A Probabilistic Label Association Algorithm for Distributed Labeled Multi-Bernoulli Filtering

Thomas Kropfreiter and Franz Hlawatsch

Institute of Telecommunications, TU Wien, Vienna, Austria
E-mail: {thomas.kropfreiter, franz.hlawatsch}@tuwien.ac.at

***Abstract*—We consider a distributed labeled multi-Bernoulli (LMB) filter that uses the generalized covariance intersection technique for fusing the local LMB distributions. A critical aspect of such filters is to correctly associate labeled Bernoulli components describing the same object at different sensors. Here, we improve on previously proposed association schemes by introducing a *probabilistic* framework and algorithm for object (label) association. Instead of enforcing a hard association, we propose to compute association probabilities and use them in the fusion of the LMB posterior distributions. To develop our probabilistic label association scheme, we first derive a formulation of the fused multiobject distribution that involves a label association distribution. We then show that approximating the label association distribution by the product of its marginals results in a fused multiobject distribution that is again of LMB type. An efficient LMB fusion algorithm is finally obtained by using a belief propagation scheme for fast approximate marginalization and a Gaussian approximation. Simulation results demonstrate that the resulting distributed LMB filter outperforms a state-of-the-art method using hard label association.**

***Index Terms*—Multiobject tracking, multitarget tracking, labeled multi-Bernoulli filter, belief propagation, random finite set.**

I. Introduction

Multiobject tracking aims to estimate the time-dependent number and states of multiple objects from noisy and cluttered sensor measurements [1]–[4]. Many applications require tracking methods that maintain track continuity, i.e., estimate entire trajectories of consecutive object states [1], [2], [5]–[8]. A powerful and efficient method maintaining track continuity is the labeled multi-Bernoulli (LMB) filter [7], [8], which models the multiobject state by a labeled random finite set (RFS) of the LMB type. In an LMB RFS, each object is represented by a labeled Bernoulli component.

In multiobject tracking, the use of multiple sensors can enable successful operation also in challenging scenarios with many closely spaced objects, low signal-to-noise ratio, strong clutter, and/or low detection probability. Of particular interest in decentralized sensor network scenarios are distributed LMB filter methods that avoid the use of a central processing unit. Here, each sensor runs a local LMB filter and is able to communicate with a set of neighboring sensors.

For distributed fusion of statistical information among the sensors, most distributed RFS-based tracking methods use the *generalized covariance intersection* (GCI) technique [9], also known as exponential mixture density [10] and Kullback-Leibler average [11]. GCI is a suboptimal technique that fuses the posterior probability density functions (pdfs) of neighboring sensors. Two distributed LMB filters using GCI fusion of LMB posterior pdfs were recently proposed in [12], [13]. GCI fusion of LMB pdfs is challenging because it is a priori unclear which labeled Bernoulli component of one sensor (representing an object) corresponds to which labeled Bernoulli component of another sensor. In [12], it is assumed that all the local posterior pdfs are defined on the same set of labels, and Bernoulli components with identical labels are matched. However, this assumption is rarely satisfied in practice. In [13], the labeled Bernoulli components of different sensors are matched by minimizing a "label inconsistency indicator." However, in more challenging scenarios, this can still result in a significant percentage of incorrect matching events and, thus, in a poor tracking performance.

Here, we propose a GCI-based fusion method for LMB multiobject pdfs that avoids a hard association of the labeled Bernoulli components of neighboring sensors and instead uses a soft (i.e., probabilistic) association. In our approach, label association probabilities are computed and used in the fusion of the multiobject pdfs. To develop this probabilistic association scheme, we first derive a formulation of the fused posterior multiobject pdf that involves a label association probability mass function (pmf). We then show that approximating this label association pmf by the product of its marginals results in a fused posterior multiobject distribution that is again of LMB type. Next, inspired by [14], we use a belief propagation (BP) scheme to perform a fast approximate marginalization of the label association pmf. Finally, for a reduction of communication and computation requirements, we use Gaussian approximations for the spatial pdfs involved in the various multiobject pdfs. Simulation results demonstrate that a distributed LMB filter using our fusion algorithm significantly outperforms the state-of-the-art method of [13] (which uses a hard label association scheme) and can perform close to the centralized multisensor LMB filter based on the iterated-corrector approach [2], [7].

We will use the following notation. Vectors are denoted by small boldface letters (e.g., $\boldsymbol{x}$) and finite sets by capital letters (e.g., X). Randomness is indicated by a sans serif font, such

This work was supported by the Austrian Science Fund (FWF) under grant P 32055-N31.

as in x or X. We write pdfs as $f(\cdot)$ and pmfs as $p(\cdot)$. Finally, $\mathcal{N}(\cdot;\boldsymbol{\mu},\boldsymbol{\Sigma})$ designates a Gaussian pdf with mean vector $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$.

The remainder of this paper is organized as follows. Section II reviews the LMB RFS and the LMB filter. Section III introduces pairwise (two-sensor) LMB fusion with label association. Section IV develops a probabilistic label association framework for pairwise fusion. Section V presents a fast approximate algorithm for probabilistic label association based on BP, and Section VI describes an efficient Gaussian implementation. Section VII extends the pairwise fusion algorithm to achieve distributed networkwide fusion in a decentralized sensor network. Section VIII demonstrates the performance of a distributed LMB filter using the proposed fusion algorithm.

II. LMB RFS and LMB Filter

We first review the LMB RFS and some general facts about the LMB filter. In a labeled RFS $\mathsf{X}=\{(\mathsf{x}_1,\mathsf{l}_1),\ldots,(\mathsf{x}_\mathsf{n},\mathsf{l}_\mathsf{n})\}$, each element is a random tuple of the form $(\mathsf{x},\mathsf{l})\in\mathbb{R}^{N_x}\times\mathbb{L}$, where $\mathbb{L}$ is a countable set of labels [5], [6]. Note that also the cardinality of the set, $|\mathsf{X}|=\mathsf{n}$, is random. Let us denote by $L_X\triangleq\{l_1,\ldots,l_n\}$ the set of all labels of a realization X of X. The statistics of a labeled RFS X can be described by its multiobject pdf $f(X)$ [5]. An LMB RFS X, in particular, is characterized by a set of existence probabilities r_l and "spatial pdfs" $f_l(\boldsymbol{x})$, i.e., $\{(r_l,f_l(\boldsymbol{x}))\}_{l\in\mathbb{L}^*}$, with some finite label set $\mathbb{L}^*\subset\mathbb{L}$. The multiobject pdf of the LMB RFS X evaluated for a labeled finite set X with label set $L_X\subset\mathbb{L}$ is given by [7]

$$f(X)=\Delta(X)\,w(L_X)\prod_{(\boldsymbol{x},l)\in X}1_{\mathbb{L}^*}(l)\,f_l(\boldsymbol{x}). \tag{1}$$

Here, $\Delta(X)=1$ if the labels of X are distinct and $\Delta(X)=0$ otherwise; $1_{\mathbb{L}^*}(l)=1$ if $l\in\mathbb{L}^*$ and $1_{\mathbb{L}^*}(l)=0$ otherwise; and

$$w(L)\triangleq\left(\prod_{l\in L}1_{\mathbb{L}^*}(l)\,r_l\right)\prod_{l'\in\mathbb{L}^*\setminus L}(1-r_{l'}), \tag{2}$$

for any $L\subseteq\mathbb{L}$.

The LMB filter is based on the LMB RFS. The purpose of a (single-sensor) LMB filter [7] is to estimate a labeled multiobject state X_k from measurements $\mathsf{Z}_{1:k}\triangleq(\mathsf{Z}_1,\ldots,\mathsf{Z}_k)$ for all times $k=1,2,\ldots$. Here, X_k is modeled as an LMB RFS and Z_k as an unlabeled RFS. To estimate X_k, the LMB filter calculates (an approximation to) the posterior multiobject pdf $f(X_k|Z_{1:k})$, based on suitable state-transition and measurement models. The posterior pdf $f(X_k|Z_{1:k})$ is of the LMB form (1), parametrized by a k-dependent set of existence probabilities and spatial pdfs $\{(r_{k,l},f_{k,l}(\boldsymbol{x}_k))\}_{l\in\mathbb{L}_k^*}$ with $\mathbb{L}_k^*\subset\mathbb{L}_k$. Here, $\mathbb{L}_k$ is the overall set of labels specified at time k for the considered multiobject tracking problem. The calculation of $f(X_k|Z_{1:k})$ is done in a time-recursive manner by converting, at each time k, the previous LMB parameter set $\{(r_{k-1,l},f_{k-1,l}(\boldsymbol{x}_{k-1}))\}_{l\in\mathbb{L}_{k-1}^*}$ into the current LMB parameter set $\{(r_{k,l},f_{k,l}(\boldsymbol{x}_k))\}_{l\in\mathbb{L}_k^*}$ while using the current measurement Z_k. The corresponding conversion relations can be found in [7].

Here, we will consider a *distributed multisensor LMB filter* that operates in a network of S sensors $s\in\{1,\ldots,S\}$. Each sensor node s acquires a sequence of local measurements $Z_{1:k}^{(s)}$ and runs a *local LMB filter*; furthermore, it is able to communicate with a set of neighboring sensor nodes. The LMB posterior pdf at sensor s, $f(X_k|Z_{1:k}^{(s)})$, is parametrized by $\{(r_{k,l}^{(s)},f_{k,l}^{(s)}(\boldsymbol{x}_k))\}_{l\in\mathbb{L}_k^{(s)*}}$ with $\mathbb{L}_k^{(s)*}\subset\mathbb{L}_k^{(s)}$. It is convenient to express each label $l\in\mathbb{L}_k^{(s)}$ as an index triple (s,k',m), where k' is the time when the Bernoulli component was created by the local LMB filter at sensor s and m is an index that distinguishes different Bernoulli components created by the local LMB filter at the same time k'.

III. Pairwise LMB Fusion with Label Association

In what follows, we consider a fixed time k. Accordingly, to simplify the notation, the time index k will be suppressed. Let us consider two sensors $s\in\{1,2\}$. The sequences of measurements observed by these sensors up to time k will be briefly denoted as $Z^{(1)}$ and $Z^{(2)}$, respectively. The LMB posterior multiobject pdfs of the two sensors at time k, $f(X|Z^{(1)})$ and $f(X|Z^{(2)})$, are characterized by the Bernoulli parameter sets $\{(r_l^{(1)},f_l^{(1)}(\boldsymbol{x}))\}_{l\in\mathbb{L}^{(1)*}}$ and $\{(r_l^{(2)},f_l^{(2)}(\boldsymbol{x}))\}_{l\in\mathbb{L}^{(2)*}}$, respectively. Here, the label sets $\mathbb{L}^{(1)*}$ and $\mathbb{L}^{(2)*}$ as well as the overall label sets $\mathbb{L}^{(1)}$ and $\mathbb{L}^{(2)}$ are trivially different because of our label representation $l=(s,k',m)$ and the fact that s is different for the two sensors. However, we emphasize that $\mathbb{L}^{(1)*}$ and $\mathbb{L}^{(2)*}$ are usually different even when some other label indexing is used. For example, they are typically different when the local LMB filters employ measurement-driven birth models and/or different strategies for pruning Bernoulli components [15].

A. Pairwise LMB Fusion

Let us study "pairwise fusion" for our two sensors, i.e., fusion of $f(X|Z^{(1)})$ and $f(X|Z^{(2)})$ into a fused multiobject pdf $\tilde{f}(X|Z^{(1)},Z^{(2)})$. (The tilde indicates the fact that $\tilde{f}(X|Z^{(1)},Z^{(2)})$ is generally different from the true posterior pdf $f(X|Z^{(1)},Z^{(2)})$.) The GCI fusion rule is given by [12]

$$\tilde{f}(X|Z^{(1)},Z^{(2)})=\frac{1}{D}\big(f(X|Z^{(1)})\big)^{\omega}\big(f(X|Z^{(2)})\big)^{1-\omega}, \tag{3}$$

with $D\triangleq\int\big(f(X|Z^{(1)})\big)^{\omega}\big(f(X|Z^{(2)})\big)^{1-\omega}\delta X$ (here, $\int\cdot\,\delta X$ denotes the set integral [6]) and some fixed $\omega\in[0,1]$. If the label sets are equal, i.e., $\mathbb{L}^{(1)*}=\mathbb{L}^{(2)*}$, then $\tilde{f}(X|Z^{(1)},Z^{(2)})$ is again an LMB pdf [12]. However, even in that case, it is likely that some objects are described by Bernoulli components with different labels in the local LMB filters at the two sensors, and thus GCI fusion according to (3) involves the matching of Bernoulli components describing different objects. Indeed, suppose temporarily that $\mathbb{L}^{(1)*}=\mathbb{L}^{(2)*}=:\mathbb{L}^*$ and consider some label $l\in\mathbb{L}^*$. The corresponding spatial pdf $\tilde{f}_l(\boldsymbol{x})$ belonging to the fused LMB pdf $\tilde{f}(X|Z^{(1)},Z^{(2)})$ is calculated from the spatial pdfs $f_l^{(1)}(\boldsymbol{x})$ and $f_l^{(2)}(\boldsymbol{x})$ belonging to $f(X|Z^{(1)})$ and $f(X|Z^{(2)})$, respectively

as [12] $\tilde{f}_l(\boldsymbol{x}) = \big(f_l^{(1)}(\boldsymbol{x})\big)^{\omega}\big(f_l^{(2)}(\boldsymbol{x})\big)^{1-\omega}/C$ with $C \triangleq \int_{\mathbb{R}^{N_x}}\big(f_l^{(1)}(\boldsymbol{x})\big)^{\omega}\big(f_l^{(2)}(\boldsymbol{x})\big)^{1-\omega}\mathrm{d}\boldsymbol{x}$. However, if $f_l^{(1)}(\boldsymbol{x})$ describes a different object than $f_l^{(2)}(\boldsymbol{x})$, then the fused spatial pdf $\tilde{f}_l(\boldsymbol{x})$ does not describe any single object, and its meaning in the fused LMB pdf $\tilde{f}(X|Z^{(1)},Z^{(2)})$ is unclear. A similar statement can be made about the fused existence probability $\tilde{r}_l$ [12]. The matching of Bernoulli components describing different objects will generally reduce the performance of the distributed LMB filter.

B. Label Association

In the following, for conceptual simplicity, we consider the case where—consistently with our label representation $l=(s,k',m)$—all the elements of $\mathbb{L}^{(1)}$ are different from all the elements of $\mathbb{L}^{(2)}$, and thus also all the elements of $\mathbb{L}^{(1)*}$ are different from all the elements of $\mathbb{L}^{(2)*}$. (We note that our development can be straightforwardly extended to more general situations.) In that case, $\tilde{f}(X|Z^{(1)},Z^{(2)})$ in (3) is zero for any given X, because in the expression (1) of $f(X|Z^{(1)})$ and $f(X|Z^{(2)})$, $1_{\mathbb{L}^{(1)*}}(l)$ or $1_{\mathbb{L}^{(2)*}}(l)$ will be zero for all $l\in L_X$. Let us, for example, adopt the viewpoint of sensor 1 and represent the labels of X, $l\in L_X$, such that $L_X\subseteq\mathbb{L}^{(1)*}$. This implies that $1_{\mathbb{L}^{(1)*}}(l)=1$ for all $l\in L_X$ and thus $f(X|Z^{(1)})$ is not zero. However, $1_{\mathbb{L}^{(2)*}}(l)=0$ for all $l\in L_X$, so that $f(X|Z^{(2)})=0$ and thus we still obtain $\tilde{f}(X|Z^{(1)},Z^{(2)})=0$.

We can, in principle, resolve this issue as follows: when evaluating $f(X|Z^{(2)})$, we first map the labels $l\in L_X\subseteq\mathbb{L}^{(1)*}$ to some labels $l'\in\mathbb{L}^{(2)*}$ so that $1_{\mathbb{L}^{(2)*}}(l')=1$. We can describe such a label mapping by a *label association vector* $\boldsymbol{a}=\big[a_{l_1}\cdots a_{l_{|\mathbb{L}^{(1)*}|}}\big]^{\mathrm{T}}$ of length $|\mathbb{L}^{(1)*}|$ and with entries $a_l\in\mathbb{L}^{(2)*}$ for all $l\in\mathbb{L}^{(1)*}$. For any $l\in\mathbb{L}^{(1)*}$, $a_l\in\mathbb{L}^{(2)*}$ indicates that the object with state $(\boldsymbol{x},l)$ tracked at sensor 1 is associated with the object with state $(\boldsymbol{x},a_l)$ tracked at sensor 2. For a one-to-one mapping, we require that different labels at sensor 1 are associated with different labels at sensor 2, which means that all entries a_l of $\boldsymbol{a}$ have to be different; this will be referred to as an *admissible* association. Using $\boldsymbol{a}$, we now modify the GCI fusion rule (3) according to

$$\tilde{f}_{\boldsymbol{a}}(X|Z^{(1)},Z^{(2)}) \triangleq \frac{1}{D_{\boldsymbol{a}}}\big(f(X|Z^{(1)})\big)^{\omega}\big(f(X_{\boldsymbol{a}}|Z^{(2)})\big)^{1-\omega}, \quad (4)$$

with $D_{\boldsymbol{a}}\triangleq\int\big(f(X|Z^{(1)})\big)^{\omega}\big(f(X_{\boldsymbol{a}}|Z^{(2)})\big)^{1-\omega}\delta X$. Here, $X_{\boldsymbol{a}}\triangleq\big\{(\boldsymbol{x}_1,a_{l_1}),\ldots,(\boldsymbol{x}_n,a_{l_n})\big\}$ for any realization $X=\big\{(\boldsymbol{x}_1,l_1),\ldots,(\boldsymbol{x}_n,l_n)\big\}$. That is, the label mapping defined by $\boldsymbol{a}$ changes each label $l_j\in L_X\subseteq\mathbb{L}^{(1)*}$ into a label $a_{l_j}\in\mathbb{L}^{(2)*}$, for $j=1,\ldots,n$. Note that now, in the expression (1) of $f(X|Z^{(1)})$ and $f(X_{\boldsymbol{a}}|Z^{(2)})$, we have, respectively, $1_{\mathbb{L}^{(1)*}}(l)=1$ and $1_{\mathbb{L}^{(2)*}}(a_l)=1$ for all $l\in L_X$, and thus $\tilde{f}_{\boldsymbol{a}}(X|Z^{(1)},Z^{(2)})$ is no longer zero in general.

More specifically, using (1) and the fact that $1_{\mathbb{L}^{(1)*}}(l)=1_{\mathbb{L}^{(2)*}}(a_l)=1$ for all $l\in L_X$, the fused pdf in (4) becomes

$$\begin{aligned}&\tilde{f}_{\boldsymbol{a}}(X|Z^{(1)},Z^{(2)})\\&\propto\Delta(X)\big(w^{(1)}(L_X)\big)^{\omega}\left(\prod_{(\boldsymbol{x},l)\in X}\big(f_l^{(1)}(\boldsymbol{x})\big)^{\omega}\right)\\&\quad\times\Delta(X_{\boldsymbol{a}})\big(w^{(2)}(L_{X_{\boldsymbol{a}}})\big)^{1-\omega}\prod_{(\boldsymbol{x}',l')\in X_{\boldsymbol{a}}}\big(f_{l'}^{(2)}(\boldsymbol{x}')\big)^{1-\omega}, \quad (5)\end{aligned}$$

with (cf. (2))

$$w^{(1)}(L_X)\triangleq\left(\prod_{l\in L_X}r_l^{(1)}\right)\prod_{l'\in\mathbb{L}^{(1)*}\setminus L_X}\big(1-r_{l'}^{(1)}\big), \quad (6)$$

$$w^{(2)}(L_{X_{\boldsymbol{a}}})\triangleq\left(\prod_{l\in L_{X_{\boldsymbol{a}}}}r_l^{(2)}\right)\prod_{l'\in\mathbb{L}^{(2)*}\setminus L_{X_{\boldsymbol{a}}}}\big(1-r_{l'}^{(2)}\big). \quad (7)$$

Moreover, because $\boldsymbol{a}$ is an admissible label association vector, $\Delta(X_{\boldsymbol{a}})$ is equal to $\Delta(X)$. Hence, we can rewrite (5) as

$$\begin{aligned}\tilde{f}_{\boldsymbol{a}}(X|Z^{(1)},Z^{(2)})&\propto\Delta(X)\big(w^{(1)}(L_X)\big)^{\omega}\big(w^{(2)}(L_{X_{\boldsymbol{a}}})\big)^{1-\omega}\\&\quad\times\prod_{(\boldsymbol{x},l)\in X}\big(f_l^{(1)}(\boldsymbol{x})\big)^{\omega}\big(f_{a_l}^{(2)}(\boldsymbol{x})\big)^{1-\omega},\end{aligned}$$

or, more compactly,

$$\tilde{f}_{\boldsymbol{a}}(X|Z^{(1)},Z^{(2)})\propto\Delta(X)w_{\boldsymbol{a}}(L_X)\prod_{(\boldsymbol{x},l)\in X}f_{l,a_l}(\boldsymbol{x}), \quad (8)$$

with the spatial pdfs

$$f_{l,a_l}(\boldsymbol{x})\triangleq\frac{1}{D_{l,a_l}}\big(f_l^{(1)}(\boldsymbol{x})\big)^{\omega}\big(f_{a_l}^{(2)}(\boldsymbol{x})\big)^{1-\omega} \quad (9)$$

and the weights

$$w_{\boldsymbol{a}}(L_X)\triangleq\big(w^{(1)}(L_X)\big)^{\omega}\big(w^{(2)}(L_{X_{\boldsymbol{a}}})\big)^{1-\omega}\prod_{l\in L_X}D_{l,a_l}, \quad (10)$$

where $D_{l,a_l}\triangleq\int_{\mathbb{R}^{N_x}}\big(f_l^{(1)}(\boldsymbol{x})\big)^{\omega}\big(f_{a_l}^{(2)}(\boldsymbol{x})\big)^{1-\omega}\mathrm{d}\boldsymbol{x}$. From (8), we conclude that $\tilde{f}_{\boldsymbol{a}}(X|Z^{(1)},Z^{(2)})$ is the pdf of a generalized LMB (GLMB) pdf [2].

IV. Probabilistic Label Association

For an LMB fusion leading to good performance of the resulting distributed LMB filter, the label mapping introduced above should be such that each mapped label of sensor 1 equals the specific label of sensor 2 that describes the *same* object. Because the label association vector $\boldsymbol{a}$ defining this "correct" mapping is unknown, we hereafter model $\boldsymbol{a}$ as a random vector $\mathbf{a}$ and perform an implicit (soft, probabilistic) estimation of $\mathbf{a}$. We refer to this approach as "probabilistic label association," by analogy to the probabilistic data association performed in many multiobject tracking methods to implicitly associate measurements with objects [1].

A. Label Association Distribution

Let us define the extended label association vector $\bar{\boldsymbol{a}}=\big[\bar{a}_{l_1}\cdots\bar{a}_{l_{|\mathbb{L}^{(1)*}|}}\big]^{\mathrm{T}}$ of length $|\mathbb{L}^{(1)*}|$ (equal to the length of $\boldsymbol{a}$) with entries $\bar{a}_l\in\mathbb{L}^{(2)*}\cup\{0\}$ for $l\in\mathbb{L}^{(1)*}$. Here, $\bar{a}_l\in\mathbb{L}^{(2)*}$ indicates—as before—that the object with state $(\boldsymbol{x},l)$ tracked at sensor 1 is associated with the object with state $(\boldsymbol{x},\bar{a}_l)$ tracked at sensor 2, and $\bar{a}_l=0$ indicates that an object with label l does not exist (i.e., $(\boldsymbol{x},l)\notin X$) according to the tracking performed at sensor 1. Note that the latter case also implies that no

object with state $(\boldsymbol{x}, l)$ is associated with any object tracked at sensor 2. We denote by $\bar{\mathcal{A}}$ the set of all admissible $\bar{\boldsymbol{a}}$, i.e., of all $\bar{\boldsymbol{a}}$ whose nonzero entries $\bar{a}_l$ are different. Furthermore, for any $L \subseteq \mathbb{L}^{(1)*}$, we define $\varphi(\bar{\boldsymbol{a}}, L)$ to be 1 for all admissible $\bar{\boldsymbol{a}}$ such that $\bar{a}_l \in \mathbb{L}^{(2)*}$ for $l \in L$ and $\bar{a}_l = 0$ for $l \in \mathbb{L}^{(1)*} \backslash L$, and to be 0 otherwise. Then, using the definitions of $w_{\boldsymbol{a}}(L_X)$ in (10) and, in turn, of $w^{(1)}(L_X)$ in (6) and $w^{(2)}(L_{X_{\boldsymbol{a}}})$ in (7), it can be shown that Eq. (8) can be rewritten in terms of $\bar{\boldsymbol{a}}$ as

$$\tilde{f}_{\bar{\boldsymbol{a}}}(X|Z^{(1)}, Z^{(2)}) \propto \Delta(X)\varphi(\bar{\boldsymbol{a}}, L_X) w_{\bar{\boldsymbol{a}}} \prod_{(\boldsymbol{x},l)\in X} f_{l,\bar{a}_l}(\boldsymbol{x}). \quad (11)$$

Here, the weights $w_{\bar{\boldsymbol{a}}}$ are given by

$$w_{\bar{\boldsymbol{a}}} = \prod_{l \in \mathbb{L}^{(1)*}} \beta_{l,\bar{a}_l}, \quad \bar{\boldsymbol{a}} \in \bar{\mathcal{A}}, \quad (12)$$

with the "label association weights"

$$\beta_{l,\bar{a}_l} \triangleq \begin{cases} \dfrac{\left(r_l^{(1)}\right)^{\omega}\left(r_{\bar{a}_l}^{(2)}\right)^{1-\omega} D_{l,\bar{a}_l}}{\left(1-r_{\bar{a}_l}^{(2)}\right)^{1-\omega}}, & \bar{a}_l \in \mathbb{L}^{(2)*}, \\ \left(1-r_l^{(1)}\right)^{\omega}, & \bar{a}_l = 0, \end{cases} \quad (13)$$

where $D_{l,\bar{a}_l} \triangleq \int_{\mathbb{R}^{N_x}} \left(f_l^{(1)}(\boldsymbol{x})\right)^{\omega} \left(f_{\bar{a}_l}^{(2)}(\boldsymbol{x})\right)^{1-\omega} \mathrm{d}\boldsymbol{x}$. Furthermore, the spatial pdfs $f_{l,\bar{a}_l}(\boldsymbol{x})$ are given by (9) with a_l replaced by $\bar{a}_l$. (Note that $f_{l,0}(\boldsymbol{x})$ does not occur in (11) because $\bar{a}_l = 0$ implies $(\boldsymbol{x}, l) \notin X$.) Differently from the weights $w_{\boldsymbol{a}}(L_X)$ in our earlier expression (8), the weights $w_{\bar{\boldsymbol{a}}}$ do not depend on L_X; however, the factor $\varphi(\bar{\boldsymbol{a}}, L_X)$ ensures that (11) is still equivalent to (8).

Just as $\boldsymbol{a}$, we hereafter consider $\bar{\boldsymbol{a}}$ as random (denoted $\bar{\mathbf{a}}$). Our probabilistic label association method is based on the idea of interpreting the weights $w_{\bar{\boldsymbol{a}}}$ in (11) and (12) as the probability distribution (pmf) of $\bar{\mathbf{a}}$. More precisely, we define the pmf of $\bar{\mathbf{a}}$ as

$$p(\bar{\boldsymbol{a}}) = \begin{cases} \alpha\, w_{\bar{\boldsymbol{a}}}, & \bar{\boldsymbol{a}} \in \bar{\mathcal{A}}, \\ 0, & \text{otherwise}, \end{cases} \quad (14)$$

for all $\bar{\boldsymbol{a}} \in (\mathbb{L}^{(2)*} \cup \{0\})^{|\mathbb{L}^{(1)*}|}$, where $\alpha = 1/\sum_{\bar{\boldsymbol{a}} \in \bar{\mathcal{A}}} w_{\bar{\boldsymbol{a}}}$. We can then rewrite (11) as

$$\tilde{f}(X, \bar{\boldsymbol{a}}|Z^{(1)}, Z^{(2)}) = \Delta(X)\varphi(\bar{\boldsymbol{a}}, L_X) p(\bar{\boldsymbol{a}}) \prod_{(\boldsymbol{x},l)\in X} f_{l,\bar{a}_l}(\boldsymbol{x}). \quad (15)$$

In fact, it can be verified that integrating/summing the right-hand side of (15) with respect to X and $\bar{\boldsymbol{a}}$ yields 1. Accordingly, expression (15) defines the joint pdf/pmf of the random variables X and $\bar{\mathbf{a}}$, whereas $\tilde{f}_{\bar{\boldsymbol{a}}}(X|Z^{(1)}, Z^{(2)})$ in (11) is the pdf of the random variable X parametrized by the nonrandom variable $\bar{\boldsymbol{a}}$. We can furthermore write the joint pdf/pmf of X and $\bar{\mathbf{a}}$ as $\tilde{f}(X, \bar{\boldsymbol{a}}|Z^{(1)}, Z^{(2)}) = \tilde{f}(X|\bar{\boldsymbol{a}}, Z^{(1)}, Z^{(2)}) p(\bar{\boldsymbol{a}})$, with the conditional fused posterior pdf

$$\tilde{f}(X|\bar{\boldsymbol{a}}, Z^{(1)}, Z^{(2)}) = \Delta(X)\varphi(\bar{\boldsymbol{a}}, L_X) \prod_{(\boldsymbol{x},l)\in X} f_{l,\bar{a}_l}(\boldsymbol{x}). \quad (16)$$

Following this interpretation, we can obtain the unconditional fused posterior pdf, to be denoted $\tilde{f}(X|Z^{(1)}, Z^{(2)})$, as

$$\tilde{f}(X|Z^{(1)}, Z^{(2)}) = \sum_{\bar{\boldsymbol{a}} \in \bar{\mathcal{A}}} \tilde{f}(X|\bar{\boldsymbol{a}}, Z^{(1)}, Z^{(2)}) p(\bar{\boldsymbol{a}}).$$

Here, we can extend the summation set $\bar{\mathcal{A}}$ to $(\mathbb{L}^{(2)*} \cup \{0\})^{|\mathbb{L}^{(1)*}|}$, because by (14) there is $p(\bar{\boldsymbol{a}}) = 0$ for $\bar{\boldsymbol{a}} \in (\mathbb{L}^{(2)*} \cup \{0\})^{|\mathbb{L}^{(1)*}|} \backslash \bar{\mathcal{A}}$. We thus obtain, using (16),

$$\tilde{f}(X|Z^{(1)}, Z^{(2)}) = \Delta(X) \sum_{\bar{\boldsymbol{a}} \in (\mathbb{L}^{(2)*} \cup \{0\})^{|\mathbb{L}^{(1)*}|}} \varphi(\bar{\boldsymbol{a}}, L_X) p(\bar{\boldsymbol{a}}) \times \prod_{(\boldsymbol{x},l)\in X} f_{l,\bar{a}_l}(\boldsymbol{x}). \quad (17)$$

B. LMB Approximation

The pdf $\tilde{f}(X|Z^{(1)}, Z^{(2)})$ in (17) is no longer of the LMB form (1); instead, it is the pdf of a GLMB RFS [2]. Therefore, we approximate it by an LMB pdf. Following [8], this is achieved by approximating the label association pmf $p(\bar{\boldsymbol{a}})$ by the product of its marginals, i.e.,

$$p(\bar{\boldsymbol{a}}) \approx \prod_{l \in \mathbb{L}^{(1)*}} p(\bar{a}_l), \quad \bar{\boldsymbol{a}} \in (\mathbb{L}^{(2)*} \cup \{0\})^{|\mathbb{L}^{(1)*}|}. \quad (18)$$

Here, $p(\bar{a}_l) = \sum_{\bar{\boldsymbol{a}}^{\sim l} \in (\mathbb{L}^{(2)*} \cup \{0\})^{|\mathbb{L}^{(1)*}|-1}} p(\bar{\boldsymbol{a}})$, where $\bar{\boldsymbol{a}}^{\sim l}$ denotes the vector $\bar{\boldsymbol{a}}$ with the lth entry, $\bar{a}_l$, removed. Inserting (18) into (17) yields an approximation $\tilde{\tilde{f}}(X|Z^{(1)}, Z^{(2)}) \approx \tilde{f}(X|Z^{(1)}, Z^{(2)})$ that is given by

$$\tilde{\tilde{f}}(X|Z^{(1)}, Z^{(2)}) = \Delta(X) \sum_{\bar{\boldsymbol{a}} \in (\mathbb{L}^{(2)*} \cup \{0\})^{|\mathbb{L}^{(1)*}|}} \varphi(\bar{\boldsymbol{a}}, L_X) \times \left(\prod_{l' \in \mathbb{L}^{(1)*}} p(\bar{a}_{l'}) \right) \prod_{(\boldsymbol{x},l)\in X} f_{l,\bar{a}_l}(\boldsymbol{x}).$$

Next, splitting $\prod_{l' \in \mathbb{L}^{(1)*}} p(\bar{a}_{l'})$ as $\left(\prod_{l' \in \mathbb{L}^{(1)*} \backslash L_X} p(\bar{a}_{l'})\right) \times \prod_{l \in L_X} p(\bar{a}_l)$ and evaluating $\varphi(\bar{\boldsymbol{a}}, L_X)$ yields

$$\tilde{\tilde{f}}(X|Z^{(1)}, Z^{(2)}) = \Delta(X) \left(\prod_{l' \in \mathbb{L}^{(1)*} \backslash L_X} p(\bar{a}_{l'} = 0) \right) \times \sum_{\bar{\boldsymbol{a}} \in (\mathbb{L}^{(2)*})^{|\mathbb{L}^{(1)*}|}} \prod_{(\boldsymbol{x},l)\in X} p(\bar{a}_l) f_{l,\bar{a}_l}(\boldsymbol{x}).$$

Using the fact that $\sum_{\bar{\boldsymbol{a}} \in (\mathbb{L}^{(2)*})^{|\mathbb{L}^{(1)*}|}} = \sum_{\bar{a}_1 \in \mathbb{L}^{(2)*}} \cdots \sum_{\bar{a}_{|\mathbb{L}^{(1)*}|} \in \mathbb{L}^{(2)*}}$, we finally obtain

$$\tilde{\tilde{f}}(X|Z^{(1)}, Z^{(2)}) = \Delta(X) \left(\prod_{l' \in \mathbb{L}^{(1)*} \backslash L_X} p(\bar{a}_{l'} = 0) \right) \times \prod_{(\boldsymbol{x},l)\in X} \sum_{\bar{a}_l \in \mathbb{L}^{(2)*}} p(\bar{a}_l) f_{l,\bar{a}_l}(\boldsymbol{x}). \quad (19)$$

Comparing this expression with[1] (1), one can easily verify that $\tilde{\tilde{f}}(X|Z^{(1)}, Z^{(2)})$ is an LMB pdf parametrized by

[1]Note that, differently from (1), expression (19) does not include the factor $1_{\mathbb{L}^{(1)*}}(l)$. This is because in (19), $L_X \subseteq \mathbb{L}^{(1)*}$ and thus always $1_{\mathbb{L}^{(1)*}}(l) = 1$, whereas in (1), $L_X \subseteq \mathbb{L}$.

$\big\{\big(\rho_l, f_l(\boldsymbol{x})\big)\big\}_{l\in\mathbb{L}^{(1)*}}$, with existence probabilities

$$\rho_l = \sum_{\bar{a}_l\in\mathbb{L}^{(2)*}} p(\bar{a}_l) \tag{20}$$

and spatial pdfs

$$f_l(\boldsymbol{x}) = \frac{1}{\rho_l}\sum_{\bar{a}_l\in\mathbb{L}^{(2)*}} p(\bar{a}_l)\, f_{l,\bar{a}_l}(\boldsymbol{x}), \tag{21}$$

where $l\in\mathbb{L}^{(1)*}$. Eqs. (20) and (21) complete the formulation of our pairwise LMB fusion scheme using soft label association. The input to this scheme are the Bernoulli parameter sets $\big\{\big(r_l^{(1)}, f_l^{(1)}(\boldsymbol{x})\big)\big\}_{l\in\mathbb{L}^{(1)*}}$ and $\big\{\big(r_l^{(2)}, f_l^{(2)}(\boldsymbol{x})\big)\big\}_{l\in\mathbb{L}^{(2)*}}$, and the output is the fused Bernoulli parameter set $\big\{\big(\rho_l, f_l(\boldsymbol{x})\big)\big\}_{l\in\mathbb{L}^{(1)*}}$, with ρ_l given by (20) and $f_l(\boldsymbol{x})$ given by (21). Here, $p(\bar{a}_l)$ in (20), (21) is calculated by marginalizing $p(\bar{\boldsymbol{a}})$ in (14), and $f_{l,\bar{a}_l}(\boldsymbol{x})$ in (21) is given by (9).

In a practical implementation, Bernoulli components $l\in\mathbb{L}^{(1)*}$ that do not have a plausible association with any Bernoulli component $l'\in\mathbb{L}^{(2)*}$ at sensor 2 can be excluded from the fusion procedure. The plausibility of an association can be measured by $D_{l,l'} = \int_{\mathbb{R}^{N_x}} \big(f_l^{(1)}(\boldsymbol{x})\big)^{\omega}\big(f_{l'}^{(2)}(\boldsymbol{x})\big)^{1-\omega}\mathrm{d}\boldsymbol{x}$. Thus, we do not fuse Bernoulli components $l\in\mathbb{L}^{(1)*}$ such that $D_{l,l'} < \gamma_{\mathrm{F}}$ for all $l'\in\mathbb{L}^{(2)*}$, with a positive threshold γ_{F}.

V. Fast BP-Based Algorithm

Next, we present a fast algorithm for approximate marginalization of $p(\bar{\boldsymbol{a}})$. This algorithm is inspired by the belief propagation (BP) based algorithm for probabilistic data association proposed in [14].

Inserting (12) into (14), we can express $p(\bar{\boldsymbol{a}})$ as

$$p(\bar{\boldsymbol{a}}) \propto \Psi(\bar{\boldsymbol{a}}) \prod_{l\in\mathbb{L}^{(1)*}} \beta_{l,\bar{a}_l}, \quad \bar{\boldsymbol{a}}\in(\mathbb{L}^{(2)*}\cup\{0\})^{|\mathbb{L}^{(1)*}|}, \tag{22}$$

where $\Psi(\bar{\boldsymbol{a}})=1$ if $\bar{\boldsymbol{a}}\in\bar{\mathcal{A}}$ and $\Psi(\bar{\boldsymbol{a}})=0$ otherwise. By analogy to [14], we introduce the alternative label association vector $\boldsymbol{b} = \big[b_{l_1}\cdots b_{l_{|\mathbb{L}^{(2)*}|}}\big]^{\mathrm{T}}$ of length $|\mathbb{L}^{(2)*}|$ and with entries $b_l\in\mathbb{L}^{(1)*}\cup\{0\}$ for $l\in\mathbb{L}^{(2)*}$. Here, $b_l\in\mathbb{L}^{(1)*}$ indicates that the object with state $(\boldsymbol{x},l)$ tracked at sensor 2 is associated with the object with state $(\boldsymbol{x},b_l)$ tracked at sensor 1, and $b_l=0$ indicates that an object with label l does not exist according to the tracking performed at sensor 2. The latter case also implies that no object with state $(\boldsymbol{x},l)$ is associated with any object tracked at sensor 1. Thus, $\boldsymbol{b}$ is a description of the label associations that is analogous to $\bar{\boldsymbol{a}}$ but "viewed from the other sensor." We can reformulate $p(\bar{\boldsymbol{a}})$ in (22) in terms of $\bar{\boldsymbol{a}}$ and $\boldsymbol{b}$ as

$$p(\bar{\boldsymbol{a}},\boldsymbol{b}) \propto \Psi(\bar{\boldsymbol{a}},\boldsymbol{b}) \prod_{l\in\mathbb{L}^{(1)*}} \beta_{l,\bar{a}_l}, \tag{23}$$

for $\bar{\boldsymbol{a}}\in(\mathbb{L}^{(2)*}\cup\{0\})^{|\mathbb{L}^{(1)*}|}$ and $\boldsymbol{b}\in(\mathbb{L}^{(1)*}\cup\{0\})^{|\mathbb{L}^{(2)*}|}$, with the admissibility constraint factor $\Psi(\bar{\boldsymbol{a}},\boldsymbol{b})\in\{0,1\}$ given by

$$\Psi(\bar{\boldsymbol{a}},\boldsymbol{b}) = \prod_{l\in\mathbb{L}^{(1)*}}\prod_{l'\in\mathbb{L}^{(2)*}} \psi_{l,l'}(\bar{a}_l, b_{l'}). \tag{24}$$

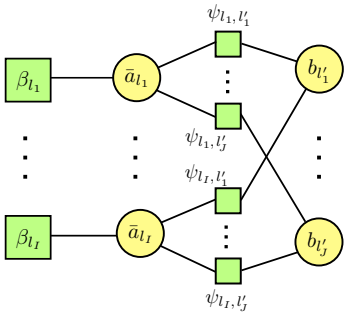

Fig. 1: Factor graph representing the factorization of $p(\bar{\boldsymbol{a}},\boldsymbol{b})$ in (23), (24). Variable nodes are depicted as circles and factor nodes as squares. The shorthands $\beta_l \triangleq \beta_{l,\bar{a}_l}$, $\psi_{l,l'} \triangleq \psi_{l,l'}(\bar{a}_l, b_{l'})$, $I=|\mathbb{L}^{(1)*}|$, and $J\triangleq|\mathbb{L}^{(2)*}|$ are used.

Here, $\psi_{l,l'}(\bar{a}_l, b_{l'}) = 0$ if either $\bar{a}_l = l'$ and $b_{l'}\neq l$ or $\bar{a}_l\neq l'$ and $b_{l'}=l$, and $\psi_{l,l'}(\bar{a}_l, b_{l'})=1$ otherwise.

The above reformulation of $p(\bar{\boldsymbol{a}})$ in terms of $\bar{\boldsymbol{a}}$ and $\boldsymbol{b}$ allows us to devise an efficient algorithm for calculating accurate approximations to the marginal association pmfs $p(\bar{a}_l)$. First, the factorization (23), (24) is represented by the corresponding factor graph [16], which is shown in Fig. 1. Then, we execute the BP algorithm, aka sum-product algorithm [16], on this factor graph. In BP iteration $p\in\{1,\ldots,P\}$, a message $\zeta_{l\to l'}^{[p]}$ is passed from variable node "$\bar{a}_l$" via factor node "$\psi_{l,l'}(\bar{a}_l, b_{l'})$" to variable node "$b_{l'}$", and a message $\nu_{l'\to l}^{[p]}$ is passed from variable node "$b_{l'}$" via factor node "$\psi_{l,l'}(\bar{a}_l, b_{l'})$" to variable node "$\bar{a}_l$". Similarly to [14], one obtains the following expressions of these messages:

$$\zeta_{l\to l'}^{[p]} = \frac{\beta_{l,l'}}{\beta_{l,0} + \sum_{\lambda\in\mathbb{L}^{(2)*}\setminus\{l'\}} \beta_{l,\lambda}\,\nu_{\lambda\to l}^{[p-1]}},$$

$$\nu_{l'\to l}^{[p]} = \frac{1}{1 + \sum_{\lambda\in\mathbb{L}^{(1)*}\setminus\{l\}} \zeta_{\lambda\to l'}^{[p]}},$$

for all $l\in\mathbb{L}^{(1)*}$ and $l'\in\mathbb{L}^{(2)*}$. The recursion established by these two equations is initialized by $\nu_{l'\to l}^{[0]}=1$. After the final iteration $p=P$, approximations to the marginal association pmfs $p(\bar{a}_l)$, $l\in\mathbb{L}^{(1)*}$ are provided by the beliefs $\tilde{p}(\bar{a}_l)$ at the respective variable nodes "$\bar{a}_l$" in Fig. 1. These are obtained as

$$\tilde{p}(\bar{a}_l{=}l') = \begin{cases} \beta_{l,l'}\,\nu_{l'\to l}^{[P]}/C_l, & l'\in\mathbb{L}^{(2)*}, \\ \beta_{l,l'}/C_l, & l'=0, \end{cases} \tag{25}$$

where $C_l \triangleq \beta_{l,0} + \sum_{\lambda\in\mathbb{L}^{(2)*}} \beta_{l,\lambda}\,\nu_{\lambda\to l}^{[P]}$.

VI. Gaussian Implementation

For a reduction of communication and computation requirements, we next assume Gaussian models for the spatial pdfs of the local LMB RFSs, i.e.,

$$f_l^{(1)}(\boldsymbol{x}) = \mathcal{N}\big(\boldsymbol{x};\boldsymbol{\mu}_l^{(1)},\boldsymbol{\Sigma}_l^{(1)}\big), \quad l\in\mathbb{L}^{(1)*},$$

$$f_l^{(2)}(\boldsymbol{x}) = \mathcal{N}\big(\boldsymbol{x};\boldsymbol{\mu}_l^{(2)},\boldsymbol{\Sigma}_l^{(2)}\big), \quad l\in\mathbb{L}^{(2)*}.$$

Accordingly, the local LMB parameter sets are now given by $\big\{\big(r_l^{(s)}, \boldsymbol{\mu}_l^{(s)}, \boldsymbol{\Sigma}_l^{(s)}\big)\big\}_{l\in\mathbb{L}^{(s)*}}$ for $s\in\{1,2\}$. In what follows, we will present the resulting expressions of $f_{l,\bar{a}_l}(\boldsymbol{x})$, $\beta_{l,\bar{a}_l}$, and $f_l(\boldsymbol{x})$, where $l\in\mathbb{L}^{(1)*}$.

We recall that the spatial pdf $f_{l,\bar{a}_l}(\boldsymbol{x})$ is calculated from $f_l^{(1)}(\boldsymbol{x})$ and $f_{\bar{a}_l}^{(2)}(\boldsymbol{x})$ via (9). As was shown in [17], if $f_l^{(1)}(\boldsymbol{x})$ and $f_{\bar{a}_l}^{(2)}(\boldsymbol{x})$ are Gaussian as assumed above, then $f_{l,\bar{a}_l}(\boldsymbol{x})$ is also Gaussian, i.e.,

$$f_{l,\bar{a}_l}(\boldsymbol{x}) = \mathcal{N}\big(\boldsymbol{x}; \boldsymbol{\mu}_{l,\bar{a}_l}, \boldsymbol{\Sigma}_{l,\bar{a}_l}\big), \tag{26}$$

with mean vector and covariance matrix given by

$$\boldsymbol{\mu}_{l,\bar{a}_l} = \boldsymbol{\Sigma}_{l,\bar{a}_l}\big(\omega\,\boldsymbol{\Sigma}_l^{(1)-1}\boldsymbol{\mu}_l^{(1)} + (1-\omega)\,\boldsymbol{\Sigma}_{\bar{a}_l}^{(2)-1}\boldsymbol{\mu}_{\bar{a}_l}^{(2)}\big), \tag{27}$$

$$\boldsymbol{\Sigma}_{l,\bar{a}_l} = \big(\omega\,\boldsymbol{\Sigma}_l^{(1)-1} + (1-\omega)\,\boldsymbol{\Sigma}_{\bar{a}_l}^{(2)-1}\big)^{-1}. \tag{28}$$

It can furthermore be shown that the normalization factor $D_{l,\bar{a}_l} = \int_{\mathbb{R}^{N_x}} \big(f_l^{(1)}(\boldsymbol{x})\big)^{\omega}\big(f_{\bar{a}_l}^{(2)}(\boldsymbol{x})\big)^{1-\omega}\mathrm{d}\boldsymbol{x}$ in (13) is given by

$$D_{l,\bar{a}_l} = \gamma_l\,\kappa_{\bar{a}_l}\,\mathcal{N}\Big(\boldsymbol{\mu}_l^{(1)}; \boldsymbol{\mu}_{\bar{a}_l}^{(2)}, \frac{1}{\omega}\,\boldsymbol{\Sigma}_l^{(1)} + \frac{1}{1-\omega}\,\boldsymbol{\Sigma}_{\bar{a}_l}^{(2)}\Big). \tag{29}$$

Here, $\gamma_l = \sqrt{\det(\frac{2\pi}{\omega}\,\boldsymbol{\Sigma}_l^{(1)})/\big(\det(2\pi\boldsymbol{\Sigma}_l^{(1)})\big)^{\omega}}$ and $\kappa_{\bar{a}_l} = \sqrt{\det(\frac{2\pi}{1-\omega}\,\boldsymbol{\Sigma}_{\bar{a}_l}^{(2)})/\big(\det(2\pi\boldsymbol{\Sigma}_{\bar{a}_l}^{(2)})\big)^{1-\omega}}$, and $\mathcal{N}\big(\boldsymbol{\mu}_l^{(1)}; \boldsymbol{\mu}_{\bar{a}_l}^{(2)}, \frac{1}{\omega}\,\boldsymbol{\Sigma}_l^{(1)} + \frac{1}{1-\omega}\,\boldsymbol{\Sigma}_{\bar{a}_l}^{(2)}\big)$ denotes the positive number obtained by evaluating the Gaussian pdf $\mathcal{N}\big(\boldsymbol{x}; \boldsymbol{\mu}_{\bar{a}_l}^{(2)}, \frac{1}{\omega}\,\boldsymbol{\Sigma}_l^{(1)} + \frac{1}{1-\omega}\,\boldsymbol{\Sigma}_{\bar{a}_l}^{(2)}\big)$ at $\boldsymbol{x} = \boldsymbol{\mu}_l^{(1)}$. An expression of the label association weight $\beta_{l,\bar{a}_l}$ is then obtained by inserting (29) into (13).

Finally, according to (21) and (26), the fused spatial pdf $f_l(\boldsymbol{x})$ is a Gaussian mixture pdf. We next approximate it by the Gaussian pdf

$$\bar{f}_l(\boldsymbol{x}) = \mathcal{N}\big(\boldsymbol{x}; \bar{\boldsymbol{\mu}}_l, \bar{\boldsymbol{\Sigma}}_l\big), \quad l\in\mathbb{L}^{(1)*},$$

whose mean $\bar{\boldsymbol{\mu}}_l$ and covariance matrix $\bar{\boldsymbol{\Sigma}}_l$ are taken to be equal to those of $f_l(\boldsymbol{x})$. Using (21) and (26), one obtains [18]

$$\bar{\boldsymbol{\mu}}_l = \frac{1}{\rho_l}\sum_{\bar{a}_l\in\mathbb{L}^{(2)*}} p(\bar{a}_l)\,\boldsymbol{\mu}_{l,\bar{a}_l}, \tag{30}$$

$$\bar{\boldsymbol{\Sigma}}_l = \frac{1}{\rho_l}\sum_{\bar{a}_l\in\mathbb{L}^{(2)*}} p(\bar{a}_l)\big(\boldsymbol{\Sigma}_{l,\bar{a}_l} + (\boldsymbol{\mu}_{l,\bar{a}_l} - \bar{\boldsymbol{\mu}}_l)(\boldsymbol{\mu}_{l,\bar{a}_l} - \bar{\boldsymbol{\mu}}_l)^{\mathrm{T}}\big), \tag{31}$$

for $l\in\mathbb{L}^{(1)*}$. Here, ρ_l is given by (20), $p(\bar{a}_l)$ is the marginal label association probability (cf. (18)), and $\boldsymbol{\mu}_{l,\bar{a}_l}$ and $\boldsymbol{\Sigma}_{l,\bar{a}_l}$ are given by (27) and (28), respectively. Within this Gaussian approximation, the fused LMB parameter set is $\big\{\big(\rho_l, \bar{\boldsymbol{\mu}}_l, \bar{\boldsymbol{\Sigma}}_l\big)\big\}_{l\in\mathbb{L}^{(1)*}}$. We note that in the final algorithm (see Section VII), $p(\bar{a}_l)$ is approximated by the belief $\tilde{p}(\bar{a}_l)$ in (25).

VII. Distributed Networkwide Fusion

In the formulation of our pairwise fusion method, we used sensor $s=1$ as a "reference sensor" for fusing the posterior pdfs of sensors $s=1$ and $s=2$. As a consequence, the fused quantities ρ_l, $p(\bar{a}_l)$, $\bar{\boldsymbol{\mu}}_l$, $\bar{\boldsymbol{\Sigma}}_l$, and $\bar{f}_l(\boldsymbol{x})$ are defined for $l\in\mathbb{L}^{(1)*}$, i.e., the underlying label set is $\mathbb{L}^{(1)*}$. In a distributed implementation, each sensor $s\in\{1,2\}$ runs its own instance of the pairwise fusion method, using its own label set $\mathbb{L}^{(s)*}$ as the reference label set (while, of course, still taking into account the local LMB parameters of the respective other sensor). This implies that the fused LMB parameter sets calculated at the two sensors will be different. Let $\big\{\big(\rho_l^{(s)}, \bar{\boldsymbol{\mu}}_l^{(s)}, \bar{\boldsymbol{\Sigma}}_l^{(s)}\big)\big\}_{l\in\mathbb{L}^{(s)*}}$ denote the fused LMB parameter set calculated at sensor $s\in\{1,2\}$; this should not be confused with the original local LMB parameter set $\big\{\big(r_l^{(s)}, \boldsymbol{\mu}_l^{(s)}, \boldsymbol{\Sigma}_l^{(s)}\big)\big\}_{l\in\mathbb{L}^{(s)*}}$.

The proposed fusion algorithm—still considering pairwise fusion—is now obtained by replacing in (20), (30), and (31) the marginal association pmfs $p(\bar{a}_l)$ by the beliefs $\tilde{p}(\bar{a}_l)$ given by (25). The fused LMB parameter set of sensor $s\in\{1,2\}$ using this additional approximation will be denoted by $\big\{\big(\tilde{\rho}_l^{(s)}, \tilde{\boldsymbol{\mu}}_l^{(s)}, \tilde{\boldsymbol{\Sigma}}_l^{(s)}\big)\big\}_{l\in\mathbb{L}^{(s)*}}$. Note that calculation of this fused LMB parameter set at sensor s presupposes that the original local LMB parameter set of the respective other sensor is available at sensor s. This means that each sensor s has to transmit its local LMB parameter set $\big\{\big(r_l^{(s)}, \boldsymbol{\mu}_l^{(s)}, \boldsymbol{\Sigma}_l^{(s)}\big)\big\}_{l\in\mathbb{L}^{(s)*}}$ to its fusion partner.

So far, we considered the pairwise fusion of the local LMB parameter sets of two sensors. This pairwise fusion can be used to achieve *networkwide* fusion in a connected network of $S\geq 2$ sensors $s\in\{1,\ldots,S\}$ via the following iterative procedure [17]. Let $N_s\subseteq\{1,\ldots,S\}\backslash\{s\}$ denote the set of "neighbors" of sensor s, i.e., the set of sensors with which sensor s is able to communicate. In the first iteration, each sensor s transmits its local LMB parameter set $\big\{\big(r_l^{(s)}, \boldsymbol{\mu}_l^{(s)}, \boldsymbol{\Sigma}_l^{(s)}\big)\big\}_{l\in\mathbb{L}^{(s)*}}$ to its neighbors $s'\in N_s$. Then, each sensor s performs pairwise fusion sequentially (recursively) with each of the LMB parameter sets it received from its neighbors, in an arbitrary order. That is, the local LMB parameter set is fused with that of an arbitrary sensor $s'\in N_s$, the LMB parameter set resulting from this pairwise fusion is fused with that of an arbitrary sensor $s''\in N_s\backslash\{s'\}$, etc. Let $\big\{\big(\tilde{\rho}_l^{(s,1)}, \tilde{\boldsymbol{\mu}}_l^{(s,1)}, \tilde{\boldsymbol{\Sigma}}_l^{(s,1)}\big)\big\}_{l\in\mathbb{L}^{(s)*}}$—the superscript "1" indicates the first iteration—denote the LMB parameter set resulting from this sequence of $|N_s|$ successive pairwise fusion steps. In the second iteration, the sequence of $|N_s|$ pairwise fusion steps is repeated but with $\big\{\big(\tilde{\rho}_l^{(s,1)}, \tilde{\boldsymbol{\mu}}_l^{(s,1)}, \tilde{\boldsymbol{\Sigma}}_l^{(s,1)}\big)\big\}_{l\in\mathbb{L}^{(s)*}}$ substituted for the original local LMB parameter set $\big\{\big(r_l^{(s)}, \boldsymbol{\mu}_l^{(s)}, \boldsymbol{\Sigma}_l^{(s)}\big)\big\}_{l\in\mathbb{L}^{(s)*}}$ and the fusion results of the other sensors substituted for their original local LMB parameter sets; this requires another round of parameter transmissions between neighboring sensors. An analogous sequence of $|N_s|$ pairwise fusions is performed at each sensor also in each subsequent iteration.

As an alternative to this recursive algorithm, a gossip algorithm [19] may be used. In each iteration of the gossip algorithm, the current LMB parameter sets of randomly chosen pairs of communicating sensors are fused. This is initialized by the original local LMB parameter sets.

VIII. Simulation Results

We consider two simulation scenarios, briefly referred to as SC1 and SC2, which are inspired by [12] and [8]. In

both scenarios, the region of interest (ROI) is $[-150, 150] \times [-150, 150]$. We simulated ten (SC1) and twenty (SC2) objects during 200 time steps. The object trajectories were randomly chosen in each simulation run. The objects appear at various times before time step 40 (SC1) or 90 (SC2) and at randomly chosen positions in the area $[-50, 50] \times [-50, 50]$, and they disappear at various times after time step 150. The object states consist of two-dimensional position and velocity, i.e., $\boldsymbol{x}_k = [x_{k,1}\ x_{k,2}\ \dot{x}_{k,1}\ \dot{x}_{k,2}]^{\mathrm{T}}$. They evolve according to the nearly constant velocity motion model [20, Sec. 6.3.2] with an independent and identically distributed (iid) Gaussian driving noise process of variance 10^{-3}. A realization of the object trajectories for SC1 is shown in Fig. 2(a).

There are two sensors in SC1 and five in SC2. The sensor positions are $\boldsymbol{p}^{(1)} = [-50\ 0]^{\mathrm{T}}$ and $\boldsymbol{p}^{(2)} = [50\ 0]^{\mathrm{T}}$ and, in SC2, additionally $\boldsymbol{p}^{(3)} = [0\ 0]^{\mathrm{T}}$, $\boldsymbol{p}^{(4)} = [0\ 50]^{\mathrm{T}}$, and $\boldsymbol{p}^{(5)} = [0\ -50]^{\mathrm{T}}$. The communication links between the sensors are shown in Fig. 2. The object-generated measurements conform to the nonlinear range-bearing measurement model $\boldsymbol{z}_k = \left[\rho(\boldsymbol{x}_k)\ \phi(\boldsymbol{x}_k)\right]^{\mathrm{T}} + \boldsymbol{v}_k$. Here, $\rho(\boldsymbol{x}_k) \triangleq \|\boldsymbol{x}'_k - \boldsymbol{p}^{(s)}\|$, where $\boldsymbol{x}'_k \triangleq [x_{k,1}\ x_{k,2}]^{\mathrm{T}}$ denotes the object position and $\boldsymbol{p}^{(s)} = [p_1^{(s)}\ p_2^{(s)}]^{\mathrm{T}}$ the position of sensor s, and $\phi(\boldsymbol{x}_k) \triangleq \tan^{-1}\left(\frac{x_{k,2} - p_2^{(s)}}{x_{k,1} - p_1^{(s)}}\right)$. Furthermore, $\boldsymbol{v}_k$ is iid Gaussian measurement noise with independent entries with standard deviations $\sigma_\rho = 2$ and $\sigma_\phi = 1°$. The detection probability is chosen as 0.7 on the entire ROI. The clutter measurements are distributed uniformly on the ROI (in polar coordinates), with mean number equal to 10 (SC1) or 50 (SC2).

We consider a distributed LMB filter, briefly referred to as S-DLMB, that employs the proposed fusion algorithm with probabilistic (soft) label association using the BP and Gaussian approximations described in Sections V and VI, respectively. For SC2 (five sensors), S-DLMB uses the iterative networkwise extension of the algorithm—described in Section VII—with five fusion iterations. We compare the performance of S-DLMB with that of the distributed LMB filter proposed in [13], which uses GCI fusion with a hard label association algorithm and will be referred to as H-DLMB.[2] Furthermore, we also consider a centralized multisensor LMB filter, referred to as CLMB, that is based on the iterated-corrector scheme [2], [7]. CLMB and the local LMB filters involved in S-DLMB and H-DLMB use a particle implementation of the fast LMB filter proposed in [8]. In all filters, each spatial pdf is represented by 500 particles, and Bernoulli components with an existence probability below 10^{-3} are pruned after each filter update step. In S-DLMB and H-DLMB, the particle representation of each spatial pdf is further approximated by a Gaussian pdf after each filter update step, with the Gaussian parameters given by the sample mean and sample covariance of the particles. Then, S-DLMB executes the fusion algorithm described in Section VI and H-DLMB the fusion algorithm described in [13]. The fusion parameter ω is 0.5 in both cases. After the fusion step, 500 particles are drawn from the fused Gaussian pdf, and the local LMB filters execute the next filtering step. The threshold γ_{F} used in S-DLMB as described in Section IV-B, and also used in a similar way in H-DLMB, is 10^{-20}. The remaining simulation parameters equal those in [8].

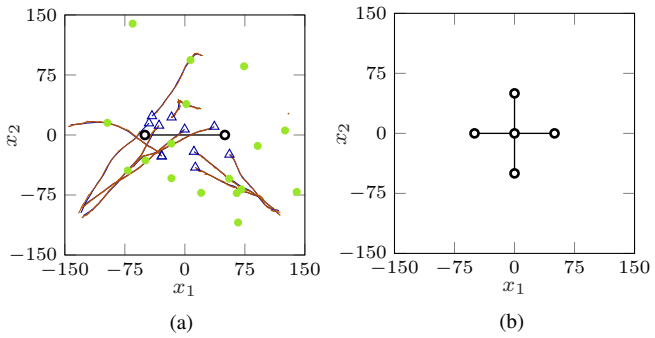

Fig. 2: (a) Example of the object trajectories (represented by blue lines, starting positions indicated by blue triangles) in simulation scenario SC1, as well as the corresponding estimates obtained with the proposed S-DLMB filter (represented by red lines). The positions of the two sensors are indicated by black circles. The green bullets show the measurements of the second sensor at time $k = 100$. (b) Sensors and communication links in simulation scenario SC2.

For SC1, Fig. 2(a) shows an example of the estimated trajectories obtained with S-DLMB at the second sensor (with position $\boldsymbol{p}^{(2)} = [50\ 0]^{\mathrm{T}}$); those obtained at the first sensor are similar. One can see that the estimated trajectories closely match the true trajectories.

For a quantitative assessment and comparison of the performance of the three filters, we computed the mean Euclidean distance based optimal subpattern assignment (MOSPA) metric [21] with cutoff parameter $c = 20$, order $p = 2$, and averaging over 1000 simulation runs and all the sensors. Fig. 3 shows the results for SC1 and SC2. The peaks in Fig. 3 correspond to object appearance and disappearance; note that several objects can appear or disappear at the same time. It is seen that S-DLMB almost always significantly outperforms H-DLMB; furthermore, it performs similarly to CLMB in SC1 and poorer than CLMB in SC2. These results show that the proposed soft label association fusion is a significant improvement over the hard label association fusion employed by H-DLMB, and the resulting LMB filter performance can come close to the performance of the centralized LMB filter based on the iterated-corrector approach.

Table I shows the average runtime per time (k) step of MATLAB implementations of S-DLMB, H-DLMB, and CLMB executed on an Intel quad core i7-6600U CPU. One can see that S-DLMB has the highest complexity, followed by H-DLMB and CLMB. However, we note that the lower complexity of H-DLMB—compared to S-DLMB—comes at the cost of a significantly poorer MOSPA performance.

[2] We do not show the results of the distributed LMB filter proposed in [12], because the LMB fusion performed by that filter uses the same label set for all the sensors, i.e., $\mathbb{L}^{(s)*} = \mathbb{L}^*$ for all s, and always matches Bernoulli components with equal labels. This is not compatible with our label indexing system, which always describes an object by different labels at different sensors. However, we note that with any other label indexing system, too, there is a high probability that an object is described by different labels at different sensors. This will generally result in a poor tracking performance of filters matching Bernoulli components with equal labels.

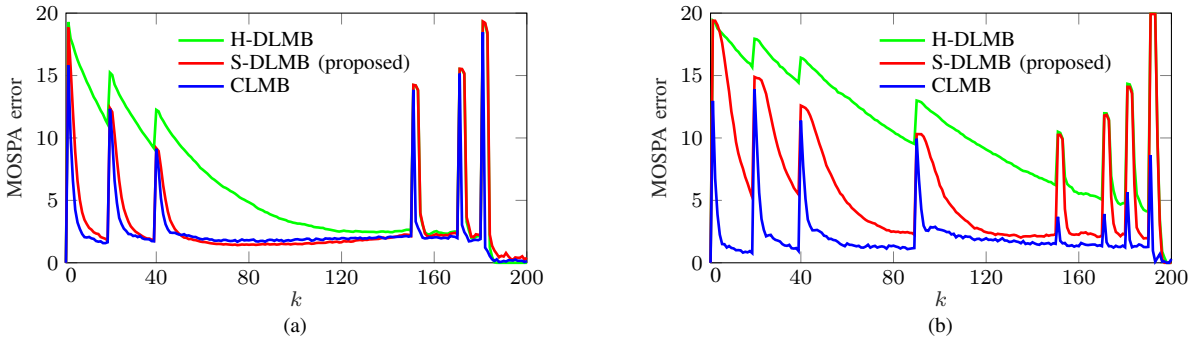

Fig. 3: MOSPA error versus time k for (a) SC1 and (b) SC2.

Filter	SC1	SC2
S-DLMB (proposed)	0.1072 s	2.2554 s
H-DLMB	0.0470 s	0.6682 s
CLMB	0.0186 s	0.4428 s

TABLE I: Measured average runtime per time step.

The communication requirements of S-DLMB and H-DLMB are generally similar. Indeed, for both S-DLMB and H-DLMB, in each fusion iteration, each local LMB filter broadcasts to its neighbors one set of Gaussian parameters per Bernoulli component.

IX. Conclusion

We proposed a labeled multi-Bernoulli (LMB) fusion method based on a new paradigm of probabilistic (soft) label association. An important aspect of LMB fusion is to correctly associate labeled Bernoulli components describing the same object at different sensors. Instead of enforcing a hard association of labels, our method calculates label association probabilities and uses them in the LMB fusion process. The proposed probabilistic association scheme was derived by reformulating the fused posterior multiobject distribution in a way that involves a label association distribution, and by exploiting the fact that an approximation of the label association distribution by the product of its marginals results in a fused multiobject distribution of LMB type. Based on this theoretical framework, we obtained an efficient LMB fusion algorithm by using a belief propagation technique for computing the marginal label association distributions and a Gaussian approximation for the spatial distributions. A distributed LMB filter using the proposed LMB fusion algorithm was demonstrated to outperform a state-of-the-art method using hard label association.

References

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

[2] R. P. S. Mahler, *Advances in Statistical Multisource-Multitarget Information Fusion*. Boston, MA, USA: Artech House, 2014.

[3] S. Challa, M. R. Morelande, D. Musicki, and R. Evans, *Fundamentals of Object Tracking*. Cambridge, UK: Cambridge University Press, 2011.

[4] F. Meyer, T. Kropfreiter, J. L. Williams, R. A. Lau, F. Hlawatsch, P. Braca, and M. Z. Win, "Message passing algorithms for scalable multitarget tracking," *Proc. IEEE*, vol. 106, no. 2, pp. 221–259, Feb. 2018.

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

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

[7] S. Reuter, B.-T. Vo, B.-N. Vo, and K. Dietmayer, "The labeled multi-Bernoulli filter," *IEEE Trans. Signal Process.*, vol. 62, no. 12, pp. 3246–3260, Jun. 2014.

[8] T. Kropfreiter, F. Meyer, and F. Hlawatsch, "A fast labeled multi-Bernoulli filter using belief propagation," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 56, no. 3, Jun. 2020.

[9] D. Clark, S. Julier, R. P. S. Mahler, and B. Ristić, "Robust multi-object sensor fusion with unknown correlations," in *Proc. SSPD-10*, London, UK, Sep. 2010.

[10] M. Uney, D. Clark, and S. J. Julier, "Distributed fusion of PHD filters via exponential mixture densities," *IEEE J. Sel. Top. Signal Process.*, vol. 7, no. 3, pp. 521–531, Jun. 2013.

[11] G. Battistelli, L. Chisci, C. Fantacci, A. Farina, and B.-N. Vo, "Average Kullback-Leibler divergence for random finite sets," in *Proc. FUSION-15*, Washington, DC, USA, Jul. 2015, pp. 1359–1366.

[12] C. Fantacci, B.-N. Vo, B.-T. Vo, G. Battistelli, and L. Chisci, "Robust fusion for multisensor multiobject tracking," *IEEE Signal Process. Lett.*, vol. 25, no. 5, pp. 640–644, May 2018.

[13] S. Li, G. Battistelli, L. Chisci, W. Yi, B. Wang, and L. Kong, "Computationally efficient multi-agent multi-object tracking with labeled random finite sets," *IEEE Trans. Signal Process.*, vol. 67, no. 1, pp. 260–275, Jan. 2019.

[14] J. Williams and R. Lau, "Approximate evaluation of marginal association probabilities with belief propagation," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 50, no. 4, pp. 2942–2959, Oct. 2014.

[15] A. Buonviri, M. York, K. LeGrand, and J. Meub, "Survey of challenges in labeled random finite set distributed multi-sensor multi-object tracking," in *Proc. IEEE AEROCONF-19*, Big Sky, MT, USA, Mar. 2019.

[16] F. R. Kschischang, B. J. Frey, and H.-A. Loeliger, "Factor graphs and the sum-product algorithm," *IEEE Trans. Inf. Theory*, vol. 47, no. 2, pp. 498–519, Feb. 2001.

[17] G. Battistelli, L. Chisci, C. Fantacci, A. Farina, and A. Graziano, "Consensus CPHD filter for distributed multitarget tracking," *IEEE J. Sel. Top. Signal Process.*, vol. 7, no. 3, pp. 508–520, Jun. 2013.

[18] D. P. Malladi and J. L. Speyer, "A new approach to multiple model adaptive estimation," in *Proc. IEEE CDC-97*, vol. 4, San Diego, CA, USA, Dec. 1997, pp. 3460–3467.

[19] A. G. Dimakis, S. Kar, J. M. F. Moura, M. G. Rabbat, and A. Scaglione, "Gossip algorithms for distributed signal processing," *Proc. IEEE*, vol. 98, no. 11, pp. 1847–1864, Nov. 2010.

[20] Y. Bar-Shalom, X.-R. Li, and T. Kirubarajan, *Estimation with Applications to Tracking and Navigation*. New York, NY, USA: Wiley, 2002.

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