



Novel Multi-Object Filtering Approach for Space Situational Awareness

Emmanuel Delande*

Heriot-Watt University, Edinburgh, Scotland EH14 4AS, United Kingdom

Carolyn Frueh†

Purdue University, West Lafayette, Indiana 47907

and

Jose Franco,‡ Jérémie Houssineau,§ and Daniel Clark¶

Heriot-Watt University, Edinburgh, Scotland EH14 4AS, United Kingdom

DOI: 10.2514/1.G002067

Surveillance activities with ground-based assets in the context of space situational awareness are particularly challenging. The observation process is indeed hindered by short observation arcs, limited observability, missed detections, measurement noise, and contamination by clutter. This paper exploits a recent estimation framework for stochastic populations for space situational awareness surveillance scenarios. This framework shares the flexibility of the finite set statistics framework in the modeling of a dynamic population of objects and the representation of all the sources of uncertainty in a single coherent probabilistic framework and the intuitive approach of traditional track-based techniques to describe individual objects and maintain track continuity. We present a recent multi-object filtering solution derived from this framework, the filter for distinguishable and independent stochastic populations, and propose a bespoke implementation of the multitarget tracking algorithm for a space situational awareness surveillance activity. The distinguishable and independent stochastic populations filter is tested on a surveillance scenario involving two ground-based Doppler radars in a challenging environment with significant measurement noise, limited observability, missed detections, false alarms, and no a priori knowledge about the number and the initial states of the objects in the scene. The tracking algorithm shows good performance in initiating tracks from object-generated observations and in maintaining track custody throughout the scenario, even when the objects are outside of the sensors' fields of view, despite the challenging conditions of the surveillance scenario.

Nomenclature

α^y	=	probability of existence of previously detected target with observation path y
$c(H)$	=	probability of existence of hypothesis H
H	=	hypothesis, i.e., subset of compatible tracks in Y
(H, n)	=	multitarget configuration, describing a possible composition of population $\mathcal{X}$
$H_{t t-1}, H_t$	=	set of all hypotheses, propagated by the filter for distinguishable and independent stochastic populations (DISP)
$\ell(z, x)$	=	likelihood of association between target with state x and observation with state z
$P_{t t-1}, P_t$	=	law of whole population of targets, propagated by the distinguishable and independent stochastic populations filter
P^a	=	law of population of appearing targets
$p_d(x)$	=	probability of detection of target with state x

$p_{fa}(z)$	=	probability that observation with state z is false alarm
p^y	=	probability distribution of previously detected target with obs. path y
p^ϕ	=	probability distribution of each yet-to-be-detected target
$w(H, n)$	=	joint probability of existence of targets in configuration (H, n)
$\mathcal{X}$	=	population of targets
X	=	target state space
$\bar{X}$	=	target state space augmented with empty state ψ
x	=	target state (e.g., position and velocity coordinates)
$Y_{t t-1}, Y_t$	=	set of all observation paths, propagated by distinguishable and independent stochastic populations filter
y	=	observation path or track characterized by said observation path
Z	=	set of observations
$\underline{Z}$	=	observation space
$\bar{Z}$	=	set of observations augmented with empty observation ϕ
z	=	observation (e.g., radar attributable)
μ_n^ϕ	=	multitarget configuration measure, describing states of n yet-to-be-detected targets
$\mu^{H,n}$	=	multitarget configuration measure, describing states of targets in configuration (H, n)
ρ^ϕ	=	cardinality distribution describing number of yet-to-be-detected targets
ϕ	=	empty observation
ϕ	=	empty observation path
ψ	=	empty state (target)

Presented at the Proceedings of the AAS Space Flight Mechanics Meeting, Williamsburg, VA, 11–15 January 2015; received 29 March 2016; revision received 20 May 2017; accepted for publication 20 May 2017; published online Open Access 17 July 2017. Copyright © 2017 by the American Institute of Aeronautics and Astronautics, Inc. All rights reserved. All requests for copying and permission to reprint should be submitted to CCC at www.copyright.com; employ the ISSN 0731-5090 (print) or 1533-3884 (online) to initiate your request. See also AIAA Rights and Permissions www.aiaa.org/randp.

*Research Associate, School of Engineering and Physical Sciences; E.D.Delande@hw.ac.uk.

†Assistant Professor, School of Aeronautics and Astronautics; cfrueh@purdue.edu.

‡Ph.D. Student, School of Engineering and Physical Sciences; jf139@hw.ac.uk.

§Research Associate, School of Engineering and Physical Sciences; j.houssineau@hw.ac.uk.

¶Associate Professor, School of Engineering and Physical Sciences; D.E.Clark@hw.ac.uk.

I. Introduction

MAINTEINING a catalog of objects of interest in orbit around the Earth is one of the major challenges within the context of space situational awareness (SSA). Typically, SSA surveillance

activities focus on the observation of a number of objects of interest in a surveillance scene, i.e., the near-Earth space containing the orbits on which the objects evolve, in order to estimate their individual characteristics of interest or states (e.g., their position and velocity coordinates in some reference frame). Since the number of orbiting objects is unknown and time varying, a surveillance activity in the context of SSA can be cast as a multi-object estimation problem of which the salient features are familiar to the tracking community. The operator possesses partial information on the dynamical behavior of the objects of interest, and the observations (or measurements) produced by sensors are the primary source of information on which the catalog of objects is maintained and updated. In the general case, the observation process is hindered by a limited sensor field of view (FOV), measurement noise, missed detections when objects in the sensor FOV do not produce observations, and, conversely, false alarms when collected observations do not originate from orbiting objects. In addition, the origin object of each collected observation is usually unknown.

Some specific features of the SSA surveillance activities characterize them as particularly challenging multi-object estimation problems. The sensor FOVs cover a very limited fraction of the surveillance scene at any given time, and meaningful information on identified objects may have to be propagated over long periods of time without the availability of up-to-date observations. The typical sensors employed in the context of SSA do not provide information on all the components of an object state; for example, telescopes do not provide information on the range and range rate of the detected objects. Among other issues, this fact makes the initial orbit determination problem conspicuously more challenging. The complexity of orbital mechanics, highly nonlinear in nature and perturbed by conservative and nonconservative forces that depend on parameters that may not be considered in the object state (e.g., the object shape, attitude, or area-to-mass ratio), makes the modeling of appropriate dynamical models for the orbiting objects particularly challenging. On the bright side, the specifics of the orbital mechanics allow for the construction of various dynamical models offering a balance between simplicity and accuracy depending on the number of physical phenomena involved in its design (e.g., solar pressure, and the effect of the moon). It thus offers some leeway in the modeling of the multi-object estimation problem in order to reach a satisfying tradeoff between the available computational resources and the required accuracy in the construction of the catalog.

Quantifying the uncertainty in object estimates maintained by a catalog provides an important advantage to various operations, some inherent to the tracking activity itself, such as the prediction of an orbital trajectory, and some peripheral, such as the assessment of the probability of collision between two objects. For this reason, casting the SSA estimation problem in a Bayesian paradigm has become a topic of growing interest but remains relatively unexplored so far. Besides some nontraditional approaches in the context of tracking [1,2], it has been addressed chiefly through either track-based [3,4] or set-based approaches [5–11]. Popular track-based solutions include the multiple hypothesis tracking (MHT) and joint probabilistic density association (JPDA) filters and follow an intuitive construction in which sequences of observations that may represent the data originating from a single specific object are maintained as tracks. They do not maintain a probabilistic description of the dynamical evolution of the population of objects and rely on heuristics and expert knowledge in order to determine, for example, at which point a stream of observations is assumed to be sufficient evidence for the creation of new track, or at which point a track is considered lost. Set-based solutions [12,13], on the other hand, approach the multi-object estimation problem in a holistic manner and incorporate all the sources of uncertainty (size of the population of objects, individual states, observation process, etc.) in a unified probabilistic framework. The construction of specific multi-object tracking algorithms follows a principled derivation exploiting the finite set statistics (FISST) framework [13] and relies on specific modeling choices for the sources of uncertainty rather than heuristics. Popular set-based solutions such as the probability hypothesis density [14] or cardinalized probability hypothesis density [15] filters

represent the population of objects with a single random object known as a random finite set [13], the size (the number of objects) and elements (the individual states) of which are both random. They maintain a collective description of the whole population of objects and avoid an explicit (and costly) data association step between objects and collected observations but do not maintain track continuity since they do not propagate specific information on any object (i.e., individual tracks). Recent developments [16] within the FISST framework, however, aim at augmenting the set representation with a unique labeling of targets in order to maintain individual information on objects; applications in the context of SSA can be found in [9,17,18].

This paper exploits the recent estimation framework for stochastic populations [19] for SSA applications. Fully probabilistic in nature, it incorporates all the sources of uncertainty in a unified coherent probabilistic representation. Its main originality lies in the nature of the information representing each object, either described as an unidentified member of a larger population (e.g., any piece of debris from a cloud, before its observation) or as an individual track as soon as specific information is available (e.g. the piece of debris from the cloud that produced a given observation). This estimation framework thus shares the flexibility of the FISST framework for the modeling of populations of objects, for example, to translate the knowledge of the operator regarding the number of new objects entering in the scene, and shares the intuitive approach of traditional track-based techniques to describe individual objects and maintain track continuity. In this paper, we explore the suitability of the distinguishable and independent stochastic populations (DISP) filter [19,20] for a multi-object estimation problem in a SSA context. Exploiting the flexibility and adaptability of the DISP structure, we propose a bespoke implementation in which single-object time prediction and single-object/single-measurement data update components are adapted to the specific challenges of SSA problems.

The paper is organized as follows. The detailed construction of the generic DISP filter is covered in Sec. II, followed by the design and implementation of a multitarget DISP tracking algorithm for a problem in Sec. III. The algorithm is illustrated on a simulated scenario in Sec. IV, and a brief discussion on further developments to the DISP filter for SSA applications follows in Sec. V. Section VI concludes.

II. Multi-Object Estimation with DISP Filter

This section describes the DISP filter [19,20], designed to handle multitarget detection and tracking problems in challenging environments including targets evolving in proximity from each other in the presence of missed detections, false alarms, etc. For the sake of clarity and generality, the argument in this section is largely developed within the context of a generic multi-object estimation problem, while the practical implementation for a SSA surveillance activity follows in Sec. III.

A. General Principle

The objects of interest form a population $\mathcal{X}$, and they follow trajectories unknown to the operator throughout the scenario (i.e., unknown orbits). An object of the population $\mathcal{X}$ is not necessarily within the surveillance scene throughout the whole scenario (it may appear in the surveillance scene following a launch, leave the surveillance scene for deep space, etc.), and thus the number of objects in the scene, unknown to the operator, is also time varying. Each object currently in the surveillance scene is described by a state x in some (single-)target state space $X \subseteq \mathbb{R}^d$ describing physical characteristics of interest to the operator. The target state space typically comprises at least the object's position and velocity coordinates in some reference frame, but the DISP filter imposes little restriction on its nature. Conversely, an object outside of the surveillance scene is described by the empty state ψ , added to the target state space X to form the augmented state space $\tilde{X} = \{\psi\} \cup X$.

It is assumed that the only source of information about the specific objects from this population are the observations collected through some sensor from the beginning of the scenario onward, although the operator may possess some general information on the population of interest through expertise and past experience (regarding the

dynamical behavior of the objects, the density of objects in the surveillance scene, etc.). For the remainder of the paper, the time flow of the scenario is indexed by some variable t marking the instants during which observations are collected by the operator. Note that the time lapse between two consecutive time steps does not need to be constant throughout the scenario, and neither does the sensor profile characterizing the observation process at a given time step (measurement noise, probability of detection, etc.).

As any filter derived from the estimation framework for stochastic populations, the DISP filter offers a principled solution to a multi-object joint detection and tracking problem within the Bayesian paradigm, relying on a few well-defined assumptions. The principal assumptions are given as follows:

- 1) Assumption A1 is that the objects of the population $\mathcal{X}$ are independent.
- 2) Assumption A2 is that the observations are produced independently by the sensor.
- 3) Assumption A3 is that an object produces at most one observation per scan (if not, it is not detected).
- 4) Assumption A4 is that an observation originates from at most one object (if not, it is a false alarm).
- 5) Assumption A5 is that objects are immediately detected upon entering the surveillance scene.

Assumptions A2–A4 are specifically related to the single-sensor/multitarget problem, in which the objects are observed at any time step t by at most one sensor and are relatively common in this context. Note in particular that objects outside of the sensor FOV are, of course, almost never detected (i.e., with probability zero).

Assumption A5 implies that the DISP filter does not describe objects that have entered the surveillance scene since the beginning of the scenario and have not been detected at least once, and neither does it propagate information about these yet-to-be-detected objects. Since we assume that little prior information is available on these objects in the context of this paper, this assumption appears reasonable. Note that it does not preclude the initialization of the filter with information on objects that were present before the beginning of the scenario, for example, through a preexisting catalog (discussed in more details in Sec. II.E.1). A more general solution can be derived without requiring Assumption A5, but doing so makes the resulting filter significantly more complex to design and implement [19], and it is left out of the scope of this paper.

The DISP filter proposes a probabilistic representation of the population $\mathcal{X}$ on two different levels:

- 1) On the individual level, it estimates the state of each potential object, assuming that it exists.
- 2) On the population level, it estimates the composition of the population of objects; that is, it assesses the joint existence of the potential objects.

Section II.B focuses on the individual level and introduces the notion of track. Section II.C then focuses on the population level.

B. Target Representation

A potential object of the population identified by the filter is called a target, and the information maintained by the filter on some target state is represented with a probability distribution p on $\bar{\mathcal{X}}$. In particular, the scalar $p(\psi)$ denotes the probability that the target is currently outside of the surveillance scene and is called the target probability of absence. Conversely, the scalar

$$\int_{\mathcal{X}} p(x) dx = 1 - p(\psi) \quad (1)$$

denotes the probability of presence of the target, i.e., the probability that the target is currently within the surveillance scene.

1. Detected Targets

An important consequence of Assumptions A3 and A4 is that an observation, unless it is a false alarm, provides specific information on a single object of the population $\mathcal{X}$, and no other object. Thus, any target that has been detected at least once can be characterized by its

observation path [21], that is, the stream of observations it has produced across time. Specific information can then be maintained by the filter on such a target, which is thus called a distinguishable target [19].

For any time step $t \geq 0$ relevant to the scenario, the current set of observations collected by the operator is denoted by Z_t . We add the empty observation ϕ , to account for missed detections, to form the augmented observation set $\bar{Z}_t = Z_t \cup \{\phi\}$. Given some arbitrary time $t \geq 0$, an observation path [21] is a sequence of t observations

$$y = \left(\underbrace{\phi, \dots, \phi}_{t_\bullet - 1 \text{ times}}, z_{t_\bullet}, z_{t_\bullet+1}, \dots, z_t \right) \quad (2)$$

where $z_{t_\bullet} \in Z_{t_\bullet}$ and $z_{t'} \in \bar{Z}_{t'}$, $t_\bullet < t' \leq t$. It denotes a stream of observations that (potentially) originates from an object of the population $\mathcal{X}$ that appeared at time $t_\bullet$ and produced the observations $z_{t_\bullet}, \dots, z_t$ ever since. Following Assumption A5, targets are detected upon entering the scene, and that is why the first observation $z_{t_\bullet}$ is necessarily nonempty. There is no such restriction, however, on the subsequent observations.

The set of all possible observation paths at some given time t is denoted by Y_t . Figure 1 illustrates the composition of the set of observation paths Y_4 , for a given sequence of collected observation sets $Z_1, \dots, Z_4$.

Given in an observations path y , the DISP filter maintains a probabilistic description of the corresponding target through a track, indexed by y and represented by a probability distribution p^y on $\bar{\mathcal{X}}$. From now on, without ambiguity, the notation y may be used interchangeably to denote a distinguishable target, the corresponding observation path, or the corresponding track.

Then, the probability that the target y is in some region $B \subseteq \bar{\mathcal{X}}$, provided that the target exists, is given by the scalar

$$\int_B p^y(x) dx \quad (3)$$

2. Appearing Targets

Assumption A5 implies that the population of yet-to-be-detected objects is reduced to the objects that have just entered the surveillance scene. We can make use of the empty observation path at time $t - 1$, denoted ϕ_{t-1} and defined as

$$\phi_{t-1} = \left(\underbrace{\phi, \dots, \phi}_{t-1 \text{ times}} \right) \quad (4)$$

to identify the population of appearing (and yet-to-be-detected) targets at time t . Note that there is no ambiguity with the observation paths $y \in Y_{t-1}$ that are used to identify tracks, for the latter are of the form of Eq. (2) and thus contain at least one nonempty observation.

Since the appearing objects have not been detected yet, no specific information is available on any of them through collected observations. In consequence, the DISP filter only maintains information representing collectively the whole population of appearing targets, which are thus called indistinguishable targets [19]. We assume that the operator does possess some knowledge on




Fig. 1 The observation paths at time $t = 4$, given a sequence of collected measurements. The five possible observation paths are listed on the right.

the number of objects having entered the scene between time steps $t-1$ and t , described by some cardinality distribution $\rho^{\phi_{t-1}}$ on $\mathbb{N}$. They also possess some knowledge on the whereabouts of the entering objects (some orbit ranges, for example, are more populated than others), described by a common probability distribution $p^{\phi_{t-1}}$ on $\bar{X}$, representing each of the appearing targets since they are indistinguishable. Note that these targets are by definition in the scene, and thus $p^{\phi_{t-1}}(\psi) = 0$.

The multitarget configuration measure $\mu_n^{\phi_{t-1}}$, describing the case in which there are exactly n indistinguishable targets, is formally given by [19]

$$\mu_n^{\phi_{t-1}} = n\delta[p^{\phi_{t-1}}] \quad (5)$$

where $\delta[p]$ is the Dirac measure at point p . The population of appearing targets is then described by the law

$$P^a = \sum_{n \geq 0} \rho^{\phi_{t-1}}(n) \delta[\mu_n^{\phi_{t-1}}] \quad (6)$$

Example 1: Suppose that the population of appearing targets at time t is described with some given cardinality distribution $\rho^{\phi_{t-1}}$ and probability distribution $p^{\phi_{t-1}}$:

1) The probability that there is a single appearing target and it is in some region $B \subseteq X$ is given by the scalar

$$\rho^{\phi_{t-1}}(1) \left[\int_B p^{\phi_{t-1}}(x) dx \right] \quad (7)$$

2) The probability that either there is a single appearing target and it is some region $B \subseteq X$ or there are two appearing targets, one in B and the other one in some region $B' \subseteq X$, is given by the scalar

$$\begin{aligned} & \rho^{\phi_{t-1}}(1) \left[\int_B p^{\phi_{t-1}}(x) dx \right] \\ & + \rho^{\phi_{t-1}}(2) \left\{ \left[\int_B p^{\phi_{t-1}}(x) dx \right] \left[\int_{B'} p^{\phi_{t-1}}(x) dx \right] \right\} \end{aligned} \quad (8)$$

Note that, since the targets are indistinguishable from each other, the notion of target labeling (even arbitrary) is irrelevant in this context. An important consequence is that the evaluation of the event described by Eq. (8) does not involve permutations between the regions B and B' and the corresponding members of the population. As will be seen later, it spares the computation of unnecessary permutations in the filtering process, notably during the data association step (see Sec. II.E for more details). Note also that the law P^a describes the event in which the population is empty through the scalar $\rho^{\phi_{t-1}}(0)$. For this reason, the quantity $1 - \rho^{\phi_{t-1}}(0)$ is also called the probability of existence of the population.

C. Population Representation

On the population level, the DISP filter assesses the composition of the population $\mathcal{X}$ by providing the joint probability of existence of every combination of tracks, including none, with every possible number of appearing targets, including none as well. We have seen in the previous section that the law $P^{\phi_{t-1}}$ of the population of appearing targets describes the joint probability of existence of any number of appearing targets through the cardinality distribution $\rho^{\phi_{t-1}}$; let us then focus on the joint probability of existence of the tracks.

It is important to note that, while every possible observation path $y \in Y_t$, the construction of which follows Eq. (2), is a valid characterization of a previously detected object, not every subset $Y \subseteq Y_t$ forms a valid characterization of the population of previously detected objects. Indeed, any two targets of which the observation paths share a nonempty observation may not exist simultaneously without violating Assumption A4 and are thus incompatible. The subsets of Y_t composed of pairwise compatible tracks are called hypotheses and form a set H_t maintained by the DISP filter.

In particular, the empty set $\emptyset \in H_t$ is the empty hypothesis representing the absence of previously detected targets.

Example 2: Suppose that, at time step $t = 4$, the measurement sets collected so far are as illustrated in Fig. 1.

The pairs of incompatible observation paths are 1) (y, y'') , $(y, y^{(4)})$, and $(y'', y^{(4)})$, sharing measurement z_4 ; 2) (y', y'') , sharing measurement z_2 ; and 3) $(y''', y^{(4)})$, sharing measurement $z_{2'}$.

Therefore, the hypotheses in H_t are 1) the empty set $\emptyset$; 2) the singletons $\{y\}$, $\{y'\}$, $\{y''\}$, $\{y'''\}$, and $\{y^{(4)}\}$; 3) the two-element sets $\{y, y'\}$, $\{y, y'''\}$, $\{y', y'''\}$, $\{y', y^{(4)}\}$, and $\{y'', y'''\}$; and 4) the three-element set $\{y, y', y'''\}$.

The joint probabilities of existence of the subset of tracks $Y \subseteq Y_t$ thus reduce to the probabilities of existence of the hypotheses $H \in H_t$. The probability of existence $c_t(H)$ of some hypothesis $H \in H_t$ denotes the likelihood that the tracks $y \in H$ represent the previously detected objects of the population of interest $\mathcal{X}$, and we have

$$\sum_{H \in H_t} c_t(H) = 1 \quad (9)$$

In particular, the scalar $c_t(\emptyset)$ denotes the probability that no objects of the population have been detected so far; in other words, that all the observations collected so far are false alarms.

The joint probability of existence of the whole population at time t is thus given by the weights

$$w_t(H, n) = c_t(H) \rho_t^{\phi_t}(n), \quad H \in H_t, \quad n \in \mathbb{N} \quad (10)$$

and from Eq. (9), we have

$$\sum_{H \in H_t} \sum_{n \geq 0} w_t(H, n) = 1 \quad (11)$$

D. Filter Output

From Secs. II.B and II.C, it follows that the DISP filter maintains a full probabilistic representation of the population of objects $\mathcal{X}$, maintaining specific information on potential objects through individual distributions and assessing the joint existence of every possible composition of potential objects. More formally, the population $\mathcal{X}$ is described at time t with the law [19]

$$P_t = \sum_{H \in H_t} \sum_{n \geq 0} w_t(H, n) \delta[\mu_t^{H,n}] \quad (12)$$

where the multitarget configuration measure $\mu_t^{H,n}$ is defined as

$$\mu_t^{H,n} = \sum_{y \in H} \delta[p_t^y] + n\delta[p_t^{\phi_t}] \quad (13)$$

While the multitarget configuration composed of the tracks given by some hypothesis H and some number n of appearing targets is completely described by the measure $\mu_t^{H,n}$, it can also be characterized by the tuple (H, n) without ambiguity. In the rest of the paper, both characterizations will be used interchangeably.

Note that a track $y \in Y_t$ may belong to several hypotheses $H \in H_t$; if we marginalize the probability of existence of y over all the possible mutually compatible tracks, we get the scalar

$$\alpha_t^y = \sum_{\substack{H \in H_t \\ H \ni y}} c_t(H) \quad (14)$$

which denotes the credibility or probability of existence of the target y . It is a scalar between 0 and 1; the closer it is to 1, the more likely it is that the observation path y represents the evidence of an object of the population $\mathcal{X}$. Note the difference with the probability of presence given by Eq. (1), which denotes the probability that the target y is still in the surveillance scene, provided that it exists.

Example 3: Suppose that the current time step is $t = 5$, but the observation set Z_5 is not available yet, and the observation sets collected so far are as illustrated by Fig. 1:

1) The probability that there are neither appearing nor detected targets is given by the scalar

$$w_{5|4}(\emptyset, 0) \quad (15)$$

2) The probability that either a) there are no appearing targets and there is a single detected target that produced the observation path y and lies within some region $B \subseteq X$ or b) there is a single appearing target, lying in B , with no detected targets is given by the scalar

$$w_{5|4}(\{y\}, 0) \int_B p_{5|4}^y(x) dx + w_{5|4}(\emptyset, 1) \int_B p_{5|4}^{\phi_5}(x) dx \quad (16)$$

3) The probability that either there are no appearing targets and there are two detected targets, one produced the observation path y and the other one produced the observation path y' , which have both left the surveillance scene, or there are two appearing targets, both lying within some region $B \subseteq X$, and there is a single detected target that produced the observation path y and that lies within some region $B' \subseteq X$ is given by the scalar

$$w_{5|4}(\{y, y'\}, 0) [p_{5|4}^y(\psi) p_{5|4}^{y'}(\psi)] + w_{5|4}(\{y\}, 2) \left\{ \left[\int_B p_{5|4}^{\phi_5}(x) dx \right]^2 \left[\int_{B'} p_{5|4}^y(x) dx \right] \right\} \quad (17)$$

4) The probability that there are three appearing targets, one lying within some region $B \subseteq X$ and two within some region $B' \subseteq X$, and there are two detected targets, one that produced the observation path y and either lies within B or has left the surveillance scene and the other that produced the observation path y' and lies within B' is given by the scalar

$$w_{5|4}(\{y, y'\}, 3) \left[\int_B p_{5|4}^{\phi_5}(x) dx \right] \left[\int_{B'} p_{5|4}^{\phi_5}(x) dx \right]^2 \times \left[p_{5|4}^y(\psi) + \int_B p_{5|4}^y(x) dx \right] \left[\int_{B'} p_{5|4}^{y'}(x) dx \right] \quad (18)$$

We see in the previous examples that the law of the population P_t provides a rich description of the population $\mathcal{X}$ and can be exploited to assess events of various complexity regarding the composition of the population and the states of the identified objects. In this paper, we will focus on the display of the most probable configuration of the population (see Sec. II.F), but other exploitation routes will be briefly discussed in Sec. V.

E. Filtering Equations

The filtering equations maintain and propagate a probabilistic representation of the population of interest $\mathcal{X}$ within a Bayesian paradigm. They cover two roles:

- 1) On the population level, they propagate the composition of the population through the set of tracks Y_t , the set of hypotheses H_t , and the associated weights w_t .
- 2) On the individual level, they propagate the single-target probability distributions p_t^y .

Following Secs. II.A–II.C, the data flow of the DISP filter can be depicted as in Fig. 2.

1. Initialization (Time t_0)

The origin of time $t = 0$ corresponds to the date since which observations have been collected and stored for the purpose of maintaining a catalog of orbiting objects. The operator may thus possess some knowledge describing the objects of the population of interest $\mathcal{X}$ that were detected before the initial time $t_0 \geq 0$ of the surveillance scenario; this is typically the case if the initial input of the filter is fed with a catalog of previously identified objects. If the




Fig. 2 Data flow of the DISP filter (time step t). The superscript d denotes the population of previously detected, hence distinguishable, targets. The superscript a denotes the population of appearing and yet-to-be-detected targets, still indistinguishable from one another.

catalog does provide, say, a collection of three compatible tracks with observations paths y, y', y'' , then the initial set of hypotheses can be set as a singleton $H_{t_0-1} = \{\{y, y', y''\}\}$; that is, the tracks y, y' , and y'' almost surely exist. If no catalog is available, then the initial set of hypotheses can be reduced to the empty set; that is, the initial population of previously identified targets is almost surely empty.

2. Prediction Step (Time t)

The prediction step of the DISP filter is straightforward. Since the observation set Z_t is not available yet, there is no additional evidence regarding the existence of targets, and thus the information on the composition of the population $\mathcal{X}$ (i.e., the observation paths Y_{t-1} , the hypotheses H_{t-1} , and the associated weights w_{t-1}) remains unchanged.

On the individual level, however, the probability distribution p_{t-1}^y of each track $y \in Y_{t-1}$ is transformed to its predicted form $p_{t|t-1}^y$ as

$$\begin{cases} p_{t|t-1}^y(x) = \int_X p_{s,t}(x') \hat{m}_{t,t-1}(x, x') p_{t-1}^y(x') dx', & x \in X, \\ p_{t|t-1}^y(\psi) = \int_X [1 - p_{s,t}(x')] p_{t-1}^y(x') dx' + p_{t-1}^y(\psi) \end{cases} \quad (19)$$

where $\hat{m}_{t,t-1}$ denotes the transition kernel representing the evolution of the state of the objects within the scene since time $t - 1$ (e.g., the motion model along an orbit) and $p_{s,t}$ denotes their probability of survival in the scene.** Both functions are model parameters that reflect the knowledge of the operator regarding the dynamical behavior of the objects.

Note that the single-target prediction mechanism on the state space X , as described in Eq. (19), follows the usual Bayesian prediction step in single-object tracking problems. For example, if the probability of survival $p_{s,t}$ is uniform and set to 1, the transition kernel $\hat{m}_{t,t-1}$ is linear, and the prior distribution on X induced by p_{t-1}^y is Gaussian, then the prediction equation on the state space X reduces to the prediction step of the Kalman filter [22]. We can also see from Eq. (19) that the predicted probability of absence $p_{t|t-1}^y(\psi)$ is no smaller than its prior value $p_{t-1}^y(\psi)$, since no new evidence of the target presence is available yet through new observations.

3. Update Step (Time t)

The core of the update step consists in the data association step, in which potential sources of observations are matched with the newly collected observation set Z_t , and the likelihood of each association is assessed. The potential sources of observations are 1) the previously detected objects, represented by the tracks; 2) the appearing objects, represented by a population of indistinguishable targets; and 3) the false alarm generators. For the purpose of data association, a false alarm generator is modeled for each collected observation $z \in Z_t$ through a false alarm function $p_{fa,t}$: the generator produced the measurement z with probability $p_{fa,t}(z)$ (in which case z is a false alarm), or it produced no observation with probability $1 - p_{fa,t}(z)$ (in which case z originates from an object of the population $\mathcal{X}$). The modeling of the false alarm function is, of course, highly dependent on the type of sensor considered in the scenario.

**See Delande, E. D., Houssineau, J., and Clark, D. E., "A Filter for Distinguishable and Independent Populations," arXiv:1501.04671, 2015.




Fig. 3 Data association, for a hypothesis $H \in H_{t-1}$ and a given number n of appearing targets. Note the absence of permutation involving the appearing targets, since they are indistinguishable.

Assume a possible multitarget configuration (H, n) , in which $H \in H_{t-1}$ describes the tracks representing the previously detected targets and $n \in \mathbb{N}$ is the number of appearing targets. An admissible data association d matches the three sources of observations described previously to the collected observations $z \in Z_t$ and the empty observation ϕ (see Fig. 3). The set of all admissible data associations is the set (see footnote **)

$$\text{Adm}_{Z_t}(H, n) = \left\{ (H_d, Z_d, Z_a, \nu) \mid H_d \subseteq H, Z_d \subseteq Z_t, Z_a \subseteq Z_t \setminus Z_d, |Z_a| = n, \nu \in S(H_d, Z_d) \right\} \quad (20)$$

where H_d denotes the tracks that are detected at the present time step, Z_d denotes the observations associated to these tracks, Z_a denotes the observations associated to the n appearing targets, and ν denotes the bijective function associating detected tracks to observations in Z_d .

Each association scheme $a = (H, n, d)$, in which $d \in \text{Adm}_{Z_t}(H, n)$ is an admissible data association for the predicted multitarget configuration (H, n) , is then assessed by the scalar (see footnote **)

$$v_t^a = v_{d,t}^a \times v_{md,t}^a \times v_{fa,t}^a \quad (21)$$

where

$$\begin{aligned} v_{d,t}^a &= \left[\prod_{z \in Z_d} \int_X \ell_t(z, x) p_{d,t}(x) p_{t|t-1}^{\phi_{t-1}}(x) dx \right] \\ &\times \left\{ \prod_{y \in H_d} \int_X \ell_t[\nu(y), x] p_{d,t}(x) p_{t|t-1}^y(x) dx \right\}, \\ v_{md,t}^a &= \prod_{y \in H \setminus H_d} \left\{ p_{t|t-1}^y(\psi) + \int_X [1 - p_{d,t}(x)] p_{t|t-1}^y(x) dx \right\}, \\ v_{fa,t}^a &= \left\{ \prod_{z \in Z_d \cup Z_a} [1 - p_{fa,t}(z)] \right\} \left[\prod_{z \in Z_t \setminus (Z_d \cup Z_a)} p_{fa,t}(z) \right] \end{aligned} \quad (22)$$

where $v_{d,t}^a$, $v_{md,t}^a$, and $v_{fa,t}^a$ denote the weights of the data association related to the targets that are detected this time step, those who are not, and the clutter generators, respectively. Also, $p_{d,t}$ denotes the probability of detection, and ℓ_t denotes the measurement likelihood function, the modeling of which is highly dependent on the type of sensor considered in the scenario.

Each association scheme $a = (H, n, d)$ leads to the construction of a unique updated hypothesis $\hat{H} \in H_t$ given by

$$\hat{H} = \left[\bigcup_{y \in H_d} \{y: \nu(y)\} \right] \cup \left[\bigcup_{y \in H \setminus H_d} \{y: \phi\} \right] \cup \left[\bigcup_{z \in Z_a} \{\phi_{t-1}: z\} \right] \quad (23)$$

where $:$ is the concatenation operator, i.e., $(e_1, \dots, e_n): e = (e_1, \dots, e_n, e)$. Using Bayes's rule, the probability of existence $c_t(\hat{H})$ is found to be (see footnote **)

$$c_t(\hat{H}) = \frac{w_{t|t-1}(H, n) v_t^a}{\sum_{H' \in H_{t-1}} \sum_{n' \geq 0} w_{t|t-1}(H', n') \sum_{d' \in \text{Adm}_{Z_t}(H', n')} v_t^{a'}} \quad (24)$$

The updated observation paths, as shown in Eq. (23), can be of three different forms. An updated observation path of the form $y: z$, where $y \in Y_{t-1}$ and $z \in Z_t$, corresponds to a predicted track y that was detected and produced observation z . Following Bayes's rule, the probability distribution of the track is updated as follows (see footnote **):

$$\begin{cases} p_t^{y:z}(x) = \frac{p_{d,t}(x) \ell_t(z, x) p_{t|t-1}^y(x)}{\int_X p_{d,t}(x') \ell_t(z, x') p_{t|t-1}^y(x') dx'}, & x \in X, \\ p_t^{y:z}(\psi) = 0 \end{cases} \quad (25)$$

As for the prediction step in Eq. (19), the single-measurement/single-target update mechanism on the state space X given by Eq. (19) follows the usual Bayesian update step in single-object tracking problems. If the probability of detection $p_{d,t}$ is uniform and set to 1, the measurement likelihood function ℓ_t is linear, and the predicted distribution on X induced by $p_{t|t-1}^y$ is Gaussian, then the update equation on the state space X reduces to the update step of the Kalman filter [22].

Note that the probability of presence of the track $y: z$ in the scene rises to 1; that is, if there is an object in the population $\mathcal{X}$ that produced the observation path $y: z$, then this object is in the almost surely scene. Indeed, only objects that are currently in the surveillance scene can be detected by the sensor; therefore, an object that has just been detected is in the surveillance scene.

An updated observation path of the form $y: \phi$, where $y \in Y_{t-1}$, corresponds to a predicted track y that was not detected this time step. Following Bayes's rule, the probability distribution of the track is updated as follows (see footnote **):

$$\begin{cases} p_t^{y:\phi}(x) = \frac{[1 - p_{d,t}(x)] p_{t|t-1}^y(x)}{p_{t|t-1}^y(\psi) + \int_X [1 - p_{d,t}(x')] p_{t|t-1}^y(x') dx'}, & x \in X, \\ p_t^{y:\phi}(\psi) = \frac{p_{t|t-1}^y(\psi)}{p_{t|t-1}^y(\psi) + \int_X [1 - p_{d,t}(x)] p_{t|t-1}^y(x) dx} \end{cases} \quad (26)$$

Note that the denominator in Eq. (26) is no greater than 1; i.e., $p_t^{y:\phi}(\psi)$ is no smaller than $p_{t|t-1}^y(\psi)$. Indeed, since the target was not detected this time step, no new evidence on its presence in the scene is available, and the probability of presence of the target is nonincreasing.

Finally, a new observation path of the form $\phi_{t-1}: z$, where $z \in Z_t$, denotes an appearing target that was detected and produced observation z . The update process is similar to the previously detected targets in Eq. (25) and yields (see footnote **)

$$\begin{cases} p_t^{\phi_{t-1}:z}(x) = \frac{p_{d,t}(x) \ell_t(z, x) p_{t|t-1}^{\phi_{t-1}}(x)}{\int_X p_{d,t}(x') \ell_t(z, x') p_{t|t-1}^{\phi_{t-1}}(x') dx'}, & x \in X, \\ p_t^{\phi_{t-1}:z}(\psi) = 0 \end{cases} \quad (27)$$

Note that, for the same reason as for detected targets in Eq. (25), the probability of presence of appearing targets is always 1.

F. Maximum a Posteriori Estimate of Population (Time t)

We shall focus on the production of the maximum a posteriori (MAP) estimate of the posterior population, i.e., the most probable multitarget configuration that can be extracted from the propagated law P_t of the population $\mathcal{X}$. By construction, it is given by the hypothesis H_t^* with the highest probability of existence, i.e.,

$$H_t^* = \arg \max_{H \in H_t} c_t(H) \quad (28)$$

It is not always desirable to display all the tracks $y \in H_t^*$, especially for a cautious operator who wishes to focus on credible targets of which the existence can be guaranteed with a certain level of confidence. Since the probability of existence of individual tracks (14) can be extracted from the filter's output as well, candidate tracks may be prevented from being displayed unless their probability of existence exceeds some confirmation threshold τ_c , e.g., $\tau_c = 0.95$ if the operator wishes to focus on tracks of which the existence is guaranteed by the filter with 95% of confidence. To prevent flickering in the display, a lower deconfirmation threshold τ_d can be added to design a simple hysteretic model.

Once the tracks $Y_t^* \subseteq H_t^*$ to be displayed have been selected, one must produce a MAP estimate of the state of each selected track $y \in Y_t^*$. The DISP filter does not impose any restriction on the nature of the single-target probability distribution p_t^y , and in practical scenarios, the single-target MAP extraction procedure depends heavily on whether the single-target distributions are maintained through sequential Monte Carlo (SMC) or Kalman filtering (KF) methods. A procedure adapted to a typical SSA scenario is proposed in Sec. III.D.

G. Filtering Approximations

From the structure of the DISP filter, it can be seen that neither tracks nor hypotheses are ever discarded (unless their probability of existence drops to zero), and thus the filter maintains in its most general form all the possible multitarget configurations given the collected observations since the beginning of the scenario. An obvious consequence is that the number of hypotheses $|H_t|$ and tracks $|Y_t|$ grows dramatically across time, increasing the computational and memory load of the filtering process. Fortunately, the full probabilistic nature of the filtering framework allows the operator to design intuitive and principled approximations in order to discard selected pieces of information.

Since the confidence in hypotheses and tracks is directly available through their probability of existence [see Eqs. (9) and (14)], either a) discarding the hypotheses and/or tracks of which the confidence drops below a given threshold, or b) discarding the hypotheses and/or tracks with lowest confidence until a given threshold on their number is met, are sensible and straightforward approximations.

Measurement gating is a very popular method in tracking algorithms involving data association, which aims at discarding unlikely associations between a specific target and a specific observation. It is an effective approximation, for it lowers considerably the number of generated tracks with limited consequences on the quality of the estimation, and its implementation in the context of the DISP filter is straightforward. The associations schemes (H, n, d) generated in the data update step, such that the weight of the association between an observation z and a parent track y in Eq. (22) fails to meet a given threshold, are simply discarded in the construction of the updated hypotheses.

While the DISP filter can handle the simultaneous appearance of any number of targets in the data update mechanism, allowing so significantly increases the number of generated hypotheses. Limiting the number of simultaneous appearances is a reasonable approximation in most practical situations because it will simply stage the creations of tracks along several successive time steps when an unusually large number of objects appear simultaneously but would not affect the quality of the tracks in the long term in a significant way. Its implementation for the DISP filter is straightforward: the multitarget configurations (H, n) in which n exceeds a given threshold are simply discarded in the construction of the updated hypotheses. This approximation can be lifted during the initial step, in which the occurrence of simultaneous appearances is more likely than in the subsequent time steps; indeed, all the targets within the initial FOV are bound to appear simultaneously (provided that they are detected) as soon as the sensor is switched on.

III. Design and Implementation of SSA Scenario

This section describes the design and implementation of a simulated SSA scenario, processed by the DISP filter described in Sec. II. Section III.A focuses on the generation of the ground truth, and Sec. III.B focuses on the production of pseudoobservations from a Doppler radar. Section III.C then focuses on the implementation of the DISP filter. Finally, Sec. III.D describes the procedure to extract the most probable multitarget configuration from the DISP filter's output.

A. Surveillance Scene and Objects

The surveillance scene is set as the region of space spanning 70,000 km in every direction from the Earth center (space fixed), in order to include all objects of interest evolving on the usual orbits in the near-Earth space. An object is described by a six-component vector state x in some single-target state space $X \subseteq \mathbb{R}^6$, describing its coordinates in the Earth-centered inertial (ECI) frame.

The start epoch is 53159.5 days (modified Julian date). The interval between two time steps is constant and set to 50 s, and the scenario lasts 1000 time steps so that the propagation duration is approximately 14 h. The ground truth trajectories are computed from the initial orbital elements using a Runge–Kutta seven-eights numerical integration with Earth gravitational expansion up to order and degree 12, third-body perturbations of the sun and moon, direct solar radiation pressure, and Earth shadow passages. For the Earth shadow passages, a multicone model has been employed. A mixture of flat plate models and cannonball shapes has been used. Details on the propagation models can be found in [23].

B. Sensors and Pseudoobservations

We consider two ground-based Doppler radars, both placed at a latitude of 20° north, one at a longitude of 135° west and the other one at a longitude of −135° west, and both pointed at the zenith. The FOV are identical and define a radarlike sensing region, and the probability of detection $p_{d,t}$ is assumed uniform across the sensor FOV.

Each observation produced by a sensor is a four-component vector z in some observation space $Z \subseteq \mathbb{R}^4$, corresponding to the components range ρ , range rate $\dot{\rho}$, azimuth θ , and elevation φ in the topocentric local horizon frame centered on the sensor. Observations originating from objects are noisy; the measurement noise is modeled as an uncorrelated zero-mean Gaussian noise on the four components of the observation vector.

False alarms occur independently in each cell of a resolution grid of the observation space (see Table 1 for the details of the parameterization). The probability of false alarm $p_{fa,t}$ is assumed identical in each cell.

The two sensors are not synchronized and produce observations alternatively at each time step. Thus, the implemented scenario can be cast within the single-sensor version of the DISP presented in Sec. II since Assumption A3, constraining the observation process to

Table 1 Parameterization of the Doppler radar

Detection profile	
Field of view, range ρ	50 m $\leq \rho \leq 45 \times 10^6$ m
Field of view, azimuth θ	$-8^\circ \leq \theta \leq 8^\circ$
Field of view, elevation φ	$-45^\circ \leq \varphi \leq 45^\circ$
Probability of detection $p_{d,t}$	0.9
Resolution profile	
Cell resolution, range ρ	1000 m
Cell resolution, azimuth θ	1°
Cell resolution, elevation φ	1°
Cell resolution, range rate $\dot{\rho}$	100 ms ^{−1}
Probability of false alarm $p_{fa,t}$	1.54×10^{-10} (2 f.a. on avg.)
Noise profile	
Standard deviation, range ρ	1000 m
Standard deviation, azimuth θ	1°
Standard deviation, elevation φ	1°
Standard deviation, range rate $\dot{\rho}$	100 ms ^{−1}

produce at most one observation per object, is not violated. Note that the DISP filtering framework can accommodate for any number of synchronized sensors, but this is out of the scope of this paper.

C. SMC-DISP Filter

Note that the construction of the DISP filter in Sec. II does not impose restrictions on the nature of the single-object probability distributions. The form of these distributions has little consequence on the structure of the filter, as it only affects the evaluation of the integrals on the single-object state space $\mathbf{X}$ in Eqs. (19), (22), (25–27).

The two widely used implementation techniques for handling single-object probability distributions are KF-based [24] and SMC-based [25] approaches. The orbit dynamics being highly nonlinear, the latter techniques seem more adapted to capturing the targets' motion behavior and representing the uncertainties in the targets' state alongside their orbit, typically assuming bananalike shapes. We have opted for a particle filter implementation inspired by the recent disparity space approach [26,27].

1. Prediction Step and Orbit Propagation

As explained in Sec. II.E.2, the prediction step of the DISP filter is straightforward, as it does not involve any changes in the composition of the set of observation paths $\mathbf{Y}_{t-1}$, the set of hypotheses $\mathbf{H}_{t-1}$, nor their corresponding weights w_{t-1} . The prediction step focuses on the propagation of the posterior distributions p_{t-1}^y of the tracks $y \in \mathbf{Y}_{t-1}$ following Eq. (19).

Assume that some track $y \in \mathbf{Y}_{t-1}$ has a posterior probability distribution p_{t-1}^y on the augmented state space $\bar{\mathbf{X}}$ approximated by a collection of weighted particles $\{(w_{t-1}^{(j)}, \tilde{x}_{t-1}^{(j)})\}_{j \in J_{t-1}}$ on the state space $\mathbf{X}$, i.e.,

$$\begin{cases} p_{t-1}^y(x) \simeq \sum_{j \in J_{t-1}} w_{t-1}^{(j)} \delta_{\tilde{x}_{t-1}^{(j)}}(x), & x \in \mathbf{X}, \\ p_{t-1}^y(\psi) \simeq 1 - \sum_{j \in J_{t-1}} w_{t-1}^{(j)} \end{cases} \quad (29)$$

Each particle $\tilde{x}_{t-1}^{(j)}$ represents a possible state for the track y , provided that it exists and that it is still in the surveillance scene. To propagate the probability distribution p_{t-1}^y , every particle $\tilde{x}_{t-1}^{(j)}$ is propagated alongside its orbit to a new value $\tilde{x}_t^{(j)}$ through a simplified propagation model, using the simple two-body problem via the Shepperd matrix [28,29] implementation. We have

$$\tilde{x}_t^{(j)} = \mathbf{S}|_{\tilde{x}_{t-1}^{(j)}} \tilde{x}_{t-1}^{(j)} + \mathbf{n}|_{\tilde{x}_{t-1}^{(j)}}, \quad j \in J_{t-1} \quad (30)$$

where $\mathbf{S}|_{\tilde{x}_{t-1}^{(j)}}$ is the Shepperd transition matrix evaluated in $\tilde{x}_{t-1}^{(j)}$ and $\mathbf{n}|_{\tilde{x}_{t-1}^{(j)}}$ is the process noise accounting for the mismatches between the simplified model exploited by the filter and the higher-fidelity model exploited for the generation of the ground truth trajectories (see Sec. III.A). The process noise follows a Gaussian distribution with zero mean and noise characteristics that are state independent, i.e.

$$\mathbf{n}|_{\tilde{x}_{t-1}^{(j)}} \sim \mathcal{N}(\cdot; 0, \mathbf{Q}_t), \quad j \in J_{t-1} \quad (31)$$

where

$$\mathbf{Q}_t = \Delta_{t-1,t} \text{diag}([10 \ 10 \ 10 \ 1 \ 1 \ 1]) \quad (32)$$

where $\Delta_{t-1,t}$ is the duration, in seconds, between time step $t-1$ and t .

The objects of the population are assumed to evolve on closed orbits; thus, no spontaneous disappearances of objects within the surveillance scene are expected. The probability of survival $p_{s,t}$ is thus set to 1 in the surveillance scene; however, predicted particles $\tilde{x}_t^{(j)}$ falling outside of the state space $\mathbf{X}$ reach the empty state ψ (typically, when the track is getting close to the edges of the surveillance scene), and their weight is naturally added to the prior probability of absence $p_{t-1}^y(\psi)$. That is, the predicted probability distribution p_{t-1}^y on the augmented state space $\bar{\mathbf{X}}$ is given by

$$\begin{cases} p_{t|t-1}^y(x) \simeq \sum_{j \in \tilde{J}_{t-1}} w_{t-1}^{(j)} \delta_{\tilde{x}_t^{(j)}}(x), & x \in \mathbf{X}, \\ p_{t|t-1}^y(\psi) \simeq 1 - \sum_{j \in \tilde{J}_{t-1}} w_{t-1}^{(j)} \end{cases} \quad (33)$$

where $\tilde{J}_{t-1} = \{j \in J_{t-1} | \tilde{x}_t^{(j)} \in \mathbf{X}\}$.

2. Update Step and Orbit Determination

The construction of the updated observation paths $\mathbf{Y}_t$ and hypotheses $\mathbf{H}_t$, explained in detail in Sec. II.E.3, is independent of the specifics of the scenario and the implementation choices for the single-object probability distributions. On the other hand, the implementation of the single-object update mechanisms for 1) the tracks that are not detected (26), 2) those that are detected (25), and 3) the newborn targets (27) is obviously SMC specific and shall be detailed in the following paragraphs. For the remainder of the section, $y \in \mathbf{Y}_{t-1}$ denotes a track with a predicted probability distribution $p_{t|t-1}^y$ approximated by a collection of weighted particles $\{(w_{t-1}^{(j)}, \tilde{x}_t^{(j)})\}_{j \in \tilde{J}_{t-1}}$ on the state space $\mathbf{X}$, as shown in Eq. (33).

a. *Undetected Track.* From the update Eq. (26), the posterior distribution $p_t^{y:\phi}$ can be approximated as follows:

$$\begin{cases} p_t^{y:\phi}(x) \simeq \sum_{j \in \tilde{J}_{t-1}} \frac{[1 - p_{d,t}(\tilde{x}_t^{(j)})] w_{t-1}^{(j)}}{p_{t-1}^y(\psi) + \sum_{i \in \tilde{J}_{t-1}} [1 - p_{d,t}(\tilde{x}_t^{(i)})] w_{t-1}^{(i)}} \delta_{\tilde{x}_t^{(j)}}(x), & x \in \mathbf{X}, \\ p_t^{y:\phi}(\psi) \simeq \frac{p_{t-1}^y(\psi)}{p_{t-1}^y(\psi) + \sum_{i \in \tilde{J}_{t-1}} [1 - p_{d,t}(\tilde{x}_t^{(i)})] w_{t-1}^{(i)}} \end{cases} \quad (34)$$

To avoid particle degeneracy [30], it may be necessary to resample the updated particle set $\{(\tilde{w}_{t-1}^{(j)}, \tilde{x}_t^{(j)})\}_{j \in \tilde{J}_{t-1}}$ given by Eq. (34) to produce the final particle set $\{(w_t^{(j)}, x_t^{(j)})\}_{j \in J_t}$ approximating the posterior distribution $p_t^{y:\phi}$.

b. *Detected Track.* From the update, Eq. (25), the posterior distribution $p_t^{y:z}$ can be approximated as follows:

$$\begin{cases} p_t^{y:z}(x) \simeq \sum_{j \in \tilde{J}_{t-1}} \frac{p_{d,t}(\tilde{x}_t^{(j)}) \mathcal{L}_t(z, \tilde{x}_t^{(j)}) w_{t-1}^{(j)}}{\sum_{i \in \tilde{J}_{t-1}} p_{d,t}(\tilde{x}_t^{(i)}) \mathcal{L}_t(z, \tilde{x}_t^{(i)}) w_{t-1}^{(i)}} \delta_{\tilde{x}_t^{(j)}}(x), & x \in \mathbf{X}, \\ p_t^{y:z}(\psi) = 0 \end{cases} \quad (35)$$

In SSA scenarios, it is frequent that tracks remain unobserved for long periods of time before being detected again, typically, when the target has rotated along Earth and reenters the sensor FOV. In this case, the set of particles approximating the parent track (33) is often spread along an orbit following a bananalike shape, and the SMC update (35) fails at capturing the posterior distribution $p_t^{y:z}$ with a reasonable number of particles. Inspired by recent works on multi-object tracking applications with cameras [26,27], we propose an alternative update process in the extended observation space (defined in the following) rather than the target state space $\mathbf{X}$ [31].

Since a Doppler radar provides information on range ρ , range rate $\dot{\rho}$, azimuth θ , and elevation φ (see Sec. III.B), an observation vector $z = (\rho, \theta, \varphi, \dot{\rho})$ can be seen as four components of a point in the six-dimension spherical frame centered on the sensor, i.e., the topocentric local horizon frame of the sensor or extended observation space. One can then note the following:

1) The projection of the probability distribution $p_{t|t-1}^y$ in the extended observation space can be realistically approximated as a multivariate Gaussian distribution, even though the probability distribution $p_{t|t-1}^y$ in the native ECI reference frame may not.

2) Because of the specifics of the sensor model, the observation process in the extended observation space is linear with added Gaussian noise.




Fig. 4 Principle of Kalman data update in extended observation space (illustration in two dimensions). The blue distribution in the reference space **A** corresponds to some predicted probability distribution $p_{t|t-1}^y$, and the red one corresponds to the updated probability distribution $p_t^{y:z}$, where $z \in Z_t$ is a range-only measurement depicted, with uncertainty, in the observation space **C**.

In consequence, the data update of the probability distribution $p_{t|t-1}^y$ with observation $z \in Z_t$ can be approximated through a simple Kalman filter update [22] in the extended observation space. The proposed procedure can be summarized as follows:

- 1) Transform the weighted particle set $\{(p_{d,t}(\tilde{x}_t^{(j)})w_{t-1}^{(j)}, \tilde{x}_t^{(j)})\}_{j \in \tilde{J}_{t-1}}$, approximating the predicted probability distribution $p_{t|t-1}^y$ rescaled with the sensor probability of detection, to the extended observation space.
- 2) Approximate the resulting weighted particle set into a multivariate Gaussian distribution on the extended observation space.
- 3) Update the resulting Gaussian distribution with likelihood $\ell_t(z, \cdot)$, using a simple Kalman update [22] step.
- 4) Sample the resulting Gaussian distribution into a particle set, with uniform weights, on the extended observation space.
- 5) Transform the resulting particle set back to the target state space **X**.

The resulting particle set $\{(w_t^{(j)}, x_t^{(j)})\}_{j \in J_t}$ is thus an approximation of the posterior distribution $p_t^{y:z}$ of the daughter track $y:z$. The procedure is illustrated in a two-dimensional space (for the sake of clarity) on Fig. 4.

c. Initial Orbit Determination. A comparison of Eqs. (25) and (27) shows that the updated probability distribution $p_{t|t-1}^{y:z}$ of a newborn track can be constructed with a similar approach as shown previously for a daughter track $y:z$ resulting from a track detection. The key difference is that the initialization of a newborn track depends on the prior knowledge of the operator regarding the state of appearing objects given by the probability distribution $p_{t|t-1}^{\phi_{t-1}}$.

Since a Doppler radar provides no information on the angular rates $\dot{\theta}$ and $\dot{\varphi}$ of the detected objects (see Sec. III.B), it is of critical importance to initialize the probability distribution of the newborn tracks with an appropriate amount of uncertainty on the unknown

angular rates in order to constrain the range of orbits on which the newborn target is evolving, until observations produced in subsequent times will (hopefully) allow for a refinement of the orbit determination.

We assume that the prior information on appearing objects $p_{t|t-1}^{\phi_{t-1}}$ is noninformative regarding their position; that is, every point in the portion of the surveillance scene (position only) delimited by the sensor FOV may be the origin state of a newborn target with equal probability. Considerations on the orbital energy constraint state, however, that for given position coordinates (x, y, z) in the ECI frame not every velocity coordinates $(\dot{x}, \dot{y}, \dot{z})$ form an admissible state, that is, a state belonging to a physical orbit around Earth. More specifically, given the position and range rate coordinates of an appearing object provided by the Doppler radar through an observation $z = (\rho, \theta, \varphi, \dot{\rho})$, an admissible region for the unobserved angular rates $(\dot{\theta}, \dot{\varphi})$ can be produced to determine constraints on the physical values of the initial orbit [32,33]. Assuming that the observed object is on a closed orbit around the Earth, the energy constraint

$$\frac{1}{2} |\dot{\mathbf{r}}|^2 - \frac{\mu}{|\mathbf{r}|} \leq 0 \quad (36)$$

where μ is the Earth gravitational parameter and $\mathbf{r}$ is the object's geocentric position vector, needs to hold. The position and velocity $\mathbf{r}$ and $\dot{\mathbf{r}}$ can then be expressed via the position vector of the topocenter $\mathbf{r}_s$, the observed range ρ , and the observed range rate $\dot{\rho}$,

$$\begin{aligned} \mathbf{r} &= \mathbf{r}_s + \rho \mathbf{p}_r \\ \dot{\mathbf{r}} &= \dot{\mathbf{r}}_s + \dot{\rho} \mathbf{p}_r + \rho \dot{\theta} \mathbf{p}_\theta + \rho \dot{\varphi} \mathbf{p}_\varphi \end{aligned} \quad (37)$$

where $\mathbf{p}_r$, $\mathbf{p}_\theta$, and $\mathbf{p}_\varphi$ are the unit vectors in the topocentric local horizon spherical reference frame associated with the measurement

$(\rho, \theta, \varphi, \dot{\rho})$. Substituting Eq. (37) into Eq. (36) yields the new form

$$\alpha_1 \dot{\theta}^2 + \alpha_2 \dot{\varphi}^2 + \alpha_3 \dot{\theta} + \alpha_4 \dot{\varphi} + a_5 \leq 0 \quad (38)$$

where the parameters α_i are as follows:

$$\begin{aligned} \alpha_1 &= \rho^2 \cos^2 \varphi, \quad \alpha_4 = \frac{\rho \dot{r}_s \cdot \rho_\varphi}{2}, \\ \alpha_2 &= \rho^2, \quad \alpha_5 = \dot{\rho}^2 + 2\dot{\rho} \dot{r}_s \cdot \rho_r + |\dot{r}_s|^2 - \frac{2\mu}{\sqrt{\rho^2 + 2\rho \dot{r}_s \cdot \rho_r + |\dot{r}_s|^2}}, \\ \alpha_3 &= \frac{\rho \dot{r}_s \cdot \rho_\theta}{2} \end{aligned} \quad (39)$$

From Eq. (38), the boundaries of the admissible region for the unknown parameters (θ, φ) are found to be [34]

$$\begin{aligned} \dot{\theta} &= \frac{\alpha_3}{\alpha_1} + \sqrt{\frac{\alpha_3^2}{\alpha_1^2} + \frac{\alpha_4^2}{\alpha_1 \alpha_2} - \frac{\alpha_5}{\alpha_1}} \cos \phi \\ \dot{\varphi} &= \frac{\alpha_4}{\alpha_2} + \sqrt{\frac{\alpha_3^2}{\alpha_1^2} + \frac{\alpha_4^2}{\alpha_1 \alpha_2} - \frac{\alpha_5}{\alpha_1}} \sin \phi \end{aligned} \quad (40)$$

with $\phi \in [0, 2\pi)$.

The initial probability distribution $p_{t-1}^{\phi_{t-1}:z}$ of a newborn track $\phi_{t-1}:z$ can then be determined following a procedure similar to that explained previously for detected tracks and illustrated in Fig. 4:

1) Determine a multivariate Gaussian distribution in the extended observation space, using observation z and the sensor observation noise to fill the observed components, and fill the unobserved components so that the admissible region (40) is adequately covered; for example, select means $\hat{\theta} = (\alpha_3/\alpha_1)$, $\hat{\varphi} = (\alpha_4/\alpha_2)$, and standard deviations $\sigma_\theta = \sqrt{(\alpha_3^2/\alpha_1^2) + (\alpha_4^2/\alpha_1 \alpha_2) - (\alpha_5/\alpha_1)}$, $\sigma_\varphi = \sqrt{(\alpha_3^2/\alpha_1^2) + (\alpha_4^2/\alpha_2^2) - (\alpha_5/\alpha_2)}$.

2) Sample the resulting Gaussian distribution into a particle set, with uniform weights, on the extended observation space.

3) Transform the resulting particle set to the target state space $\mathbf{X}$.

The resulting particle set $\{(w_t^{(j)}, x_t^{(j)})\}_{j \in J_t}$ is thus an approximation of the posterior distribution $p_t^{\phi_{t-1}:z}$ of the newborn track $\phi_{t-1}:z$.

D. Track Display Procedure

The procedure exposed in Sec. II.F allows for the selection of the candidate tracks $Y_t^* \subseteq Y_t$ for display, all members of the most probable multitarget configuration. We shall now propose a procedure to extract the MAP state x_t^* of the posterior distribution p_t^y of some candidate track $y \in Y_t^*$, i.e., the most probable state of the object of the population $\mathcal{X}$ represented by y , provided that it exists and is currently in the scene.

Let $\{(w_t^{(j)}, x_t^{(j)})\}_{j \in J_t}$ be the set of weighted particles approximating the posterior distribution p_t^y on the augmented state space $\tilde{\mathbf{X}}$. The MAP x_t^* is naturally given by the particle with the highest weight, i.e.,

$$x_t^* = x_t^{(j^*)}, \quad j^* = \arg \max_j \{w_t^{(j)}\}_{j \in J_t} \quad (41)$$

The data update mechanisms for undetected tracks and detected tracks presented in Sec. III.C.2 are such that the particle weights in the posterior distribution $\{(w_t^{(j)}, x_t^{(j)})\}_{j \in J_t}$ have close values. In consequence, the highest particle weight is often shared by several particles, and producing the MAP using Eq. (41) becomes unreliable in practice.

Following reasoning similar to that shown in Sec. III.C.2, the spherical reference frame corresponding to the ECI frame (i.e., sharing the same origin and axes) can be seen as a candidate frame in which the posterior distribution p_t^y can be approximated as a multivariate Gaussian distribution, and a substitute to the MAP given by Eq. (41) can thus be produced through the following procedure:

1) Transform the coordinates of the weighted particle set $\{(w_t^{(j)}, x_t^{(j)})\}_{j \in J_t}$, approximating the posterior distribution p_t^y , into the spherical reference frame.

2) Approximate the resulting weighted particle set into a multivariate Gaussian distribution on the spherical reference frame.

3) Produce the mean (equal to the MAP) of the resulting Gaussian distribution.

4) Transform the coordinates of the resulting point into the ECI frame.

The procedure is illustrated in a two-dimensional space (for the sake of clarity) on Fig. 5. Note that this method does not produce the MAP, in the general case (the red circle does not match the blue circle in the illustrated example), but it proves to be more consistent than the method extracting the exact MAP for the reasons explained previously. Note that the extraction of the most probable multitarget configuration has no effect on the filtering process, since the DISP filter propagates the full probabilistic description of the population through the law P_r .

IV. Simulation Results

In this section, we illustrate the DISP filter on a simulated SSA scenario exploiting the design and implementation techniques presented in Sec. III.

A. Parameterization of SSA Scenario

1. Ground Truth

We consider a SSA scenario as depicted in Fig. 6, with five objects evolving on different orbits. All the objects are present in the surveillance scene at the beginning of the scenario, and their orbits are contained within the surveillance scene. The initial orbital elements and area-to-mass ratios of the five objects are listed in Table 2.

This scenario presents several challenges. Object 1 is on an inclined geostationary orbit and remains in one sensor FOV




Fig. 5 Display of the most probable state for some track $y \in Y_t$ (illustration in two dimensions). The blue circle is the MAP given by the particle approximation of the posterior distribution p_t^y (blue points). The blue cross is the mean state given by the particle approximation. The MAP (and mean) of the Gaussian approximation of the posterior distribution in the polar reference frame (red circle and cross, on the right-hand side) can be transformed back to the Cartesian reference frame (red circle and cross, on the left-hand side) to produce an approximation of the MAP state of track y .




Fig. 6 Illustration of the scenario. Initial positions (left figures) of the five objects are marked by stars, and the last positions (right figures) are marked by crosses. The gray region denotes the sensor FOV.

throughout the whole duration of the scenario. Object 2 is also on a retrograde inclined geostationary orbit. Objects 1 and 2 cross each other's paths when they are covered by one of the sensors. Since a Doppler radar does not provide angular rates, the observations produced by these two objects are expected to be close in value. Objects 3 and 4 are on low Earth orbits and do not stay in the sensor FOV for long periods of time. Object 5 is on a high-eccentricity geostationary transfer orbit. It is initially in one of the sensor FOV, leaves it after a few times steps, and reenters it once during the scenario.

Table 2 Parameterization of the scenario (object 1 is on an inclined geostationary orbit (GEO), object 2 is on a retrograde inclined GEO, objects 3 and 4 are on low Earth orbits (LEO), and object 5 is on a high-eccentricity geostationary transfer orbit (GTO))

Object number	1	2	3	4	5
Orbit type	GEO	GEO	LEO	LEO	GTO
Right ascension of the ascending node Ω , deg	61	50	20	30	60
Inclination i , deg	10	170	2	3	11.3
Argument of perigee ω , deg	349	30	311	250	351
Semi-major axis a , km	42164	42164	16495	17490	26352
Eccentricity e	0.012	0.01	0.01	0.3	0.6
True anomaly ν , deg	82	20	80	50	80
Area-to-mass ratio, m ² /km	0.02	0.6	0.02	0.02	0.02

2. Doppler Radar

For the sake of simplicity, the two sensors are given the same profile, defined as in Table 1.

3. SMC-DISP Filter

The details of the SMC implementation of the DISP filter are given in Table 3.

4. Output Extraction

The candidate tracks for display are selected using a simple hysteretic model based on the probability of existence, as explained in Sec. II.D. The confirmation threshold τ_c is set to 0.999, and the deconfirmation threshold τ_d is set to 0.8.

B. Results and Discussion

1. Methodology

One of the key abilities of a Bayesian multi-object filter is to deliver an estimation with the appropriate level of uncertainty. In usual single-object tracking problems, the Cramér-Rao lower bound [35] provides a measure of the minimum uncertainty reachable by an unbiased estimator for given sensor capabilities. This concept, however, does not extend easily to the multi-object estimation problem with unknown data association, missed detections, false alarms, etc.

Table 3 Parameterization of the SMC-DISP filter

Filtering (population level)	
Set of hypotheses at initialization	$\emptyset$
Number of appearing targets ($t = 0$)	Depends on Z_0
Number of appearing targets ($t > 0$)	$\rho_{t t-1}^{\phi_{t-1}}(0) = 1 - 10^{-6}$ $\rho_{t t-1}^{\phi_{t-1}}(1) = 10^{-6}$
Filtering (individual level)	
Process noise for target motion model	$\Delta_{t-1,t} \text{diag}([10 \ 10 \ 10 \ 1 \ 1 \ 1])$
Number of particles per track	200
Approximation (information loss)	
Max number of hypotheses	30
Min. prob. of existence for hypotheses	10^{-30}
Min. prob. of existence for tracks	10^{-7}
Measurement gating (confidence level)	95%

To provide a reference for the evaluation of the DISP filter, we designed the unrealistic yet informative Oracle filter, modeled as a DISP filter fed with the correct data association at each update step, that is, with the origin of each collected observation. The Oracle filter thus propagates only one hypothesis, corresponding to the correct multitarget configuration. Note that it does not provide the ground truth of the scenario, for its estimation relies on the same modeling choices as the DISP filter and in particular on the same level of uncertainty regarding the dynamical behavior of the objects of interest, the same measurement noise in the observation process, the missed detections, etc.

To facilitate the analysis of the results, the output of the filter is displayed separately for each of the five objects involved in the scenario. By construction, each track maintained by the Oracle filter corresponds to an identified object. There is no such straightforward correspondence in the case of the DISP filter, since the data association is unknown and part of the estimation problem. For the sake of analysis, each track produced by the DISP filter is associated to the closest object (in the Euclidean position subspace) at its time of birth (corresponding to the natural association that would be given by a human operator).

For each object, the analysis of the DISP filter is illustrated through the number of tracks associated to this object (in Fig. 7), the quality of the tracking through the distance between the object and the closest




Fig. 7 Cardinality error averaged over 22 MC runs. Blue zones correspond to time periods in which the object lies within the sensor FOV. Dark lines refer to the Oracle filter, and bright lines refer to the DISP filter (GEO, geostationary orbit; LEO, low Earth orbits; GTO, geostationary transfer orbit).

associated track (in Fig. 8), and the number of swappings between the associated tracks (in Fig. 9).

Both filters are compared on the same input data, i.e., the same pseudoobservations collected from the simulated sensors as explained in Sec. III.B. The results shown in this section are averaged over 22 Monte Carlo (MC) runs of the scenario described in Sec. IV.A, each run being based on the same ground truth but generating different pseudoobservations.

2. Cardinality Error

Figure 7 illustrates the cardinality error for each object, penalizing the absence of a track or the presence of redundant tracks associated to a given object.

By construction, the Oracle filter produces a track upon the first detection of each object and maintains a single track until the end of




Fig. 8 Track quality averaged over 22 MC runs. Blue zones corresponds to time periods in which the object lies within a sensor FOV. The Euclidean distance between an object and the MAP estimate of the corresponding track is depicted in green, and the Mahalanobis distance between an object and the spatial distribution of the corresponding track is depicted in red. Dark lines refer to the Oracle filter, and bright lines refer to the DISP filter (GEO, geostationary orbit; LEO, low Earth orbits; GTO, geostationary transfer orbit).




Fig. 9 Track swaps (averaged over 22 MC runs). Blue zones correspond to time periods in which the object lies within a sensor FOV. A track swap occurs whenever there is a change in association between an object and the closest track maintained by the DISP filter (GEO, geostationary orbit; LEO, low Earth orbits; GTO, geostationary transfer orbit).

the scenario. Thus, its cardinality error is one before the first detection of each object and drops to zero upon first detection. The DISP filter, on the other hand, must discriminate between false alarms and object-generated observations in order to assess the appearance of new objects in the sensor FOV; in consequence, the initial cardinality error of 1 does not drop instantaneously to zero following the first detection of each object. The figure shows that in roughly 20% of the runs the DISP filter failed at detecting object 3 when it entered in the sensor FOV for the first time, though it detected it the second time. It also shows that in roughly 30% of the runs the DISP filter failed at detecting object 4, resulting in a nonzero cardinality error throughout the scenario. Note that the sensor coverage of object 4 was particularly poor, for the object never stayed for an extended and continuous period of time within reach of the sensors. In these conditions, the initial orbit determination procedure was not always successful at producing a stable track. We also see on this figure that the DISP filter did not always identify object 5 upon reentry in the sensor FOV around time step 560: in roughly 70% of the runs, the initial track was not updated with the new observations, and another track was created instead, resulting in a cardinality error. This situation will be further discussed in the next paragraph.

3. Track Quality

Figure 8 illustrates the quality of the tracks maintained for each object.

The Euclidean distance between an object and the MAP estimate of the corresponding track provides an intuitive picture of the quality of the tracks. Overall, the tracks produced by the DISP filter and the Oracle filter show similar quality, the latter filter providing a baseline for the targeted performance on this particular scenario. It must be noted that quality of tracks degrades quickly once the object has left the sensor FOV, regardless of the filter employed. This suggests that either the initial observation windows are too short for the initialization of a stable track or the prediction model implemented through the Shepperd's matrix is not accurate enough to propagate the track along a stable orbit, or both. We see in particular that object 5 stays outside of the sensor FOV for over 500 time steps following its first detection. The discrepancies between the propagated track and the object are significant, and upon object 5's reentry in the sensor FOV, the DISP ponders the possibility that the newly collected observations are originating from a sixth object, unobserved so far. In roughly 70% of the runs, the existence of a sixth object is assessed as more likely, and an additional track is created (and increases the cardinality error, as discussed in the previous paragraph).

The Mahalanobis distance between an object state and the spatial distributions of the corresponding track illustrates the uncertainty carried by the filters regarding the estimation [36]. A value exceeding the $n\%$ threshold of a six-dimensional chi-square distribution indicates that the object lies outside of the region containing $n\%$ of the mass of the spatial distribution of the track. The relatively low values of the Mahalanobis distances observed throughout the scenario for all the objects show that both filters are consistent and show no sign of overconfidence. In other words, the quality of the tracks degrades quickly when the objects evolve through blind regions with no sensor coverage, but the filters know so. This illustrates the ability of the DISP to provide an appropriate level of uncertainty in its estimation of the multitarget configuration on this particular scenario, in which the observability is relatively poor (see the sensor parameters in Table 1).

4. Track Swaps

Figure 9 illustrates the number of track swaps during the scenario, for each object.

When the DISP maintains several tracks for a given object, a track swap occurs whenever the association between the object and the closest of these tracks, measured with the Euclidean distance on the position coordinates, changes. As discussed in the two previous paragraphs, track swaps typically occurs whenever the DISP (wrongly) assumes that a

sixth object is present in the scene, due to the growing discrepancies between object 5 and the initial track maintained by the DISP filter. We see on this figure that the track swaps are limited to the period following the creation of the additional track. The fact that no track swaps occur in the subsequent times suggests that the two tracks are not competing; the initial track has been definitively put aside in favor of the more recent one, which was created closer to the new observations through a new initial orbit determination procedure.

The fact that the cardinality error does not drop to zero in the subsequent times (see object 5 in Fig. 7) shows that the redundant track remains displayed by the DISP filter, highlighting a challenging situation common to SSA scenarios in the presence of poor sensor coverage. Once the two tracks produced from object 5 have left the sensor FOV, no information is available for the filter to correct its estimation and reassess their probability of existence until they reenter a sensor FOV and eventually discard the redundant track.

5. Computational Complexity

The DISP maintains the joint probabilities of existence of all the possible combinations of compatible tracks that could potentially represent the true multitarget configuration in the surveillance scene, or hypotheses. The bulk of the computational complexity originates from the number of hypotheses maintained and the generation of the admissible data associations (20). The data association is highly combinatorial in nature, though reduced in practice by measurement gating of which the efficacy depends in particular on whether false alarms fall close enough to the objects to induce ambiguous data associations. The computational complexity of the DISP filter thus rises dramatically with the number of objects and false alarms in a manner that remains difficult to quantify in the general case, perhaps similarly to the MHT filter. The simulations in this paper, run on a laptop with an unoptimized MATLAB® code and some embedded C code, required several hours per Monte Carlo run. A natural extension of this work to larger-scale scenarios run in realistic computational times is discussed in Sec. V, through a principled approximation of the DISP filter of which the complexity is linear in the number of objects and observations.

V. Further Work

An important feature of the DISP filter is the well-separated mechanisms estimating the composition of the population $\mathcal{X}$ and the individual states of its identified members. In particular, the structure of the filter is flexible and easily adaptable to different modeling choices relating to the single-object time prediction step or the single-object/single-measurement data update step. Essentially, single-object mechanisms are modeling blocks that can be replaced without altering the general structure of the filter with respect to the composition of the tracks, the hypotheses, and the related probabilities of existence.

The results presented in Sec. IV show that the quality of the track estimates drop quickly once they leave the sensor FOV and the filtering process relies solely on the prediction model (19) to estimate the motion of the objects along their orbits. In this paper, we implement a linearization of the orbital mechanics in the six-dimensional Cartesian target state space $\mathbf{X}$ through the Shepperd transition matrix with additional state-independent noise (33), but alternative solutions exist. The DISP filter imposes few constraints on the nature of the individual target state space $\mathbf{X}$, and the modeling of the targets' dynamical behavior should be explored in other spaces that are more suited to SSA tracking problems. Exploiting a target space describing the orbital elements of the objects rather than their position and velocity coordinates in a Cartesian reference frame is a promising lead, since the orbital mechanics would be more naturally described and modeled in such a space. This modeling change would have little impact on the implementation of the filter, save the prediction equation (19) itself. In particular, the same data update procedure could be applied, with an adapted mapping between the target state space and the extended sensor space (steps 1 and 5 in Fig. 4).

Provided that an initial orbit determination procedure is available and Assumptions A2–A4 hold, sensors of different nature can be incorporated into the filtering framework presented in this paper with a bespoke sensor profile (observation space $\mathbf{Z}$, shape of the FOV, probability of detection $p_{d,t}$, etc.). Again, the nature of the sensor has little impact on the structure of the filter, save the evaluation of the integrals in Eqs. (24–27). To that respect, the modeling of a telescope providing information on angles and angular rates has been recently explored in [31] for a subsequent incorporation to the DISP filter.

On a broader perspective, the fully probabilistic nature of the DISP filter has yet to be exploited for SSA applications, and the law of the population P_t propagated by the filter can be exploited beyond the extraction the most probable multitarget configuration exploited shown in Sec. II.D. Recent developments within the estimation framework for stochastic populations [19], such as target classification [37] or information-theoretic policies for closed-loop sensor management [38], are fully relevant to the context of SSA and directly applicable to the DISP filter. The regional statistics for multi-object filters [39], estimating the number of objects with associated uncertainty in any desired region of the state space, could be exploited in the context of a SSA scenario to provide, for example, a dynamical assessment of the density of objects on particular orbits. Originally developed within the FISST framework, the regional statistics could be adapted to any filter derived from the estimation framework for stochastic populations [19] as well, including the DISP filter.

Maintaining a catalog of orbiting objects will involve tracking a large number of targets, and the complexity of the DISP filter increases dramatically with the number of objects involved in the scenario. The work presented in this paper paves the way for the implementation of the filter for hypothesised and independent stochastic populations (HISP) [40]^{††} for SSA scenarios, a principled approximation of the DISP filter of which the complexity is linear in the number of objects and observations and that is particularly adapted to tracking problems such as SSA scenarios in which the data association is moderately ambiguous. Since the single-object mechanisms are the same as in the DISP filter, the specific developments proposed in Sec. III can be directly incorporated to the HISP filter for SSA scenarios [41]. The two filtering solutions shall be compared to analyze the cost in terms of accuracy induced by the greater efficiency of the HISP filter in the context of SSA surveillance activities.

VI. Conclusions

This paper presents the distinguishable and independent stochastic populations (DISP) filter, a recent multi-object joint detection/tracking algorithm derived from the estimation framework for stochastic populations, in the context of wide-area surveillance for a space situational awareness (SSA) application. It shows that a Bayesian filtering framework can be exploited in order to maintain a probabilistic description of a population of orbiting objects, assessing the composition of the population and the individual states of its identified members, with associated uncertainty.

The DISP filter is illustrated on a multi-object surveillance scenario involving five objects of interest evolving on different orbits, with no prior information available on their number nor their individual states, observed by two sensors with limited observability and field of view and hindered by missed detections, measurement noise, and false alarms. In this scenario, the DISP filter proves to be responsive in the creation of tracks following their first detection and robust in maintaining track custody despite the challenging conditions.

A few leads are provided for the improvement of the tracking capabilities within the scope of the DISP filter, including an exploitation of the orbital elements to capture the orbital mechanics and the incorporation of a wider range of sensor models suitable to SSA applications. The adaptation of this work to large scale scenarios is also discussed through the implementation of the hypothesised and independent stochastic populations filter, a principled approximation

of the DISP filter of which the computational complexity is linear in the number of objects and observations.

Acknowledgments

This work was supported by the Engineering and Physical Sciences Research Council (EPSRC) grants EP/K014277/1 and EP/J015180/1, the EPSRC/Dstl University Defence Research Collaboration (UDRC) in Signal Processing, the Dstl funding contract DSTL/AGR/00363/01 (ED-TIN10), and the AFSOR European Office of Aerospace Research and Development Grant 13-3030. Carolin Frueh was supported through the Vison and roBOTics (VIBOT) Erasmus Mundus Visiting Scholar programme.

References

- [1] Frueh, C., and Schildknecht, T., "Object Image Linking of Objects in Near Earth Orbits in the Presence of Cosmics," *Advances in Space Research*, Vol. 49, No. 3, 2012, pp. 594–602. doi:10.1016/j.asr.2011.10.021
- [2] Frueh, C., *Identification of Space Debris*, Shaker Verlag, Aachen, Germany, 2011.
- [3] Kelec, T., Jah, M. K., and DeMars, K. J., "Application of a Multiple Hypothesis Filter to near GEO High Area-to-Mass Ratio Space Objects State Estimation," *Acta Astronautica*, Vol. 81, No. 2, 2012, pp. 435–444. doi:10.1016/j.actaastro.2012.08.006
- [4] Singh, N., Poore, A., Sheaff, C., Aristoff, J., and Jah, M. K., "Multiple Hypothesis Tracking (MHT) for Space Surveillance: Results and Simulation Studies," *Advanced Maui Optical and Space Surveillance Technologies Conference*, Maui Economic Development Board, Maui, HI, Sept. 2013, p. E16.
- [5] Kreucher, C. M., Kastella, K. D., and Hero, A. O., III, "Tracking Multiple Targets Using a Particle Filter Representation of the Joint Multitarget Probability Density," *Signal and Data Processing of Small Targets 2003, Proceedings of SPIE*, Vol. 5204, SPIE, San Diego, CA, Jan. 2004, p. 258.
- [6] DeMars, K. J., Hussein, I. I., Frueh, C., Jah, M. K., and Erwin, R. S., "Multiple-Object Space Surveillance Tracking Using Finite-Set Statistics," *Journal of Guidance, Control, and Dynamics*, Vol. 38, No. 9, 2015, pp. 1741–1756. doi:10.2514/1.G000987
- [7] Jia, B., Pham, K., Blasch, E., Shen, D., Wang, Z., and Chen, G., "Cooperative Space Object Tracking Using Consensus-Based Filters," *Proceedings of the 17th International Conference on Information Fusion*, IEEE Publ., Piscataway, NJ, July 2014, pp. 1–8.
- [8] Gehly, S., Jones, B. A., and Axelrad, P., "An AEGIS-CPHD Filter to Maintain Custody of GEO Space Objects with Limited Tracking Data," *Advanced Maui Optical and Space Surveillance Technologies Conference*, Maui Economic Development Board, Maui, HI, Sept. 2014, p. 34.
- [9] Jones, B. A., and Vo, B.-N., "A Labelled Multi-Bernoulli Filter for Space Object Tracking," *2015 AAS/AIAA Spaceflight Mechanics Meeting*, American Astronautical Soc. Paper 15-413, Washington, D.C., Jan. 2015.
- [10] Cheng, Y., DeMars, K. J., Frueh, C., and Jah, M. K., "Gaussian Mixture PHD Filter for Space Object Tracking," *2013 AAS/AIAA Spaceflight Mechanics Meeting*, American Astronautical Soc. Paper 13-242, Washington, D.C., Feb. 2013.
- [11] Cheng, Y., Frueh, C., and DeMars, K. J., "Comparisons of PHD Filter and CPHD Filter for Space Object Tracking," *AAS/AIAA Astrodynamics Specialist Conference*, American Astronautical Soc. Paper 13-770, Washington, D.C., Aug. 2013.
- [12] Goodman, I. R., Mahler, R. P. S., and Nguyen, H.-T., *Mathematics of Data Fusion*, Kluwer, Dordrecht, The Netherlands, 1997.
- [13] Mahler, R. P. S., *Statistical Multisource-Multitarget Information Fusion*, Artech House, Norwood, MA, 2007.
- [14] Mahler, R. P. S., "Multitarget Bayes Filtering via First-Order Multitarget Moments," *IEEE Aerospace and Electronic Systems Magazine*, Vol. 39, No. 4, 2003, pp. 1152–1178. doi:10.1109/TAES.2003.1261119
- [15] Mahler, R. P. S., "PHD Filters of Higher Order in Target Number," *IEEE Aerospace and Electronic Systems Magazine*, Vol. 43, No. 4, 2007, pp. 1523–1543. doi:10.1109/TAES.2007.4441756
- [16] Vo, B.-T., and Vo, B.-N., "Labeled Random Finite Sets and Multi-Object Conjugate Priors," *IEEE Transactions on Signal Processing*, Vol. 61, No. 13, 2013, pp. 3460–3475. doi:10.1109/TSP.2013.2259822

^{††}See Houssineau, J., and Clark, D. E., "Multi-Target Filtering with Linearised Complexity," arXiv:1404.7408v2, 2016.

- [17] Jones, B. A., Bryant, D. S., Vo, B.-T., and Vo, B.-N., "Challenges of Multi-Target Tracking for Space Situational Awareness," *Proceedings of the 18th International Conference on Information Fusion*, IEEE Publ., Piscataway, NJ, July 2015, pp. 1278–1285.
- [18] Vo, B.-N., and Vo, B.-T., "Random Finite Set Multi-Target Trackers: Stochastic Geometry for Space Situational Awareness," *Proceedings of SPIE, Signal Processing, Sensor Fusion, and Target Recognition XXIV*, Vol. 9474, SPIE, Baltimore, MD, May 2015, Paper 94740H.
- [19] Houssineau, J., "Representation and Estimation of Stochastic Populations," Ph.D. Thesis, Heriot-Watt Univ., Edinburgh, U.K., 2015.
- [20] Delande, E. D., Frueh, C., Houssineau, J., and Clark, D. E., "Multi-Object Filtering for Space Situational Awareness," *2015 AAS/AIAA Spaceflight Mechanics Meeting*, American Astronautical Soc. Paper 15-406, Washington, D.C., Jan. 2015.
- [21] Pace, M., and Del Moral, P., "Mean-Field PHD Filters Based on Generalized Feynman-Kac Flow," *IEEE Journal of Selected Topics in Signal Processing*, Vol. 7, No. 3, 2013, pp. 484–495. doi:10.1109/JSTSP.2013.2250909
- [22] Kalman, R. E., "A New Approach to Linear Filtering and Prediction Problems," *Journal of Basic Engineering*, Vol. 82, No. 1, 1960, pp. 35–45. doi:10.1115/1.3662552
- [23] Frueh, C., Kelec, T., and Jah, M. K., "Coupled Orbit-Attitude Dynamics of High Area-to-Mass Ratio (HAMR) Objects: Influence of Solar Radiation Pressure, Earth's Shadow and the Visibility in Light Curves," *Celestial Mechanics and Dynamical Astronomy*, Vol. 117, No. 4, 2013, pp. 385–404. doi:10.1007/s10569-013-9516-5
- [24] Alspach, D. L., and Sorenson, H. W., "Nonlinear Bayesian Estimation Using Gaussian Sum Approximations," *IEEE Transactions on Automatic Control*, Vol. 17, No. 4, 1972, pp. 439–448. doi:10.1109/TAC.1972.1100034
- [25] Doucet, A., de Freitas, N., and Gordon, N., *Sequential Monte Carlo Methods in Practice*, Statistics for Engineering and Information Science, Springer-Verlag, Berlin, 2001.
- [26] Houssineau, J., Ivekovic, S., and Clark, D. E., "Disparity Space: A Parameterisation for Bayesian Triangulation from Multiple Cameras," *Proceedings of the 15th International Conference on Information Fusion*, IEEE Publ., Piscataway, NJ, July 2012, pp. 1734–1740.
- [27] Houssineau, J., Clark, D. E., Ivekovic, S., Lee, C. S., and Franco, J., "A Unified Approach for Multi-Object Triangulation, Tracking and Camera Calibration," *IEEE Transactions on Signal Processing*, Vol. 64, No. 11, 2016, pp. 2934–2948. doi:10.1109/TSP.2016.2523454
- [28] Shepperd, S. W., "Universal Keplerian State Transition Matrix," *Celestial Mechanics*, Vol. 35, No. 2, 1985, pp. 129–144. doi:10.1007/BF01227666
- [29] Goodyear, W. H., "Completely General Closed-Form Solution for Coordinates and Partial Derivative of the Two-Body Problem," *Astronomical Journal*, Vol. 70, No. 3, 1965, pp. 189–192. doi:10.1086/109713
- [30] Doucet, A., Godsill, S. J., and Andrieu, C., "On Sequential Monte Carlo Sampling Methods for Bayesian Filtering," *Statistics and Computing*, Vol. 10, No. 3, 2000, pp. 197–208. doi:10.1023/A:1008935410038
- [31] Franco, J., Delande, E. D., Frueh, C., Houssineau, J., and Clark, D. E., "A Spherical Co-ordinate Space Parameterisation for Orbit Estimation," *IEEE 2016 Aerospace Conference*, IEEE Publ., Piscataway, NJ, March 2016, pp. 1–12.
- [32] Tommei, G., Milani, A., and Rossi, A., "Orbit Determination of Space Debris: Admissible Regions," *Celestial Mechanics and Dynamical Astronomy*, Vol. 97, No. 4, 2007, pp. 289–304. doi:10.1007/s10569-007-9065-x
- [33] Farnocchia, D., Tommei, G., Milani, A., and Rossi, A., "Innovative Methods of Correlation and Orbit Determination for Space Debris," *Celestial Mechanics and Dynamical Astronomy*, Vol. 107, Nos. 1–2, 2010, pp. 169–185. doi:10.1007/s10569-010-9274-6
- [34] DeMars, K. J., and Jah, M. K., "Probabilistic Initial Orbit Determination using Radar Returns," *AAS/AIAA Astrodynamics Specialists Conference*, American Astronautical Soc. Paper 13-704, Washington, D.C., Aug. 2013.
- [35] Rao, C. R., "Information and the Accuracy Attainable in the Estimation of Statistical Parameters," *Bulletin of Calcutta Mathematical Society*, Vol. 37, No. 3, 1945, pp. 81–91.
- [36] Mahalanobis, P. C., "On the Generalized Distance in Statistics," *Proceedings of the National Institute of Sciences of India*, Vol. 2, No. 1, 1936, pp. 49–55.
- [37] Pailhas, Y., Houssineau, J., Delande, E. D., Petillot, Y., and Clark, D. E., "Tracking Underwater Objects Using Large MIMO Sonar Systems," *2nd International Conference and Exhibition on Underwater Acoustics*, European Acoustics Association (EEA), Rhodes, Greece, June 2014, pp. 1033–1040.
- [38] Delande, E. D., Houssineau, J., and Clark, D. E., "Performance Metric in Closed-Loop Sensor Management for Stochastic Populations," *Sensor Signal Processing for Defence (SSPD)*, IEEE Publ., Piscataway, NJ, Sept. 2014, pp. 1–5.
- [39] Delande, E. D., Uney, M., Houssineau, J., and Clark, D. E., "Regional Variance for Multi-Object Filtering," *IEEE Transactions on Signal Processing*, Vol. 62, No. 13, 2014, pp. 3415–3428. doi:10.1109/TSP.2014.2328326
- [40] Houssineau, J., Clark, D. E., and Del Moral, P., "A Sequential Monte Carlo Approximation of the HISP Filter," *2015 23rd European Signal Processing Conference (EUSIPCO)*, IEEE Publ., Piscataway, NJ, Aug. 2015, pp. 1251–1255.
- [41] Delande, E. D., Houssineau, J., Franco, J., Frueh, C., and Clark, D. E., "A New Multi-Target Tracking Algorithm for a Large Number of Orbiting Objects," *2017 AAS/AIAA Spaceflight Mechanics Meeting*, American Astronautical Soc., Washington, D.C., Feb. 2017.