Sensor management with regional statistics for the PHD filter

Marian Andrecki, Emmanuel D. Delande, Jérémie Houssineau, and Daniel E. Clark
School of Engineering & Physical Sciences, Heriot-Watt University, Edinburgh, UK

***Abstract*—This paper investigates a sensor management scheme that aims at minimising the regional variance in the number of objects present in regions of interest whilst performing multi-target filtering with the Probability Hypothesis Density (PHD) filter. The experiments are conducted in a simulated environment with groups of targets moving through a scene in order to inspect the behaviour of the manager. The results demonstrate that computing the variance in the number of objects in different regions provides a viable means of increasing situational awareness where complete coverage is not possible. A discussion follows, highlighting the limitations of the PHD filter and discussing the applicability of the proposed method to alternative available approaches in multi-object filtering.**

I. INTRODUCTION

Exploring sensor management problems becomes a topic of increasing interest as the controllability of modern sensor systems advances. In the context of Bayesian estimation for target detection/tracking, the information acquired by the operator on the object of interests - the targets - is provided by one or several sensors observing the surveillance scene. If the operator has some degree of freedom in the dynamical control of these sensors (e.g. orientation of camera, selection of radar scanning modes), its efficiency can be improved by solving a sensor management problem through finding a suitable sensor policy, i.e., a principled decision-making procedure that will dynamically select a sensor control based on the acquired information in order to achieve some surveillance objective.

Approaches to the sensor management problem are extremely varied, and their applicability depend on the nature of the information propagated by the tracking algorithm (see [1] for a recent account on the topic). The regional statistics provide first- and second-order information on the target activity in any region of the surveillance space, i.e., the expected number of targets lying within the said region and associated uncertainty [2], [3]. Because they produce a principled and meaningful quantification of the estimated target activity in any desired region of the surveillance space, the regional statistics provide grounds for a simple sensor policy aiming at focusing the surveillance activity on regions of critical interest to the operator – say, the surroundings of a facility.

Naturally available to any multi-object filter maintaining a probabilistic representation of the population of targets, the regional statistics for filtering solutions stemming from the Finite Set Statistics (FISST) framework [4] are derived in [2], [5], [6] for the PHD filter [7] and the Cardinalized Probability Hypothesis Density (CPHD) filter [8], and discussed in [3] for the multi-Bernoulli filter [4]. Applicability to traditional track-based filters such as the Multiple Hypothesis Tracking (MHT) and Joint Probabilistic Density Association (JPDA) approaches [9] is less straightforward and will be briefly discussed in this paper.

This paper aims at illustrating the concept of regional statistics as a basis for a sensor policy in the context of the PHD filter, the simplest solution derived from the FISST framework [4]. Section II provides a brief description of the PHD filter with regional statistics (a detailed construction can be found in [2]). Section III describes the surveillance activity on which the sensor policy exploiting regional statistics will be illustrated, and simulation results are presented in Section IV. Section V analyses the limited applicability of the PHD filter in this context and discusses leads for further development.

II. PHD FILTERING WITH REGIONAL STATISTICS

A. Multi-object filtering and regional statistics

The FISST [8] is a filtering framework in which the acquired information about the multi-target state, i.e., the states of all the targets currently living in the surveillance scene, is represented by a random object Φ called a Random Finite Set (RFS). A realization $\varphi = \{x_1, \dots, x_n\}$ of Φ is a *set* of individual target states x_i in a (single-)target state space $\mathbf{X} \subset \mathbb{R}^d$, which is a d-dimensional space describing the targets' physical characteristics of interest to the operator (position, velocity, etc.). As with any random object, a RFS is described by its probability distribution P_Φ, defined on the process state space $\mathcal{X}$, i.e., the space of all the finite sets of points in $\mathbf{X}$.

The multitarget Bayes filter [4] integrates the Bayesian paradigm within the FISST framework and provides filtering equations for the propagation of the multi-target probability distribution P_Φ, but it is intractable in the general case. Practical filters rely on assumptions on the multi-target RFS Φ leading to more tractable equations propagating *reduced information* on Φ. As with any random variables, statistical moments can be defined on any RFS Φ to produce a reduced, yet meaningful description out of the full probability distribution P_Φ. The *first-order moment measure*[1] μ_Φ plays a central

M. Andrecki, E. D. Delande, J. Houssineau, and D. E. Clark are with the School of Engineering & Physical Sciences, Heriot-Watt University, Edinburgh, UK (e-mails: ma804@hw.ac.uk, E.D.Delande@hw.ac.uk, J.Houssineau@hw.ac.uk and D.E.Clark@hw.ac.uk).

[1] In this paper we follow the approach in [2] and describe RFSs with probability measures handled with standard integrals, rather than with multi-object densities handled with set integrals [8], for the measure-theoretic formulation is required for the expression of the regional statistics.

role in the construction of RFS-based filters and is propagated by both PHD [7] and CPHD [8] filters. The *centred second moment* or *variance* var_Φ was introduced more recently in the context of the PHD and CPHD [2], [5], [6], and the multi-Bernoulli [3] filters.

Given some RFS Φ and some arbitrary[2] region $B \subseteq \mathbf{X}$, the regional statistics $\mu_\Phi(B), \mathrm{var}_\Phi(B)$ are defined in [2] and can be interpreted as follows:

- $\mu_\Phi(B)$ is the expected number of targets within B,
- $\mathrm{var}_\Phi(B)$ quantifies the spread, around its expected value, of the estimated number of targets within B.

B. The PHD filter with regional statistics

For any time step k, we denote by $\Phi_{k|k-1}$ (respectively (resp.) Φ_k) the predicted (resp. posterior) RFS representing the current multi-target configuration, given the measurements collected up to time step $k-1$ (and k). The first-order moment measure of the predicted (and posterior) RFS is denoted by $\mu_{k|k-1}$ (resp. μ_k). Further denoting by $\mu_{\mathrm{b},k}$ (and $\mu_{\mathrm{fa},k}$) the first-order moment measure (resp. density) of the newborn target RFS $\Phi_{\mathrm{b},k}$ (and false alarm RFS $\Psi_{\mathrm{fa},k}$), the filtering equations of the PHD filter are given by [7]

$$\mu_{k|k-1}(\mathrm{d}y) = \int p_{\mathrm{s},k}(x) f_{\mathrm{s},k}(\mathrm{d}y|x)\mu_{k-1}(\mathrm{d}x) + \mu_{\mathrm{b},k}(\mathrm{d}y), \quad (1)$$

$$\mu_k(\mathrm{d}y) = \mathrm{C}_k(y|Z_k)\mu_{k|k-1}(\mathrm{d}y), \quad (2)$$

with

$$\mathrm{C}_k(y|Z_k) = g_k(\phi|y) + \sum_{z\in Z_k} \frac{g_k(z|y)}{\mu_{\mathrm{fa},k}(z) + \int g_k(z|x)\mu_{k|k-1}(\mathrm{d}x)}, \quad (3)$$

where g_k is the extended sensor likelihood function given by

$$\begin{cases} g_k(\phi|\cdot) &= 1 - p_{\mathrm{d},k}(\cdot), \\ g_k(z|\cdot) &= p_{\mathrm{d},k}(\cdot)\ell_k(z|\cdot), \end{cases} \quad (4)$$

and where

- $p_{\mathrm{s},k}$ is the target probability of survival,
- $f_{\mathrm{s},k}$ is the Markov transition kernel,
- $p_{\mathrm{d},k}$ is the sensor probability of detection,
- ℓ_k is the sensor likelihood function,
- Z_k is the set of collected observations.

The posterior regional statistics $\mu_k(B), \mathrm{var}_k(B)$ are given by [2]

$$\mu_k(B) = \mu_\phi(B) + \sum_{z\in Z_k} \frac{\mu_z(B)}{\mu_{\mathrm{fa},k}(z) + \mu_z(\mathbf{X})}, \quad (5)$$

$$\mathrm{var}_k(B) = \mu_\phi(B) + \sum_{z\in Z_k} \frac{\mu_z(B)}{\mu_{\mathrm{fa},k}(z) + \mu_z(\mathbf{X})}\left(1 - \frac{\mu_z(B)}{\mu_{\mathrm{fa},k}(z) + \mu_z(\mathbf{X})}\right), \quad (6)$$

where, for any region $B \subseteq \mathbf{X}$,

$$\begin{cases} \mu_\phi(B) &= \int_B g_k(\phi|x)\mu_{k|k-1}(\mathrm{d}x), \\ \mu_z(B) &= \int_B g_k(z|x)\mu_{k|k-1}(\mathrm{d}x). \end{cases} \quad (7)$$

[2]Provided that is it measurable, see [2] for more details.

An analysis of the posterior regional statistics for the PHD filter is given in [5].

From now on, the newborn target RFS $\Phi_{\mathrm{b},k}$ (resp. the false alarm RFS $\Psi_{\mathrm{fa},k}$) is assumed Poisson [4] with rate $\lambda_{\mathrm{b},k}$ (resp. $\lambda_{\mathrm{fa},k}$). Also, time subscripts will be omitted when there is no ambiguity.

III. Sensor management problem

A. Surveillance activity

The surveillance scene is 2D region of the physical space, in which targets of interest are expected to enter the scene from the west (left edge), and leave through the east. Little is known by the operator about the population of incoming targets. The general surveillance objective is to estimate as accurately as possible the size of the population within specific regions of the scene at every time step. The sensor coverage is very limited (see Section III-B), and the surveillance is primarily focussed on three regions of the scene (dotted rectangles in Figure 1). Note that the identification and classification of individual targets are out of scope in this context, for the sole purpose of this surveillance activity is to provide a first estimate of the amount of incoming target activity. This situation could correspond to the first task in a border surveillance activity, where target identification, classification and accurate state estimation would be carried on by another sensor system further east.

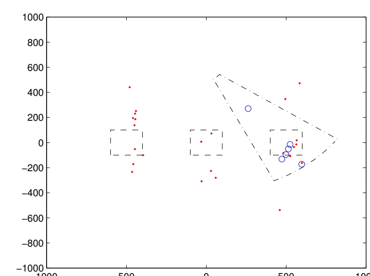

Fig. 1: Example of scenario. The focus regions are denoted by dotted rectangles, and the current field of view is denoted by a dotted triangle. The targets are red dots, and the current detections are blue circles.

From now on, B_1 (resp. B_2, B_3) will denote the westernmost (resp. centred, easternmost) focus region within the surveillance area.

B. Sensor

The only available sensor is a radar providing range and bearing, positioned north of the scene. The sensor can be oriented towards one, and only one, of the three focus regions at any time. Note that the sensor's field of view (dotted triangle int Figure 1) is not limited to the focus region the sensor is oriented at, but the three possible field of views do not overlap. Note that the observation process is usually noisy, for *a*) the targets lying within the field of view may be miss-detected, and *b*) the sensor may produce spurious measurements (false alarms), and *c*) Gaussian noise is added to the bearing and range measurements.

C. Sensor control policy

The regional statistics computed in the three focus regions provide a natural tool to quantify the quality of the estimation, for the variance gives the uncertainty in the level of target activity *as maintained by the filter*. The straightforward policy that we explore in this paper is therefore to focus the sensor, at each time step, on the region whose variance is the *highest*.

It is important to note from the PHD filtering equations (1), (2) that the posterior variance (6) of the multi-target process is *not* propagated across time, but merely extracted at each time step (this limitation of the PHD filter is discussed in Section V).

The proposed sensor control policy is given, at any time k, as follows:

1) Compute posterior variances $\mathrm{var}_{k-1}(B_i)$, $1 \leq i \leq 3$;
2) Focus sensor towards region $\arg\max \mathrm{var}_{k-1}(B_i)$.

IV. Simulation results

A. Implementation

The simulated targets appear in the scene in groups, whose size is Poisson with mean 15. Newborn targets appear along the western edge of the scene, move eastward and leave the scene from its eastern edge.

In order to illustrate the value of a sensor control policy, a scarcity in sensing resources was necessary so that the options of *exploitation* – keeping the sensor focussed on the same region to refine the estimation of the target population currently looked at – and *exploration* – orienting the sensor to another region to learn about a different population – were competitive. To this end, the sensor is modelled with poor measurement accuracy (12 m on range, 2.86° on bearing), and either low probability of detection and low false alarm rate ($p_\mathrm{d} = 0.6$, $\lambda_\mathrm{fa} = 1$), or higher probability of detection and but higher false alarm rate as well ($p_\mathrm{d} = 0.8$, $\lambda_\mathrm{fa} = 10$).

A target's state is described by its position and velocity coordinates in the surveillance scene. The PHD filter is implemented with Gaussian mixtures [10], and the regional statistics are computed in the three focus regions using (5), (6). The target behaviour is modelled with a near constant velocity model, with a slight noise to account for the model mismatches, and initial velocity towards the east.

A critical component of the filter parametrisation is the target birth model. Recall that little is known by the operator regarding the population of incoming targets. Introducing an overestimated number of potential targets in the scene through the birth model proved an incentive for the sensor to explore regularly the foremost region B_1 in order to prevent incoming waves from being missed by the operator. The influence of the intensity of the birth RFS in Eq. (1) on the performance of the filter is further discussed in Section IV-C; the default value for the birth rate, used in the experiment in Section IV-B, is $\lambda_\mathrm{b} = 1$ (e.g. on average, one target is expected to enter the scene at each time step).

B. Comparison of naive/variance-based control policies

This scenario aims at illustrating the behaviour of the variance-based sensor control policy proposed in Section III-C compared to a "naive" sensor control policy sequentially scanning the three focus regions. A single wave of targets goes through the surveillance scene, unknown to the operator. In this situation, the naive controller is expected to waste sensing resources on void regions, while the variance-based controller should focus either on the region where the wave of targets is (exploitation), or focus on the first region to look for potential incoming targets (exploration).

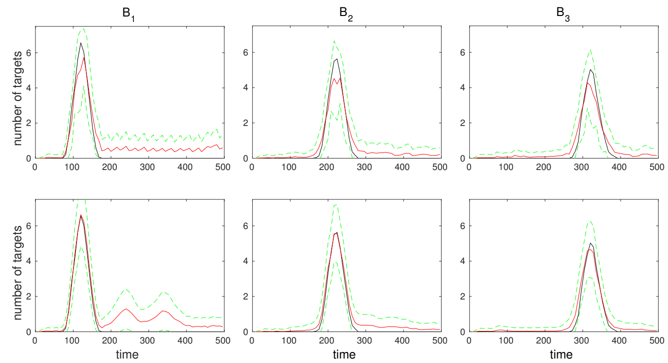

Fig. 2: Estimated target activity for naive (top) and variance-based (bottom) controllers ($p_\mathrm{d} = 0.6$, $\lambda_\mathrm{fa} = 1$). True target number is in black, mean target number is in red, +/- standard deviation is in green.

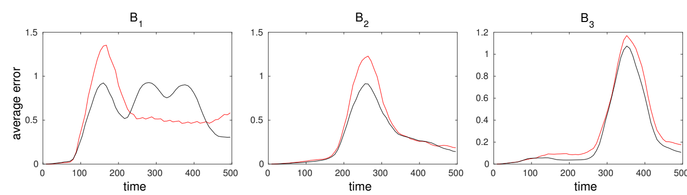

Fig. 3: Absolute error in target number ($p_\mathrm{d} = 0.6$, $\lambda_\mathrm{fa} = 1$). Naive controller is in red, variance-based controller is in black.

Figures 2 and 3 depict the results of the scenario with low probability of detection and low false alarm rate, averaged over 30 Monte Carlo runs. Overall, the variance-based controller performs better than the naive controller, for in the former case the sensor will tend to focus longer on a region while it is crossed by the wave of targets. Note that the target number is overestimated in the foremost region B_1, regardless of the chosen controller, once the wave of true targets has passed (from time step 180 onwards). Recall from Section IV-A that the intensity of the birth model must be significant enough to introduce "uncertainty" in the composition of the incoming population and provide an incentive to check the foremost region regularly enough for the detection of potential incoming targets. It creates a bias in the estimated target number in that particular region, especially for the variance-based controller during the later stages of the scenario where the sensor spreads its attention on the three regions. This point is further discussed in Section V. While the bias is clear in the foremost region B_1, it seems to disappear in the following two regions, for the false alarm rate is low enough to prevent a gross overestimation of the populations once the initial bias is corrected through successive observations.

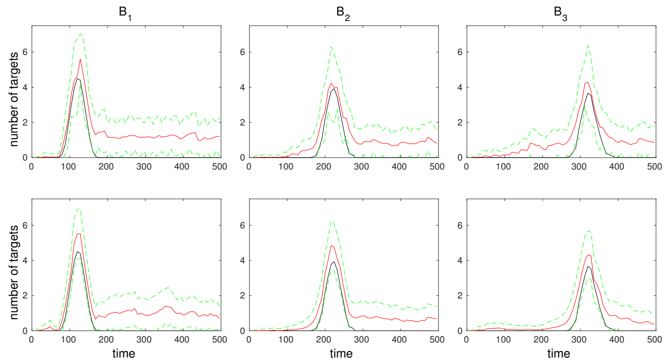

Fig. 4: Estimated target activity for naive (top) and variance-based (bottom) controllers ($p_{\mathrm{d}} = 0.8$, $\lambda_{\mathrm{fa}} = 10$). True target number is in black, mean target number is in red, +/- standard deviation is in green.

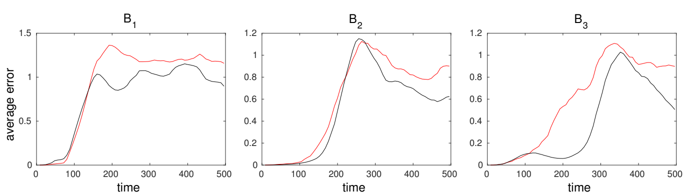

Fig. 5: Absolute error in target number ($p_{\mathrm{d}} = 0.8$, $\lambda_{\mathrm{fa}} = 10$). Naive controller is in red, variance-based controller is in black.

Figures 4 and 5 depict the results of the scenario with higher probability of detection and higher false alarm rate, averaged over 20 Monte Carlo runs. Regardless of the controller, the higher false alarm rate seems to induce an overestimation of the target activity that is not fully corrected by the filter. The variance-based controller still performs slightly better than the naive controller, especially in regions B_2 and B_3 when they are void of true targets. The variance-based controller is likely to focus on a region while the uncertainty in the target number is high enough, and in the process discards previous target evidence as false alarms if not backed by new measurements. On the other hand, a constant swiping between the different regions may not allow the filter to acquire enough evidence in a region to discriminate between false alarms and true targets.

C. Effect of the intensity of the birth model

This scenario aims at illustrating the dependence of the variance-based controller to the intensity of the birth model. Several groups of targets are spawned throughout the scenario, with a random time interval between two successive waves (the number and frequency of waves varied with the different Monte Carlo runs). The variance-based controller is tested with three different intensities for the birth model: $\lambda_{\mathrm{b}} = 0.05, 1, 20$. In all three cases, the probability of detection is set to $p_{\mathrm{d}} = 0.8$ and the false alarm rate to $\lambda_{\mathrm{fa}} = 1$.

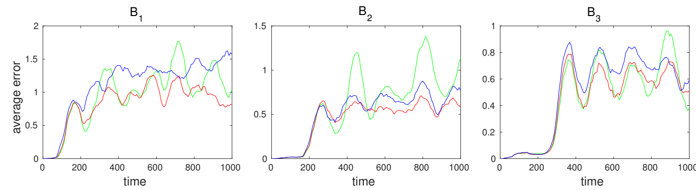

Fig. 6: Absolute error in target number ($p_{\mathrm{d}} = 0.8$, $\lambda_{\mathrm{fa}} = 1$). $\lambda_{\mathrm{b}} = 0.05$ is in green, $\lambda_{\mathrm{b}} = 1$ is in red, $\lambda_{\mathrm{b}} = 20$ is in blue.

Figure 6 depicts the results of the scenario, averaged over 20 Monte Carlo runs. Note that the three controllers perform similarly in the estimation of the first wave of targets hitting the foremost region B_1: at this point, the first wave receives full attention of the controllers as there are no other populations to focus on. Once the first wave of targets leave region B_1, the behaviour of the three controllers present noticeable differences.

The controller with low intensity birth model (green in Figure 6) has little incentive to explore the foremost region B_1. It tends to focus on region B_1 until hit by the first wave of targets, and then follows the first wave as it progresses eastward through regions B_2 and B_3. In overall it performs better than the other controllers in the estimation of the first wave, but tends to miss some of the subsequent incoming waves, resulting in large spikes in the estimation error. The errors are less significant in the easternmost region B_3, suggesting that by the time the subsequent waves have reach B_3 they have usually been detected. In overall this controller seems to performs the worst as it is prone to miss incoming waves.

The controller with high intensity birth model (blue in Figure 6) has strong incentive to check regularly on the foremost region B_1, since a large number of incoming targets are expected. All the target waves are detected when they hit the foremost region, and are subsequently tracked in the following regions; however, the controller spends little time focussing on a given wave a target as the exploration of the foremost region B_1 requires significant attention.

Overall, the controller with medium intensity birth model (red in Figure 6) seems to perform best, as it is more likely to balance the need for exploration and tracking of confirmed target waves properly.

V. Conclusion and further development

This paper proposes a novel sensor control policy for RFS-based filters based on the recently developed regional statistics. We exploit the concept of regional variance in the target number [2], which provides a principled quantification of the uncertainty in the estimation of the target activity in any desired region of the surveillance scene. It is implemented for the PHD filter [7] illustrated on a scenario simulating a border surveillance activity. The posterior regional statistics have a remarkably simple expression in the context of the PHD filter, and they can be implemented without incurring a significant computational cost. The inherent simplicity of the underlying filter, however, limits the applicability of the regional statistics for sensor management in complex scenarios.

The discussions in Section IV showed the critical importance of the birth model in this sensor management scenario. Because the sensor coverage is very limited and the targets are not detected upon entering the scene, the birth model must be designed in order to balance the need for exploitation of previously confirmed targets and exploration of the surveillance scene for the detection of new individuals. The structure of the PHD filter is such that the birth model is only described by its intensity μ_{b}, and therefore no second-order information is available for a more refined modelling of the incoming targets. A birth model with a large enough intensity seems necessary to provide incentive for the controller to check the activity in the foremost region B_1 regularly enough, but it induces an undesirable bias leading to an overestimation of the target activity in this region. An interesting lead to follow would be to model the newborn targets with a RFS with a regional variance *larger* than its regional mean, such as an over-dispersed Poisson point process [11], and derive an appropriate RFS-based filter able to capture its features. An obvious limitation of the PHD filtering framework, in the context of sensor management for the estimation of regional target activity, is that the uncertainty on the filtered estimate is not propagated across time – recall that the regional variance is *extracted* from the posterior probability distribution, yet *not* propagated to the next time step. Since the regional statistics can be produced through a principled procedure for any multi-target filter providing a probabilistic representation of the population of targets [2], a natural follow-up to this paper would be to design a similar controller for more recent filtering solutions propagating higher-order information. Promising candidates would be the Hypothesised and Independent Stochastic Populations (HISP) filter [12], whose computational complexity is similar to the PHD filter, or the more involved Distinguishable and Independent Stochastic Populations (DISP) filter [13] for more challenging scenarios.

On a broader perspective, alternative statistical tools to the regional statistics should be explored for the design of more generic control policies. Indeed, the applicability and relevance of the regional statistics for more traditional track-based approaches such as the MHT or JPDA approaches [9] remain unclear, for they do not maintain a probabilistic description of the whole population of targets. The recent information-theoretic and track-based control policy [14], although designed for the filtering solutions stemming from the novel estimation framework for stochastic populations [15], should be easily adaptable to the traditional track-based filters and could provide grounds for a generic sensor management solution.

Acknowledgements

This work was supported by the Engineering and Physical Sciences Research Council (EPSRC) Platform Grant (EP/J015180/1), the MOD University Defence Research Centre on Signal Processing (UDRC) Phase 2 (EP/K014227/1). Jeremie Houssineau is supported by Dstl task ED TIN 2-3.

References

[1] A. O. Hero III, "Sensor Management: Past, Present, and Future," *IEEE Sensor Journals*, vol. 11, no. 12, pp. 3064–3075, Dec. 2011.

[2] E. Delande, M. Uney, J. Houssineau, and D. E. Clark, "Regional Variance for Multi-Object Filtering," *IEEE Transactions on Signal Processing*, vol. 62, no. 13, pp. 3415 – 3428, 2014.

[3] E. Delande, J. Houssineau, and D. E. Clark, "Regional variance in target number: analysis and application for multi-Bernoulli point processes," in *Data Fusion & Target Tracking Conference (DF&TT 2014): Algorithms & Applications, 10th IET*, Apr. 2014.

[4] R. P. S. Mahler, *Statistical Multisource-Multitarget Information Fusion*. Artech House, 2007.

[5] E. Delande, J. Houssineau, and D. E. Clark, "PHD filtering with localised target number variance," in *Signal Processing, Sensor Fusion, and Target Recognition XXII, Proceedings of SPIE*, Apr. 2013, p. 87450E.

[6] ——, "Localised variance in target number for the Cardinalized Hypothesis Density Filter," in *Information Fusion, Proceedings of the 16th International Conference on*, Jul. 2013.

[7] R. P. S. Mahler, "Multitarget Bayes Filtering via First-Order Multitarget Moments," *Aerospace and Electronic Systems, IEEE Transactions on*, vol. 39, no. 4, pp. 1152–1178, Oct. 2003.

[8] ——, "PHD Filters of Higher Order in Target Number," *Aerospace and Electronic Systems, IEEE Transactions on*, vol. 43, no. 4, pp. 1523–1543, Oct. 2007.

[9] W. D. Blair and Y. Bar-Shalom, Eds., *Multitarget-Multisensor Tracking: Applications and Advances (Volume III)*. Artech House, 2000.

[10] B.-N. Vo and W.-K. Ma, "The Gaussian Mixture Probability Hypothesis Density Filter," *Signal Processing, IEEE Transactions on*, vol. 54, no. 11, pp. 4091–4104, Nov. 2006.

[11] D. J. Daley and D. Vere-Jones, *An Introduction to the Theory of Point Processes*, 2nd ed. Springer, 2003, vol. 1.

[12] J. Houssineau and D. E. Clark, "Hypothesised filter for independent stochastic populations," 2014, arXiv:1404.7408.

[13] E. Delande, J. Houssineau, and D. E. Clark, "A filter for distinguishable and independent populations," arXiv:1501.04671.

[14] ——, "Performance metric in closed-loop sensor management for stochastic populations," in *Sensor Signal Processing for Defence (SSPD)*, 2014.

[15] J. Houssineau, "Representation and estimation of stochastic populations," Ph.D. dissertation, Heriot Watt University, 2015.