

Multi-target filtering with linearised complexity

Jeremie Houssineau and Daniel E. Clark

Abstract—An algorithm for the estimation of multiple targets from partial and corrupted observations is introduced based on the concept of partially-distinguishable multi-target system. It combines the advantages of engineering solutions like multiple hypothesis tracking (MHT) with the rigour of point-process-based methods. It is demonstrated that under intuitive assumptions and approximations, the complexity of the proposed multi-target estimation algorithm can be made linear in the number of tracks and linear in the number of observations, while naturally preserving distinct tracks for detected targets, unlike point-process-based methods.

Index Terms—Multi-target tracking, Partial information

I. INTRODUCTION

MULTI-target tracking refers to the problem of estimating the number and state of targets in a dynamical multi-target system via partial and uncertain observations corrupted by detection failures and false alarms. The difficulty of this task is further amplified by the fact that targets in the system appear and disappear at unknown times. The main challenge with multi-target tracking (MTT) is the absence of a priori information about *data association*, i.e. about the association between the received observations and the different targets being estimated. This aspect makes the problem highly combinatorial in nature so that approximations are required for systems made of a large number of targets and/or for which the data association problem is not easily resolved.

The existing solutions in the field of MTT can be divided into two classes of methods. One class consists of “classical” methods [2], [1], [22], [30] that are based on practical generalisations of single-target estimation. Their strength lies in their ability to distinguish the targets and to naturally characterise each of them. The other class of methods comprises approaches based on simple point processes, such as the probability hypothesis density (PHD) filter introduced by [20] or one of the many variants of it [29], [35], [23]. These methods successfully propagate global statistics about the system of interest and integrate false alarms and appearance of targets in a principled way. However, they do not naturally propagate specific information about targets because of the point process assumption of indistinguishability. One of the attempts to overcome this limitation [33], [32] uses marked point processes.

More recently, a new paradigm for modelling systems of targets with uncertain cardinality and state has been introduced in [16], which has been combined in [13] with a novel representation of uncertainty [14], [15] allowing for modelling partial information. This paradigm embeds the concept of

partially-distinguishable systems which is useful in MTT, e.g. for jointly representing never-detected targets (indistinguishable) and previously-detected ones (distinguishable). It also enables the computation of more diverse types of statistics [16] about multi-target systems than with point processes [9], which are helpful, for instance, for sensor management [6]. By considering the usual assumptions of MTT, a first algorithm following from this paradigm has been introduced in [13, Chapt. 3] and detailed in [7], [5] and is referred to as the filter for distinguishable and independent stochastic populations (DISP filter). This filter can be derived without relying on approximations but unsurprisingly displays high computational complexity, similarly to the MHT. In this article, an additional MTT algorithm, called the hypothesised filter for independent stochastic population (HISP filter), is derived and its efficiency is demonstrated on simulated data. The objective is to have recourse to some intuitive approximations in order to lower the complexity.

The structure of the article is as follows: the recursion of the HISP filter is given in section II and in a more practical way in section III. This is followed by an introduction of the main approximation in section IV, showing how the computational complexity is brought down to linear in the number of tracks and in the number of observations. A discussion of the connections with existing algorithms can be found in section V. Details of the implementation are given in section VI and performance is demonstrated on simulated data in section VII.

II. THE HISP FILTER

The specific modelling enabling the derivation of the HISP filter is given in section II-B and section II-C, followed by a detailed presentation of the successive steps of the HISP filter’s recursion in section II-D and section II-E.

Without loss of generality, the time is indexed by the set $\mathbb{T} \doteq \mathbb{N}$, where \mathbb{N} is the set of natural numbers. For any $t \in \mathbb{T}$, the state and observation spaces, denoted $\mathbf{X} = \mathbf{X}^\bullet \cup \{\psi_\sharp, \psi_b\}$ and $\mathbf{Z} = \mathbf{Z}^\bullet \cup \{\phi\}$ respectively, are defined as the union of the spaces of interest \mathbf{X}^\bullet and \mathbf{Z}^\bullet together with the empty states ψ_\sharp and ψ_b as well as the empty observation ϕ . It is common to consider $\mathbf{X}^\bullet = \mathbb{R}^d$ and $\mathbf{Z}^\bullet = \mathbb{R}^{d'}$ with the integers d and d' usually verifying $d > d'$. The empty state ψ_\sharp is attributed to targets that are not in the area of interest, modelled by \mathbf{X}^\bullet at time t , but which might appear in it at any subsequent time. The state ψ_b is attributed to *false-alarm generators* which are objects in the field of view of the sensor that are not of direct interest but which might interfere with the observation of the targets¹. Note that the symbol “ \sharp ” will generally refer to targets whereas “ b ” will refer to the false alarm. The empty observation ϕ is attributed to objects that

¹this construction has been previously proposed in [29]

J. Houssineau is with the Department of Statistics and Applied Probability, National University of Singapore, 117546, SG. Email: stahje@nus.edu.sg

D.E. Clark is with the School of Electrical and Physical Sciences, Heriot-Watt University, Edinburgh EH14 4AS, UK. Email: D.E.Clark@hw.ac.uk.

are not actually observed at a given observation time. The use of empty state/observation is common in the target-tracking literature [3]. Because of the nature of the spaces \mathbf{X} and \mathbf{Z} , the integral $\int f(\mathbf{x})d\mathbf{x}$ of an integrable function f over \mathbf{X} is understood as

$$\int f(\mathbf{x})d\mathbf{x} = \int_{\mathbf{X}^\bullet} f(\mathbf{x})d\mathbf{x} + f(\psi_\#) + f(\psi_b),$$

and, similarly, $\int g(\mathbf{z})d\mathbf{z} = \int_{\mathbf{Z}^\bullet} g(\mathbf{z})d\mathbf{z} + g(\phi)$ for any integrable function g on \mathbf{Z} .

Consider the hypothesis that a distribution p on \mathbf{X} represents a true target. We define ϵ as the random variable on $\{0, 1\}$ describing the fact that this hypothesis is true (1) or false (0) and introduce φ as the virtual state occupied by a target whose existence is assumed by a false hypothesis. The conditional distribution $P(\cdot | \epsilon)$ of the target therefore verifies

$$P(\cdot | \epsilon = 0) = \delta_\varphi \quad \text{and} \quad P(\cdot | \epsilon = 1) = p,$$

which induces a distribution \bar{p} on the extended state space $\bar{\mathbf{X}} \doteq \mathbf{X} \cup \{\varphi\}$ defined as

$$\begin{aligned} \bar{p}(\mathbf{x}) &= P(\mathbf{x} | \epsilon = 0)\Pr(\epsilon = 0) + P(\mathbf{x} | \epsilon = 1)\Pr(\epsilon = 1) \\ &= (1 - w)\delta_\varphi(\mathbf{x}) + wp(\mathbf{x}) \end{aligned} \quad (1)$$

for any $\mathbf{x} \in \bar{\mathbf{X}}$, with $w = \Pr(\epsilon = 1)$ and with δ_φ the Dirac function at φ such that $\delta_\varphi(\mathbf{x})$ equals to 1 if $\mathbf{x} = \varphi$ and 0 otherwise (well defined since φ is assumed to be an isolated point). This extensions is particularly useful for proving results since it combines the *probability of existence* w and the law p into a single distribution \bar{p} on $\bar{\mathbf{X}}$.

A. Sensor modelling

The sensor is understood as a *finite-resolution* sensor which can only generates observations in a finite partition Π_t of \mathbf{Z}^\bullet at time $t \in \mathbb{T}$. With each *observation cell* in Π_t is associated a unique index and the set of all these indices is denoted Z'_t . A set $\{A_t^z : z \in Z_t\}$ of subsets of Π_t with $Z_t \subseteq Z'_t$ is made available by the sensor at each time step $t \in \mathbb{T}$ and corresponds to the actual observation of the targets in the system. It is also helpful to consider another subset $A_t^\phi = \{\phi\}$ that does not correspond to the given data but which will be associated with undetected targets. The set of subsets corresponding to the observations at time t is $\{A_t^z : z \in \bar{Z}_t\}$, where $\bar{Z}_t = Z_t \cup \{\phi\}$.

Although this sensor modelling is not the most usual in the target-tracking literature, many physical sensors, such as radars, sonars or cameras, are actually finite-resolution sensors. This approach will also motivate the introduction of a special form of likelihood in section III-D.

B. Single-target modelling

1) *Transition*: Three (sub-)transition functions q_t^ι from \mathbf{X} to itself indexed by $\iota \in \{\alpha, \pi, \omega\}$ are introduced in order to model the motion as well as the appearance and disappearance of targets between times $t - 1$ and t . These functions are said to be sub-transitions because $q_t^\iota(\cdot | \mathbf{x})$ is not a probability distribution in general but is assumed to verify $\int q_t^\iota(\mathbf{x}' | \mathbf{x})d\mathbf{x}' \leq 1$. Since there is no possible transition between the subset $\mathbf{X}^\bullet \cup \{\psi_\#\}$ describing the targets and the

point ψ_b describing the false-alarm generators, it holds for any $\iota \in \{\alpha, \pi, \omega\}$ that

$$q_t^\iota(\mathbf{x}' | \psi_b) = q_t^\iota(\psi_b | \mathbf{x}) = 0$$

for any $\mathbf{x}, \mathbf{x}' \neq \psi_b$. These transition functions can then be characterised as follows:

i) q_t^α models the appearance of a target, i.e. the transition from $\psi_\#$ to \mathbf{X}^\bullet , so that $q_t^\alpha(\cdot | \mathbf{x}) = \mathbf{0}$ for any $\mathbf{x} \neq \psi_\#$, where $\mathbf{0}$ is the function equal to 0 everywhere. It is assumed that a target cannot appear and disappear during one time step so $q_t^\alpha(\psi_\# | \psi_\#) = 0$.

ii) q_t^π models targets' dynamics, i.e. transitions from \mathbf{X}^\bullet to \mathbf{X}^\bullet or from $\psi_\#$ or ψ_b to themselves, so that for any $\mathbf{x} \in \mathbf{X}^\bullet$ it holds that

$$q_t^\pi(\psi_\# | \mathbf{x}) = 0 \quad \text{and} \quad q_t^\pi(\psi_\# | \psi_\#) = q_t^\pi(\psi_b | \psi_b) = 1.$$

iii) q_t^ω models targets' disappearance, that is a transition from \mathbf{X}^\bullet to $\psi_\#$, so that for any $\mathbf{x}, \mathbf{x}' \in \mathbf{X}^\bullet$ it holds that

$$q_t^\omega(\mathbf{x}' | \mathbf{x}) = 0 \quad \text{and} \quad q_t^\omega(\cdot | \psi_\#) = q_t^\omega(\cdot | \psi_b) = \mathbf{0}.$$

The transition functions q_t^π and q_t^ω are additionally assumed to verify

$$q_t^\omega(\psi_\# | \mathbf{x}) + \int q_t^\pi(\mathbf{x}' | \mathbf{x})d\mathbf{x}' = 1 \quad (2)$$

for any $\mathbf{x} \in \mathbf{X}$, since a target with state \mathbf{x} in \mathbf{X}^\bullet can be either be propagated to \mathbf{X}^\bullet with probability $p_t^\pi(\mathbf{x}) \doteq \int q_t^\pi(\mathbf{x}' | \mathbf{x})d\mathbf{x}'$ or can disappear and be moved to $\psi_\#$. It is convenient to model the appearance of targets with a transition function, e.g. if a rate of appearance is available then the probability of appearance $w_t^\alpha \doteq \int q_t^\alpha(\mathbf{x}' | \psi_\#)d\mathbf{x}'$ will depend on the duration of the considered time step, which is natural for a transition-related quantity. A graphical representation of the transfer of probability mass induced by the different transition functions is shown in fig. 1.

In order to use these transitions in the HISP filter, they have to be extended to include the point φ as well. For any $\iota \in \{\alpha, \pi, \omega\}$ the transition function \bar{q}_t^ι from $\bar{\mathbf{X}}$ to $\bar{\mathbf{X}}$ is defined as

$$\bar{q}_t^\iota(\mathbf{x}' | \mathbf{x}) = \begin{cases} q_t^\iota(\mathbf{x}' | \mathbf{x}) & \text{if } \mathbf{x}, \mathbf{x}' \in \mathbf{X} \\ 1 - \int q_t^\iota(\mathbf{x}'' | \mathbf{x})d\mathbf{x}'' & \text{if } \mathbf{x} \in \mathbf{X} \text{ and } \mathbf{x}' = \varphi \\ 1 & \text{if } \mathbf{x} = \varphi \text{ and } \mathbf{x}' = \varphi. \end{cases} \quad (3)$$

This can be considered as the natural extension of a transition function from \mathbf{X} to $\bar{\mathbf{X}}$ to the corresponding extended spaces. The transition \bar{q}_t^α has a *multiplicity* n_t^α , meaning that it is used exactly n_t^α times. It is assumed that the maximum number of appearing targets n_t^α is larger than the number of observations at time t so that all observations might correspond to appearing targets. One possibility is to consider as many possibly appearing targets as there are resolution cells in the sensor. The transition \bar{q}_t^π and \bar{q}_t^ω have unconstrained multiplicities so that their number simply adapt to the number of targets.²

²This can be made rigorous using the concept of partial information [14], [15] as discussed in [13]

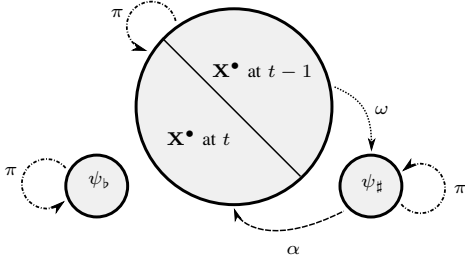


Fig. 1: Relation between the subsets of \mathbf{X} between times $t-1$ and t induced by the transition functions. The lines describe the transfer of probability mass from one subset to another, e.g. q_t^α transfers mass from $\psi_\#$ to \mathbf{X}^\bullet at time t but to and from no other subsets.

2) *Likelihood*: Let ℓ_t^d be a likelihood from $\bar{\mathbf{X}}$ to \mathbf{Z} describing the possible detection of targets and verifying

$$\ell_t^d(\phi | \psi_\#) = \ell_t^d(\phi | \varphi) = 1.$$

The probability for a target with state $\mathbf{x} \in \mathbf{X}^\bullet$ to generate an observation is therefore $p_t^d(\mathbf{x}) \doteq \int_{\mathbf{Z}} \ell_t^d(\mathbf{z} | \mathbf{x}) d\mathbf{z}$. The state ψ_b is the only point of $\bar{\mathbf{X}}$ at which the likelihood ℓ_t^d does not integrate to 1 and we assume instead that $\ell_t^d(\cdot | \psi_b) = \mathbf{0}$ to ensure that false-alarm generator cannot be updated with it.

For each observation cell indexed by $z \in Z_t^z$, a likelihood ℓ_t^z from $\bar{\mathbf{X}}$ to \mathbf{Z} is introduced such that $\ell_t^z(\cdot | \psi_b) = w_t^{b,z} \mathbf{1}_{A_t^z}$, with $\mathbf{1}_A$ the indicator of A , such that $\ell_t^z(\phi | \psi_b) = 1 - w_t^{b,z}$, with $w_t^{b,z} \in [0, 1]$ the probability for a false alarm to be generated in the observation cell z , and such that $\ell_t^z(\cdot | \mathbf{x}) = \mathbf{0}$ for any $\mathbf{x} \neq \psi_b$. Each likelihood in $\{\ell_t^z : z \in Z_t^z\}$ is then assumed to be used exactly once, i.e. each of them has multiplicity one. This model ensures that no more than one false alarm will appear in each resolution cell. If false alarms are independently identically distributed, a single likelihood function can be used to represent them all, with a multiplicity of $|Z_t^z|$. The set $Z_t^z \cup \{d\}$ gathers all the possible indices for likelihoods.

Although this way of describing the dynamics and observation models is more sophisticated than usual, it allows for modelling a closed system where targets merely change state when appearing and disappearing and where all observations come from some objects, either a target or a false-alarm generator.

C. Multi-target modelling

1) *Standard MTT observation model*: It is assumed that each subset A_t^z , $z \in Z_t^z$, is an observation which corresponds to one object, a target or a false-alarm generator. It is also assumed that objects cannot be related to several observations.

2) *Multi-target configuration*: In general, the represented multi-target system might have some distinguishable parts, e.g. the previously-detected targets, and some indistinguishable parts, e.g. the never-detected targets. This aspect makes challenging the expression of the joint law of all these targets. We consider instead *multi-target configurations* of the form $\bar{\mathcal{P}} = \{(\bar{p}_i, n_i)\}_{i \in \mathbb{I}}$ where \bar{p}_i is the law on $\bar{\mathbf{X}}$ for the target(s) with index i and n_i is the associated multiplicity, i.e. there are n_i targets with law \bar{p}_i in the considered configuration. It

is assumed without loss of generality that $\bar{p}_i \neq \bar{p}_j$ for any $i, j \in \mathbb{I}$ such that $i \neq j$. If we assume that $\{\bar{p}_i\}_{i \in \mathbb{I}}$ contains all the possible single-target laws, then varying n_i gives all the possible configurations. It remains to define a suitable index set \mathbb{I} to make sure that the associated multi-target configurations are suitable for the problem at hand. Note that we allow the multiplicity of the laws $\delta_{\psi_\#}$, δ_{ψ_b} and δ_φ to take values in the extended set of natural numbers $\bar{\mathbb{N}} = \mathbb{N} \cup \{\infty\}$.

Since φ does not represent an actual target state, care must be taken when considering aspects related to distinguishability. Indeed, targets are indistinguishable if and only if their respective laws are equal when conditioned on the targets being in \mathbf{X} , i.e. the targets represented by the random variables X and X' with respective laws \bar{p} and \bar{p}' on $\bar{\mathbf{X}}$ are indistinguishable if and only if

$$\bar{p}(\mathbf{x} | X \in \mathbf{X}) = \bar{p}'(\mathbf{x} | X' \in \mathbf{X}) \quad (4)$$

for all $\mathbf{x} \in \mathbf{X}$.

Example 1. If \mathbb{I} is a singleton then the multi-target configuration $\bar{\mathcal{P}} = \{(\bar{p}, N)\}$ is equivalent to the symmetric multi-target law P on $\bar{\mathbf{X}}^N$ characterised for any $\mathbf{x}_1, \dots, \mathbf{x}_N \in \bar{\mathbf{X}}$ by

$$\begin{aligned} P(\mathbf{x}_1, \dots, \mathbf{x}_{n_i}) &= \bar{p}(\mathbf{x}_1) \dots \bar{p}(\mathbf{x}_N) \\ &= \prod_{i=1}^N \left((1-w)\delta_\varphi(\mathbf{x}_i) + wp(\mathbf{x}_i) \right) \end{aligned}$$

In this specific case, P is symmetrical in its arguments and is related to the law of a multi-Bernoulli random finite set (RFS). Alternatively, if $\mathbb{I} = \{1, \dots, N\}$ and $n_i = 1$ for all $i \in \mathbb{I}$ then $\bar{\mathcal{P}}$ is equivalent to the multi-target law Q on $\bar{\mathbf{X}}^N$ characterised for any $\mathbf{x}_1, \dots, \mathbf{x}_N \in \bar{\mathbf{X}}$ by

$$\begin{aligned} Q(\mathbf{x}_1, \dots, \mathbf{x}_N) &= \bar{p}_1(\mathbf{x}_1) \dots \bar{p}_N(\mathbf{x}_N) \\ &= \prod_{i=1}^N \left((1-w_i)\delta_\varphi(\mathbf{x}_i) + w_i p_i(\mathbf{x}_i) \right), \end{aligned}$$

where each law \bar{p}_i has been expressed as in eq. (3), which is similar to the law of a labelled multi-Bernoulli RFS [33]. The law P does not distinguish any target whereas Q distinguishes all targets. The objective is then to use other multi-target configurations which allow for partial distinguishability.

Remark 1. Even if labelled RFS were generalised to permit several labels to be equal, the RFS representation would not allow for a representation of indistinguishable targets almost surely at an isolated point state such as $\psi_\#$ since a RFS is a simple point process and therefore does not allow for multiplicity strictly greater than 1 (both the state and the label would be equal in the case of indistinguishable targets at point $\psi_\#$). Handling non-simple, partially-labelled point processes is a complex task, and multi-target configurations alleviate much of this complexity by appropriately representing the hierarchical structure of partially-distinguishable multi-target system. A more general discussion about the different approaches for target tracking is provided in section V.

3) *Observation path*: With the standard MTT observation model, the targets are made distinguishable as soon as they are detected since each observation corresponds to no more than one target. Also, one of the main characteristics of the induced single-target laws is the corresponding *observation path* (the law is basically the Bayesian posterior law of one target state given the associated observations). For this reason, we consider the set $\bar{\mathbb{O}}_t$, defined as

$$\bar{\mathbb{O}}_t \doteq \bar{Z}_0 \times \cdots \times \bar{Z}_t,$$

so that a sequence of observation $\mathbf{o}_t \in \bar{\mathbb{O}}_t$ takes the form $\mathbf{o}_t = (\phi, \dots, \phi, z_{t_+}, \dots, z_{t_-}, \phi, \dots, \phi)$ with t_+ and t_- the time of appearance and disappearance of the target, with $z_t \in \bar{Z}_t$ for any $t \in \{t_+, \dots, t_-\}$ and with $z_{t_+}, z_{t_-} \neq \phi$. The empty observation path $(\phi, \dots, \phi) \in \bar{\mathbb{O}}_t$ is denoted ϕ_t .

4) *Simplifying procedures*: In order to reduce the number of terms in the considered multi-target configurations, two simplifying procedures are considered. They rely on *mixing* two or more elements of a given multi-target configuration $\bar{\mathcal{P}}$ on $\bar{\mathbf{X}}$, say the elements with index i and j in the associated index set \mathbb{I} , by defining a new index k based on i and j with $n_k = n_i + n_j$ and with the single-target law \bar{p}_k defined as

$$\bar{p}_k(\mathbf{x}) = \frac{n_i \bar{p}_i(\mathbf{x}) + n_j \bar{p}_j(\mathbf{x})}{n_i + n_j}.$$

for any \mathbf{x} in $\bar{\mathbf{X}}$. The index set after mixing i and j is then defined as $(\mathbb{I} \setminus \{i, j\}) \cup \{k\}$. The two simplifying procedures can then be formalised as

S.1 Mixing of the appearing targets with the never-detected ones

S.2 Mixing with $\delta_{\psi_\#}$ of the laws of never-detected targets that disappeared

Simplification S.1 fixes the number of groups of indistinguishable targets to one: the group of never-detected targets. The number of groups would otherwise grow by one every time step, with the appearance of additional targets. Simplification S.1 is justified by the fact that the distribution of appearing targets is often uninformative and constant in time, so that the difference between the appearing targets at time t and the ones who appeared at time $t-1$ and who have not been detected at time t can often be neglected. Similarly, Simplification S.2 is associated with the fact that never-detected disappeared targets are often irrelevant in MTT. Their single-target law is the form $w\delta_{\psi_\#} + (1-w)\delta_\varphi$ for some $w \in [0, 1]$ that is often close to 0, so that forcing the mixing with the single-target law $\delta_{\psi_\#}$ only incurs a small information loss.

5) *Indexing of single-target laws*: It is assumed that the only source of specific information at time $t \in \mathbb{T}$ lies in the observations made before time t . The targets' laws after prediction can then be indexed by the set $\mathbb{I}_{t|t-1}^\#$ of triplets $(\#, T, \mathbf{o})$ such that T is either the empty set or a non-empty interval $[\cdot, t_-]$ of $\{0, \dots, t\}$ with unknown starting time and such that the observation time \mathbf{o} in \mathbb{O}_{t-1} satisfies $\mathbf{o}_{t'} \neq \phi \implies t' \leq t_-$. An interval $[\cdot, t_-]$ is considered instead of a more standard given interval $\{t_+, \dots, t_-\}$ because newly appeared targets will be mixed with never-detected targets which implies that the time of appearance is forgotten. Although writing

$t \in [\cdot, t_-]$ is equivalent to $t \leq t_-$, the expression $[\cdot, t_-]$ is preferred since it allows for using \emptyset for targets that did not appear yet. The index set $\mathbb{I}_{t|t-1}^\#$ can be interpreted as follows: predicted single-target laws are distinguished by their interval of presence in \mathbf{X}^\bullet up to time t and by their observation path in $\bar{\mathbb{O}}_{t-1}$. Simplification S.2 implies that the elements of the form $(\#, [\cdot, t_-], \phi_t)$ with $t_- < t$ are not included in $\mathbb{I}_{t|t-1}^\#$. The set $\mathbb{I}_t^\#$ indexing single-target laws at time t after the update is defined similarly but with the observation path in $\bar{\mathbb{O}}_t$, i.e. with the observation up to time t , rather than up to time $t-1$.

False-alarm generators are at ψ_b almost-surely and are assumed to give inconsistent observations, which corresponds to indices of the form $(b, \emptyset, \mathbf{o})$ with \mathbf{o} containing a single non-empty observation. The set indexing the false-alarm generators at time t after update is then

$$\mathbb{I}_t^b = \bigcup_{0 \leq t' \leq t} \{(b, \emptyset, (\phi, \dots, \phi, z_{t'}, \phi, \dots, \phi)) : z_{t'} \in \bar{Z}_{t'}\}.$$

The sets indexing the multi-target configurations after prediction and update are then defined as

$$\mathbb{I}_{t|t-1} \doteq \mathbb{I}_{t|t-1}^\# \cup \mathbb{I}_{t-1}^b \quad \text{and} \quad \mathbb{I}_t \doteq \mathbb{I}_t^\# \cup \mathbb{I}_t^b.$$

The single-target laws, e.g. after the update, can then be indexed as follows:

- a target that is still present at time t and that has only been detected during the current time step has index $(\#, [\cdot, t], (\phi, \dots, \phi, z))$ for some $z \in \bar{Z}_t$,
- a target that has not appeared yet has index $(\#, \emptyset, \phi_t)$ (which would also be an element of \mathbb{I}_t^b in the absence of the symbols $\#$ and b), and
- to each index $(\#, T, \mathbf{o}) \in \mathbb{I}_{t|t-1}^\#$, corresponds $|\bar{Z}_t|$ indices in $\mathbb{I}_t^\#$, that is one index $(\#, T, \mathbf{o} \times z)$ for each $z \in \bar{Z}_t$, with $\mathbf{o} \times z$ denoting the concatenation of the sequence \mathbf{o} with the element z

It is useful to introduce two more symbols to represent some specific parts of the multi-target system: the targets with an index $(\#, T, \mathbf{o}) \in \mathbb{I}_t^\#$ that

- (m) have been previously detected, so that $\mathbf{o} \neq \phi$,
- (u) are in the state space but have never been detected, so that $T \ni t$ and $\mathbf{o} = \phi$,

Because of Simplification S.1, there is a single element in the set \mathbb{I}_t^u , which we denote \mathbf{i}_t^u (similarly $\mathbb{I}_{t|t-1}^u = \{\mathbf{i}_{t|t-1}^u\}$). For the sake of compactness, single-target laws and multiplicities with index \mathbf{i}_t^u or $\mathbf{i}_{t|t-1}^u$ will simply bear the superscript u .

We will first consider multi-target configurations on $\bar{\mathbf{X}}$ in the following sections, before reformulating some of the equations of the filter with multi-target configurations on \mathbf{X} . In both cases, an element of the underlying collections will be referred to as an *hypothesis*.

The following multi-target configuration represents the false-alarm generator, the yet-to-appear targets as well as some erroneous hypotheses:

$$\bar{\mathcal{P}}_t^{\psi, \varphi} \doteq \{(\delta_{\psi_b}, n_t^i)\}_{i \in \mathbb{I}_t^b} \cup \{(\delta_{\psi_\#}, n_{\psi_\#}), (\delta_\varphi, n_\varphi)\},$$

where $n_{\psi_b} = \infty$ and $n_t^i = n_{\psi_b} = \infty$ when $i = (b, \emptyset, \phi_t)$. This means that there is an infinite number of targets that

are not currently in the area of interest (the state space \mathbf{X}^\bullet) but which might enter it at a later time and there also are an infinite number of potential false-alarm generators (but the likelihoods corresponding to the actual generation of false alarms come in finite number). Setting $n_{\psi_\#}$ and n_{ψ_b} equal to ∞ is a simplifying assumption as it implies that these cardinalities will never change during the scenario since subtracting any finite number to account for target appearance or the generation of false alarms will not affect them. The law δ_φ only serves as a representation of erroneous single-target laws, so that the value of n_φ is irrelevant and its time dependency is omitted. The initial multi-target configuration $\bar{\mathcal{P}}_0$ is defined as

$$\bar{\mathcal{P}}_0 = \bar{\mathcal{P}}_{0|-1} \doteq \{(\bar{p}_0^u, n_0^u)\} \cup \bar{\mathcal{P}}_0^{\psi, \varphi}$$

with $n_0^u = n_0^\alpha$, $\bar{p}_0^u(\mathbf{x}) = \bar{q}_0^\alpha(\mathbf{x} | \psi_\#)$ for any $\mathbf{x} \in \mathbf{X}$.

D. Prediction

The number of never-detected targets at time $t-1$ after the update is denoted n_{t-1}^u . The multi-target configuration $\bar{\mathcal{P}}_{t-1}$ after the update at time $t-1$ is assumed to have the following form:

$$\bar{\mathcal{P}}_{t-1} = \{(\bar{p}_{t-1}^i, 1)\}_{i \in \mathbb{I}_{t-1}^m} \cup \{(\bar{p}_{t-1}^u, n_{t-1}^u)\} \cup \bar{\mathcal{P}}_{t-1}^{\psi, \varphi}. \quad (5)$$

Let $\sigma_{t|t-1} : \mathbb{I}_{t-1} \times \{\alpha, \pi, \omega\} \rightarrow \mathbb{I}_{t|t-1}$ be the one-to-one mapping giving an index in $\mathbb{I}_{t|t-1}$ for each pair of indices in $\mathbb{I}_{t-1} \times \{\alpha, \pi, \omega\}$ corresponding to an hypothesis at time $t-1$ and a transition function. This mapping can be defined as

$$\sigma_{t|t-1} : ((a, T, \mathbf{o}), \iota) \mapsto \begin{cases} (a, T \cup \{t\}, \mathbf{o}) & \text{if } \iota \in \{\alpha, \pi\} \\ (a, T, \mathbf{o}) & \text{otherwise,} \end{cases}$$

with a begin either $\#$ or b , since the presence of an target is extended to time t with the transitions q_t^α and q_t^π only. The index $\sigma_{t|t-1}(i, \iota)$ is simply the index of the hypothesis obtained when predicting the previous hypothesis i with the transition ι . The prediction can now be expressed as in the following theorem, which proof is given in the appendix, and which is based on the following approximation.

A.1 The hypotheses formed by predicting each single-target law in $\mathbb{I}_{t-1}^\#$ with the transitions q_t^π and q_t^ω are independent.

Theorem 1. Under A.1, the multi-target configuration $\bar{\mathcal{P}}_{t|t-1}$ after prediction to time t is characterised by

$$\bar{\mathcal{P}}_{t|t-1} = \{(\bar{p}_{t|t-1}^i, 1)\}_{i \in \mathbb{I}_{t|t-1}^m} \cup \{(\bar{p}_{t|t-1}^u, n_{t|t-1}^u)\} \cup \bar{\mathcal{P}}_{t|t-1}^{\psi, \varphi},$$

where $n_t^u = n_{t-1}^u + n_t^\alpha$ is the predicted number of never-detected targets and where the predicted single-target law $\bar{p}_{t|t-1}^i$ is defined on $\bar{\mathbf{X}}$ for any index $i \in \mathbb{I}_{t|t-1}$ with $(\mathbf{k}, \iota) = \sigma_{t|t-1}^{-1}(i)$ as

$$\bar{p}_{t|t-1}^i(\mathbf{x}) = \frac{n_{t-1}^u \int \bar{q}_t^\pi(\mathbf{x} | \mathbf{x}') \bar{p}_{t-1}^u(\mathbf{x}') d\mathbf{x}' + n_t^\alpha \bar{q}_t^\alpha(\mathbf{x} | \psi_\#)}{n_{t-1}^u + n_t^\alpha}$$

if $\mathbf{k} = \mathbf{i}_{t|t-1}^u$ and as $\bar{p}_{t|t-1}^i(\mathbf{x}) = \int \bar{q}_t^\iota(\mathbf{x} | \mathbf{x}') \bar{p}_{t-1}^{\mathbf{k}}(\mathbf{x}') d\mathbf{x}'$ otherwise.

Note that $\sigma_{t|t-1}^{-1}(i)$ simply recovers the indices of the previous hypothesis and transition that lead to i . It appears that the predicted configuration $\bar{\mathcal{P}}_{t|t-1}$ takes the same form as the posterior configuration $\bar{\mathcal{P}}_{t-1}$ at time $t-1$ expressed in eq. (5). Although the single-target laws of disappeared targets are not very informative, they can be useful in practice since the scalar $\int \bar{q}_t^\omega(\psi_\# | \mathbf{x}) \bar{p}_{t-1}^{\mathbf{k}}(\mathbf{x}) d\mathbf{x}$ gives the credibility of the hypothesis that the target with index \mathbf{k} at time $t-1$ disappeared between $t-1$ and t .

E. Update

Let $\sigma_t : \mathbb{I}_{t|t-1} \times (Z_t' \cup \{d\}) \times (Z_t' \cup \{\varphi\}) \rightarrow \mathbb{I}_t$ be the one-to-one mapping describing the connections of all possible combinations of prior single-target law, likelihood and observation with the indices in \mathbb{I}_t be defined as

$$\sigma_t : ((a, T, \mathbf{o}), s, z) \mapsto (a, T, \mathbf{o} \times z).$$

Note that, by construction, many combinations will result in a posterior with all its probability mass on φ , such as with applying the likelihood ℓ_t^d to a target with state ψ_b . To facilitate the statement of the update, let

$$L_t^{s,z}(\mathbf{x}) = \int \mathbf{1}_{A_t^z}(\mathbf{z}) \ell_t^s(\mathbf{z} | \mathbf{x}) d\mathbf{z}$$

be the probability that an object with state $\mathbf{x} \in \mathbf{X}$ has generated an observation in the cell $z \in Z_t'$ under the likelihood ℓ_t^s with $s \in Z_t' \cup \{d\}$. Similarly, we denote $p_t^{\mathbf{k},s}(z)$ as the marginal likelihood of the observation z when the prior $\bar{p}_{t|t-1}^{\mathbf{k}}$, $\mathbf{k} \in \mathbb{I}_{t|t-1}$, is restricted to \mathbf{X} , that is

$$p_t^{\mathbf{k},s}(z) = \int_{\mathbf{X}} L_t^{s,z}(\mathbf{x}) \bar{p}_{t|t-1}^{\mathbf{k}}(\mathbf{x}) d\mathbf{x},$$

and $\bar{p}_t^{\mathbf{k},s}(z) = p_t^{\mathbf{k},s}(z) + L_t^{s,z}(\varphi) \bar{p}_{t|t-1}^{\mathbf{k}}(\varphi)$ as the extended marginal likelihood. Note that the value of s is known once \mathbf{k} is fixed, so that it is possible to omit it and write $\bar{p}_t^{\mathbf{k}}(z)$ instead of $\bar{p}_t^{\mathbf{k},s}(z)$. The next theorem, proved in the appendix, relies on the following approximation.

A.2 The hypotheses formed by updating each single-target law in $\mathbb{I}_{t|t-1}$ by each observation in \bar{Z}_t are independent.

Theorem 2. Under A.2, the multi-target configuration $\bar{\mathcal{P}}_t$ after the update at time t is characterised by

$$\bar{\mathcal{P}}_t = \{(\bar{p}_t^i, 1)\}_{i \in \mathbb{I}_t^m} \cup \{(\bar{p}_t^u, n_t^u)\} + \bar{\mathcal{P}}_t^{\psi, \varphi}, \quad (6)$$

where the posterior single-target law \bar{p}_t^i with index $i \in \mathbb{I}_t^\#$, denoting $(\mathbf{k}, s, z) \doteq \sigma_t^{-1}(i)$, is characterised for any $\mathbf{x} \in \mathbf{X}$ by

$$\bar{p}_t^i(\mathbf{x}) = w_t^i \frac{L_t^{s,z}(\mathbf{x}) \bar{p}_{t|t-1}^{\mathbf{k}}(\mathbf{x})}{p_t^{\mathbf{k},s}(z)},$$

and by $\bar{p}_t^i(\varphi) = 1 - w_t^i$, where w_t^i is defined as

$$w_t^i = \frac{w_{\text{ex}}^{\mathbf{k},z} p_t^{\mathbf{k},s}(z)}{\sum_{z' \in \bar{Z}_t} w_{\text{ex}}^{\mathbf{k},z'} \bar{p}_t^{\mathbf{k},s}(z')} \quad (7a)$$

or, equivalently, as

$$w_t^i = \frac{w_{\text{ex}}^{\mathbf{k},z} p_t^{\mathbf{k},s}(z)}{\sum_{\mathbf{k}' \in \mathbb{I}_{t|t-1}} w_{\text{ex}}^{\mathbf{k}',z} \bar{p}_t^{\mathbf{k}'}(z)}, \quad (7b)$$

and where $w_{\text{ex}}^{\mathbf{k},z}$ is the multi-target marginal likelihood corresponding to the update by the observations in $Z_t \setminus \{z\}$ of the targets with index in

$$\mathbb{I}_{t|t-1}^{\mathbf{k}} \doteq (\mathbb{I}_{t|t-1}^{\text{m}} \setminus \{\mathbf{k}\}) \cup \{u, b\},$$

and with multiplicity $n_{t|t-1}^u - \mathbf{1}_u(\mathbf{k})$ for the never-detected targets.

Less formally, the marginal likelihood $w_{\text{ex}}^{\mathbf{k},z} \in [0, 1]$ can be understood as the assessment of the compatibility between the predicted laws and the collection of observations at the current time excluding the/a target with index \mathbf{k} and the observation z . Note that we consider $Z_t \setminus \{z\}$ rather than $\bar{Z}_t \setminus \{z\}$ in the definition of $w_{\text{ex}}^{\mathbf{k},z}$ since the empty observation ϕ might be associated with an arbitrary number of targets, i.e. it is not because one target is not detected that other targets have to be detected.

An important feature of the HISP filter can already be highlighted: an a posteriori probability of detection failure can be computed through eq. (7a) when $z = \phi$. Also, an a posteriori probability for an observation to be a false alarm is obtained when considering $\mathbf{k} \in \mathbb{I}_{t|t-1}^b$. Theorem 2 reveals the fact that the collection of marginalised single-target target laws $\{\bar{p}_t^i\}_{i \in \mathbb{I}_t^{\#}}$ can be seen as single-target filters in interaction, where the weights of the filters are $\{w_t^i\}_{i \in \mathbb{I}_t^{\#}}$.

III. SUMMARY OF THE HISP FILTER RECURSION

The extension of the probability distributions and transition functions to $\bar{\mathbf{X}}$ makes the proofs of the results easier, however, from an implementation viewpoint, an hypothesis is best described by a probability distribution p on \mathbf{X} together with a scalar $w \in [0, 1]$ describing the probability of the corresponding hypothesis to represent a target. This is the approach taken in [8], which is given here for completeness.

For any $i \in \mathbb{I}_t^{\#}$, the scalar $w_t^i \in [0, 1]$ and the probability distribution p_t^i on \mathbf{X} are defined by

$$w_t^i \doteq \int_{\mathbf{X}} \bar{p}_t^i(\mathbf{x}') d\mathbf{x}' \quad \text{and} \quad p_t^i(\mathbf{x}) = \bar{p}_t^i(\mathbf{x}) / w_t^i,$$

for any $\mathbf{x} \in \mathbf{X}$. The multi-target configuration can now be equivalently defined as $\mathcal{P}_{t-1} = \{(p_{t-1}^i, w_{t-1}^i, n_{t-1}^i)\}_{i \in \mathbb{I}_{t-1}}$.

A. Approximations

Approximations **A.1** and **A.2** are of the same nature: they allow hypotheses for *co-existing*, i.e. for being part of a joint multi-target law, when it should not be the case. For instance, a target cannot simultaneously remain in \mathbf{X}^\bullet and disappear, and it cannot generate several observations under the standard MTT assumptions. However, allowing for the co-existence of all hypotheses has for consequence the reduction of the number of multi-target configurations to one. There are two aspects in the HISP filter that mitigate the effects of these approximations: a) although all hypotheses are considered at once, each of them is allowed to be false (represented by φ), and b) the propagation of the distinguishability enables an efficient track extraction to be devised, in which the standard assumptions of MTT can be enforced.

B. Prediction

The hypothesis with index $\mathbf{k} = (\#, [\cdot, t-1], \mathbf{o}) \in \mathbb{I}_{t-1}^{\#}$ can be propagated with the transition q_t^π and yield the hypothesis $i = (\#, [\cdot, t], \mathbf{o}) \in \mathbb{I}_{t|t-1}^{\#}$ with weight and law

$$w_{t|t-1}^i = w_{t-1}^{\mathbf{k}} \iint q_t^\pi(\mathbf{x} | \mathbf{x}') p_{t-1}^{\mathbf{k}}(\mathbf{x}') d\mathbf{x}' d\mathbf{x}$$

$$p_{t|t-1}^i(\mathbf{x}) \propto \int q_t^\pi(\mathbf{x} | \mathbf{x}') p_{t-1}^{\mathbf{k}}(\mathbf{x}') d\mathbf{x}',$$

or it can disappear when propagated with the transition q_t^ω , in which case the index in $\mathbb{I}_{t|t-1}$ of the corresponding hypothesis remains equal to \mathbf{k} and the associated weight and law are

$$w_{t|t-1}^{\mathbf{k}} = w_{t-1}^{\mathbf{k}} \int q_t^\omega(\psi_{\#} | \mathbf{x}') p_{t-1}^{\mathbf{k}}(\mathbf{x}') d\mathbf{x}'$$

$$p_{t|t-1}^{\mathbf{k}} = \delta_{\psi_{\#}}.$$

For never-detected targets, it holds that $n_{t|t-1}^u = n_{t-1}^u + n_{t-1}^\alpha$ and

$$p_{t|t-1}^u(\mathbf{x}) = \frac{n_{t-1}^u \int q_t^\pi(\mathbf{x} | \mathbf{x}') p_{t-1}^{\mathbf{k}}(\mathbf{x}') d\mathbf{x}' + n_{t-1}^\alpha p_t^\alpha(\mathbf{x})}{n_{t-1}^u + n_{t-1}^\alpha}$$

$$w_{t|t-1}^u = \frac{n_{t-1}^u w_{t-1}^{\mathbf{k}} + n_{t-1}^\alpha w_t^\alpha}{n_{t-1}^u + n_{t-1}^\alpha},$$

with $p_t^\alpha = q_t^\alpha(\cdot | \psi_{\#})$ the distribution of the appearing targets.

C. Update

The hypothesis with index $\mathbf{k} = (\#, [\cdot, t-1], \mathbf{o}) \in \mathbb{I}_{t-1}^{\text{m}}$ can be updated with the observation $z \in \bar{Z}_t$ via the likelihood ℓ_t^d and yield the hypothesis $i = (\#, [\cdot, t], \mathbf{o} \times z)$ with weight given by eq. (7) and law

$$p_t^i(\mathbf{x}) \propto L_t^{d,z}(\mathbf{x}) p_{t|t-1}^{\mathbf{k}}(\mathbf{x}).$$

The marginal likelihoods $p_t^{\mathbf{k},d}(z)$ and $\bar{p}_t^{\mathbf{k},d}(z)$ appearing in eq. (7) can be expressed in this case as

$$p_t^{\mathbf{k},d}(z) = w_{t|t-1}^{\mathbf{k}} \int L_t^{d,z}(\mathbf{x}) p_{t|t-1}^{\mathbf{k}}(\mathbf{x}) d\mathbf{x}$$

$$\bar{p}_t^{\mathbf{k},d}(z) = p_t^{\mathbf{k},d}(z) + \mathbf{1}_\varphi(z)(1 - w_{t|t-1}^{\mathbf{k}}).$$

Other hypotheses can be specified from theorem 2 in the same way such as the ones corresponding to false alarms. For instance, the posterior probability for the observation $z \in \bar{Z}_t$ to be a false alarm can be computed from eq. (7b) as

$$w_t^i = \frac{w_{\text{ex}}^{b,z} w_t^{b,z}}{\sum_{\mathbf{k}' \in \mathbb{I}_{t|t-1}} w_{\text{ex}}^{\mathbf{k}',z} \bar{p}_t^{\mathbf{k}'}(z)},$$

with $i = (b, \emptyset, \phi_{t-1} \times z)$. This quantity is not usually computed directly, but it will appear to be crucial in the proposed track-extraction procedure.

D. Alternative sensor modelling

In some situations, it is simpler to assume that observations can be any point \mathbf{z} of the observation space \mathbf{Z}^\bullet rather than a resolution cell. For instance, the shape of the resolution cells can be approximated by a Gaussian function of the form

$$f_{\mathbf{z}}(\mathbf{z}') = \exp\left(-\frac{1}{2}(\mathbf{z}' - \mathbf{z})^T \Sigma^{-1}(\mathbf{z}' - \mathbf{z})\right),$$

where Σ approximates the extent of the corresponding resolution cell. In this case, the analogue of the function $L_t^{\mathbf{k},z}$ takes the form

$$L_t(\mathbf{z}, \mathbf{x}) = \int f_{\mathbf{z}}(\mathbf{z}') \ell_t^{\mathbf{k}}(\mathbf{z}' | \mathbf{x}) d\mathbf{z}'.$$

If the noise in the propagation of the signal emitted by the target and received by the sensor is negligible when compared to the size of the resolution cells then it holds that $L(\mathbf{z}, \mathbf{x}) = f_{\mathbf{z}}(H(\mathbf{x}))$ where H is the observation function, that is,

$$L_t(\mathbf{z}, \mathbf{x}) = \exp\left(-\frac{1}{2}(H(\mathbf{x}) - \mathbf{z})^T \Sigma^{-1}(H(\mathbf{x}) - \mathbf{z})\right).$$

The expression of L_t is very close to the one of a standard likelihood function, except that there is no normalising constant ($L(\cdot, \mathbf{x})$ has maximum 1 for all $\mathbf{x} \in \mathbf{X}$). This approach can also be justified through a direct modelling of uncertainty [14], [15] which has connections with Dempster-Shafer theory [10], [28].

IV. APPROXIMATION OF THE MULTI-TARGET MARGINAL LIKELIHOOD

Theorem 2 is based on the yet-to-be-defined multi-target marginal likelihood $w_{\text{ex}}^{\mathbf{k},z}$ which value is needed for all pairs (\mathbf{k}, z) in $\mathbb{I}_{t|t-1} \times \bar{Z}_t$. However, the computation of these marginal likelihoods comes at the cost of a high complexity which, if performed exactly, would significantly reduce the efficiency of the proposed method. Instead, we formulate two possible approximations which are related to the ‘‘sparsity’’ of the scenario, either from the viewpoint of the targets or from the viewpoint of the observations: for all $\mathbf{k} \in I$ with $I \subseteq \mathbb{I}_{t|t-1}^m$ and for all $z \in Z$ with $Z \subseteq Z_t$, we assume that

- A.3** $p_t^{\mathbf{k},d}(z)p_t^{\mathbf{k}',d}(z) \approx 0$ for any $\mathbf{k}' \in I$ such that $\mathbf{k} \neq \mathbf{k}'$, or
A.4 $p_t^{\mathbf{k},d}(z)p_t^{\mathbf{k},d}(z') \approx 0$ for any $z' \in Z$ such that $z \neq z'$.

Considering **A.3** for a given I and a given Z is equivalent to assuming that two single-target laws with index in I are unlikely to be successfully updated (in terms of marginal likelihood) with the same observation $z \in Z$. Approximation **A.4** is the counterpart of **A.3**, for which two observations in Z are unlikely to successfully update the same single-target law $\bar{p}_t^{\mathbf{k}}$ with $\mathbf{k} \in I$. These two approximations allow for factorising the expression of $w_{\text{ex}}^{\mathbf{k},z}$. The results will be given using **A.3**, the analogue with **A.4** follows a very similar path.

Theorem 3. *For any $\mathbf{k} \in \mathbb{I}_{t|t-1}$ and any $z \in \bar{Z}_t$, applying **A.3** to the subsets $\mathbb{I}_{t|t-1}^{\mathbf{k}}$ and $Z_t \setminus \{z\}$, the scalar $w_{\text{ex}}^{\mathbf{k},z}$ can be factorised as follows*

$$w_{\text{ex}}^{\mathbf{k},z} = C'_t(\mathbf{k}, z) \prod_{\mathbf{k}' \in \mathbb{I}_{t|t-1}^m \setminus \{\mathbf{k}\}} \left[\bar{p}_t^{\mathbf{k}',d}(\phi) + \sum_{z' \in Z_t \setminus \{z\}} \frac{p_t^{\mathbf{k}',d}(z')}{C_t^{\mathbf{u},b}(z')} \right]$$

where

$$C_t^{\mathbf{u},b}(z) = \frac{p_t^{\mathbf{u},d}(z)}{\bar{p}_t^{\mathbf{u},d}(\phi)} + \frac{p_t^{\mathbf{b},z}(z)}{\bar{p}_t^{\mathbf{b},z}(\phi)}, \quad (8)$$

and where

$$C'_t(\mathbf{k}, z) = [\bar{p}_t^{\mathbf{u},d}(\phi)]^{n_t^{\mathbf{u}} - 1_{\mathbf{u}}(\mathbf{k})} \times \left[\prod_{z \in Z_t \setminus Z_b} \bar{p}_t^{\mathbf{b},z}(\phi) \right] \left[\prod_{z \in Z_t \setminus \{z\}} C_t^{\mathbf{u},b}(z) \right]$$

with Z_b equal to $\{z\}$ when $\mathbf{k} = \mathbf{b}$ and \emptyset otherwise.

Theorem 3 is a direct consequence of lemma 1 stated in the appendix. Note that the marginal likelihood $\bar{p}_t^{\mathbf{b},z}(\phi)$ used in the theorem is simply equal to $\ell_t^z(\phi | \psi_b)$. An important property of the HISP filter that appears in theorem 3 is that all the terms $w_{\text{ex}}^{\mathbf{k},z}$ can be computed with a complexity of order $\mathcal{O}(|\mathbb{I}_{t|t-1}| |Z_t|)$, as demonstrated in algorithm 1. The computation of all the terms $C'_t(\mathbf{k}, z)$ has a lower complexity ($\mathcal{O}(|Z_t|)$) and is not detailed.

Algorithm 1 Computation of $w_{\text{ex}}^{\mathbf{k},z}$ for all $(\mathbf{k}, z) \in \mathbb{I}_{t|t-1} \times \bar{Z}_t$

```

for  $\mathbf{k} \in \mathbb{I}_{t|t-1}$  do
  for  $z \in Z_t$  do
     $u_{\mathbf{k},z} \leftarrow p_t^{\mathbf{k},d}(z) / C_t^{\mathbf{u},b}(z)$ 
  end for
   $w_{\mathbf{k},\phi} \leftarrow \bar{p}_t^{\mathbf{k},d}(\phi)$ 
  for  $z \in Z_t$  do
     $w_{\mathbf{k},\phi} \leftarrow w_{\mathbf{k},\phi} + u_{\mathbf{k},z}$ 
  end for
  for  $z \in Z_t$  do
     $w_{\mathbf{k},z} \leftarrow w_{\mathbf{k},\phi} - u_{\mathbf{k},z}$ 
  end for
end for
for  $z \in \bar{Z}_t$  do
   $W_z \leftarrow 1$ 
  for  $\mathbf{k} \in \mathbb{I}_{t|t-1}$  do
     $W_z \leftarrow W_z w_{\mathbf{k},z}$ 
  end for
end for
for  $\mathbf{k} \in \mathbb{I}_{t|t-1}$  do
  for  $z \in \bar{Z}_t$  do
     $w_{\text{ex}}^{\mathbf{k},z} \leftarrow C'_t(\mathbf{k}, z) W_z / w_{\mathbf{k},z}$ 
  end for
end for

```

V. RELATION WITH OTHER WORKS

In this section, the relation between the proposed approach and other MTT techniques is discussed.

Labelled multi-Bernoulli (LMB) filter: Distinguishing targets has been made possible with point processes by using labelling [33], [32]. The LMB filter [25] follows as an approximation and is close in principle to the proposed approach. However, point processes have been built on the assumption that the targets are indistinguishable [4, p. 124] and labelling is usually meant to represent characteristics of the target that do not evolve in time instead of representing a target identity. These facts do not make the use of labels for distinguishing targets straightforward and specific techniques have to be used to prevent the natural symmetrisation of point-process laws. The objective with the proposed framework is to build on a natural representation of partially-distinguishable multi-target systems [16] which is based on a constructive approach leading in the independent case to a point process on the space of probability measures, and reducing to a multi-target configuration in the specific case of the HISP filter. From a practical point of view, the labelled multi-Bernoulli filter

computational complexity can be, in a worst-case scenario, as high as with non-approximated techniques, whereas the HISP filter's complexity is linear.

Poisson multi-Bernoulli filter: The idea of separating the never-detected targets from the detected ones has been proposed in [36], [37] where detected and never-detected targets are respectively represented by a Poisson point process and by a (mixture of) multi-Bernoulli point process(es). The use of a Poisson distribution for appearing targets offers a practical advantage when no upper bound is known for the associated cardinality, as opposed to the case of a finite-resolution sensor considered here where the number of appearing targets cannot exceed the number of resolution cells. Since it is based on unlabelled point processes, the Poisson multi-Bernoulli approach does not allow for a principled track extraction or for any post-processing requiring targets to be distinguishable, such as classification, unlike the proposed MTT algorithm [24]. To reduce the mixture of multi-Bernoulli point processes obtained after the update, [37] proposes to select the multi-Bernoulli distribution that minimises the Kullback-Leibler divergence with the mixture. This approach is well suited to unlabelled point processes where there is no question of distinguishability.

Many other approaches exist in the target-tracking literature, e.g. [11] considers random finite sets of trajectories.

VI. IMPLEMENTATION

Although the complexity is linear in the number of considered single-target laws and in the number of observations, specific approximations have to be used in practice to limit the computational cost and the number of propagated laws while ensuring that track extraction can be efficiently applied.

Let $\mathbb{J}_{t-1} \subseteq \mathbb{I}_{t-1}^\#$ contain the indices that have been retained up to time $t-1$ and denote $\mathbb{J}_t \subseteq \mathbb{I}_t^\#$ the set of indices obtained at time t after applying prediction and update to \mathbb{J}_{t-1} (in the same way \mathbb{I}_t is deduced from \mathbb{I}_{t-1}). Allowing different hypotheses to share the same single-target law, we denote $\bar{\mathbb{I}}_{t-1}$ the partition of \mathbb{J}_{t-1} characterising this aspect at time $t-1$, that is, the index $\mathbf{k} \in \bar{\mathbb{I}}_{t-1}$ of a single-target law $p_{t-1}^{\mathbf{k}}$ is actually defined as the set containing the indices of the hypotheses described by the law $p_{t-1}^{\mathbf{k}}$. The objective with this approach is to allow for the merging of single-target laws while keeping distinct the other characteristics of the involved hypotheses. In this way, a large number of hypotheses can be propagated with a reasonable computational cost (which is mostly determined by the number of single-target laws to predict and update).

Let \mathbb{L}_t be the propagated version of the partition $\bar{\mathbb{I}}_{t-1}$. Once again, an index $\mathbf{k} \in \mathbb{L}_t$ corresponds to a single-target law toward which several indices in \mathbb{J}_t can point. The weight $w_t^{\mathbf{k}}$ associated with the single-target law $\mathbf{k} \in \mathbb{L}_t$ is the sum of the weights of the hypotheses which relies on it, i.e. $w_t^{\mathbf{k}} = \sum_{i \in \mathbf{k}} w_t^i$ (which can be larger than 1, as opposed to hypotheses' weight).

Pruning: Some hypotheses' weight will be very close to 0 so that their probability of existence is low enough to discard them. The actual index set is then a subset $\tilde{\mathbb{J}}_t$ of \mathbb{J}_t . This *pruning* procedure is characterised by

A.5 The set $\tilde{\mathbb{J}}_t$ is the subset of \mathbb{J}_t containing indices i such that $w_t^i > \tau$ only.

The procedure affects the indices of single-target laws, however we also denote by $\bar{\mathbb{L}}_t$ the corresponding partition of $\tilde{\mathbb{J}}_t$. Single-target laws will be automatically discarded when their index becomes empty (as a subset of $\bar{\mathbb{J}}_t$).

Merging: Some of the single-target laws will, in practice, be too close to each other to justify propagating them separately, a partition $\bar{\mathbb{L}}_t$ of $\tilde{\mathbb{J}}_t$ can be introduced to group the laws that are alike. This is the *merging* procedure characterised in the Gaussian case as follows (denoting $\mathcal{N}(\mathbf{m}, \mathbf{V})$ the Gaussian distribution with mean \mathbf{m} and variance \mathbf{V} and assuming $p_t^{\mathbf{k}} = \mathcal{N}(\mathbf{m}_t^{\mathbf{k}}, \mathbf{V}_t^{\mathbf{k}})$ for any $\mathbf{k} \in \bar{\mathbb{L}}_t$).

A.6 The partition $\bar{\mathbb{L}}_t$ of $\tilde{\mathbb{J}}_t$ is introduced recursively as:

- (i) Define $\bar{\mathbb{L}}_t$ on $K = \emptyset$ as the empty partition
- (ii) Find the index $\mathbf{k} = \operatorname{argmax}_{j \in \bar{\mathbb{L}}_t \setminus K} w_t^j$ corresponding to the single-target law with highest weight among the ones that have not already been merged and define K' as the set containing any index $j \in \bar{\mathbb{L}}_t$ such that the Mahalanobis distance [19] between $\mathcal{N}(\mathbf{m}_t^{\mathbf{k}}, \mathbf{V}_t^{\mathbf{k}} + \mathbf{V}_t^j)$ and \mathbf{m}_t^j is strictly less than τ' (recalling that each index $j \in K'$ is itself a subset of $\tilde{\mathbb{J}}_t$)
- (iii) Let \mathbf{k}' be the union of the cells in K' and let $p_t^{\mathbf{k}'}$ be characterised by its mean and variance as

$$\mathbf{m}_t^{\mathbf{k}'} = \frac{1}{\sum_{\mathbf{k} \in K'} w_t^{\mathbf{k}}} \sum_{\mathbf{k} \in K'} w_t^{\mathbf{k}} \mathbf{m}_t^{\mathbf{k}}$$

$$\mathbf{V}_t^{\mathbf{k}'} = \frac{1}{\sum_{\mathbf{k} \in K'} w_t^{\mathbf{k}}} \times \sum_{\mathbf{k} \in K'} w_t^{\mathbf{k}} (\mathbf{V}_t^{\mathbf{k}} + (\mathbf{m}_t^{\mathbf{k}'} - \mathbf{m}_t^{\mathbf{k}})(\mathbf{m}_t^{\mathbf{k}'} - \mathbf{m}_t^{\mathbf{k}})^T)$$

- (iv) Extend $\bar{\mathbb{L}}_t$ to $K \cup \mathbf{k}'$ by letting \mathbf{k}' be a cell of the partition
- (v) Redefine K as $K \cup \mathbf{k}'$ and return to step (ii) until $K = \bar{\mathbb{L}}_t$

The two indexed families of interest are then $\{w_t^i\}_{i \in \tilde{\mathbb{J}}_t}$ and $\{p_t^{\mathbf{k}}\}_{\mathbf{k} \in \bar{\mathbb{L}}_t}$. These approximations are usual for handling Gaussian mixtures [26] but can be applied here for any implementation of the filter by adapting the considered distance, e.g. the Hellinger distance [12].

Although the reduction of the number of hypotheses does not induce a computational gain as significant as the reduction of the number of single-target laws, there is still some interest in mixing hypotheses that are very similar, especially for long and/or complex scenarios. To decide when hypotheses are similar enough to be mixed, we consider a time window $T = \{t-l, \dots, t\}$ at time t for some lag l and require the corresponding observation paths to be the same over the window T . If there is a subset I of hypotheses' indices verifying this assumption and if it holds that $w_t^I = \sum_{i \in I} w_t^i \leq 1$ then these hypotheses can be mixed: the observation path of the hypothesis with highest weight can be kept, e.g. for display purposes, and the weight of the new hypothesis is w_t^I . If the laws associated with each hypothesis being mixed are distinct then the resulting hypothesis is associated with a weighted mixture of the corresponding laws (although this does not typically happen when $l > 1$). We denote by \mathbb{J}_t the set of hypotheses at time t resulting from this mixing procedure.

Track extraction: As mentioned in previous sections, a posterior probability for an observation to be a false alarm is computed and the result is stored as an hypothesis for the purpose of track extraction. The hypotheses corresponding to disappeared targets are kept for the same reasons. The track extraction also operates on the time window T , so that these hypotheses can be discarded once the time of their last observation falls out of this window. Finally, in order to perform track extraction, one can solve the following optimisation problem:

$$\operatorname{argmax}_{I \subseteq \mathbb{J}_t} \prod_{i \in I} w_t^i \quad (9)$$

subject to:

- 1) the union of all observation paths over the time window T must contain all the observations over this window,
- 2) the observations paths in I must be *compatible*: each observation cannot be used more than once.

The solution to this problem is the same as the one for

$$\operatorname{argmax}_{I \subseteq \mathbb{J}_t} \sum_{i \in I} \log w_t^i$$

with the same constraints, since all w_t^i are strictly positive. The latter problem can however be solved by linear programming. Constraint 1 justifies the fact that false alarms and disappeared targets are kept as hypotheses; the solution to eq. (9) would not be meaningful otherwise. Constraint 2 ensures that the assumptions of standard MTT are satisfied. The only parameter for track extraction is the size of the time window T which is practically appealing since it is easy to interpret and tune.

Remark 2. To ensure that previously displayed tracks do not disappear when they have not been detected over the time window T , the corresponding observations can be kept even when their time is prior to T , i.e. observations corresponding to confirmed tracks are held longer in order to improve the results with a limited impact on the computational time.

The track extraction procedure proposed in this section is only one among many possible. The fact that the HISP filter provides distinct hypotheses enables the introduction of tailored extraction methods depending on the application and computational resources at hand. The procedure proposed in this section is considered since it is one of the simplest that utilises the structure of the filter as opposed, for instance, to selecting single-target laws based on their weight.

VII. SIMULATIONS

The performance of the HISP filter is compared against the PHD and cardinalised PHD (CPHD) filters [21] as well as the LMB filter. Note that because of its hierarchical nature, the HISP filter can be easily implemented using any Bayesian filtering technique for each single-target law, e.g. Sequential Monte Carlo (SMC) as in [17] or Kalman filter (KF).

We consider a sensor placed at the centre of the 2-D Cartesian plane that delivers range and bearing observations every $\Delta = 4\text{s}$ during 200s, i.e. the time index set is $\mathbb{T} \doteq \{0, \dots, 50\}$ with the actual time being $4t$ for any $t \in \mathbb{T}$. The size of the

resolution cells of this sensor is $1^\circ \times 15\text{m}$. Considering small fixed random error and bias error, the standard deviation of the observations is $\sigma_r = 6.2\text{m}$ for the range and $\sigma_\theta = 4.5\text{mrad}$ for the bearing, for a signal-to-noise ratio (SNR) of 3dB and $\sigma_r = 4.87\text{m}$ and $\sigma_\theta = 3.5\text{mrad}$ for a SNR of 5dB. The range r is in $[50\text{m}, 500\text{m}]$ and the bearing θ is in $(-\pi, \pi]$. For the comparison with the (C)PHD and LMB filters to be possible, point observations are generated according to the standard observation model with the standard deviations given above, instead of using the resolution cells. The alternative sensor modelling of section III-D is thus used for the HISP filter.

The scenario comprises 5 targets which motion is driven by a nearly-constant velocity model so that the random variable X_t representing the state of a target in \mathbf{X}^\bullet at time t given its state \mathbf{x}_{t-1} at the previous time verifies $X_t \sim \mathcal{N}(F\mathbf{x}_{t-1}, \mathbf{P})$ with

$$F = \begin{bmatrix} 1 & 0 & \Delta & 0 \\ 0 & 1 & 0 & \Delta \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad \mathbf{P} = \sigma^2 \begin{bmatrix} \Delta^3/3 & 0 & \Delta^2/2 & 0 \\ 0 & \Delta^3/3 & 0 & \Delta^2/2 \\ \Delta^2/2 & 0 & \Delta & 0 \\ 0 & \Delta^2/2 & 0 & \Delta \end{bmatrix}$$

with $\sigma^2 = 0.05\text{m}^2/\text{s}^4$. All targets are present at the beginning of the scenario and Targets 2 to 5 never spontaneously disappear whereas Target 1 disappear at 160s in Case 1 below and does not disappear in Cases 2 and 3. The scenario is depicted in fig. 2a. Note that Targets 2 and 3 are crossing around $t = 120\text{s}$.

We consider a KF implementation of the HISP filter based on A.3, A.5 and A.6 and similarly for the LMB filter. In this implementation, the detected and never-detected hypotheses are updated through eq. (7b) and eq. (7a) respectively. For the (C)PHD filter, a Gaussian mixture implementation³ [31], [34] with a confirmation threshold $\tau_c = 0.9$ is considered, i.e. all the terms in the Gaussian mixture with a weight exceeding τ_c are declared as confirmed tracks. The non-linearity of the observation model is dealt with by an extended Kalman filter. To reduce the computational cost, pruning (with parameter $\tau = 10^{-5}$) and merging (with parameter $\tau' = 4$) are carried out on the collection of posterior single-target laws or on the Gaussian mixture, depending on the filter. The probability for a target with state $\mathbf{x} \in \mathbf{X}^\bullet$ at time $t-1$ of remaining within \mathbf{X}^\bullet at time t is set to $p_\pi \doteq p_t^\pi(\mathbf{x}) = 0.999$.

In the considered scenarios, the mean number of appearing targets m_α is set to 10^{-2} per time step. Targets are assumed to be detected upon appearance in the PHD and HISP filters, the corresponding distribution is induced by the observation as in [18] with the velocity components being initialised as Gaussian with mean 0 and standard deviation 3.5m/s ; the associated weight is w_t^α which is defined as $m_\alpha/|Z_t|$ for any $t \in \mathbb{T}$. The CPHD and LMB filters are initialised with a Gaussian mixture whose terms are centred around the location of appearance of the targets, with a standard deviation of 50m on the position, and with the other parameters being the same as for the other filters. The average number of false alarms per time step is denoted n_b . The probability of detection is constant across the state space and through time, so that it is

³The code of the CPHD and LMB filters was downloaded from Ba Tuong Vo's page (http://ba-tuong.vo-au.com/rfs_tracking_toolbox_beta.7z)

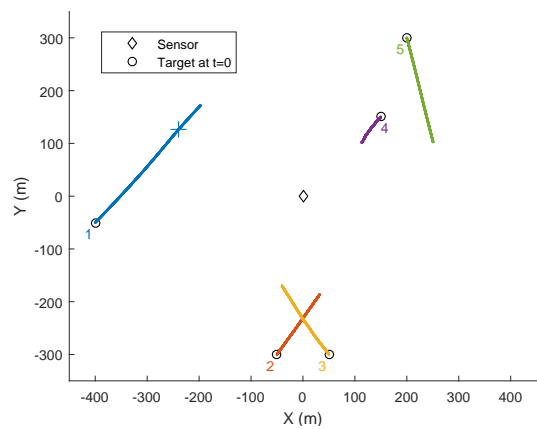
denoted $p_d \doteq p_t^d(\mathbf{x})$ for any $\mathbf{x} \in \mathbf{X}^\bullet$ and any time t . From the given characteristics of the sensor and for a given value of p_d , we deduce the probability for a single observation cell to produce a false alarm and we denote it w_b . The approximate value of n_b can then be deduced directly from the number of observations cells. We proceed to the performance assessment on three different scenarios.

Case 1: High probability of detection (5dB): We set $p_d = 0.995$ so that $w_b = 7.67 \times 10^{-3}$ and $n_b \approx 83$. The HISP track-extraction window is set to a length of 5. The OSPA distance [27] depicted in fig. 2b is based on a cutoff of 100 and a 2-norm and is averaged over 100 Monte Carlo (MC) runs. Even though the estimation problem is not challenging with these parameters, there is a noticeable difference of performance between the two types of filters. This is mainly caused by the additional weighting term $w_{\text{ex}}^{k,z}$ of the HISP filter which allows for a better discrimination between likely and unlikely hypotheses and which reduces the effects of association uncertainty on the overall performance. The CPHD filter takes the longest time to react to the disappearance of Target 1. The performance of the PHD filter is reduced when Targets 2 and 3 cross whereas the performance of the other filters is not affected.

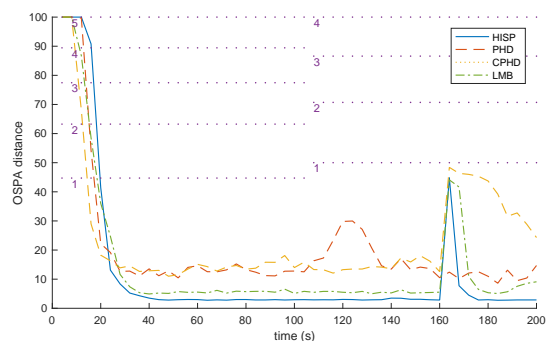
Case 2: Low probability of detection (3dB): We set $p_d = 0.5$ so that $w_b = 1.34 \times 10^{-3}$ and $n_b \approx 15$. The HISP track-extraction window is set to a length of 6 since there might be many consecutive detection failures. The average OSPA distance is shown in fig. 2c. The OSPA distance for the HISP filter is the lowest at all time steps. Due to the combination of a low probability of detection and the uncertainty on the association, the OSPA distance for the HISP filter increases when Targets 3 and 4 cross. The performance of the HISP filter in this case is mainly explained by the fact that it computes an a posteriori probability of detection, so that the prior probability, $p_d = 0.5$ here, has a lower impact on the final result when compared to the other methods.

Case 3: High probability of false alarms (3dB): In this case, we set $p_d = 0.8$ so that $w_b = 1.54 \times 10^{-2}$ and $n_b \approx 167$. The HISP track-extraction window is set to a length of 3 in order to alleviate the computational cost. The average OSPA distance is shown in fig. 2d. The PHD filter, which is known to be robust to high numbers of false alarms, behaves slightly better than in Case 2. The CPHD and LMB filters react significantly faster to target appearance than the HISP filter but this might be related to the more informative birth process they use. The CPHD filter is more prone to the creation of false tracks that impede its performance in the longer run.

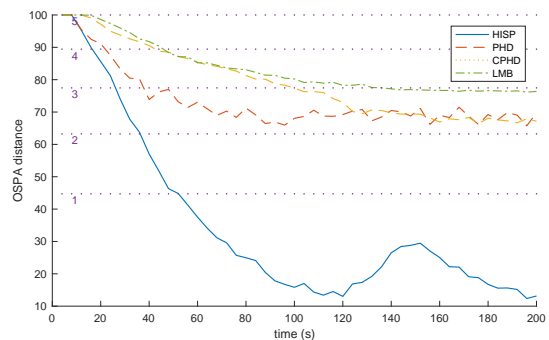
Parameter sensitivity: The HISP filter displays a high sensitivity to some parameters when compared to the PHD filter. In particular, and as shown in fig. 3, the value of the probability p_π has an important impact on the behaviour of the filter: $p_\pi = 1$ implies that if an hypothesis is almost-surely correct then it will be displayed at all following time steps, alternatively, if $p_\pi \leq p_d$ then hypotheses stop to be considered as tracks as soon as a detection failure happens. Conversely, the behaviour of the PHD filter is nearly independent of p_π , so that this filter does not actually allow for taking the knowledge about persistence of targets into account. The



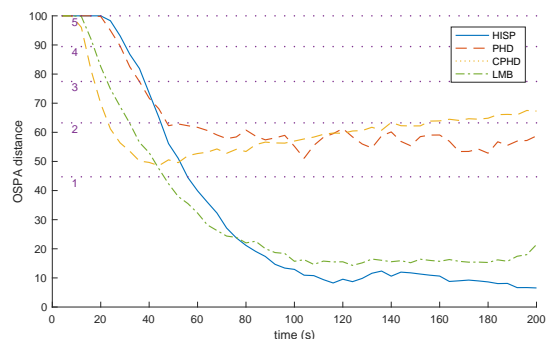
(a) A realisation of the target trajectories (blue cross: location of Target 1 when it disappears in Case 1.)



(b) Case 1: $p_d = 0.995$ and $n_b \approx 83$.



(c) Case 2: $p_d = 0.5$ and $n_b \approx 15$.



(d) Case 3: $p_d = 0.8$ and $n_b \approx 167$.

Fig. 2: OSPA distance in Cases 1-3 (b-d) on the scenario (a) over 100 MC runs. (HISP filter: solid line – PHD: dashed line – the dotted line numbered n represents the OSPA for a cardinality-only error of n)

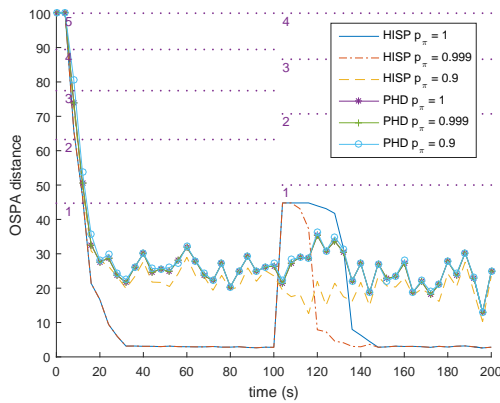


Fig. 3: OSPA distance versus time for different values of p_π , averaged over 50 MC runs.

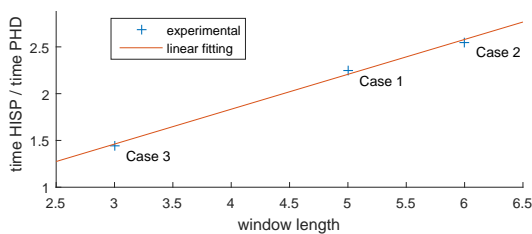


Fig. 4: Relative computational time of the HISP filter w.r.t. the PHD filter for Cases 1 to 3, averaged over 100 MC runs.

scenario considered in fig. 3 is a slightly modified version of the one considered above for Case 1, with a probability of detection $p_d = 0.9$, with $n_b = 10$ and with Target 1 disappearing at 100s rather than at 160s.

Computational time: Although both the PHD and HISP filters have a linear complexity, the computational cost for the HISP filter tends to be higher than for the PHD filter, especially when the time window used in track extraction is large. The ratio between the measured computational times of the two filters is displayed in fig. 4 and shows that the relation between this ratio and the length of the time window appears to be linear, although this is only based on 3 data points.

CONCLUSION

A new MTT algorithm, called the HISP filter, has been derived and detailed. When studying this filter, it appeared that there is more than one way of using the update equations and that there are different possible approximations as well as diverse applicable modelling alternatives. In this sense, the HISP filter can be seen as a general and computationally-affordable way of approaching the problem of MTT. The HISP filter allows for characterising each hypothesis separately thus giving a local picture of the underlying multi-target problem while controlling the level of approximation. Its efficiency has been compared with the performance of the PHD, CPHD and LMB filters and the results show that the HISP filter outperforms these algorithms in several cases with varying probabilities of detection and statistics for the false alarms.

REFERENCES

- [1] Y. Bar-Shalom. *Tracking and data association*. Academic Press Professional, Inc., 1987.
- [2] S. S. Blackman. *Multiple-target tracking with radar applications*, volume 1. Artech House, Norwood, 1986.
- [3] F. Caron, P. Del Moral, A. Doucet, and M. Pace. On the conditional distributions of spatial point processes. *Advances in Applied Probability*, 43(2):301–307, 2011.
- [4] D. Daley and D. Vere-Jones. *An Introduction to The Theory of Point Processes*. Springer, 2nd edition, 2003.
- [5] E. Delande, C. Früh, J. Franco, J. Houssineau, and D. E. Clark. Novel multi-object filtering approach for space situational awareness. *Accepted to: Journal of Guidance, Control, and Dynamics; Special Issue on Space Domain Awareness*, 2017.
- [6] E. Delande, J. Houssineau, and D. E. Clark. Performance metric in closed-loop sensor management for stochastic populations. In *IEEE Sensor Signal Processing for Defence (SSPD)*, 2014.
- [7] E. Delande, J. Houssineau, and D. E. Clark. Multi-object filtering with stochastic populations. *arXiv preprint arXiv:1501.04671v2*, 2016.
- [8] E. Delande, J. Houssineau, J. Franco, C. Früh, and D. E. Clark. A new multi-target tracking algorithm for a large number of orbiting objects. In *27th AAS/AIAA Space Flight Mechanics Meeting*, 2017.
- [9] E. Delande, M. Uney, J. Houssineau, and D. E. Clark. Regional variance for multi-object filtering. *IEEE T on Signal Processing*, 62(13):3415–3428, 2014.
- [10] A. P. Dempster. A generalization of Bayesian inference. *Journal of the Royal Statistical Society. Series B (Methodological)*, 1968.
- [11] Á. F. García-Fernández, L. Svensson, and M. R. Morelande. Multiple target tracking based on sets of trajectories. *arXiv preprint arXiv:1605.08163*, 2016.
- [12] E. Hellinger. Neue Begründung der Theorie quadratischer Formen von unendlichvielen Veränderlichen. *Journal für die reine und angewandte Mathematik*, 136:210–271, 1909.
- [13] J. Houssineau. *Representation and estimation of stochastic populations*. PhD thesis, Heriot-Watt University, 2015.
- [14] J. Houssineau and A. Bishop. Smoothing and filtering with a class of outer measures. *arXiv preprint arXiv:1704.01233*, 2017.
- [15] J. Houssineau and D. E. Clark. Bayesian data assimilation based on a family of outer measures. *arXiv preprint arXiv:1611.02989*, 2016.
- [16] J. Houssineau and D. E. Clark. On a representation of partially-distinguishable populations. *arXiv:1608.00723*, 2016.
- [17] J. Houssineau, D. E. Clark, and P. Del Moral. A sequential Monte Carlo approximation of the HISP filter. In *European Signal Processing Conference (EUSIPCO)*, 2015.
- [18] J. Houssineau and D. Laneuville. PHD filter with diffuse spatial prior on the birth process with applications to GM-PHD filter. In *13th Conference on Information Fusion (FUSION)*, 2010.
- [19] P. C. Mahalanobis. On the generalized distance in statistics. *Proceedings of the National Institute of Sciences (Calcutta)*, 2:49–55, 1936.
- [20] R. P. S. Mahler. Multitarget Bayes filtering via first-order multitarget moments. *Aerospace and Electronic Systems, IEEE T on*, 39(4):1152–1178, 2003.
- [21] R. P. S. Mahler. PHD filters of higher order in target number. *Aerospace and Electronic Systems, IEEE T on*, 43(4):1523–1543, 2007.
- [22] K. Okuma, A. Taleghani, N. De Freitas, J. J. Little, and D. G. Lowe. A boosted particle filter: Multitarget detection and tracking. In *European Conference on Computer Vision*. Springer, 2004.
- [23] M. Pace and P. Del Moral. Mean-field PHD filters based on generalized Feynman-Kac flow. *Journal of selected topics in signal processing, special issue on multi-target tracking*, 2013.
- [24] Y. Pailhas, J. Houssineau, Y. Petillot, and D. E. Clark. Tracking with MIMO sonar systems: applications to harbour surveillance. *IET Radar, Sonar & Navigation*, 2016.
- [25] S. Reuter, B.-T. Vo, B.-N. Vo, and K. Dietmayer. The labeled multi-Bernoulli filter. *IEEE T on Signal Processing*, 62(12):3246–3260, 2014.
- [26] D. J. Salmond. Mixture reduction algorithms for target tracking in clutter. In *OE/LASE'90*. Int. Society for Optics and Photonics, 1990.
- [27] D. Schuhmacher, B.-T. Vo, and B.-N. Vo. A consistent metric for performance evaluation of multi-object filters. *IEEE T on Signal Processing*, 56(8):3447–3457, 2008.
- [28] G. Shafer. *A mathematical theory of evidence*, volume 1. Princeton university press Princeton, 1976.
- [29] S. S. Singh, B.-N. Vo, A. Baddeley, and S. Zuyev. Filters for spatial point processes. *SIAM Journal on Control and Optimization*, 48(4):2275–2295, 2009.

- [30] J. Vermaak, S. J. Godsill, and P. Perez. Monte Carlo filtering for multi target tracking and data association. *IEEE T on Aerospace and Electronic systems*, 41(1):309–332, 2005.
- [31] B.-N. Vo and W.-K. Ma. The Gaussian mixture probability hypothesis density filter. *IEEE T on Signal Processing*, 54(11):4091–4104, 2006.
- [32] B.-N. Vo, B.-T. Vo, and D. Phung. Labeled random finite sets and the Bayes multi-target tracking filter. *IEEE T on Signal Processing*, 62(24):6554–6567, 2014.
- [33] B.-T. Vo and B.-N. Vo. Labeled random finite sets and multi-object conjugate priors. *IEEE T on Signal Processing*, 61(13):3460–3475, 2013.
- [34] B.-T. Vo, B.-N. Vo, and A. Cantoni. Analytic implementations of the cardinalized probability hypothesis density filter. *IEEE T on Signal Processing*, 55(7):3553–3567, 2007.
- [35] B.-T. Vo, B.-N. Vo, and A. Cantoni. The cardinality balanced multi-target multi-Bernoulli filter and its implementations. *IEEE T on Signal Processing*, 57(2):409–423, 2009.
- [36] J. L. Williams. Hybrid Poisson and multi-Bernoulli filters. In *15th International Conference on Information Fusion (FUSION)*, 2012.
- [37] J. L. Williams. An efficient, variational approximation of the best fitting multi-Bernoulli filter. *IEEE T on Signal Processing*, 63(1):258–273, 2015.

APPENDIX

Proof of theorem 1. The principle of the proof is as follows: in general, all combinations of transition functions must be applied to the joint law of all hypotheses at time $t - 1$. Without the considered assumptions on the transitions, this would generate a large number of multi-target configurations at time t . However, even when making use of the considered modelling, each single-target law which has its probability mass on \mathbf{X}^\bullet at time $t-1$ can be predicted using two transitions: q_t^π and q_t^ω . In order to obtain a single predicted multi-target configuration, we introduce a new transition q_t defined as $q_t(\mathbf{x}' | \mathbf{x}) = q_t^\pi(\mathbf{x}' | \mathbf{x}) + q_t^\omega(\mathbf{x}' | \mathbf{x})$ for any $\mathbf{x}, \mathbf{x}' \in \mathbf{X}$ as well as its extension \bar{q}_t which is defined as in eq. (3). As a consequence, we only consider the transitions \bar{q}_t^α and \bar{q}_t .

The transition \bar{q}_t is the only one which can be used to predict the law $\delta_{\psi_\#}$ without yielding δ_φ , so that there is no viable alternative. The transition \bar{q}_t^α is applied to a fixed number n_t^α of targets with law $\delta_{\psi_\#}$ and the other targets with this law can only be propagated with \bar{q}_t , so that there is only one possibility. The multiplicity n_φ of the law δ_φ is irrelevant and is not estimated. For the same reasons, targets that are almost-surely in the state space \mathbf{X}^\bullet can only be propagated with \bar{q}_t so that the joint law $P_{t|t-1}^\#$ induced by the hypotheses in $\mathbb{I}_{t|t-1}^\#$ is

$$P_{t|t-1}^\#(\mathbf{x}_1, \dots, \mathbf{x}_N) = \prod_{j=1}^{n_{t-1}^u} \left[\int \bar{q}_t(\mathbf{x}_j | \mathbf{x}') \bar{p}_{t-1}^u(\mathbf{x}') d\mathbf{x}' \right] \\ \times \prod_{\mathbf{k} \in \mathbb{I}_{t-1}^m} \left[\int \bar{q}_t(\mathbf{x}_{n(i)} | \mathbf{x}') \bar{p}_{t-1}^{\mathbf{k}}(\mathbf{x}') d\mathbf{x}' \right],$$

where $N = n_{t-1}^u + |\mathbb{I}_{t-1}^m| = |\mathbb{I}_{t|t-1}^\#|$ and where n is any one-to-one mapping between the sets \mathbb{I}_{t-1}^m and $\{n_{t-1}^u + 1, \dots, N\}$. However, when marginalising all single-target laws but one, say $i \in \mathbb{I}_{t-1}^m$, and when considering the restriction to \mathbf{X}^\bullet or to $\psi_\#$, the prediction of the single-target law with index \mathbf{k} via \bar{q}_t^π or \bar{q}_t^ω can be recovered through

$$\bar{p}_{t|t-1}^{\mathbf{k}, \omega}(\psi_\#) \doteq \int \delta_{\psi_\#}(\mathbf{x}_{n(\mathbf{k})}) P_{t|t-1}^\#(\mathbf{x}_1, \dots, \mathbf{x}_N) d(\mathbf{x}_1, \dots, \mathbf{x}_N)$$

and $\bar{p}_{t|t-1}^{\mathbf{k}, \omega}(\varphi) = 1 - \bar{p}_{t|t-1}^{\mathbf{k}, \omega}(\psi_\#)$, and through

$$\bar{p}_{t|t-1}^{\mathbf{k}, \pi}(\mathbf{x}) \doteq \int \delta_{\mathbf{x}}(\mathbf{x}_{n(\mathbf{k})}) P_{t|t-1}^\#(\mathbf{x}_1, \dots, \mathbf{x}_N) d(\mathbf{x}_1, \dots, \mathbf{x}_N)$$

for any $\mathbf{x} \in \mathbf{X}^\bullet$ and $\bar{p}_{t|t-1}^{\mathbf{k}, \pi} = \int_{\mathbf{X}^\bullet} \bar{p}_{t|t-1}^{\mathbf{k}, \pi}(\mathbf{x}) d\mathbf{x}$. It can indeed be verified that

$$\bar{p}_{t|t-1}^{\mathbf{k}, \omega} = \int \bar{q}_t^\omega(\cdot | \mathbf{x}) \bar{p}_{t-1}^{\mathbf{k}}(\mathbf{x}) d\mathbf{x} \\ \bar{p}_{t|t-1}^{\mathbf{k}, \pi} = \int \bar{q}_t^\pi(\mathbf{x} | \mathbf{x}) \bar{p}_{t-1}^{\mathbf{k}}(\mathbf{x}) d\mathbf{x}$$

The propagation of the single-target law \bar{p}_{t-1}^u of the never-detected targets via \bar{q}_t^π or \bar{q}_t^ω can be similarly recovered; for instance the law defined by

$$\bar{p}_{t|t-1}^{u, \pi}(\mathbf{x}) \doteq \int P_{t|t-1}^\#(\mathbf{x}, \mathbf{x}_2, \dots, \mathbf{x}_N) d(\mathbf{x}_2, \dots, \mathbf{x}_N)$$

and by $\bar{p}_{t|t-1}^{u, \pi}(\varphi) = 1 - \int \bar{p}_{t|t-1}^{u, \pi}(\mathbf{x}) d\mathbf{x}$ is indeed equal to the common law of the never-detected targets predicted via \bar{q}_t^π . Thus, **A.1** yields a multi-target configuration of the form

$$\{(\bar{p}_{t|t-1}^i, 1)\}_{i \in \mathbb{I}_{t|t-1}^m} \cup \{(\bar{p}_{t|t-1}^{u, \pi}, n_{t-1}^u)\} \\ \cup \{(\bar{p}_{t|t-1}^{u, \omega}, n_{t-1}^u)\} \cup \{(\bar{p}_t^\alpha, n_t^\alpha)\} \cup \bar{P}_{t-1}^{\psi_\#, \varphi}$$

where $\bar{p}_t^\alpha = \bar{q}_t^\alpha(\cdot | \psi_\#)$. The targets that disappear during prediction are then mixed by Simplification S.2 and hence return to $\bar{P}_{t-1}^{\psi_\#, \varphi}$. Although the law $\int \bar{q}_t^\omega(\cdot | \mathbf{x}) \bar{p}_{t-1}^{\mathbf{k}}(\mathbf{x}) d\mathbf{x}$ might have some probability mass on φ , the mixing with an infinite number of laws $\delta_{\psi_\#}$ will make this negligible. Then, by Simplification S.1, the law $\bar{p}_{t|t-1}^u$ results from the mixing of the n_{t-1}^u never-detected targets with law $\bar{p}_{t|t-1}^{u, \pi}$ and the n_t^α appearing targets with law \bar{p}_t^α , so that $n_{t|t-1}^u = n_{t-1}^u + n_t^\alpha$ also follows. This concludes the proof of the theorem. \square

The proof of theorem 2 requires the introduction of some additional notations. Let \mathcal{A}_t be the set of subsets of $\mathbb{I}_t^\#$ describing all the possible associations of hypotheses in $\mathbb{I}_{t|t-1}^\#$ with observations in \bar{Z}_t such that hypotheses do not share non-empty observations and let $\mathcal{A}_t(\mathbf{k}, z)$ be the subset of \mathcal{A}_t made of sets such that the hypothesis $\mathbf{k} \in \mathbb{I}_{t|t-1}$ is associated with $z \in \bar{Z}_t$. Let w_t be the DISP filter's posterior probability mass function on $\mathcal{A}_t \times \mathbb{N}$ (see [13, Corollary 3.4]) corresponding to the HISP filter's predicted multi-target configuration $\bar{P}_{t|t-1}$ and defined for any $I \in \mathcal{A}_t$ as

$$w_t(I, n_{t|t-1}^u - n_I) = [\bar{p}_{t|t-1}^{u, d}(\phi)]^{n_{t|t-1}^u - n_I} \left[\prod_{z \in Z_I^b} \bar{p}_{t|t-1}^{b, z}(z) \right] \\ \times \left[\prod_{z \in Z_t^b \setminus Z_I^b} \bar{p}_{t|t-1}^{b, z}(\phi) \right] \left[\prod_{\substack{(\mathbf{k}, s, z) \in \sigma_t^{-1}[I] \\ (\mathbf{k}, z) \neq (u, \phi)}} \bar{p}_{t|t-1}^{\mathbf{k}, d}(z) \right],$$

with $n_I \leq |Z_t|$ the number of targets in I detected for the first time and with $Z_I^b \subseteq Z_t$ the subset of observations considered as false alarms in I . The proof of the expression of w_t is out of the scope of the present article; however, the scalar $w_t^I = w_t(I, n_{t|t-1}^u - n_I)$ can be seen to be the marginal likelihood associated with the update of the hypotheses selected by I

with the corresponding observations and with the failure of the detection of the others. It follows that

$$W_t = \sum_{I \in \mathcal{A}_t} w_t^I \quad (10)$$

is the full multi-target marginal likelihood.

Proof of theorem 2. For any $\mathbf{i} \in \mathbb{I}_t$ with $(\mathbf{k}, s, z) \doteq \sigma_t^{-1}(\mathbf{i})$, the posterior marginal law $\bar{p}_t^{\mathbf{i}}$ on $\bar{\mathbf{X}}$ can be characterised by

$$\bar{p}_t^{\mathbf{i}}(\mathbf{x}) = \sum_{I \in \mathcal{A}_t(\mathbf{k}, z)} \frac{w_t^I}{W_t \bar{p}_{t|t-1}^{\mathbf{k}, s}(z)} L_t^{s, z}(\mathbf{x}) p_{t|t-1}^{\mathbf{k}}(\mathbf{x}), \quad (11)$$

for any $\mathbf{x} \in \mathbf{X}$, and by $\bar{p}_t^{\mathbf{i}}(\varphi) = 1 - \int_{\mathbf{X}} \bar{p}_t^{\mathbf{i}}(\mathbf{x}) d\mathbf{x}$. Defining

$$w_{\text{ex}}^{\mathbf{k}, z} = \sum_{I \in \mathcal{A}_t(\mathbf{k}, z)} w_t^I,$$

we can rewrite eq. (11) as

$$\bar{p}_t^{\mathbf{i}}(\mathbf{x}) = W_t^{-1} w_{\text{ex}}^{\mathbf{k}, z} L_t^{\mathbf{k}, z}(\mathbf{x}) p_{t|t-1}^{\mathbf{k}}(\mathbf{x})$$

for any $\mathbf{x} \in \mathbf{X}$, so that it remains to prove that the marginal likelihood W_t can be equivalently expressed as the denominator of either eq. (7a) or eq. (7b). It is sufficient to verify that for any $z \in \bar{Z}_t$ and any $\mathbf{k} \in \mathbb{I}_{t|t-1}$, it holds that

$$\bigcup_{\mathbf{k}' \in \mathbb{I}_{t|t-1}} \mathcal{A}_t(\mathbf{k}', z) = \bigcup_{z' \in \bar{Z}_t} \mathcal{A}_t(\mathbf{k}, z') = \mathcal{A}_t$$

since the constraint ‘‘hypothesis \mathbf{k} is associated with observation z ’’ is loosen by taking the union over all \mathbf{k} in $\mathbb{I}_{t|t-1}$ or over all z in Z_t (informally, this constraint becomes ‘‘an object is associated with z ’’ or ‘‘ \mathbf{k} is associated with some possibly-empty observations’’, which is always true under the standard MTT assumptions). It follows that

$$W_t = \sum_{I \in \mathcal{A}_t} w_t^I = \sum_{\mathbf{k}' \in \mathbb{I}_{t|t-1}} w_{\text{ex}}^{\mathbf{k}', z} \bar{p}_t^{\mathbf{k}'}(z) = \sum_{z' \in \bar{Z}_t} w_{\text{ex}}^{\mathbf{k}, z'} \bar{p}_t^{\mathbf{k}}(z').$$

The expression eq. (6) of the approximated multi-target configuration $\bar{\mathcal{P}}_t$ if then formed under **A.2**. \square

Theorem 3 is a corollary of the following lemma, where \mathcal{A}_t^{m} refers to the subset of \mathcal{A}_t where all hypotheses have been detected at least once up to time $t-1$.

Lemma 1. *Considering A.3 for the sets \mathbb{I}_t^{m} and Z_t , the multi-target marginal likelihood W_t can be factorised as*

$$W_t = C_t \prod_{\mathbf{k} \in \mathbb{I}_{t|t-1}^{\text{m}}} \left[\bar{p}_{t|t-1}^{\mathbf{k}, \text{d}}(\phi) + \sum_{z \in Z_t} \frac{p_{t|t-1}^{\mathbf{k}, \text{d}}(z)}{C_t^{\text{u}, \text{b}}(z)} \right],$$

where

$$C_t = [\bar{p}_{t|t-1}^{\text{u}, \text{d}}(\phi)]^{n_{t|t-1}^{\text{u}}} \left[\prod_{z \in Z_t^{\text{u}}} \bar{p}_{t|t-1}^{\text{b}, z}(\phi) \right] \left[\prod_{z \in Z_t} C_t^{\text{u}, \text{b}}(z) \right].$$

Alternatively, considering **A.4** for \mathbb{I}_t^{m} and Z_t , it holds that

$$W_t = C_t^{\phi} \prod_{z \in Z_t} \left[C_t^{\text{u}, \text{b}}(z) + \sum_{\mathbf{k} \in \mathbb{I}_{t|t-1}^{\text{m}}} \frac{p_{t|t-1}^{\mathbf{k}, \text{d}}(z)}{\bar{p}_{t|t-1}^{\mathbf{k}, \text{d}}(\phi)} \right],$$

where the constant C_t^{ϕ} is the joint probability for all the targets in the system to be undetected at time t , defined as

$$C_t^{\phi} = [\bar{p}_{t|t-1}^{\text{u}, \text{d}}(\phi)]^{n_{t|t-1}^{\text{u}}} \left[\prod_{z \in Z_t^{\text{u}}} \bar{p}_{t|t-1}^{\text{b}, z}(\phi) \right] \left[\prod_{\mathbf{k} \in \mathbb{I}_{t|t-1}^{\text{m}}} \bar{p}_{t|t-1}^{\mathbf{k}, \text{d}}(\phi) \right].$$

Proof. We first rewrite the multi-target marginal likelihood w_t^I in a suitable way, for any $I \in \mathcal{A}_t$. For this purpose let $J_I^{\text{m}, \text{d}}$ be the subset of $\mathbb{I}_{t|t-1}^{\text{m}}$ corresponding to targets with index in I that have been detected at time t , let $Z_I^{\text{m}, \text{d}}$ be the subset of Z_t containing the corresponding observations, let $\sigma_I : J_I^{\text{m}, \text{d}} \rightarrow Z_I^{\text{m}, \text{d}}$ be the one-to-one mapping describing this identification, and let Z_I^{u} and Z_I^{b} be the subsets of Z_t containing the observations associated with the never-detected targets and with false alarms respectively, then it holds that

$$Z_t = Z_I^{\text{m}, \text{d}} \uplus Z_I^{\text{u}} \uplus Z_I^{\text{b}},$$

and w_t^I can be expressed as

$$w_t^I = C_t^{\phi} \times \left[\prod_{(\mathbf{k}, z) \in \text{Gr}(\sigma_I)} \frac{p_{t|t-1}^{\mathbf{k}, \text{d}}(z)}{\bar{p}_{t|t-1}^{\mathbf{k}, \text{d}}(\phi)} \right] \left[\prod_{z \in Z_I^{\text{u}}} \frac{p_{t|t-1}^{\text{u}, \text{d}}(z)}{\bar{p}_{t|t-1}^{\text{u}, \text{d}}(\phi)} \right] \left[\prod_{z \in Z_I^{\text{b}}} \frac{\bar{p}_{t|t-1}^{\text{b}, z}(z)}{\bar{p}_{t|t-1}^{\text{b}, z}(\phi)} \right]$$

where $\text{Gr}(\sigma_I) \doteq \{(\mathbf{k}, \sigma_I(\mathbf{k})) : \mathbf{k} \in J_I^{\text{m}, \text{d}}\}$ is the graph of σ_I . We can proceed to the second step of the proof by rewriting W_t as follows

$$W_t = C_t^{\phi} \sum_{I \in \mathcal{A}_t^{\text{m}}} \left[\prod_{(\mathbf{k}, z) \in \text{Gr}(\sigma_I)} \frac{p_{t|t-1}^{\mathbf{k}, \text{d}}(z)}{\bar{p}_{t|t-1}^{\mathbf{k}, \text{d}}(\phi)} \right] \times \left[\sum_{Z_{\text{u}}, Z_{\text{b}} \subseteq Z_t : Z_{\text{u}} \uplus Z_{\text{b}} = Z_t - Z_I^{\text{m}, \text{d}}} \left[\prod_{z \in Z_{\text{u}}} \frac{p_{t|t-1}^{\text{u}, \text{d}}(z)}{\bar{p}_{t|t-1}^{\text{u}, \text{d}}(\phi)} \right] \left[\prod_{z \in Z_{\text{b}}} \frac{\bar{p}_{t|t-1}^{\text{b}, z}(z)}{\bar{p}_{t|t-1}^{\text{b}, z}(\phi)} \right] \right].$$

We conclude by noticing that the sum over Z_{u} and Z_{b} has a binomial form and can thus be factorised, so that

$$W_t = C_t^{\phi} \sum_{I \in \mathcal{A}_t^{\text{m}}} \left[\prod_{(\mathbf{k}, z) \in \text{Gr}(\sigma_I)} \frac{p_{t|t-1}^{\mathbf{k}, \text{d}}(z)}{\bar{p}_{t|t-1}^{\mathbf{k}, \text{d}}(\phi)} \right] \left[\prod_{z \in Z_t - Z_I^{\text{m}, \text{d}}} C_t^{\text{u}, \text{b}}(z) \right], \\ = C_t^{\phi} \left[\prod_{z \in Z_t} C_t^{\text{u}, \text{b}}(z) \right] \\ \times \sum_{I \in \mathcal{A}_t^{\text{m}}} \left[\prod_{(\mathbf{k}, s, z) \in \sigma_t^{-1}[I]} \frac{p_{t|t-1}^{\mathbf{k}, \text{d}}(z)}{\bar{p}_{t|t-1}^{\mathbf{k}, \text{d}}(\phi) C_t^{\text{u}, \text{b}}(z)} \right].$$

We finally notice that

$$\prod_{\mathbf{k} \in \mathbb{I}_{t|t-1}^{\text{m}}} \left[\bar{p}_{t|t-1}^{\mathbf{k}, \text{d}}(\phi) + \sum_{z \in Z_t} \frac{p_{t|t-1}^{\mathbf{k}, \text{d}}(z)}{C_t^{\text{u}, \text{b}}(z)} \right] \\ = \left[\prod_{\mathbf{k} \in \mathbb{I}_{t|t-1}^{\text{m}}} \bar{p}_{t|t-1}^{\mathbf{k}, \text{d}}(\phi) \right] \sum_{I \in \mathcal{A}_t^{\text{m}}} \left[\prod_{(\mathbf{k}, s, z) \in \sigma_t^{-1}[I]} \frac{p_{t|t-1}^{\mathbf{k}, \text{d}}(z)}{\bar{p}_{t|t-1}^{\mathbf{k}, \text{d}}(\phi) C_t^{\text{u}, \text{b}}(z)} \right]$$

where the product on the r.h.s. has been developed and the obtained terms have been simplified using **A.3**, which completes the proof of the first part of the lemma. The alternative expression of W_t under **A.4** can be proved in a similar way. \square