Consensus Labeled Random Finite Set Filtering for Distributed Multi-Object Tracking

Claudio Fantacci, Ba-Ngu Vo, Ba-Tuong Vo, Giorgio Battistelli and Luigi Chisci

arXiv:1501.01579v2 [cs.SY] 9 Jun 2016

Abstract

This paper addresses distributed multi-object tracking over a network of heterogeneous and geographically dispersed nodes with sensing, communication and processing capabilities. The main contribution is an approach to distributed multi-object estimation based on labeled Random Finite Sets (RFSs) and dynamic Bayesian inference, which enables the development of two novel consensus tracking filters, namely a Consensus Marginalized δ-Generalized Labeled Multi-Bernoulli and Consensus Labeled Multi-Bernoulli tracking filter. The proposed algorithms provide fully distributed, scalable and computationally efficient solutions for multi-object tracking. Simulation experiments via Gaussian mixture implementations confirm the effectiveness of the proposed approach on challenging scenarios.

Index Terms

RFS, FISST, labeled multi-object Bayes filter, multi-object tracking, sensor networks, consensus.

I. Introduction

Multi-Object Tracking (MOT) involves the on-line estimation of an unknown and time-varying number of objects and their individual trajectories from sensor data [1]–[8]. In a multiple object scenario, the sensor observations are affected by misdetection (e.g., occlusions, low radar cross section, etc.) and false alarms (e.g., observations from the environment, clutter, etc.), which is further compounded by association uncertainty, i.e. it is not known which object generated which measurement. The key challenges in MOT

Claudio Fantacci was with the Dipartimento di Ingegneria dell'Informazione, Università di Firenze, Florence, 50139, Italy, on leave at the Curtin University, Bentley, WA, 6102, Australia, in the period January-July 2014. He is currently with the Istituto Italiano di Tecnologia (IIT), Genoa, Italy (e-mail: claudio.fantacci@unifi.it, claudio.fantacci@iit.it).

Ba-Ngu Vo and Ba-Tuong Vo are with the Department of Electrical and Computer Engineering, Curtin University, Bentley, WA, 6102, Australia (e-mail: ba-ngu.vo@curtin.edu.au, ba-tuong@curtin.edu.au).

Giorgio Battistelli and Luigi Chisci are with the Dipartimento di Ingegneria dell'Informazione, Università di Firenze, Florence, 50139, Italy (e-mail: giorgio.battistelli@unifi.it, luigi.chisci@unifi.it).

include *detection uncertainty*, *clutter*, and *data association uncertainty*. Numerous multi-object tracking algorithms have been developed in the literature and most of these fall under the three major paradigms of *Multiple Hypothesis Tracking* (MHT) [6], [9], *Joint Probabilistic Data Association* (JPDA) [4], and *Random Finite Set* (RFS) [7].

Recent advances in wireless sensor technology inspired the development of large sensor networks consisting of radio-interconnected nodes (or agents) with sensing, communication and processing capabilities [10]. The main goal of such a net-centric sensing paradigm is to provide a more complete picture of the environment by combining information from many individual nodes (usually with limited observability) using a suitable *information fusion* procedure, in a way that is *scalable* (with the number of nodes), *flexible* and *reliable* (i.e. *resilient* to failures) [10]. Reaping the benefits of a sensor network calls for distributed architectures and algorithms in which individual agents can operate with neither central fusion node nor knowledge of the information flow in the network [11].

The wide applicability of MOT together with the emergence of net-centric sensing motivate the investigation of *Distributed Multi-Object Tracking* (DMOT). Scalability with respect to network size, lack of a fusion center as well as knowledge of the network topology call for a *consensus* approach to achieve a collective information fusion over the network [11]–[20]. In fact, consensus has recently emerged as a powerful tool for distributed computation over networks [11], [12], including parameter/state estimation [13]–[20]. Furthermore, a robust (possibly suboptimal) information fusion procedure is needed to combat the data incest problem that causes *double counting* of information. To this end, *Chernoff fusion* [21], [22], also known as *Generalized Covariance Intersection* [23], [24] (that encompasses *Covariance Intersection* [25], [26]) or *Kullback-Leibler average* [20], [27], is adopted to fuse multi-object densities computed by various nodes of the network. Furthermore, it was proven in [26] for the single-object case, and subsequently in [28] for the multi-object case, that Chernoff fusion is inherently immune to the double counting of information, thereby justifying its use in a distributed setting wherein the nodes operate without knowledge about their common information.

While the challenges in MOT are further compounded in a distributed architecture, the notion of *multi-object probability density* in the RFS formulation enables consensus for distributed state estimation to be directly applied to multi-object systems [27]–[31]. Indeed, a robust and tractable multi-object fusion solution based on Kullback-Leibler averaging, together with the *Consensus Cardinalized Probability Hypothesis Density* (CPHD) filter have been proposed in [27]. However, this RFS-based filtering solution does not provide estimates of the object trajectories and suffers from the so-called "*spooky effect*" [32]. Note that one of the original intents of the RFS formulation is to propagate the distribution of the set of tracks via the use of labels, see [5, p. 135, pp. 196-197], [7, p. 506]. However, this capability was

overshadowed by the popularity of unlabeled RFS-based filters such as PHD, CPHD, and multi-Bernoulli [7], [33]–[36].

This paper proposes the first consensus DMOT algorithms based on the recently introduced labeled RFS formulation [37]. This formulation admits a tractable analytical MOT solution called the *δ-Generalized Labeled Multi-Bernoulli* (δ-GLMB) filter [38] that does not suffer from the "*spooky effect*", and more importantly, outputs trajectories of objects. Furthermore, efficient approximations that preserve key summary statistics such as the *Marginalized δ-GLMB* (Mδ-GLMB) and the *Labeled Multi-Bernoulli* (LMB) filters have also been developed [39], [40]. In this paper, it is shown that the Mδ-GLMB and LMB densities are algebraically closed under Kullback-Leibler averaging, and novel consensus DMOT Mδ-GLMB and LMB filters are developed.

The rest of the paper is organized as follows. Section II presents notation, the network model, and background on Bayesian filtering, RFSs, and distributed estimation. Section III presents the Kullback-Leibler average based fusion rules for Mδ-GLMB and LMB densities. Section IV describes the multi-object Bayesian recursion with labeled RFSs and presents the novel *Consensus Mδ-GLMB* and *Consensus LMB* filters with *Gaussian Mixture* (GM) implementation. Section V provides a performance evaluation of the proposed DMOT filters via simulated case studies. Concluding remarks and perspectives for future work are given in Section VI.

II. Background and problem formulation

A. Notation

Throughout the paper, we use the standard inner product notation $\langle f, g \rangle \triangleq \int f(x)\, g(x) dx$, and the multi-object exponential notation $h^X \triangleq \prod_{x \in X} h(x)$, where h is a real-valued function, with $h^{\varnothing} = 1$ by convention [7]. The cardinality (or number of elements) of a finite set X is denoted by $|X|$. Given a set S, $1_S(\cdot)$ denotes the indicator function of S, $\mathcal{F}(S)$ the class of finite subsets of S, and S^i the ith-fold Cartesian product of S with the convention $S^0 = \{\varnothing\}$. We introduce a generalization of the Kronecker delta that takes arguments such as sets, vectors, etc., i.e.

$$\delta_Y(X) \triangleq \begin{cases} 1, \text{ if } X = Y \\ 0, \text{ otherwise} \end{cases},$$

A Gaussian *Probability Density Function* (PDF) with mean μ and covariance Σ is denoted by $\mathcal{N}(\cdot; \mu, \Sigma)$. Vectors are represented by lowercase letters, e.g. x, $\mathbf{x}$, while finite sets are represented by uppercase letters, e.g. X, $\mathbf{X}$; spaces are represented by blackboard bold letters e.g. $\mathbb{X}$, $\mathbb{Z}$, $\mathbb{L}$.

Given PDFs p, q and a scalar $\alpha > 0$, the information fusion $\oplus$ and weighting operators $\odot$ [20], [27], [41] are defined as:

$$(p \oplus q)(x) \triangleq \frac{p(x)\, q(x)}{\langle p, q \rangle}, \tag{1}$$

$$(\alpha \odot p)(x) \triangleq \frac{[p(x)]^{\alpha}}{\langle p^{\alpha}, 1 \rangle}. \tag{2}$$

For any PDFs p, q, h, and positive scalars α, β, the fusion and weighting operators satisfy:

P.A	$(p \oplus q) \oplus h$	$=$	$p \oplus (q \oplus h) = p \oplus q \oplus h$
P.B	$p \oplus q$	$=$	$q \oplus p$
P.C	$(\alpha\, \beta) \odot p$	$=$	$\alpha \odot (\beta \odot p)$
P.D	$1 \odot p$	$=$	p
P.E	$\alpha \odot (p \oplus q)$	$=$	$(\alpha \odot p) \oplus (\alpha \odot q)$
P.F	$(\alpha + \beta) \odot p$	$=$	$(\alpha \odot p) \oplus (\beta \odot q)$

B. Network model

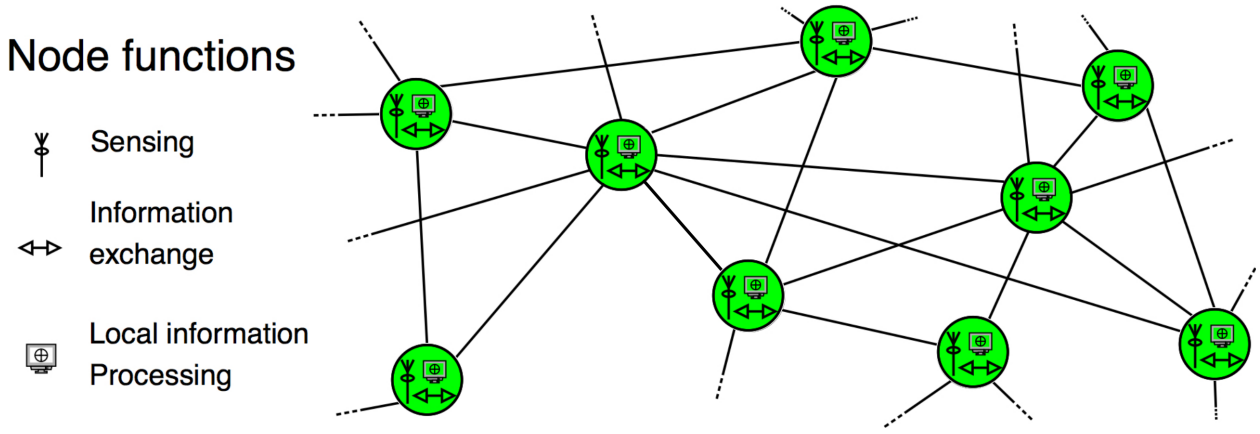

Fig. 1. Network model

The network considered in this work consists of heterogeneous and geographically dispersed nodes having processing, communication and sensing capabilities as depicted in Fig. 1. From a mathematical viewpoint, the network is described by a directed graph $\mathcal{G} = (\mathcal{N}, \mathcal{A})$ where $\mathcal{N}$ is the set of nodes and $\mathcal{A} \subseteq \mathcal{N} \times \mathcal{N}$ the set of arcs, representing *links* (or *connections*). In particular, $(i, j) \in \mathcal{A}$ if node j can receive data from node i. For each node $j \in \mathcal{N}$, $\mathcal{N}^{(j)} \triangleq \{i \in \mathcal{N} : (i, j) \in \mathcal{A}\}$ denotes the set of in-neighbours (including j itself), i.e. the set of nodes from which node j can receive data.

Each node performs local computations, exchanges data with the neighbors and gathers measurements (e.g., angles, distances, Doppler shifts, etc.) of objects present in the *surrounding environment* (or

surveillance area). The network of interest has no *central fusion node* and its agents operate without knowledge of the network topology.

We are interested in networked estimation algorithms that are scalable with respect to network size, and permit each node to operate without knowledge of the dependence between its own information and the information from other nodes.

C. Distributed Single-Object Filtering and Fusion

For single-object filtering, the problem of propagating information throughout a sensor network $(\mathcal{N}, \mathcal{A})$ with neither central fusion node nor knowledge of the network topology can be formalized as follows.

The system model is described by the following Markov transition density and measurement likelihood functions

$$f_{k|k-1}(x_k|x_{k-1}) \tag{3}$$

$$g_k^{(i)}(z_k^{(i)}|x_k)\,, i \in \mathcal{N} \tag{4}$$

The measurement at time k is a vector $z_k = \left(z_k^{(1)}, \ldots, z_k^{(|\mathcal{N}|)}\right)$ of measurements from all $|\mathcal{N}|$ sensors, which are assumed to be conditionally independent given the state. Hence the likelihood function of the measurement z_k is given by

$$g_k(z_k|x_k) = \prod_{i \in \mathcal{N}} g_k^{(i)}(z_k^{(i)}|x_k)\,. \tag{5}$$

Let $p_{k|k-1}(\cdot|z_{1:k-1})$ denote the prediction density of the state at time k given $z_{1:k-1} \triangleq (z_1, \ldots, z_{k-1})$, and similarly $p_k(\cdot|z_{1:k})$ the posterior density of the state at time k given $z_{1:k} \triangleq (z_1, \ldots, z_k)$. For simplicity we omit the dependence on the measurements and write the prediction and posterior densities respectively as $p_{k|k-1}$ and p_k.

In a *centralized* setting, i.e. when the central node has access to all measurements, the solution of the state estimation problem is given by the Bayesian filtering recursion starting from a suitable initial prior p_0:

$$p_{k|k-1}(x_k) = \left\langle f_{k|k-1}(x_k|\,\cdot\,)\,, p_{k-1} \right\rangle\,, \tag{6}$$

$$p_k(x_k) = \left(g_k(z_k|\,\cdot\,) \oplus p_{k|k-1}\right)(x_k)\,. \tag{7}$$

On the other hand, in a *distributed* setting each agent $i \in \mathcal{N}$ updates its own posterior density $p_k^{(i)}$ by appropriately fusing the available information provided by the subnetwork $\mathcal{N}^{(i)}$ (including node i). Thus, central to networked estimation is the capability to fuse the posterior densities provided by different nodes in a mathematically consistent manner. In this respect, the information-theoretic notion of *Kullback-Leibler Average* (KLA) provides a consistent way of fusing PDFs [20].

Given the PDFs $p^{(i)}$, $i \in \mathcal{I}$, and *normalized non-negative weights* $\omega^{(i)}$ (i.e. non-negative weights that sum up to 1), $i \in \mathcal{I}$, the weighted *Kullback-Leibler Average* (KLA) $\overline{p}$ is defined as

$$\overline{p} = \arg\min_{p} \sum_{i \in \mathcal{I}} \omega^{(i)} D_{KL}\left(p \parallel p^{(i)}\right) \tag{8}$$

where

$$D_{KL}\left(p \parallel p^{(i)}\right) = \int p\left(x\right) \log\left(\frac{p(x)}{p^{(i)}(x)}\right) dx \tag{9}$$

is the *Kullback-Leibler Divergence* (KLD) of $p^{(i)}$ from p. In [20] it is shown that the weighted KLA in (8) is the *normalized weighted geometric mean* of the PDFs, i.e.

$$\overline{p}(x) = \frac{\prod_{i \in \mathcal{I}} \left[p^{(i)}(x)\right]^{\omega^{(i)}}}{\int \prod_{i \in \mathcal{I}} \left[p^{(i)}(x)\right]^{\omega^{(i)}} dx} \triangleq \bigoplus_{i \in \mathcal{I}} \left(\omega^{(i)} \odot p^{(i)}\right)(x)\,. \tag{10}$$

Indeed, (10) defines the well-known Chernoff fusion rule [21], [22]. Note that in the *unweighted KLA* $\omega^{(i)} = 1/|\mathcal{I}|$, i.e.

$$\overline{p} = \bigoplus_{i \in \mathcal{I}} \frac{1}{|\mathcal{I}|} \odot p^{(i)}. \tag{11}$$

Remark 1. The weighted KLA of Gaussians is also Gaussian [20]. More precisely, let $(\Phi, q) \triangleq (\Sigma^{-1}, \Sigma^{-1}\mu)$ denote the *information (matrix-vector) pair* associated with $\mathcal{N}\left(\cdot; \mu, \Sigma\right)$, then the *information pair* $(\overline{\Phi}, \overline{q})$ of the KLA $\overline{p}(\cdot) = \mathcal{N}\left(\cdot; \overline{\mu}, \overline{\Sigma}\right)$ is the *weighted arithmetic mean* of the information pairs $(\Phi^{(i)}, q^{(i)})$ of $p^{(i)}(\cdot) = \mathcal{N}\left(\cdot; \mu^{(i)}, \Sigma^{(i)}\right)$. This is indeed the well-known *Covariance Intersection* fusion rule [25].

Having reviewed the fusion of PDFs via KLA, we next outline distributed computation of the KLA via consensus.

D. Consensus on PDFs

The idea behind consensus is to reach a *collective agreement* (over the entire network), by allowing each node $i \in \mathcal{N}$ to iteratively update and pass its local information to neighbouring nodes [11]. Such repeated local operations provide a mechanism for propagating information throughout the whole network. In the context of this paper, consensus is used (at each time step k) to perform distributed computation of the collective unweighted KLA of the posterior densities $p_k^{(i)}$ over all nodes $i \in \mathcal{N}$.

Given the consensus weights $\omega^{(i,j)} \in [0, 1]$ relating agent i to its in-neighbour nodes $j \in \mathcal{N}^{(i)}$, satisfying $\sum_{j \in \mathcal{N}^{(i)}} \omega^{(i,j)} = 1$, suppose that, at time k, each agent i starts with the posterior $p_k^{(i)}$ as the initial iterate $p_{k,0}^{(i)}$, and computes the nth consensus iterate by

$$p_{k,n}^{(i)} = \bigoplus_{j \in \mathcal{N}^{(i)}} \left(\omega^{(i,j)} \odot p_{k,n-1}^{(j)}\right) \tag{12}$$

Then, using the properties of the operators $\oplus$ and $\odot$, it can be shown that [20]

$$p_{k,n}^{(i)} = \bigoplus_{j\in\mathcal{N}} \left(\omega_n^{(i,j)} \odot p_k^{(j)}\right) \tag{13}$$

where $\omega_n^{(i,j)}$ is the (i,j)-th entry of the square matrix Ω^n, and Ω is the consensus matrix with (i,j)-th entry given by $\omega^{(i,j)}1_{\mathcal{N}^{(i)}}(j)$ (it is understood that $p_k^{(j)}$ is omitted from the fusion whenever $\omega_n^{(i,j)}=0$). Notice that (13) expresses the local PDF in each node i at consensus iteration n as a weighted geometric mean of the initial local PDFs of all nodes. More importantly, it was shown in [11], [12] that if the consensus matrix Ω is primitive (i.e. with all non-negative entries and such that there exists an integer m such that Ω^m has all positive entries) and doubly stochastic (all rows and columns sum up to 1), then for any $i,j\in\mathcal{N}$

$$\lim_{n\to\infty} \omega_n^{(i,j)} = \frac{1}{|\mathcal{N}|}. \tag{14}$$

In other words, if the consensus matrix is primitive and doubly stochastic, then the consensus iterate of each node approaches the collective unweighted KLA of the posterior densities over the entire network as the number of consensus steps tends to infinity [16], [20].

A necessary condition for Ω to be primitive [16] is that the associated network $\mathcal{G}$ be strongly connected, i.e. for any pair of nodes i, $j\in\mathcal{N}$ there exists a directed path from i to j and vice versa. This condition is also sufficient when $\omega^{(i,j)}>0$ for all $i\in\mathcal{N}$ and $j\in\mathcal{N}^{(i)}$. Further, when $\mathcal{G}$ is undirected (i.e. whenever node i receives information from node j, it also sends information to j), choosing the *Metropolis weights*

$$\omega^{(i,j)} = \begin{cases} \dfrac{1}{1+\max\{|\mathcal{N}^{(i)}|,|\mathcal{N}^{(j)}|\}}, & i\in\mathcal{N},\ j\in\mathcal{N}^{(i)}\backslash\{i\} \\ 1-\sum_{j\in\mathcal{N}^{(i)}\backslash\{i\}}\omega^{(i,j)}, & i\in\mathcal{N},\ j=i \end{cases} \tag{15}$$

ensures that Ω is also doubly stochastic [12], [16].

In most tracking applications, the number of objects is unknown and varies with time, while measurements are subjected to misdetection, false alarms and association uncertainty. This more general setting can be conveniently addressed by a rigorous mathematical framework for dealing with multiple objects. Such a framework is reviewed next, followed by the extension of the consensus methodology to the multi-object realm.

E. Labeled Random Finite Sets

The RFS formulation of MOT provides the notion of *multi-object probability density* (for an unknown number of objects) [42] that conceptually allows direct extension of the consensus methodology to multi-object systems. Such a notion of multi-object probability density is not available in the MHT or JPDA approaches [1], [2], [4], [6], [9].

From a Bayesian estimation viewpoint the multi-object state is naturally represented as a finite set, and subsequently modeled as an RFS [34]. In this paper, unless otherwise stated we use the *Finite Set STatistics* (FISST) notion of integration/density to characterize RFSs [7]. While not a probability density [7], the FISST density is equivalent to a probability density relative to an unnormalized distribution of a Poisson RFS [42].

Let $\mathbb{L}$ be a discrete space, and $\mathcal{L}: \mathbb{X}\times\mathbb{L} \to \mathbb{L}$ be the projection defined by $\mathcal{L}((x,\ell)) = \ell$. Then $\mathcal{L}(\mathbf{x})$ is called the label of the point $\mathbf{x} \in \mathbb{X}\times\mathbb{L}$, and a finite subset $\mathbf{X}$ of $\mathbb{X}\times\mathbb{L}$ is said to have *distinct labels* if and only if $\mathbf{X}$ and its labels $\mathcal{L}(\mathbf{X}) = \{\mathcal{L}(\mathbf{x}) : \mathbf{x} \in \mathbf{X}\}$ have the same cardinality. We define the *distinct label indicator* of $\mathbf{X}$ as $\Delta(\mathbf{X}) \triangleq \delta_{|\mathbf{X}|}(|\mathcal{L}(\mathbf{X})|)$.

A *labeled RFS* is an RFS over $\mathbb{X}\times\mathbb{L}$ such that each realization has distinct labels. These distinct labels provide the means to identify trajectories or tracks of individual objects since a track is a time-sequence of states with the same label [37]. The distinct label property ensures that at any time no two points can share the same label, and hence no two trajectories can share any common point in the extended space $\mathbb{X}\times\mathbb{L}$. Hereinafter, symbols for labeled states and their distributions are bolded to distinguish them from unlabeled ones, e.g. $\mathbf{x}$, $\mathbf{X}$, $\boldsymbol{\pi}$.

1) Generalized Labeled Multi-Bernoulli (GLMB): A *GLMB* [37] is a labeled RFS with state space $\mathbb{X}$ and (discrete) label space $\mathbb{L}$ distributed according to

$$\boldsymbol{\pi}(\mathbf{X}) = \Delta(\mathbf{X}) \sum_{\xi\in\Xi} w^{(\xi)}(\mathcal{L}(\mathbf{X})) \left[p^{(\xi)}\right]^{\mathbf{X}} \tag{16}$$

where Ξ is a given discrete index set, each $p^{(\xi)}(\cdot,\ell)$ is a PDF on $\mathbb{X}$, and each $w^{(\xi)}(L)$ is non-negative with

$$\sum_{\xi\in\Xi} \sum_{L\in\mathcal{F}(\mathbb{L})} w^{(\xi)}(L) = 1\,. \tag{17}$$

Each term in the mixture (16) consists of: a weight $w^{(\xi)}(\mathcal{L}(\mathbf{X}))$ that only depends on the labels $\mathcal{L}(\mathbf{X})$ of the multi-object state $\mathbf{X}$; a multi-object exponential $\left[p^{(\xi)}\right]^{\mathbf{X}}$ that depends on the entire multi-object state.

The cardinality distribution and intensity function (which is also the first moment) of a GLMB are respectively given by

$$\Pr(|X|=n) = \sum_{\xi\in\Xi} \sum_{I\in\mathcal{F}(\mathbb{L})} \delta_n(|I|)\, w^{(\xi)}(I), \tag{18}$$

$$v(x,\ell) = \sum_{\xi\in\Xi} p^{(\xi)}(x,\ell) \sum_{I\in\mathcal{F}(\mathbb{L})} 1_I(\ell)\, w^{(\xi)}(I). \tag{19}$$

The GLMB is often written in the so-called δ-GLMB form by using the identity $w^{(\xi)}(J) = \sum_{I\in\mathcal{F}(\mathbb{L})} w^{(\xi)}(I)\,\delta_I(J)$

$$\boldsymbol{\pi}(\mathbf{X}) = \Delta(\mathbf{X}) \sum_{(I,\xi)\in\mathcal{F}(\mathbb{L})\times\Xi} w^{(\xi)}(I)\,\delta_I(\mathcal{L}(\mathbf{X}))\left[p^{(\xi)}\right]^{\mathbf{X}} \tag{20}$$

For the standard multi-object system model that accounts for thinning, Markov shifts and superposition, the GLMB family is a conjugate prior, and is also closed under the Chapman-Kolmogorov equation [37]. Moreover, the GLMB posterior can be tractably computed to any desired accuracy in the sense that, given any $\epsilon > 0$, an approximate GLMB within ϵ from the actual GLMB in L_1 distance, can be computed (in polynomial time) [38].

2) Marginalized δ-GLMB (Mδ-GLMB): An Mδ-GLMB [39] is a special case of a GLMB with $\Xi = \mathcal{F}(\mathbb{L})$ and density:

$$\boldsymbol{\pi}(\mathbf{X}) = \Delta(\mathbf{X}) \sum_{I\in\mathcal{F}(\mathbb{L})} \delta_I(\mathcal{L}(\mathbf{X}))\,w(I)\,[p(\cdot;I)]^{\mathbf{X}} \tag{21}$$

$$= \Delta(\mathbf{X}) w(\mathcal{L}(\mathbf{X}))\,[p(\cdot;\mathcal{L}(\mathbf{X}))]^{\mathbf{X}}\,. \tag{22}$$

An Mδ-GLMB is completely characterized by the parameter set $\{(w(I)\,, p(\cdot;I)) : I \in \mathcal{F}(\mathbb{L})\}$, and for compactness we use the abbreviation $\boldsymbol{\pi} = \{(w(I)\,, p(\cdot;I))\}_{I\in\mathcal{F}(\mathbb{L})}$ for its density.

In [39], an Mδ-GLMB of the form (21) was proposed to approximate a δ-GLMB of the form (20), by marginalizing (summing) over the discrete space Ξ, i.e. setting

$$w(I) = \sum_{\xi\in\Xi} w^{(\xi)}(I)\,, \tag{23}$$

$$p(x,\ell;I) = \frac{1_I(\ell)}{w(I)} \sum_{\xi\in\Xi} w^{(\xi)}(I)\,p^{(\xi)}(x,\ell)\,. \tag{24}$$

Moreover, using a general result from [43] it was shown that such Mδ-GLMB approximation mimimizes the KLD from the δ-GLMB while preserving the first moment and cardinality distribution [39]. The Mδ-GLMB approximation was used to develop a multi-sensor MOT filter that is scalable with the number of sensors [39].

3) Labeled Multi-Bernoulli (LMB): An LMB [37] is another special case of a GLMB with density

$$\boldsymbol{\pi}(\mathbf{X}) = \Delta(\mathbf{X})\,[1-r]^{\mathbb{M}\setminus\mathcal{L}(\mathbf{X})}\,[1_{\mathbb{M}}\,r]^{\mathcal{L}(\mathbf{X})}\,p^{\mathbf{X}} \tag{25}$$

An LMB is completely characterized by the (finite) parameter set $\{(r(\ell)\,, p(\cdot,\ell)) : \ell \in \mathbb{M}\}$, where $\mathbb{M} \subseteq \mathbb{L}$, $r(\ell) \in [0,1]$ is the *existence probability* of the object with label ℓ, and $p(\cdot,\ell)$ is the PDF (on $\mathbb{X}$) of the object's state. For convenience we use the abbreviation $\boldsymbol{\pi} = \{(r(\ell)\,, p(\cdot,\ell))\}_{\ell\in\mathbb{M}}$ for the density of an LMB. In [40], an approximation of a δ-GLMB by an LMB with matching unlabeled first moment was proposed together with an efficient MOT filter known as the LMB filter.

III. Information Fusion with labeled RFS

In this section, it is shown that the Mδ-GLMB and LMB densities are algebraically closed under KL averaging, i.e. the KLAs of Mδ-GLMBs and LMBs are respectively Mδ-GLMB and LMB. In particular we derive closed form expressions for KLAs of Mδ-GLMBs and LMBs, which are then used to develop consensus fusion of Mδ-GLMB and LMB posterior densities.

A. Multi-Object KLA

The concept of probability density for the multi-object state allows direct extension of the KLA notion to multi-object systems [27].

Given the labeled multi-object densities $\boldsymbol{\pi}^{(i)}$ on $\mathcal{F}(\mathbb{X}\times\mathbb{L})$, $i \in \mathcal{I}$, and the normalized non-negative weights $\omega^{(i)}$, $i \in \mathcal{I}$ (i.e. non-negative weights that sum up to 1):

1) The weighted KLA: $\overline{\boldsymbol{\pi}}$ is defined by

$$\overline{\boldsymbol{\pi}} \triangleq \arg\min_{\boldsymbol{\pi}} \sum_{i\in\mathcal{I}} \omega^{(i)} D_{KL}\left(\boldsymbol{\pi} \parallel \boldsymbol{\pi}^{(i)}\right) \tag{26}$$

where

$$D_{KL}\left(\boldsymbol{\pi} \parallel \boldsymbol{\pi}^{(i)}\right) \triangleq \int \boldsymbol{\pi}(\mathbf{X}) \log\left(\frac{\boldsymbol{\pi}(\mathbf{X})}{\boldsymbol{\pi}^{(i)}(\mathbf{X})}\right) \delta\mathbf{X} \tag{27}$$

is the KLD of $\boldsymbol{\pi}^{(i)}$ from $\boldsymbol{\pi}$ [7], [35], and the integral is the FISST *set integral* defined for any function f on $\mathcal{F}(\mathbb{X}\times\mathbb{L})$ by

$$\int f(\mathbf{X})\delta\mathbf{X} = \sum_{i=0}^{\infty} \frac{1}{i!} \sum_{(\ell_1,...,\ell_i)\in\mathbb{L}^i} \int f(\{(x_1,\ell_1),\ldots,(x_i,\ell_i)\})\, d(x_1,...,x_i). \tag{28}$$

Note that the integrand $f(\mathbf{X})$ has unit of $K^{-|\mathbf{X}|}$, where K is the unit of hyper-volume on $\mathbb{X}$. For compactness, the inner product notation $\langle f, g\rangle$ will be used also for the set integral $\int f(\mathbf{X})g(\mathbf{X})\delta\mathbf{X}$, when $g(\mathbf{X})$ has unit independent of $|\mathbf{X}|$.

2) The normalized weighted geometric mean: is defined by

$$\bigoplus_{i\in\mathcal{I}} \left(\omega^{(i)} \odot \boldsymbol{\pi}^{(i)}\right) = \frac{\prod_{i\in\mathcal{I}} \left[\boldsymbol{\pi}^{(i)}(\mathbf{X})\right]^{\omega^{(i)}}}{\int \prod_{i\in\mathcal{I}} \left[\boldsymbol{\pi}^{(i)}(\mathbf{X})\right]^{\omega^{(i)}} \delta\mathbf{X}}, \tag{29}$$

Note that since the exponents $\omega^{(i)}$, $i \in \mathcal{I}$, sum up to unity, the product in the numerator of (29) has unit of $K^{-|\mathbf{X}|}$, and the set integral in the denominator of (29) is well-defined and unitless. Hence, the normalized weighted geometric mean (29), originally proposed by Mahler in [23] as the multi-object Chernoff fusion rule, is well-defined.

Similar to the single object case, the weighted KLA is given by the normalized weighted geometric mean.

Theorem 1. *[27] - Given multi-object densities $\boldsymbol{\pi}^{(i)}$, $i \in \mathcal{I}$, and normalized non-negative weights $\omega^{(i)}$, $i \in \mathcal{I}$,*

$$\arg\min_{\boldsymbol{\pi}} \sum_{i\in\mathcal{I}} \omega^{(i)} D_{KL}\Big(\boldsymbol{\pi} \,\|\, \boldsymbol{\pi}^{(i)}\Big) = \bigoplus_{i\in\mathcal{I}} \Big(\omega^{(i)} \odot \boldsymbol{\pi}^{(i)}\Big). \tag{30}$$

Note that the label space $\mathbb{L}$ has to be the same for all the densities $\boldsymbol{\pi}^{(i)}$ for the KLA to be well-defined. In [28, Theorem 5.1], it has been mathematically proven that, due to the weight normalization $\sum_i \omega^{(i)} = 1$, the weighted geometric mean (30) ensures immunity to the double counting of information irrespective of the unknown common information in the densities $\boldsymbol{\pi}^{(i)}$.

In [27], it was shown that Poisson and *independently identically distributed cluster* (IID-cluster) RFSs are algebraically closed under KL averaging. While the GLMB family is algebraically closed under the Bayes recursion for the standard multi-object system model and enjoys a number of useful analytical properties, it is not algebraically closed under KL averaging. Nonetheless, there are versatile subfamilies of the GLMBs that are algebraically closed under KL averaging.

B. Weighted KLA of Mδ-GLMB Densities

The following result states that the KLA of Mδ-GLMB densities is also an Mδ-GLMB density. The proof is provided in Appendix A.

Proposition 1. *Given Mδ-GLMB densities $\boldsymbol{\pi}^{(i)} = \{(w^{(i)}(I), p^{(i)}(\cdot; I))\}_{I\in\mathcal{F}(\mathbb{L})}$, $i \in \mathcal{I}$, and normalized non-negative weights $\omega^{(i)}$, $i \in \mathcal{I}$, the normalized weighted geometric mean $\overline{\boldsymbol{\pi}}$, and hence the KLA, is an Mδ-GLMB given by:*

$$\overline{\boldsymbol{\pi}} = \{(\overline{w}(L), \overline{p}(\cdot; L))\}_{L\in\mathcal{F}(\mathbb{L})} \tag{31}$$

where

$$\overline{w}(L) = \frac{\prod_{i\in\mathcal{I}} \Big(w^{(i)}(L)\Big)^{\omega^{(i)}} \left[\int \prod_{i\in\mathcal{I}} \Big(p^{(i)}(x,\cdot;L)\Big)^{\omega^{(i)}} dx\right]^L}{\sum_{J\subseteq\mathbb{L}} \prod_{i\in\mathcal{I}} \Big(w^{(i)}(J)\Big)^{\omega^{(i)}} \left[\int \prod_{i\in\mathcal{I}} \Big(p^{(i)}(x,\cdot;J)\Big)^{\omega^{(i)}} dx\right]^J} \tag{32}$$

$$\overline{p}(\cdot,\ell;L) = \frac{\prod_{i\in\mathcal{I}} \Big(p^{(i)}(\cdot,\ell;L)\Big)^{\omega^{(i)}}}{\int \prod_{i\in\mathcal{I}} \Big(p^{(i)}(x,\ell;L)\Big)^{\omega^{(i)}} dx}\,. \tag{33}$$

Remark 2. The component $(\overline{w}(L), \overline{p}(\cdot; L))$ of the KLA Mδ-GLMB can be rewritten as

$$\overline{w}(L) \propto \prod_{i\in\mathcal{I}} \Big(w^{(i)}(L)\Big)^{\omega^{(i)}} \left[\int \prod_{i\in\mathcal{I}} \Big(p^{(i)}(x,\cdot;L)\Big)^{\omega^{(i)}} dx\right]^{L} \tag{34}$$

$$\overline{p}(\cdot; L) = \bigoplus_{i\in\mathcal{I}} \Big(\omega^{(i)} \odot p^{(i)}(\cdot; L)\Big) \tag{35}$$

where (35) is indeed the Chernoff fusion rule for the single-object PDFs [25]. Note also from (34) and (35) that each Mδ-GLMB component $(\overline{w}(L), \overline{p}(\cdot; L))$ can be independently determined. Thus, the overall fusion procedure is fully parallelizable.

C. Weighted KLA of LMB Densities

The following result states that the KLA of LMB densities is also an LMB density. The proof is provided in Appendix B.

Proposition 2. *Given LMB densities $\boldsymbol{\pi}^{(i)} = \{(r^{(i)}(\ell), p^{(i)}(\cdot,\ell))\}_{\ell\in\mathbb{L}}$, $i \in \mathcal{I}$, and normalized non-negative weights $\omega^{(i)}$, $i \in \mathcal{I}$, the normalized weighted geometric mean $\overline{\boldsymbol{\pi}}$, and hence the KLA, is an LMB given by:*

$$\overline{\boldsymbol{\pi}} = \{(\overline{r}(\ell), \overline{p}(\cdot,\ell))\}_{\ell\in\mathbb{L}} \tag{36}$$

where

$$\overline{r}(\ell) = \frac{\displaystyle\int \prod_{i\in\mathcal{I}} \Big(r^{(i)}(\ell) p^{(i)}(x,\ell)\Big)^{\omega^{(i)}} dx}{\displaystyle\prod_{i\in\mathcal{I}} \Big(1 - r^{(i)}(\ell)\Big)^{\omega^{(i)}} + \int \prod_{i\in\mathcal{I}} \Big(r^{(i)}(\ell) p^{(i)}(x,\ell)\Big)^{\omega^{(i)}} dx} \tag{37}$$

$$\overline{p}(\cdot,\ell) = \bigoplus_{i\in\mathcal{I}} \Big(\omega^{(i)} \odot p^{(i)}(\cdot,\ell)\Big) \,. \tag{38}$$

Remark 3. Similar to the KLA of Mδ-GLMBs, (38) is indeed the Chernoff fusion rule for the single-object PDFs [25]. Note from (37) and (38) that each LMB component $(\overline{r}(\ell), \overline{p}(\cdot,\ell))$ can be independently determined. Thus, the overall fusion procedure is fully parallelizable.

D. Consensus Fusion for Labeled RFSs

Consider a sensor network $\mathcal{N}$ with multi-object density $\boldsymbol{\pi}^{(i)}$ at each node i, and non-negative consensus weights $\omega^{(i,j)}$ relating node i to nodes $j \in \mathcal{N}^{(i)}$, such that $\sum_{j\in\mathcal{N}^{(i)}} \omega^{(i,j)} = 1$. The global KLA over

the entire network can be computed in a distributed and scalable way by using the consensus algorithm [20], [27, Section III.A]. Starting with $\boldsymbol{\pi}_0^{(i)} = \boldsymbol{\pi}^{(i)}$, each node $i \in \mathcal{N}$ carries out the consensus iteration

$$\boldsymbol{\pi}_n^{(i)} = \bigoplus_{j \in \mathcal{N}^{(i)}} \left(\omega^{(i,j)} \odot \boldsymbol{\pi}_{n-1}^{(j)} \right) . \tag{39}$$

As shown in [27, Section III-B], the consensus iteration (39)—which is the multi-object counterpart of equation (12)—enjoys some nice convergence properties. In particular, if the consensus matrix is primitive and doubly stochastic, the consensus iterate of each node in the network converges to the global unweighted KLA of the multi-object posterior densities as n tends to infinity. Convergence analysis for the multi-object case follows along the same line as in [16], [20] since $\mathcal{F}(\mathbb{X} \times \mathbb{L})$ is a metric space [7]. In practice, the iteration is stopped at some finite n. Further, as pointed out in [28, Remark 1], the consensus iterations (39) always generate multi-object densities $\boldsymbol{\pi}_n^{(i)}$ that mitigate double counting irrespectively of the number n of iterations.

Starting with δ-GLMBs, the consensus iteration (39) always returns Mδ-GLMBs, moreover the Mδ-GLMB parameter set can be computed by the *Mδ-GLMB fusion rules* (34) and (35). Similarly, for LMBs the consensus iteration (39) always returns LMBs whose parameter set can be computed by the *LMB fusion rules* (37) and (38). The fusion rules (35) and (38) involve consensus of single-object PDFs.

A typical choice for representing each single-object density is a *Gaussian Mixture* (GM) [44], [45]. In this case, the fusion rules (35) and (38) involve exponentiation and multiplication of GMs where the former, in general, does not provide a GM. Hence, in order to preserve the GM form, a suitable approximation of the GM exponentiation has to be devised. The in-depth discussion and efficient implementation proposed in [27, Section III.D] for generic GMs can also be applied to the location PDF fusion (35) and (38). Considering, for the sake of simplicity, the case of two GMs

$$p_i(x) = \sum_{j=1}^{N_i} \alpha_{i,j} \mathcal{N}\left(x; \mu_{i,j}, P_{i,j}\right)$$

for $i \in \{a, b\}$, (35) and (38) can be approximated as follows:

$$\overline{p}(x) = \frac{\sum_{j=1}^{N_a} \sum_{k=1}^{N_b} \overline{\alpha}_{jk} \mathcal{N}\left(x; \overline{\mu}_{jk}, \overline{P}_{jk}\right)}{\sum_{j=1}^{N_a} \sum_{k=1}^{N_b} \overline{\alpha}_{jk}} \tag{40}$$

where

$$\overline{P}_{jk} = \left[\omega P_{a,j}^{-1} + (1-\omega)P_{b,k}^{-1}\right]^{-1} \tag{41}$$

$$\overline{\mu}_{jk} = P_{jk}\left[\omega P_{a,j}^{-1}\mu_{a,j} + (1-\omega)P_{b,k}^{-1}\mu_{b,k}\right] \tag{42}$$

$$\overline{\alpha}_{jk} = \alpha_{a,j}^{\omega}\,\alpha_{b,k}^{1-\omega}\beta(\omega, P_{a,j})\,\beta(1-\omega, P_{b,k})\,\mathcal{N}\left(\mu_{a,j}-\mu_{b,k}; 0, \frac{P_{a,j}}{\omega}+\frac{P_{b,k}}{1-\omega}\right) \tag{43}$$

$$\beta(\omega, P) \triangleq \frac{\left[\det\left(2\pi P\omega^{-1}\right)\right]^{\frac{1}{2}}}{\left[\det(2\pi P)\right]^{\frac{\omega}{2}}} \tag{44}$$

The fusion (40) can be extended to $|\mathcal{N}| \geq 2$ agents by sequentially applying the pairwise fusion rule (40)-(44) $|\mathcal{N}|-1$ times. By the associative and commutative properties of multiplication, the ordering of pairwise fusions is irrelevant. Notice that (40)-(44) amounts to performing a Chernoff fusion on any possible pair formed by a Gaussian component of agent a and a Gaussian component of agent b. Moreover, the coefficient $\overline{\alpha}_{jk}$ of the resulting (fused) component includes a factor $\mathcal{N}\left(\mu_{a,j}-\mu_{b,k}; 0, \omega^{-1}P_{a,j}+(1-\omega)^{-1}P_{b,k}\right)$ that measures the separation of the two fusing components $(\mu_{a,j}, P_{a,j})$ and $(\mu_{b,k}, P_{b,k})$. The approximation (40)-(44) is reasonably accurate for well-separated Gaussian components but might easily deteriorate in presence of closely located components. In this respect, merging of nearby components before fusion has been exploited in [27] to mitigate the problem. Further, a more accurate, but also more computationally demanding, approximation has been proposed in [46].

The other common approach for approximating a single object PDF $p(\cdot)$ is via *particles*, i.e. weighted sums of Dirac delta functions, which can address non-linear, non-Gaussian dynamics and measurements as well as non-uniform field of view. However, computing the KLA requires multiplying together powers of relevant PDFs, which cannot be performed directly on weighted sums of Dirac delta functions. While this problem can be addressed by further approximating the particle PDFs by continuous PDFs (e.g. GMs) using techniques such as kernel density estimation [30], least square estimation [47], [48] or parametric model approaches [49], such approximations increase the in-node computational burden. Moreover, the local filtering steps are also more resource demanding compared to a GM implementation. Hence, at this developmental stage, it is more efficient to work with GM approximations.

IV. Consensus DMOT

In this section, we present two novel fully distributed and scalable multi-object tracking algorithms based on Propositions 1 and 2 along with consensus [11], [12], [16], [20] to propagate information throughout the network.

A. Bayesian Multi-Object Filtering

We begin this section with the Bayes MOT filter that propagates the multi-object posterior/filtering density. In this formulation the multi-object state is modeled as a labeled RFS in which a label is an ordered pair of integers $\ell = (k, i)$, where k is the *time of birth*, and $i \in \mathbb{N}$ is a unique index to distinguish objects born at the same time. The label space for objects born at time k is $\mathbb{L}_k = \{k\} \times \mathbb{N}$. An object born at time k has, therefore, state $\mathbf{x} \in \mathbb{X} \times \mathbb{L}_k$. Hence, the label space for objects at time k (including those born prior to k), denoted as $\mathbb{L}_{0:k}$, is constructed recursively by $\mathbb{L}_{0:k} = \mathbb{L}_{0:k-1} \cup \mathbb{L}_k$ (note that $\mathbb{L}_{0:k-1}$ and $\mathbb{L}_k$ are disjoint). A multi-object state $\mathbf{X}$ at time k, is a finite subset of $\mathbb{X} \times \mathbb{L}_{0:k}$. For convenience, we denote $\mathbb{L}_- \triangleq \mathbb{L}_{0:k-1}$, $\mathbb{B} \triangleq \mathbb{L}_k$, and $\mathbb{L} \triangleq \mathbb{L}_- \cup \mathbb{B}$.

Let $\boldsymbol{\pi}_k$ denote the *multi-object filtering density* at time k, and $\boldsymbol{\pi}_{k|k-1}$ the *multi-object prediction density* (for compactness, the dependence on the measurements is omitted). Then, starting from $\boldsymbol{\pi}_0$, the *multi-object Bayes recursion* propagates $\boldsymbol{\pi}_k$ in time according to the following update and prediction [7], [35]

$$\boldsymbol{\pi}_{k|k-1}(\mathbf{X}_k) = \left\langle \boldsymbol{f}_{k|k-1}(\mathbf{X}_k|\,\cdot), \boldsymbol{\pi}_{k-1}(\,\cdot) \right\rangle , \tag{45}$$

$$\boldsymbol{\pi}_k(\mathbf{X}_k) = \left(g_k(Z_k|\,\cdot) \oplus \boldsymbol{\pi}_{k|k-1}(\,\cdot)\right)(\mathbf{X}_k) , \tag{46}$$

where $\boldsymbol{f}_{k|k-1}(\,\cdot\,|\,\cdot\,)$ is the *multi-object transition density* from time $k-1$ to time k, and $g_k(\,\cdot\,|\,\cdot\,)$ is the *multi-object likelihood function* at time k. The multi-object likelihood function encapsulates the underlying models for detections and false alarms while the multi-object transition density encapsulates the underlying models of motion, birth and death. The multi-object filtering (or posterior) density captures all information on the number of objects, and their states [7].

Note that the recursions (45)-(46) are the multi-object counterpart of (6)-(7), which admit a closed form solution, under the standard multi-object system model, known as the GLMB filter [37] (see also [38] for implementation details). However, the GLMB family is not closed under KL averaging. Consequently, we look towards approximations such as the Mδ-GLMB and LMB filters for analytic solutions to DMOT.

B. The Mδ-GLMB Filter

In the following we outline the prediction and update steps for the Mδ-GLMB filter. Additional details can be found in [38].

1) Mδ-GLMB Prediction: Given the previous multi-object state $\mathbf{X}_{k-1}$, each state $(x_{k-1}, \ell_{k-1}) \in \mathbf{X}_{k-1}$ either continues to exist at the next time step with probability $P_S(x_{k-1}, \ell_{k-1})$ and evolves to a new state

(x_k, ℓ_k) with probability density $f_{k|k-1}(x_k|x_{k-1}, \ell_{k-1})$, or dies with probability $1 - P_S(x_{k-1}, \ell_{k-1})$. The set of new objects born at the next time step is distributed according to the LMB

$$\boldsymbol{f}_B(\mathbf{X}) = \Delta(\mathbf{X}) \left[1 - r_B\right]^{\mathbb{B}\backslash\mathcal{L}(\mathbf{X})} \left[1_{\mathbb{B}}\, r_B\right]^{\mathcal{L}(\mathbf{X})} \left[p_B\right]^{\mathbf{X}}. \tag{47}$$

It is assumed that $p_B(\cdot, l) \neq p_B(\cdot, j)$ when $l \neq j$. Note that $\boldsymbol{f}_B(\mathbf{X}) = 0$ if $\mathbf{X}$ contains any element $\mathbf{x}$ with $\mathcal{L}(\mathbf{x}) \notin \mathbb{B}$. The multi-object state at the next time $\mathbf{X}$ is the superposition of surviving objects and new born objects, and the multi-object transition density can be found in [37, Subsection IV.D].

Remark 4. The LMB birth model assigns unique labels to objects in the following sense. Consider two objects born at time k with kinematic states x and y. In birth models such as labeled Poisson [37], x could be assigned label $(k, 1)$ and y label $(k, 2)$, i.e. the multi-object state is $\{(x, (k, 1)), (y, (k, 2))\}$, or conversely x assigned label $(k, 2)$ and y label $(k, 1)$, i.e. the multi-object state is $\{(x, (k, 2)), (y, (k, 1))\}$. Such non-uniqueness arises because the kinematic state of an object is generated independently of the label. This does not occur in the LMB model because an object with label ℓ, has kinematic state generated from $p_B(\cdot, \ell)$. If kinematic states x and y are drawn respectively from $p_B(\cdot, (k, 1))$ and $p_B(\cdot, (k, 2))$, then the labeled multi-object state is uniquely $\{(x, (k, 1)), (y, (k, 2))\}$.

Given the Mδ-GLMB multi-object posterior density $\boldsymbol{\pi}_{k-1} = \{(w_{k-1}(I), p_{k-1}(\cdot; I))\}_{I \in \mathcal{F}(\mathbb{L}_-)}$, the multi-object prediction density is the Mδ-GLMB $\boldsymbol{\pi}_{k|k-1} = \{(w_{k|k-1}(I), p_{k|k-1}(\cdot; I))\}_{I \in \mathcal{F}(\mathbb{L})}$, where

$$w_{k|k-1}(I) = \left[1 - r_B\right]^{\mathbb{B}\backslash I} \left[1_{\mathbb{B}}\, r_B\right]^{I \cap \mathbb{B}} w_S^{(I)}(I \cap \mathbb{L}_-) \tag{48}$$

$$p_{k|k-1}(x, \ell; I) = 1_{\mathbb{B}}(\ell) p_B(x, \ell) + 1_{\mathbb{L}_-}(\ell) p_S(x, \ell; I) \tag{49}$$

$$w_S^{(I)}(L) = \left[\overline{P}_S^{(I)}\right]^L \sum_{J \supseteq L} \left[1 - \overline{P}_S^{(I)}\right]^{J-L} w_{k-1}(J) \tag{50}$$

$$p_S(x, \ell; I) = \frac{\left\langle P_S(\cdot, \ell) f_{k|k-1}(x|\cdot, \ell), p_{k-1}(\cdot, \ell, I) \right\rangle}{\overline{P}_S^{(I)}(\ell)} \tag{51}$$

$$\overline{P}_S^{(I)}(\ell) = \left\langle P_S(\cdot, \ell), p_{k-1}(\cdot, \ell; I) \right\rangle. \tag{52}$$

2) Mδ-GLMB Update: Given a multi-object state $\mathbf{X}_k$, each state $(x_k, \ell_k) \in \mathbf{X}_k$ is either detected with probability $P_D(x_k, \ell_k)$ and generates a measurement z with likelihood $g_k(z|x_k, \ell_k)$, or missed with probability $1 - P_D(x_k, \ell_k)$. The multi-object observation $Z_k = \{z_1, \ldots, z_{|Z_k|}\}$ is the superposition of the detected points and Poisson clutter with intensity function κ. Assuming that, conditional on $\mathbf{X}_k$, detections are independent, and that clutter is independent of the detections, the multi-object likelihood is given by [37, Subsection IV.D]

$$g_k(Z_k|\mathbf{X}_k) \propto \sum_{\theta \in \Theta(\mathcal{L}(\mathbf{X}_k))} \left[\psi_{Z_k}(\cdot; \theta)\right]^{\mathbf{X}_k} \tag{53}$$

where $\Theta(I)$ is the set of mappings $\theta : I \to \{0, 1, \ldots, |Z_k|\}$, such that $\theta(i) = \theta(i') > 0$ implies $i = i'$, and

$$\psi_{Z_k}(x, \ell; \theta) = \begin{cases} \dfrac{P_D(x, \ell)\, g_k(z_{\theta(\ell)}|x, \ell)}{\kappa(z_{\theta(\ell)})}, & \text{if } \theta(\ell) > 0 \\ 1 - P_D(x, \ell), & \text{if } \theta(\ell) = 0 \end{cases}.$$

Note that an association map θ specifies which tracks generated which measurements, i.e. track ℓ generates measurement $z_{\theta(\ell)} \in Z_k$, with undetected tracks assigned to 0. The condition "$\theta(i) = \theta(i') > 0$ implies $i = i'$", means that, at any time, a track can generate at most one measurement, and a measurement can be assigned to at most one track.

Given the Mδ-GLMB multi-object prediction density $\boldsymbol{\pi}_{k|k-1} = \{(w_{k|k-1}(I), p_{k|k-1}(\cdot; I))\}_{I \in \mathcal{F}(\mathbb{L})}$, the M$\delta$-GLMB updated density is given by $\boldsymbol{\pi}_k = \{(w_k(I), p_k(\cdot; I))\}_{I \in \mathcal{F}(\mathbb{L})}$, where

$$w_k(I) = \sum_{\theta \in \Theta(I)} w_k^{(I,\theta)}, \tag{54}$$

$$p_k(x, \ell; I) = \frac{1_I(\ell)}{w_k(I)} \sum_{\theta \in \Theta(I)} w_k^{(I,\theta)} p_k^{(\theta)}(x, \ell; I) \tag{55}$$

$$w_k^{(I,\theta)} \propto w_{k|k-1}(I) \left[\overline{\psi}_{Z_k}^{(I,\theta)}(\cdot)\right]^I \tag{56}$$

$$\overline{\psi}_{Z_k}^{(I,\theta)}(\ell) = \left\langle p_{k|k-1}(\cdot, \ell; I), \psi_{Z_k}(\cdot, \ell; \theta) \right\rangle \tag{57}$$

$$p_k^{(\theta)}(x, \ell; I) = \frac{p_{k|k-1}(x, \ell; I) \psi_{Z_k}(x, \ell; \theta)}{\overline{\psi}_{Z_k}^{(I,\theta)}(\ell)}. \tag{58}$$

Note that the exact multi-object posterior density is not a Mδ-GLMB, but a δ-GLMB. The Mδ-GLMB update approximates the posterior density by a Mδ-GLMB that preserves the posterior PHD and cardinality distribution.

A tractable suboptimal multi-object estimate can be obtained from the posterior Mδ-GLMB $\boldsymbol{\pi}_k = \{(w_k(I), p_k(\cdot; I))\}_{I \in \mathcal{F}(\mathbb{L})}$ as follows: first determine the maximum a-posteriori cardinality estimate N^* from

$$\Pr(|X| = n) = \sum_{I \in \mathcal{F}(\mathbb{L})} \delta_n(|I|)\, w_k(I); \tag{59}$$

then determine the label set I^* with highest weight $w_k(I^*)$ among those with cardinality N^*; and finally determine the expected values of the kinematic states from $p_k(\cdot, \ell; I^*)$, $\ell \in I^*$. Alternatively, one can determine the set I^* of labels with the N^* highest existence probabilities $\sum_{I \in \mathcal{F}(\mathbb{L})} 1_I(\ell) w_k(I)$; and then the expected values of the kinematic states from $\sum_{I \in \mathcal{F}(\mathbb{L})} 1_I(\ell) p_k(\cdot, \ell; I)$, $\ell \in I^*$.

In addition to the generality of the tracking solution, the consideration of label-dependent P_S and P_D are useful in some applications. For instance, in live cell microscopy the survival probability of a cell is also a function of its age. This can be accommodated by a label-dependent P_S because the label contains

the time of birth and the age of a labeled state can be determined by subtracting the time of birth from the current time. In some trackers, a track is considered to begin when it is first detected, by convention. In this case, label-dependent P_D would be able to capture this assumption, since the label provides the time of birth.

C. Consensus Mδ-GLMB Filter

This subsection details the *Consensus* Mδ-GLMB filter using a Gaussian mixture implementation. Each node $i \in \mathcal{N}$ of the network operates autonomously at each sampling interval k, starting from its own previous estimates of the multi-object distribution $\boldsymbol{\pi}^{(i)}$, with PDFs $p(\cdot, \ell; I)$, $\forall\, \ell \in I$, $I \in \mathcal{F}(\mathbb{L})$, represented by Gaussian mixtures, and producing, at the end of N consensus iterations, its new consensus multi-object distribution $\boldsymbol{\pi}^{(i)} = \boldsymbol{\pi}_N^{(i)}$.

The steps of the Consensus Mδ-GLMB filter over the network $\mathcal{N}$ are given as follows.

1) Each node $i \in \mathcal{N}$ locally performs an Mδ-GLMB prediction and update. The details of these two procedures are described in the previous subsections.
2) At each consensus step, node i transmits its data to neighbouring nodes $j \in \mathcal{N}^{(i)} \backslash \{i\}$. Upon receiving data from its neighbours, node i carries out the fusion rule of Proposition 1 over its in-neighbours $\mathcal{N}^{(i)}$, i.e. performs (39) using information from $\mathcal{N}^{(i)}$. A merging operation for each of the PDFs is applied to reduce the joint communication-computation burden for the next consensus step. This procedure is repeatedly applied for a chosen number $N \geq 1$ of consensus steps.
3) After consensus, an estimate of the multi-object state is obtained via the procedure described in Table II.

The operations executed locally by each node $i \in \mathcal{N}$ of the network are summarized in Table I.

We point out that in the algorithm of Table I, each single object state lives in the space $\mathbb{X} \times \mathbb{L}_{0:k}$ and the multi-object state space at time k, $\mathcal{F}(\mathbb{X} \times \mathbb{L}_{0:k})$, is the same for all nodes. This is fully consistent with the fusion rule of Theorem 1 and the fusion rule for Mδ-GLMBs.

In implementation, each component of the Mδ-GLMB, also known as a hypothesis, is indexed by an element of $\mathcal{F}(\mathbb{L}_{0:k})$, i.e., a set of labels. Since the cardinality of the label space $\mathbb{L}_{0:k}$ increases with time, each node performs a pruning of the hypotheses (for instance by removing those with low weights so that the total number of hypotheses never exceed a fixed number $I_{\max}$). Hence, at time k each node i has a density containing at most $I_{\max}$ components with indices in $\mathbb{I}_k^{(i)} \subseteq \mathcal{F}(\mathbb{L}_{0:k})$, and the weights of the remaining components, i.e. those in $\mathcal{F}(\mathbb{L}_{0:k}) \setminus \mathbb{I}_k^{(i)}$, are set to zero. As a result, when the densities of nodes i and j are fused, only the common components, i.e those belonging to $\mathbb{I}_k^{(i)} \cap \mathbb{I}_k^{(j)}$ have non-zero weights.

TABLE I

CONSENSUS MARGINALIZED δ-GLMB FILTER

procedure CONSENSUS Mδ-GLMB(NODE i, TIME k)

LOCAL PREDICTION ▷ See subsection IV-B1

LOCAL UPDATE ▷ See subsection IV-B2

MARGINALIZATION ▷ See eqs. (54)-(58)

for $n = 1, \ldots, N$ **do**

INFORMATION EXCHANGE

FUSION OVER $\mathcal{N}^{(i)}$ ▷ See eqs. (34) and (35)

MERGING ▷ See [44, Table II, Section III.C]

end for

ESTIMATE EXTRACTION ▷ See algorithm in Table II

end procedure

TABLE II

Mδ-GLMB ESTIMATE EXTRACTION

INPUT: $\boldsymbol{\pi}_k = \{(w_k(I), p_k(\cdot; I))\}_{I \in \mathcal{F}(\mathbb{L})}$

OUTPUT: $\mathbf{X}^*$

for $c = 1, \ldots,$ **do**

$$\rho(c) = \sum_{I \in \mathcal{F}(\mathbb{L})} \delta_c(|I|)\, w_k(I)$$

end for

$$N^* = \arg\max_c \rho(c)$$

$$I^* = \arg \max_{I \in \mathcal{F}_{N^*}(\mathbb{L})} w_k^{(I)}$$

$$\mathbf{X}^* = \left\{(x^*, \ell^*) \colon \ell^* \in I^*,\, x^* = \arg\max_x p_k(x, \ell^*; I^*)\right\}$$

Remark 5. In [50] it has been shown that the centralized Mδ-GLMB filter features linear complexity in the number of sensors. As far as consensus Mδ-GLMB is concerned, each node has to carry out local prediction and local update, whose computational complexity is clearly independent of the number of nodes (sensors), and the consensus task (i.e. repeated KLA fusion over the subset of in-neighbors) which requires in the order of $N\, I_{max}\, d$ computations, d being the node in-degree.

D. The LMB Filter

As suggested by its name, the LMB filter propagates an LMB multi-object posterior density forward in time [40]. It is an approximation of the δ-GLMB filter [37], [38].

1) LMB Prediction: Given the LMB multi-object posterior density $\boldsymbol{\pi}_{k-1} = \{(r_{k-1}(\ell), p_{k-1}(\cdot,\ell))\}_{\ell\in\mathbb{L}_-}$, the multi-object prediction density is the LMB [40]

$$\boldsymbol{\pi}_{k|k-1} = \{(r_S(\ell), p_S(\cdot,\ell))\}_{\ell\in\mathbb{L}_-} \cup \{(r_B(\ell), p_B(\cdot,\ell))\}_{\ell\in\mathbb{B}} \tag{60}$$

where

$$r_S(\ell) = \langle P_S(\cdot,\ell), p_{k-1}(\cdot,\ell)\rangle\, r_{k-1}(\ell) \tag{61}$$

$$p_S(\cdot,\ell) = \frac{\langle P_S(\cdot,\ell) f_{k|k-1}(x|\cdot,\ell), p_{k-1}(\cdot,\ell)\rangle}{\langle P_S(\cdot,\ell), p_{k-1}(\cdot,\ell)\rangle} \tag{62}$$

2) LMB Update: Given the LMB multi-object prediction density $\boldsymbol{\pi}_{k|k-1} = \{(r_{k|k-1}(\ell), p_{k|k-1}(\cdot,\ell))\}_{\ell\in\mathbb{L}}$, the LMB updated density is given by $\boldsymbol{\pi}_k = \{(r_k(\ell), p_k(\cdot,\ell))\}_{\ell\in\mathbb{L}}$, where

$$r_k(\ell) = \sum_{(I,\theta)\in\mathcal{F}(\mathbb{L})\times\Theta(I)} 1_I(\ell) w_k^{(\theta)}(I) \tag{63}$$

$$p_k(x,\ell) = \frac{1}{r_k(\ell)} \sum_{(I,\theta)\in\mathcal{F}(\mathbb{L})\times\Theta(I)} 1_I(\ell) w_k^{(\theta)}(I)\, p_k^{(\theta)}(x,\ell) \tag{64}$$

$$w_k^{(\theta)}(I) \propto \left[\overline{\psi}_{Z_k}^{(\theta)}\right]^I \left[1 - r_{k|k-1}(\cdot)\right]^{\mathbb{L}\setminus I} \left[1_{\mathbb{L}}\, r_{k|k-1}\right]^I \tag{65}$$

$$p_k^{(\theta)}(x,\ell) = \frac{p_{k|k-1}(x,\ell)\, \psi_{Z_k}(x,\ell;\theta)}{\overline{\psi}_{Z_k}^{(\theta)}(\ell)} \tag{66}$$

$$\overline{\psi}_{Z_k}^{(\theta)}(\ell) = \left\langle p_{k|k-1}(\cdot,\ell), \psi_{Z_k}(\cdot,\ell;\theta)\right\rangle . \tag{67}$$

Note that the exact multi-object posterior density is not an LMB, but a GLMB. The LMB update approximates the GLMB posterior by an LMB that matches the unlabeled PHD. The reader is referred to [40] for an efficient implementation of the LMB filter.

E. Consensus LMB Filter

This subsection describes the *Consensus* LMB filter using a Gaussian mixture implementation. The steps of the Consensus LMB filter are the same as the Consensus Mδ-GLMB tracking filter described in section IV-C, with the LMB prediction and update in place of those of the Mδ-GLMB. Each node $i \in \mathcal{N}$ of the network operates autonomously at each sampling interval k, starting from its own previous estimates of the multi-object distribution $\boldsymbol{\pi}^{(i)}$, with PDFs $p(\cdot,\ell)$, $\forall\, \ell \in \mathbb{L}$, represented by Gaussian mixtures, and producing, at the end of N consensus iterations, its new consensus multi-object distribution $\boldsymbol{\pi}^{(i)} = \boldsymbol{\pi}_N^{(i)}$. The operations executed locally by each node $i \in \mathcal{N}$ of the network are summarized in Table III.

TABLE III

CONSENSUS LMB FILTER

procedure CONSENSUS LMB(NODE i, TIME k)

LOCAL PREDICTION ▷ See subsection IV-D1

LOCAL UPDATE ▷ See subsection IV-D2

for $n = 1, \ldots, N$ **do**

INFORMATION EXCHANGE

FUSION OVER $\mathcal{N}^{(i)}$ ▷ See eqs. (37) and (38)

MERGING ▷ See [44, Table II, Section III.C]

end for

ESTIMATE EXTRACTION ▷ See algorithm in Table IV

end procedure

TABLE IV

LMB ESTIMATE EXTRACTION

INPUT: $\boldsymbol{\pi}_k = \{r_k(\ell), p_k(\cdot)\}_{\ell \in \mathbb{L}}$, N^*

OUTPUT: $\mathbf{X}^*$

for $c = 1, \ldots, N^*$ **do**

$$\rho(c) = \sum_{I \in \mathcal{F}(\mathbb{L})} \delta_c(|I|) \prod_{\ell \in \mathbb{L} \setminus I} (1 - r_k(\ell)) \prod_{\ell \in I} r_k(\ell)$$

end for

$C^* = \arg\max_c \rho(c)$

$\hat{\mathbb{L}} = \varnothing$

for $c^* = 1, \ldots, C^*$ **do**

$$\mathbb{L}^* = \mathbb{L}^* \cup \arg \max_{\ell \in \mathbb{L} \setminus \mathbb{L}^*} r_k(\ell)$$

end for

$$\mathbf{X}^* = \left\{ (x^*, \ell^*) \colon \ell^* \in \mathbb{L}^*,\ x^* = \arg\max_x p_k(x, \ell^*) \right\}$$

Similar to Mδ-GLMB Consensus, each node performs pruning of the components (or hypotheses) so as to cap their numbers. For an LMB, each component is indexed by an element of $\mathbb{L}_k^{(i)}$, i.e. a label. Hence, for LMB Consensus, at time k each node i has a density containing at most $I_{\max}$ components indexed by $\mathbb{L}_k^{(i)} \subseteq \mathbb{L}_{0:k}$, and the weights of those components indexed by $\mathbb{L}_{0:k} \setminus \mathbb{L}_k^{(i)}$ are set to zero. Table V summarizes the information exchanged among the nodes for both Mδ-GLMB and LMB trackers.

TABLE V

INFORMATION EXCHANGED AT TIME INTERVAL k BY AGENTS $i \in \mathcal{N}$ FOR BOTH Mδ-GLMB AND LMB TRACKERS. IT IS ASSUMED THAT THE EXCHANGED INFORMATION ARE REPRESENTED WITH 4 BYTES FLOATING POINT VARIABLES.

Tracker	**Information exchanged**	**Total bytes exchanged**
Mδ-GLMB	$\forall\, I \in \mathbb{I}_k^{(i)}$: $\omega_k^{(i)}$, $p_k^{(i)}(x, \ell; I)$	$4 \sum_{I \in \mathbb{I}_k^{(i)}} (1 + (4 + 10)\lvert I \rvert)$
LMB	$\forall\, \ell \in \mathbb{L}_k^{(i)}$: $r_k^{(i)}(\ell)$, $p_k^{(i)}(x, \ell)$	$4(1 + (4 + 10)\lvert \mathbb{L}_k^{(i)} \rvert)$

V. PERFORMANCE EVALUATION

To assess the performance of the proposed consensus multi-object tracking filters, we consider a 2-D multi-object tracking scenario over a surveillance area of $50 \times 50\, km^2$, wherein the sensor network depicted in Fig. 2 is deployed. The scenario consists of 5 objects as shown in Fig. 3. The proposed trackers are also compared with the Consensus CPHD filter of [27] which, however, does not provide tracks, and with the *centralized* Mδ-GLMB filter [39] which recursively processes all the measurements collected by the nodes, thus providing a performance reference for the distributed filters.

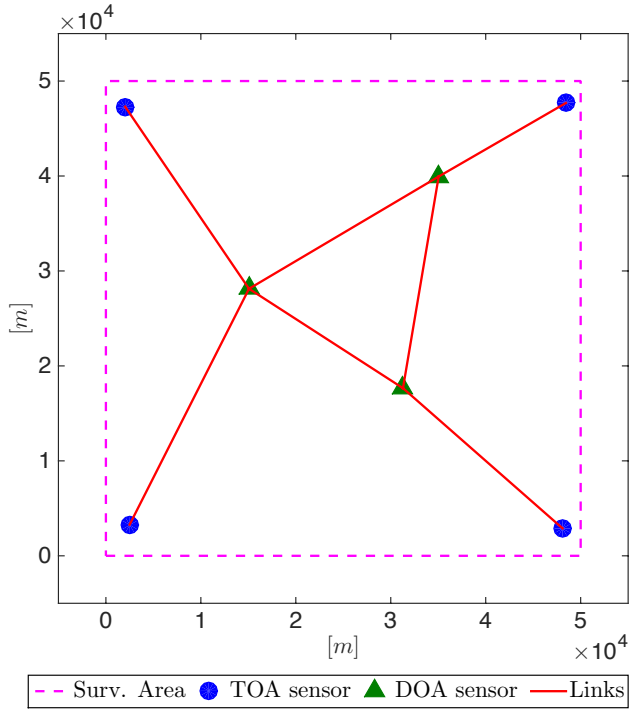

Fig. 2. Network with 7 sensors: 4 TOA and 3 DOA.

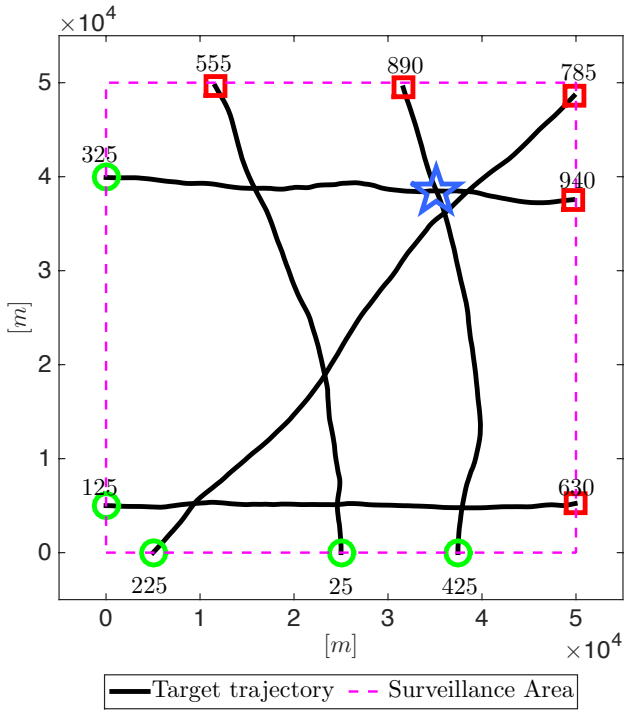

Fig. 3. Target trajectories considered in the simulation experiment. The start/end point for each trajectory is denoted, respectively, by •\■. The ★ indicates a rendezvous point.

The kinematic object state is denoted by $x = [p_x,\, \dot{p}_x,\, p_y,\, \dot{p}_y]^\top$, i.e. the planar position and velocity. The motion of objects is modeled by the filters according to the Nearly-Constant Velocity (NCV) model [1]–[4]: $\mathcal{N}(x_k; Fx_{k-1}, Q)$, where

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

$\sigma_w = 5\, m/s^2$ and the sampling interval is $T_s = 5\, s$. Objects pass through the surveillance area and partial prior information of the object birth locations is assumed. Accordingly, a 10-component LMB RFS $\boldsymbol{\pi}_B = \{(r_B(\ell)\,, p_B(\cdot, \ell))\}_{\ell \in \mathbb{B}}$ is used to model the birth process. Table V gives a detailed summary of such components. The aim of using such a birth process is to cover all possible locations where objects appear, but also locations where no objects are present or born. In this way, it is also tested the algorithm's ability of ruling out possible false objects, arising in wrong birth locations, that are generated by clutter measurements.

The sensor network considered in this example (see Fig. 2) consists of 4 *range-only* (Time Of Arrival, TOA) and 3 *bearing-only* (Direction Of Arrival, DOA) sensors characterized by the following

TABLE VI

COMPONENTS OF THE LMB RFS BIRTH PROCESS AT A GIVEN TIME k

$$r(\ell) = 0.09,\ p_B(x,\ell) = \mathcal{N}(x;\ m_B(\ell)\,,P_B)$$

$$P_B = \mathrm{diag}\big(10^6, 10^4, 10^6, 10^4\big)$$

ℓ	$m_B(\ell)$
$(k,\ 1)$	$[0,\ 0,\ 40000,\ 0]^\top$
$(k,\ 2)$	$[0,\ 0,\ 25000,\ 0]^\top$
$(k,\ 3)$	$[0,\ 0,\ 5000,\ 0]^\top$
$(k,\ 4)$	$[5000,\ 0,\ 0,\ 0]^\top$
$(k,\ 5)$	$[25000,\ 0,\ 0,\ 0]^\top$
$(k,\ 6)$	$[36000,\ 0,\ 0,\ 0]^\top$
$(k,\ 7)$	$[50000,\ 0,\ 15000,\ 0]^\top$
$(k,\ 8)$	$[50000,\ 0,\ 40000,\ 0]^\top$
$(k,\ 9)$	$[40000,\ 0,\ 50000,\ 0]^\top$
$(k,\ 10)$	$[10000,\ 0,\ 50000,\ 0]^\top$

measurement functions:

$$h^{(i)}(x) = \begin{cases} \angle[\big(p_x - x^{(i)}\big) + j\,\big(p_y - y^{(i)}\big)], & \text{DOA} \\ \sqrt{\big(p_x - x^{(i)}\big)^2 + \big(p_y - y^{(i)}\big)^2}, & \text{TOA} \end{cases}$$

where $(x^{(i)}, y^{(i)})$ represents the known position of sensor i. The standard deviation of DOA and TOA measurement noises are taken respectively as $\sigma_{DOA} = 1^\circ$ and $\sigma_{TOA} = 100\,m$. Each sensor has a uniform clutter spatial distribution over the surveillance area. Due to the non linearity of the sensor models, the *Unscented Kalman Filter* (UKF) [51] is used to update means and covariances of the Gaussian mixture components.

Three different scenarios with various Poisson clutter rates λ_c and constant detection probabilities P_D are considered:

- **High SNR**: $\lambda_c = 5$, $P_D = 0.99$. These parameters were used in [27] and, therefore, will be used as a first comparison test.
- **Low SNR**: $\lambda_c = 15$, $P_D = 0.99$. These parameters emulate a realistic a scenario characterized by high clutter rate.

- **Low $\mathbf{P_D}$**: $\lambda_c = 5$, $P_D = 0.7$. These parameters test the distributed algorithms in the presence of severe misdetection.

Multi-object tracking performance is evaluated in terms of the *Optimal SubPattern Assignment* (OSPA) metric [52] with Euclidean distance, i.e. $p = 2$, and cutoff $c = 600\,m$. The reported metric is averaged over 100 Monte Carlo trials for the same object trajectories but different, independently generated, clutter and measurement noise realizations. The duration of each simulation trial is fixed to $1000\,s$ (200 samples).

The Consensus Mδ-GLMB and the Consensus LMB filters are limited to 3000 hypotheses and are coupled with the *parallel CPHD look ahead strategy* described in [37], [38]. The CPHD filter is similarly limited to the same number of components through pruning and merging of mixture components [45].

The parameter setting used in [27] for the Consensus CPHD filter has been adopted for the present simulation campaigns. In particular, the survival probability is $P_S = 0.99$; the maximum number of Gaussian components is $N_{max} = 25$; the merging threshold is $\gamma_m = 4$; the truncation threshold is $\gamma_t = 10^{-4}$; the extraction threshold is $\gamma_e = 0.5$; the birth intensity function is the PHD of the LMB RFS of Table V.

$N = 1$ and $N = 3$ consensus steps have been considered for the simulations. The choice $N = 1$ is clearly the most critical one for tracking performance due to the minimal amount of information exchanged during consensus, but at the same time the most parsimonious in terms of data communication load. On the other hand, the choice $N = 3$ (the diameter of the network, i.e. the maximum distance between any two nodes in the network) allows to show the benefits of performing multiple consensus steps in terms of performance gain. The case $N = 1$ is certainly the most interesting one for comparing the various multi-object consensus filters as it highlights the capability of the proposed fusion technique to provide satisfactory results with little information exchanged and fused; the case $N = 3$ is interesting to understand how many consensus steps are needed to achieve comparable performance to the centralized setting (i.e. the Mδ-GLMB filter) where measurements from all the sensors are recursively processed by a single tracker.

A. High SNR

Figs. 4, 5 and 6 display the statistics (mean and standard deviation) of the estimated number of objects obtained, respectively, with the Consensus CPHD, the Consensus LMB and the Consensus Mδ-GLMB filters. Observe that all three distributed algorithms estimate the object cardinality accurately, with the Consensus Mδ-GLMB exhibiting the best cardinality estimate (with least variance). Note that the difficulties introduced by the rendezvous point (e.g. merged or lost tracks) are correctly addressed by

all three distributed algorithms. Performing $N = 3$ consensus steps provides a significant improvement only to the cardinality estimation of the Consensus CPHD filter.

Fig. 7 shows the OSPA distance for the three algorithms. Compared to Consensus CPHD, the improved localization performance of the Consensus LMB and the Consensus Mδ-GLMB is attributed to two factors: (a) the "spooky effect" [32] causes the Consensus CPHD filter to temporarily drop tracks which are subjected to missed detections and to declare multiple estimates for existing tracks in place of the dropped tracks, and (b) the two tracking filters are generally able to better localize objects due to a more accurate propagation of the posterior density. Note that Consensus LMB and Consensus Mδ-GLMB filters exhibit similar performance since the additional approximation in the LMB filter (see (63)-(64)) is not significant in high SNR. Multiple consensus steps provide a remarkable performance gain in terms of state estimation error. As it can be seen from Figs. 4, 5 and 6, the cardinality is, on average, correctly estimated by all sensors in all Monte Carlo trials. Thus, the main contributor to the OSPA error reduction is the state estimation error. The object births and deaths are responsible for the peaks of the OSPA error in the distributed algorithms. The peaks are not present in the Mδ-GLMB because it makes use of all measurements provided by the sensors at each time interval k, while the distributed trackers only use the local measurements and require a few fusion steps in order to properly estimate the states of the objects. It is worth noticing that with $N = 3$ the OSPA error of the Consensus LMB and Consensus Mδ-GLMB after each peak is very close to the one of the Mδ-GLMB.

B. Low SNR

Figs. 8 and 9 display the statistics (mean and standard deviation) of the estimated number of objects obtained, respectively, with the Consensus CPHD and the Consensus Mδ-GLMB. Observe that these two distributed filters estimate the object cardinality accurately, with the Consensus Mδ-GLMB exhibiting again better cardinality estimate (with lower variance).

Note that the Consensus LMB filter fails to track the objects. The problem is due to the approximation of the GLMB posteriors by LMBs, which becomes more severe with low SNR. In particular, each local tracker fails to properly capture the existence probability of the tracks due to three main factors: (a) no local observability, (b) high clutter rate and (c) loss of the full posterior cardinality distribution after the LMB approximation. Having low existence probabilities, the extraction of the tracks fails even if the single object densities are correctly propagated in time.

Fig. 10 shows the OSPA distance for the current scenario. As in the previous case study, the Consensus Mδ-GLMB filter outperforms the Consensus CPHD filter. The same conclusion as in the previous case (High SNR) can be drawn for multiple consensus steps.

C. Low P_D

Fig. 11 displays the statistics (mean and standard deviation) of the estimated number of objects obtained with the Consensus Mδ-GLMB. It is worth noting that in this very challenging scenario with $P_D = 0.7$, the only working distributed algorithm is indeed the Consensus Mδ-GLMB filter, and that it exhibits good performance in terms of the average number of estimated objects. Fig. 12 shows the OSPA distance for the current scenario.

The benefit of using multiple consensus steps is particularly stressed by this simulation setting. As it can be seen from Figs. 11 and 12, there is a remarkable improvement in both cardinality and state estimation error. Further, once the peaks in the OSPA reduce, the error is comparable to the (centralized) Mδ-GLMB filter.

VI. Conclusions

In this paper, we have presented fully distributed multi-object tracking solutions over a sensor network using labeled RFSs. Consensus algorithms have been developed for fully distributed and scalable fusion of information collected from the multiple heterogeneous and geographically dispersed sensors. The proposed consensus algorithms are based on the notion of Kullback-Leibler averaging of the local multi-object probability densities. Efficient Gaussian mixture implementations have been successfully tested on realistic multi-object tracking scenarios. Possible topics for future work are to consider sensors with different field-of-view and to investigate distributed measurement-driven object initialization.

VII. Appendix A

Proof of Proposition 1:

Let $\eta^{(L)}(\ell) \triangleq \int \prod_{i\in\mathcal{I}} (p^{(i)}(x,\ell;L))^{\omega^{(i)}} dx$, and note from the definitions of $\overline{w}(L)$ and $\overline{p}(\cdot,L)$ in Proposition 1 that

$$\overline{w}(L) = \frac{\prod_{i\in\mathcal{I}} (w^{(i)}(L))^{\omega^{(i)}} [\eta^{(L)}]^L}{\sum_{J\in\mathcal{F}(\mathbb{L})} \prod_{i\in\mathcal{I}} (w^{(i)}(J))^{\omega^{(i)}} [\eta^{(J)}]^J} \tag{68}$$

$$\eta^{(L)}(\ell)\, \overline{p}(x,\ell;L) = \prod_{i\in\mathcal{I}} p^{(i)}(x,\ell;L) \tag{69}$$

Using the form (22) for Mδ-GLMB densities we have

$$\begin{aligned}\prod_{i\in\mathcal{I}}(\boldsymbol{\pi}^{(i)}(\mathbf{X}))^{\omega^{(i)}} &= \Delta(\mathbf{X})\prod_{i\in\mathcal{I}}(w^{(i)}(\mathcal{L}(\mathbf{X})))^{\omega^{(i)}}\prod_{i\in\mathcal{I}}\left([p^{(i)}(\cdot)]^{\mathbf{X}}\right)^{\omega^{(i)}}\\ &= \Delta(\mathbf{X})\prod_{i\in\mathcal{I}}(w^{(i)}(\mathcal{L}(\mathbf{X})))^{\omega^{(i)}}\left[\prod_{i\in\mathcal{I}}(p^{(i)}(\cdot))^{\omega^{(i)}}\right]^{\mathbf{X}}\\ &= \Delta(\mathbf{X})\prod_{i\in\mathcal{I}}(w^{(i)}(\mathcal{L}(\mathbf{X})))^{\omega^{(i)}}[\eta^{(\mathcal{L}(\mathbf{X}))}]^{\mathcal{L}(\mathbf{X})}[\overline{p}(\cdot;\mathcal{L}(\mathbf{X}))]^{\mathbf{X}} \quad (70)\\ &= \sum_{J\in\mathcal{F}(\mathbb{L})}\prod_{i\in\mathcal{I}}(w^{(i)}(J))^{\omega^{(i)}}[\eta^{(J)}]^{J}\Delta(\mathbf{X})\delta_{J}(\mathcal{L}(\mathbf{X}))[\overline{p}(\cdot;J)]^{\mathbf{X}} \quad (71)\end{aligned}$$

where (70) follows by substituting (69).

Integrating (71), applying Lemma 3 of [37, Section III.B], and noting that $\int \overline{p}(\cdot,\ell,J)\,dx = 1$ gives

$$\begin{aligned}\int\prod_{i\in\mathcal{I}}(\boldsymbol{\pi}^{(i)}(\mathbf{X}))^{\omega^{(i)}}\delta\mathbf{X} &= \sum_{J\in\mathcal{F}(\mathbb{L})}\prod_{i\in\mathcal{I}}(w^{(i)}(J))^{\omega^{(i)}}[\eta^{(J)}]^{J}\int\Delta(\mathbf{X})\delta_{J}(\mathcal{L}(\mathbf{X}))[\overline{p}(\cdot;J)]^{\mathbf{X}}\delta\mathbf{X}\\ &= \sum_{J\in\mathcal{F}(\mathbb{L})}\prod_{i\in\mathcal{I}}(w^{(i)}(J))^{\omega^{(i)}}[\eta^{(J)}]^{J}\sum_{L\in\mathcal{F}(\mathbb{L})}\delta_{J}(L)\\ &= \sum_{J\in\mathcal{F}(\mathbb{L})}\prod_{i\in\mathcal{I}}(w^{(i)}(J))^{\omega^{(i)}}[\eta^{(J)}]^{J}. \quad (72)\end{aligned}$$

Dividing (70) by (72), and using (68) yields

$$\frac{\prod_{i\in\mathcal{I}}\left(\boldsymbol{\pi}^{(i)}(\mathbf{X})\right)^{\omega^{(i)}}}{\int\prod_{i\in\mathcal{I}}\left(\boldsymbol{\pi}^{(i)}(\mathbf{X})\right)^{\omega^{(i)}}\delta\mathbf{X}} = \Delta(\mathbf{X})\overline{w}(\mathcal{L}(\mathbf{X}))\,[\overline{p}(\cdot;\mathcal{L}(\mathbf{X}))]^{\mathbf{X}}$$

which is an Mδ-GLMB with parameter set $\{(\overline{w}(L),\overline{p}(\cdot;L))\}_{L\in\mathcal{F}(\mathbb{L})}$. Finally, using the equivalence between the KLA and the normalized geometric mean in Theorem 1 completes the proof.

VIII. Appendix B

Proof of Proposition 2:

Let

$$\eta(\ell) \triangleq \int\prod_{i\in\mathcal{I}}(p^{(i)}(x,\ell))^{\omega^{(i)}}dx \quad (73)$$

$$\widetilde{q}(\ell) \triangleq \prod_{i\in\mathcal{I}}(1-r^{(i)}(\ell))^{\omega^{(i)}} \quad (74)$$

$$\widetilde{r}(\ell) \triangleq \eta(\ell)\prod_{i\in\mathcal{I}}(r^{(i)}(\ell))^{\omega^{(i)}} \quad (75)$$

Note from the definitions of $\overline{r}(\ell)$ and $\overline{p}(x,\ell)$ in Proposition 2 that

$$\eta(\ell)\,\overline{p}(x,\ell) = \prod_{i\in\mathcal{I}}(p^{(i)}\,(x,\ell))^{\omega^{(i)}}, \tag{76}$$

$$\overline{r}(\ell) = \frac{\widetilde{r}(\ell)}{\widetilde{q}(\ell)+\widetilde{r}(\ell)}\,, 1-\overline{r}(\ell) = \frac{\widetilde{q}(\ell)}{\widetilde{q}(\ell)+\widetilde{r}(\ell)}\,. \tag{77}$$

Using the form (25) for LMB densities we have

$$\begin{aligned}
\prod_{i\in\mathcal{I}}(\boldsymbol{\pi}^{(i)}\,(\mathbf{X}))^{\omega^{(i)}} &= \Delta(\mathbf{X})\prod_{i\in\mathcal{I}}\Big([1-r^{(i)}]^{\mathbb{L}\backslash\mathcal{L}(\mathbf{X})}[1_{\mathbb{L}}\,r^{(i)}]^{\mathcal{L}(\mathbf{X})}\Big)^{\omega^{(i)}}\prod_{i\in\mathcal{I}}[(p^{(i)})^{\mathbf{X}}]^{\omega^{(i)}}\\
&= \Delta(\mathbf{X})\prod_{i\in\mathcal{I}}\Big([1-r^{(i)}]^{\mathbb{L}\backslash\mathcal{L}(\mathbf{X})}[1_{\mathbb{L}}\,r^{(i)}]^{\mathcal{L}(\mathbf{X})}\Big)^{\omega^{(i)}}\left[\prod_{i\in\mathcal{I}}(p^{(i)})^{\omega^{(i)}}\right]^{\mathbf{X}}\\
&= \Delta(\mathbf{X})\left[\prod_{i\in\mathcal{I}}(1-r^{(i)})^{\omega^{(i)}}\right]^{\mathbb{L}\backslash\mathcal{L}(\mathbf{X})}\left[1_{\mathbb{L}}\prod_{i\in\mathcal{I}}(r^{(i)})^{\omega^{(i)}}\right]^{\mathcal{L}(\mathbf{X})}\eta^{\mathcal{L}(\mathbf{X})}\overline{p}^{\mathbf{X}}\\
&= \Delta(\mathbf{X})\,[\widetilde{q}]^{\mathbb{L}\backslash\mathcal{L}(\mathbf{X})}\left[1_{\mathbb{L}}\left(\eta(\cdot)\prod_{i\in\mathcal{I}}(r^{(i)})^{\omega^{(i)}}\right)\right]^{\mathcal{L}(\mathbf{X})}\overline{p}^{\mathbf{X}}\\
&= \Delta(\mathbf{X})\,[\widetilde{q}]^{\mathbb{L}\backslash\mathcal{L}(\mathbf{X})}\,[1_{\mathbb{L}}\,\widetilde{r}]^{\mathcal{L}(\mathbf{X})}\overline{p}^{\mathbf{X}}
\end{aligned} \tag{78}$$

where the substitutions (76), (74) and (75) have been performed.

Integrating (78), applying Lemma 3 of [37, Section III.B] and noting that $\int\overline{p}(x,\cdot)dx = 1$ gives

$$\begin{aligned}
\int\prod_{i\in\mathcal{I}}(\boldsymbol{\pi}^{(i)}(\mathbf{X}))^{\omega^{(i)}}\delta\mathbf{X} &= \sum_{L\in\mathcal{F}(\mathbb{L})}[\widetilde{q}]^{\mathbb{L}\backslash L}\,[\widetilde{r}]^{L}\\
&= [\widetilde{q}+\widetilde{r}]^{\mathbb{L}}
\end{aligned} \tag{79}$$

where in the last step we applied the Binomial Theorem [53]

$$\sum_{L\subseteq\mathbb{L}} g^{\mathbb{L}\backslash L}f^{L} = [g+f]^{\mathbb{L}}.$$

Dividing (78) by (79), and using (77) yields

$$\begin{aligned}
\frac{\prod_{i\in\mathcal{I}}\left(\boldsymbol{\pi}^{(i)}\,(\mathbf{X})\right)^{\omega^{(i)}}}{\int\prod_{i\in\mathcal{I}}\left(\boldsymbol{\pi}^{(i)}\,(\mathbf{X})\right)^{\omega^{(i)}}\delta\mathbf{X}} &= \Delta(\mathbf{X})\frac{(\widetilde{q})^{\mathbb{L}\backslash\mathcal{L}(\mathbf{X})}\,(1_{\mathbb{L}}\,\widetilde{r})^{\mathcal{L}(\mathbf{X})}}{(\widetilde{q}+\widetilde{r})^{\mathbb{L}}}\overline{p}^{\mathbf{X}}\\
&= \Delta(\mathbf{X})\frac{(\widetilde{q})^{\mathbb{L}\backslash\mathcal{L}(\mathbf{X})}\,(1_{\mathbb{L}}\,\widetilde{r})^{\mathcal{L}(\mathbf{X})}}{(\widetilde{q}+\widetilde{r})^{\mathbb{L}\backslash\mathcal{L}(\mathbf{X})}\,(\widetilde{q}+\widetilde{r})^{\mathcal{L}(\mathbf{X})}}\overline{p}^{\mathbf{X}}\\
&= \Delta(\mathbf{X})\left(\frac{\widetilde{q}}{\widetilde{q}+\widetilde{r}}\right)^{\mathbb{L}\backslash\mathcal{L}(\mathbf{X})}\left(1_{\mathbb{L}}\,\frac{\widetilde{r}}{\widetilde{q}+\widetilde{r}}\right)^{\mathcal{L}(\mathbf{X})}\overline{p}^{\mathbf{X}}\\
&= \Delta(\mathbf{X})\,(1-\overline{r})^{\mathbb{L}\backslash\mathcal{L}(\mathbf{X})}\,(1_{\mathbb{L}}\,\overline{r})^{\mathcal{L}(\mathbf{X})}\,\overline{p}^{\mathbf{X}}
\end{aligned}$$

which is an LMB with parameter set $\{(\overline{r}(\ell), \overline{p}(\ell))\}_{\ell \in \mathcal{F}(\mathbb{L})}$. Finally, using the equivalence between the KLA and the normalized geometric mean in Theorem 1 completes the proof.

REFERENCES

[1] A. Farina and F. A. Studer, *Radar data processing, vol. I: introduction and tracking*. Research Studies Press, Letchworth, Hertfordshire, England, 1985.

[2] A. Farina and F. A. Studer, *Radar data processing, vol. II: advanced topics and applications*. Research Studies Press, Letchworth, Hertfordshire, England, 1986.

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

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

[5] I. Goodman, R. Mahler, and H. Nguyen, *Mathematics of Data Fusion*. Kluwer Academic Publishers, 1997.

[6] S. Blackman and R. Popoli, *Design and analysis of modern tracking systems*. Artech House, Norwood, MA, USA, 1999.

[7] R. Mahler, *Statistical multisource-multitarget information fusion*. Artech House, Norwood, MA, USA, vol. 685, 2007.

[8] B.-N. Vo, M. Mallick, Y. Bar-Shalom, S. Coraluppi, R. Osborne III, R. Mahler, and B.-T. Vo, "Multitarget tracking," *Wiley Encyclopedia of Electrical and Electronics Engineering*, Wiley, Sept. 2015.

[9] D. Reid, "An algorithm for tracking multiple targets," *IEEE Trans. Automatic Control*, vol. 24, no. 6, pp. 843-854, 1979.

[10] Special issue on Distributed Signal Processing in Sensor Networks, *IEEE Signal Processing Magazine*, vol. 23, no. 4, 2006.

[11] R. Olfati-Saber, J.A. Fax and R. Murray, "Consensus and cooperation in networked multi-agent systems," *Proc. of the IEEE*, vol. 95, no. 1, pp. 215-233, 2007.

[12] L. Xiao, S. Boyd and S. Lall, "A scheme for robust distributed sensor fusion based on average consensus," *Proc. 4th Int. Symp. on Information Processing in Sensor Networks*, pp. 63-70, Los Angeles, CA, 2005.

[13] R. Olfati-Saber, "Distributed Kalman filtering for sensor networks," *Proc. 46th IEEE Conf. on Decision and Control*, pp. 5492-5498, 2007.

[14] M. Kamgarpour and C. Tomlin, "Convergence properties of a decentralized Kalman filter," *Proc. 47th IEEE Conf. on Decision and Control*, pp. 3205-3210, 2008.

[15] R. Carli, A. Chiuso, L. Schenato, and S. Zampieri, "Distributed Kalman filtering based on consensus strategies," *IEEE Journal on Selected Areas in Communications*, vol. 26, pp. 622-633, 2008.

[16] G.C. Calafiore and F. Abrate, "Distributed linear estimation over sensor networks," *Int. J. of Control*, vol. 82, no. 5, pp. 868-882, 2009.

[17] S.S. Stankovic, M.S. Stankovic, and D.M. Stipanovic, "Consensus based overlapping decentralized estimation with missing observations and communication faults," *Automatica*, vol. 45, no. 6, pp. 1397-1406, 2009.

[18] F.S. Cattivelli and A.H. Sayed, "Diffusion strategies for distributed Kalman filtering and smoothing," *IEEE Trans. on Automatic Control*, vol. 55, pp. 2069-2084, 2010.

[19] M. Farina, G. Ferrari-Trecate and R. Scattolini, "Distributed moving horizon estimation for linear constrained systems," *IEEE Trans. on Automatic Control*, vol. 55, no. 11, pp. 2462-2475, 2010.

[20] G. Battistelli, and L. Chisci, "Kullback-Leibler average, consensus on probability densities, and distributed state estimation with guaranteed stability," *Automatica*, vol. 50, no. 3, pp. 707-718, 2014.

[21] T.M. Cover and J.A. Thomas. *Elements of information theory*. John Wiley & Sons, Hoboken, NJ, USA, 2012.

[22] K.C. Chang, C.-Y. Chong and S. Mori, "Analytical and computational evaluation of scalable distributed fusion algorithms," *IEEE Trans. on Aerospace and Electronic Systems*, vol. 46, no. 4, pp. 2022-2034, 2010.

[23] R. Mahler, "Optimal/robust distributed data fusion: a unified approach," *Proc. of SPIE*, vol. 4052 Signal Processing, Sensor Fusion, and Target Recognition IX, Orlando, FL, USA, April 2000.

[24] M.B. Hurley, "An information-theoretic justification for covariance intersection and its generalization", *Proc. 5th Int. Conf. on Information Fusion*, pp. 505-511, 2002.

[25] S.J. Julier and J.K. Uhlmann, "A non-divergent estimation algorithm in the presence of unknown correlations," *Proc. of the IEEE American Control Conf. (ACC 1997)*, vol. 4, pp. 2369-2373, 1997.

[26] S.J. Julier, "Fusion without independence," *IET Seminar on Target Tracking and Data Fusion: Algorithms and Applications*, pp. 3-4, Birmingham, UK, 2008.

[27] 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.

[28] G. Battistelli, L. Chisci, C. Fantacci, A. Farina, R. Mahler, "Distributed fusion of multitarget densities and consensus PHD/CPHD filters", *SPIE 9474 - Signal Processing, Sensor/Information Fusion, and Target Recognition*, XXIV, 94740E, 2015.

[29] R. Mahler, "Toward a theoretical foundation for distributed fusion," in *Distributed Data Fusion for Network-Centric Operations* (D. Hall, M. Liggins II, C.-Y. Chong, and J. Llinas, eds.), ch. 8, CRC Press, Boca Raton, FL, USA, 2012.

[30] 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.

[31] G. Battistelli, L. Chisci, C. Fantacci, N. Forti, A. Farina and A. Graziano, "Distributed peer-to-peer multitarget tracking with association-based track fusion," *Proc. of the 17th Int. Conf. on Information Fusion (FUSION'14)*, Salamanca, Spain, 7-10 July, 2014.

[32] D. Franken, M. Schmidt and M. Ulmke, ""Spooky action at a distance in the cardinalized probability hypothesis density filter," *IEEE Trans. on Aerospace and Electronic Systems*, vol. 45, no. 4, pp. 1657-1664, 2009.

[33] B.-T. Vo, B.-N. Vo, and A. Cantoni, "The cardinality balanced multi-target multi-Bernoulli filter and its implementations," *IEEE Trans. Signal Processing*, vol. 57, no. 2, pp. 409–423, 2009.

[34] B.-N. Vo, B.-T. Vo, N.-T. Pham, and D. Suter, "Joint detection and estimation of multiple objects from image observations," *IEEE Trans. on Signal Processing*, vol. 58, no. 10, pp. 5129–5241, 2010.

[35] R. Mahler, "Multi-target Bayes filtering via first-order multi-target moments," *IEEE Trans. on Aerospace and Electronic Systems*, vol. 39, no. 4, pp. 1152-1178, 2003.

[36] R. Mahler, "PHD filters of higher order in target number," *IEEE Trans. Aerospace & Electronic Systems*, vol. 43, no. 3, July 2007.

[37] B.-T. Vo and B.-N. Vo, "Labeled random finite sets and multi-object conjugate priors," *IEEE Trans. on Signal Processing*, vol. 61, no. 13, pp. 3460-3475, 2013.

[38] B.-N. Vo, B.-T. Vo and D. Phung, "Labeled random finite sets and the Bayes multi-target tracking filter", *IEEE Trans. on Signal Processing*, vol. 62, no. 24, pp. 6554-6567, 2014.

[39] C. Fantacci, B.-T. Vo, F. Papi and B.-N Vo, "The marginalized δ-GLMB filter", *preprint available online at arXiv:1501.00926*, 2015.

[40] S. Reuter, B.-T. Vo, B.-N. Vo and K. Dietrmayer, "The labeled multi-Bernoulli filter," *IEEE Trans. on Signal Processing*, vol. 62, no.12, pp. 3246-3260, 2014.

[41] G. Battistelli, L. Chisci and C. Fantacci, "Parallel consensus on likelihoods and priors for networked nonlinear filtering," *IEEE Signal Processing Letters*, vol. 21, no. 7, pp. 787-791, 2014.

[42] B.-N. Vo, S. Singh and A. Doucet, "Sequential Monte Carlo methods for multi-target filtering with random finite Sets," *IEEE Trans. on Aerospace and Electronic Systems*, vol. 41, no. 4, pp. 1224-1245, 2005.

[43] F. Papi, B.-N. Vo, B.-T. Vo, C. Fantacci and M. Beard, "Generalized labeled multi-Bernoulli approximation of multi-object densities," *IEEE Trans. on Signal Processing*, vol. 63, no. 20, pp. 5487-5497, 2015.

[44] B.-N. Vo and W.K. Ma, "The Gaussian mixture probability hypothesis density filter," *IEEE Trans. on Signal Processing*, vol. 54, no. 11, pp. 4091-4104, 2006.

[45] B.-T. Vo, B.-N. Vo and A. Cantoni, "Analytic implementations of the cardinalized probability hypothesis density filter", *IEEE Trans. on Signal Processing*, vol. 55, no. 7, pp. 3553-3567, 2007.

[46] M. Günay, U. Orguner and M. Demirekler, "Approximate Chernoff fusion of Gaussian mixture using sigmapoints", *Proc. 17th Int. Conf. on Information Fusion*, Salamanca, Spain, 2014.

[47] O. Hlinka, O. Slučiak, F. Hlawatsch, P.M. Djurić and M. Rapp, "Likelihood consensus and its application to particle filtering," *IEEE Trans. on Signal Processing*, vol. 60, no. 8, pp. 4334-4349, 2012.

[48] O. Hlinka, F. Hlawatsch and P.M. Djuric, "Consensus-based distributed particle filtering with distributed proposal adaptation," *IEEE Trans. on Signal Processing*, vol. 62, no. 12, pp. 3029-3041, 2014.

[49] M. Coates, "Distributed particle filters for sensor networks," *Proceedings of the 3rd International Symposium on Information Processing in Sensor Networks (IPSN '04)*, pp. 99-107, New York, NY, USA, 2004.

[50] C. Fantacci and F. Papi, "Scalable Multi-Sensor Multi-Target Tracking using the Marginalized δ-GLMB Density", *IEEE Signal Processing Letters*, DOI:10.1109/LSP.2016.2557078, in final production stage with IEEE Publishing Operations.

[51] S.J. Julier and J.K. Uhlmann, "Unscented filtering and nonlinear estimation," *Proc. of the IEEE*, vol. 92, pp. 401-422, 2004.

[52] 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.

[53] M. Abramowitz and I. A. Stegun, *Handbook of mathematical functions: with formulas, graphs, and mathematical tables*. National Bureau of Standards, Washington, DC, 1964. Republished by Courier Dover Publications, 2012. Available free online at http://people.math.sfu.ca/∼cbm/aands/abramowitz_and_stegun.pdf.

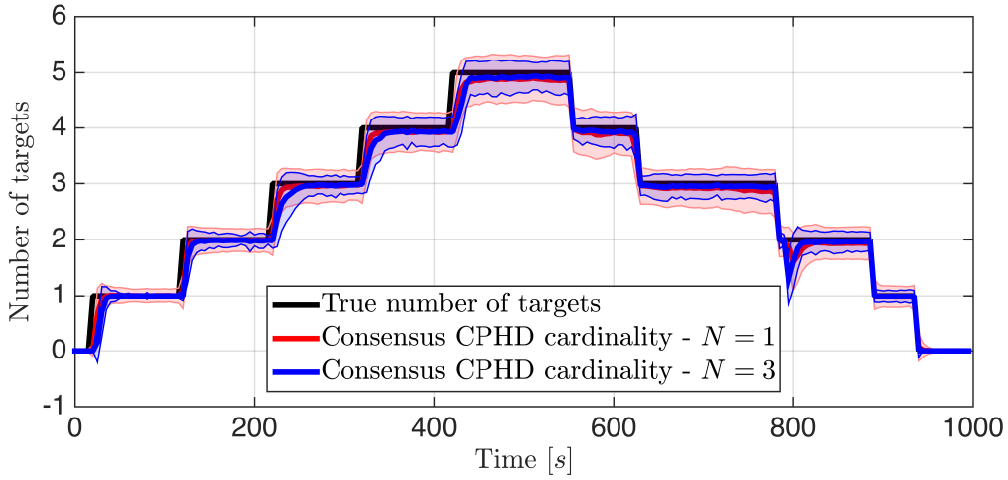

Fig. 4. Cardinality statistics for Consensus CPHD filter under high SNR.

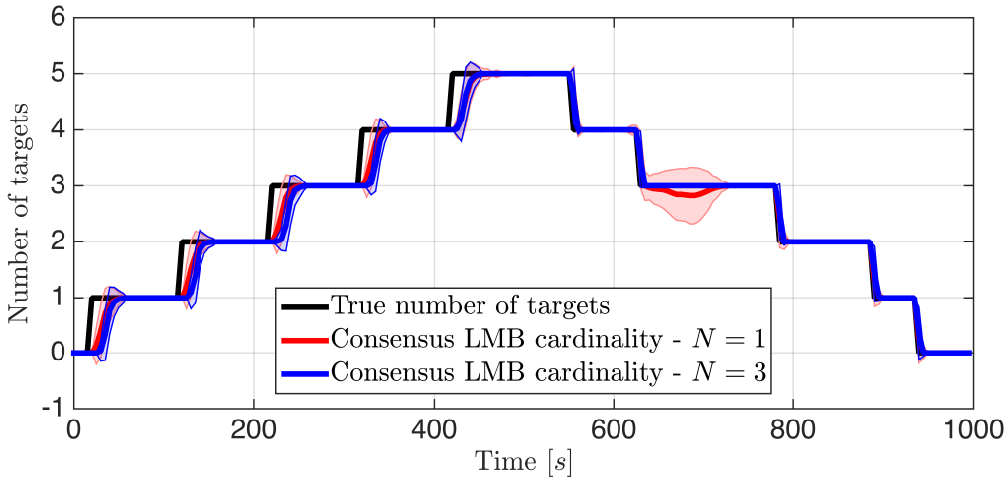

Fig. 5. Cardinality statistics for Consensus LMB tracker under high SNR.

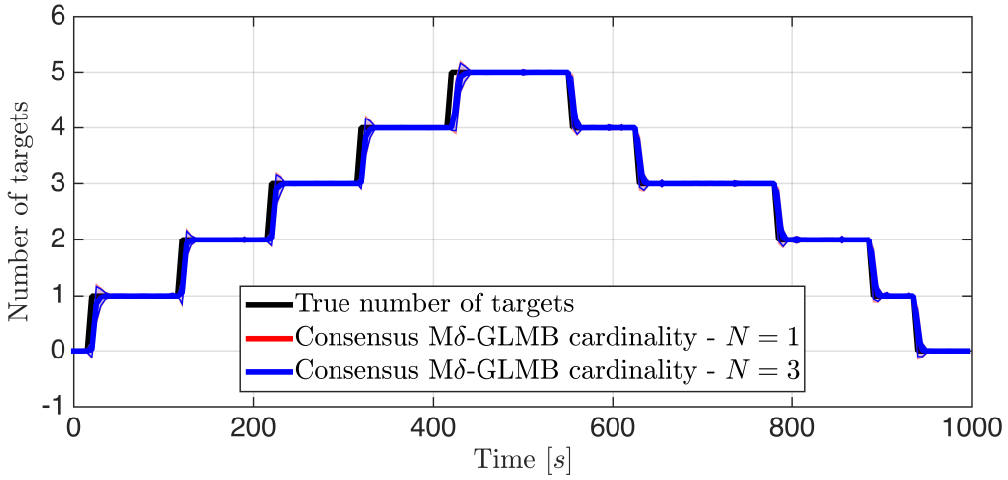

Fig. 6. Cardinality statistics for Consensus Mδ-GLMB tracker under high SNR.

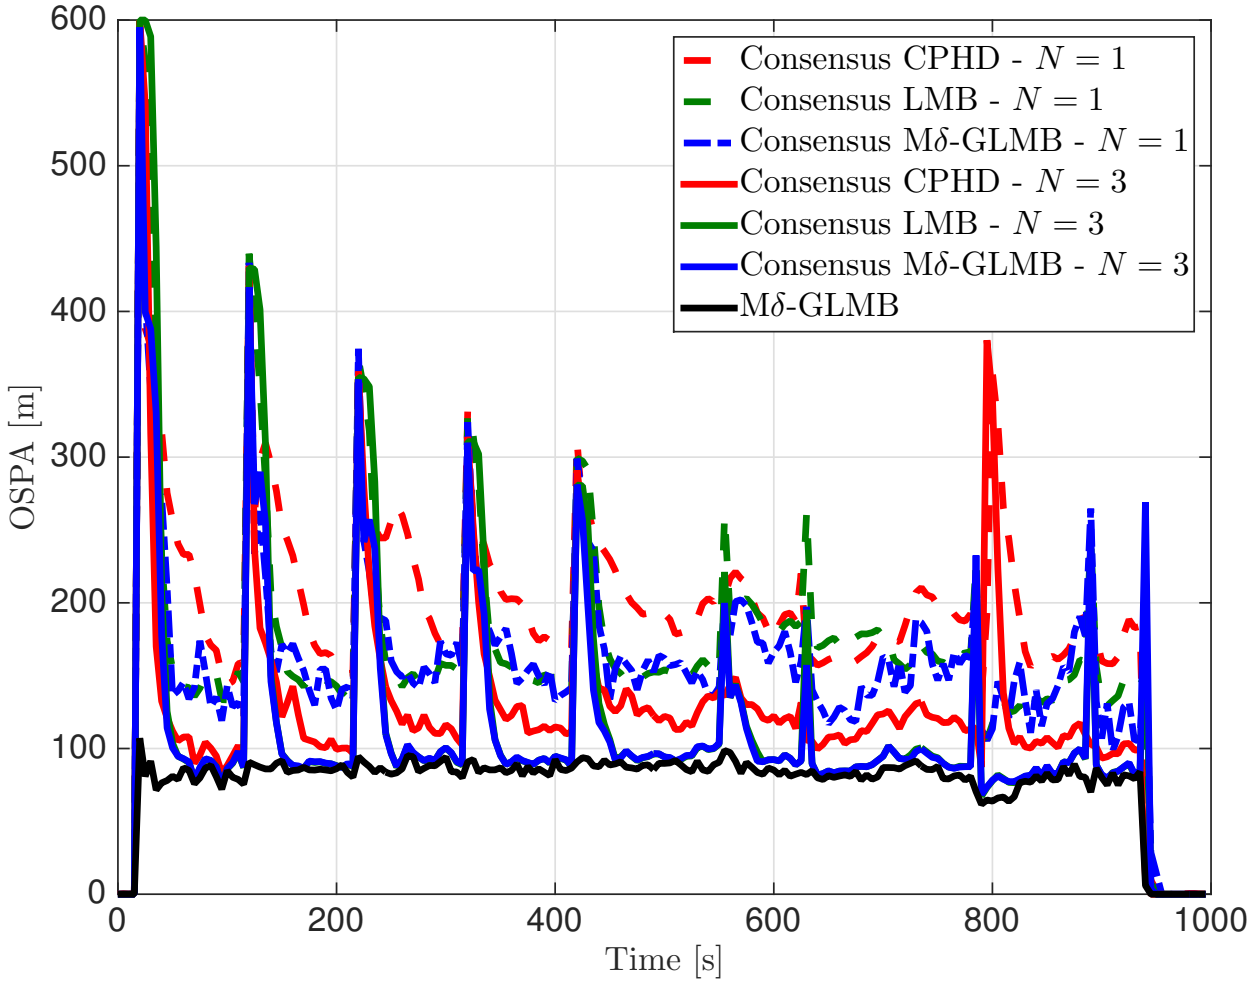

Fig. 7. OSPA distance ($c = 600\,[m]$, $p = 2$) under high SNR.

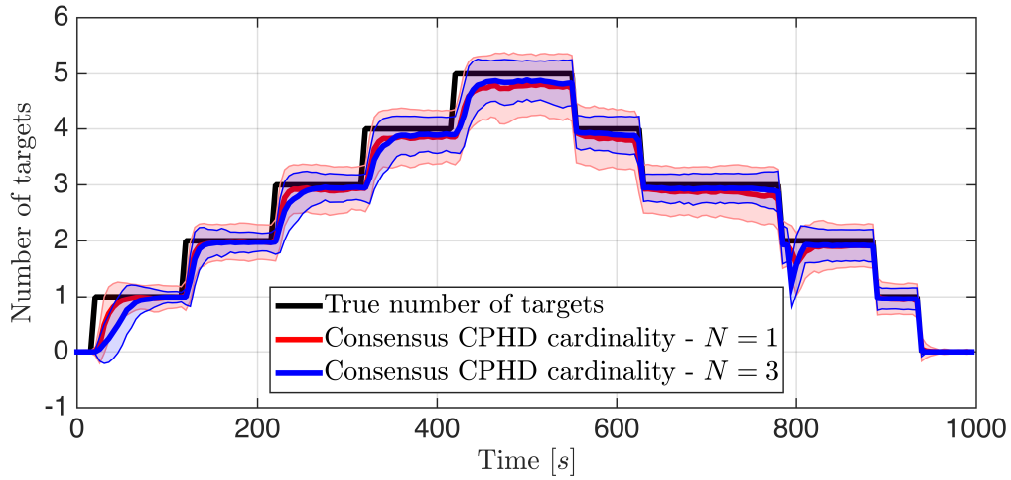

Fig. 8. Cardinality statistics for Consensus CPHD filter under low SNR.

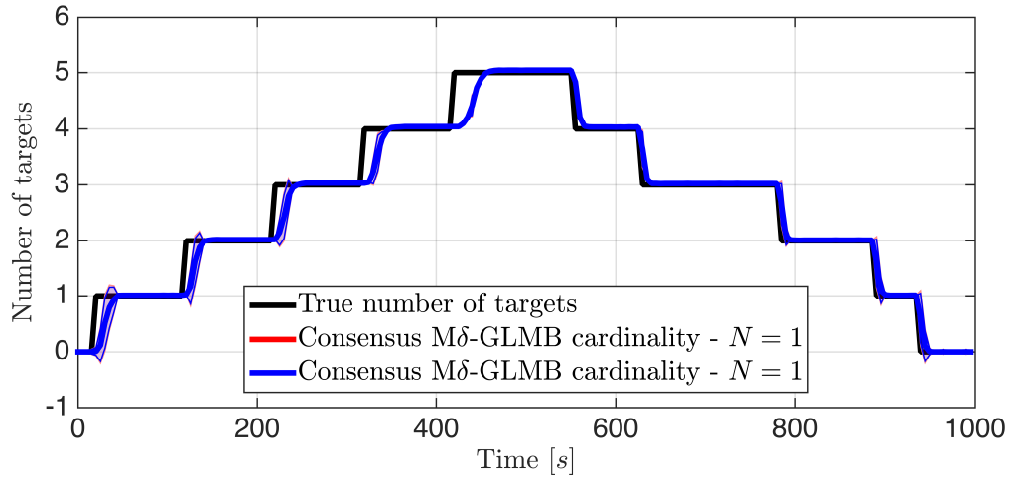

Fig. 9. Cardinality statistics for Consensus Mδ-GLMB tracker under low SNR.

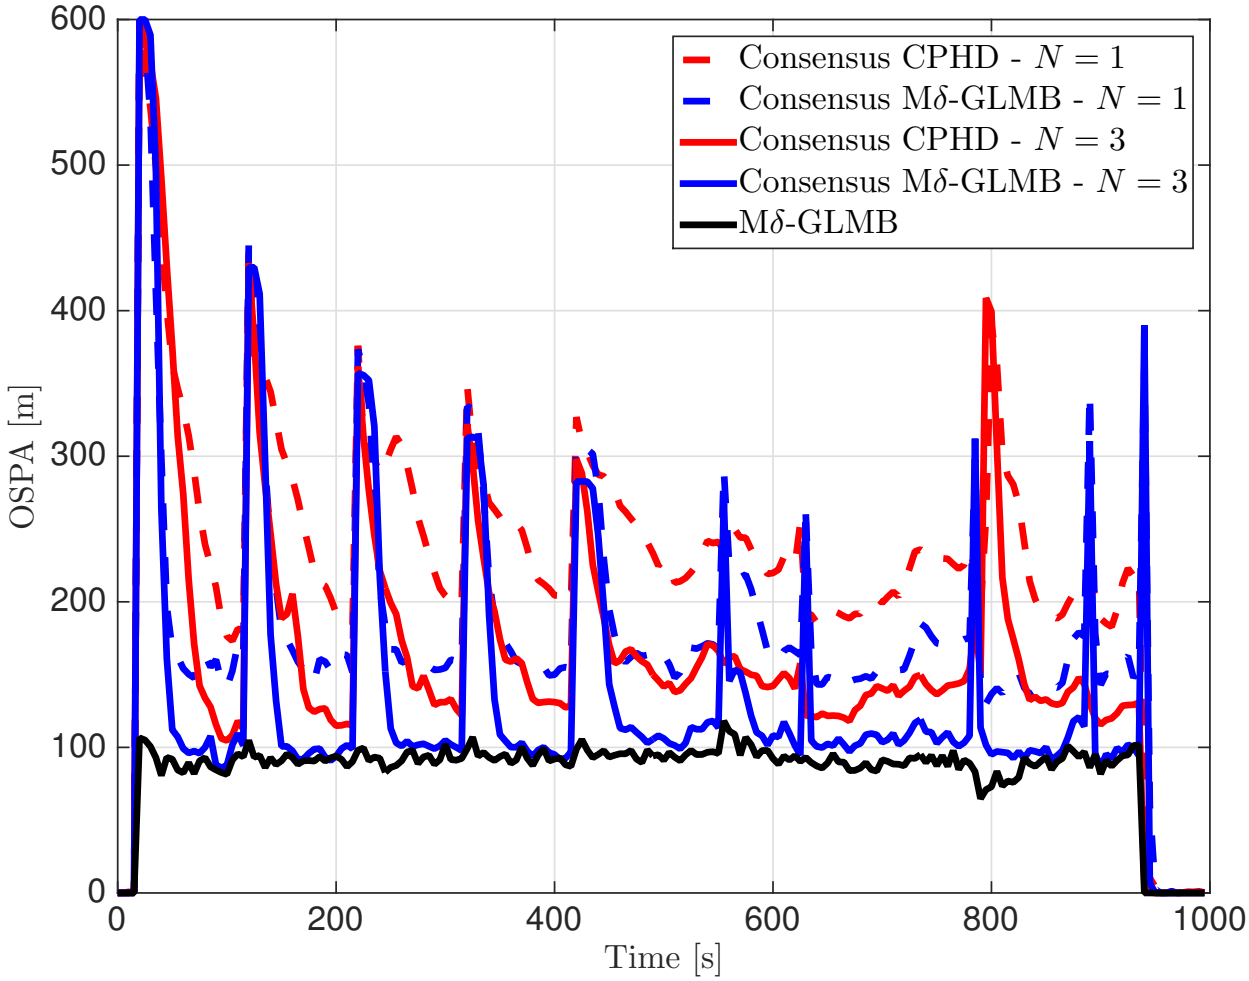

Fig. 10. OSPA distance ($c = 600\,[m]$, $p = 2$) under low SNR.

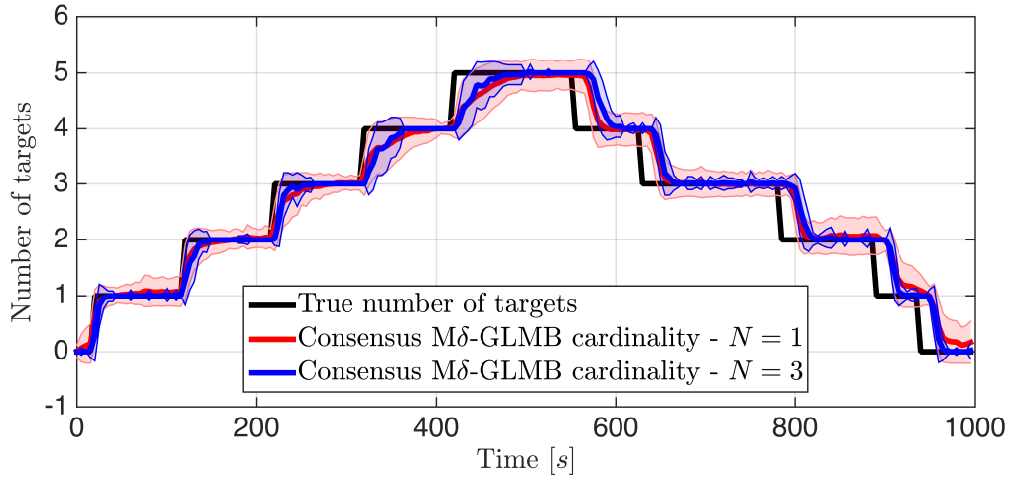

Fig. 11. Cardinality statistics for Consensus Mδ-GLMB tracker under low P_D.

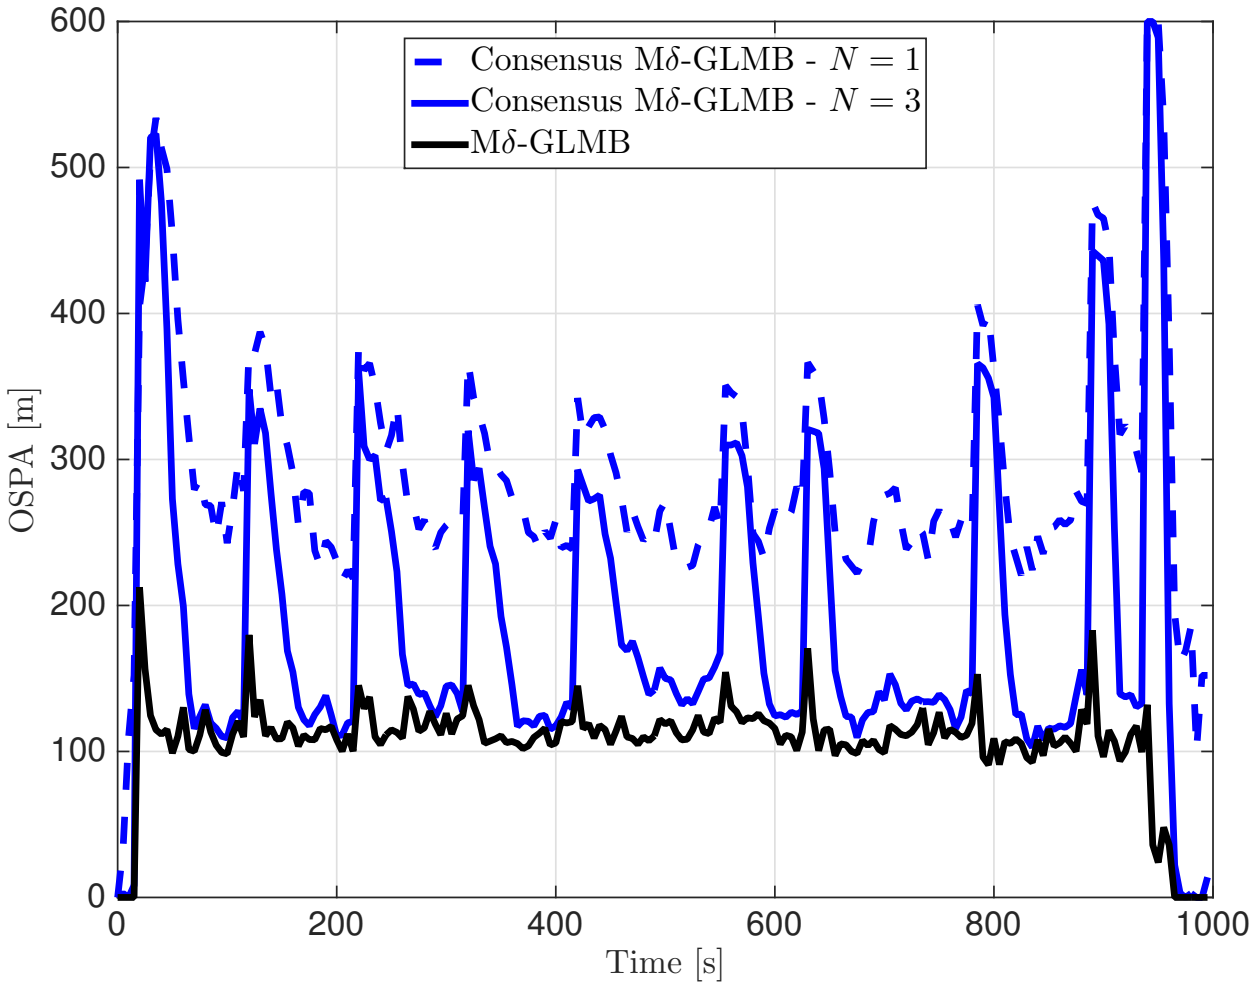

Fig. 12. OSPA distance ($c = 600\,[m]$, $p = 2$) under low P_D.