

MULTI-OBJECT FILTERING FOR SPACE SITUATIONAL AWARENESS

E. D. Delande^{*}; C. Frueh[†]; J. Houssineau[‡]; D. E. Clark[§]

The sources of information for Earth orbiting space objects are passive or active observations by ground based or space based sensors. The majority of tracking and detection of new objects for contemporary applications is performed by ground based sensors, which are mainly optical or radar based. Challenges lie in the long times spans between very short observation arcs, large number of objects, brightness variations during observations, occlusions, crossing targets and clutter. Traditionally, tracking problems for space situational awareness have been approached through heuristics-based techniques such as Multiple Hypothesis Tracking (MHT). More recently, solutions derived from Finite Set Statistics (FISST) have been used, such as the Probability Hypothesis Density (PHD) filter, that describe the targets at the population rather than individual level. The recent mathematical framework for the estimation of stochastic populations combines the advantages of previous approaches by propagating specific information on targets (i.e. tracks) whenever appropriate, and by avoiding heuristics through its fully probabilistic nature. This paper presents the first application to space situational awareness problems of the novel filter for Distinguishable and Independent Stochastic Populations (DISP), a tracking algorithm derived from this framework, through a multi-object surveillance scenario involving a ground based Doppler radar and five crossing objects in all orbital regions. The preliminary results show that, despite the sensor's limited observability and constrained field of view, the DISP filter is able to detect the orbiting objects entering the field of view and maintain individual information on them, with associated uncertainty, even once they have left the field of view.

INTRODUCTION

Due to the increasing number and diversity of objects orbiting around Earth, maintaining an up-to-date catalogue of the objects evolving on every orbit has become a topic of growing interest, but a problem of growing complexity as well. The orbits may be observed by a range of different sensor systems, either ground-based or space-based, with specific capabilities and characteristics. Producing a catalogue of the orbiting objects based on the measurements collected by these sensors is a challenging and exciting field of active research, for they have limited abilities, such as:

- Their field of view is constrained (i.e. only a fraction of space can be covered at any time);
- The data capture process is imprecise (i.e. the collected measurements are noisy);
- Spurious measurements and missed detections occur;
- The state of the objects is only partially observable (e.g. the full position and velocity components cannot be inferred from a single measurement).

^{*}Research Associate, Heriot Watt University, School of Engineering & Physical Sciences, Edinburgh, EH14 4AS, United Kingdom, E.D.Delande@hw.ac.uk

[†]Assistant Professor, School of Aeronautics and Astronautics, Purdue University, 701 W. Stadium Ave., West Lafayette, IN 47907-2045, USA, AIAA Member, cfrueh@purdue.edu/carolin.frueh@gmail.com

[‡]PhD student, Heriot Watt University, School of Engineering & Physical Sciences, Edinburgh, EH14 4AS, United Kingdom, jh207@hw.ac.uk

[§]Associate Professor, Heriot Watt University, School of Engineering & Physical Sciences, Edinburgh, EH14 4AS, United Kingdom, D.E.Clark@hw.ac.uk

On top of these issues related to every single sensor, the fusion of the information collected by several sensors – especially if they provide measurements of different nature (e.g. optical: astrometric, infrared, spectral or filter measurements, radar: astrometric, astrometric Doppler, wide or narrow beam, characterization measurements, laser ranging etc.) – in a single coherent representation of the orbiting objects, while out of the scope of this paper, is a challenging problem in itself.

The construction of an up-to-date catalogue of the orbiting objects, fed by the measurements collected sequentially from the sensor system, is referred to as a *multi-object detection and tracking problem*. It can be decomposed as follows:

- How many objects are there? (*cardinality* problem);
- Where/what is each object? (*state estimation* problem);
- Which measurement goes with which object? (*data association* problem).

Several approaches can be adopted to solve multi-object detection and tracking problems. The solution presented in this paper is set within the broad class of Bayesian estimation problems, whose objective is to provide the operator with a *probabilistic* representation of the multiple objects, accounting for the *uncertainties* in the information collected from the sensors (due to inaccuracies, missed detections, etc.), in the operator’s knowledge about the objects’ behaviour (including, most notably, the orbital motion model).

The objective of this paper is to illustrate a novel multi-object estimation solution, the DISP filter,¹ on a challenging scenario in the context of space situational awareness.

The first section presents the Bayesian estimation framework in a general way, and provides a brief description of the well-established track-based^{2,3,4,5} and Random Finite Set (RFS)-based^{6,7,8} approaches. It then introduces an alternative approach, the novel estimation framework for stochastic populations, from which the DISP filter is derived. The next section introduces the DISP filter, and provides the outlines of its construction. The following section presents the multi-object surveillance scenario, and details the implementation of the DISP in the context of of this scenario. Preliminary results are presented and commented in the next section. The last section concludes and discusses further developments of the DISP filter.

MULTI-OBJECT BAYESIAN ESTIMATION

Surveillance problem

The focus of the surveillance problem is a population \mathcal{X} of individuals (e.g. satellites, space debris), expected to evolve in some bounded region of the physical space called the *surveillance scene*. The surveillance scene has to be discriminated from the field of view (FOV) of the sensor that is used to perform the actual survey. The surveillance scene is the whole space of interest, in our case, the complete Near Earth space around the Earth, from the Earth surface well beyond the geostationary ring in all directions. An individual of the surveyed population enters or comes to existence in the scene* at some unknown time t_{\bullet} , and may leave it (or disappear, object death, by decay or when docking on another space craft and becoming effectively one object) at some unknown time t_{\circ} later on during the scenario. While an individual evolves in the scene, it possesses some physical and measurable characteristics described via a *state* \mathbf{x} , belonging to some target state space \mathbf{X} that denotes physical and measurable characteristics of the individual which are unknown, but of interest to the operator (position, velocity, etc.). In order to allow for the description of individuals that are located *outside* of the scene (e.g. prior to coming into existence or being born through launch or collision), the state space is extended with the empty state ψ to form the augmented state space $\bar{\mathbf{X}} = \{\psi\} \cup \mathbf{X}$,

The aim of the surveillance problem is then to determine, at any time t relevant to the surveillance activity, the state of each individual of the population \mathcal{X} in the augmented state space $\bar{\mathbf{X}}$, i.e. *a*) whether the individual is in the scene and, *b*) if so, the value of its characteristics of interest.

*“Entering the scene” should be interpreted in a broad sense in this context. It is not restricted to individuals crossing the physical boundaries of the scene and may also cover, for example, individuals appearing spontaneously anywhere in the scene following a collision in the sense of object birth.

Bayesian estimation: general principle

In the target tracking community, the Bayesian framework is largely known and exploited for sequential estimation problems, in which some system of interest – usually, one or several moving individuals of strategic interest to the operator – is observed through some sensor that produces *sequences* of measurements across time. The aim of Bayesian *filters**, such as Kalman filter variants^{9,10,11} and particle filters,¹² is then to provide the operator with a probabilistic representation of the system of interest, dynamically updated with the measurements collected from the sensor so far (see Figure 1).

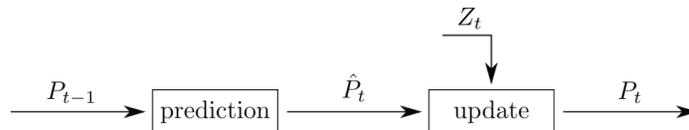


Figure 1: Data flow of a Bayesian estimator (time step t).

The Bayesian filters share a common structure depicted by Figure 1. At some time $t > 0$ relevant to the scenario, where a collection of measurements Z_t produced by the sensor system becomes available to the operator:

1. Some *prior* information P_{t-1} is propagated by the filter from past times;
2. During the *time update step*, the *predicted* information \hat{P}_t is produced by the filter, based on the knowledge of the operator regarding dynamic behaviour of the individuals (usually through a motion model);
3. During the *data update step*, the *posterior* information P_t is produced by the filter, as a correction of the predicted information through the exploitation of the newly available collection of measurements Z_t .

Numerous works in target detection/tracking problems, spanning general Bayesian estimation framework illustrated in Figure 1, have been proposed; essentially, the choice of a filter aims at shaping the nature of the propagated information P_t , tailored to the objective of the surveillance problem – i.e. what is of interest to the operator? – and the constraints imposed by the available resources – i.e. what can the operator afford to achieve their objective?.

Modern developments in Bayesian estimation for target detection/tracking focussed first on the single-object case, for which the cardinality problem is avoided – the operator *knows* that there is one and only one individual in the surveillance scene – and the data association problem is significantly simpler – unless it is spurious, a measurement necessarily originates from the only individual in the scene. The bulk of the developments in single-object Bayesian estimation thus focused on the single-object prediction and update steps. Well-established solutions typically include Gaussian-based approaches such as the Kalman filter,⁹ and later on its extended¹⁰ and unscented¹¹ extensions, and the broad class of particle-based approaches.¹² Note that single-object Bayesian estimation techniques form a constitutive element of multi-object Bayesian filters that follows a track-based approach, and will play an important role in the implementation of the DISP filter. The interest of the tracking community has also broadened to multi-object scenarios. Broadly speaking, the well-established solutions follow either a track-based or a RFS-based approach, described in the following paragraphs.

Track-based filtering

The earlier developments in multi-object filtering followed a track-based approach, in which specific information on each potential individual of the population \mathcal{X} – or *target* – is maintained through a *track*. Each track

*In the context of Bayesian estimation, the term “filter” refers to an estimator that maintains a probabilistic representation of some observed system’s state.

describes the stream of measurements presumed to have been produced by the corresponding target across time, and maintains a description of the target's current state (position, velocity, etc.) through a single-object probability distribution on the target state space.

The inherent advantage of track-based approaches is that the multi-object probabilistic representation maintained by the filter essentially breaks down to a collection of single-object probability representations accounting for each potential individual. Its intuitive construction naturally allows the direct exploitation of single-object techniques for the time and data update of each target (see Figure 1). In addition, a track history – i.e. a sequence of the estimated states throughout the past times – is available for each target, from which trajectories can be readily displayed.

Track-based approaches are relatively challenging to implement into practical algorithms due to the complexity of the data association step. Pioneering works include the construction of the well-established MHT filter, in which every prior track is matched with every current measurement in order to produce a corresponding updated track; in other words, every possible sequence of measurements is maintained as possible evidence of an underlying individual, with a corresponding probability distribution.^{2,3,4} Provided that the number of objects to be tracked is known by the operator, the Joint Probabilistic Density Association (JPDA) filter provides a lighter but less accurate approach, in which the likelihood of each measurement/track pair is estimated and weighted accordingly; the probability distribution of each predicted track is then updated as a weighted average of all the relevant measurement/track single-object data updates.⁵ Track-based approaches typically include mechanisms relying on heuristics (e.g. track creation/deletion, track pruning, etc.) that must be tuned accordingly to the specifics of the surveillance problem (sensor characteristics, target behaviour, etc.).

RFS-based filtering

The exploitation of early results in point process theory^{13,14} for multi-object detection/tracking problems started with Mahler's seminal papers on the PHD⁶ and the Cardinalized Probability Hypothesis Density (CPHD)⁷ filters. RFS-based filters focused on the *population* of individuals as a whole; most notably, the information on *all* the target states is represented by the filter as a *single* random object, namely a RFS, whose size (number of targets) and elements (target states) are both random. Recent works have exploited this framework in the context of space situational awareness.^{15,16}

The RFS-based filters all proceed from the rigorous application of derivation rules from the FISST framework;⁸ they are fully probabilistic in nature, providing two distinct advantages. First, the derivation of the filters being principled, the amount of fine tuning required for the implementation of practical algorithms is limited, and their performance relies primarily on the modelling choices. Second, the principled production of meaningful statistical tools from the filters' outputs is readily available, for example to enrich the information transmitted to the operator or to compare the performance of RFS-based filters. For example, the estimated number of individuals, with uncertainty, that currently live in an arbitrary region of the surveillance scene can be extracted in a principled way from the RFS-based filters.^{17,18}

The elements of a set being unordered the RFS representation does not allow for a natural identification of a *specific* individual within the population of targets. A positive consequence is that the costly, but usually necessary, data association step in the construction of track-based filter is usually spared in the derivation of RFS-based filters; in other words, the whole population of targets is updated *once* with the set of all the measurements Z_t *taken as a whole*, without explicitly considering all the possible measurement-target pairs. On the other hand, this representation precludes the natural propagation of individual information on a specific target, i.e. the production of *tracks*. Recent works^{19,20} propose the introduction of labelled RFSs in order to address this shortcoming within the FISST framework.

Estimation framework for stochastic populations

The estimation framework for stochastic populations,²¹ from which the DISP filter is derived,¹ proposes an alternative approach to the track-based and RFS-based solutions. Broadly speaking, its development was shaped by the following objectives:

- It must be fully probabilistic in nature;
- It must provide a suitable representation for individual information about a specific target (i.e. a track), *whenever appropriate*.

In its most general form, the resulting estimation framework is designed to provide tailored solutions to more complex estimation problems whose scope goes beyond multi-object joint detection/tracking; recent exploitations include information-based sensor management for surveillance,²² and target classification, detection and tracking for harbour surveillance.²³ In particular, this estimation framework provides grounds for a natural extension of the DISP filter in order to approach other estimation problems encountered in the context of space situational awareness, such as object classification.

The second requirement above was motivated by the shortcomings of the RFS-based representation mentioned earlier on in the paper, namely the lack of a natural distinction between the targets composing the estimated population. This novel estimation framework²¹ proposes two exclusive probabilistic representations for a target, depending on its nature:

- An *indistinguishable* target belongs to a population whose members have not been identified (yet) by the operator through individual information;
- A *distinguishable* target has been identified by the operator through individual information.

In the context of a space surveillance scenario, a typical example of an indistinguishable target would be a piece of debris, within a larger cloud, following a recent collision of known objects and before the outcome of the collision has been observed: propagated information about the objects prior to their collision, and some knowledge about orbital mechanics may allow a probabilistic description of the cloud of debris *as a whole*, but not of any specific piece of debris. On the other hand, provided that each measurement produced by the sensor system stems from *at most* one object, any detected target, i.e. associated to at least one measurement collected by the operator, becomes distinguishable, and can be appropriately described by a *track*.

THE DISP FILTER

This section provides a general description of the filter and presents the necessary concepts before an illustration in the context of a space surveillance problem. The construction of the DISP filter, exploiting notations and concepts from the measure theory,^{24,25} is discussed in more details in recent works.^{22,1}

Outline and data flow

The DISP filter aims at providing an optimal solution, in the Bayesian sense, to the multi-object surveillance problem under a few assumptions.¹ The main ones are as follows:

1. The individuals behave independently from one another;
2. At each time step, each individual in the scene produces *at most* one measurement (if none, it has a *missed-detection*);
3. At each time step, each measurement stems from *at most* one individual in the scene (if none, it is a *false alarm*);
4. Individuals are immediately detected by the sensor upon entering the surveillance scene.

The second and third assumptions* above imply that a stream of measurements across time characterises a *single* target without ambiguity, and that a target becomes *distinguishable* upon its first association to a

*These assumptions are standard in multi-target detection and tracking problems. They are not, however, necessary within the estimation framework for stochastic populations. Relaxing any one of these two assumptions would lead to the derivation of more involved filtering solutions, out of scope in this study.

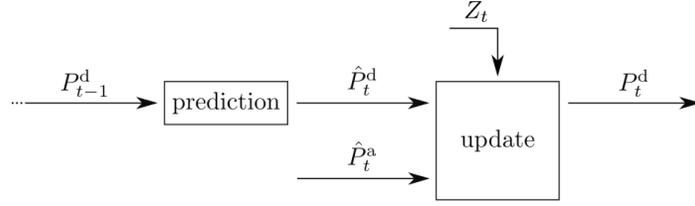


Figure 2: Data flow of the DISP filter (time step t). The superscript d denotes a population of previously detected, hence distinguishable, targets. The superscript a denotes a population of appearing targets, still indistinguishable from one another.

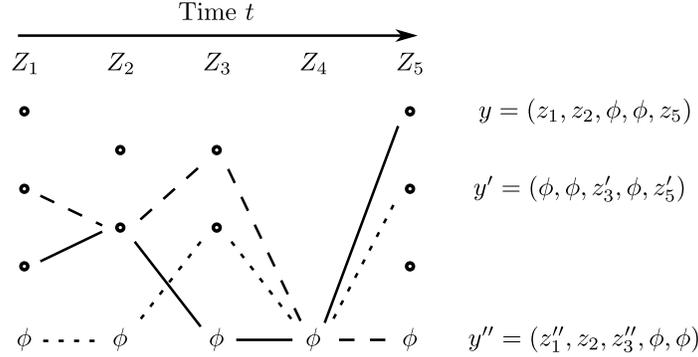


Figure 3: An example of a few possible observation paths at time $t = 6$, before the availability of the measurement set Z_6 . The target characterised by y' entered the scene at time 3 and was detected through the measurement z'_3 , was miss-detected at time 4 and produced the measurement z'_5 at time 5.

collected measurement (see Figure 3). The fourth assumption implies that, for estimation purposes, the individuals are assumed to have appeared in the surveillance scene at the time of their first detection. An immediate consequence is that the DISP filter is unable to represent the population of individuals that have *already* entered the surveillance scene but have *not* (yet) entered the sensor's field of view.

Since *a*) individuals are not represented by targets prior to their their detection, and *b*) targets becomes distinguishable as soon as they are associated to a measurement, it follows that the DISP filter does not propagate information on population of indistinguishable targets (see data flow in Figure 2).

Input (time t)

From the discussion above and the data flow in Figure 2, it follows that the DISP filter propagates information from past times through some probability P_{t-1}^d , on a population of *distinguishable* targets, every one of them identified by its own *track*.

Tracks

A target is identified through a sequence of observations presumed to have been produced by the underlying individual across past times. We denote by $\bar{Z}_t = \{\phi\} \cup Z_t$ the measurement set at time t augmented with the empty observation ϕ , in order to account for miss-detections. Since individuals are assumed detected at their time of appearance in the scene t_\bullet , a target can be characterised before the update at time t by the *observation path*²⁶

$$y = (\phi, \dots, \phi, z_{t_\bullet}, z_{t_\bullet+1}, \dots, z_{t-1}) \quad (1)$$

where $z_{t_\bullet} \in Z_{t_\bullet}$ and $z_{t'} \in \bar{Z}_{t'}$ for any $t_\bullet < t' \leq t-1$. We implicitly assume that individuals with state ψ are never detected by the sensor, that is, individuals can only be detected while they lie in the scene. We denote by Y_{t-1} the set of all possible observation paths at time t , before the availability of the current measurement set Z_t (see Figure 3).

Any target represented in the input population is therefore identified by its own observation path y , and the information regarding the current state of the (potential) individual it represents is described by a probability distribution p_{t-1}^y on the augmented state space $\bar{\mathbf{X}}$. From the probability distribution p_{t-1}^y we can extract the following information regarding the corresponding target:

- The scalar $p_{t-1}^y(\mathbf{X}) = \int 1_{\mathbf{X}}(x)p_{t-1}^y(dx)$ is the probability that the target is still in the scene, also called the *probability of presence* of the target*;
- Its restriction to the state space describes the state of the target *provided that the underlying individual is still in the scene*, also called the *spatial distribution* of the target.

A *track* is then defined by a pair with an observation path $y \in Y_{t-1}$ and the corresponding probability distribution p_{t-1}^y . Because a target is identified through the corresponding observation path y and completely characterised by the corresponding track (y, p_{t-1}^y) , the terms *target*, *observation path*, and *track* may be used interchangeably when there is no ambiguity.

Hypotheses

Recall from the previous sections that any measurement stems from *at most* one individual. Therefore, the possible configuration of individuals that have entered the scene since the beginning of the scenario are described by all the subsets of observation paths $h \subseteq Y_t$, called *hypotheses*, such that any pair of distinct observation paths belonging to a common hypothesis $y, y' \in h$ do not share a common measurement, unless it is the empty one. In the situation illustrated in Figure 3, for example, the subsets $\emptyset, \{y\}, \{y'\}, \{y''\}, \{y, y'\}, \{y', y''\}$ are hypotheses but the subsets $\{y, y''\}, \{y, y', y''\}$ are not, for the observation paths y and y'' share the measurement z_2 and thus *at most* one of them can represent a valid individual of the surveyed population \mathcal{X} . The set of all hypotheses propagated from the previous time step is denoted by H_{t-1} , and the DISP filter maintains a cardinality distribution c_{t-1} on the hypotheses such that

$$\sum_{h \in H_{t-1}} c_{t-1}(h) = 1. \quad (2)$$

The weight $c_{t-1}(h)$, where $h \in H_{t-1}$, denotes the *credibility* of the hypothesis h , i.e. the probability that the individuals of the population \mathcal{X} , that have entered the surveillance scene so far, have produced the sequences of measurements denoted by the observation paths $y \in h$. It is also called the *probability of existence* of the hypothesis h .

Note that a given observation path $y \in Y_{t-1}$ may belong to several hypotheses $h \in H_{t-1}$. For every observation path $y \in Y_{t-1}$ the scalar

$$\alpha_{t-1}^y = \sum_{\substack{h \in H_{t-1} \\ h \ni y}} c_{t-1}(h) \quad (3)$$

denotes the *credibility* or *probability of existence* of target y . It is a scalar between 0 and 1; the closer α_{t-1}^y is to 1, the more likely it is that some individual of the population \mathcal{X} has produced the sequence of measurements given by the observation path $y \in Y_{t-1}$. Note the difference with the probability of presence $p_{t-1}^y(\mathbf{X})$, which denotes the probability that the individual represented by the target is still in the scene, *provided that it exists*.

The prior information, propagated by the DISP filter from the previous time step (see Figure 2) can then be formally described by the product measure¹

$$P_{t-1}^d = \sum_{h \in H_{t-1}} c_{t-1}(h) \bigotimes_{y \in h} p_{t-1}^y. \quad (4)$$

*Its complementary to one, i.e. $p_{t-1}^y(\psi)$, is the *probability of absence* of the target.

Prediction step (time t)

The prediction step of the DISP filter is quite straightforward. Since the measurement set Z_t has not been collected yet, the possible observation paths Y_{t-1} and the corresponding hypotheses H_{t-1} remain unchanged. The probability distribution p_{t-1}^y of each track $y \in Y_{t-1}$ is transformed to its predicted form \hat{p}_t^y as follows:

$$\begin{cases} \hat{p}_t^y(dx') = \int_{\mathbf{X}} p_{s,t-1}(x) \hat{m}_{t-1,t}(x, dx') p_{t-1}^y(dx), & x' \in \mathbf{X}, \\ \hat{p}_t^y(\psi) = \int_{\mathbf{X}} [1 - p_{s,t-1}(x)] p_{t-1}^y(dx) + p_{t-1}^y(\psi), \end{cases} \quad (5)$$

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

Note that the predicted probability of presence $\hat{p}_t^y(\mathbf{X}) = 1 - \hat{p}_t^y(\psi)$ of a given track $y \in Y_{t-1}$ is smaller or equal to its prior value $p_{t-1}^y(\mathbf{X})$. Indeed, only a new detection may increase the evidence of an individual's presence in the scene, and new measurements are unavailable at this point.

The predicted information (see Figure 2) can then be formally described by the product measure¹

$$\hat{P}_t^d = \sum_{h \in H_{t-1}} c_{t-1}(h) \bigotimes_{y \in h} \hat{p}_t^y. \quad (6)$$

Update step (time t)

In the update step, the measurement set Z_t collected from the sensor is processed in order to update the information \hat{P}_t^d maintained so far by the DISP filter (see Figure 2).

Appearing individuals

The individuals appearing in the surveillance are described by a population of *indistinguishable* targets, for none has been associated to a measurement yet and thus no specific information is available on any of them. Based on the knowledge of the operator regarding the frequency of appearance of individuals within the surveillance scene, and their whereabouts, this population is described by:

- Some cardinality function \hat{c}_t^a , describing the number of individuals that have appeared in the scene since the previous time step;
- Some probability distribution \hat{p}_t^a on the state space \mathbf{X} , *collectively* describing the initial state of *all* the appearing individuals.

Recall from the previous sections that the appearing individuals are immediately detected upon arrival in the scene and thus the mass of the probability distribution \hat{p}_t^a is concentrated within the sensor's field of view. The information on appearing targets (see Figure 2) can then be formally described by the product measure¹

$$\hat{P}_t^a = \sum_{n \in \mathbb{N}} \hat{c}_t^a(n) (\hat{p}_t^a)^{\otimes n}. \quad (7)$$

Data association

The core of the update step consists in the *data association step*, in which potential sources of measurements are matched with the newly available measurement set Z_t and the likelihood of each association is assessed. The potential sources of measurements are:

1. The individuals that entered the scene in an previous step, represented by \hat{P}_t^d ;

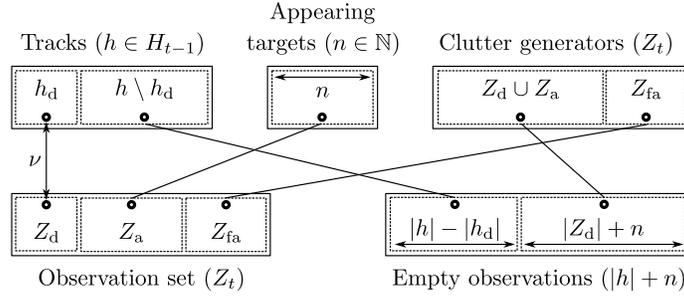


Figure 4: Data association, for a hypothesis $h \in H_{t-1}$ and given number n of appearing individuals.

2. The individuals that have just entered the scene, represented by \hat{P}_t^a ;
3. The false alarm generators.

For the purpose of data association, a false alarm generator is modelled for each measurement $z \in Z_t$ 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 individual of the population \mathcal{X}). Note that the modelling of the false alarm function is highly dependent on the type of sensor.

For every hypothesis $h \in H_{t-1}$ and every possible number of appearing individuals $n \in \mathbb{N}$, a valid data association \mathbf{h} then associates the three sources of measurements – the distinguishable targets $y \in h$, the n appearing targets, and the $|Z_t|$ clutter generators – to the collected measurements $z \in Z_t$ and the empty observation ϕ (see Figure 4). A valid association \mathbf{h} is then an element of the set¹

$$\text{Adm}_{Z_t}(h, n) = \left\{ (h_d, Z_d, Z_a, \nu) \mid h_d \subseteq h, Z_d \subseteq Z_t, Z_a \subseteq Z_t \setminus Z_d, |Z_a| = n, \nu \in \mathbf{S}(h_d, Z_d) \right\}, \quad (8)$$

where h_d designates the targets that are detected at the present time step, Z_d the measurements associated to these detected targets, Z_a the measurements associated to the n appearing targets, and ν the bijective function associating detected targets to measurements in Z_d .

Each triplet $\mathbf{a} = (h, n, \mathbf{h})$, where $\mathbf{h} \in \text{Adm}_{Z_t}(h, n)$, then represents a specific *association scheme* described by the scalar¹

$$P_t^{\mathbf{a}} = P_{d,t}^{\mathbf{a}} \times P_{md,t}^{\mathbf{a}} \times P_{fa,t}^{\mathbf{a}}, \quad (9)$$

where

$$\begin{aligned} P_{d,t}^{\mathbf{a}} &= \left[\prod_{z \in Z_a} \int_{\mathbf{X}} \ell_t(z, x) p_{d,t}(x) \hat{p}_t^{\mathbf{a}}(dx) \right] \left[\prod_{y \in h_d} \int_{\mathbf{X}} \ell_t(\nu(y), x) p_{d,t}(x) \hat{p}_t^y(dx) \right], \\ P_{md,t}^{\mathbf{a}} &= \prod_{y \in h \setminus h_d} \left[\hat{p}_t^y(\psi) + \int_{\mathbf{X}} [1 - p_{d,t}(x)] \hat{p}_t^y(dx) \right], \\ P_{fa,t}^{\mathbf{a}} &= \left[\prod_{z \in Z_d \cup Z_a} [1 - p_{fa,t}(z)] \right] \left[\prod_{z \in Z_t \setminus (Z_d \cup Z_a)} p_{fa,t}(z) \right], \end{aligned} \quad (10)$$

where $P_{d,t}^{\mathbf{a}}$, $P_{md,t}^{\mathbf{a}}$ and $P_{fa,t}^{\mathbf{a}}$ denote the joint probability of the data association of the detected targets, the miss-detected targets, and the clutter generators, respectively, $p_{d,t}$ denotes the probability of detection, and ℓ_t the measurement likelihood function. As for the probability of false alarm function, the modelling of the probability of detection and likelihood functions are heavily dependent on the type of sensor used for a specific surveillance problem.

Hypothesis update

For a given population $h \in H_{t-1}$ of predicted targets, a given number $n \in \mathbb{N}$ of appearing individuals, some

valid association $\mathbf{h} \in \text{Adm}_{Z_t}(h, n)$ leads to the construction of a *unique* updated hypothesis \hat{h} , composed of the updated observation paths of the form

$$\hat{h} = \left[\bigcup_{y \in h_d} \{y:\nu(y)\} \right] \cup \left[\bigcup_{y \in h \setminus h_d} \{y:\phi\} \right] \cup \left[\bigcup_{z \in Z_a} \{\phi_{t-1}:z\} \right], \quad (11)$$

where ϕ_{t-1} denotes the “empty” observation path, i.e. the sequence of $t - 1$ empty measurements*. Using Bayes’ rule, its probability of existence $c_t(\hat{h})$ is found to be[†]

$$c_t(\hat{h}) = \frac{\hat{c}_t^a(n)c_{t-1}(h)P_t^a}{\sum_{n' \geq 0} \sum_{h' \in H_{t-1}} \hat{c}_t^a(n')c_{t-1}(h') \sum_{\mathbf{h}' \in \text{Adm}_{Z_t}(h', n')} P_t^a}, \quad (12)$$

Note that the number of updated hypotheses $|H_t|$ is finite, for the number of predicted hypotheses $|H_{t-1}|$ is finite, the number of observations $|Z_t|$ is assumed finite, and the set of admissible associations $\text{Adm}_{Z_t}(h, n)$ given by Eq. (8) is empty if $n > |Z_t|$.

Track update

Following the construction of an updated hypothesis $\hat{h} \in H_t$ given by Eq. (11), the updated observation paths it contains 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 which was detected and produced measurement z . Following Bayes’ rule, the probability distribution of the track is updated as follows:¹

$$\begin{cases} p_t^{y:z}(dx') = \frac{p_{d,t}(x')\ell_t(z, x')\hat{p}_t^y(dx')}{\int_{\mathbf{X}} p_{d,t}(x)\ell_t(z, x)\hat{p}_t^y(dx)}, & x' \in \mathbf{X}, \\ p_t^{y:z}(\psi) = 0. \end{cases} \quad (13)$$

Note that the probability of presence in the scene raises to one[†]. Indeed, only individuals in the surveillance scene can be detected by the sensor; therefore, a target that has just been detected is necessarily 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 miss-detected this time step. Following Bayes’ rule, the probability distribution of the track is updated as follows:¹

$$\begin{cases} p_t^{y:\phi}(dx') = \frac{[1 - p_{d,t}(x')]\hat{p}_t^y(dx')}{\hat{p}_t^y(\psi) + \int_{\mathbf{X}} [1 - p_{d,t}(x)]\hat{p}_t^y(dx)}, & x' \in \mathbf{X}, \\ p_t^{y:\phi}(\psi) = \frac{\hat{p}_t^y(\psi)}{\hat{p}_t^y(\psi) + \int_{\mathbf{X}} [1 - p_{d,t}(x)]\hat{p}_t^y(dx)}. \end{cases} \quad (14)$$

Note that the denominator in Eq. 14 is no greater than one; thus $p_t^{y:\phi}(\psi)$ is no smaller than $\hat{p}_t^y(\psi)$; in other words, the probability of presence of the target is non-increasing. Indeed, no fresh evidence on the target presence is available since it was miss-detected this time step.

Finally, a new observation path of the form $\phi_{t-1}:z$, where $z \in Z_t$, denotes an appearing target which was detected and produced measurement z . The update process is similar to the previously detected targets in Eq. (13) and yields:¹

$$\begin{cases} p_t^{\phi_{t-1}:z}(dx') = \frac{p_{d,t}(x')\ell_t(z, x')\hat{p}_t^a(dx')}{\int_{\mathbf{X}} p_{d,t}(x)\ell_t(z, x)\hat{p}_t^a(dx)}, & x' \in \mathbf{X}, \\ p_t^{\phi_{t-1}:z}(\psi) = 0. \end{cases} \quad (15)$$

Note that, for the same reason as detected targets in Eq. (13), the probability of presence of appearing targets is always one.

*“:” is the concatenation operator, i.e. $(e_1, \dots, e_n):e = (e_1, \dots, e_n, e)$.

[†]Recall that $p_t^{y:z}(\mathbf{X}) = 1 - p_t^{y:z}(\psi)$.

Posterior information

The posterior information (see Figure 2) can then be formally described by the product measure¹

$$P_t^d = \sum_{h \in H_t} c_t(h) \bigotimes_{y \in h} p_t^y. \quad (16)$$

Maximum a Posteriori (MAP) of the multi-object configuration

The DISP filter being fully probabilistic in nature, the posterior information P_t^d representing all the individuals of the population \mathcal{X} that have entered the surveillance scene can be exploited in various ways. We shall focus here on the production of MAP estimate, i.e. the most probable configuration of the population of these individuals.

Since the maintained hypotheses $h \in H_t$ denote exclusive multi-object configurations representing the population of individuals, the most probable configuration must be chosen as the hypothesis h_t^* with the highest credibility, i.e.

$$h_t^* = \arg \max_{h \in H_t} c_t(h). \quad (17)$$

Then, the operator is not necessarily interested in being transmitted *all* of the tracks $y \in h_t^*$. Some may be tentative tracks with a low probability of existence that have not been confirmed yet through a consistent stream of measurements, and a cautious operator may not be interested in displaying such targets. Consider some track of the form $y:z \in h_t^*$. In order to prevent flickering in the display of targets, a simple rule, based on the probability of existence computed in Eq. (3), can be designed as follows:

1. If the parent track y was not displayed, do not display $y:z$ unless its probability of existence $\alpha_t^{y:z}$ exceeds some confirmation threshold t_c ;
2. If the parent track y was displayed, keep displaying $y:z$ if its probability of existence $\alpha_t^{y:z}$ exceeds some *lower* de-confirmation threshold $t_d < t_c$.

Setting the values of the thresholds t_u , t_c should be left to the operator based on their expertise and objectives for a specific surveillance problem; they need not and should not be tuned according to the specifics of the implementation (notably the chosen approach for the implementation of single-object probability distributions), the nature of the sensor, or any other parameter of the filtering process that is not the primary concern of the operator.

IMPLEMENTATION OF THE MULTI-OBJECT SURVEILLANCE PROBLEM

Surveillance problem

Individuals

The DISP filter is applied to the problem of Earth orbiting space object tracking, where five targets evolve on different orbits. The surveillance scene is a portion of space centred on Earth, spanning 70 000 km in every direction, and the targets are characterised by their six-variable state, consisting of their position (x, y, z) and velocity $(\dot{x}, \dot{y}, \dot{z})$ elements. The initial orbital elements and area to mass ratios of the five targets are listed in Table 1. The start epoch is 53159.5 MJD, the propagation duration is 10 hours, and the interval between two time steps is constant and set to 100 s. The orbits have been created from the initial orbital elements utilizing a Runge-Kutta 7/8 numerical integration. The Earth gravitational field has been taken into account up to order and degree 12, third body perturbations of Sun and Moon as well as direct radiation pressure, a spherical object shape was assumed.

This scenario presents several challenges (see Figure 5). Individual 1 is the “test target”; it is on a geostationary orbit and remains in the sensor’s field of view throughout the whole duration of the scenario. Individual 2 is also on a geostationary orbit, on a retrograde orbit. Individuals 1 and 2 cross each other’s paths when they are covered by the sensor. Since the measurements do not provide angular rates, the measurements produced by these two targets are expected to be close in value. Individual 3 and 4 are on medium

Target no.	a (km)	e	i ($^\circ$)	Ω ($^\circ$)	ω ($^\circ$)	ν ($^\circ$)	AMR (m^2/kg)
1	42164.0	0.012	10.0	61.0	349.0	82.0	0.02
2	42164.0	0.010	170.0	50.0	30.0	20.0	0.60
3	26595.0	0.010	2.0	20.0	311.0	80.0	0.02
4	27495.0	0.30	3.0	30.0	250.0	50.0 <td 0.02	
5	42164.0	0.80	11.3	60.0	351.0	80.0	0.02

Table 1: Semi-major axis, eccentricity, inclination, right ascension of the ascending node, argument of perigee, true anomaly and area-to-mass ratio.

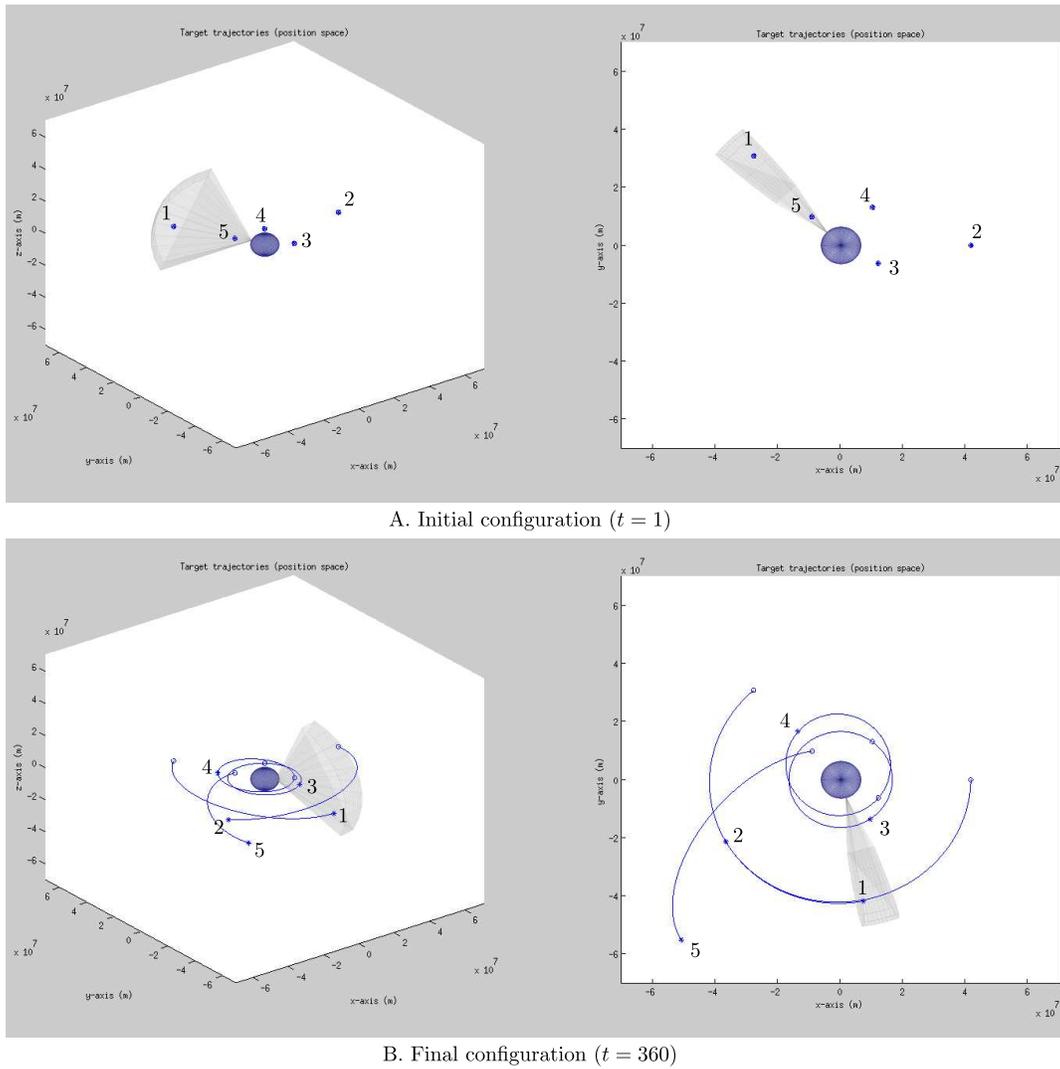


Figure 5: Illustration of the scenario. Initial positions of the five individuals are marked by circles, an last positions by stars. The grey region denotes the sensor's field of view.

Table 2: Sensor model: cell resolution and noise profile

cell resolution				noise (standard deviation)			
range r	azimuth θ	elev. φ	range rate \dot{r}	range r	azimuth θ	elev. φ	range rate \dot{r}
1000 m	1°	1°	100 m s ⁻¹	1000 m	1°	1°	100 m s ⁻¹

Earth orbits, and do not stay in the sensor’s field of view for long. Individual 4, however, re-enters the sensor’s field of view once during the scenario. Individual 5 is on a geostationary high eccentricity transition orbit, about to leave the vicinity of Earth. It is initially in the sensor’s field of view, but leaves it after four time steps and never re-enters it.

Sensor

A Doppler radar is placed on the northern Hemisphere, at a latitude of 20°; the sensor’s field of view thus rotates in the reference frame. The sensor’s field of view defines a radar-like sensing region of about 45×10^6 m deep, 15° wide, in Earth fixed coordinates; it is fixed pointed at the zenith and covers the inclinations from -45° to 45° . The sensor provides range and range rate, as well as azimuth and elevation measurements every 100 seconds when the object is in the field of view. The information on the angular rates is not observed, hence only a subset of the full state is available through a single measurement. At most one false alarm occurs per sensor resolution cell and per scan, occurring with equal probability in each cell. An average of 2 false alarms is produced per scan. The probability of detection of the sensor is set to 0.9 throughout the scenario. The characteristics of the observation noise are given in Table 2.

Prediction step in the filter and orbital mechanics

At this point, it is important to note that the construction of the DISP filter does not impose any restriction on the single-object probability distributions. An important consequence is that the implementation choices regarding these distributions have minimal impact upon the design of the practical detection/tracking algorithm: its structure remains identical, and the only changes will be in the evaluation of the integrals on the single-object state space \mathbf{X} in Eqs (5), (10), (13), (14), and (15).

Two broad classes of implementation techniques are widely popular to represent store and maintain single-object probability distributions: particle filters,¹² and Gaussian mixtures.²⁷ For this first illustration of the DISP filter in a space surveillance problem, we are exploiting a relatively simple orbital motion model that could be adapted to the more restrictive Gaussian mixture representation. However, with the objective of an enrichment of the orbital motion model in mind, we have opted for a particle filter implementation, able to accommodate highly non-linear motion models, inspired from the recent disparity space approach.^{28,29} The number of particles per track is constant and set to 400.

In this scenario, then, each track $y \in Y_{t-1}$ maintained by the filter is propagated along its orbit during the prediction step (see Eq. 5 via the Shepperd matrix,^{30,31} with additional noise to account for the approximations, applied to each particle composing the prior distribution p_{t-1}^y . No target disappearance is expected in this scenario; thus, the probability of survival $p_{s,t}$ is set to almost one ($1 - 10^{-3}$) in the surveillance scene but drops to zero along its physical edges.

Update step and orbit determination

One of the major challenges in space surveillance problems is that the single-object probability distributions in the Cartesian reference frame may not be accurately parametrised into a simple analytical form (e.g. Gaussian distribution), for both the orbital mechanics and the sensor observation process are highly non-linear in nature.

However, since the sensor observation noise is modelled as a multivariate Gaussian distribution with uncorrelated noise on each component $r, \theta, \varphi, \dot{r}$, it may be expected that the same single-object probability distributions, once projected into the augmented sensor space (a full spherical frame centred on the sensor), assume a simpler form that can be approximated by a multivariate Gaussian distribution. Inspired by recent

works on multi-object joint detection/tracking problems with cameras exploiting this principle,²⁹ we apply in this paper a similar technique to a Doppler radar sensor.

Measurement update

The measurement update (13) is implemented as illustrated in Figure 6 (depicted in 2D for the sake of clarity):

1. The particle states of the predicted distribution \hat{p}_t^y are transformed into the augmented sensor space (full 3D spherical coordinates);
2. The resulting distribution is approximated as a Gaussian distribution, in the same space;
3. The Gaussian distribution is updated, with a classic Kalman filter, using the measurement in the sensor space (3D spherical coordinates without angular rates);
4. New particles are sampled from the updated Gaussian distribution, still in the augmented sensor space;
5. The particles states are transformed back into the reference space and yield the updated distribution $p_t^{y:z}$.

Initial Orbit determination

Due to the structure of the data association mechanism of the DISP filter (see Figure 4), the probability distribution $p_t^{\phi_{t-1}:z}$ of a newborn track must be determined upon a *single* measurement $z \in Z_t$ (see Eq. (15)). In the context of this scenario, no prior information is assumed by the operator regarding the orbits on which the individuals evolve. For this reason, the initial distribution of the appearing individuals \hat{p}_t^a is uninformative regarding the *position* of the individuals within the sensor's current field of view.

Given the position and range rate coordinates of an appearing individual provided by the Doppler radar through some measurement $(r, \theta, \varphi, \dot{r})$, an admissible region for the unknown angular rates $(\dot{\theta}, \dot{\varphi})$ can be shaped from internal energy constraints.^{32,33} Since the appearing individual is on a real orbit, its internal energy is non-positive, i.e.

$$\frac{1}{2}|\dot{\mathbf{r}}|^2 - \frac{\mu}{|\mathbf{r}|} \leq 0, \quad (18)$$

where μ is the Earth mass times the gravitational constant and \mathbf{r} is the geocentric inertial position vector of the individual. The vectors \mathbf{r} , $\dot{\mathbf{r}}$ can then be expressed as follows:

$$\begin{cases} \mathbf{r} &= \mathbf{r}_s + r\boldsymbol{\rho}_r \\ \dot{\mathbf{r}} &= \dot{\mathbf{r}}_s + \dot{r}\boldsymbol{\rho}_r + r\dot{\theta}\boldsymbol{\rho}_\theta + r\dot{\varphi}\boldsymbol{\rho}_\varphi, \end{cases} \quad (19)$$

where \mathbf{r}_s is the geocentric inertial position vector of the sensor, and $\boldsymbol{\rho}_r, \boldsymbol{\rho}_\theta, \boldsymbol{\rho}_\varphi$ the unit vectors associated to the measurement $(r, \theta, \varphi, \dot{r})$.

Substituting (19) into (18) yields the new form

$$\alpha_1\dot{\theta}^2 + \alpha_2\dot{\varphi}^2 + \alpha_3\dot{\theta} + \alpha_4\dot{\varphi} + a_5 \leq 0, \quad (20)$$

where the parameters α_i are as follows:

$$\begin{cases} \alpha_1 &= r^2 \cos^2 \varphi \\ \alpha_2 &= r^2 \\ \alpha_3 &= \frac{r\dot{\mathbf{r}}_s \cdot \boldsymbol{\rho}_\theta}{2} \\ \alpha_4 &= \frac{r\dot{\mathbf{r}}_s \cdot \boldsymbol{\rho}_\varphi}{2} \\ \alpha_5 &= \dot{r}^2 + 2\dot{r}\dot{\mathbf{r}}_s \cdot \boldsymbol{\rho}_r + |\dot{\mathbf{r}}_s|^2 - \frac{2\mu}{\sqrt{r^2 + 2r\mathbf{r}_s \cdot \boldsymbol{\rho}_r + |\mathbf{r}_s|^2}}. \end{cases} \quad (21)$$

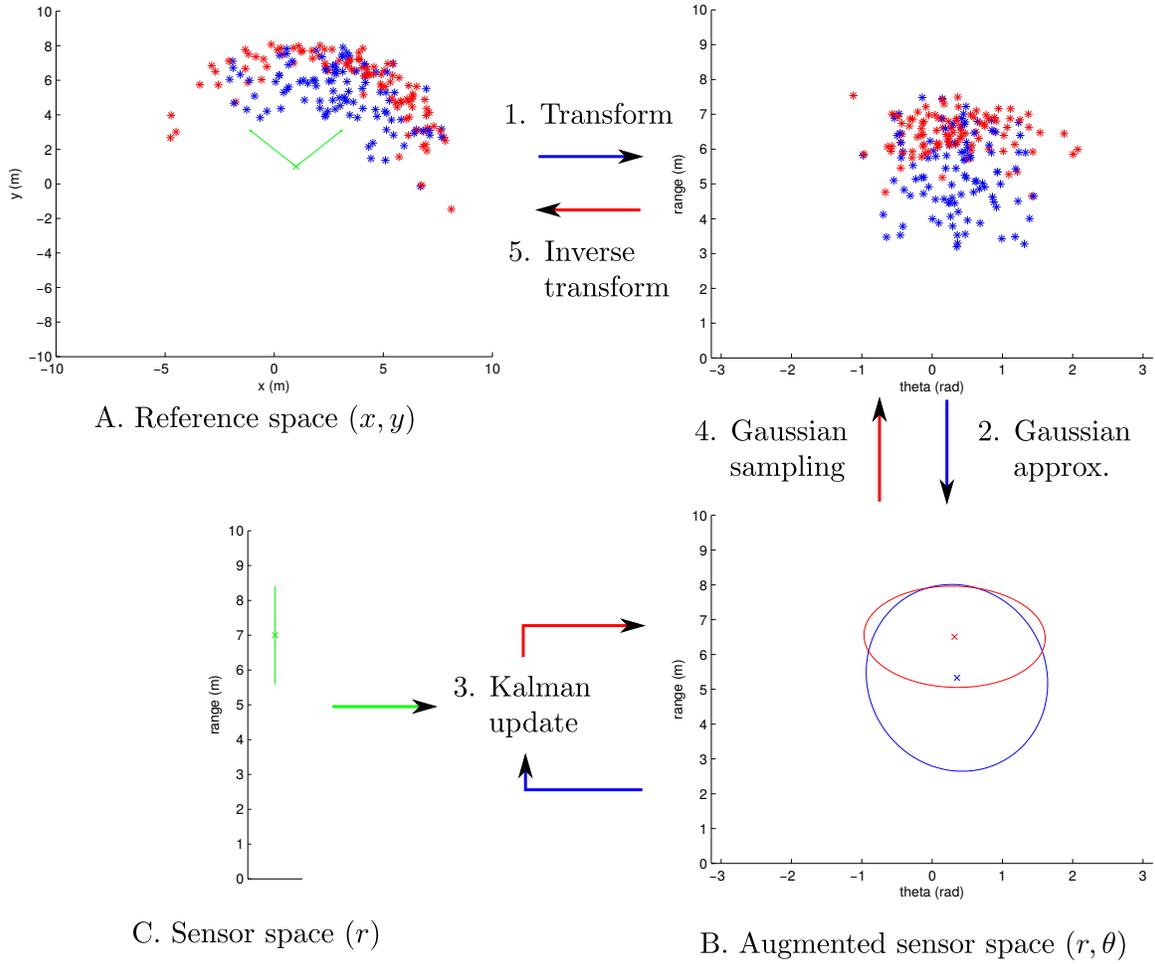


Figure 6: Principle of Kalman data update in augmented sensor space (illustration in 2D). The blue distribution in the reference space A corresponds to some predicted probability distribution \hat{p}_t^y , and the red one to the updated probability distribution $p_t^{y:z}$, where $z \in Z_t$ is a range-only measurement depicted, with uncertainty, in the sensor space C.

From (20), the boundaries of the admissible region for the unknown parameters $(\hat{\theta}, \hat{\varphi})$ are found to be:³⁴

$$\begin{cases} \hat{\theta} &= \frac{\alpha_3}{\alpha_1} + \sqrt{\frac{\alpha_3^2}{\alpha_1^2} + \frac{\alpha_4^2}{\alpha_1 \alpha_2} - \frac{\alpha_5}{\alpha_1}} \cos \phi \\ \hat{\varphi} &= \frac{\alpha_4}{\alpha_2} + \sqrt{\frac{\alpha_3^2}{\alpha_1 \alpha_2} + \frac{\alpha_4^2}{\alpha_2^2} - \frac{\alpha_5}{\alpha_2}} \sin \phi, \end{cases} \quad (22)$$

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

Following the measurement update process illustrated in Figure 6, the mean and variance parameters on the components $(\hat{\theta}, \hat{\varphi})$ of the initial Gaussian distribution of an appearing individual (red ellipse in the augmented sensor space B) are thus chosen such that the sampled particle distribution populates the admissible region adequately (in particular, the standard deviations on the $(\hat{\theta}, \hat{\varphi})$ components are set to the values of the corresponding square root terms in Eq. (22)).

Filter output: MAP estimates of the single-object probability distributions p_t^y

In the general case, the posterior probability distributions p_t^y of the tracks cannot be parametrised into a simple analytical form in the Cartesian reference frame – most notably, if the track has not been detected for a long time, the particles are usually spread along the estimated orbit in a “banana shape” – and the extraction of a MAP estimate for displaying purposes is thus rather challenging. Projected in the full 3D spherical frame centred on the sensor, however, these probability distributions may assume a simpler form that may be adequately approximated by a Gaussian. In order to extract these MAP estimates, we followed the same approach as depicted in Figure 6:

1. The particle states of the posterior distribution p_t^y are transformed into the augmented sensor space;
2. The resulting distribution is approximated as a Gaussian distribution, in the same space;
3. The mean of the Gaussian distribution, transformed back in the Cartesian reference space, yields the MAP estimate of the track y .

Approximations for the practical implementation

From the structure of the DISP filter it can be seen that no information loss is incurred during the filtering process, in the sense that *a)* the data association step produces a track for every possible observation path, and an hypothesis for every possible combination of compatible observation paths, and, *b)* neither tracks nor hypotheses, regardless of their credibility, are ever discarded by the filter. 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. However, the probabilistic nature of the filtering framework allows for the operator to make principled approximations in order to reduce the complexity of the algorithm to available resources.

Discarding hypotheses and/or tracks with low credibility is perhaps the most intuitive and straightforward approximation to implement. In the context of this scenario where the density of false alarm is relatively scarce, it is efficient to discard tentative tracks that were initiated through a false alarm, but remained within the sensor’s field of view without being confirmed by subsequent coherent measurements. Recall that the credibility of an hypothesis $h \in H_t$ (respectively a track $y \in Y_t$) is directly given by its probability of existence $c_t(h)$ (respectively c_t^y), and a simple threshold on these quantities provides an efficient way to curtail the growth in their number. In this scenario, the threshold on existence for hypotheses (respectively tracks) is set to 10^{-30} (respectively 10^{-7}). A complementary approach consists in controlling the number of hypotheses and/or tracks following each time step by discarding those with the lowest credibility. Considering the small number of targets involved in this scenario, the maximum number of hypotheses to propagate is set to 30.

Another option is to discard tracks representing individuals that are likely to have left the surveillance scene, for they may be of little interest to the operator. Recall that the mass of the probability distribution of some track $y \in Y_t$ concentrated in the single-object state space, i.e. $p_t^y(\mathbf{X})$, is the probability of presence of the underlying target in the scene; therefore, tracks that are a likely to be absent from the scene can be discarded through a simple threshold on this quantity (set to 10^{-5} in this scenario).

Although the simultaneous appearance of several individuals in the sensor’s field of view is covered by the data association mechanism – the number of appearing individuals in Eq. (8) may be as high as the number of collected measurements – limiting the number of appearances to one per time step is a reasonable approximation in this context that reduces significantly the number of hypotheses generated during the data update step. This approximation is nonetheless lifted during the initial step, for all the targets within the initial field of view are bound to appear simultaneously (provided that they are detected) as soon as the sensor is “switched on”.

SIMULATION RESULTS

The preliminary results of the multi-object surveillance problem are depicted in Figure 7. A first general comment is that, despite the challenging conditions regarding the observability – limited field of view and relatively low probability of detection – the DISP filter proves able, on this scenario, to initiate and propagate a track for each individual that has already entered the field of view.

As expected, individual 1 is estimated with the best accuracy since its state is updated consistently throughout the whole scenario. The estimation of individuals 2 and 4 show similar performances: the time period during which the individual remains in the field of view is seemingly too short to allow for an appropriate correction of the state initialisation provided by the orbit determination technique. Individual 3 also shows a quick degradation of the state estimation once the individual has left the field of view; the estimation is quickly then corrected once the individual hits the field of view a second time; however, the discrepancies between the estimated track and the individual remains at a significant level, and the filter eventually assigns the remaining stream of measurements to a *new* track. The estimation of individual 5 shows a similar degradation of performance than individuals 2 and 4. Around time step 300, the uncertainty regarding the position of the individual is large enough for a lot of particles to have left the surveillance scene: the probability of presence of the corresponding track quickly drops to negligible levels, and the track is eventually discarded by the filter.

An important fact that can be drawn from these results is that the DISP filter does maintain tracks whose Mahalanobis distances to the corresponding individuals, taking into account the *uncertainty* on the filter in its estimation, remain low throughout the scenario. In other words, the quality of the estimation quickly degrades once the individuals have left the sensor’s field of view, *but the filter is aware of this degradation*. This suggests that the orbital motion model used in the prediction step, the sole source of information about individuals that are away from the sensor’s field of view, is *relatively inaccurate but consistent* in the context of this scenario.

CONCLUSION AND FURTHER DEVELOPMENTS

This paper presents the DISP filter, a novel multi-object joint detection/tracking algorithm derived from the very recently developed estimation framework for stochastic populations, in the context of wide area surveillance for space situational awareness. It is illustrated on a challenging multi-object surveillance problem involving five targets orbiting Earth on five different trajectories, observed by a single sensor with limited coverage.

In this scenario, the preliminary results suggest that the DISP filter is able to initialise and maintain specific information (e.g. tracks) on all the individuals that have already entered the sensor’s field of view, despite the challenging conditions. The degradation in quality of the estimation, which occurs once the individuals have left the field of view, is acknowledged and maintained by the filter through an accurate estimate of uncertainty.

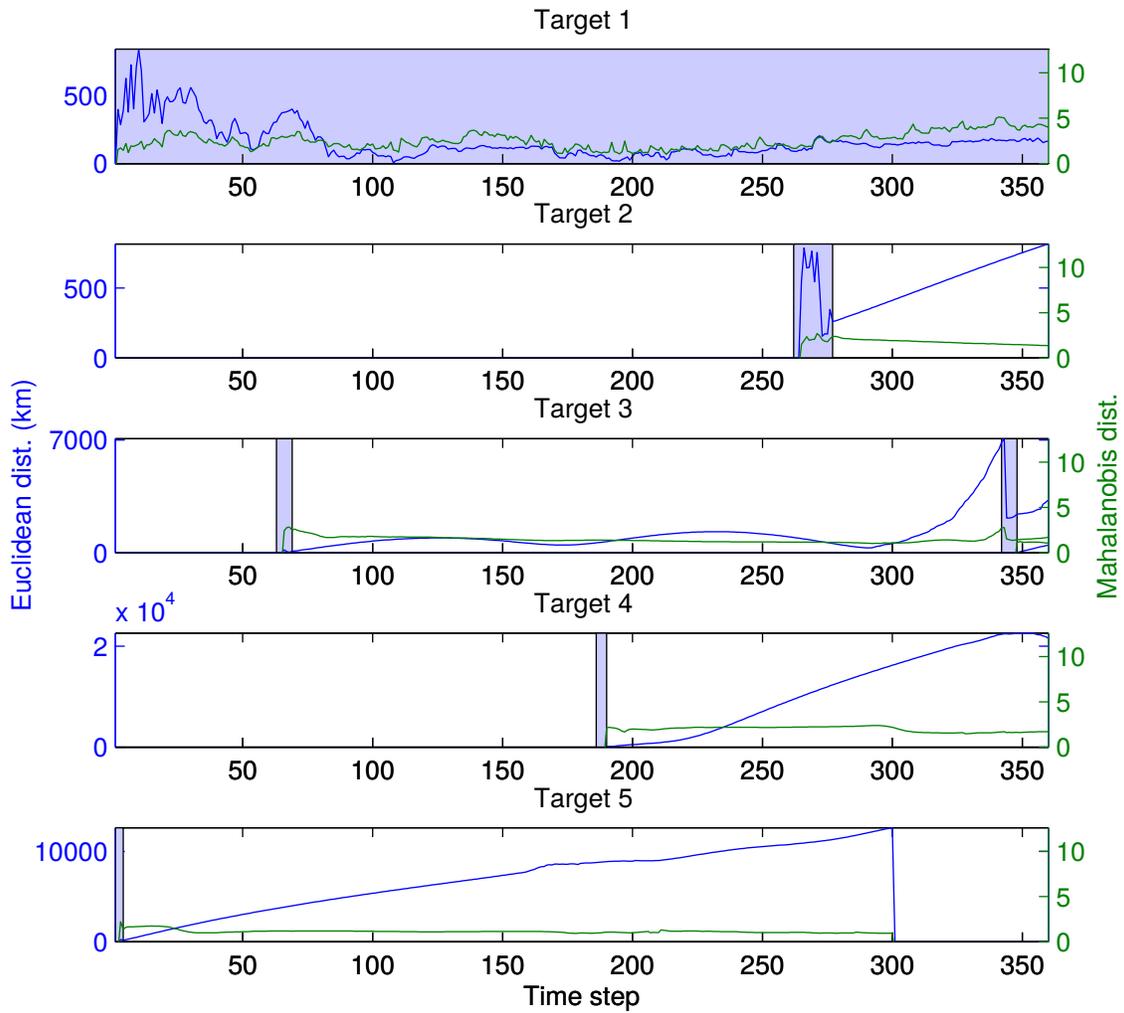


Figure 7: Simulation results (one figure per individual). Light blue zones corresponds to time periods where the individual lies within the sensor’s field of view. The Euclidean distance (position) between an individual and the MAP estimate of the corresponding track is depicted in blue, while the Mahalanobis distance (full state) between an individual and the corresponding track is depicted in green. The maximum value on the right axis (Mahalanobis distance) corresponds to the 95% confidence interval.

A natural route lead for future research would involve the incorporation of an alternative orbital motion model in the Bayesian prediction step, in order to improve the performance of the filter in the estimation of targets while they are outside sensor's field of view.

ACKNOWLEDGMENTS

This work was supported by the USAF EOARD Grant 13-3030, the Engineering and Physical Sciences Research Council (EPSRC) Platform Grant (EP/J015180/1), the EPSRC Grant number EP/K014277/1 and the UK MoD University Defence Research Collaboration in Signal Processing, and UK MoD funding under contract DSTL/AGR/00363/01. Carolin has been supported through the VIBOT Erasmus Mundus Visiting Scholar programme. Jeremie Houssineau has a PhD scholarship sponsored by DCNS and a tuition fee scholarship sponsored by Heriot-Watt University.

REFERENCES

- [1] E. Delande, J. Houssineau, and D. E. Clark, "A filter for distinguishable and independent populations," arXiv:1501.04671.
- [2] Y. Bar-Shalom, "Tracking methods in a multitarget environment," *Automatic Control, IEEE Transactions on*, Vol. 23, Aug. 1978, pp. 618 – 626.
- [3] D. Reid, "An Algorithm for Tracking Multiple Targets," *Automatic Control, IEEE Transactions on*, Vol. 24, Dec. 1979, pp. 843–854.
- [4] S. S. Blackman and R. Popoli, *Design and analysis of modern tracking systems*. Artech House, 1999.
- [5] W. D. Blair and Y. Bar-Shalom, eds., *Multitarget-Multisensor Tracking: Applications and Advances (Volume III)*. Artech House, 2000.
- [6] R. P. S. Mahler, "Multitarget Bayes Filtering via First-Order Multitarget Moments," *Aerospace and Electronic Systems, IEEE Transactions on*, Vol. 39, Oct. 2003, pp. 1152–1178.
- [7] R. P. S. Mahler, "PHD Filters of Higher Order in Target Number," *Aerospace and Electronic Systems, IEEE Transactions on*, Vol. 43, Oct. 2007, pp. 1523–1543.
- [8] R. P. S. Mahler, *Statistical Multisource-Multitarget Information Fusion*. Artech House, 2007.
- [9] R. E. Kalman, "A New Approach to Linear Filtering and Prediction Problems," *ASME–Journal of Basic Engineering*, Vol. 82, No. D, 1960, pp. 32–45.
- [10] B. D. O. Anderson and J. B. Moore, *Optimal Filtering*. Prentice-Hall, 1979.
- [11] S. Julier and J. K. Uhlmann, "Unscented filtering and nonlinear estimation," *Proceedings of the IEEE*, Vol. 92, Mar. 2004, pp. 401 – 422.
- [12] A. Doucet, N. d. Freitas, and N. Gordon, *Sequential Monte Carlo Methods in Practice*. Statistics for Engineering and Information Science, Springer, 2001.
- [13] D. Stoyan, W. S. Kendall, and J. Mecke, *Stochastic geometry and its applications*. Wiley, 2 ed., Sept. 1995.
- [14] D. Vere-Jones and D. J. Daley, *An Introduction to the Theory of Point Processes*. Springer Series in Statistics, 1 ed., 1988.
- [15] I. I. Hussein, K. J. DeMars, C. Früh, R. S. Erwin, and M. K. Jah, "An AEGIS-FISST integrated detection and tracking approach to Space Situational Awareness," *Information Fusion, Proceedings of the 15th International Conference on*, July 2012, pp. 2065–2072.
- [16] I. I. Hussein, K. J. DeMars, C. Früh, M. K. Jah, and R. S. Erwin, "An AEGIS-FISST Algorithm for Multiple Object Tracking in Space Situational Awareness," *AIAA Guidance, Navigation, and Control Conference 2012*, Aug. 2012.
- [17] E. Delande, M. Uney, J. Houssineau, and D. E. Clark, "Regional Variance for Multi-Object Filtering," *IEEE Transactions on Signal Processing*, Vol. 62, July 2014, pp. 3415 – 3428.
- [18] E. Delande, J. Houssineau, and D. E. Clark, "Regional variance in target number: analysis and application for multi-Bernoulli point processes," *Data Fusion & Target Tracking Conference (DF&TT 2014): Algorithms & Applications, 10th IET*, Apr. 2014.
- [19] B.-T. Vo and B.-N. Vo, "Labeled Random Finite Sets and Multi-Object Conjugate Priors," *Signal Processing, IEEE Transactions on*, Vol. 61, Jul 2013, pp. 3460 – 3475.
- [20] B.-T. Vo and B.-N. Vo, "Tracking, identification, and classification with random finite sets," *Signal Processing, Sensor Fusion, and Target Recognition XXII, Proceedings of SPIE*, Apr 2013, p. 87450D.
- [21] J. Houssineau and D. E. Clark, "Hypothesised filter for independent stochastic populations," 2014, arXiv:1404.7408.

- [22] E. Delande, J. Houssineau, and D. E. Clark, "Performance metric in closed-loop sensor management for stochastic populations," *Sensor Signal Processing for Defence (SSPD)*, 2014. <http://home.eps.hw.ac.uk/jh207/> (preprint).
- [23] Y. Pailhas, J. Houssineau, E. Delande, Y. Petillot, and D. E. Clark, "Tracking underwater objects using large MIMO sonar systems," *International Conference on Underwater Acoustic*, 2014.
- [24] V. I. Bogachev, *Measure Theory*, Vol. 1. Springer, 2007.
- [25] V. I. Bogachev, *Measure Theory*, Vol. 2. Springer, 2007.
- [26] M. Pace and P. Del Moral, "Mean-field PHD filters based on generalized Feynman-Kac flow," *Selected Topics in Signal Processing, IEEE Journal of*, Vol. 7, June 2013, pp. 484 – 495.
- [27] D. L. Alspach and H. W. Sorenson, "Nonlinear Bayesian estimation using Gaussian sum approximations," *Automatic Control, IEEE Transactions on*, Vol. 17, Aug. 1972, pp. 439 – 448.
- [28] J. Houssineau, S. Ivekovic, and D. E. Clark, "Disparity space: A parameterisation for Bayesian triangulation from multiple cameras," *Information Fusion, 2012 15th International Conference on*, July 2012, pp. 1734–1740.
- [29] J. Houssineau, D. E. Clark, S. Ivekovic, C. S. Lee, and J. Franco, "A unified approach for multi-object triangulation, tracking and camera calibration," Oct. 2014. arXiv:1410.2535v1.
- [30] S. W. Shepperd, "Universal Keplerian state transition matrix," *Celestial Mechanics*, Vol. 35, 1985, pp. 129 – 144.
- [31] W. H. Goodyear, "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.
- [32] G. Tommei, A. Milani, and A. Rossi, "Orbit determination of space debris: admissible regions," *Celestial Mechanics and Dynamical Astronomy*, Vol. 97, No. 4, 2007, pp. 289–304.
- [33] D. Farnocchia, G. Tommei, A. Milani, and A. Rossi, "Innovative methods of correlation and orbit determination for space debris," *Celestial Mechanics and Dynamical Astronomy*, Vol. 107, No. 1-2, 2010, pp. 169–185.
- [34] K. J. DeMars and M. K. Jah, "Probabilistic Initial Orbit Determination using Radar Returns," No. AAS 13-704 in *Advances in the Astronautical Sciences*, 2013.