The Labeled Multi-Bernoulli Filter

Stephan Reuter*, Ba-Tuong Vo†, Ba-Ngu Vo†, Klaus Dietmayer*

***Abstract*—This paper proposes a generalization of the multi-Bernoulli filter called the labeled multi-Bernoulli filter that outputs target tracks. Moreover, the labeled multi-Bernoulli filter does not exhibit a cardinality bias due to a more accurate update approximation compared to the multi-Bernoulli filter by exploiting the conjugate prior form for labeled Random Finite Sets. The proposed filter can be interpreted as an efficient approximation of the δ-Generalized Labeled Multi-Bernoulli filter. It inherits the advantages of the multi-Bernoulli filter in respect of particle implementation and state estimation. It also inherits advantages of the δ-Generalized Labeled Multi-Bernoulli filter in that it outputs (labeled) target tracks and achieves better performance.**

***Index Terms*—Random finite set, marked point process, conjugate prior, Bayesian estimation, target tracking.**

I. Introduction

In multi-target tracking, the objective is to jointly estimate the number of trajectories and their states from a sequence of noisy and cluttered observation sets. A multi-target system is fundamentally different from a single-target system in that the number of states changes with time due to births and deaths of targets [1], [2], [3]. In addition, existing targets may or may not be detected and the sensor also receives a set of spurious measurements (clutter) not originating from any target. As a result, at each time step the measurement is a collection of detections, only some of which are generated by targets.

Multi-target tracking is an established subfield of signal processing with applications spanning many areas such as radar, sonar [1], computer vision [4], robotics [5], cell biology [6], vehicle environment perception [7], etc. The three major approaches to multi-target tracking are Multiple Hypotheses Tracking (MHT) [1], Joint Probabilistic Data Association (JPDA) [2], and the Random Finite Set (RFS) framework [3]. MHT and its variations involve the propagation of data association hypotheses in time while the JPDA approach weights individual observations by their association probabilities. MHT and JPDA are classical approaches that have dominated the field of multi-target tracking and are well documented in [1], [2]. The RFS approach is the latest development that provides a general systematic treatment of multi-target systems by modeling the multi-target state as an RFS. The centerpiece of the RFS approach is the so-called *Bayes multi-target filter* that propagates the posterior density of the multi-target state recursively in time.

Due to the numerical complexity of the Bayes multi-target filter, the Probability Hypothesis Density (PHD) [8], Cardinalized PHD (CPHD) [9], and multi-Bernoulli filters [10], [11] have been developed as approximations. The PHD and CPHD recursions propagate moments and cardinality distributions while the multi-Bernoulli filter propagates the parameters of a multi-Bernoulli distribution that approximate the posterior multi-target density. Gaussian mixture and particle implementations of these filters have been developed in [10], [12], [13], [14], [15], [16]. Note that the multi-Bernoulli filter in [10] operates on point measurements while the multi-Bernoulli filter in [11] operates on image measurements.

While the PHD and CPHD filters are synonymous with the RFS approach, the lesser known multi-Bernoulli filters are more useful in problems that require particle implementations or object individual existence probabilities. Formal convergence results for the particle multi-Bernoulli filter were established in [17]. Unlike the particle PHD and CPHD filters, the particle multi-Bernoulli filter does not require additional processing for extracting the multi-target state estimate. A number of extensions to the multi-Bernoulli filter have been considered in [18], [19], [20], [21], [22], [23]. The multi-Bernoulli filter has been applied to tracking in sensor networks [24], [25], [26], [27], [28], [29], tracking from audio and video data [30], as well as visual tracking and cell tracking [31], [6], [32]. Hybrid multi-Bernoulli and Poisson multi-target filters were also considered in [33].

This paper proposes a generalization of the multi-Bernoulli filter called the labeled multi-Bernoulli filter (LMB) that offers two important advantages. The multi-Bernoulli filter does not formally accommodate target tracks, and assumes high signal to noise ratio, i.e. high detection probability and low clutter, which is quite restrictive in real applications. On the other hand, by utilizing the labeled RFS formulation [34], [35], the proposed LMB filter formally estimates target tracks and does not require high signal to noise ratio. The proposed LMB filter also does not exhibit the so called "spooky effect" [36], a phenomenon present in the CPHD filter due its use of an i.i.d. cluster approximation, which causes a redistribution of target masses within the PHD when missed detections are encountered. In addition, we propose a dynamic grouping procedure which drastically reduces execution time by enabling parallelization, as well as an adaptive birth model that obviates the need for a static a priori specification. The proposed filter can also be interpreted as an efficient approximation of the δ-GLMB filter in [34], [35] (empirical results in this paper show orders of magnitude increase in speed) as well as an

Acknowledgement: This work is supported by the German Research Foundation (DFG) within the Transregional Collaborative Research Center SFB/TRR 62 "Companion-Technology for Cognitive Technical Systems" and by the Australian Research Council under schemes DE120102388 and FT0991854.

* S. Reuter and K. Dietmayer are with the Institute of Measurement, Control and Microtechnology, Ulm University, 89081 Ulm, Germany, e-mail: {stephan.reuter, klaus.dietmayer}@uni-ulm.de.

† B.-T. Vo and B.-N. Vo are with Department of Electrical and Computer Engineering, Curtin University, Bentley, WA 6102, Australia. E-mail: {ba-tuong.vo, ba-ngu.vo}@curtin.edu.au

improved approximation compared to the multi-Bernoulli filter in [10]. It thus inherits the advantages of the multi-Bernoulli filter with straightforward particle implementations and state estimation. It also inherits the advantages of the δ-GLMB filter in that it outputs target tracks and gives better performance. The disadvantage is the higher computational cost (than the multi-Bernoulli filter in [10]), which is at worst cubic in the number of measurements. In contrast, classical approaches such as MHT and JPDA are generally NP-hard.

The paper is organized as follows. In Section II, we review the necessary background material on RFS and the Bayes multi-target filter. Section III then establishes the basic results which underpin the time prediction and data update steps of the proposed labeled multi-Bernoulli filter. Section IV then gives a full description of the filter implementation, including grouping, gating, and adaptive birth strategies. Section V then presents numerical results and comparisons with existing RFS filters, as well as a challenging scenario with over one hundred targets. Section VI presents concluding remarks.

II. Background

This section provides a brief review of multi-Bernoulli and generalized labeled multi-Bernoulli RFS necessary for the results of this paper. For further details on mathematical proofs, intuitive interpretations, modeling, analysis and numerical implementation, we refer the reader to [34], [35] and the accompanying report [37].

A. Bayes Multi-Target Filter

The Bayes multi-target filter [3] is a rigorous extension of the Bayes (single-target) filter to multi-target states using RFSs. An RFS is a finite-set-valued random variable, i.e. it consists of a random number of points, which are unordered, and whose individual states are random (usually random vectors). Thus, instead of a single-target state vector $x \in \mathbb{X}$, a finite set of state vectors $X \subset \mathbb{X}$, is used in the Bayes multi-target filter.

If the multi-target prior is distributed according to $\pi(X)$, the multi-target posterior is given by

$$\pi(X|Z) = \frac{g(Z|X)\pi(X)}{\int g(Z|X)\pi(X)\delta X} \tag{1}$$

where the integral in the denominator is a set integral defined by

$$\int f(X)\delta X = \sum_{i=0}^{\infty} \frac{1}{i!} \int_{\mathbb{X}^i} f(\{x_1, ..., x_i\}) d(x_1, \cdots, x_i).$$

The multi-target posterior contains all information about the number of targets and their current states given the current set of measurements $Z \subset \mathbb{Z}$. The multi-target likelihood function $g(Z|X)$ models the measurement process of the sensor. In addition to the measurement noise, $g(Z|X)$ incorporates detection and data association uncertainty as well as clutter or false alarms in the following sense. Given a multi-target state X, the measurement set Z is obtained by generating detections and clutter. Each state $x \in X$ is detected with probability $p_D(\cdot)$, and conditional upon detection generates a measurement $z \in Z$ according to the single-target likelihood $g(z|x)$, while clutter is assumed to follow a Poisson process with intensity function $\kappa(\cdot)$ defined on Z. Thus $g(Z|X)$ is the probability density for a conditional observation and is completely parameterized by p_D, $g(z|x)$ and κ.

In order to update the multi-target posterior distribution with the next set of measurements, the distribution is predicted to the time of the next observation using the multi-target Chapman-Kolmogorov equation [3]

$$\pi_+(X_+) = \int f(X_+|X)\pi(X)\delta X. \tag{2}$$

The multi-target Markov transition density $f(X_+|X)$ models the evolution of the multi-target state incorporating the motion of the objects in addition to target births and deaths in the following sense. Given a multi-target state X, the predicted set X_+ is obtained by generating survivals and births. Each state $x \in X$ survives or persists with probability $p_S(\cdot)$, and conditional upon survival transitions to a new state $x_+ \in X_+$ according to the single-target transition $f(x_+|x)$, while births are assumed to follow a multi-Bernoulli process π_B specified by a set of birth probabilities and initial tracks $\{r_B^{(\ell)}, p_B^{(\ell)}(x)\}$. Thus $f(X_+|X)$ is the probability density of a transition and is completely parameterized by p_S, $f(x_+|x)$ and π_B.

B. Notation

In this paper, single-target states are denoted by small letters (e.g. x) and multi-target states are denoted by capital letters (e.g. X). In order to distinguish labeled states and distributions (see II-D) from unlabeled ones, bold face letters are used for labeled ones (e.g. $\mathbf{x}$, $\mathbf{X}$, $\boldsymbol{\pi}$). Additionally, blackboard bold letters represent spaces, e.g. the state space is represented by $\mathbb{X}$ and the measurement space by $\mathbb{Z}$. The collection of all finite subsets of $\mathbb{X}$ is denoted by $\mathcal{F}(\mathbb{X})$ and $\mathcal{F}_n(\mathbb{X})$ denotes all finite subsets with exactly n elements.

For notational convenience, several abbreviations are used throughout the paper. The inner product is denoted by

$$\langle f, g \rangle \triangleq \int f(x)g(x)dx.$$

We use the multi-object exponential notation

$$h^X \triangleq \prod_{x \in X} h(x) \tag{3}$$

for real-valued functions h, with $h^{\emptyset} = 1$ by convention.

In order to support arbitrary arguments like sets, vectors and integers, the generalized Kronecker delta function is given by

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

and the inclusion function is given by

$$1_Y(X) \triangleq \begin{cases} 1, & \text{if } X \subseteq Y \\ 0, & \text{otherwise.} \end{cases}$$

If X is a singleton, i.e. $X = \{x\}$, the notation $1_Y(x)$ is used instead of $1_Y(\{x\})$.

C. Multi-Bernoulli RFS

A Bernoulli RFS X is empty with a probability of $1-r$ or is a singleton with probability r with a spatial distribution $p(\cdot)$ on the state space $\mathbb{X}$. Its cardinality distribution follows a Bernoulli distribution with parameter r and its probability density is given by [3, p. 368]

$$\pi(X) = \begin{cases} 1-r & X = \emptyset, \\ r \cdot p(x) & X = \{x\}. \end{cases} \tag{4}$$

A multi-Bernoulli RFS X is given by the union of the M independent Bernoulli RFSs $X^{(i)}$, i.e. $X = \bigcup_{i=1}^{M} X^{(i)}$, and is completely defined by the parameter set $\{(r^{(i)}, p^{(i)})\}_{i=1}^{M}$. The probability density of a multi-Bernoulli RFS is given by [3, p. 368]

$$\pi(\{x_1,...,x_n\}) = \prod_{j=1}^{M}\left(1-r^{(j)}\right) \times \sum_{1\leq i_1\neq...,\neq i_n\leq M} \prod_{j=1}^{n} \frac{r^{(i_j)} p^{(i_j)}(x_j)}{1-r^{(i_j)}}, \tag{5}$$

The form of the density can be interpreted as follows. The sum is taken over all permutations of $n \leq M$ of the M constituent Bernoulli RFSs. The numerator of the products within the sum is the probability density that the Bernoulli components with indices $i_1,...,i_n$ generate the realizations $x_1,...,x_n$. The leading constant then cancels with the divisors inside the sum to give the probability that the leftover Bernoulli components with indices $1,...,M-i_1,...,i_n$ generate null realizations $\emptyset$. In the following, the abbreviation $\pi = \{(r^{(i)}, p^{(i)})\}_{i=1}^{M}$ is adopted for the probability density. Further, the cardinality distribution of a multi-Bernoulli RFS is given by neglecting the spatial distribution in (5):

$$\rho(n) = \prod_{j=1}^{M}\left(1-r^{(j)}\right) \sum_{1\leq i_1\neq...,\neq i_n\leq M} \prod_{j=1}^{n} \frac{r^{(i_j)}}{1-r^{(i_j)}}. \tag{6}$$

D. Labeled Multi-Bernoulli RFS

In this work, a state $x \in \mathbb{X}$ is augmented by a label $\ell \in \mathbb{L}$ in order to be able to estimate the identity or trajectory of a single target in a multi-target scenario. The labels are usually drawn from a discrete label space $\mathbb{L} = \{\alpha_i : i \in \mathbb{N}\}$, where all α_i are distinct and the index space $\mathbb{N}$ is the set of positive integers.

The following convention is adopted in respect of target labels and label spaces. Labels for individual targets are ordered pairs of integers $\ell = (k,i)$, where k is the time of birth, and $i \in \mathbb{N}$ is a unique index to distinguish new targets born at the same time. Target labels are thus unique and static. The label space for new targets born at time k, denoted as $\mathbb{L}_k$, is then $\{k\} \times \mathbb{N}$. A new target born at time k, has state $\mathbf{x} \in \mathbb{X}\times\mathbb{L}_k$. The label space for all targets at time k (survivals and births), denoted as $\mathbb{L}_{0:k}$, is constructed recursively by $\mathbb{L}_{0:k} = \mathbb{L}_{0:k-1} \cup \mathbb{L}_k$. Where no confusion arises the abbreviation $\mathbb{L} \triangleq \mathbb{L}_{0:k}$ is used for compactness.

The notion of labeled random finite sets is introduced in [35]. A labeled RFS with state space $\mathbb{X}$ and label space $\mathbb{L}$ is simply an RFS on $\mathbb{X}\times\mathbb{L}$ with distinct labels, that is, a finite-set-valued random variable on $\mathbb{X}\times\mathbb{L}$ such that all realizations of labels are unique.

The set of labels of a labeled RFS $\mathbf{X}$ is given by $\mathcal{L}(\mathbf{X}) = \{\mathcal{L}(\mathbf{x}) : \mathbf{x} \in \mathbf{X}\}$, where $\mathcal{L} : \mathbb{X}\times\mathbb{L} \to \mathbb{L}$ is the projection defined by $\mathcal{L}((x,\ell)) = \ell$. For each realization $\mathbf{X}$ of a labeled RFS, the labels must be distinct, i.e. $|\mathcal{L}(\mathbf{X})| = |\mathbf{X}|$, where $|\cdot|$ denotes the cardinality of a set. The distinct label indicator

$$\Delta(\mathbf{X}) = \delta_{|\mathbf{X}|}(|\mathcal{L}(\mathbf{X})|) \tag{7}$$

is used to ensure distinct labels. By projecting a labeled RFS from $\mathbb{X}\times\mathbb{L}$ into $\mathbb{X}$, its unlabeled version is obtained. A key characteristic of a labeled RFS is that its cardinality distribution is equivalent to the cardinality distribution of its unlabeled version.

Like a multi-Bernoulli RFS, a labeled multi-Bernoulli (LMB) RFS is completely described using a parameter set $\{(r^{(\zeta)}, p^{(\zeta)}) : \zeta \in \Psi\}$ with index set Ψ. If a single Bernoulli component $(r^{(\zeta)}, p^{(\zeta)})$ returns a non-empty set, a label $\alpha(\zeta)$ is attached, where for simplicity it is assumed that $\alpha : \Psi \to \mathbb{L}$ is a 1-1 mapping. A procedure for sampling from a labeled multi-Bernoulli RFS can be found in [35].

Since α is a 1-1 mapping, the labels of a labeled multi-Bernoulli RFS are distinct. The density of the labeled multi-Bernoulli RFS on the space $\mathbb{X}\times\mathbb{L}$ is given by

$$\boldsymbol{\pi}(\{(x_1,\ell_1),...,(x_n,\ell_n)\}) = \delta_n(|\{\ell_1,...\ell_n\}|) \prod_{\zeta\in\Psi}\left(1-r^{(\zeta)}\right) \times \prod_{j=1}^{n} \frac{1_{\alpha(\Psi)}(\ell_j) r^{(\alpha^{-1}(\ell_j))} p^{(\alpha^{-1}(\ell_j))}(x_j)}{1-r^{(\alpha^{-1}(\ell_j))}}, \tag{8}$$

and for convenience the abbreviation $\boldsymbol{\pi} = \{(r^{(\zeta)}, p^{(\zeta)})\}_{\zeta\in\Psi}$ will also be used for the density of an LMB RFS.

Although the formulation of a labeled Multi-Bernoulli RFS allows for a general mapping α for the labels, in this work we will assume that α is an identity mapping in order to simplify notations. Thus superscripts for component indices correspond directly to the label in question. The density of an LMB RFS with parameter set $\boldsymbol{\pi} = \{r^{(\ell)}, p^{(\ell)}\}_{\ell\in\mathbb{L}}$ is given more compactly by

$$\boldsymbol{\pi}(\mathbf{X}) = \Delta(\mathbf{X}) w(\mathcal{L}(\mathbf{X})) p^{\mathbf{X}} \tag{9}$$

where

$$w(L) = \prod_{i\in\mathbb{L}}\left(1-r^{(i)}\right) \prod_{\ell\in L} \frac{1_{\mathbb{L}}(\ell) r^{(\ell)}}{1-r^{(\ell)}}, \tag{10}$$

$$p(x,\ell) = p^{(\ell)}(x). \tag{11}$$

E. Generalized Labeled Multi-Bernoulli RFS

A generalized labeled multi-Bernoulli (GLMB) RFS [35] is a labeled RFS with state space $\mathbb{X}$ and (discrete) label space $\mathbb{L}$ distributed according to

$$\boldsymbol{\pi}(\mathbf{X}) = \Delta(\mathbf{X}) \sum_{c\in\mathbb{C}} w^{(c)}(\mathcal{L}(\mathbf{X})) \left[p^{(c)}\right]^{\mathbf{X}} \tag{12}$$

where $\mathbb{C}$ is a discrete index set. The weights $w^{(c)}(L)$ and the spatial distributions $p^{(c)}$ satisfy the normalization conditions

$$\sum_{L\subseteq\mathbb{L}}\sum_{c\in\mathbb{C}} w^{(c)}(L) = 1, \quad (13)$$

$$\int p^{(c)}(x,\ell)dx = 1. \quad (14)$$

The labeled multi-Bernoulli RFS distributed by (8) is a special case of the generalized labeled multi-Bernoulli RFS with

$$\begin{aligned} p^{(c)}(x,\ell) &= p^{(\ell)}(x) \\ w^{(c)}(L) &= \prod_{i\in\mathbb{L}}\left(1-r^{(i)}\right)\prod_{\ell\in L}\frac{1_{\mathbb{L}}(\ell)r^{(\ell)}}{1-r^{(\ell)}} \end{aligned}$$

comprising a single component in which case the superscript (c) is omitted.

Intuitively, and in the context of multi-target tracking, the fundamental difference between a GLMB and LMB process is that of a mixture versus single component representation. The sum over $c \in \mathbb{C}$ in the definition of the GLMB process facilitates the propagation of multiple "hypotheses", involving different sets of track labels, which arises due to data association uncertainty seen in the update step of the Bayes multi-target filter. With only a single component as in the definition of the LMB process, it is only possible to propagate the uncertainty for a single set of track labels L, although this can be exploited for computational savings.

F. δ-Generalized Labeled Multi-Bernoulli RFS

A δ-generalized labeled multi-Bernoulli (δ-GLMB) RFS [35] with state space $\mathbb{X}$ and (discrete) label space $\mathbb{L}$ is a special case of a generalized labeled multi-Bernoulli RFS with

$$\begin{aligned} \mathbb{C} &= \mathcal{F}(\mathbb{L})\times\Xi \\ w^{(c)}(L) &= w^{(I,\xi)}\delta_I(L) \\ p^{(c)} &= p^{(I,\xi)} = p^{(\xi)} \end{aligned}$$

where Ξ is a discrete space, ξ are realizations of Ξ, and I denotes a set of track labels. In target tracking applications, the discrete space Ξ typically represents the history of track to measurement associations. A δ-GLMB RFS is thus a special case of a GLMB RFS but with a particular structure on the index space which arises naturally in target tracking applications. The δ-GLMB RFS has density

$$\boldsymbol{\pi}(\mathbf{X}) = \Delta(\mathbf{X})\sum_{(I,\xi)\in\mathcal{F}(\mathbb{L})\times\Xi} w^{(I,\xi)}\delta_I(\mathcal{L}(\mathbf{X}))\left[p^{(\xi)}\right]^{\mathbf{X}}. \quad (15)$$

and its cardinality distribution is given by

$$\begin{aligned} \boldsymbol{\rho}(n) &= \sum_{(I,\xi)\in\mathcal{F}(\mathbb{L})\times\Xi}\sum_{L\in\mathcal{F}_n(\mathbb{L})} w^{(I,\xi)}\delta_I(L) \\ &= \sum_{(I,\xi)\in\mathcal{F}_n(\mathbb{L})\times\Xi} w^{(I,\xi)}. \end{aligned} \quad (16)$$

The PHD of the corresponding unlabeled RFS is given by

$$\begin{aligned} v(x) &= \sum_{(I,\xi)\in\mathcal{F}(\mathbb{L})\times\Xi}\sum_{\ell\in\mathbb{L}} p^{(\xi)}(x,\ell)\sum_{L\subseteq\mathbb{L}} 1_L(\ell)w^{(I,\xi)}\delta_I(L) \\ &= \sum_{\ell\in\mathbb{L}}\sum_{(I,\xi)\in\mathcal{F}(\mathbb{L})\times\Xi} w^{(I,\xi)}1_I(\ell)p^{(\xi)}(x,\ell). \end{aligned} \quad (17)$$

Due to the inclusion function $1_I(\ell)$, each summand of the inner sum of (17) is zero if track ℓ is not included in the set of track labels I. The inner sum thus represents the PHD of track ℓ. The existence probability for a track ℓ may be obtained by the sum of the weights of the corresponding joint label set and discrete index pairs (with the latter herein referred to as a "hypothesis")

$$r^{(\ell)} = \sum_{(I,\xi)\in\mathcal{F}(\mathbb{L})\times\Xi} w^{(I,\xi)}1_I(\ell).$$

III. Premise of the Labeled Multi-Bernoulli Filter

In [10], an approximation to the Bayes multi-target filter was proposed based on approximating the posterior as a multi-Bernoulli RFS. Track labels were assigned to individual components by post processing in order to estimate target trajectories. Although the time prediction step was exact, the data update approximated the posterior via the probability generating functional (PGFl) in two steps. The first approximation expresses the posterior PGFl as a product of legacy and updated terms which are not necessarily in multi-Bernoulli form (see eq. (18) in [10]). The second approximation chooses a multi-Bernoulli PGFl that matches the posterior cardinality mean of the first PGFl approximation (see Section III-C in [10]). This two stage approximation resulted in a computationally inexpensive multi-target tracking filter which was based on the assumption of a high signal-to-noise ratio for sensor detections. Although performance was accurate in such situations, a significant cardinality bias was observed in harsher signal environments, due mainly to the first PGFl approximation.

In [35], it was shown that the δ-GLMB density is closed under the multi-target prediction and update operations. This result was used to develop a δ-GLMB filter which formally produces tracks. Moreover the only approximation used was a direct truncation of the multi-target posterior rather than approximating the multi-target PGFl. The only disadvantage is that the tracking filter exhibits an exponential growth in the number of posterior components. Tractable techniques for truncating the posterior and prediction densities were proposed based on the k-shortest paths [38], [39] and ranked assignment algorithms [40]. The resultant δ-GLMB filter outperforms the CPHD and multi-Bernoulli filters but is more computationally expensive. See [35] for full details.

Using labeled multi-Bernoulli RFSs we can exploit the intuitive mathematical structure of the multi-Bernoulli RFS, without the weaknesses of the multi-Bernoulli filter, which are that it does not formally produce track estimates and exhibits cardinality bias. The premise of the proposed labeled multi-Bernoulli (LMB) filter is the approximation of the

predicted and posterior densities in (1)-(2) as a labeled multi-Bernoulli process given by (9). The number of components for a δ-GLMB representation grows exponentially, whereas the growth is linear for an LMB representation. Compared to the multi-Bernoulli filter, a suitable LMB approximation is much more accurate, with an exact match of the first posterior moment and hence no cardinality bias. In this section we derive the exact time prediction step and the approximate measurement update. These results are used in the following section to describe the full implementation of the labeled multi-Bernoulli filter.

A. Prediction

Note from [35] that the GLMB (and δ-GLMB) density is closed under the prediction step, albeit the number of terms in the predicted GLMB expression grows exponentially. While an LMB is a special case of a GLMB (with one term) it is not necessarily true that the prediction of an LMB density is still an LMB density. Indeed, suppose that the prior and birth are both LMBs:

$$\boldsymbol{\pi}(\mathbf{X}) = \Delta(\mathbf{X})w(\mathcal{L}(\mathbf{X}))p^{\mathbf{X}} \quad (18)$$

$$\boldsymbol{\pi}_B(\mathbf{X}) = \Delta(\mathbf{X})w_B(\mathcal{L}(\mathbf{X}))\left[p_B\right]^{\mathbf{X}} \quad (19)$$

where

$$w(L) = \prod_{i\in\mathbb{L}}\left(1-r^{(i)}\right)\prod_{\ell\in L}\frac{1_{\mathbb{L}}(\ell)r^{(\ell)}}{1-r^{(\ell)}}, \quad (20)$$

$$w_B(L) = \prod_{i\in\mathbb{B}}\left(1-r_B^{(i)}\right)\prod_{\ell\in L}\frac{1_{\mathbb{B}}(\ell)r_B^{(\ell)}}{1-r_B^{(\ell)}}, \quad (21)$$

$$p(x,\ell) = p^{(\ell)}(x), \quad (22)$$

$$p_B(x,\ell) = p_B^{(\ell)}(x), \quad (23)$$

then it follows from Proposition 8 of [35], that the prediction of an LMB is a GLMB with state space $\mathbb{X}$ and (finite) label space $\mathbb{L}_+ = \mathbb{B}\cup\mathbb{L}$ ($\mathbb{L}\cap\mathbb{B} = \varnothing$) given by

$$\boldsymbol{\pi}_+(\mathbf{X}_+) = \Delta(\mathbf{X}_+)w_+(\mathcal{L}(\mathbf{X}_+))\left[p_+\right]^{\mathbf{X}_+} \quad (24)$$

where

$$w_+(I_+) = w_B(I_+\cap\mathbb{B})w_S(I_+\cap\mathbb{L}) \quad (25)$$

$$p_+(x,\ell) = 1_{\mathbb{L}}(\ell)p_{+,S}(x,\ell) + 1_{\mathbb{B}}(\ell)p_B(x,\ell) \quad (26)$$

$$p_{+,S}(x,\ell) = \frac{\langle p_S(\cdot,\ell)f(x|\cdot,\ell), p(\cdot,\ell)\rangle}{\eta_S(\ell)} \quad (27)$$

$$\eta_S(\ell) = \langle p_S(\cdot,\ell), p(\cdot,\ell)\rangle \quad (28)$$

$$w_S(L) = [\eta_S]^L\sum_{I\supseteq L}[1-\eta_S]^{I-L}w(I) \quad (29)$$

$p_S(x,\ell)$ is the state dependent survival probability, $\eta_S(\ell)$ is the survival probability of track ℓ, and $f(x|x',\ell)$ is the single target transition density for track ℓ.

Note that $w_B(L)$ is the weight of an LMB. However, the weight of the prediction density $w_+(I_+)$ in (25) does not necessarily satisfy that of an LMB because the term $w_S(L)$ given by (29) consists of a sum over the supersets of L rather than a product over L of the form (20). Remarkably, the sum in (29) can be factorized into a product, and consequently written as the weight of an LMB, via the following Lemma.

Lemma 1

$$\frac{(1-r^{(\cdot)})^{\mathbb{L}}\eta_S^L}{(1-\eta_S)^L}\sum_{I\supseteq L}(1-\eta_S)^I\left(\frac{r^{(\cdot)}}{1-r^{(\cdot)}}\right)^I = (1-r^{(\cdot)}\eta_S)^{\mathbb{L}}\left(\frac{r^{(\cdot)}\eta_S}{1-r^{(\cdot)}\eta_S}\right)^L \quad (30)$$

Proof: Let

$$f^{(L)}(I) = (1-r^{(\cdot)})^{\mathbb{L}}\eta_S^L\frac{(1-\eta_S)^I}{(1-\eta_S)^L}\left(\frac{r^{(\cdot)}}{1-r^{(\cdot)}}\right)^I. \quad (31)$$

Then for any $\ell\in L$,

$$f^{(L\backslash\{\ell\})}(I) = \frac{1-\eta_S(\ell)}{\eta_S(\ell)}f^{(L)}(I) = \frac{r^{(\ell)}-r^{(\ell)}\eta_S(\ell)}{r^{(\ell)}\eta_S(\ell)}f^{(L)}(I) \quad (32)$$

$$f^{(L\backslash\{\ell\})}(I\backslash\{\ell\}) = \frac{1-r^{(\ell)}}{r^{(\ell)}\eta_S(\ell)}f^{(L)}(I). \quad (33)$$

Hence, adding the two expressions above gives

$$f^{(L\backslash\{\ell\})}(I)+f^{(L\backslash\{\ell\})}(I\backslash\{\ell\}) = \frac{1-r^{(\ell)}\eta_S(\ell)}{r^{(\ell)}\eta_S(\ell)}f^{(L)}(I). \quad (34)$$

Note also that for any function $g:\mathcal{F}(\mathbb{L})\to\mathbb{R}$

$$\sum_{I\supseteq\{\ell_1,...,\ell_{n-1}\}}g(I) = \sum_{I\supseteq\{\ell_1,...,\ell_n\}}[g(I)+g(I\backslash\{\ell_n\})] \quad (35)$$

since the only sets which contain $\{\ell_1,...,\ell_{n-1}\}$ are those sets I which contain $\{\ell_1,...,\ell_n\}$ and the sets $I\backslash\{\ell_n\}$.

Now the LHS of eq (30) is $\sum_{I\supseteq L}f^{(L)}(I)$ which we denote by $f(L)$. To prove the lemma we show that $f(L)$ is equal to the RHS by induction. The lemma holds for $L=\mathbb{L}$. Assume the result holds for some $L=\{\ell_1,...,\ell_n\}$. Then

$$\begin{aligned}
f(L\backslash\{\ell_n\}) &= \sum_{I\supseteq L\backslash\{\ell_n\}}f^{(L\backslash\{\ell_n\})}(I)\\
&= \sum_{I\supseteq L}\left[f^{(L\backslash\{\ell_n\})}(I)+f^{(L\backslash\{\ell_n\})}(I\backslash\{\ell_n\})\right] \quad \text{(by (35))}\\
&= \left(\frac{1-r^{(\ell_n)}\eta_S(\ell_n)}{r^{(\ell_n)}\eta_S(\ell_n)}\right)\sum_{I\supseteq L}f^{(L)}(I) \quad \text{(by 34)}\\
&= \left(\frac{1-r^{(\ell_n)}\eta_S(\ell_n)}{r^{(\ell_n)}\eta_S(\ell_n)}\right)(1-r^{(\cdot)}\eta_S)^{\mathbb{L}}\left(\frac{r^{(\cdot)}\eta_S}{1-r^{(\cdot)}\eta_S}\right)^L\\
&= (1-r^{(\cdot)}\eta_S)^{\mathbb{L}}\left(\frac{r^{(\cdot)}\eta_S}{1-r^{(\cdot)}\eta_S}\right)^{L\backslash\{\ell_n\}}
\end{aligned}$$

where the second to last line follows from the assumption. Hence the Lemma holds by induction. ■

Proposition 2 *Suppose that the multi-target posterior density is an LMB RFS with state space $\mathbb{X}$, (finite) label space $\mathbb{L}$, and parameter set $\boldsymbol{\pi} = \{r^{(\ell)}, p^{(\ell)}\}_{\ell\in\mathbb{L}}$, and that the multi-target birth model is an LMB RFS with state space $\mathbb{X}$, (finite) label space $\mathbb{B}$, and parameter set $\boldsymbol{\pi}_B = \{r_B^{(\ell)}, p_B^{(\ell)}\}_{\ell\in\mathbb{B}}$, then the*

multi-target predicted density is also an LMB RFS with state space $\mathbb{X}$ and finite label space $\mathbb{L}_+ = \mathbb{B} \cup \mathbb{L}$ (with $\mathbb{L} \cap \mathbb{B} = \varnothing$) given by the parameter set

$$\boldsymbol{\pi}_+ = \left\{\left(r_{+,S}^{(\ell)}, p_{+,S}^{(\ell)}\right)\right\}_{\ell\in\mathbb{L}} \cup \left\{\left(r_B^{(\ell)}, p_B^{(\ell)}\right)\right\}_{\ell\in\mathbb{B}}, \quad (36)$$

where

$$r_{+,S}^{(\ell)} = \eta_S(\ell) r^{(\ell)}, \quad (37)$$

$$p_{+,S}^{(\ell)} = \left\langle p_S(\cdot,\ell) f(x|\cdot,\ell), p(\cdot,\ell)\right\rangle / \eta_S(\ell), \quad (38)$$

Proof: Enumerate the labels $\mathbb{L} = \{\ell_1, ..., \ell_M\}$ and rewrite the weight $w_S(L)$ in eq (29), in the following form:

$$w_S(L) = \sum_{I\subseteq\mathbb{L}} 1_I(L) f^{(L)}(I) = \sum_{I\supseteq L} f^{(L)}(I), \quad (39)$$

$$f^{(L)}(I) = (1-r^{(\cdot)})^{\mathbb{L}} \eta_S^L \frac{(1-\eta_S)^I}{(1-\eta_S)^L} \left(\frac{r^{(\cdot)}}{1-r^{(\cdot)}}\right)^I. \quad (40)$$

Then by the above Lemma

$$w_S(L) = \prod_{i\in\mathbb{L}} \left(1 - r^{(i)}\eta_S(i)\right) \prod_{\ell\in L} \frac{1_{\mathbb{L}}(\ell) r^{(\ell)} \eta_S(\ell)}{1 - r^{(\ell)}\eta_S(l)}$$
$$= (1 - r^{(\cdot)}\eta_S)^{\mathbb{L}} \left(\frac{r^{(\cdot)}\eta_S}{1 - r^{(\cdot)}\eta_S}\right)^L. \quad (41)$$

Hence, the LMB RFS due to the weight $w_S(L)$ has components $\left(r_{+,S}^{(\ell)}, p_{+,S}^{(\ell)}\right)$, $\ell \in \mathbb{L}$, given by (37)-(38). Since the weights $w_B(I_+ \cap \mathbb{B})$ and $w_S(I_+ \cap \mathbb{L})$ have the form of an LMB weight, it follows that $w_+(I_+) = w_B(I_+\cap\mathbb{B}) w_S(I_+\cap\mathbb{L})$ also has the form of an LMB weight. Consequently the LMB components for the prediction density are given by the birth LMB components $\left\{\left(r_B^{(\ell)}, p_B^{(\ell)}\right)\right\}_{\ell\in\mathbb{B}}$ and the survival LMB components $\left\{\left(r_{+,S}^{(\ell)}, p_{+,S}^{(\ell)}\right)\right\}_{\ell\in\mathbb{L}}$ are given by (37)-(38). ■

Remark 3 *The multi-target prediction for an LMB process actually coincides with performing the prediction on the unlabeled process and interpreting the component indices as track labels. Thus to perform the LMB filter prediction it is sufficient to predict the parameters forward according to* (36) *which is identical to the prediction for the multi-Bernoulli filter. This result is used later in implementations.*

B. Update

While the LMB family is closed under the prediction operation, it is not closed under the update operation. In this section we derive an LMB approximation of the posterior multi-target density. Inspired by the multi-Bernoulli filter, we seek the LMB approximation that matches the first moment of the multi-target posterior density [10]. The key advantage of the LMB update is that it only involves one approximation of the multi-target posterior density. In contrast the multi-Bernoulli filter requires two approximations on the multi-target posterior probability generating functional [10]. Consequently, apart from producing target tracks, the proposed LMB filter is superior to the multi-Bernoulli filter.

Proposition 4 *Suppose that the multi-target predicted density is an LMB RFS with state space $\mathbb{X}$, (finite) label space $\mathbb{L}_+$, and parameter set $\boldsymbol{\pi}_+ = \{r_+^{(\ell)}, p_+^{(\ell)}\}_{\ell\in\mathbb{L}_+}$. The LMB RFS that matches exactly the first moment of the multi-target posterior density is $\boldsymbol{\pi}(\cdot|Z) = \{(r^{(\ell)}, p^{(\ell)}(\cdot))\}_{\ell\in\mathbb{L}_+}$ where*

$$r^{(\ell)} = \sum_{(I_+,\theta)\in\mathcal{F}(\mathbb{L}_+)\times\Theta_{I_+}} w^{(I_+,\theta)}(Z) 1_{I_+}(\ell), \quad (42)$$

$$p^{(\ell)}(x) = \frac{1}{r^{(\ell)}} \sum_{(I_+,\theta)\in\mathcal{F}(\mathbb{L}_+)\times\Theta_{I_+}} w^{(I_+,\theta)}(Z) 1_{I_+}(\ell) p^{(\theta)}(x,\ell), \quad (43)$$

$$w^{(I_+,\theta)}(Z) \propto w_+(I_+)[\eta_Z^{(\theta)}]^{I_+}, \quad (44)$$

$$p^{(\theta)}(x,\ell|Z) = \frac{p_+(x,\ell)\psi_Z(x,\ell;\theta)}{\eta_Z^{(\theta)}(\ell)}, \quad (45)$$

$$\eta_Z^{(\theta)}(\ell) = \left\langle p_+(\cdot,\ell), \psi_Z(\cdot,\ell;\theta)\right\rangle, \quad (46)$$

$$\psi_Z(x,\ell;\theta) = \begin{cases} \frac{p_D(x,\ell) g(z_{\theta(\ell)}|x,\ell)}{\kappa(z_{\theta(\ell)})}, & \text{if } \theta(\ell) > 0 \\ q_D(x,\ell), & \text{if } \theta(\ell) = 0 \end{cases} \quad (47)$$

and Θ_{I_+} is the space of mappings $\theta : I_+ \to \{0, 1, ..., |Z|\}$, such that $\theta(i) = \theta(i') > 0$ implies $i = i'$. Here, $p_D(x,\ell)$ is the detection probability at (x,ℓ), $q_D(x,\ell) = 1 - p_D(x,\ell)$ is the probability for missed detection at (x,ℓ), $g(z|x,\ell)$ is the single target likelihood for z given (x,ℓ), and $\kappa(\cdot)$ is the intensity of Poisson clutter.

Proof: Write the LMB prediction density in δ-GLMB form as

$$\boldsymbol{\pi}_+(\mathbf{X}) = \Delta(\mathbf{X}) \sum_{I_+\in\mathcal{F}(\mathbb{L}_+)} w_+(I_+) \delta_{I_+}(\mathcal{L}(\mathbf{X})) p_+^{\mathbf{X}}. \quad (48)$$

Using the result from [35], the multi-target posterior is a δ-GLMB RFS with state space $\mathbb{X}$ and (finite) label space $\mathbb{L}_+$ (and $\Xi = \Theta_{I_+}$) simplifies to

$$\boldsymbol{\pi}(\mathbf{X}|Z) = \Delta(\mathbf{X}) \sum_{(I_+,\theta)\in\mathcal{F}(\mathbb{L}_+)\times\Theta_{I_+}} w^{(I_+,\theta)}(Z) \delta_{I_+}(\mathcal{L}(\mathbf{X})) \left[p^{(\theta)}(\cdot|Z)\right]^{\mathbf{X}}. \quad (49)$$

The (unlabeled) PHD of the full posterior is given by

$$v(x) = \sum_{(I_+,\theta)\in\mathcal{F}(\mathbb{L}_+)\times\Theta_{I_+}} w^{(I_+,\theta)}(Z) \sum_{\ell\in I_+} p^{(\theta)}(x,\ell) \quad (50)$$

which can be interpreted as a weighted sum of the PHD of all individual tracks. Thus, the full posterior can be interpreted as containing track distributions weighted by their existence probability. Summing over all component weights for an individual track label gives the existence probability for an individual track, and similarly summing over all spatial distributions containing an individual track label gives the spatial distribution for an individual track. Comparing the above expression with the unlabeled PHD for the proposed LMB approximation, it can be easily seen that both are identical. Hence the proposed LMB approximation matches the original posterior in terms of its decomposition into individual tracks as well as in the unlabeled first moment or the PHD. ■

Remark 5 *Since the PHD mass gives the mean number of targets (mean cardinality), the mean cardinality of the LMB approximation is identical to that of the full posterior. However, the cardinality distributions themselves differ, since the cardinality distribution of the LMB approximation follows the one for a multi Bernoulli RFS* (6) *while the cardinality distribution of the full posterior is given by* (16)*. Although there are many possible choices of approximate LMB posteriors, our particular approximation admits an intuitive interpretation. Our choice best approximates the original distribution in the sense that it preserves the estimated spatial densities of each track and matches exactly the first moment.*

Remark 6 *Compared to the update step of the multi-Bernoulli filter in [10], the LMB filter update does not involve approximation of the posterior PGFl. Instead, we directly approximate the posterior multi-target density by exact moment matching. Compared to the update step of the δ-GLMB filter which propagates a large sum of multi-target exponentials, the LMB filter update only propagates one component which approximates the former sum.*

C. Discussion

We briefly illustrate via an example the net effect of an LMB approximation. Figure 1 gives a simple example of a δ-GLMB of two tracks with labels ℓ_1 and ℓ_2. The δ-GLMB represents all possible hypothesis with their corresponding weights. Figure 2 illustrates the result of conversion to an

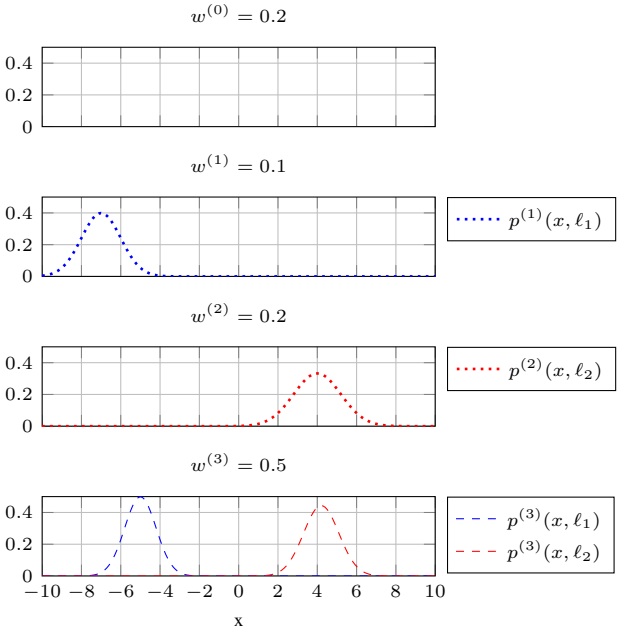

Fig. 1. δ-GLMB representation of all possible hypotheses for two track labels ℓ_1 and ℓ_2. The upper plot represents the empty state space, the plots in the middle correspond to the case that exactly one target exists, and the plot at the bottom represents the existence of both targets.

approximate LMB process. Figure 3 shows the result of converting the approximate LMB process in Figure 2 back to δ-GLMB form. Observe that the coefficients are now different

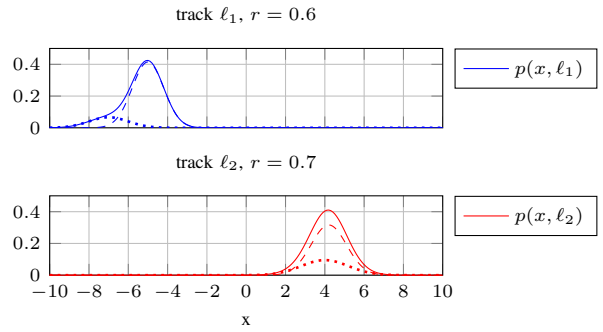

Fig. 2. LMB representation of the δ-GLMB in Figure 1. Each track is represented by a mixture of Gaussians and its existence probability.

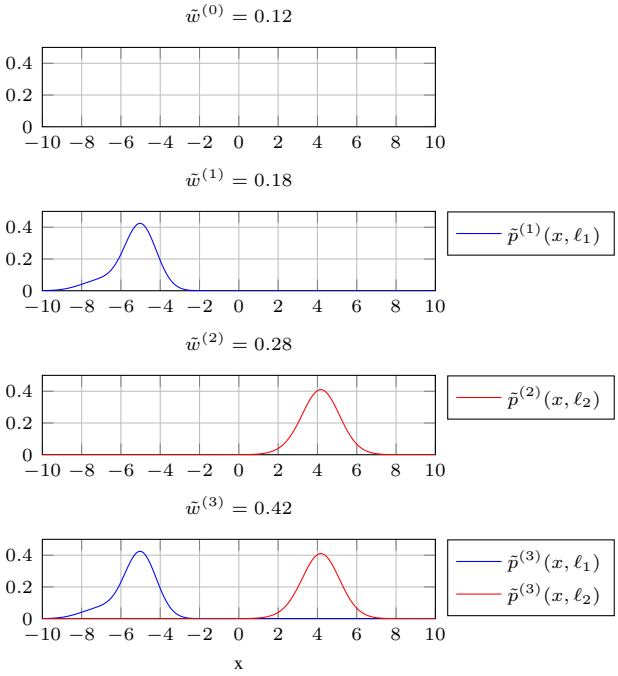

Fig. 3. δ-GLMB representation of all possible hypotheses for two track labels ℓ_1 and ℓ_2 obtained by the conversion of the LMB representation shown in Figure 2.

and hence the cardinality distribution is also different. Further, the spatial distribution for each label ℓ is identical in each of the hypotheses and is given by the weighted sum of the distributions containing the label ℓ in the original process. Both effects are due to the approximation of the original δ-GLMB to LMB form.

IV. Implementation of the Labeled Multi-Bernoulli Filter

The proposed labeled multi-Bernoulli (LMB) filter is an approximation to the full posterior recursion (49). At a conceptual level, the resultant recursion operates as shown in Figure 4 and is explained as follows.

It is possible to apply the LMB update directly. However, the LMB form further allows for the construction of so called 'groups' containing only closely spaced targets and their asso-

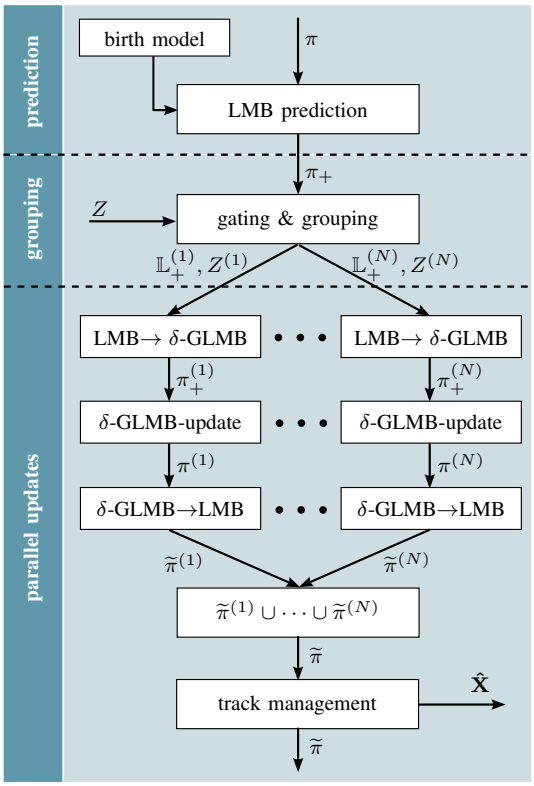

Fig. 4. LMB filter schematic.

ciated measurements, which can then be exploited for significant computational reduction as well as for parallel calculation. Target grouping is based on a standard gating procedure which also partitions the measurement set, resulting in groups of targets and measurements which can be considered statistically independent [41]. Each resultant group is usually much smaller than the entire multi-target state and measurement, and can then be updated in parallel, which is usually much simpler and faster than updating the entire multi-target state as a single group. The update for each group is performed by expanding the LMB prediction into δ-GLMB form, and performing a standard δ-GLMB update resulting in a δ-GLMB posterior. The posterior for each group is then collapsed to a matching LMB approximation after which groups are recombined and the recursion continues. This proposed implementation applies to both Gaussian Mixture and Sequential Monte Carlo based representations for the underlying single target densities. The filter is explained in full detail in the following subsections.

The advantage of the proposed LMB approximation is that it exploits the benefit of the full update, but allows principled tractable approximations which preserves the target tracks interpretation. Compared to the multi-Bernoulli filter, the LMB approximation is more expensive but more accurate with no cardinality bias. For notational convenience, we will not distinguish between a random finite set and its realization throughout this section.

In this section we first describe the basic prediction, target grouping, and parallel update steps of the LMB filter. We then briefly analyse the target grouping procedure which enables parallel execution of the update step and the error incurred in the grouping process. We also propose a basic track estimation procedure and an adaptive birth distribution strategy to ensure robustness to choice of birth models.

A. Prediction

Suppose that the posterior is an LMB distribution with state space $\mathbb{X}$ and (finite) label space $\mathbb{L}$ with parameters

$$\boldsymbol{\pi} = \left\{\left(r^{(\ell)}, p^{(\ell)}\right)\right\}_{\ell\in\mathbb{L}},$$

then the prediction to the time of the next measurement follows an LMB distribution (36) with state space $\mathbb{X}$ and (finite) label space $\mathbb{L}_+ = \mathbb{L} \cup \mathbb{B}$ given by

$$\boldsymbol{\pi}_+ = \left\{\left(r_{+,S}^{(\ell)}, p_{+,S}^{(\ell)}\right)\right\}_{\ell\in\mathbb{L}} \cup \left\{\left(r_B^{(\ell)}, p_B^{(\ell)}\right)\right\}_{\ell\in\mathbb{B}},$$

where the first LMB RFS represents the surviving labeled Bernoulli tracks of the previous time step given by (37), (38) and the second one represents the labeled multi-Bernoulli birth components which are specified a priori. For the surviving tracks, the predicted label is the same as the previous label, and the predicted existence probability and spatial distribution are reweighted by the survival probability and transition density. For the new birth tracks, the labels $\ell \in \mathbb{B}$ are new distinct labels.

B. Grouping and Gating

Given the prediction in LMB form, it is possible to apply the full δ-GLMB update directly, and then collapse back to a matching LMB approximation, as outlined in the previous section. However, it is more efficient to exploit the spatial groupings of targets and measurements which enables updates to be performed for each group in parallel. This results in a drastic reduction in computation times with a slight compromise in accuracy.

The intuition behind this additional approximation is that the predicted LMB parameters can be partitioned into mutually exclusive subsets which have negligible influence on each other during the update. This is the case when each of the partitioned label subsets is well separated spatially as determined by gating of the measurements. Thus, the proposed approximation works best when the birth distribution has reasonably small covariances. Small covariances can be ensured by assuming static birth locations or using a measurement driven birth model as described later in this section.

We now describe a simple and efficient procedure for generating a suitable partition of the predicted LMB parameters and the received measurement set. Figure 5 conceptually illustrates the proposed partitioning scheme. Target groups are formed based on targets which share at least one measurement. Note that a group may have no associated measurement.

Let $\{\mathbb{L}_+^{(1)}, \ldots, \mathbb{L}_+^{(N)}\}$ be a partition of the label set $\mathbb{L}_+ = \mathbb{L} \cup \mathbb{B}$, i.e.

$$\mathbb{L}_+ = \bigcup_{n=1}^{N} \mathbb{L}_+^{(n)}$$

where $\mathbb{L}_+^{(n)} \cap \mathbb{L}_+^{(m)} = \emptyset$ for $n \neq m$. The partitioning of the predicted labels is determined based on the proximity of

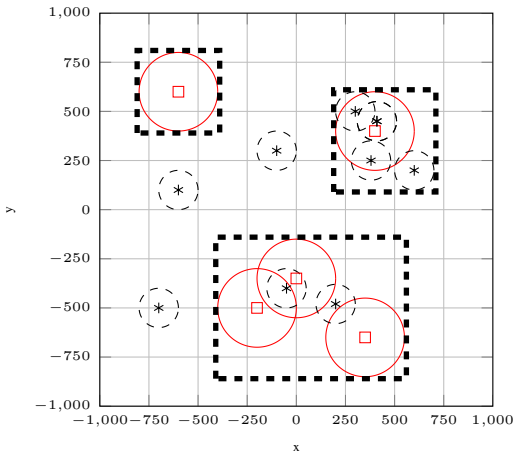

Fig. 5. Partitioning example: dashed rectangles illustrate the three partitions obtained for five tracks (red square) and nine measurements (black star).

received measurements Z which is accordingly partitioned into $\{Z^{(0)}, Z^{(1)}, \ldots, Z^{(N)}\}$, i.e.

$$Z = Z^{(0)} \cup \bigcup_{n=1}^{N} Z^{(n)}$$

and $Z^{(n)} \cap Z^{(m)} = \emptyset$ if $n \neq m$. Note, that $Z^{(0)}$ is the set of measurements which are not assigned to any target labels and $Z^{(n)}$ is the set of measurements associated with targets labels $\mathbb{L}_+^{(n)}$.

A grouping is defined as a set of pairs $\{(\mathbb{L}_+^{(1)}, Z^{(1)}), \ldots, (\mathbb{L}_+^{(N)}, Z^{(N)})\}$ and

$$\mathcal{G}^{(n)} = \left(\mathbb{L}_+^{(n)}, Z^{(n)}\right) \tag{51}$$

is referred to as a group. We describe an iterative procedure for generating a particular grouping which is used for parallel group updates. The procedure is initialized with a tentative grouping of each labeled Bernoulli component ℓ and any associated measurements which fall within its prediction gates:

$$\tilde{\mathcal{G}}^{(\ell)} = \left(\{\ell\}, \{z : d_{\text{MHD}}(\hat{z}^{(\ell)}, z) < \sqrt{\gamma}\}\right). \tag{52}$$

where $d_{\text{MHD}}(\hat{z}^{(\ell)}, z)$ is the Mahalanobis distance (MHD) between the predicted measurement for track ℓ and the received measurement $z \in Z$, and γ is the gating distance threshold calculated using the inverse Chi-squared cumulative distribution corresponding to the desired σ-gate size for gating of measurements from tracks. Then, tenative groups with common measurements

$$Z^{(i)} \cap Z^{(j)} \neq \emptyset \tag{53}$$

are merged where merging of two groups is defined by

$$\tilde{\mathcal{G}}^{(i,j)} = \left(\mathbb{L}_+^{(i)} \cup \mathbb{L}_+^{(j)}, Z^{(i)} \cup Z^{(j)}\right). \tag{54}$$

and the merging is repeated for all tentative groups until there are no common measurements. Finally, a total of N groups $\mathcal{G}^{(1)}, ..., \mathcal{G}^{(N)}$ of tracks and associated measurements is obtained. Consequently, the predicted multi-target density can be rewritten as

$$\boldsymbol{\pi}_+ = \bigcup_{i=1}^{N} \boldsymbol{\pi}_+^{(i)} \tag{55}$$

comprising the target groups

$$\boldsymbol{\pi}_+^{(i)} = \left\{\left(r_+^{(\ell)}, p_+^{(\ell)}\right)\right\}_{\ell \in \mathbb{L}_+^{(i)}} \tag{56}$$

C. Parallel Group Updates

1) Representation of the predicted LMB as δ-GLMB: Since the predicted multi-target density for each group is an LMB of the form (55), recall from section III that it is necessary to express the predicted density in δ-GLMB form, in order to apply the data update. The predicted δ-GLMB (48) for the i-th group of (target) labels and measurements $\mathcal{G}^{(i)} = \left(\mathbb{L}_+^{(i)}, Z^{(i)}\right)$ is given by

$$\boldsymbol{\pi}_+^{(i)}(\widetilde{\mathbf{X}}_+^{(i)}) = \Delta(\widetilde{\mathbf{X}}_+^{(i)}) \sum_{I_+ \in \mathcal{F}(\mathbb{L}_+^{(i)})} w_{+,i}^{(I_+)} \delta_{I_+}(\mathcal{L}(\widetilde{\mathbf{X}}_+^{(i)}))\, [p_+]^{\widetilde{\mathbf{X}}_+^{(i)}} \tag{57}$$

with

$$w_{+,i}^{(I_+)} = \prod_{\ell \in \mathbb{L}_+^{(i)}} \left(1 - r_+^{(\ell)}\right) \prod_{\ell' \in I_+} \frac{1_{\mathbb{L}_+^{(i)}}(\ell') r_+^{(\ell')}}{1 - r_+^{(\ell')}},$$

where $\widetilde{\mathbf{X}}_+^{(i)}$ is the the multi-target state for group i. Thus, a predicted δ-GLMB is determined separately for each predicted component or track in each of the groups. Due to the smaller label space within one group, the number of hypotheses across all groups is significantly smaller than for the number of hypotheses in case of a single big group.

A brute-force way to enumerate the sum is to generate all possible combinations for a set of labels $\mathbb{L}_+^{(i)}$ and cardinalities $n = 0, 1, \ldots, |\mathbb{L}_+^{(i)}|$. The number of combinations for each cardinality is given by the binomial coefficient $C(|\mathbb{L}_+^{(i)}|, n) = |\mathbb{L}_+^{(i)}|!/(n!(|\mathbb{L}_+^{(i)}| - n)!)$ and the number of combinations for a set of track labels is $2^{|\mathbb{L}_+^{(i)}|}$. Consequently explicit evaluation of all combinations is only possible for small $|\mathbb{L}_+^{(i)}|$. For larger $|\mathbb{L}_+^{(i)}|$, the sum can be approximated to its k most significant terms by use of the k-shortest paths algorithm [38], [39] without enumerating all possible terms [35]. Consequently, I_+ only consists of the most significant hypotheses. An alternative solution to enumeration and generation of the k most significant terms is a sampling based approach. For large target or component numbers, this may be a faster solution method. This involves sampling a desired number of unique label sets I_+ by drawing i.i.d. random numbers $a^{(\ell)} \sim \mathcal{U}([0,1])$ and testing for acceptance of each labeled Bernoulli track according to $I_+ = \{\ell | a^{(\ell)} < r_+^{(\ell)} \ \forall\ \ell \in \mathbb{L}_+^{(i)}\}$.

2) Parallel δ-GLMB Group Updates: The δ-GLMB update (49) for each group i is given by

$$\boldsymbol{\pi}^{(i)}(\widetilde{\mathbf{X}}^{(i)}|Z^{(i)}) = \Delta(\widetilde{\mathbf{X}}^{(i)}) \times$$

$$\sum_{(I_+,\theta)\in\mathcal{F}(\mathbb{L}_+^{(i)})\times\Theta_{I_+}} w^{(I_+,\theta)}(Z^{(i)})\delta_{I_+}(\mathcal{L}(\widetilde{\mathbf{X}}^{(i)}))\left[p^{(\theta)}(\cdot|Z^{(i)})\right]^{\widetilde{\mathbf{X}}^{(i)}} \tag{58}$$

where Θ_{I_+} represents the space of mappings $\theta : I_+ \to \{0,1,...,|Z^{(i)}|\}$, such that $\theta(\iota) = \theta(\iota') > 0$ implies $\iota = \iota'$, and

$$w^{(I_+,\theta)}(Z^{(i)}) \propto w_{+,i}^{(I_+)}[\eta_{Z^{(i)}}^{(\theta)}]^{I_+} \tag{59}$$

$$p^{(\theta)}(x,\ell|Z^{(i)}) = \frac{p_{+,i}(x,\ell)\psi_{Z^{(i)}}(x,\ell;\theta)}{\eta_{Z^{(i)}}^{(\theta)}(\ell)}, \tag{60}$$

$$\eta_{Z^{(i)}}^{(\theta)}(\ell) = \langle p_{+,i}(x,\ell), \psi_{Z^{(i)}}(\cdot,\ell;\theta)\rangle, \tag{61}$$

$$\psi_{Z^{(i)}}(x,\ell;\theta) = \begin{cases} \frac{p_D(x,\ell)p_G g(z_{\theta(\ell)}|x,\ell)}{\kappa(z_{\theta(\ell)})}, & \text{if } \theta(\ell) > 0 \\ q_{D,G}(x,\ell), & \text{if } \theta(\ell) = 0 \end{cases}. \tag{62}$$

Observe, that (62) has to incorporate the gating probability p_G [1], since small gates increase the probability of missed detection. Hence, the missed detection probability $q_{D,G}(x,\ell) = 1 - p_D(x,\ell)p_G$.

By representing the complete predicted distribution using several groups, the number of track labels $|\mathbb{L}_+^{(i)}|$ within each group is significantly lower than the total number of possible labels $|\mathbb{L}_+|$. Since the update is combinatorial, the number of components or hypotheses within each group grows exponentially in the number of track labels $|\mathbb{L}_+^{(i)}|$. Thus, for large $|\mathbb{L}_+^{(i)}|$ a truncation of the distribution (58) is required. The truncation can be realized using ranked assignment algorithm [40] which evaluates only the M most significant hypotheses without evaluating all possible solutions [35]. Since the complexity of ranked assignment algorithm is cubic, the computational costs for the evaluation of multiple groups is generally cheaper, compared to the evaluation for a single large group. Moreover, the evaluation for each group can be performed in parallel.

3) Approximation of the updated δ-GLMB as LMB: After performing the group updates across all groups $\mathcal{G}^{(i)}$ for $i = 1,...,N$ the δ-GLMB form for each group is then collapsed back to LMB form

$$\boldsymbol{\pi}^{(i)}(\cdot|Z^{(i)}) \approx \widetilde{\boldsymbol{\pi}}^{(i)}(\cdot|Z^{(i)}) = \left\{\left(r^{(\ell,i)}, p^{(\ell,i)}\right)\right\}_{\ell\in\mathbb{L}_+^{(i)}}$$

where $r^{(\ell,i)}, p^{(\ell,i)}$ are calculated according to (42)-(43) and taking the union of the approximate LMB groups given by $\widetilde{\boldsymbol{\pi}}^{(i)}$ yields the LMB approximation to the multi-target posterior

$$\boldsymbol{\pi}(\cdot|Z) \approx \widetilde{\boldsymbol{\pi}}(\cdot|Z) = \bigcup_{i=1}^{N}\left\{\left(r^{(\ell,i)}, p^{(\ell,i)}\right)\right\}_{\ell\in\mathbb{L}_+^{(i)}}. \tag{63}$$

D. Grouping Error

In the following, we show that the error obtained by partitioning and grouping before the update is negligible, i.e.

$$\boldsymbol{\pi}(\mathbf{X}|Z) \approx \boldsymbol{\pi}^{(1)}\left(\widetilde{\mathbf{X}}^{(1)}\,|\,Z^{(1)}\right)\cdots\boldsymbol{\pi}^{(N)}\left(\widetilde{\mathbf{X}}^{(N)}\,|\,Z^{(N)}\right). \tag{64}$$

Assuming that the target groups are approximately independent and since each target has distinct labels, the predicted distribution can be factored into a single product:

$$\boldsymbol{\pi}_+(\mathbf{X}_+) \approx \boldsymbol{\pi}_+^{(1)}\left(\widetilde{\mathbf{X}}_+^{(1)}\right)\cdots\boldsymbol{\pi}_+^{(N)}\left(\widetilde{\mathbf{X}}_+^{(N)}\right). \tag{65}$$

Additionally, assume that the spatial likelihood for a measurement $z \notin Z^{(i)}$ given any track in $\mathcal{G}^{(i)}$is negligible, and assume that the spatial likelihood for any measurement which is assigned to a track in another group is also negligible:

$$g(z_{\theta(\ell)}|x,\ell) \approx 0 \;\forall\; z_{\theta(\ell)} \notin Z^{(i)},\; (x,\ell) \in \widetilde{\mathbf{X}}_+^{(i)} \tag{66}$$

$$g(z_{\theta(\ell)}|x,\ell) \approx 0 \;\forall\; z_{\theta(\ell)} \in Z^{(i)},\; (x,\ell) \notin \widetilde{\mathbf{X}}_+^{(i)}. \tag{67}$$

The validity of these approximations can be ensured by choosing sufficiently large gating values. Consequently, the multi-target likelihood function can be approximately decomposed as

$$g(Z|\mathbf{X}) \approx g\left(Z^{(1)}\mid\widetilde{\mathbf{X}}^{(1)}\right)\cdots g\left(Z^{(N)}\mid\widetilde{\mathbf{X}}^{(N)}\right). \tag{68}$$

Using (65) and (68), the posterior distribution is given by

$$\begin{aligned}\boldsymbol{\pi}(\mathbf{X}|Z) &\propto \boldsymbol{\pi}_+(\mathbf{X})g(Z|\mathbf{X}) \\ &\approx \boldsymbol{\pi}_+^{(1)}\big(\widetilde{\mathbf{X}}^{(1)}\big)g\big(Z^{(1)}\big|\widetilde{\mathbf{X}}^{(1)}\big)\cdots\boldsymbol{\pi}_+^{(N)}\big(\widetilde{\mathbf{X}}^{(N)}\big)g\big(Z^{(N)}\big|\widetilde{\mathbf{X}}^{(N)}\big) \\ &= \boldsymbol{\pi}^{(1)}\big(\widetilde{\mathbf{X}}^{(1)}\mid Z^{(1)}\big)\cdots\boldsymbol{\pi}^{(N)}\big(\widetilde{\mathbf{X}}^{(N)}\mid Z^{(N)}\big).\end{aligned} \tag{69}$$

Remark 7 *Theoretically in the worst case where all targets are close together, it may not be possible to partition into smaller groups. In this case, the complexity of the update step is the same as that of the δ-GLMB filter. However, our experience with empirical data suggests that in most scenarios we can partition the targets into many groups with small numbers of targets.*

E. Track Pruning and Extraction

Since tracks are represented after update by a labeled multi-Bernoulli, track pruning can be carried out by deleting tracks with an existence probability lower than an application specific threshold, and an obvious track extraction scheme is to pick all tracks with an existence probability higher than an application specific threshold:

$$\hat{\mathbf{X}} = \left\{(\hat{x},\ell) : r^{(\ell)} > \vartheta\right\} \tag{70}$$

where $\hat{x} = \arg_x \max p^{(\ell)}(x)$. On the one hand, choosing a high value for ϑ significantly reduces the number of clutter tracks at the cost of a delayed inclusion of a new-born track. On the other hand, low values for ϑ report new-born tracks immediately at the cost of a higher number of clutter tracks. Choosing a high value for ϑ, additional care has to be taken for missed detections. In case of $p_D \approx 1$, a missed detection considerably reduces the existence probability. Consequently, one missed detection might suppress the output of a previously confirmed track with $r^{(\ell)} \approx 1$.

To mitigate this issue, a hysteresis is used, which activates the output if the maximum existence probability $r_{max}^{(\ell)}$ of a track ℓ has once exceeded an upper threshold ϑ_u and if the current existence probability $r^{(\ell)}$ is higher than a lower threshold ϑ_l:

$$\hat{\mathbf{X}} = \{(\hat{x},\ell) : r_{max}^{(\ell)} > \vartheta_u \wedge r^{(\ell)} > \vartheta_l\} \tag{71}$$

F. Adaptive Birth Distribution

Since the computational efficiency of the LMB filter depends on grouping, it is desirable then to use a birth distribution which leads to tractable groupings. One possibility is to use fixed birth locations with small spatial uncertainties, which requires background knowledge about the environment of the system. Furthermore, it is necessary to detect and initialize all tracks very close to the birth locations, since it is not possible to initialize a track after the target has left the surroundings of the birth location. In scenarios with high clutter rate, fixed birth locations are beneficial, since they minimize the incidence of false tracks.

In [42], an adaptive birth intensity for the SMC-PHD and SMC-CPHD filter is proposed. Similar to the approach in [42], the adaptive birth distribution for the LMB filter should concentrate around measurements which do not originate from existing tracks. However since the adaptive birth distribution for the LMB filter must be an LMB distribution, a direct adaptation of the approach in [42] is not straightforward. Instead we proceed as follows.

The multi-Bernoulli birth distribution $\boldsymbol{\pi}_{B,k+1}$ at the next time step $k+1$ depends on the set of measurements Z_k of the current time step and is given by

$$\boldsymbol{\pi}_{B,k+1} = \left\{ r_{B,k}^{(\ell)}(z), p_{B,k}^{(\ell)}(x|z) \right\}_{\ell=1}^{|Z_k|}. \tag{72}$$

On the one hand, a measurement which is used in all hypotheses may not lead to a new-born target due to the assumption that a measurement is generated by at most one target. On the other hand, measurements which are not included in any hypotheses are likely to create a birth distribution. Consequently, the probability for a measurement to create a birth component depends on the probability $r_{U,k}(z)$ that a measurement z is associated to a track in the hypotheses. For each measurement $z \in Z^{(i)}$ of grouping i, the association probability is given by

$$r_{U,k}(z) = \sum_{(I_+,\theta)\in\mathcal{F}(\mathbb{L}_+^{(i)})\times\Theta_{I_+}} 1_\theta(z) w_k^{(I_+,\theta)}. \tag{73}$$

Here, the inclusion function $1_\theta(z)$ ensures that the sum goes only over updated hypotheses which assigned the measurement z to one of the targets and the weight $w_k^{(I_+,\theta)}$ is given by (59). For measurements which are not associated to any of the partitions, the association probability is

$$r_{U,k}(z) = 0. \tag{74}$$

Thus, the existence probability of the Bernoulli birth distribution at time $k+1$ depending on the measurement z_k is proportional to the probability that z_k is not assigned to any target during the update at time k:

$$r_{B,k+1}(z) = \min\left(r_{B,max}\ ,\ \frac{1-r_{U,k}(z)}{\sum_{\xi\in Z_k} 1-r_{U,k}(\xi)} \cdot \lambda_{B,k+1} \right), \tag{75}$$

where $\lambda_{B,k+1}$ is the expected number of target birth at time $k+1$ and $r_{B,max} \in [0,1]$ is the maximum existence probability of a new born target. Taking the minimum in (75) is necessary, since it is possible that the second term can be larger than one for $\lambda_{B,k+1} > 1$. Generally speaking, a larger value $r_{B,max}$ results in faster track confirmation along with a higher incidence of false tracks. The opposite is also generally true, in other words a lower value $r_{B,max}$ results in slower track confirmation and less false tracks. Thus, the choice of $r_{B,max}$ is application dependent. Since the mean cardinality of a multi-Bernoulli RFS is given by the sum of the existence probabilities, the mean cardinality of new born targets is

$$\sum_{\xi\in Z_k} r_{B,k+1}(\xi) \leq \lambda_{B,k+1},$$

where equality holds if the capping using $r_{B,max}$ is never triggered.

Since (73) can be evaluated for δ-GLMB distributions, the proposed adaptive birth procedure described in this section can also be applied to δ-GLMB filter.

V. Results

The performance of the LMB filter is evaluated in two scenarios. First, SMC implementations of the LMB filter, the δ-GLMB filter, and the multi-Bernoulli filter are compared in a non-linear scenario. Then, the Gaussian mixture (GM) implementation of the LMB filter is used to track a large number of targets in scenarios with high clutter or low detection probability.

A. SMC-Implementation

The non-linear scenario used in [10], and [35], with up to 10 targets involing target birth, death, missed detections and clutter measurements is considered.

The target state $x_k = [\ \tilde{x}_k^T, \omega_k\ ]^T$, where $\tilde{x}_k = [\ p_{x,k}, \dot{p}_{x,k}, p_{y,k}, \dot{p}_{y,k}\ ]^T$ is the position and velocity vector and ω_k is the turn rate. The transition density for the coordinated turn model is given by $f_{k|k-1}(x'|x_k) = \mathcal{N}(x'; m(x_k), Q)$ where $m(x_k) = [\ [F(\omega_k)\tilde{x}_k]^T, \omega_k\ ]^T$, $Q = \mathrm{diag}([\sigma_w^2 GG^T, \sigma_u^2])$ with $\sigma_w = 15$ m/s^2, $\sigma_u = \pi/180$ rad/s,

$$F(\omega) = \begin{bmatrix} 1 & \frac{\sin\omega}{\omega} & 0 & -\frac{1-\cos\omega}{\omega} \\ 0 & \cos\omega & 0 & -\sin\omega \\ 0 & \frac{1-\cos\omega}{\omega} & 1 & \frac{\sin\omega}{\omega} \\ 0 & \sin\omega & 0 & \cos\omega \end{bmatrix}, G = \begin{bmatrix} \frac{1}{2} & 0 \\ 1 & 0 \\ 0 & \frac{1}{2} \\ 0 & 1 \end{bmatrix}.$$

The state independent survival probability of the targets is given by $p_{S,k}(x) = 0.99$. The birth is a multi Bernoulli RFS with density $\pi_B = \{(r_B^{(i)},\ p_B^{(i)})\}_{i=1}^4$, where $r_B^{(1)} = r_B^{(2)} = 0.02$, $r_B^{(3)} = r_B^{(4)} = 0.03$, $p_B^{(i)}(x) = \mathcal{N}(x; m_B^{(i)}, P_B)$, $m_B^{(1)} = [-1500, 0, 250, 0, 0]^T$, $m_B^{(2)} = [-250, 0, 1000, 0, 0]^T$, $m_B^{(3)} = [250, 0, 750, 0, 0]^T$, $m_B^{(4)} = [1000, 0, 1500, 0, 0]^T$, $P_B = \mathrm{diag}([50, 50, 50, 50, 6(\pi/180)]^T)^2$.

The sensor returns bearing and range measurements $z = [\theta, r]^T$ on the region $r \in [0, 2000]$ and $\theta = [-\pi/2, \pi/2]$. The standard deviations of the Gaussian distributed measurement noise are $\sigma_\theta = (\pi/180)$ rad and $\sigma_r = 5$ m. The range dependent detection probability is given by $p_{D,k}(x) = p_{D,max} \exp\left([p_{x,k}, p_{y,k}]^T \Sigma_D^{-1} [p_{x,k}, p_{y,k}]\right)$ with $\Sigma_D = 6000^2 I_2$, where I_2 is the two dimensional identity matrix. The clutter measurements are uniformly distributed on the

observation space and is a Poisson RFS with $\kappa_k(z) = \frac{\lambda_c}{V}\mathcal{U}(z)$, where $\mathcal{U}(\cdot)$ denotes a uniform density on the observation region and $V = \int \mathcal{U}(z)dz$.

The performance of the proposed LMB filter is compared to a δ-GLMB filter [35] and the multi-Bernoulli filter [10]. In the δ-GLMB filter, the number of components is set to 1000 and components with a weight below 10^{-5} are pruned. The LMB and multi-Bernoulli filters use $N = 1000$ particles to represent each track, the δ-GLMB filter uses $N = 1000$ particles for each track in a component. All filters use the multi-Bernoulli birth RFS given above except for the adaptive birth LMB filter which uses the adaptive birth distribution introduced in Section IV-F with $\lambda_{B,k+1|k} = \sum_{i=1}^{i=4} r_B^{(i)} = 0.12$.

Figure 6 shows the trajectories of the objects as well as the tracking result of a single sample run. This result is typical of the filter performance which successfully tracks all of the objects with no track label switching. However occasional track label switching does occur during target crossings as would be expected of any tracking filter due to the high process noise of the motion model. In addition the presence of false alarms and missed detections increases the ambiguity in the data associations during crossings.

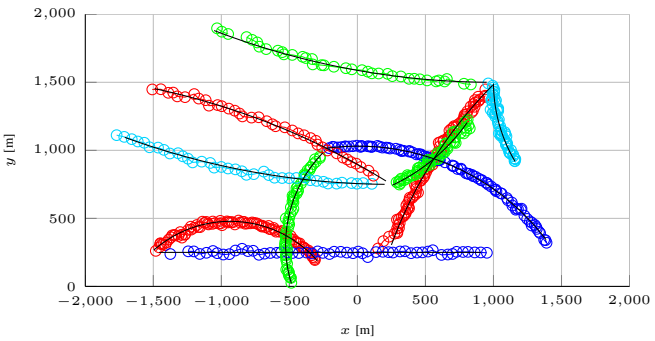

Fig. 6. Ground truth trajectories of the objects (black lines) and estimated trajectories (colored circles).

Figure 7 shows the cardinality estimates of the filters for $\lambda_c = 60$ and $p_{D,max} = 0.98$ averaged over 100 MC runs. Obviously, the multi-Bernoulli filter significantly over-estimates the cardinality, which is expected since its approximations are only valid for reasonably low clutter rates. The estimates for δ-GLMB and LMB filter are nearly identical, only for $k > 80$ the cardinality error of the LMB filter is slightly higher. The adaptive birth LMB filter shows a larger cardinality error after new targets appear, since the adaptive birth distribution needs the first measurement to set up a new-born track while the other filters already use the first measurement to update the static birth distributions. It is also apparent that dropped tracks are very rare with the LMB filter, although as previously noted track switching can occasionally occur.

The OSPA distances in Figure 8 illustrate the performance difference between the δ-GLMB and LMB filters. The LMB filter performs most of the time slightly worse than the δ-GLMB filter due to mainly the approximation within the δ-GLMB→LMB conversion. Especially in ambiguous situations, where m measurements are used to update one track, the δ-GLMB represents the hypotheses for this track by $(m+1)\cdot N$

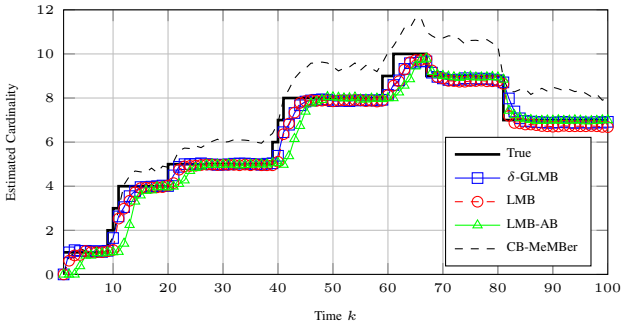

Fig. 7. Cardinality estimates versus time for SMC versions with clutter rate $\lambda_c = 60$ and $p_{D,max} = 0.98$ (averaged over 100 MC runs).

particles (m measurements plus missed detection), while the LMB filter uses only N particles. Further, we observe that the adaptive birth LMB filter has higher OSPA distances during target birth, but significantly outperforms the δ-GLMB filter for $k > 70$ since it is able to re-initialize a lost track at any position in the state space.

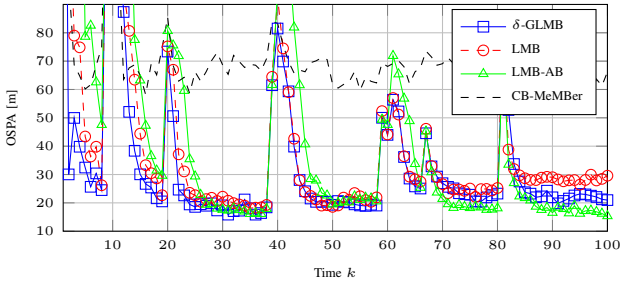

Fig. 8. OSPA distances of order $p = 1$ and cut-off $c = 300$ for SMC versions with $\lambda_c = 60$ and $p_{D,max} = 0.98$ (averaged over 100 MC runs).

Figure 9 shows the OSPA distance for the same simulation with a lower detection probability using $p_{D,max} = 0.75$ and a lower clutter rate $\lambda_c = 10$. Compared to Figure 8 the OSPA distances of all filters are slightly higher, but the differences between the performances of the investigated filters are similar.

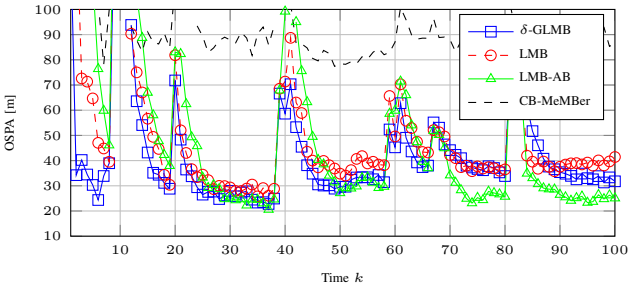

Fig. 9. OSPA distances of order $p = 1$ and cut-off $c = 300$ for SMC versions with $\lambda_c = 10$ and $p_{D,max} = 0.75$ (averaged over 100 MC runs).

B. High Target Number

In order to demonstrate the performance of the LMB filter, a scenario with a high number of targets and clutter

measurements is investigated. A total number of 150 targets appears throughout the scenario, 110 of them at $k = 1$, 20 targets at $k = 30$, and another 20 at $k = 50$. The birth position (x, y) of each target is obtained by drawing a sample of a uniform distribution which covers the surveillance region $[-1000, 1000] \times [-1000, 1000]$. The temporary birth velocities $\bar{v}_x$ and $\bar{v}_y$ are drawn from a uniform distribution on $[1.5, 15]$. In order to mitigate the effect that targets disappear due to leaving the surveillance region, the orientation of the birth velocities is adjusted by $v_x = -sgn(x) \cdot \bar{v}_x$ and $v_y = -sgn(y) \cdot \bar{v}_y$. For each target i, the time of target death is sampled from $t_{death}^{(i)} = t_{birth}^{(i)} + 20 + 10 \cdot \lfloor 15 \cdot \mathcal{U}([0,1]) \rfloor$, where $\mathcal{U}([0,1])$ is the uniform distribution on $[0, 1]$. Additionally, targets are assumed to die if they leave the surveillance region. All targets follow a constant velocity model and the standard deviation of the process noise is $\sigma_v = 5$ m/s. The survival probability is $p_S = 0.99$.

The sensor returns position measurements, and the standard deviation of the measurement noise is $\sigma_x = \sigma_y = 10$ m. In the scenario with high clutter, each target is detected with $p_D = 0.98$ and the sensor generates a mean number of $\lambda_c = 100$ clutter measurements per scan, which are uniformly distributed over the observation space. In the scenario with low detection probability, $p_D = 0.75$ and $\lambda_c = 30$ are used. The partitioning uses a 3-σ gate, i.e. $\gamma = 9$.

Since the targets appear at arbitrary positions, the adaptive birth distribution introduced in Section IV-F is used. The spatial distribution is a Gaussian distribution with mean $[x_m, 0, y_m, 0]$, where x_m and y_m are the measured x and y positions. The covariance matrix of the distribution is initialized with $\Sigma_\Gamma = 100 \cdot I_4$, where I_n is the identity matrix of size n. Within each LMB track, Gaussian components with a Mahalanobis distance $d \leq 1$ are merged. LMB tracks with an existence probability $r \leq 0.001$ are pruned after the update step. In order to reduce the number of false alarms, the parameters $\vartheta_u = 0.75$ and $\vartheta_l = 0.2$ are used for track extraction.

Figure 10 shows the cardinality estimates for both scenarios. The estimated cardinality matches the true number of targets if the number of targets is not changing for about five to ten time steps. In case of low detection probability, the filter reacts faster to new born objects due to the lower clutter rate. As expected, a minor delay in track deletion is observed in the scenario with low detection probability.

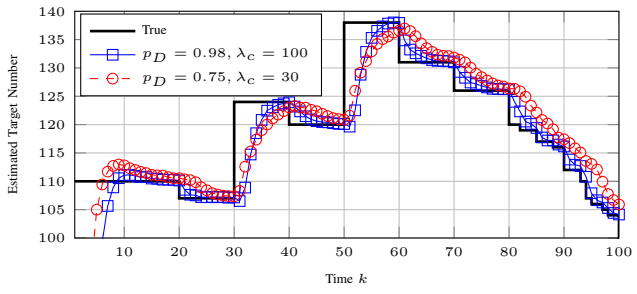

Fig. 10. Cardinality estimate of the GM implementation with adaptive birth for low and high detection probability (100 MC runs).

The OSPA distance for the scenarios is shown in Figure 11. The OSPA distance indicates that, the distance between the true state of a target and the estimate of the LMB filter is on average between 10 and 15 meters, which is only slightly more than one standard deviation of the measurement noise.

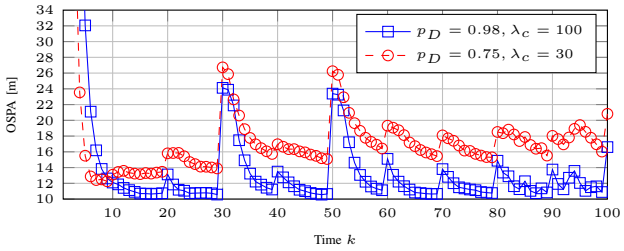

Fig. 11. OSPA distance of order $p = 1$ and cut-off $c = 100$ for the GM implementation with adaptive birth and varying detection probability (100 MC runs).

The results confirm that the LMB filter is able to handle scenarios with a large number of targets and a high clutter rate or a low detection probability. Figure 12 shows the median of the largest update times across all target groups over the 100 MC runs with $\lambda_c = 100$. The 5% and 95% quantiles are also shown to indicate the variability in the computation time. The measured time includes the hypotheses generation, the update of each group with up to 100 association hypotheses, and the conversion back to LMB form. The computation times are obtained based on an unoptimized MATLAB implementation on a Core i7-980 CPU, which leaves significant potential for further speed up if an optimized C++ implementation is used. These results suggest that a real-time capable implementation of the LMB filter is possible by parallelization of the group updates as well as of the single target prediction and the single target updates.

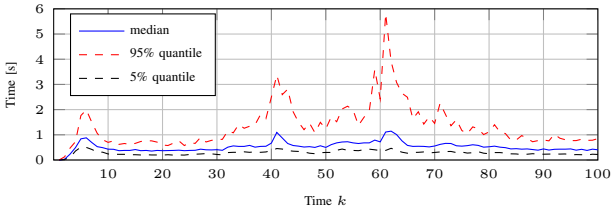

Fig. 12. Median and quantiles of the maximum computation times for the full δ-GLMB update of a group of targets (100 MC runs).

If all targets are very close together and cannot be partitioned into more than one group, the complexity of the LMB filter update is equivalent to the δ-GLMB filter. This is however an extreme case and in most scenarios the grouping procedure significantly reduces computation times.

To ensure real-time performance of the proposed LMB filter, a maximum calculation time per filter update has to be imposed independent of the current number of targets and their states. The filter update computations may be reduced using adaptive partitioning schemes, which reduce the gating distance depending on the current situation in order to ensure a maximum size for each partition.

VI. Conclusion

This paper has proposed a tractable and efficient multi-target tracking algorithm based on a labeled multi-Bernoulli RFS approximation to the multi-target posterior density. The proposed labeled multi-Bernoulli filter combines the relative strengths of the (unlabeled) multi-Bernoulli filter and the δ-generalized labeled multi-Bernoulli filter. It retains the intuitive mathematical structure of the (unlabeled) multi-Bernoulli RFSs but without its inherent weaknesses. At the same time the labeled multi-Bernoulli filter is able to exploit the accuracy of a GLMB representation without exhibiting the exponential growth in the number of posterior components of the δ-GLMB filter. Moreover, the proposed labeled multi-Bernoulli filter formally estimates tracks, is unbiased in posterior cardinality, even for harsh scenarios with low detection probability and high false alarms. The proposed filter is also amenable to paralellization through the grouping of targets. Due to the nature of the labeled multi-Bernoulli approximation which yields discrete tracks, the proposed LMB filter can easily be combined with standard post processing techniques, as suited to the scenario complexity and computational resources.

References

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

[2] Y. Bar-Shalom and T. Fortmann, *Tracking and Data Association*. Academic Press, Inc., 1988.

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

[4] I. Cox and S. Hingorani, "An efficient implementation of Reid's multiple hypothesis tracking algorithm and its evaluation for the purpose of visual tracking," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 18, no. 2, pp. 138–150, 1996.

[5] J. Mullane, B.-N. Vo, M. Adams, and B.-T. Vo, "A random-finite-set approach to Bayesian SLAM," *IEEE Transactions on Robotics*, vol. 27, no. 2, pp. 268–282, 4 2011.

[6] R. Hoseinnezhad, B.-N. Vo, B.-T. Vo, and D. Suter, "Visual tracking of numerous targets via multi-Bernoulli filtering of image data," *Pattern Recognition*, vol. 45, no. 10, pp. 3625–3635, 2012.

[7] M. Munz, M. Mählisch, and K. Dietmayer, "Generic centralized multi sensor data fusion based on probabilistic sensor and environment models for driver assistance systems," *IEEE Intelligent Transportation Systems Magazine*, vol. 2, no. 1, pp. 6–17, 2010.

[8] R. Mahler, "Multitarget Bayes filtering via first-order multitarget moments," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 39, no. 4, pp. 1152–1178, 10 2003.

[9] ——, "PHD filters of higher order in target number," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 43, no. 4, pp. 1523–1543, 10 2007.

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

[11] B.-N. Vo, B.-T. Vo, N.-T. Pham, and D. Suter, "Joint detection and estimation of multiple objects from image observations," *IEEE Transactions on Signal Processing*, vol. 58, no. 10, pp. 5129–5141, 2010.

[12] B.-N. Vo, S. Singh, and A. Doucet, "Sequential Monte Carlo methods for multitarget filtering with random finite sets," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 41, Issue 4, pp. 1224–1245, 2005.

[13] H. Sidenbladh and S.-L. Wirkander, "Tracking random sets of vehicles in terrain," in *Conference on Computer Vision and Pattern Recognition Workshop*, 2003, p. 98.

[14] T. Zajic and R. Mahler, "A particle-systems implementation of the PHD multitarget-tracking filter," in *Signal Processing, Sensor Fusion, and Target Recognition XII, SPIE, Vol. 5096*, Bellingham, WA, 2003, pp. 291–299.

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

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

[17] F. Lian, C. Li, C. Han, and H. Chen, "Convergence analysis for the SMC-MeMBer and SMC-CBMeMBer filters," *Journal of Applied Mathematics*, vol. 2012, Article ID 584140, 2012.

[18] J. Yin, J. Zhang, and J. Zhao, "The Gaussian particle multi-target multi-Bernoulli filter," in *2nd International Conference on Advanced Computer Control (ICACC)*, vol. 4, 2010, pp. 556–560.

[19] J. Yin, J. Zhang, and L. Ni, "The polynomial predictive particle MeMBer filter," in *2nd International Conference on Mechanical and Electronics Engineering (ICMEE)*, vol. 1, 2010, pp. 18–22.

[20] J. J. Yin and J. Q. Zhang, "The nonlinear multi-target multi-Bernoulli filter using polynomial interpolation," in *IEEE 10th International Conference on Signal Processing (ICSP)*, 2010, pp. 2551–2554.

[21] C. Ouyang, H. Ji, and C. Li, "Improved multi-target multi-Bernoulli filter," *IET Radar, Sonar and Navigation*, vol. 6, no. 6, pp. 458–464, 2012.

[22] V. Ravindra, L. Svensson, L. Hammarstrand, and M. Morelande, "A cardinality preserving multitarget multi-Bernoulli RFS tracker," in *Information Fusion (FUSION), 2012 15th International Conference on*, 2012, pp. 832–839.

[23] B.-T. Vo, B.-N. Vo, R. Hoseinnezhad, and R. Mahler, "Robust multi-Bernoulli filtering," *IEEE Journal of Selected Topics in Signal Processing*, vol. 7, no. 3, pp. 399–409, 2013.

[24] J. Wei and X. Zhang, "Mobile multi-target tracking in two-tier hierarchical wireless sensor networks," in *Proceedings of the IEEE Military Communications Conference*, 2009, pp. 1–6.

[25] ——, "Dynamic node collaboration for mobile multi-target tracking in two-tier wireless camera sensor networks," in *Proceedings of the IEEE Military Communications Conference*, 2009, pp. 1–7.

[26] ——, "Sensor self-organization for mobile multi-target tracking in decentralized wireless sensor networks," in *Proceedings of the IEEE Wireless Communications and Networking Conference*, 2010, pp. 1–6.

[27] ——, "Efficient node collaboration for mobile multi-target tracking using two-tier wireless camera sensor networks," in *Proceedings of the IEEE International Conference on Communications*, 2010, pp. 1–5.

[28] X. Zhang, "Adaptive control and reconfiguration of mobile wireless sensor networks for dynamic multi-target tracking," *IEEE Transactions on Automatic Control*, vol. 56, no. 10, pp. 2429–2444, 2011.

[29] J. Lee and K. Yao, "Initialization of multi-Bernoulli random-finite-sets over a sensor tree," in *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2012.

[30] R. Hoseinnezhad, B.-N. Vo, B.-T. Vo, and D. Suter, "Bayesian integration of audio and visual information for multi-target tracking using a CB-MeMBer filter," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2011, pp. 2300–2303.

[31] R. Hoseinnezhad, B.-N. Vo, D. Suter, and B.-T. Vo, "Multi-object filtering from image sequence without detection," in *Proceedings of the IEEE International Conference on Acoustics Speech and Signal Processing*, Dallas, TX, 2010, pp. 1154–1157.

[32] R. Hoseinnezhad, B.-N. Vo, and B.-T. Vo, "Visual tracking in background subtracted image sequences via multi-Bernoulli filtering," *IEEE Transactions on Signal Processing*, vol. 61, no. 2, pp. 392–397, 2013.

[33] J. Williams, "Hybrid Poisson and multi-Bernoulli filters," in *Information Fusion (FUSION), 2012 15th International Conference on*, 2012, pp. 1103–1110.

[34] B.-T. Vo and B.-N. Vo, "A random finite set conjugate prior and application to multi-target tracking," in *Proceedings of the 7th International Conference on Intelligent Sensors, Sensor Networks and Information Processing*, 12 2011, pp. 431–436.

[35] ——, "Labeled random finite sets and multi-object conjugate priors," *IEEE Transactions on Signal Processing*, vol. 61, no. 13, pp. 3460–3475, 2013.

[36] D. Fränken, M. Schmidt, and M. Ulmke, ""spooky action at a distance" in the cardinalized probability hypothesis density filter," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 45, no. 4, pp. 1657–1664, 2009.

[37] B.-N. Vo, B.-T. Vo, and D. Phung, "Labeled random finite sets and the bayes multi-target tracking filter," *pre-print*, 2013. [Online]. Available: http://arxiv.org/abs/1312.2372

[38] D. Eppstein, "Finding the k shortest paths," *SIAM Journal on Computing*, vol. 28, no. 2, pp. 652–673, 1998.

[39] J. Y. Yen, "Finding the k shortest loopless paths in a network," *Management Science*, vol. 17, no. 11, pp. 712–716, 07 1971.

[40] K. Murty, "An algorithm for ranking all the assignments in order of increasing cost," *Operations Research*, vol. 16, pp. 682–687, 1968.

[41] J. Dezert and Y. Bar-Shalom, "Joint probabilistic data association for autonomous navigation," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 29, no. 4, pp. 1275–1286, 10 1993.
[42] B. Ristic, D. Clark, B.-N. Vo, and B.-T. Vo, "Adaptive target birth intensity for PHD and CPHD filters," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 48, no. 2, pp. 1656–1668, 2012.

Stephan Reuter was born 1981 in Crailsheim, Germany. He received a Diploma degree (equivalent to M.Sc. degree) in electrical engineering from Ulm University, Germany, in 2008. Afterwards, he joined the Institute of Measurement, Control and Microtechnology in the school of Engineering and Computer Science at Ulm University as a research assistant. His main research topics are sensor data fusion, multi-object tracking, environment perception for cognitive technical systems, and sensor data processing.

Ba-Tuong Vo was born in Perth, Australia, in 1982. He received the B.Sc. degree in applied mathematics and B.E. degree in electrical and electronic engineering (with first-class honors) in 2004 and the Ph.D. degree in engineering (with Distinction) in 2008, all from the University of Western Australia. He is currently an Associate Professor in the Department of Electrical and Computer Engineering at Curtin University and a recipient of an Australian Research Council Fellowship. His primary research interests are in point process theory, filtering and estimation, and multiple object filtering. Dr. Vo is a recipient of the 2010 Australian Museum DSTO Eureka Prize for "Outstanding Science in Support of Defence or National Security".

Ba-Ngu Vo received his Bachelor degrees jointly in science and electrical engineering with first class honours from the University of Western Australia in 1994, and PhD from Curtin University in 1997.

He had held various research positions before joining the Department of Electrical and Computer Engineering at Curtin University in 2012 where he holds the position of Professor and Chair of Signal Processing.

Prof. Vo is a recipient of the Australian Research Council's inaugural Future Fellowship and the 2010 Australian Museum DSTO Eureka Prize for "Outstanding Science in Support of Defence or National Security". His research interests are Signal Processing, Systems Theory and Stochastic Geometry with emphasis on target tracking, robotics and computer vision.

Klaus Dietmayer (M'05) was born in Celle, Germany in 1962. He received his Diploma degree (equivalent to M.Sc. degree) in 1989 electrical engineering from the Technical University of Braunschweig (Germany), and the Dr.-Ing. degree (equivalent to PhD) in 1994 from the University of Armed Forces in Hamburg (Germany). In 1994 he joined the Philips Semiconductors Systems Laboratory in Hamburg, Germany as a research engineer. Since 1996 he became a manager in the field of networks and sensors for automotive applications.

In 2000 he was appointed to a professorship at Ulm University in the field of measurement and control. Currently he is Full Professor and Director of the Institute of Measurement, Control and Microtechnology in the school of Engineering and Computer Science at Ulm University. His research interests include information fusion, multi-object tracking, environment perception for advanced automotive driver assistance, and E-Mobility.

Klaus Dietmayer is member of the IEEE and the German society of engineers VDI / VDE.