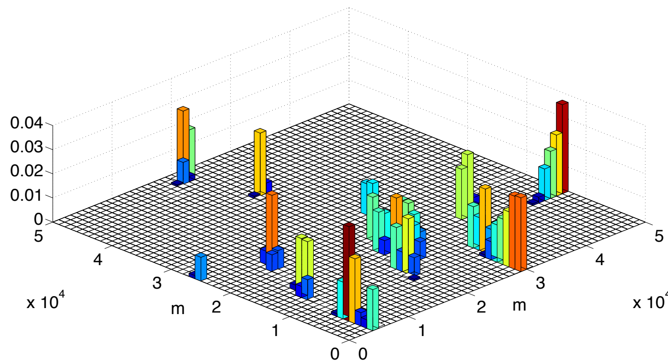

Fig. 7. Birth probabilities after the first consensus step.

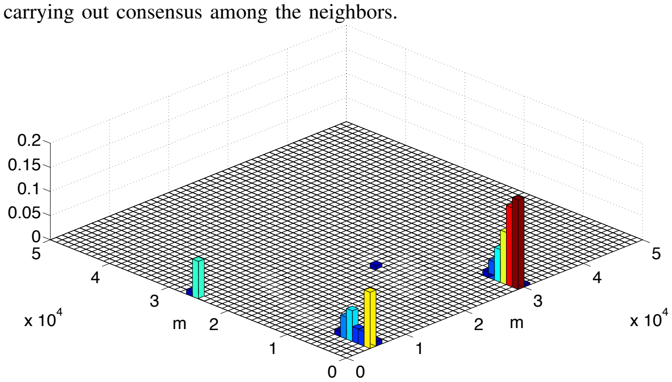

Fig. 8. Birth probabilities after the second consensus step.

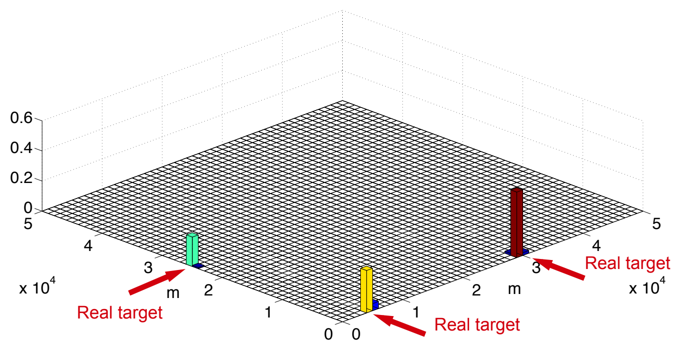

Fig. 9. Birth probabilities after the third consensus step. Three new tracks will originate from the three local maxima corresponding to real targets entering the surveillance region.

different, independently generated, clutter and measurement noise realizations. The duration of each simulation trial is fixed to $500[s]$ (100 samples). Figs. 4 and 5 compare the estimated number of targets and, respectively, OSPA (with Euclidean distance, $p = 2$, and cutoff parameter $c = 600$) for both D-CJPDA and CGM-CPHD trackers. Please notice that, for D-CJPDA, the estimated cardinality plotted in Fig. 4 is actually the number of tracks sent to track display (Fig. 1), i.e. the confirmed tracks that are associated to all neighbors. From the examination of Figs. 4 and 5, the following considerations are in order.

- D-CJPDA tends to overestimate the number of targets, i.e. to create more easily clutter-originated false tracks, while CGM-CPHD is more prone to underestimation of the target number (see Fig. 4).
- D-CJPDA is less reactive than CGM-CPHD to display new-born targets (see Fig. 4). This is due, besides the intrinsic delay of the M/N confirmation logic, also to the fact that, in the current implementation, each node displays only the confirmed tracks that are associated to all neighbors.
- In terms of OSPA (see Fig. 5) CGM-CPHD, for $L = 3$ consensus steps, provides the best performance but also D-CJPDA exhibits a good behaviour (much better than CGM-CPHD with a single consensus step), except for the unavoidable peaks due to the changes of cardinality.

To demonstrate the effectiveness of the proposed distributed track initiation procedure, Figs. 6-9 display the birth probabilities b_{1m} at node 1 (see Fig. 2) at the initial time instant (whenever there are three targets in the surveillance area) respectively before consensus and after the first, second and third consensus iteration. As it can be seen, at the end of the consensus procedure the three simulated targets are detected in the cells corresponding to their true initial positions.

A theoretical evaluation of the computational complexity of D-CJPDA and CGM-CPHD trackers is difficult to carry out. The considered implementations of both trackers have been run on the same computer and simulation scenario and it has been found that D-CJPDA required an average computation time roughly $2 \div 3$ times faster than CGM-CPHD.

VI. Conclusions

The paper has presented a novel approach, named Distributed-CJPDA (D-CJPDA), to distributed multitarget tracking over a peer-to-peer network of sensors with limited target observability. D-CJPDA relies on local Cheap Joint Probabilistic Data Association filters running in each peer node, as well as on a consensus procedure for distributed track initiation and association-based track fusion. Simulation experiments have demonstrated good performance of D-CJPDA, also in comparison with a previously introduced distributed peer-to-peer multitarget tracker based on random set theory [11]. Future work will try to improve D-CJPDA in several ways, e.g. considering distributed track confirmation/termination and adopting a majority consensus strategy (instead of the currently employed strategy that requires unanimity) in the association-based track fusion.

REFERENCES

[1] A. Farina and F. A. Studer, *Radar Data Processing, vol. I: Introduction and Tracking*. Researches Studies Press, England, John Wiley & Sons, USA, 1985.

[2] A. Farina and F. A. Studer, *Radar Data Processing, vol. II: Advanced Topics and Applications*. Researches Studies Press, England, John Wiley & Sons, USA, 1986.

[3] Y. Bar-Shalom and T. E. Fortmann, *Tracking and Data Association*. Academic Press, San Diego, CA, 1988.

[4] Y. Bar-Shalom and X. R. Li, *Multitarget-multisensor tracking: principles and techniques*. YBS Publishing, Storrs, CT, 1995.

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

[6] Y. Bar-Shalom, X. R. Li and T. Kirubarajan, *Estimation with Applications to Tracking and Navigation: Theory Algorithms and Software*. Wiley, 2004.

[7] B. Ristic, *Particle filters for random set models*. Springer, 2012.

[8] K. C. Chang, C. Y. Chong and Y. Bar-Shalom, Joint probabilistic data association in distributed sensor networks. *IEEE Trans. on Automatic Control*, vol. 31, no. 10, pp. 889–897, 1986.

[9] C. Y. Chong, S. Mori and K. C. Chang, Distributed multitarget multisensor tracking. In: *Multitarget Multisensor Tracking: Advanced Applications*, pp. 247–295, Y. Bar-Shalom Ed., Artech House, Norwood, MA, 1990.

[10] M. E. Liggins II, C. Y. Chong, I. Kadar, M. G. Alford, V. Vannicola and S. Thomopoulo, Distributed fusion architectures and algorithms for target tracking. *Proceedings of the IEEE*, vol. 85, no. 1, pp. 95–107, 1997.

[11] G. Battistelli, L. Chisci, C. Fantacci, A. Farina and A. Graziano, Consensus CPHD filter for distributed multitarget tracking. *IEEE Journal of Selected Topics in Signal Processing*, vol. 7, no. 3, pp. 508–520, 2013.

[12] M. Uney, D. E. Clark and S. J. Julier, Distributed fusion of PHD filters via exponential mixture densities. *IEEE Journal of Selected Topics in Signal Processing*, vol. 7, no. 3, pp. 521–531, 2013.

[13] G. Battistelli and L. Chisci, Kullback-Leibler average, consensus on probability densities, and distributed state estimation with guaranteed stability. *Automatica*, vol. 50, pp. 707–718, 2014.

[14] R. J. Fitzgerald, Development of practical PDA logic for multitarget tracking by microprocessor. *Proc. American Control Conference*, pp. 889–898, 1986.

[15] H. P. Moravec and A. Elfes, High resolution maps from wide angle sonar. *Proc. IEEE Int. Conf. on Robotics and Automation*, vol. 2, pp. 116–121, 1985.

[16] A. Birk and S. Carpin, Merging occupancy grid maps from multiple robots. *Proceedings of the IEEE*, vol. 94, no. 7, pp. 1384–1397, 2006.

[17] D. Elsaesser, Sensor data fusion using a probability density grid. *Proc. 10th International Conference on Information Fusion*, pp. 1–8, 2007.

[18] O. Erdinc, P. Willett and Y. Bar-Shalom, The bin-occupancy filter and its connection to the PHD filters. *IEEE Trans. on Signal Processing*, vol. 57, no. 11, pp. 4232–4246, 2009.

[19] R. Olfati-Saber, J. A. Fax and R. M. Murray, Consensus and cooperation in networked multi-agent systems. *Proceedings of the IEEE*, vol. 95, no. 1, pp. 49–54, 2007.

[20] L. Xiao, S. Boyd and S. Lall, A scheme for robust distributed sensor fusion based on average consensus. *Proc. 4th Int. Symposium on Information Processing in Sensor Networks (IPSN)*, pp. 63–70, 2005.

[21] D. E. Clark, S. J. Julier, R. Mahler and B. Ristic, Robust multi-object sensor fusion with unknown correlations. *Sensor Signal Processing for Defence (SSPD 2010)*, pp.1–5, 2010.

[22] J. L. Williams, Hybrid Poisson and multi-Bernoulli filters. *Proc. 15th Int. Conf. on Information Fusion*, pp. 1103–1100, Singapore, 2012.

[23] Y. Bar-Shalom and E. Tse, Tracking in a cluttered environment with probabilistic data association. *Automatica*, vol. 11, pp. 451–460, 1975.

[24] Y. Bar-Shalom, F. Daum and J. Huang, The probabilistic data association filter. *IEEE Control Systems Magazine*, vol. 29, no. 6, pp. 82–100, 2009.

[25] H. W. Kuhn, The Hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, vol. 2, pp. 83–97, 1955.

[26] J. Munkres, Algorithms for the assignment and transportation problems. *Journal of the Society for Industrial and Applied Mathematics*, pp. 32–38, 1957.

[27] S. J. Julier and J. K. Uhlmann, A non-divergent estimation algorithm in the presence of unknown correlations. *Proc. American Control Conference*, vol. 4, pp. 2369–2373, 1997.

[28] R. Mahler, Optimal/robust distributed data fusion: a unified approach. *Proc. SPIE 4052, Signal Processing, Sensor Fusion, and Target Recognition IX*, pp. 128–138, 2000.

[29] S. J. Julier and J. K. Uhlmann, Unscented filtering and nonlinear estimation. *Proceedings of the IEEE*, vol. 92, pp. 401–422, 2004.

[30] B. Ristic, S. Arulampalam, and N. Gordon, *Beyond the Kalman filter*. Artech House, 2004.

[31] D. Schuhmacher, B. T. Vo, and B. N. Vo, A consistent metric for performance evaluation of multi-object filters. *IEEE Trans. on Signal Processing*, vol. 56, no. 8, pp. 3447–3457, 2008.