Labeling Uncertainty in Multitarget Tracking

(article starts on next page)

Labelling uncertainty in multi-target tracking

Edson Hiroshi Aoki, Pranab K. Mandal, Lennart Svensson, Yvo Boers, and Arunabha Bagchi

Abstract—In multi-target tracking, the problem of track labelling (assigning labels to tracks) is an on-going research topic. The existing literature, however, lacks an appropriate measure of uncertainty related to the assigned labels which 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 labelling".

In this paper, we specify comprehensively the necessary assumptions for a Bayesian formulation of the Multi-Target Tracking and Labelling (MTTL) problem to be meaningful. We provide a mathematical characterization of the labelling uncertainties with clear physical interpretation. We also propose a novel labelling procedure that can be used in combination with any existing (unlabelled) MTT algorithm to obtain a Bayesian solution to the MTTL problem. One advantage of the resulting solution is that it readily provides the labelling uncertainty measures. Using the mixed labelling phenomenon in the presence of two targets as our test bed we show with simulation results that an *unlabelled* Multi-target Sequential Monte Carlo (M-SMC) algorithm which employs Sequential Importance Resampling (SIR) augmented with our labelling procedure performs much better than its "naive" extension, the *labelled* SIR M-SMC filter.

Index Terms—Multi-target tracking, Track labelling, Labelling error, Random finite set, Sequential Monte Carlo methods.

I. INTRODUCTION

The track labelling problem is perhaps just as old as the multi-target 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, e.g., 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 labelling possibility, is referred to as "mixed labelling" by Boers, Sviestins and Driessen [1]. Being well informed about the labelling uncertainty is of utmost relevance to an end-user when, for instance, a decision involving a target with a particular label is only acceptable 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 fixed number of targets to time varying number of targets due to target birth and death. While these works typically suggest methods for extracting labelled tracks from a multi-target density, they do not attempt to quantify the amount of uncertainty in the assigned labels.

Other recent works [8]–[12] have proposed quantities to be associated with the labelling uncertainty in a MTTL problem. However, the definition of these quantities rely 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 labelling 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 labelling uncertainties should therefore take this into consideration.

Throughout the article we will use as illustrative example the situation of mixed labelling as depicted in Fig. 1. A natural measure to characterize the labelling 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 labelling" 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 labelling"?

In finding the answers we consider the Bayesian formulation of the MTTL problem, based on the concept of labelled Random Finite Sets (RFS) presented in [6]. We propose labelling-related statistics with clear meaning in terms of quantities similar to conditional probability. Our starting point is that the 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 labelled tracks) so that they are not lost during the filter recursions.

We should note here that in the target tracking literature, labels are mostly used 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, though we consider the problem of labelled state estimation, it is not our primary goal to estimate trajectories of the targets. Our main goal is to estimate, at each time, the labelled 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 left-upper corner. One can also think of situations with Air Traffic Control (ATC) callsigns, which are assigned to aircraft by the ATC in order to uniquely identify them. Usually, cooperative aircrafts observed by a secondary radar periodically inform their own assigned callsign to the ATC. However, if, for example, an aircraft's transponder suddenly stops functioning, so does the information about its callsign. In this case, the ATC attempts to associate the previously assigned callsigns to targets observed using the primary radar. 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:

- 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 labelling error* with clear physical interpretation, based on the labelled multi-target posterior density;
- We present a labelling procedure using the proposed labelling uncertainty measures, that 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 article are presented in Section III and Section IV describing the proposed measure of labelling uncertainty and a new method to solve the MTTL problem, respectively. Section V contains the simulation results for labelled tracking of two closely spaced targets. Some conclusions and recommendations are given in Section VI.

NOTATION CONVENTIONS

An upper-case letter (like X) denotes a vector-valued random variable, and its lower-case counterpart (x) denotes a particular realization. An upper-case bold-faced letter (like $\mathbf{X}$) denotes a finite set-valued random variable, and its lower-case 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, $x = \left[x^{(1)}, x^{(2)}\right]$, $\mathbf{x} = \left\{x^{(1)}, x^{(2)}\right\}$.

II. THE BAYESIAN MULTI-TARGET TRACKING AND LABELLING (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. Finally in Section II-C we discuss why mixed labelling such as in Fig. 1 creates 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 *labelled multi-target state* at time k as the random finite set

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

where $X_k^{(i)} = \left[S_k^{(i)}, L_k^{(i)}\right]$ 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, a RFS density function associated with $\mathbf{X}_k$ (referred to as a *labelled RFS density*), must satisfy

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

Examples of closed-form RFS densities that satisfy (1) are the labelled Poisson RFS density, the labelled multi-Bernoulli RFS density and the generalized labelled multi-Bernoulli RFS

density, all described in [6]. Let us denote the corresponding observation as $\mathbf{Z}_k$ (also a 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 labelled state and observation processes $(\mathbf{X}_k, \mathbf{Z}_k)$ to be a first order partially observed Markov process with

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

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

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

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

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

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

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
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*

$$f\left(\mathbf{s}_k|\mathbf{x}_{k-1}\right) = f\left(\mathbf{s}_k|\mathbf{s}_{k-1}\right), \qquad (7)$$
$$and \qquad f\left(\mathbf{z}_k \left| \mathbf{x}_k\right.\right) = f\left(\mathbf{z}_k \left| \mathbf{s}_k\right.\right). \qquad (8)$$

Note that the 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 a MTTL problem any relevant multi-target RFS measurement model that can be used in an (unlabelled) MTT problem. Examples of such multi-target 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 multi-target prior and state transition densities (for the unlabelled MTT problem) can be found in [14, Chapters 13, 14].

B. Assumptions and properties of the MTTL problem

1) Non-ambiguity 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 multi-target tracking scenario.

In the introduction we have equated "labels" to "ATC callsigns", but when they cannot be observed (e.g., when transponders stopped working). Clearly, estimating ATC callsigns for aircraft with non-functioning transponders only makes sense if the aircraft transponders were functioning until some point, i.e., the aircrafts 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 callsign became unobservable.

We view the labelling 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 non-functioning transponder we assume that the model assigns a non-ambiguous label to the target, at its birth. This is possible if we assume that the support for the probability distribution of location for a new-born target does not overlap with that for other (existing or other new born) targets. With the aircrafts and ATC it is true because of the strict regulations.

On the other hand, suppose two targets appear at the same time and 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 labelling*). Then this mixed labelling will persist at all later times ([16, Section IV-C]), rendering the attempt to assign labels to location estimates futile.

Unfortunately, the labelled RFS model described in [6] will always produce total mixed labelling 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 labelling scheme 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 in here. We assume henceforth that there is no ambiguity regarding the labels of appearing targets, which means that total mixed labelling 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 the introduction, 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 which

appeared in the upper-left corner and in the context of ATC callsigns, 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 about the location at birth.

Note that if labels are indeed artificial introduction to the physical model of the unlabelled states, i.e., the unlabelled RFS model, one would expect that given the measurements, the results obtained using *labelled* and *unlabelled* RFS models will be consistent with each other. In particular, one should expect that the filtered distribution of the unlabelled states obtained using the *unlabelled* RFS model (and the observations) should coincide with the distribution of the unlabelled states, calculated by marginalizing the filtered distribution obtained using the *labelled* 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 labelling sub-problems: 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, i.e., estimating the corresponding set of unlabelled target states $\mathbf{S}_k = \left\{ S_k^{(1)}, \ldots, S_k^{(t)} \right\}$ from a sequence of observations $Z^k = (\mathbf{z}_1, \ldots, \mathbf{z}_k)$, completely disregarding the labelling (sub)problem. The tracking problem will involve the following recursions, similar to (4) – (6) but with the unlabelled states.

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

where

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

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

See, e.g., [14, Chapter 14].

Though the Bayesian tracking recursion (9) does not involve any probability distribution of labels or labelled multi-target states, i.e., does not depend on Bayesian labelling, the (sub)problem of Bayesian labelling 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 (unlabelled) tracking and labelling.

C. Why mixed labelling 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 labelling uncertainties. In most of the cases a computationally feasible approximation of the $f\left(\mathbf{x}_k \left| Z^k \right.\right)$ 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 labelling 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 labelling error tends to get underestimated. We explain this further below.

Suppose the Multi-target Sequential Monte Carlo (M-SMC) filter, presented in [17] and [14, Chapter 15] is used to obtain the posterior. These are multi-target versions of the well-known Sequential Importance Resampling Particle Filter (SIR PF) proposed by Gordon, Salmond and Smith [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 multi-target trajectory $(\mathbf{X}_0, \ldots, \mathbf{X}_j)$ to eventually (i.e., at some time step $k > j$) collapse into a single hypothesis $\left(\mathbf{x}_0^*, \ldots, \mathbf{x}_j^*\right)$, leading the particle approximation of the posterior $f(\mathbf{x}_k | Z^k)$ to be biased toward $f\left(\mathbf{x}_k \left| \mathbf{x}_0^*, \ldots, \mathbf{x}_j^*, Z^k \right.\right)$. This degeneracy will definitely have an impact on the filter estimates unless we have the "forgetting condition"

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

Condition (12) is likely going to fail in a situation such as "total mixed labelling", i.e., 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 labelling arose at some point, it will persist at all later times. Consequently, if total mixed labelling already arises at $j' < j$ then it will persist at all later times, including time k. In this case, the true posterior $f\left(\mathbf{x}_k \left| Z^k \right.\right)$ will contain mixed labelling. However, given $\mathbf{x}_j^*$, i.e., assuming unique labels for all targets at time j, $f\left(\mathbf{x}_k \left| \mathbf{x}_j^*, Z^k \right.\right)$ may not have any mixed labelling, thus significantly violating (12).

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

It is easy to see that multi-target tracking techniques based on representing the multi-target 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 multi-target state at time k implicitly assumes

hypotheses on the multi-target 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 the introduction "probability of (in)correct labelling" is mentioned loosely as a possible measure of the labelling uncertainty. 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* labelling of tracks, if the tracks themselves, being estimated values, are not correct.

At this point, many different alternative interpretations can be envisaged, such as e.g. [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. He is interested in having labels assigned to these point estimates. With this viewpoint, we define the labelling probability associated with a point estimate of the *labelled* state to be the conditional probability that the assignment of labels is correct *if* the targets were (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 *labelled* target states, as described in Section II-A. Let the RFS $\mathbf{S}$ be the corresponding *unlabelled* target states. The **labelling probability** associated with a realization $\mathbf{x} = \left\{ \left[s^{(1)}, l^{(1)} \right], \ldots, \left[s^{(t)}, l^{(t)} \right] \right\}$ of $\mathbf{X}$ is defined to be the (conditional) probability that label $l^{(i)}$ is associated with the (unlabelled) state $s^{(i)}$, $i = 1, \ldots, t$, given that (there are t targets present and) $\mathbf{S} = \mathbf{s} := \left\{ s^{(1)}, \ldots, s^{(t)} \right\}$, It is subsequently denoted as $\mathrm{L}(\mathbf{x}|\mathbf{s})$. ∎

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

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

Remark 3.2: The assumption of nonambiguity of initial labels can now be fully described in terms of the labelling probability. In particular, we assume that for any initial unlabelled state $\mathbf{s}_0$, there exists one labelled state $\mathbf{x}_0$ such that $\mathrm{L}(\mathbf{x}_0|\mathbf{s}_0) = 1$. ∎

2) Posterior labelling probability: For Bayesian labelling purposes, we are interested in the posterior version of the labelling probability, i.e., conditioned on all observations up to and including time k, given by

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

where

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

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

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

The following lemma will be useful in our later analysis.
Lemma 3.3: Under condition (8)

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

Proof From (13) we have

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

where the last equality follows from (9).

Lemma 3.3 complies with the interpretation that 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 labelling 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 less labelling uncertainties. On the other hand, once we have reached a stage of *total mixed labelling* [16], where the labeling uncertainties are large everywhere, this is no longer possible and the uncertainties will never be reduced again.

3) Labelling error: We can now use (13) to measure the labelling error in a labelled track estimate.

Definition 3.4: Let $\hat{\mathbf{s}}_k$ be the *unlabelled tracks* corresponding to a set of *labelled tracks* $\hat{\mathbf{x}}_k$. Then the *labelling 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 labelling algorithm for MTTL problems

In this section we propose a labelling procedure that can be combined with existing (unlabelled) 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 it 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 instance, 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 (RBPF) [20] and the marginal particle filter (MPF) [26].

The idea of the RBMPF is to split the state vector into two parts and only handle 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 unlabelled states. We estimate the unlabelled states by a suitable MTT algorithm and calculate the probabilities of the labels, conditiond on the unlabelled 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 labelling procedure

Note that due to the one-sided decoupling property, as described in Section II-B3, we can iteratively obtain a representation of the unlabelled multi-target posterior $f\left(\mathbf{s}_k \left| Z^k\right.\right)$ without any need to concern ourselves about labelling. 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 unlabelled tracking represents the unlabelled multi-target posterior $f\left(\mathbf{s}_k \left| Z^k\right.\right)$ by a set of weighted particles as follows:

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

Otherwise, we should be able to obtain such representation by numerically approximating (e.g. by sampling) the output of the algorithm. Choices of such unlabelled 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 labelling 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 labelled posterior as

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

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

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

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

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

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

$$= \int f(\mathbf{x}_k|\mathbf{x}_{k-1})f\left(\mathbf{x}_{k-1} \left| Z^{k-1}\right.\right) \delta\mathbf{x}_{k-1} \quad (19)$$

$$\approx \sum_{j=1}^{N_P} w_{k-1}(j) \sum_{\mathbf{x}_{k-1} \in \Pi_{k-1}(\mathbf{s}_{k-1}(j))} f(\mathbf{x}_k|\mathbf{x}_{k-1})$$

$$\times \mathrm{L}\left(\mathbf{x}_{k-1} \left|\mathbf{s}_{k-1}(j), Z^{k-1}\right.\right), \quad (20)$$

assuming that we have already computed the labelling probabilities $\mathrm{L}\left(\mathbf{x}_{k-1} \left|\mathbf{s}_{k-1}(i), Z^k\right.\right)$ 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 labelled 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 possible 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

$$\Pi_k\left(\left\{s_k^{(1)}, \ldots, s_k^{(t_k)}\right\}\right)$$

$$\approx \left\{\mathbf{x}_k \left|\mathbf{x}_k = \left\{[s_k^{(1)}, l_k^{(1)}], \ldots, [s_k^{(t_k)}, l_k^{(t_k)}]\right\}\right. \text{ and}\right.$$

$$\exists j, \mathbf{x}_{k-1} \in \Pi_{k-1}(\mathbf{s}_{k-1}(j))$$

$$\left.\text{s.t. } f(\mathbf{x}_k|\mathbf{x}_{k-1})\mathrm{L}\left(\mathbf{x}_{k-1} \left|\mathbf{s}_{k-1}(j), Z^{k-1}\right.\right) > 0\right\}. \quad (21)$$

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

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

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 labelling hypotheses to disappear prematurely. This would, in turn, lead to a sort of degeneracy, which we are trying to prevent at the first place.

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

$$\Pi_0\left(\left\{s_0^{(1)}, \ldots, s_0^{(t_0)}\right\}\right)$$

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

and
$$L(\mathbf{x}_0|\mathbf{s}_0) = f(\mathbf{x}_0)/f(\mathbf{s}_0). \qquad (23)$$

Note that according to the assumption of non-ambiguity 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 $L(\mathbf{x}_0|\mathbf{s}_0) = 1$ for $\mathbf{x}_0 \in \Pi_0(\mathbf{s}_0)$.

B. The algorithm

We now describe the (sub)algorithm for the labelling procedure to be used as a "plug-in" for an (unlabelled) 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 unlabelled posterior. For each particle $\mathbf{s}_k(i)$, the labelling algorithm computes the corresponding labelling probabilities $L\left(\mathbf{x}_k|\mathbf{s}_k(i), Z^k\right)$ using the (unlabelled) particles and labelling 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 $L(\mathbf{x}_0|\mathbf{s}_0(i)) = f(\mathbf{x}_0)/f(\mathbf{s}_0(i))$.

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 labelling probabilities $\overline{L}\left(\mathbf{x}_k|\mathbf{s}_k(i), Z^k\right)$ according to (20).
(L.3) For $\mathbf{x}_k \in \Pi_k(\mathbf{s}_k(i))$, normalize the labelling probabilities as

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

C. Computational cost of the labelling 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 labelling 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 labelling 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 instance, when a target may have disappeared a long time ago, but its corresponding label still maintain 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 labelling 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 labelling 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 labelling 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 labelling procedure, by analysing the effect of adding it as a "plug-in" to a *unlabelled* M-SMC filter. We compare the results to those from *labelled* 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 Section V-C and Section V-D.

A. Metrics for performance evaluation

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

Suppose in a simulation run $\tilde{\mathbf{s}}_k$ is the true unlabelled multi-target states (locations) at time k and $\tilde{\mathbf{x}}_k$ is the true labelled multi-target states. For any given algorithm producing a particle representation of the posterior $f\left(\mathbf{x}_k|Z^k\right)$ at each time k, we calculate the labelling probabilities $L\left(\mathbf{x}_k|\tilde{\mathbf{s}}_k, Z^k\right)$ 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 $f\left(\mathbf{x}_{k-1}|Z^{k-1}\right)$.

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

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

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

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

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

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

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

Recall further that the labelling error in $\hat{\mathbf{x}}_k$, given by, $1 - L\left(\hat{\mathbf{x}}_k|\tilde{\mathbf{s}}_k, Z^k\right)$ 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 $1 - L\left(\hat{\mathbf{x}}_k|\tilde{\mathbf{s}}_k, Z^k\right)$. 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 labelling error as

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

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 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 labelling error, $\left(1 - \text{L} \left(\hat{\mathbf{x}}_k \,\middle|\, \tilde{\mathbf{s}}_k, Z^k \right) \right)$. Note that this variance has two contributors: (1) the variability due to the varying sequence of (unlabelled) state $\mathbf{s}_k$ and observations Z^k, and (2) the variability in the estimation of posterior distribution of the labelled locations. It is however the latter type which is more relevant for us because a high variance would indicate that the calculated $\left(1 - \text{L} \left(\hat{\mathbf{x}}_k \,\middle|\, \tilde{\mathbf{s}}_k, Z^k \right) \right)$, 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 entirely due to the algorithm of obtaining the labelled states. In this analysis, we look at the standard deviation of the calculated labelling error given by

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

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 their 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 multi-target 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 only considered 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

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

The location has the form $S_k^{(i)} = \left[P_x^{(i)}, P_y^{(i)}, V_x^{(i)}, V_y^{(i)} \right]$, where $\left(P_x^{(i)}, P_y^{(i)} \right)$ is the position in Cartesian coordinates x and y and $\left(V_x^{(i)}, V_y^{(i)} \right)$ 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 only considered 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 multi-target predictive density $f \left(\mathbf{x}_k \,\middle|\, Z^{k-1} \right)$ and the posterior density $f \left(\mathbf{x}_k \,\middle|\, Z^k \right)$ are generalized multi-Bernoulli RFS densities, with analytical formulas presented in [6, Section IV].

For all scenarios, we evaluate two MTTL algorithms:

1) A "naive" M-SMC filter which 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/labelling approach which uses a M-SMC filter only for unlabelled tracking, and the labelled "plug-in" described in Section IV-B to calculate the labelling probabilities.

For both the naive and decoupled algorithms, we use 2,000 particles for the scenarios 1, 2 and 3, and 4,000 particles for scenario 4. For both filters, we use blind importance sampling, i.e., 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 labelling 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 a lower average labelling errors for all scenarios. The improvement of using our proposed labelling 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 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 labelling 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 labelling, 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, $\sigma_k^{\bar{\varepsilon}}$, of the labelling errors from a labelled 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 labelling errors with the naive (labelled) M-SMC filter are erroneous (it underestimates severely), there is no point analyzing the variance of these underestimated quantities. So we exclude the naive (unlabelled) 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 labelling error is affected by the degree of target separation (while they move closely-spaced).

The results on the standard deviation, $\sigma_k^{\bar{\varepsilon}}$, of the algorithm-suggested labelling 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 labelled locations. In our case, this means by the (unlabelled) M-SMC filter and subsequently when that is used to calculate the labelling probabilities. Hence a large $\sigma_k^{\bar{\varepsilon}}$ indicates low reliability of the augmented algorithm. It is interesting to notice in Fig. 6 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 counter intuitive because labelling 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 labelling" 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 labelling 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 Bayesian MTTL problem, presented in [6], with a discussion on the additional assumptions needed to keep the target labelling problem meaningful. A mathematical characterization of the labelling uncertainties present in an MTTL solution is provided by defining properly the quantities such as *labelling 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 does not have clear interpretation.

We have also devised a new labelling procedure that can be combined with existing (unlabelled) 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 (unlabelled) M-SMC filter is augmented with the proposed labelling procedure, it performs much better than the naive labelled 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 labelling procedure, by finding more computationally efficient ways to calculate the labelling probabilities. Naturally, it would also be interesting to try the labelling procedure with more complex observation models, such as the track-before-detect observation model.

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,

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

From (30) and the definition of set integral (see, e.g., [6, Proposition 2]) we then have (31) (see bottom of this page). Taking the sum over $l_k^{(i)}$ inside the integral (which is permitted because all the terms are nonnegative) and subsequently, using marginalization and (7), we obtain (32). Interchanging the sum over $l_{k-1}^{(j)}$ and the integral, it follows, from marginalization and the definition of set integral, that

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

which is the Chapman-Kolmogorov equation corresponding to the unlabelled RFS model.

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

$$f(\mathbf{s}_k|Z^k) = \sum_{\substack{i=1 \\ l_k^{(i)} \in \Pi}}^{t} f\left(\left\{[s_k^{(1)}, l_k^{(1)}], \ldots, [s_k^{(t)}, l_k^{(t)}]\right\} \middle| Z^k\right)$$

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

$$f\left(\left\{[s_k^{(1)}, l_k^{(1)}], \ldots, [s_k^{(t)}, l_k^{(t)}]\right\} \middle| Z^{k-1}\right)$$

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

$$= \frac{f(\mathbf{z}_k|\mathbf{s}_k)}{f\left(\mathbf{z}_k \middle| Z^{k-1}\right)} f(\mathbf{s}_k|Z^{k-1})$$

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

REFERENCES

[1] Boers, Y., Sviestins, E., and Driessen, J. N., "Mixed labelling in multi-target particle filtering," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 46, no. 2, pp. 792-802, 2010.

[2] Salmond, D. J., Fisher, D., and Gordon, N. J., "Tracking and identification for closely spaced objects in clutter," in *Proc. European Control Conf.*, 1997, pp. 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 Trans. Signal Process.*, vol. 54, no. 9, pp. 3291-3304, 2006.

[4] Morelande, M., Kreucher, C., and Kastella, K., "A Bayesian approach to multiple target detection and tracking," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 55, no. 5, pp. 1589-1604, 2007.

[5] García-Fernández, A. and Grajal, J., "Multitarget tracking using the Joint Multitrack Probability Density," in *Proc. 12th International Conference on Information Fusion*, Seattle, WA, 2009, pp. 595-602.

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

[7] García-Fernández, A., Grajal, J., and Morelande, M., "Two-layer particle filter for multiple target detection and tracking," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 49, no. 3, pp. 1569-1588, 2013.

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

[9] García-Fernández, Morelande, M., and A., Grajal, J., "Particle filter for extracting target label information when targets move in close proximity," in *Proc. 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 *Proc. 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 *Proc. 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 *Proc. 15th International Conference on Information Fusion*, Singapore, 2012, pp. 2370-2377.

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

[14] Mahler, R., Statistical Multisource-Multitarget Information Fusion. Noorwood, 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, Oct. 2013.

[16] Aoki, E. H., Boers, Y., Svensson, L., Mandal, P. K. and A. Bagchi, "A Bayesian solution to multi-target tracking problems with mixed labelling," Department of Applied Mathematics, University of Twente, Enschede, The Netherlands, Memorandum 2036, July 2014. [Online]. Available: 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 Trans. Aerosp. Electron. Syst.*, vol. 41, no. 4, pp. 1224-1245, 2005.

[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*, vol. 140, no. 2, pp. 107-113, 1993.

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

[20] Andrieu C. and Doucet, A., "Particle filtering for partially observed Gaussian state space models," *J. Royal Stat. Soc. B*, vol. 64, pp. 827-836, 2002.

[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 *Proc. 15th IFAC Symp. 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 University Press, 2011.

[23] García-Fernández, Morelande, M., and Grajal, J., "Bayesian sequential track formation," *IEEE Trans. Signal Process.*, vol. 62, no. 24, pp. 6366-6379, 2014.

[24] Lindsten, F., Schön, T. B. and Svensson, L., "A non-degenerate Rao-Blackwellised particle filter for estimating static parameters in dynamical models," in *Proc. 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 *Proc. 15th International Conference of Information Fusion*, Singapore, 2012, pp. 98-105.

[26] Klaas, M., de Freitas, N., and Doucet, A., "Toward practical N^2 Monte Carlo: the marginal particle filter," in *Proc. 21th Conference Annual*

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

$$f\left(\left\{[s_{k-1}^{(1)}, l_{k-1}^{(1)}], \ldots, [s_{k-1}^{(m)}, l_{k-1}^{(m)}]\right\} \middle| Z^{k-1}\right) d(s_{k-1}^{(1)}, \ldots, s_{k-1}^{(m)}) \quad (31)$$

$$= \sum_{m=0}^{\infty} \frac{1}{m!} \sum_{\substack{l_{k-1}^{(j)} \in \Pi \\ 1 \le j \le m}} \int_{(\mathbb{R}^n)^m} f\left(\mathbf{s}_k \middle| \{s_{k-1}^{(1)}, \ldots, s_{k-1}^{(m)}\}\right)$$

$$f\left(\left\{[s_{k-1}^{(1)}, l_{k-1}^{(1)}], \ldots, [s_{k-1}^{(m)}, l_{k-1}^{(m)}]\right\} \middle| Z^{k-1}\right) d(s_{k-1}^{(1)}, \ldots, s_{k-1}^{(m)}) \quad (32)$$

Conference on Uncertainty in Artificial Intelligence (UAI-05). Arlington, Virginia: 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, NY: John Wiley & Sons, 2001, ch. 6.

Edson Hiroshi Aoki received his M.Sc. degree in Mechanical-Aeronautical Engineering from Instituto Tecnológico de Aeronáutica, Brazil, in 2007 and his Ph.D. degree in Applied Mathematics from 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 15th International Symposium on Advanced Intelligent Systems, Kitakyushu, Japan, 2014.

His industry experience with applied statistics started in 2007, when he was a Product Development Enginner at the C4I2SR group of Embraer S.A., Brazil, focusing on tracking, simulation and other defence applications. Currently he is a Data Scientist at Teralytics Pte. Ltd., Singapore, working on applying statistical methods and data analytics, aided by distributed computing, on processing large datasets for purposes such as urban planning, transportation, consumer targeting and profiling.

Pranab Kumar Mandal is currently an Assistant Professor at the University of Twente. He received the masters degree in Statistics from the Indian Statistical Institute in 1992, and the PhD degree in Statistics from the University of North Carolina at Chapel Hill in 1997. After his PhD he held a visiting position at the Michigan State University (USA) and a postdoctoral reseracher 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 the M.S. degree in electrical engineering in 1999 and the Ph.D. degree in 2004, both from Chalmers University of Technology, Gothenburg, Sweden.

He is currently Associate Professor at the Signal Processing group, again 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 with serving term 2011-2013.

Arunabha Bagchi was born in Calcutta, India, in September 1947. He received the M.Sc. degree in Applied Mathematics from Calcutta University in 1969 and his MS and Ph.D. degrees in Engineering from the University of California, Los Angeles, in 1970 and 1974, respectively.

He has been with the University of Twente from 1974 to 2012, when he retired, with 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 mainly lies 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, SUNY Stony Brook, Northeastern University and the Indian Statistical Institute.
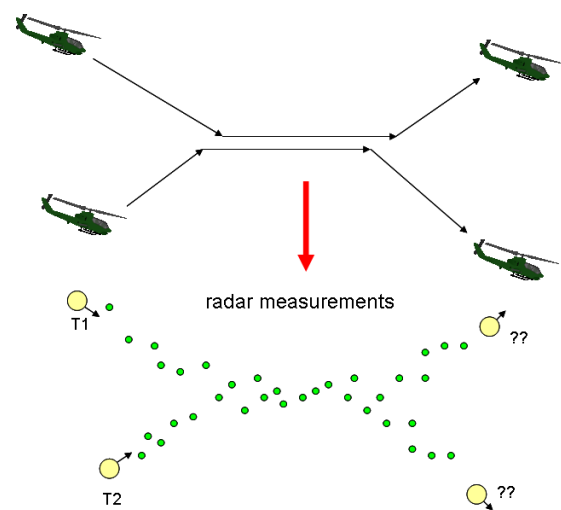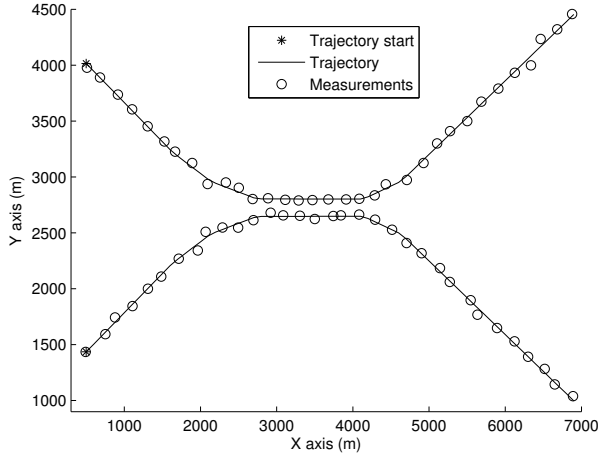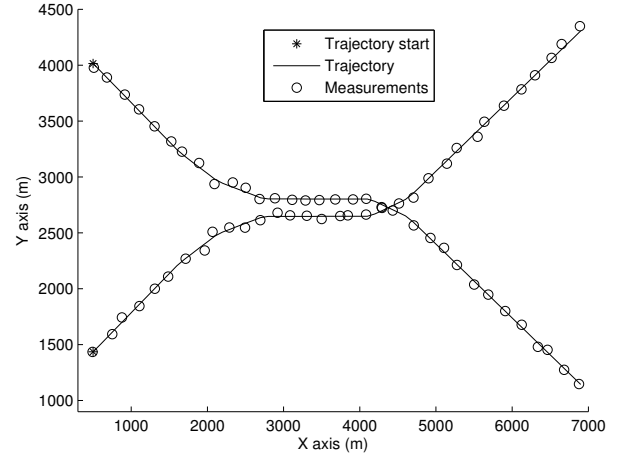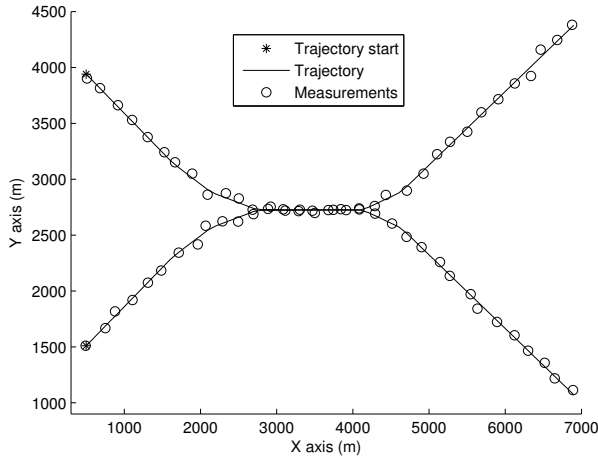

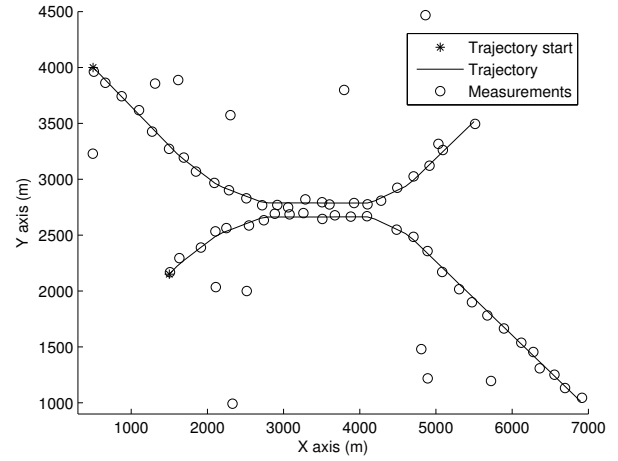Fig. 1. Situation where assignment of labels to tracks is ambiguous

(a) Scenario 1

(b) Scenario 2

(c) Scenario 3

(d) Scenario 4

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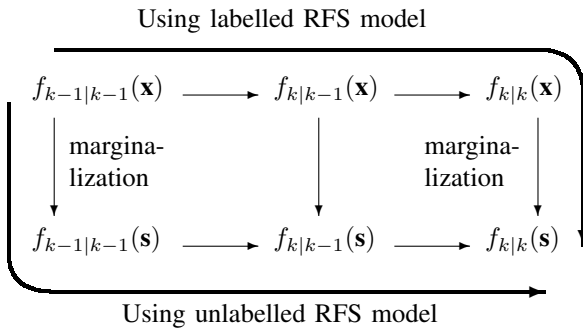

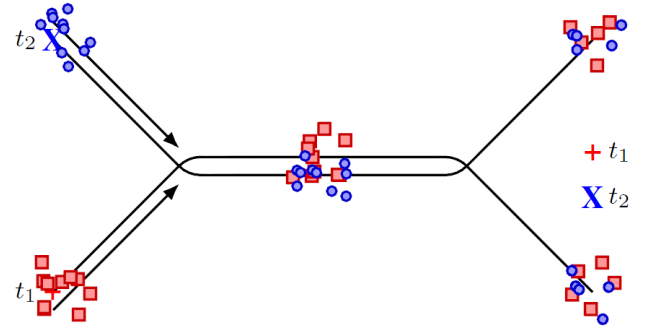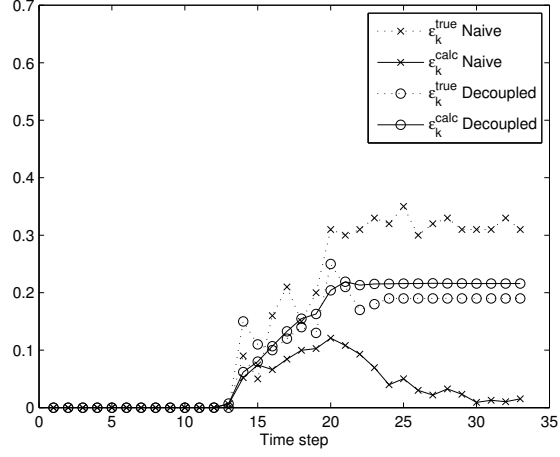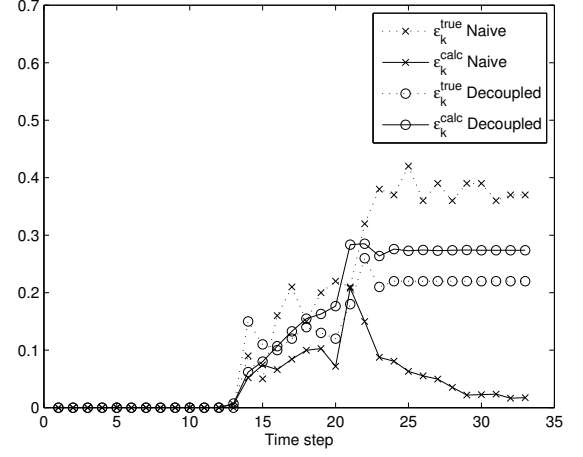Fig. 2. Different ways of obtaining posterior $f_{k|k}(\mathbf{s})$ from prior $f_{k-1|k-1}(\mathbf{x})$
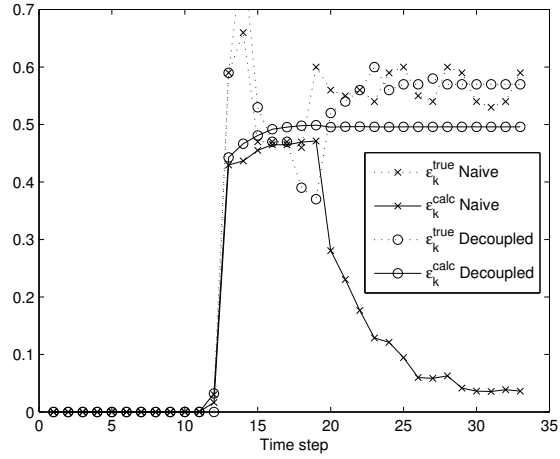

Fig. 3. Particle representation of the multi-target distribution in a situation where mixed labelling occurs (source: [10]). The squares and circles mark the possible locations of each target in terms of particles. "+" and "X" denotes the MMSE estimates.
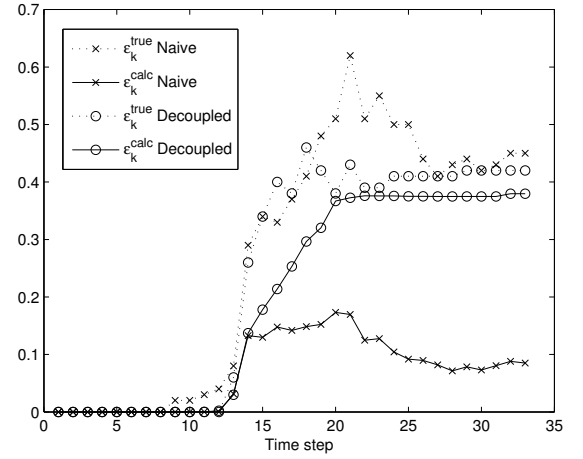
(a) Scenario 1
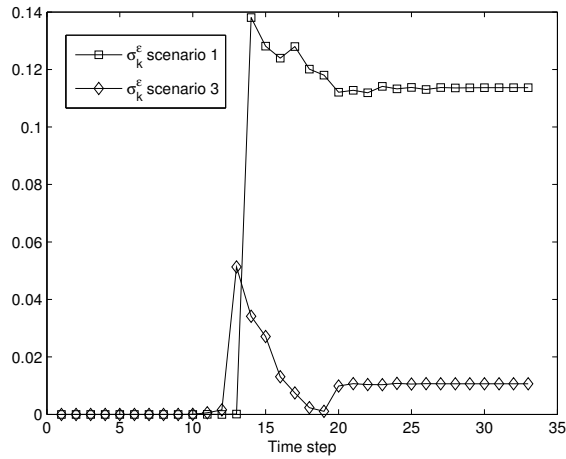
(b) Scenario 2

(c) Scenario 3

(d) Scenario 4

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

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