Labeling Uncertainty in Multitarget Tracking

EDSON HIROSHI AOKI
PRANAB K. MANDAL
University of Twente
Enschede, The Netherlands

LENNART SVENSSON
Chalmers University
Gothenburg, Sweden

YVO BOERS
Thales Nederland B.V.
Hengelo, The Netherlands

ARUNABHA BAGCHI
University of Twente
Enschede, The Netherlands

In multitarget tracking, the problem of track labeling (assigning labels to tracks) is an ongoing research topic. The existing literature, however, lacks an appropriate measure of uncertainty related to the assigned labels that has a sound mathematical basis as well as clear practical meaning to the user. This is especially important in a situation where well separated targets move in close proximity with each other and thereafter separate again; in such a situation, it is well known that there will be confusion on target identities, also known as "mixed labeling." In this paper, we specify comprehensively the necessary assumptions for a Bayesian formulation of the multitarget tracking and labeling (MTTL) problem to be meaningful. We provide a mathematical characterization of the labeling uncertainties with clear physical interpretation. We also propose a novel labeling procedure that can be used in combination with any existing (unlabeled) MTT algorithm to obtain a Bayesian solution to the MTTL problem. One advantage of the resulting solution is that it readily provides the labeling uncertainty measures. Using the mixed labeling phenomenon in the presence of two targets as our test bed, we show with simulation results that an *unlabeled* multitarget sequential Monte Carlo (M-SMC) algorithm that employs sequential importance resampling (SIR) augmented with our labeling procedure performs much better than its "naive" extension, the *labeled* SIR M-SMC filter.

Manuscript received February 21, 2015; revised November 7, 2015, December 14, 2015; released for publication December 17, 2015.

DOI. No. 10.1109/TAES.2016.140613.

Preliminary versions of the contents of this paper have appeared in the conference papers "A Bayesian look at the optimal track labelling problem" presented at the 9th IET Data Fusion & Target Tracking Conference (DF&TT'12) and "The Rao-Blackwellized marginal M-SMC filter for Bayesian multi-target tracking and labelling" presented at the 15th International Conference of Information Fusion (FUSION 2012). This paper contains essential corrections and updates to the contents of the aforementioned papers along with new numerical results.

The research leading to these results has received funding from the EU's Seventh Framework Programme under grant agreement 238710. The research has been carried out in the MC IMPULSE project: https://mcimpulse.isy.liu.se. The research has also been supported by the Netherlands Organisation for Scientific Research (NWO) under the Casimir program, contract 018.003.004.

Refereeing of this contribution was handled by M. Idan.

Authors' addresses: E. H. Aoki, TÜV SÜD PSB Pte. Ltd., 1 Science Park Drive, Singapore 118221; P. K. Mandal, Department of Applied Mathematics, TW/EWI, University of Twente, PO Box 217, 7500 AE, Enschede, The Netherlands; L. Svensson, Department of Signals and Systems, Chalmers University of Technology, 412 96 Gothenburg, Sweden; Y. Boers, J. J. van Deinselaan 440, 7535BT Enschede, The Netherlands; A. Bagchi, Department of Applied Mathematics, TW/EWI, University of Twente, PO Box 217, 7500 AE, Enschede, The Netherlands. Corresponding author is Edson H. Aoki, E-mail: (e.h.aoki@gmail.com).

I. INTRODUCTION

The track labeling problem is perhaps just as old as the multitarget tracking (MTT) problem itself. In the display of a radar operator, it is often necessary not only to display the estimated locations (what we refer to as *the tracks*) of the multiple objects but also to attribute a unique label to each track. Ideally, a label should consistently be associated with the same real-world object (*target*), enhancing the situational awareness of, for example the radar operator.

In practice, the feasibility of maintaining this label-to-target consistency depends on the observability conditions. One situation where this consistency is frequently lost is when the well-separated targets move in close proximity to each other. In this case, even after the separation, the measurements and initial information may not allow us to precisely determine which target is which (as illustrated in Fig. 1 with two targets). Therefore, if required to make a hard decision to assign labels to the estimated locations, the tracking system will frequently make wrong choices.

This situation, where the available information allow for more than one labeling possibility, is referred to as "mixed labeling" by Boers et al. [1]. Being well informed about the labeling uncertainty is of utmost relevance to an end user when, for example, a decision involving a target with a particular label is acceptable only if we have high confidence in the label. It is therefore interesting and of great importance to characterize and report these uncertainties.

The idea of obtaining target identities using a probabilistic approach has been known for some time and has received its due attention in the literature (e.g., [2–7]). These works consider situations ranging from a fixed number of targets to a time-varying number of targets due to target birth and death. While these works typically suggest methods for extracting labeled tracks from a multitarget density, they do not attempt to quantify the amount of uncertainty in the assigned labels.

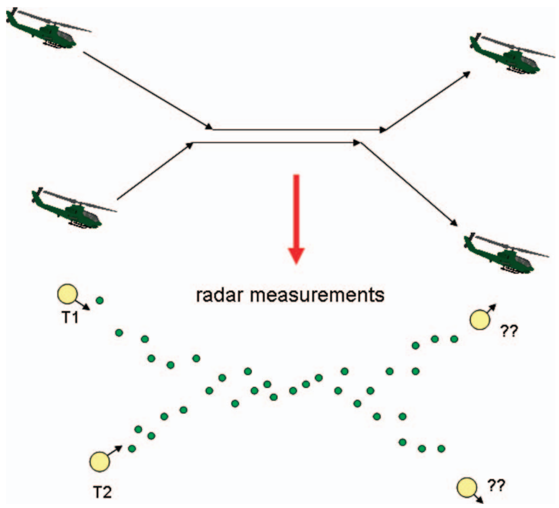

Fig. 1. Situation where assignment of labels to tracks is ambiguous.

Other recent works [8–12] have proposed quantities to be associated with the labeling uncertainty in a multitarget tracking and labeling (MTTL) problem. However, the definition of these quantities relies on abstract concepts, such as decomposition of densities into weighted sums and permutations of the state vector, that tend to make them hard to interpret. They are also based on restrictive assumptions, such as linear-Gaussian target dynamics or being defined for only two targets or assuming the number of targets to be known and time invariant.

Furthermore, even if one defines a suitable labeling uncertainty measure, the uncertainty might be lost/underestimated when it is calculated based on particle filters or multiple hypotheses due to the degeneracy phenomenon present in these algorithms. This weakness has been noticed by [1, 8–12] and is further explained in section II-C. Practical implementations of the labeling uncertainties should therefore take this into consideration.

Throughout this paper, we will use as illustrative example the situation of mixed labeling as depicted in Fig. 1. A natural measure to characterize the labeling uncertainties in Fig. 1 could be the probability that the assignment of labels to the tracks is incorrect, in other words, that a track swap has occurred. It is, however, not completely clear what the exact meaning of "probability of incorrect labeling" is. After all, the tracks are only (point) estimates of the true target states, and they almost never coincide. If the tracks themselves are not "correct," what shall we understand by "correct labeling"?

In finding the answers, we consider the Bayesian formulation of the MTTL problem, based on the concept of labeled random finite sets (RFS) presented in [6]. We propose labeling-related statistics with clear meaning in terms of quantities similar to conditional probability. Our starting point is that an end user, for example, the radar operator, will prefer point estimates of the target locations rather than the whole posterior probability distribution of the locations. The user would like to assign labels to these point estimates of the locations. Furthermore, based on an intuition that resembles [8–12], we explicitly make the proposed labeling uncertainties part of our density approximation (for the labeled tracks) so that they are not lost during the filter recursions.

We should note here that in the target tracking literature, labels are used mostly as means to extract trajectories of the targets. Recently, the authors in [13] moved away from this approach of artificial label and considered the problem of estimating the trajectories directly. In our work, although we consider the problem of labeled state estimation, it is not our primary goal to estimate trajectories of the targets. Our main goal is to estimate, at each time, the labeled states so that we can associate a current target to its "location at birth." In the context of Fig. 1, for example, if a current target is assigned the label "T1," then we can infer that it originated in the upper-left corner. One can also think of situations with air traffic control (ATC) call signs that are assigned to aircraft by the ATC in order to uniquely identify them. Usually, cooperative aircraft observed by a secondary radar periodically inform their own assigned call sign to the ATC. However, if, for example, an aircraft's transponder suddenly stops functioning, so does the information about its call sign. In this case, the ATC attempts to associate the previously assigned call signs to targets observed using the primary radar. In this way, the operator can relate a current target to the previously existing targets in terms of their last known position.

The contributions of this paper are the following:

- We complement the formulation of the MTTL problem in [6] by stating additional necessary assumptions for the problem we consider to be meaningful.
- We provide a *mathematical description of the labeling error* with clear physical interpretation based on the labeled multitarget posterior density.
- We present a labeling procedure using the proposed labeling uncertainty measures, which can be used to augment existing MTT algorithms to obtain a complete solution to the joint MTTL problem. This procedure avoids the degeneracy in labels that typically arises in MTTL filtering algorithms based on particles or hypotheses.

The organization of this paper is as follows. In section II, we review the Bayesian formulation of the MTTL problem given in [6] and complement it to formulate the problem we consider. The other contributions of this paper are presented in sections III and IV, describing the proposed measure of labeling uncertainty and a new method to solve the MTTL problem, respectively. Section V contains the simulation results for labeled tracking of two closely spaced targets. Some conclusions and recommendations are given in section VI.

A. Notation Conventions

An uppercase letter (like X) denotes a vector-valued random variable, and its lowercase counterpart (x) denotes a particular realization. An uppercase boldface letter (like $\mathbf{X}$) denotes a finite set-valued random variable, and its lowercase counterpart ($\mathbf{x}$) denotes the corresponding realization. Vector entries and set elements have superscripts containing their indexes, and vectors are always row vectors (written horizontally), such as $x = [x^{(1)}, x^{(2)}]$, $\mathbf{x} = \{x^{(1)}, x^{(2)}\}$.

II. THE BAYESIAN MTTL PROBLEM

In this section, we present the mathematical formulation of the Bayesian MTTL problem that we consider. The formulation (section II-A) follows the one given in [6] but with a couple of extra assumptions. These assumptions are elaborated further in section II-B. In our opinion, these assumptions, though quite intuitive, have not been discussed in detail in the existing literature. We also present an important property of the considered Bayesian MTTL problem, namely, the one-sided decoupling property, which will play a central role in the derivation of our proposed algorithm in section IV-C. Finally, in section II-C, we discuss why mixed labeling, such as in Fig. 1, creates an extra problem.

In what follows, we assume that the reader has basic familiarity with the concepts of finite set statistics (FISST), such as random finite sets and the corresponding density functions (see, e.g., [14]).

A. Mathematical Formulation

Let us assume that the single-target state vector (composed of entries such as position, velocity, etc., which we will henceforth refer to simply as *location*) assumes values in $\mathbb{R}^n$ and that a label to be assigned to a location may assume values in a discrete set Π. We then define the *labeled multitarget state* at time k as the random finite set

$$\mathbf{X}_k = \left\{X_k^{(1)}, \cdots, X_k^{(T_k)}\right\},$$

where $X_k^{(i)} = [S_k^{(i)}, L_k^{(i)}]$ with locations $S_k^{(i)} \in \mathbb{R}^n$ and labels $L_k^{(i)} \in \Pi$. Clearly, no two single-target states can have the same label if the labels are to be useful as target identifiers. As a result, an RFS density function associated with $\mathbf{X}_k$ (referred to as a *labeled RFS density*) must satisfy

$$f\left(\left\{\left[s_k^{(1)}, l_k^{(1)}\right], \cdots, \left[s_k^{(t_k)}, l_k^{(t_k)}\right]\right\}\right) = 0,$$
$$\text{if } \exists\, i, j \in \{1, \cdots, t_k\} \text{ s.t. } i \neq j,\ l_k^{(i)} = l_k^{(j)}. \quad (1)$$

Examples of closed-form RFS densities that satisfy (1) are the labeled Poisson RFS density, the labeled multi-Bernoulli RFS density, and the generalized labeled multi-Bernoulli RFS density, all described in [6]. Let us denote the corresponding observation as $\mathbf{Z}_k$ (also an RFS) and the sequence of all observations available until and including time k by Z^k.

As typical in the literature, in this paper we assume the labeled state and observation processes $(\mathbf{X}_k, \mathbf{Z}_k)$ to be a first-order partially observed Markov process with

$$f\left(\mathbf{x}_k \left|\mathbf{x}_{k-1}, Z^{k-1}\right.\right) = f(\mathbf{x}_k|\,\mathbf{x}_{k-1}) \quad (2)$$

and

$$f\left(\mathbf{z}_k \left|\mathbf{x}_k, Z^{k-1}\right.\right) = f(\mathbf{z}_k|\,\mathbf{x}_k). \quad (3)$$

The exact formulas for the *multitarget state transition function* $f(\mathbf{x}_k|\mathbf{x}_{k-1})$, the *multitarget likelihood densities* $f(\mathbf{z}_k|\mathbf{x}_k)$, and the *initial multitarget prior* $f(\mathbf{x}_0)$ depend on the assumptions of the scenario (see, e.g., [6]).

The *multitarget posterior* $f(\mathbf{x}_k|Z^k)$ can be calculated recursively as (see, e.g., [6])

$$f\left(\mathbf{x}_k\middle|Z^k\right) = \frac{f(\mathbf{z}_k|\mathbf{x}_k)\, f\left(\mathbf{x}_k\left|Z^{k-1}\right.\right)}{f\left(\mathbf{z}_k\left|Z^{k-1}\right.\right)}, \quad (4)$$

where

$$f\left(\mathbf{x}_k\left|Z^{k-1}\right.\right) = \int f(\mathbf{x}_k|\mathbf{x}_{k-1}) f\left(\mathbf{x}_{k-1}\left|Z^{k-1}\right.\right) \delta\mathbf{x}_{k-1} \quad (5)$$

and

$$f(\mathbf{z}_k|Z^{k-1}) = \int f(\mathbf{z}_k|\mathbf{x}_k) f(\mathbf{x}_k|\,Z^{k-1})\delta\mathbf{x}_k, \quad (6)$$

with the integrals being *set integrals* (see, e.g., [14, Section 9.3.2]).

We assume further that *the new targets are assigned unambiguous labels at the time of their appearances* (as explained further in section II-B1) and the following assumption.

Assumption (L) *The labels affect neither the kinematic states of the target nor the generated observation corresponding to those kinematic states. In particular, we assume that*

$$f\left(\mathbf{s}_k|\,\mathbf{x}_{k-1}\right) = f\left(\mathbf{s}_k|\,\mathbf{s}_{k-1}\right) \quad (7)$$

and

$$f\left(\mathbf{z}_k\,|\,\mathbf{x}_k\right) = f\left(\mathbf{z}_k\,|\,\mathbf{s}_k\right). \quad (8)$$

Note that conditions (7) and (8) are not explicitly assumed in [6]. Neither are they automatically satisfied by all the models considered there. For example, (8) will be violated if in the observation model the detection probability $p_D([s, l])$ (see [6, section IV-C]) depends on the label l. Also, in the multi-Bernoulli RFS model, if the survival probability $p_S([s, l])$ (see [6, section IV-D]) depends on the label l, then (7) will not be satisfied. Conditions (7) and (8) are, however, consistent with the definition of label used in this paper (see section II-B2).

Note further that condition (8) allows us to use in an MTTL problem any relevant multitarget RFS measurement model that can be used in an (unlabeled) MTT problem. Examples of such multitarget RFS measurement models are the point measurement model described in [14, chapter 12] and the track-before-detect measurement model described in [15]. Closed-form expressions for the multitarget prior and state transition densities (for the unlabeled MTT problem) can be found in [14, chapters 13, 14].

B. Assumptions and Properties of the MTTL Problem

1) *Nonambiguity of Initial Labels*: The complete mathematical description of the assumption is given later

in Remark 3.2, once other necessary quantities are introduced. Here we provide an intuitive explanation.

As in [6], we too consider a label to be a placeholder for a target's identity, which cannot be observed and which is to be estimated along with the target locations in a multitarget tracking scenario.

In section I, we have equated "labels" with "ATC call signs" when they cannot be observed (e.g., when transponders stopped working). Clearly, estimating ATC call signs for aircraft with nonfunctioning transponders makes sense only if the aircraft transponders were functioning until some point; that is, the aircraft were attributed unique labels in the past. In some sense, the moment of transponder failure can be considered as the birth of the target because from that point on, the call sign became unobservable.

We view the labeling problem to be similar to this situation, where one associates, to a current (estimated) target, one of the previously assigned labels (generally assigned when a target is detected for the first time) or a new one to indicate the birth of a new target. Furthermore, keeping similarity with the nonfunctioning transponder, we assume that the model assigns a nonambiguous label to the target at its birth. This is possible if we assume that the support for the probability distribution of location for a newborn target does not overlap with that for other (existing or other newborn) targets. With the aircraft and ATC, it is true because of the strict regulations.

On the other hand, suppose that two targets appear at the same time and that no matter where they appear, the model assigns to each target, say, label A with probability 0.5 and label B with probability 0.5 (i.e., enter into so-called *total mixed labeling*). Then this mixed labeling will persist at all later times [16, section IV-C], rendering futile the attempt to assign labels to location estimates.

Unfortunately, the labeled RFS model described in [6] will always produce total mixed labeling if two (or more) new targets appear at the same time instant. When one needs to deal with more than one target appearing simultaneously, other labeling schemes can be envisaged to circumvent this problem, for example, by partitioning the surveillance space into small grids and attributing labels according to the time and grid the target appears in. However, we do not go into this aspect here. We assume henceforth that there is no ambiguity regarding the labels of appearing targets, which means that total mixed labeling is at least avoided at the time of appearance.

2) *Interpretation of Labels*: When labels are used only to estimate trajectories, their values at one particular time instant do not carry any useful information. They exist *solely* to connect target states at different time points. This, however, does not hold in our case. In conformity with the nonambiguous initial labels (see section II-B1) and as mentioned in section I, we note that a label carries some information about the location, at birth, of the target it is attached to. In the context of Fig. 1, the label "T1" refers to the target that appeared in the upper-left corner,

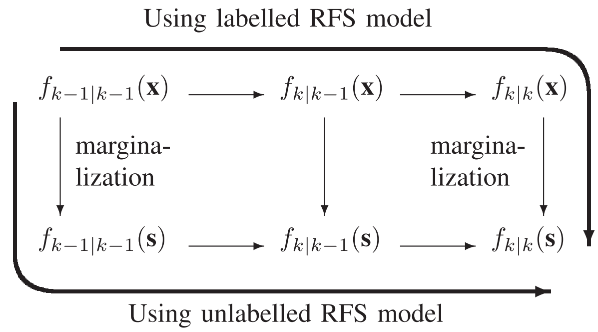

Fig. 2. Different ways of obtaining posterior $f_{k|k}(\mathbf{s})$ from prior $f_{k-1|k-1}(\mathbf{x})$.

and in the context of ATC call signs, a label refers to its last known position.

Despite this, we consider the label to be an artificial addition to the physical/kinematic model of the targets in the sense that labels do not carry any information about the target that may influence the measurements it generate or the transition mechanism of the *locations*.

Looking at it from the reverse side, the measurements do not provide any direct information about the labels. Information about labels are obtained only via the estimated location of the targets (on the basis of the measurements) and combining this with the knowledge of transition mechanism to infer the location at birth.

Note that if labels are indeed artificial introduction to the physical model of the unlabeled states (i.e., the unlabeled RFS model), one would expect that given the measurements, the results obtained using *labeled* and *unlabeled* RFS models will be consistent with each other. In particular, one should expect that the filtered distribution of the unlabeled states obtained using the *unlabeled* RFS model (and the observations) should coincide with the distribution of the unlabeled states, calculated by marginalizing the filtered distribution obtained using the *labeled* RFS model (see Fig. 2).

As shown in the appendix, Assumption (L) of section II-A and, in particular, conditions (7) and (8) indeed imply this consistency, corroborating the viewpoint that the labels are artificial introduction to the physical model.

3) *One-Sided Decoupling Between the Tracking and Labeling Subproblems*: In view of the discussion in section II-B2 and in particular from Fig. 2, it is clear that for the considered Bayesian MTTL problem, one can solve the (sub)problem of tracking, that is, estimating the corresponding set of unlabeled target states $\mathbf{S}_k = \{S_k^{(1)}, \cdots, S_k^{(t)}\}$, from a sequence of observations $Z^k = (\mathbf{z}_1, \ldots, \mathbf{z}_k)$, completely disregarding the labeling (sub)problem. The tracking problem will involve the following recursions, similar to (4)–(6) but with the unlabeled states:

$$f(\mathbf{s}_k | Z^k) = \frac{f(\mathbf{z}_k|\mathbf{s}_k) f(\mathbf{s}_k | Z^{k-1})}{f(\mathbf{z}_k | Z^{k-1})}, \tag{9}$$

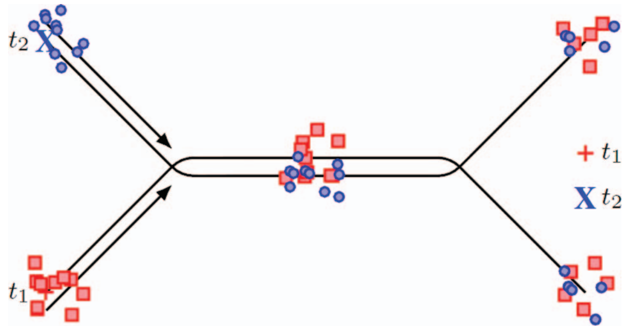

Fig. 3. Particle representation of the multitarget distribution in a situation where mixed labeling occurs (source: [10]). The squares and circles mark the possible locations of each target in terms of particles. "+" and "X" denote the MMSE estimates.

where

$$f(\mathbf{s}_k|\, Z^{k-1}) = \int f(\mathbf{s}_k|\, \mathbf{s}_{k-1}) f\left(\mathbf{s}_{k-1}|\, Z^{k-1}\right) \delta\mathbf{s}_{k-1} \quad (10)$$

and

$$f(\mathbf{z}_k|\, Z^{k-1}) = \int f(\mathbf{z}_k|\, \mathbf{s}_k) f(\mathbf{s}_k|\, Z^{k-1}) \delta\mathbf{s}_k. \quad (11)$$

See, for example, [14, chapter 14].

Although the Bayesian tracking recursion (9) does not involve any probability distribution of labels or labeled multitarget states, that is, does not depend on Bayesian labeling, the (sub)problem of Bayesian labeling will depend on the solution of Bayesian tracking due to the interpretation provided in section II-B2. We refer to this property as one-sided decoupling between the two subproblems of (unlabeled) tracking and labeling.

C. Why Mixed Labeling Makes Bayesian MTTL More Challenging

It is tempting to believe that approximating recursion (4), by itself, gives a practical solution to the complete MTTL problem. However, the way the recursion is implemented plays a major role in providing the correct statistics about the labeling uncertainties. In most cases, a computationally feasible approximation of the $f(\mathbf{x}_k|Z^k)$ is used (see [14, chapter 15]). This being an approximation may, in turn, deter us from obtaining accurate estimates of the statistics of interest. Consider again the situation depicted in Fig. 1. If one implements the Bayes recursion using a particle filter (PF), then the mixed labeling manifests itself by particle clouds corresponding to each target getting mixed, as shown in Fig. 3. In this case, due to the inherent resampling mechanism in the PF method, the actual labeling error tends to get underestimated. We explain this further below.

Suppose the multitarget sequential Monte Carlo (M-SMC) filter, presented in [17] and [14, chapter 15], is used to obtain the posterior. These are multitarget versions of the well-known sequential importance resampling particle filter (SIR PF) proposed by Gordon et al. [18] and Kitagawa [19].

As a SIR PF, the M-SMC filter suffers from the well-known degeneracy phenomenon described in [20–22]. For any given time j, the resampling mechanism will cause the hypotheses on the multitarget trajectory $(\mathbf{X}_0, \ldots, \mathbf{X}_j)$ to eventually (i.e., at some time step $k > j$) collapse into a single hypothesis $(\mathbf{x}_0^*, \cdots, \mathbf{x}_j^*)$, leading the particle approximation of the posterior $f(\mathbf{x}_k|Z^k)$ to be biased toward $f(\mathbf{x}_k|\mathbf{x}_0^*, \cdots, \mathbf{x}_j^*, Z^k)$. This degeneracy will definitely have an impact on the filter estimates unless we have the "forgetting condition"

$$f(\mathbf{x}_k|Z^k) \approx f\left(\mathbf{x}_k \left|\mathbf{x}_0^*, \cdots, \mathbf{x}_j^*, Z^k\right.\right). \quad (12)$$

Condition (12) is likely going to fail in a situation such as "total mixed labeling," that is, where, according to the true posterior distribution, for any given location, all possible labeling assignments are equally probable. It is argued in [16, section IV-C] that if a total mixed labeling arose at some point, it would persist at all later times. Consequently, if total mixed labeling already arises at $j' < j$, then it will persist at all later times, including time k. In this case, the true posterior $f(\mathbf{x}_k|Z^k)$ will contain mixed labeling. However, given $\mathbf{x}_j^*$, that is, assuming unique labels for all targets at time j, $f(\mathbf{x}_k|\mathbf{x}_j^*, Z^k)$ may not have any mixed labeling, thus significantly violating (12).

The M-SMC filter has therefore a tendency to "forget" the mixed labeling that exists in the true posterior density, leading to the underestimation of the labeling errors. This is what the authors in [1] observed empirically through the analysis of the SIR PF mechanism.

It is easy to see that multitarget tracking techniques based on representing the multitarget posterior as some sort of set of hypotheses and periodically pruning low-probability hypotheses will generally suffer from a similar degeneracy phenomenon. This will happen, in particular, if each hypothesis on the multitarget state at time k implicitly assumes hypotheses on the multitarget state at past times $0, \ldots, k-1$, like the approach presented in [6]. Since low-probability hypotheses are periodically pruned, information about past trajectories will eventually collapse into a single hypothesis.

III. MEASURE OF LABELING UNCERTAINTY IN BAYESIAN MTTL

In section I, the "probability of (in)correct labeling" is mentioned loosely as a possible measure of the labeling uncertainty. A proper way of formalizing this will be to define the *event* "correct labeling" and consider the probability of the (complementary) event. However, as mentioned in the same discussion, it is not obvious how to define *correct* labeling of tracks if the tracks themselves, being estimated values, are not correct.

At this point, many different alternative interpretations can be envisaged, such as [23]. We adopt a conditional approach. From the point of view of an end user, say, the radar operator, there already exist satisfactory solutions to the MTT problem at hand based on which point estimate of the target locations are on the display. The end user is interested in having labels assigned to these point estimates. With this viewpoint, we define the labeling probability associated with a point estimate of the *labeled* state to be the conditional probability that the assignment of labels is correct *if* the targets are (truly) located at the estimated positions. The precise definition is as follows.

1) *The Labelling (Un)Certainty*:

DEFINITION 3.1 Consider a RFS $\mathbf{X}$, a finite set of *labeled* target states, as described in section II-A. Let the RFS $\mathbf{S}$ be the corresponding *unlabeled* target states. The *labeling probability* associated with a realization $\mathbf{x} = \{[s^{(1)}, l^{(1)}], \cdots, [s^{(t)}, l^{(t)}]\}$ of $\mathbf{X}$ is defined to be the (conditional) probability that label $l^{(i)}$ is associated with the (unlabeled) state $s^{(i)}$, $i = 1, \ldots, t$, given that (there are t targets present and) $\mathbf{S} = \mathbf{s} := \{s^{(1)}, \ldots, s^{(t)}\}$. It is subsequently denoted as $\mathrm{L}(\mathbf{x}|\mathbf{s})$. ■

It can be shown (see, e.g., [16, lemma 3.5]) that the labeling probability is given by

$$\mathrm{L}(\mathbf{x}|\,\mathbf{s}) = \frac{f\left(\left\{\left[s^{(1)}, l^{(1)}\right], \cdots, \left[s^{(t)}, l^{(t)}\right]\right\}\right)}{\sum\limits_{\substack{\tilde{l}^{(i)}\in\Pi \\ 1\le i\le t}} f\left(\left\{[s^{(1)}, \tilde{l}^{(1)}], \ldots, [s^{(t)}, \tilde{l}^{(t)}]\right\}\right)} = \frac{f(\mathbf{x})}{f(\mathbf{s})}.$$

REMARK 3.2 The assumption of nonambiguity of initial labels can now be fully described in terms of the labeling probability. In particular, we assume that for any initial unlabeled state $\mathbf{s}_0$, there exists one labeled state $\mathbf{x}_0$ such that $\mathrm{L}(\mathbf{x}_0 \mid \mathbf{s}_0) = 1$. ■

2) *Posterior Labeling Probability*: For Bayesian labeling purposes, we are interested in the posterior version of the labeling probability, that is, conditioned on all observations up to and including time k, given by

$$\mathrm{L}\left(\mathbf{x}_k|\,\mathbf{s}_{k,} Z^k\right) = \frac{f\left(\mathbf{x}_k|\,Z^k\right)}{f\left(\mathbf{s}_k|\,Z^k\right)} = \frac{f\left(\mathbf{x}_k|\,Z^k\right)}{\sum\limits_{\mathbf{x}_k\in\Pi_k(\mathbf{s}_k)} f\left(\mathbf{x}_k|\,Z^k\right)}, \quad (13)$$

where

$$\Pi_k\left(\left\{s_k^{(1)}, \cdots, s_k^{(t_k)}\right\}\right) \triangleq \left\{\mathbf{x}_k \,\middle|\, \mathbf{x}_k = \{[s_k^{(1)}, l_k^{(1)}], \ldots, [s_k^{(t_k)}, l_k^{(t_k)}]\},\ f(\mathbf{x}_k|Z^k) > 0\right\}. \quad (14)$$

Note that (13) defines a discrete probability distribution on $\Pi_k(\mathbf{s}_k)$, that is, over the possible values of $\mathbf{x}_k$ formed by assigning labels to $\mathbf{s}_k$, because

$$\sum_{\mathbf{x}_k\in\Pi_k(\mathbf{s}_k)} \mathrm{L}\left(\mathbf{x}_k|\,\mathbf{s}_k, Z^k\right) = 1. \quad (15)$$

The following lemma will be useful in our later analysis.

LEMMA 3.3 Under condition (8),

$$\mathrm{L}\left(\mathbf{x}_k|\,\mathbf{s}_{k,} Z^k\right) = \frac{f\left(\mathbf{x}_k|\,Z^{k-1}\right)}{f\left(\mathbf{s}_k|\,Z^{k-1}\right)} \left[= \mathrm{L}\left(\mathbf{x}_k\,|\mathbf{s}_k, Z^{k-1}\right)\right]. \quad (16)$$

PROOF From (13), we have

$$\begin{aligned}
\mathrm{L}(\mathbf{x}_k|\mathbf{s}_k, Z^k) &= \frac{f(\mathbf{x}_k|Z^k)}{f\left(\mathbf{s}_k\,|Z^k\right)} = \frac{\frac{f(\mathbf{z}_k|\mathbf{x}_k)f\left(\mathbf{x}_k\,|Z^{k-1}\right)}{f\left(\mathbf{z}_k\,|Z^{k-1}\right)}}{f(\mathbf{s}_k|Z^k)} && \text{[from (4)]}\\
&= \frac{f\left(\mathbf{z}_k\,|\mathbf{s}_k\right) f\left(\mathbf{x}_k\,\left|Z^{k-1}\right.\right)}{f\left(\mathbf{z}_k\,\left|Z^{k-1}\right.\right) f\left(\mathbf{s}_k\,\left|Z^k\right.\right)} && \text{[using (8)]}\\
&= \frac{f\left(\mathbf{x}_k\,\left|Z^{k-1}\right.\right)}{\frac{f\left(\mathbf{z}_k|Z^{k-1}\right)f\left(\mathbf{s}_k|Z^k\right)}{f(\mathbf{z}_k|\mathbf{s}_k)}} = \frac{f\left(\mathbf{x}_k\,\left|Z^{k-1}\right.\right)}{f\left(\mathbf{s}_k\,\left|Z^{k-1}\right.\right)},
\end{aligned}$$

where the last equality follows from (9). ■

Lemma 3.3 complies with the interpretation that an observation provides information about label only via the location estimate (see section II-B2). Hence, given "locations" $\mathbf{s}_k$, the corresponding observation $\mathbf{z}_k$ does not carry any extra information about the labels at time k.

In other words, measurements cannot reduce the labeling uncertainties for some given locations $\mathbf{s}_k$. However, they may still reduce our overall (i.e., when locations are not given) uncertainties in the labels if they happen to tell us that $\mathbf{s}_k$ is more likely to belong to a region with fewer labeling uncertainties. On the other hand, once we have reached a stage of *total mixed labeling* [16], where the labeling uncertainties are large everywhere, this is no longer possible, and the uncertainties will never be reduced again.

3) *Labeling Error:* We can now use (13) to measure the labeling error in a labeled track estimate.

DEFINITION 3.4 Let $\hat{\mathbf{s}}_k$ be the *unlabeled tracks* corresponding to a set of *labeled tracks* $\hat{\mathbf{x}}_k$. Then the *labeling error* associated with $\hat{\mathbf{x}}_k$ is defined to be $1 - \mathrm{L}(\hat{\mathbf{x}}_k|\hat{\mathbf{s}}_k, Z^k)$. ■

IV. A LABELING ALGORITHM FOR MTTL PROBLEMS

In this section, we propose a labeling procedure that can be combined with existing (unlabeled) MTT algorithms to provide a complete MTTL solution. The resulting solution has the advantage that it makes the statistics proposed in section III readily available and does not suffer from the degeneracy phenomenon described in section II-C.

Recall from section II-C that when labels are part of the state, the degeneracy phenomenon of the SIR PF and similar algorithms creates an extra problem for MTTL in, for example, the situation depicted in Fig. 1. In the literature, the Rao-Blackwellized marginal particle filter (RBMPF) (see [24, 25]) has been successfully applied to counter PF degeneracy for the joint state and parameter estimation problem. The algorithm is essentially a combination of the Rao-Blackwellized particle filter [20] and the marginal particle filter [26].

The idea of the RBMPF is to split the state vector into two parts and handle only the part that is less likely to violate the "forgetting condition" in (12) using particles while attempting to express the conditional distribution of the other states exactly.

We apply the same idea to the MTTL problem by decoupling the labels and the unlabeled states. We estimate the unlabeled states by a suitable MTT algorithm and calculate the probabilities of the labels, conditioned on the unlabeled states, in a deterministic manner. By doing so, we prevent degeneracy in the MTTL solution that may arise due to the labels (e.g., when the solution involves pruning of hypothesis on labels).

The proposed method is described in detail in section IV-A. The corresponding algorithm is presented in section IV-B. Finally, we discuss the computational aspects in section IV-C.

A. Derivation of the Labeling Procedure

Note that due to the one-sided decoupling property, as described in section II-B3, we can iteratively obtain a representation of the unlabeled multitarget posterior $f(\mathbf{s}_k|Z^k)$ without any need to concern ourselves with labeling. Labeling can be done at a complementary step.

We will henceforth assume that $f(\mathbf{s}_k|Z^k)$ can be effectively approximated using state-of-the-art MTT techniques. We further assume that the chosen MTT technique for unlabeled tracking represents the unlabeled multitarget posterior $f(\mathbf{s}_k|Z^k)$ by a set of weighted particles as follows:

$$f\left(\mathbf{s}_k\left|Z^k\right.\right) \approx \sum_{i=1}^{N_P} w_k(i)\delta_{\mathbf{s}_k(i)}(\mathbf{s}_k). \tag{17}$$

Otherwise, we should be able to obtain such representation by numerically approximating (e.g., by sampling) the output of the algorithm. Choices of such unlabeled tracking algorithms include the M-SMC filters in [17] and [14, chapter 15], the hypotheses-based algorithm proposed in [6] (though we should discard the labels generated by the algorithm, as we are using another labeling scheme), and, for the track-before-detect measurement model, the Markov chain Monte Carlo (MCMC) MTT algorithm proposed in [15].

By combining (13) with (17), we can express the labeled posterior as

$$f\left(\mathbf{x}_k\left|Z^k\right.\right) \approx \sum_{i=1}^{N_P} w_k(i)\mathrm{L}\left(\mathbf{x}_k\left|\mathbf{s}_k(i), Z^k\right.\right)\delta_{\mathbf{s}_k(i)}(\mathbf{s}_k). \tag{18}$$

A general expected value $E[g(\mathbf{X}_k)\,|Z^k]$ can thus be approximated as

$$E\left[g(\mathbf{X}_k)|Z^k\right] \approx \sum_{i=1}^{N_P} w_k(i) \sum_{\mathbf{x}_k\in\Pi_k(\mathbf{s}_k(i))} g(\mathbf{x}_k)\mathrm{L}\left(\mathbf{x}_k|\mathbf{s}_k(i), Z^k\right),$$

where $\Pi_k(\cdot)$ is as given in (14).

From (18), it is clear that the additional quantity that our labeling algorithm should compute is the labeling probability $\mathrm{L}(\mathbf{x}_k|\mathbf{s}_k(i), Z^k)$. We develop a recursive algorithm to compute these probabilities by making use of (16). Noting that the denominator $f(\mathbf{s}_k|Z^{k-1})$ in (16) does not depend on the labels; we can approximate the labeling probabilities as

$$\mathrm{L}\left(\mathbf{x}_k\left|\mathbf{s}_k, Z^k\right.\right) \propto f\left(\mathbf{x}_k\left|Z^{k-1}\right.\right) =: \overline{\mathrm{L}}\left(\mathbf{x}_k\left|\mathbf{s}_k, Z^k\right.\right)$$
$$= \int f\left(\mathbf{x}_k\left|\mathbf{x}_{k-1}\right.\right) f\left(\mathbf{x}_{k-1}\left|Z^{k-1}\right.\right)\delta\mathbf{x}_{k-1} \tag{19}$$

$$\approx \sum_{j=1}^{N_P} w_{k-1}(j) \sum_{\mathbf{x}_{k-1}\in\Pi_{k-1}(\mathbf{s}_{k-1}(j))} f(\mathbf{x}_k|\mathbf{x}_{k-1}) \times \mathrm{L}\left(\mathbf{x}_{k-1}\left|\mathbf{s}_{k-1}(j), Z^{k-1}\right.\right), \tag{20}$$

assuming that we have already computed the labeling probabilities $\mathrm{L}(\mathbf{x}_{k-1}|\mathbf{s}_{k-1}(j), Z^k)$ at time $(k-1)$, for $j = 1, \ldots, N_p$ and $\mathbf{x}_{k-1} \in \Pi_{k-1}(\mathbf{s}_{k-1}(j))$.

It should be noted that the summations in (20) need to be computed for every particle $\mathbf{s}_k(i)$ and its labeled version $\mathbf{x}_k \in \Pi_k(\mathbf{s}_k(i))$ at time k. Depending on the number of targets present and the number of used particles, this may become computationally intractable. One way to reduce the burden is to set the terms inside the second sum in (20) with negligible contribution to zero.

This is equivalent to approximating the sets $\Pi_k(\mathbf{s}_k)$ in a specific way. First, note from (14), (18), and (20) that $\Pi_k(\mathbf{s}_k)$ may be approximated as

$$\Pi_k\left(\left\{s_k^{(1)}, \ldots, s_k^{(t_k)}\right\}\right) \approx \Big\{\mathbf{x}_k \,\Big|\, \mathbf{x}_k = \left\{\left[s_k^{(1)}, l_k^{(1)}\right], \ldots, \left[s_k^{(t_k)}, l_k^{(t_k)}\right]\right\} \text{ and } \exists\, j, \mathbf{x}_{k-1} \in \Pi_{k-1}(\mathbf{s}_{k-1}(j)) \text{ s.t. } f(\mathbf{x}_k|\mathbf{x}_{k-1})\mathrm{L}\left(\mathbf{x}_{k-1}|\mathbf{s}_{k-1}(j), Z^{k-1}\right) > 0\Big\}. \tag{21}$$

In the approximation above, one can use a higher threshold τ_k instead of 0 and use the condition

$$f(\mathbf{x}_k|\mathbf{x}_{k-1})\mathrm{L}\left(\mathbf{x}_{k-1}|\mathbf{s}_{k-1}(j), Z^{k-1}\right) > \tau_k.$$

Then the number of terms in the sums over $\Pi_k(\mathbf{s}_k)$ will be reduced. One has to be very careful, though, in selecting a threshold. A large threshold can cause labeling hypotheses to disappear prematurely. This would, in turn, lead to a sort of degeneracy that we are trying to prevent in the first place.

To initialize the recursion for $\mathrm{L}(\mathbf{x}_k|\mathbf{s}_k, Z^k)$, the quantities $\Pi_0(\mathbf{s}_0)$ and $\mathrm{L}(\mathbf{x}_0|\mathbf{s}_0)$ are obtained as follows:

$$\Pi_0\left(\left\{s_0^{(1)}, \cdots, s_0^{(t_0)}\right\}\right) \approx \left\{\mathbf{x}_0 \,\middle|\, \mathbf{x}_0 = \left\{\left[s_0^{(1)}, l_0^{(1)}\right], \ldots, \left[s_0^{(t_0)}, l_0^{(t_0)}\right]\right\},\ f(\mathbf{x}_0) > 0\right\} \tag{22}$$

and

$$\mathrm{L}\left(\mathbf{x}_0\left|\mathbf{s}_0\right.\right) = f\left(\mathbf{x}_0\right)/f\left(\mathbf{s}_0\right). \tag{23}$$

Note that according to the assumption of nonambiguity of initial labels (see section II-B1 and Remark 3.2), given $\mathbf{s}_0$, $\Pi_0(\mathbf{s}_0)$ will be a singleton set, and $\mathrm{L}(\mathbf{x}_0|\mathbf{s}_0) = 1$ for $\mathbf{x}_0 \in \Pi_0(\mathbf{s}_0)$.

B. The Algorithm

We now describe the (sub)algorithm for the labeling procedure to be used as a "plug-in" to an (unlabeled) MTT algorithm. The latter is assumed to generate, at every time step k, a particle representation $\{\mathbf{s}_k(i), w_k(i)\}_{i=1}^{N_P}$ of the unlabeled posterior. For each particle $\mathbf{s}_k(i)$, the labeling algorithm computes the corresponding labeling probabilities $\mathrm{L}(\mathbf{x}_k|\mathbf{s}_k(i), Z^k)$ using the (unlabeled) particles and labeling probabilities of previous time $(k-1)$.

Initialization: (for each $\mathbf{s}_0(i)$, $i = 1, \ldots, N_P$)

- Set $\Pi_0(\mathbf{s}_0(i))$ according to (22), by taking $\mathbf{s}_0 = \mathbf{s}_0(i)$.
- For each $\mathbf{x}_0 \in \Pi_0(\mathbf{s}_0(i))$, set $\mathrm{L}(\mathbf{x}_0|\mathbf{s}_0(i)) = f(\mathbf{x}_0)/f(\mathbf{s}_0(i))$.

At every time step k: (for each $\mathbf{s}_k(i)$, $i = 1, \ldots, N_P$)

(L.1) Obtain $\Pi_k(\mathbf{s}_k(i))$ by taking $\mathbf{s}_k = \mathbf{s}_k(i)$ in (21).

(L.2) For each $\mathbf{x}_k \in \Pi_k(\mathbf{s}_k(i))$, calculate the unnormalized labeling probabilities $\overline{\mathrm{L}}(\mathbf{x}_k|\mathbf{s}_k(i), Z^k)$ according to (20).

(L.3) For $\mathbf{x}_k \in \Pi_k(\mathbf{s}_k(i))$, normalize the labeling probabilities

$$\mathrm{L}\left(\mathbf{x}_k \left|\mathbf{s}_k(i), Z^k\right.\right) = \frac{\overline{\mathrm{L}}\left(\mathbf{x}_k \left|\mathbf{s}_k(i), Z^k\right.\right)}{\sum_{\tilde{\mathbf{x}}_k \in \Pi_k(\mathbf{s}_k(i))} \overline{\mathrm{L}}\left(\tilde{\mathbf{x}}_k \left|\mathbf{s}_k(i), Z^k\right.\right)}.$$

C. Computational Cost of the Labeling Procedure

If we consider a constant number of targets t, there are t possible labels. Hence, given a location $\mathbf{s}_{k-1}(i)$, the corresponding $|\Pi_{k-1}(\mathbf{s}_{k-1}(i))|$ can be as high as $t!$. Then, from (20), it follows that the worst-case complexity of calculating a single labeling probability for a single particle–label combination $\mathbf{x}_k$ is $O(N_P t!)$ and $O(N_P^2 (t!)^2)$ to compute all labeling probabilities for all particles. Needless to say, this computational cost can be prohibitive if we have large number of targets.

The computational problem is aggravated when we consider target births and deaths, where in fact $\Pi_k(\mathbf{s}_k)$ may grow with time, for example, when a target disappeared a long time ago but its corresponding label still maintains a nonzero probability of existence. This may happen if other targets whose identities have been confused with the "dead" target still exist.

Therefore, without additional approximations, the labeling procedure presented here would be unsuitable for large-scale MTTL problems. However, for problems of tracking a small group of targets in a situation like Fig. 1 and individually identifying the targets after separation, the labeling procedure is suitable. The algorithm has also good parallelization properties: steps (L.1)–(L.3) can be fully parallelized by letting each (parallel) computing node process a single labeling hypothesis $\mathbf{x}_k(i)$ with a computational complexity of $O(N_P\, t!)$.

V. NUMERICAL RESULTS

In this section, we present the simulation results for the proposed labeling procedure by analyzing the effect of adding it as a "plug-in" to an *unlabeled* M-SMC filter. We compare the results to those from a *labeled* M-SMC filter, which estimates labels as part of the single-state state.

We start by explaining the metrics we use to evaluate the algorithms. Subsequently, in section V-B, the considered scenarios are described. The comparison results are presented in sections V-C and V-D.

A. Metrics for Performance Evaluation

We recall that though the overall goal of MTTL is to obtain labeled tracks, the focus of this paper is on the labeling part. Subsequently, we compare different algorithms on the basis of the labels the algorithms will assign to a set of given *unlabeled states*. In the simulation examples, the natural choice of the unlabeled states is the synthetic location values used in the simulation.

Suppose that in a simulation run $\tilde{\mathbf{s}}_k$ represents the true unlabeled multitarget states (locations) at time k and that $\tilde{\mathbf{x}}_k$ represents the true labeled multitarget states. For any given algorithm producing a particle representation of the posterior $f(\mathbf{x}_k|Z^k)$ at each time k, we calculate the labeling probabilities $\mathrm{L}(\mathbf{x}_k|\tilde{\mathbf{s}}_k, Z^k)$ for $\mathbf{x}_k \in \Pi_k(\tilde{\mathbf{s}}_k)$ by calculating first the unnormalized versions from (19), where the set integral is evaluated using the particle representation of $f(\mathbf{x}_{k-1}|Z^{k-1})$.

Subsequently, we compare the point estimate $\hat{\mathbf{x}}_k$, given by

$$\hat{\mathbf{x}}_k = \arg \max_{\mathbf{x}_k \in \Pi_k(\tilde{\mathbf{s}}_k)} \mathrm{L}\left(\mathbf{x}_k \left|\tilde{\mathbf{s}}_k, Z^k\right.\right), \tag{24}$$

to the true labeled state $\tilde{\mathbf{x}}_k$ using the hit-or-miss metric:

$$\begin{aligned}\epsilon(\hat{\mathbf{x}}_k, \tilde{\mathbf{x}}_k) &\equiv \epsilon\Big(\Big\{\big[\hat{s}_k^{(1)}, \hat{l}_k^{(1)}\big], \ldots, \big[\hat{s}_k^{(t_k)}, \hat{l}_k^{(t_k)}\big]\Big\}, \\ &\qquad \Big\{\big[\tilde{s}_k^{(1)}, \tilde{l}_k^{(1)}\big], \ldots, \big[\tilde{s}_k^{(t_k)}, \tilde{l}_k^{(t_k)}\big]\Big\}\Big) \\ &\triangleq \begin{cases} 1, & \exists\, i, j \text{ s.t. } \hat{s}_k^{(i)} = \tilde{s}_k^{(j)} \text{ and } \hat{l}_k^{(i)} \neq \tilde{l}_k^{(j)} \\ 0, & \text{otherwise;} \end{cases}\end{aligned} \tag{25}$$

that is, the value of the metric is 1 if there is at least one incorrectly assigned label and 0 otherwise. Naturally, the metric is statistically relevant only when averaged over a sufficient number of Monte Carlo runs. Thus, we define for an MTTL algorithm the *observed average labeling error* at time k as

$$\varepsilon_k^{\text{true}} = \frac{1}{N_R} \sum_{i=1}^{N_R} \epsilon^{(i)}, \tag{26}$$

where N_R is the number of Monte Carlo runs and $\epsilon^{(i)}$ is calculated according to (25) for the ith Monte Carlo run.

Recall further that the labeling error in $\hat{\mathbf{x}}_k$, given by $1 - \mathrm{L}(\hat{\mathbf{x}}_k|\tilde{\mathbf{s}}_k, Z^k)$, represents the (conditional) probability that the true labels $(L_k^{(1)}, \ldots, L_k^{(t_k)})$ associated with the locations $(\tilde{s}_k^{(1)}, \ldots, \tilde{s}_k^{(t_k)})$ are different from $(\hat{l}_k^{(1)}, \ldots, \hat{l}_k^{(t_k)})$. In other words, $\epsilon(\hat{\mathbf{x}}_k, \tilde{\mathbf{x}}_k)$ of (25) can be considered as a realization/observation from the Bernoulli distribution with success probability $1 - \mathrm{L}(\hat{\mathbf{x}}_k|\tilde{\mathbf{s}}_k, Z^k)$. Thus, when averaged over a series of Monte Carlo runs, these two quantities should be close to each other.

We thus also calculate the *algorithm-suggested* average labeling error as

$$\varepsilon_k^{\text{calc}} = \frac{1}{N_R}\sum_{i=1}^{N_R}\left(1 - \mathrm{L}\left(\hat{\mathbf{x}}_k(i)\,\middle|\,\tilde{\mathbf{s}}_k(i), Z^k(i)\right)\right), \quad (27)$$

where i stands for the Monte Carlo run number. As argued above, the quantities $\varepsilon_k^{\text{calc}}$ and $\varepsilon_k^{\text{true}}$ should not have a large difference. Otherwise, it will indicate inconsistency between the observation and the expectation, both derived using the same algorithm.

Another related measure of importance is the variance of the labeling error, $(1 - \mathrm{L}(\hat{\mathbf{x}}_k|\tilde{\mathbf{s}}_k, Z^k))$. Note that this variance has two contributors: 1) the variability due to the varying sequence of (unlabeled) state $\mathbf{s}_k$ and observations Z^k and 2) the variability in the estimation of posterior distribution of the labeled locations. It is, however, the latter type that is more relevant for us because a high variance would indicate that the calculated $(1 - \mathrm{L}(\hat{\mathbf{x}}_k|\tilde{\mathbf{s}}_k, Z^k))$ is unreliable. In order to observe this latter variance, we perform a second analysis, this time running the algorithm many times on a *fixed* sequence of measurements Z^k (and, necessarily, for fixed sequence of locations $\mathbf{s}_k$). Then the variance in these estimated errors will be due entirely to the algorithm of obtaining the labeled states. In this analysis, we look at the standard deviation of the calculated labeling error given by

$$\sigma_k^{\varepsilon} = \sqrt{\frac{1}{N_R}\sum_{i=1}^{N_R}\left(\left(1 - \mathrm{L}\left(\hat{\mathbf{x}}_k(i)\,\middle|\,\tilde{\mathbf{s}}_k, Z^k\right)\right) - \varepsilon_k^{\text{calc}}\right)^2}. \quad (28)$$

B. Simulation Scenarios

In our analysis, we consider the following four scenarios:

1) Two targets approach each other, move closely spaced for a while, and separate.
2) Two targets approach each other, move closely spaced for a while, and separate, crossing paths.
3) Two targets approach each other, move coalesced for a while, and separate.
4) Two targets approach each other, move closely spaced for a while, and separate. However, one of the targets appears later but well before they come close, and the other disappears soon after the separation.

The trajectory of the targets and the simulated measurements in one Monte Carlo run is shown in Fig. 4. The targets (as well as the time) move from left to right. The multitarget measurement model $f(\mathbf{z}_k|\mathbf{s}_k)$ is taken to be the detection-type measurement model described in [14, section 12.3]. Missed detections and false alarms are considered only in the last scenario (with target birth and death), with probability of detection 0.95 and uniform clutter density of $2 \cdot 10^{-7}$ per unit of area. The single-measurement, single-target likelihood function is given by

$$p\left(z_k^{(i)}\,\middle|\,s_k^{(j)}\right) = \mathcal{N}\left(z_k^{(i)}; \left[p_{\mathrm{x}}^{(j)}, p_{\mathrm{y}}^{(j)}\right], \begin{bmatrix} 676 & 0 \\ 0 & 676 \end{bmatrix}\right). \quad (29)$$

The location has the form $S_k^{(i)} = [P_{\mathrm{x}}^{(i)}, P_{\mathrm{y}}^{(i)}, V_{\mathrm{x}}^{(i)}, V_{\mathrm{y}}^{(i)}]$, where $(P_{\mathrm{x}}^{(i)}, P_{\mathrm{y}}^{(i)})$ is the position in Cartesian coordinates x and y and $(V_{\mathrm{x}}^{(i)}, V_{\mathrm{y}}^{(i)})$ corresponds to the velocities. The single-target state transition model corresponds to the popular discretized white noise acceleration model described in [27], with $T = 2$ as the interval between observations and $\sigma^2 = 676$ as the power spectral density of the process noise.

In all scenarios, we assume perfect knowledge of the targets' initial positions as well as their time of appearance (for the appearing target in scenario 4). The possibility of target death is considered only in scenario 4, with the probability of target survival at each time step assumed to be constant and equal to 0.95. With the given assumptions, the multitarget predictive density $f(\mathbf{x}_k|Z^{k-1})$ and the posterior density $f(\mathbf{x}_k|Z^k)$ are generalized multi-Bernoulli RFS densities, with analytical formulas presented in [16, section IV].

For all scenarios, we evaluate two MTTL algorithms:

1) A "naive" M-SMC filter that attempts to estimate labels as part of the single-target state without preventing the degeneracy phenomenon described in section II-C;
2) A decoupled tracking/labeling approach that uses an M-SMC filter only for unlabeled tracking and the labeled "plug-in" described in section IV-B to calculate the labeling probabilities.

For both the naive and the decoupled algorithms, we use 2000 particles for scenarios 1, 2, and 3, and 4000 particles for scenario 4. For both filters, we use blind importance sampling; that is, we use $f(\mathbf{x}_k|\mathbf{x}_{k-1})$ as proposal density for the naive M-SMC filter and $f(\mathbf{s}_k|\mathbf{s}_{k-1})$ for the decoupled algorithm. For the calculation of the average errors, we have used $N_R = 100$ Monte Carlo runs.

C. Results for Monte Carlo Runs With Varying Sequence of Measurements

The results from the Monte Carlo simulation with $\mathbf{s}_k$ and $\mathbf{Z}_k$ being regenerated at each Monte Carlo run are shown in Fig. 5. The *observed* and *algorithm-suggested* average labeling errors, $\varepsilon_k^{\text{true}}$ and $\varepsilon_k^{\text{calc}}$, respectively, are plotted for both the naive M-SMC filter and the decoupled MTTL algorithms. In terms of the observed errors, $\varepsilon_k^{\text{true}}$, we see that the decoupled approach provides lower average labeling errors for all scenarios. The improvement of using our proposed labeling procedure is much more significant in scenarios 1 and 2, where the separation between the targets was larger (and hence ambiguity in label-to-location association was lower).

In terms of the algorithm-suggested errors, $\varepsilon_k^{\text{calc}}$, we see that for the naive M-SMC filter, after the targets separate, it decreases with time. This indicates that the algorithm becomes increasingly confident in the

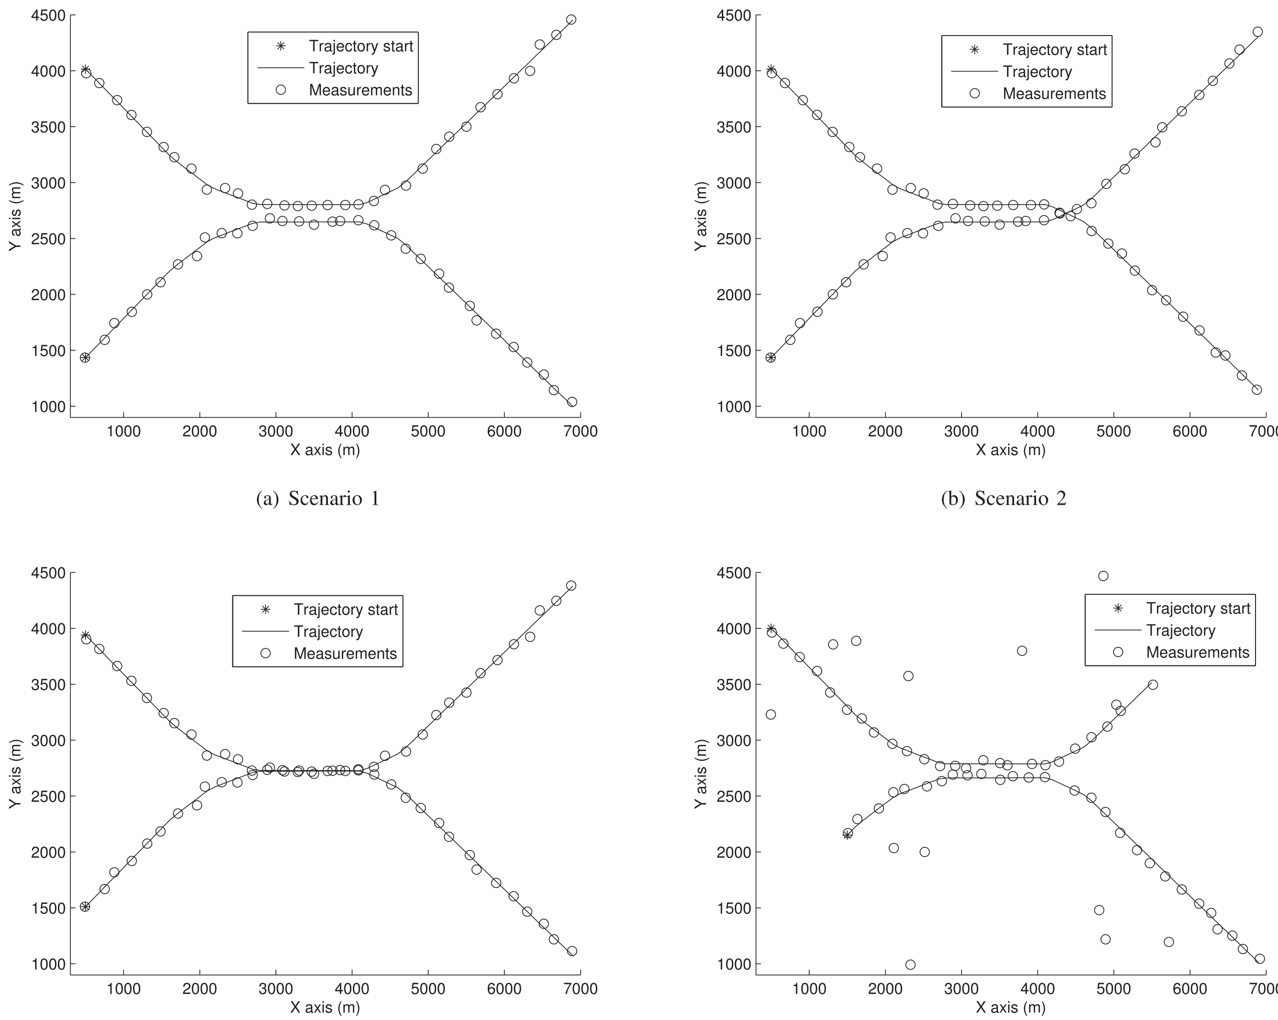

(a) Scenario 1 (b) Scenario 2 (c) Scenario 3 (d) Scenario 4

Fig. 4. Multitarget simulation scenarios. Plots show the trajectory of the targets, moving from left to right, and the measurements in a particular MC run. The x-axis can also be considered as time axis (naturally, with a different scale).

correctness of the assigned labels, while the observed errors, $\varepsilon_k^{\text{true}}$, are much higher. This clearly shows that the algorithm underestimates the true labeling uncertainty. This can be attributed to the degeneracy inherent in the naive M-SMC filter.

The decoupled algorithm, on the other hand, exhibits far more consistency between the observed and calculated errors, as $\varepsilon_k^{\text{calc}}$ remains constant over time after the targets separate. This is consistent with the theoretical behavior of $f(\mathbf{x}_k|Z^k)$ for this type of scenario, namely, the persistence of mixed labeling, as described in [16, section IV-C].

D. Results for Monte Carlo Runs With Fixed Sequence of Measurements

In this section, we analyze the standard deviation, σ_k^{ε}, of the labeling errors from a labeled MTTL solution. For this analysis, as mentioned in section V-A, we perform the simulations for fixed sequence of $\mathbf{s}_k$ and $\mathbf{Z}_k$. Since in section V-C we have established that the estimated labeling errors with the naive (labeled) M-SMC filter are erroneous (severely underestimated) there is no point analyzing the variance of these underestimated quantities. So we exclude the naive (unlabeled) M-SMC filter and consider only the decoupled algorithm in our analysis. Furthermore, we consider only scenarios 1 and 3 to see how the estimation of labeling error is affected by the degree of target separation (while they move closely spaced).

The results on the standard deviation, σ_k^{ε}, of the algorithm-suggested labeling errors are shown in Fig. 6. Recall that with fixed observation series, the variation in $\varepsilon_k^{\text{calc}}$ is caused solely by the estimation of posterior distribution of the labeled locations. In our case, this means by the (unlabeled) M-SMC filter and subsequently when that is used to calculate the labeling probabilities. Hence, a large σ_k^{ε} indicates low reliability of the augmented algorithm. It is interesting to notice in Fig. 6

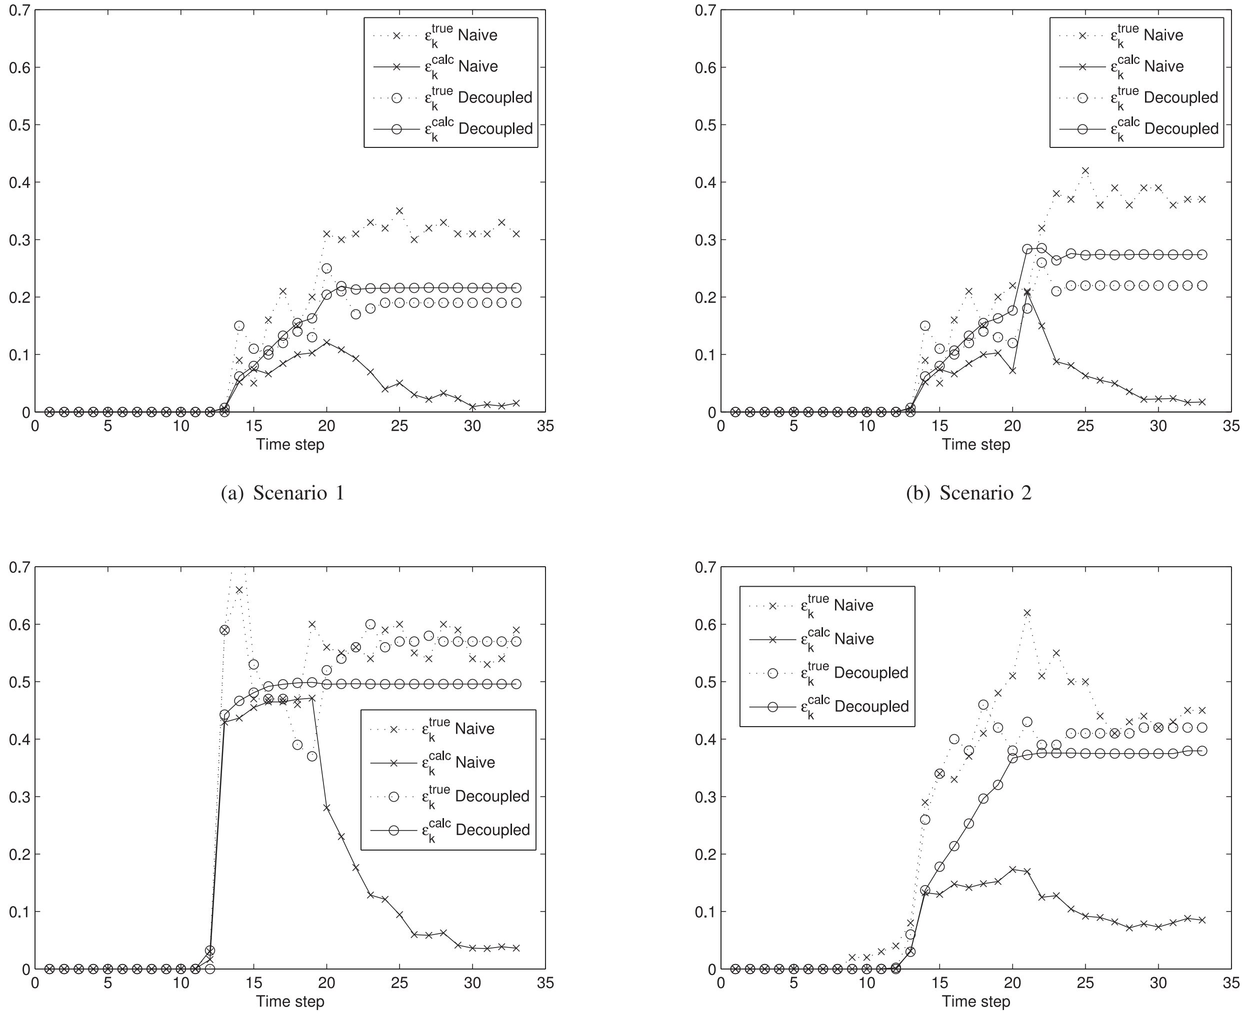

(a) Scenario 1

(b) Scenario 2

(c) Scenario 3

(d) Scenario 4

Fig. 5. Comparison of naive labeled M-SMC filter and decoupled M-SMC filter (varying Z^k).

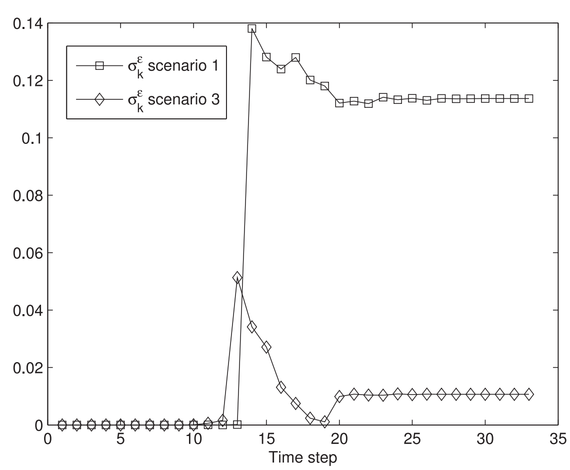

Fig. 6. MC variation in estimated labeling error probabilities for the decoupled MTTL algorithm.

that the variance is higher when targets came close but not very close than when the targets actually coalesce. At first sight, this may seem counterintuitive because labeling should be easier when the targets are more separated, and so variability in error probabilities should be less.

However, this is not so if we realize that when the targets coalesce and move at this state for some time, then "total mixed labeling" will appear. In other words, according to the true posterior distribution, given any possible location of the targets, all possible label assignments will become equally probable, and it will continue to be like that at all later times (see, e.g., [16, section IV-C]). A good algorithm will reflect this by having/estimating the labeling errors to be almost constant (in our case 0.5), which will lead to smaller variance.

VI. CONCLUSIONS AND RECOMMENDATIONS

In this paper, we have complemented the introduction to the Bayesian MTTL problem, presented in [6], with a discussion on the additional assumptions needed to keep the target labeling problem meaningful. A mathematical characterization of the labeling uncertainties present in an

MTTL solution is provided by defining properly the quantities such as *labeling probability* and *labelling error*. These quantities have clear practical interpretations (i.e., meaningful to the user of the system) rather than being only abstract mathematical quantities. The existing literature either lacks these sort of uncertainty measures or these measures do not have clear interpretation.

We have also devised a new labeling procedure that can be combined with existing (unlabeled) MTT algorithms to provide a complete solution to the Bayesian MTTL problem. The resulting solution avoids the degeneracy that may appear due to the labels (e.g., when the solution involves pruning of hypothesis on labels).

Numerical examples show that when the (unlabeled) M-SMC filter is augmented with the proposed labeling procedure, it performs much better than the naive labeled M-SMC filter, which estimates labels as part of the single-target state.

In terms of theoretical research, an interesting topic of future work would be to devise different ways of generating labels that have better capabilities of assigning unambiguous labels to appearing targets than in [6] and therefore could be applied to more general scenarios. In terms of practical research, it is worth investigating possible improvement of the computational performance of our proposed labeling procedure by finding more computationally efficient ways to calculate the labeling probabilities. Naturally, it would also be interesting to try the labeling procedure with more complex observation models, such as the track-before-detect observation model.

ACKNOWLEDGMENT

The authors would like to thank Ba Ngu-Vo for his valuable comments and suggestions. Authors also thank the anonymous reviewers for their helpful comments in improving the paper.

APPENDIX

To see that the conditions (7) and (8) indeed lead to the consistency criteria mentioned in section II-B2 (Fig. 2), consider $\mathbf{s}_k = \{s_k^{(1)}, \ldots, s_k^{(t)}\}$. From marginalization and (5), we have the following:

$$
\begin{aligned}
f(\mathbf{s}_k|Z^{k-1})
&= f\left(\left\{s_k^{(1)}, \ldots, s_k^{(t)}\right\} \middle| Z^{k-1}\right) \\
&= \sum_{\substack{l_k^{(i)} \in \Pi \\ 1\le i \le t}} f\left(\left\{\left[s_k^{(1)}, l_k^{(1)}\right], \ldots, \left[s_k^{(t)}, l_k^{(t)}\right]\right\} \middle| Z^{k-1}\right) \\
&= \sum_{\substack{l_k^{(i)} \in \Pi \\ 1\le i \le t}} \int f\left(\left\{\left[s_k^{(1)}, l_k^{(1)}\right], \ldots, \left[s_k^{(t)}, l_k^{(t)}\right]\right\} \middle| \mathbf{x}_{k-1}\right) \\
&\quad \times f\left(\mathbf{x}_{k-1} \middle| Z^{k-1}\right) \delta \mathbf{x}_{k-1}.
\end{aligned} \tag{30}
$$

From the definition of set integral (see, e.g., [6, proposition 2], we then have (31) below. Furthermore taking the sum over $l_k^{(i)}$ in (31) inside the integral (which is permitted because all the terms are nonnegative) and subsequently using marginalization and (7), we obtain

$$
\begin{aligned}
&f(\mathbf{s}_k|Z^{k-1}) \\
&= \sum_{\substack{l_k^{(i)} \in \Pi \\ 1\le i \le t}} \sum_{m=0}^{\infty} \frac{1}{m!} \sum_{\substack{l_{k-1}^{(j)} \in \Pi \\ 1\le j \le m}} \int_{(\mathbb{R}^n)^m} f\bigg(\left\{\left[s_k^{(1)}, l_k^{(1)}\right], \ldots, \right. \\
&\quad \left. \left[s_k^{(t)}, l_k^{(t)}\right]\right\} \bigg| \left\{\left[s_{k-1}^{(1)}, l_{k-1}^{(1)}\right], \ldots, \left[s_{k-1}^{(m)}, l_{k-1}^{(m)}\right]\right\} \bigg) \\
&\quad \times f\left(\left\{\left[s_{k-1}^{(1)}, l_{k-1}^{(1)}\right], \ldots, \left[s_{k-1}^{(m)}, l_{k-1}^{(m)}\right]\right\} \middle| Z^{k-1}\right) \\
&\quad \times d(s_{k-1}^{(1)}, \ldots, s_{k-1}^{(m)})
\end{aligned} \tag{31}
$$

$$
\begin{aligned}
&= \sum_{m=0}^{\infty} \frac{1}{m!} \sum_{\substack{l_{k-1}^{(j)} \in \Pi \\ 1\le j \le m}} \int_{(\mathbb{R}^n)^m} f\left(\mathbf{s}_k \middle| \left\{s_{k-1}^{(1)}, \ldots, s_{k-1}^{(m)}\right\}\right) \\
&\quad \times f\left(\left\{\left[s_{k-1}^{(1)}, l_{k-1}^{(1)}\right], \ldots, \left[s_{k-1}^{(m)}, l_{k-1}^{(m)}\right]\right\} \middle| Z^{k-1}\right) \\
&\quad \times d\left(s_{k-1}^{(1)}, \ldots, s_{k-1}^{(m)}\right).
\end{aligned} \tag{32}
$$

Interchanging the sum over $l_{k-1}^{(j)}$ and the integral, it follows, from marginalization and the definition of set integral, that

$$
\begin{aligned}
&f(\mathbf{s}_k|Z^{k-1}) \\
&\quad = \sum_{m=0}^{\infty} \frac{1}{m!} \int_{(\mathbb{R}^n)^m} f\left(\mathbf{s}_k | \{s_{k-1}^{(1)}, \ldots, s_{k-1}^{(m)}\}\right) \\
&\quad \times f\left(\{s_{k-1}^{(1)}, \ldots, s_{k-1}^{(m)}\} | Z^{k-1}\right) d(s_{k-1}^{(1)}, \ldots, s_{k-1}^{(m)}) \\
&= \int f(\mathbf{s}_k|\mathbf{s}_{k-1}) f(\mathbf{s}_{k-1}|Z^{k-1}) \delta \mathbf{s}_{k-1},
\end{aligned}
$$

which is the Chapman–Kolmogorov equation corresponding to the unlabeled RFS model.

Furthermore, from (4) and (8), we have

$$
\begin{aligned}
&f(\mathbf{s}_k|Z^k) \\
&= \sum_{\substack{l_k^{(i)} \in \Pi \\ 1\le i \le t}} f\left(\left\{[s_k^{(1)}, l_k^{(1)}], \ldots, [s_k^{(t)}, l_k^{(t)}]\right\} \middle| Z^k\right) \\
&= \sum_{\substack{l_k^{(i)} \in \Pi \\ 1\le i \le t}} \frac{f\left(\mathbf{z}_k \middle| \left\{[s_k^{(1)}, l_k^{(1)}], \ldots, [s_k^{(t)}, l_k^{(t)}]\right\}\right)}{f\left(\mathbf{z}_k \middle| Z^{k-1}\right)} \\
&\qquad \times f\left(\left\{[s_k^{(1)}, l_k^{(1)}], \ldots, [s_k^{(t)}, l_k^{(t)}]\right\} \middle| Z^{k-1}\right) \\
&= \sum_{\substack{l_k^{(i)} \in \Pi \\ 1\le i \le t}} \frac{f(\mathbf{z}_k|\mathbf{s}_k)}{f\left(\mathbf{z}_k \middle| Z^{k-1}\right)} f\left(\{[s_k^{(1)}, l_k^{(1)}], \ldots, [s_k^{(t)}, l_k^{(t)}]\} | Z^{k-1}\right) \\
&= \frac{f(\mathbf{z}_k|\mathbf{s}_k)}{f\left(\mathbf{z}_k \middle| Z^{k-1}\right)} f(\mathbf{s}_k|Z^{k-1}),
\end{aligned}
$$

which is the measurement update equation for the unlabeled RFS model.

REFERENCES

[1] Boers, Y., Sviestins, E., and Driessen, J. N.
Mixed labelling in multitarget particle filtering.
IEEE Transactions on Aerospace and Electronic Systems, **46**, 2 (2010), 792–802.

[2] Salmond, D. J., Fisher, D., and Gordon, N. J.
Tracking and identification for closely spaced objects in clutter.
in *Proceedings of the European Control Conference*, Brussels, Belgium, 1997, 2973–2978.

[3] Ma, W.-K., Vo, B.-N., Singh, S., and Baddeley, A.
Tracking an unknown time-varying number of speakers using TDOA measurements: a random finite set approach.
IEEE Transactions on Signal Processing, **54**, 9 (2006), 3291–3304.

[4] Morelande, M., Kreucher, C., and Kastella, K.
A Bayesian approach to multiple target detection and tracking.
IEEE Transactions on Aerospace and Electronic Systems, **55**, 5 (2007), 1589–1604.

[5] García-Fernández, A., and Grajal, J.
Multitarget tracking using the joint multitrack probability density.
In *Proceedings of the 12th International Conference on Information Fusion*, Seattle, WA, 2009, 595–602.

[6] Vo, B.-T., and Vo, B.-N.
Labeled random finite sets and multi-object conjugate priors.
IEEE Transactions on Signal Processing, **61**, 13 (2013), 3460–3475.

[7] García-Fernández, A., Grajal, J., and Morelande, M.
Two-layer particle filter for multiple target detection and tracking.
IEEE Transactions on Aerospace and Electronic Systems, **49**, 3 (2013), 1569–1588.

[8] Blom, H., and Bloem, E.
Permutation invariance in Bayesian estimation of two targets that maneuver in and out formation flight.
In *Proceedings of the 12th International Conference on Information Fusion*, Seattle, WA, 2009, 1296–1303.

[9] García-Fernández, A., Morelande, M., and Grajal, J.,
Particle filter for extracting target label information when targets move in close proximity.
In *Proceedings of the 14th International Conference on Information Fusion*, Chicago, IL, 2011.

[10] Crouse, D., Willett, P., Svensson, L., Svensson, D., and Guerriero, M.
The set MHT.
In *Proceedings of the 14th International Conference on Information Fusion*, Chicago, IL, 2011.

[11] Blom, H. A. P., and Bloem, E. A.
Decomposed particle filtering and track swap estimation in tracking two closely spaced targets.
In *Proceedings of the 14th International Conference on Information Fusion*, Chicago, IL, 2011.

[12] Georgescu, R., Willett, P., Svensson, L., and Morelande, M.
Two linear complexity particle filters capable of maintaining target label probabilities for targets in close proximity.
In *Proceedings of the 15th International Conference on Information Fusion*, Singapore, 2012, 2370–2377.

[13] Svensson, L., and Morelande, M.
Target tracking based on estimation of sets of trajectories.
In *Proceedings of the 17th International Conference on Information Fusion*, Salamanca, Spain, 2014.

[14] Mahler, R.
Statistical Multisource-Multitarget Information Fusion.
Norwood, MA: Artech House, 2007.

[15] Bocquel, M.
Random finite sets in multi-target tracking—efficient sequential MCMC implementation. Ph.D. dissertation, University of Twente, Enschede, The Netherlands, 2013.

[16] Aoki, E. H., Boers, Y., Svensson, L., Mandal, P. K., and Bagchi, A. A Bayesian solution to multi-target tracking problems with mixed labelling. Department of Applied Mathematics, University of Twente, Enschede, The Netherlands, Memorandum 2036, Jul. 2014 [Online] http://eprints.eemcs.utwente.nl/24915.

[17] Vo, B.-N., Singh, S., and Doucet, A.
Sequential Monte Carlo methods for multitarget filtering with random finite sets.
IEEE Transactions on Aerospace and Electronic Systems, **41**, no. 4 (2005), 1224–1245.

[18] Gordon, N. J., Salmond, D. J., and Smith, A. F. M.
Novel approach to non-linear/non-Gaussian Bayesian state estimation.
Radar and Signal Processing, IEE Proceedings F, **140**, no. 2 (1993), 107–113.

[19] Kitagawa, G.
A Monte Carlo filtering and smoothing method for non-Gaussian nonlinear state space models.
In *Proceedings of the 2nd US-Japan Joint Seminar on Statistical Time Series Analysis*, Honolulu, HI, 1993, 110–131.

[20] Andrieu, C., and Doucet, A.
Particle filtering for partially observed Gaussian state space models.
Journal of the Royal Statistical Society Series B, **64** (2002), 827–836.

[21] Kantas, N., Doucet, A., Singh, S. S., and Maciejowski, J. M.
An overview of sequential Monte Carlo methods for parameter estimation on general state space models.
In *Proceedings of the 15th IFAC Symposium on System Identification (SYSID)*, Saint-Malo, France, 2009.

[22] Doucet, A., and Johansen, A. M.
Tutorial on particle filtering and smoothing: fifteen years later.
In *The Oxford Handbook of Nonlinear Filtering*, D. Crisan and B. Rozovskii, Eds. Oxford: Oxford University Press, 2011, pp. 656–704.

[23] García-Fernández, A., Morelande, M., and Grajal, J.
Bayesian sequential track formation.
IEEE Transactions on Signal Processing, **62**, no. 24 (2014), 6366–6379.

[24] Lindsten, F., Schön, T. B., and Svensson, L.
A non-degenerate Rao-Blackwellised particle filter for estimating static parameters in dynamical models.
In *Proceedings of the 16th IFAC Symposium on System Identification (SYSID)*, Brussels, Belgium, 2012.

[25] Aoki, E. H., Boers, Y., Svensson, L., Mandal, P. K., and Bagchi, A.
SMC methods to avoid self-resolving for online Bayesian parameter estimation.
In *Proceedings of the 15th International Conference of Information Fusion*, Singapore, 2012, 98–105.

[26] Klaas, M., de Freitas, N., and Doucet, A.
Toward practical N^2 Monte Carlo: the marginal particle filter.
In *Proceedings of the 21th Conference Annual Conference on Uncertainty in Artificial Intelligence (UAI-05).* Arlington, VA: AUAI Press, 2005, pp. 308–315.

[27] Bar-Shalom, Y., Li, X. R., and Kirubarajan, T.
Estimation With Applications to Tracking and Navigation.
New York: John Wiley & Sons, 2001, ch. 6.

Edson Hiroshi Aoki received his M.Sc. degree in mechanical-aeronautical engineering from the Instituto Tecnológico de Aeronáutica, Brazil, in 2007 and his Ph.D. degree in applied mathematics from the University of Twente, The Netherlands, in 2013. Together with Dr. Shaohui Foong, Dushyanth Madhavan, and Prof. Yew Long Lo, he received best application award at the 7th International Conference on Soft Computing and Intelligent Systems and the 15th International Symposium on Advanced Intelligent Systems, Kitakyushu, Japan, in 2014.

He has a long history of industry experience with applied statistics, data analytics, and data fusion. He worked at the C4I2SR group of Embraer S.A., Brazil, at Teralytics Pte. Ltd., Singapore, and he is currently a lead data scientist in the Digital Service group of TÜV SÜD Asia Pacific Pte. Ltd., Singapore.

Pranab Kumar Mandal is currently an assistant professor at the University of Twente, The Netherlands. He received his master's degree in statistics from the Indian Statistical Institute in 1992 and his Ph.D. degree in statistics from the University of North Carolina at Chapel Hill in 1997. After receiving his Ph.D., he held a visiting position at Michigan State University and a postdoctoral researcher position at EURANDOM in The Netherlands. His research interests include mathematical statistics and (nonlinear) filtering, in particular, particle filtering with applications to statistical signal processing, system identification, and financial mathematics.

Lennart Svensson was born in Älvängen, Sweden, in 1976. He received his M.S. degree in electrical engineering in 1999 and his Ph.D. degree in 2004, both from Chalmers University of Technology, Gothenburg, Sweden.

He is currently associate professor at the Signal Processing Group at Chalmers University of Technology. His main research interests include machine learning and Bayesian inference in general and nonlinear filtering and tracking in particular.

Yvo Boers received his M.Sc. degree in applied mathematics from Twente University, The Netherlands, in 1994 and his Ph.D. degree in electrical engineering from the Technical University Eindhoven, The Netherlands, in 1999.

He has been with Thales Nederland B.V. since 1999, and as of 2012 he is on permanent disability leave. His research interests are in the areas of detection, (particle) filtering, target tracking, sensor networks, and control. He was an NWO-Casimir research fellow in the period 2008–2011 at the University of Twente in the field of distributed sensor systems.

Together with Hans Driessen, Dr. Boers received the best paper award at the FUSION 2006 conference in Florence, Italy. He has coedited several special issues for different journals. He was an associate editor for the International Society of Information Fusion (ISIF) journal, *Advances on Information Fusion*, during 2005–2008 and an elected member of the Board of Directors for the ISIF, serving the 2011–2013 term.

Arunabha Bagchi was born in Calcutta, India, in September 1947. He received his M.Sc. degree in applied mathematics from Calcutta University in 1969 and his M.S. and Ph.D. degrees in engineering from the University of California, Los Angeles (UCLA), in 1970 and 1974, respectively.

He has been with the University of Twente, The Netherlands, from 1974 to 2012, when he retired, with his last position as professor of applied mathematics. He has been the founder and head of the FELab (Financial Engineering Laboratory) of the University of Twente. His research interest lies mainly in particle filtering, distributed sensor network and financial engineering. He is the author of *Optimal Control of Stochastic Systems* (Prentice Hall International, 1993) and *Stackelberg Differential Games in Economic Models* (LCCIS 64, Springer-Verlag, 1984).

Dr. Bagchi has been associate editor of *IEEE Transactions on Automatic Control and of Automatica*. He was a Fulbright scholar in 1998–1999 and has been visiting professor at UCLA; the State University of New York, Stony Brook; Northeastern University; and the Indian Statistical Institute.