THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Joint Registration and Fusion of an Infra-Red Camera and Scanning Radar in a Maritime Context

Citation for published version:

Cormack, D, Schlangen, I, Hopgood, J & Clark, D 2019, 'Joint Registration and Fusion of an Infra-Red Camera and Scanning Radar in a Maritime Context', *IEEE Transactions on Aerospace and Electronic Systems*. <https://doi.org/10.1109/TAES.2019.2929974>

Digital Object Identifier (DOI):

[10.1109/TAES.2019.2929974](https://doi.org/10.1109/TAES.2019.2929974)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

IEEE Transactions on Aerospace and Electronic Systems

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Joint Registration and Fusion of an Infra-Red Camera and Scanning Radar in a Maritime Context

David Cormack, Isabel Schlangen, *Member, IEEE*, James R. Hopgood, *Member, IEEE*
and Daniel E. Clark, *Senior Member, IEEE*

Abstract—The number of nodes in sensor networks is continually increasing, and maintaining accurate track estimates inside their common surveillance region is a critical necessity. Modern sensor platforms are likely to carry a range of different sensor modalities, all providing data at differing rates, and with varying degrees of uncertainty. These factors complicate the fusion problem as multiple observation models are required, along with a dynamic prediction model. However, the problem is exacerbated when sensors are not registered correctly with respect to each other, i.e. if they are subject to a static or dynamic bias. In this case, measurements from different sensors may correspond to the same target, but do not correlate with each other when in the same Frame of Reference (FoR), which decreases track accuracy. This paper presents a method to jointly estimate the state of multiple targets in a surveillance region, and to correctly register a radar and an Infrared Search and Track (IRST) system onto the same FoR to perform sensor fusion. Previous work using this type of parent-offspring process has been successful when calibrating a pair of cameras, but has never been attempted on a heterogeneous sensor network, nor in a maritime environment. This article presents results on both simulated scenarios and a segment of real data that show a significant increase in track quality in comparison to using incorrectly calibrated sensors or single-radar only.

Index Terms—Sensor fusion, registration, PHD filter, radar, infrared, calibration, tracking, maritime

I. INTRODUCTION

A. State-of-the-Art and Problem Outline

WITH many advances in sensor suites in recent years, tracking targets from multiple aspects using a range of different sensor modalities is now possible. Sensor fusion is a mature and sophisticated technology that looks to automate the process of combining a number of heterogeneous sources of information. By combining these sources of different information, the result should, in some sense, be better than what would have been possible if the data from each sensor was used individually [1], [2]. Maritime navigation radars typically

take a number of seconds to perform a full sweep of the region they are observing, whereas optical systems can potentially have much faster update rates, making the measurements asynchronous. By exploiting the high update rate and fusing these image measurements with the radar, target tracks can be updated and maintained more frequently [3].

The calibration of tracking systems [4] through estimation of their model parameters is an important prerequisite before such systems are deployed out in the field. In particular, when fusing multiple sensor measurements, it is necessary to consider sensor calibration or registration into the same Frame of Reference (FoR). If, in the real world, the set of sensors is incorrectly calibrated, any fusion in the multi-sensor multi-target tracking algorithm could possibly lead to total failure through loss of useful tracking information. Possible calibration or registration errors could come from a number of different sources. These could include incorrect calibration during sensor manufacture, incorrect alignment during installation and setup, and potentially even uncontrollable factors such as harsh weather, sensor drift (Global Positioning System (GPS) drift or Inertial Measurement Unit (IMU) inaccuracies), or misalignment through platform vibrations. Such drift might possibly be dynamic, resulting in the need of repeated, time-consuming re-calibration, or it could be impossible to calibrate the system to a global frame altogether, e.g. in GPS-denied environments. In such cases, it would be advantageous to calibrate sensors relative to each other in an automatic manner to avoid loss in track accuracy due to incorrect data fusion.

Previous work in the fusion of radar and optical imagery data has been shown to be successful when tracking manoeuvring targets [5] and also for the application of avian monitoring [6]. In these articles, the registration problem is either not accounted for, or is treated as a separate process before fusion occurs. Methods for solving the sensor registration problem have been shown in [7]–[10] which use pseudo-measurement approaches to estimate biases in sensor networks, and in [11] which uses deep learning to determine appropriate registration parameters. Machine learning-type methods require vast amounts of realistic training data in order to give reliable results, furthermore these articles only present results on datasets with biases around one order of magnitude smaller than those treated in this article.

The joint method presented below is flexible and allows for varying types of multiple target tracking algorithm to be used, such as Joint Probabilistic Data Association (JPDA) [12], [13], Multiple Hypothesis Tracking (MHT) [14], [15], and Belief Propagation (BP) and message passing [16], [17].

This work was supported by the Engineering and Physical Sciences Research Council (EPSRC) Grant number EP/S000631/1 and the MOD University Defence Research Collaboration (UDRC) in Signal Processing.

The work of D. Cormack is supported by Leonardo MW Ltd., Edinburgh, EH5 2XS, U.K.

D. Cormack is with the School of Engineering and Physical Sciences, Heriot-Watt University, Edinburgh, EH14 4AS, U.K. (email: drc9@hw.ac.uk)

I. Schlangen is with Fraunhofer FKIE, Wachtberg, Germany (email: isabel.schlangen@fkie.fraunhofer.de)

J.R. Hopgood is with the Institute for Digital Communications, University of Edinburgh, Edinburgh, EH9 3FG, U.K. (email: james.hopgood@ed.ac.uk)

D.E. Clark is with Department CITI, Telecom-SudParis, 9, rue Charles Fourier 91011, EVRY Cedex, France (e-mail: daniel.clark@telecom-sudparis.eu)

Fig. 1. An example of the sensor registration problem for a simple two-dimensional case, using an IRST and a radar that are co-located. Both sensors detect the target and give accurate measurements in their own FoR. However, when the measurements are projected into a common FoR, a relative angular bias ϕ is identified which needs to be accounted for during fusion.

For this work, the now commonly-used Probability Hypothesis Density (PHD) filter [18]–[20] is implemented. The PHD filter holds a number of advantages when dealing with multiple-target distributions, such as a low computational cost, and the ability to correctly estimate clutter and target populations even with large variance in the number of objects. The low computational cost is attractive in this defence application, as time-critical decisions need to be taken often. Other more expensive algorithms such as the Generalized Labelled Multi-Bernoulli (GLMB) filter [21] may not operate quickly enough.

B. Proposed Method and Contributions

The solution outlined in this article is motivated by a specific class of PHD filters that are based on hierarchical point process models, called *single-cluster filters* [22]. In these filters, the target population is regarded as a *single group with one or several common parameters* which are hidden and have to be estimated. In order to do so, the problem is modelled in two inter-dependent layers: The high-level process, also called the parent process, estimates the hidden parameter(s); and the low-level process, also called the offspring process, estimates the target states depending on those parameters.

Suitable models for sea clutter were identified to deal with highly fluctuating numbers of sea spikes [23], which motivated the usage of offspring processes with different false alarm models. Therefore, two types of filter will be used in this work. First, we consider the PHD filter [18], where the number of targets and false alarms are assumed to be Poisson distributed and only the first moment is propagated. Second, we consider a recent development in this field, the Panjer PHD filter [24] that can, by assuming the underlying distribution is Panjer, propagate both the mean and variance of the process. Both filters can be integrated in the single-cluster framework using a filter-dependent multi-object likelihood (MOL) [22], [25] which serves as a quality measure for the estimation provided by the parent process.

The single-cluster method has been successfully applied to Simultaneous Localisation and Mapping (SLAM) problems [26], sensor calibration [4], [27] and sensor drift estimation [28], [29], but to our knowledge, this is the first article that addresses the registration problem in a heterogeneous sensor network within a defence or surveillance context. More specifically, this work presents a method for estimating and tracking multiple targets from a maritime radar and an Infrared Search and Track (IRST) system like those in Fig. 3, while jointly registering the sensors onto the same FoR for fusion. Here, the FoR can be chosen arbitrarily and it will be shown that the relative bias is estimated correctly irrespective of that choice. One possibility is to assume that one sensor is perfectly calibrated onto the WGS84 coordinate system and a second, colocated sensor shows an angular displacement as shown in Fig. 1. In fact, by estimating the relative (angular or translational) displacement between sensors, fusion can be successfully performed even if the calibration onto a world FoR is impossible, e.g. in GPS-denied environments.

The contributions of this paper include:

- 1) a multi-target tracking (MTT) technique that incorporates sensor calibration in a joint manner, in contrast to existing techniques which solve tracking and calibration separately by using pseudo-measurements [7]–[10]; the proposed technique also avoids the computationally expensive data association problem found in some joint fusion and registration approaches, such as [30];
- 2) a new predictive model for tracking changes in the dynamic sensor configuration which is included in the parent process computations (Sec. II-A); it also incorporates non-uniform sampling information to overcome the asynchronous aspect of the sensor network (Sec. II-B);
- 3) non-linear observation models respecting the MTT output states being given in Cartesian coordinates and measurements in polar coordinates; therefore, implementations of two multi-object estimation algorithms are introduced in this paper to solve the non-linear fusion problem, including a novel Extended Kalman Filter (EKF) version of the Panjer filter (Sec. II-B);
- 4) a comprehensive set of simulations (Sec. III-D), involving a typically challenging tracking scenario where target trajectories cross one another; moreover, the simulations shown in this article consider a much larger angular offset, by an order of magnitude, compared to those simulated in [8], [31].

C. Paper Organization

The remainder of this paper is organized as follows: Sec. II gives an overview of the joint estimation problem, with modelling and implementation information given in Sec. III; results are shown in Sec. IV and conclusions are drawn in Sec. V.

II. MULTIPLE TARGET TRACKING AND FUSION

The main challenge in this work is that the offset angle between the radar and the IRST system, ϕ , is unknown and must be estimated recursively along with the target states based

Fig. 2. Flowchart of the joint registration and fusion process. Asynchronous measurements are sent to the buffer when ready, and then used in the offspring layer for tracking. When IRST measurements are received, the MOL $\hat{\ell}_k^i = \hat{\ell}_k(\phi_k^i | Z_k)$ is calculated to update the sensor registration parameter(s).

on imperfect data. For example, consider the scenario shown in Fig. 1 where a slow moving maritime target follows the trajectory shown in the dark thick line. Both sensors generate accurate measurements in their own FoR; the registration error is only apparent once the measurements are projected onto a common FoR such as WGS84. The radar produces range-azimuth measurements represented by the crosses which are assumed to be close to the true target trajectory.¹ However, the azimuth-only IRST measurements shown as the light lines radiating from the origin contain a systematic angular offset, ϕ , that is to be estimated in order to perform accurate sensor fusion. The following sections will provide detailed information on the work flow presented in Fig. 2.

A. Parent Process: Sensor Registration

The parent process is responsible for estimating the potentially time-varying sensor registration angle, ϕ , in order to project the measurements of one sensor into the FoR of the other sensor for fusion purposes. The densities involved in the parent process are denoted with $\hat{\cdot}$ for the rest of this article. A Bayes recursion is used to propagate the posterior density, $\hat{P}_k(\phi)$, for the calibration angle:

$$\hat{P}_{k|k-1}(\phi) = \int \hat{f}_{k|k-1}(\phi|\phi') \hat{P}_{k-1}(\phi') d\phi', \quad (1a)$$

$$\hat{P}_k(\phi|Z_k) = \frac{\hat{\ell}_k(\phi|Z_k) \hat{P}_{k|k-1}(\phi)}{\int \hat{\ell}_k(\phi'|Z_k) \hat{P}_{k|k-1}(\phi') d\phi'} \quad (1b)$$

where $\hat{\ell}_k(\phi|Z_k)$ is a multi-object likelihood (MOL) function, and $\hat{f}_{k|k-1}(\phi|\phi')$ is the transition density. Note that the

¹In Fig. 1, the world FoR in which the ground truth is assumed to be located coincides with the radar's FoR.

MOL, or registration parameter likelihood, is different from a multitarget association likelihood like the one used in the offspring process below as it describes the plausibility of a given sensor registration parameterisation ϕ based on a set of measurements Z_k ; more details can be found in [4], [26]. The multi-object/multi-measurement association likelihood, on the other hand, measures the association of the target states given the sensor measurements.

Using a particle filter approach, the registration parameter to be estimated at time-step k is represented with a particle distribution ϕ_k^i for $1 \leq i \leq N$, where each i represents a different sensor geometry. Each particle has a corresponding weight w_k^i and an underlying set of multi-target estimation statistics θ_k^i that are dependent on the offspring process. The particle weight w_k^i essentially encapsulates the belief that registration parameter ϕ_k^i best represents the true sensor calibration. With the particles fixed to a grid, the particle weights can be computed using grid-based methods [32, p. 9]. Using this method, the recursion for predicting and updating the parent process weights is

$$w_{k|k-1}^i = \sum_{j=1}^N w_{k-1}^j \hat{f}_{k|k-1} \left(\phi_k^i | \phi_k^j \right) \quad (2a)$$

$$w_k^i = \frac{w_{k|k-1}^i \hat{\ell}_k(\phi_k^i | Z_k)}{\sum_{j=1}^N w_{k|k-1}^j \hat{\ell}_k(\phi_k^j | Z_k)} \quad (2b)$$

where $\hat{f}_{k|k-1}(\phi_k^i | \phi_k^j) = \hat{f}_{k|k-1}^{i-j}$ is a discrete density over the difference in angle-indices and $\hat{\ell}_k(\phi_k^i | Z_k)$ is the MOL function evaluated for a given sensor geometry at ϕ_k^i . Equation (2a) can be written as the convolution of weights w_{k-1}^j and a kernel $\hat{f}_{k|k-1}^{i-j}$. The transition is modelled as a perturbation with a discretised wrapped Gaussian distribution. This is approximated using a finite-support shifted binomial distribution where $u \sim B(n, p)$ and $i = j + u - n/2$ is the relationship between the predicted angle-index, i , and the particle angle-index, j . Heuristically, the values $n = 6, p = 0.5$ are used, such that Equation (2a) can in practice be calculated as the convolution of w_{k-1} and the kernel $B(6, 0.5) = \{0.0156; 0.0938; 0.2344; 0.3125; 0.2344; 0.0938; 0.0156\}$. The parent process is initialized with a flat prior distribution, where all sensor geometries are equally likely.

The information contained inside θ_k^i at time-step k is dependent on the type of offspring filter that is chosen, i.e. it either just contains the intensity, μ_k , of the multi-target distribution in case of the PHD filter, or both μ_k and the variance var_k in case of the Panjer filter. The filters and their respective MOL functions are described below.

B. Offspring Process: Sensor Fusion

Let $\mathcal{X}$ and $\mathcal{Z}$ denote the state and measurement spaces, respectively. The offspring process estimates the time-varying multi-object state with n_k targets, $\psi \in \mathcal{X}^{n_k}$, dependent on a certain sensor configuration ϕ . This process is assumed to evolve with a Markov transition function $f_{k|k-1}(\psi|\psi')$ which in this case will follow the near constant velocity (NCV) motion model [33], with a dynamic value for state transition

Fig. 3. Sensor setup for collecting real data with the IRST system in the foreground and radar in the background. © Crown copyright, 2019.

time Δ_k . This dynamic value is calculated at the measurement buffer stage as shown in Fig. 2, using the equation $\Delta_k = t_k - t_{k-1}$ where t_k is the known current measurement time, and t_{k-1} is the previous measurement time. The multi-measurement/multi-target likelihood $l_k(\psi|\phi, Z_k)$ describes the association likelihood of targets and measurements. The following Bayes recursion is used to propagate the law P_k of the target process at time k :

$$P_{k|k-1}(\psi|\phi) = \int f_{k|k-1}(\psi|\psi') P_{k-1}(\psi'|\phi) d\psi' \quad (3a)$$

$$P_k(\psi|\phi, Z_k) = \frac{l_k(\psi|\phi, Z_k) P_{k|k-1}(\psi|\phi)}{\int l_k(\psi'|\phi, Z_k) P_{k|k-1}(\psi'|\phi) d\psi'} \quad (3b)$$

For this work, the offspring process will be fulfilled using either the original PHD or the Panjer PHD filter as described below, and concrete model choices are described in Sec. III-B. When a measurement arrives from the reference sensor, θ_k^i are updated for all $i = 1, \dots, N$, and when a measurement arrives from an uncalibrated sensor, both θ_k^i and the registration parameter weights w_k^i in (2b) are updated using the MOL in Eqns (9) or (13), respectively.

1) *PHD Filter*: The PHD filter is now a mature method of performing multi-target tracking. It was first developed in 2003 [18], [19] and the common Gaussian mixture implementation was introduced in 2006 [20]. A full derivation of the filter can be found in [18], [20]. This first-order filter propagates only the mean μ_k of the point process and assumes the predicted number of targets and false alarms are both Poisson distributed [25]. The implementation used in this work can be found in [20, Table IV].

2) *Panjer PHD Filter*: The Panjer PHD filter was first introduced in [24] as a useful extension to the first-order PHD filter defined earlier. This process propagates both the mean and variance of the point process, using the assumption that the number of predicted targets and the false alarms are Panjer distributed. The Panjer distribution is characterised by two parameters α and β which closely correspond to its mean and variance.

Case 1: $0 < \{\alpha, \beta\} \in \mathbb{R}^+ \times \mathbb{R}^+$: Represents a negative binomial distribution, where the variance is greater than the mean. This would be useful, for example, in situations where

there may be sudden large influxes of measurements such as strong returns on a rough sea.

Case 2: $0 > \{\alpha, \beta\} \in \mathbb{Z}^- \times \mathbb{R}^-$: Represents a binomial distribution where the variance is less than the mean. This could be used in a situation where a very consistent number of false alarms is expected, e.g. a static or slow-changing environment.

Case 3: $\{\alpha, \beta\} \rightarrow \infty$: The limit case where the ratio stays constant, resulting in the Poisson distribution. This represents the standard false alarm model used in many pieces of target tracking literature.

Having these three different cases available gives more flexibility in modelling the number of targets and false alarms. A full mathematical derivation and pseudo-code is available elsewhere [24], and is omitted here.

3) *Extended Kalman Filter PHD and Panjer filters*: An EKF version of both offspring filters is required to overcome the non-linearity between the Cartesian state space and the Polar observation space. EKFs use a Jacobian matrix to linearize the non-linear function around the current state estimate. This Jacobian is found by performing partial differentiation of the observation model equations with respect to the variables in the state vector, such that the elements of the matrix are

$$J_{pq} = \frac{\partial h_p}{\partial X_q}, \quad p \in \{r, \phi\}, q \in \{1, \dots, 4\} \quad (4)$$

where h_p is the conventional Cartesian to Polar transformation

$$h_r = r = \sqrt{x^2 + y^2}, \quad h_\phi = \phi = \tan^{-1}(x, y),$$

and where $\tan^{-1}(x, y)$ is the four-quadrant inverse tangent function and $X = [x, \dot{x}, y, \dot{y}]^T$. During an offspring update with radar measurements, the Jacobian matrix used is

$$J_R = \begin{bmatrix} \frac{x}{\sqrt{x^2 + y^2}} & 0 & \frac{y}{\sqrt{x^2 + y^2}} & 0 \\ \frac{-y}{x^2 + y^2} & 0 & \frac{x}{x^2 + y^2} & 0 \end{bmatrix} \quad (5)$$

and for an IRST update, the Jacobian becomes

$$J_I = \begin{bmatrix} \frac{-y}{x^2 + y^2} & 0 & \frac{x}{x^2 + y^2} & 0 \end{bmatrix}. \quad (6)$$

C. The multi-object likelihood functions

In this section, the PHD and the Panjer version of the MOL are revisited, and a short version of their derivations is given in a supplementary document. More detailed explanations are found in [22] and [34]. Note that we will omit an implicit time index k in all notations for the sake of brevity.

1) *Notations*: In the following, all equations referring to the PHD filter or Poisson distributed phenomena are marked with the symbol $\bullet$ and any reference to the Panjer filter or the Panjer distribution is indicated with the symbol $\circ$. Dependent on the sensor configuration ϕ , we can write down the spatial distribution of the false alarms as $s_c(\cdot|\phi)$. Furthermore, suppose that $\mu_{pr}^{\bullet/\circ}(\cdot|\phi)$ is the predicted intensity of the PHD or Panjer filter, $l(x|z)$ denotes the single-target single-measurement association likelihood, and $p_d(\cdot|\phi)$ is the (state-dependent) probability of target detection. For arbitrary regions $B \subseteq \mathcal{X}$, the association terms $\mu_z^{\bullet/\circ}(\cdot|\phi)$ are given by

$$\mu_z^{\bullet/\circ}(B|\phi) = \int_B p_d(x) l(x|z) \mu_{pr}^{\bullet/\circ}(x) dx. \quad (7)$$

2) *The PHD filter likelihood* [22]: For the original PHD filter, the number of false alarms is assumed to be Poisson distributed, in which case $\lambda_c^\bullet$ shall denote the Poisson false alarm rate and $\mu_c^\bullet(\cdot|\phi) = \lambda_c^\bullet s_c^\bullet(\cdot|\phi)$ is its intensity function (cf. Eq. (28) in the supplementary document).

Theorem II.1 (MOL of the PHD filter [22]). *Given the short-hand notation*

$$\mu_d^\bullet(\mathcal{X}) = \int_{\mathcal{X}} p_d(x|\phi) \mu_{pr}^\bullet(x|\phi) dx, \quad (8)$$

the likelihood function of the PHD filter for a given sensor state ϕ is found to be

$$\hat{\ell}_\bullet(\phi|Z) = \frac{\prod_{z \in Z} [\mu_c^\bullet(z|\phi) + \mu_d^\bullet(\mathcal{X}|\phi)]}{\exp \left[\int_{\mathcal{Z}} \mu_c^\bullet(z|\phi) dz + \mu_d^\bullet(\mathcal{X}) \right]}. \quad (9)$$

3) *The Panjer filter likelihood* [25], [34]: Let us define the Pochhammer symbol or rising factorial $x_{n\uparrow}$ by

$$x_{n\uparrow} := \prod_{i=0}^{n-1} (x+i), \quad x_{0\uparrow} := 1. \quad (10)$$

As the name suggests, the Panjer filter assumes Panjer distributed false alarms, in which case we write α_c, β_c for the two Panjer clutter parameters and s_c° for the spatial distribution of the false alarms.

Theorem II.2 (MOL of the Panjer PHD filter). *Write $\alpha = \alpha_{pr}$, $\beta = \beta_{pr}$ and s_{pr}° for the Panjer parameters and the spatial distribution of the predicted process, and let*

$$F_{d,\phi} = 1 + \frac{1}{\beta} \int_{\mathcal{X}} p_d(x|\phi) s_{pr}^\circ(x|\phi) dx, \quad (11)$$

$$F_c = 1 + \frac{1}{\beta_c} \quad (12)$$

be two auxiliary functions for a given sensor state ϕ . The multi-object likelihood function of the Panjer PHD filter for ϕ is found to be

$$\begin{aligned} \hat{\ell}_\circ(\phi|Z) = & \sum_{j=0}^{|Z|} \frac{\alpha_{j\uparrow}}{\beta^j} \frac{(\alpha_c)_{(|Z|-j)\uparrow}}{(\beta_c + 1)^{|Z|-j}} F_{d,\phi}^{-\alpha-j} F_c^{-\alpha_c - |Z|-j} \\ & \cdot \sum_{\substack{Z' \subseteq Z \\ |Z'|=j}} \prod_{z \in Z'} \mu_z^\circ(\mathcal{X}|\phi) \prod_{z' \in Z \setminus Z'} s_c^\circ(z|\phi). \end{aligned} \quad (13)$$

III. MODELLING, DATA, AND SCENARIOS

A. Implementation

Much like the parent process in [35], the parent process in this application is represented with a one-dimensional even spread of particles. In this case, particles are distributed on a fixed grid between angles $\pm 10^\circ$ from the centre of the field of view. The number of particles N has been chosen to be 201 such that the angular resolution is 0.1° between consecutive filters. With this even spread of particles, and a consistent test between $\pm 10^\circ$, there is no need for a particle resampling step [36] in this algorithm, reducing the computational effort required. In this work we only consider estimating one parameter, however several parameters could be estimated simultaneously if necessary at the cost of using

Algorithm 1 Joint Sensor Registration and Fusion

Input: Set of particles $\{\phi_{k-1}^i, w_{k-1}^i, \theta_{k-1}^i\}_{i=1}^N$
Set of measurements Z_k

procedure PREDICTION

for $1 \leq i \leq N$ **do**

$w_{k|k-1}^i = \text{ParentPrediction}(w_{k-1}^i)$ $\triangleright$ Eq. (2a)

$\theta_{k|k-1}^i = \text{OffspringPrediction}(\theta_{k-1}^i)$ $\triangleright$ Eq. (3a)

end for

end procedure

procedure UPDATE

for $1 \leq i \leq N$ **do**

if Z_k from reference sensor **then**

$\theta_k^i = \text{OffspringUpdate}(\theta_{k|k-1}^i, Z_k)$ $\triangleright$ Eq. (3b)

$w_k^i = w_{k|k-1}^i$

else if Z_k from uncalibrated sensor **then**

$\theta_k^i = \text{OffspringUpdate}(\theta_{k|k-1}^i, Z_k)$ $\triangleright$ Eq. (3b)

$w_k^i = \text{ParentUpdate}(\theta_{k|k-1}^i, w_{k|k-1}^i)$ $\triangleright$ Eq. (2b)

end if

end for

end procedure

Output: Set of particles $\{\phi_k^i, w_k^i, \theta_k^i\}_{i=1}^N$

more particles. A high-level pseudocode for this solution is given in Algorithm 1.

Both offspring processes are implemented using a Gaussian-Mixture representation of the target population like in [20], [24], but using an EKF as described earlier. A measurement-driven birth process will be used as described in [37], and the Hellinger distance is used for component merging [38]; components with low weights below a threshold τ_{prune} are removed at every update.

B. Model Definitions

1) *Dynamical Model:* The measurement buffer will have access to the raw range-bearing radar measurements and the bearing-only IRST measurements that are recorded at a given iteration k and physical time t_k . The MTT routine will be performed using a 4-D Cartesian state vector with elements

$$\mathbf{x}_k = [x_k \ \dot{x}_k \ y_k \ \dot{y}_k]^\top \quad (14)$$

where x_k, y_k are the x and y positions of a target, and $\dot{x}_k, \dot{y}_k$ are the x and y velocities of a target. For this maritime surveillance-based scenario, it is assumed that each and every target follows a near-constant velocity (NCV) model [39], [40] that is described by

$$\mathbf{x}_k = \mathbf{F}_k \mathbf{x}_{k-1} + \mathbf{w}_k \quad (15)$$

where $\mathbf{F}_k$ is the state transition matrix

$$\mathbf{F}_k = \begin{bmatrix} 1 & \Delta_k & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & \Delta_k \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad \Delta_k = t_k - t_{k-1}, \quad (16)$$

and $\mathbf{w}_k$ represents zero-mean white Gaussian process noise with covariance

$$\mathbf{Q}_k = \begin{bmatrix} q\Delta_k^3/3 & q\Delta_k^2/2 & 0 & 0 \\ q\Delta_k^2/2 & q\Delta_k & 0 & 0 \\ 0 & 0 & q\Delta_k^3/3 & q\Delta_k^2/2 \\ 0 & 0 & q\Delta_k^2/2 & q\Delta_k \end{bmatrix} \quad (17)$$

and q is the acceleration noise value in both the X and Y directions.

2) *Measurement Models*: The radar measurement model is defined as

$$\mathbf{z}_k^R = h^R(\mathbf{x}_k) + \eta_k^R, \quad (18)$$

with

$$h^R(\mathbf{x}_k) = \begin{bmatrix} r_k \\ \phi_k \end{bmatrix} = \begin{bmatrix} \sqrt{x_k^2 + y_k^2} \\ \tan^{-1}(x_k, y_k) \end{bmatrix}, \quad (19)$$

where $r_k > 0$, $\tan^{-1}(x_k, y_k)$ is the four-quadrant inverse tangent function, and the resulting ϕ_k lies within $[0, 2\pi)$. The additive noise term η_k^R is defined by

$$\eta_k^R \sim \mathcal{N}(\eta_k^R; \mathbf{0}, \text{diag}(\sigma_{r_r}^2, \sigma_{\phi_r}^2)) \quad (20)$$

where σ_{r_r} and σ_{ϕ_r} are the radar's range and azimuth standard deviations respectively.

The camera measurement model is described by

$$\mathbf{z}_k^C = h^C(\mathbf{x}_k) + \eta_k^C, \quad (21)$$

where

$$h^C(\mathbf{x}_k) = \phi_k = \tan^{-1}(x_k, y_k). \quad (22)$$

The additive noise term η_k^C is defined by

$$\eta_k^C \sim \mathcal{N}(\eta_k^C; \mathbf{0}, \sigma_{\phi_c}^2) \quad (23)$$

where σ_{ϕ_c} is the camera azimuth standard deviation.

C. Sensors

The maritime navigation radar used as a part of the data collection was a Kelvin Hughes SharpEye system [41], shown in the background of Fig. 3. It uses a number of different radar techniques to detect a wide-range of targets in a maritime environment. During the trial, the system was operating in a full 360° sweep mode, with a full scan taking approximately 2.5 seconds to complete. The low-profile antenna gives an azimuthal beam-width of less than 1° at the 3 dB point. Techniques such as monopulse to further improve the angular accuracy of the radar were not in use; accurate target localisation inside the beam was not possible.

The IRST system, shown in the foreground of Fig. 3, was a research platform which limited the amount of information available. IRST is one method for detecting and tracking targets that give off infra-red signatures. The wavelength of this sensor tends to be shorter than that of a conventional maritime radar, thus giving much better angular resolution. One drawback however is that it can be affected by atmospheric conditions and the weather, both shortening its effective range. The reasons for performing sensor fusion in this case are to exploit the better angular accuracy in the IRST measurements and the high update rate available from the IRST sensor.

Fig. 4. Trajectories used in simulated scenario, with sensors co-located at the origin (0,0).

D. Simulations

In order to test the algorithm using simulated data, a challenging scenario involving crossing target trajectories has been created. Two targets move around inside the IRST field-of-view and their trajectories cross around 60 iterations into the scenario. Crossing targets often make the data association problem difficult for tracking algorithms; a sensor with a slow update rate will not generate estimates quickly enough, and therefore there will be a loss in track resolution. Sensor properties and update rates are close to those defined in the accompanying data sheets and manuals [41].

In this first scenario, we consider two potential false alarm models for a maritime tracking scenario; the common Poisson model, and also a negative binomial model [42]. The latter will allow for a large variance in the number of false alarms generated from crests of waves or from rough sea conditions [23]. Simulations have been created to test the filters on each of those models. The parameters used for the simulations and the filters are shown in Table I and are typical for multiple target tracking simulations using PHD filters [20], [24]; any variation in parameters for different experiments will be stated later. For the first two experiments (Sec. IV-A and IV-B) where the IRST is calibrated onto the radar FoR, a misalignment of 3° has been simulated in the IRST measurements. Due to the different characteristics and sampling rates of radar and an infrared sensor, a third simulated scenario is considered (Sec. IV-C) where the radar measurements contain an angular bias and the IRST measurements have no bias. This alternative scenario will show that it is possible to still estimate this angular bias, even when exploiting the infrequent radar measurements, rather than the high-frequency IRST measurements.

E. Real Scenario

In order to further test the algorithm developed, a short segment of data from a real scenario was used. Both sensors were located on a sea wall at Fort Blockhouse in Gosport overlooking Portsmouth harbour. The trial coordinators had instructed a number of instrumented targets to be present during the trial. However, there was a large amount of background

TABLE I
TRACKING PARAMETERS

Quantity	Symbol	Sim Value	Real Value
Detection Probability	p_d	0.99	0.6
Survival Probability	p_s	0.95	0.95
Pruning Threshold	τ_{prune}	0.001	0.001
Merging Threshold	τ_{merge}	0.8	0.6
Extraction Threshold	τ_{extract}	0.5	0.5
False Alarm Rates	λ_r, λ_c	2, 5	5, 10
False Alarm Variance	$\text{var}_r, \text{var}_c$	10, 50	N/A
Birth Intensity	μ_b	1	0.01
Acceleration Noise (ms^{-2})	q	1	3
Radar Meas. Noise (m, deg)	$\sigma_{r_r}, \sigma_{\phi_r}$	5, 0.06	3.873, 0.059
IRST Meas. Noise (deg)	σ_{ϕ_c}	0.01	0.016

traffic such as ferries and cargo vessels passing through, which is actually an advantage: having more targets present in the scene is preferable when attempting to calibrate using this method as more measurement-to-track associations can be made, and therefore increase the MOL.

In the segment of real data used, a target crosses the field of view at approximately 2 kilometres away from the sensors. The segment lasts for approximately 50 seconds and then the target disappears from view. For the sensor fusion aspect, a subset of the IRST measurement sets are used to improve the update rate, but not so much as to heavily rely on bearings-only track updates; a typically challenging problem in MTT. Approximately three sets of IRST measurements are used between consecutive radar scans.

IV. RESULTS

All results have been averaged over 50 Monte Carlo runs, and the registration angle estimate is taken as the Maximum A Posteriori (MAP) estimate of the parent likelihood function. The tracked output from the PHD filter and Panjer PHD filter will be compared to the simulated trajectories using the Optimal Subpattern Assignment (OSPA) distance [43]. The OSPA distance is a combination of a cardinality error and a localisation error between two sets X and Y with cardinalities² m and n , which is widely used to determine accuracy in multi-target tracking systems. It is given by [43, Eq. (3)]

$$d_p^{(c)}(X, Y) = \left[\min_{\pi \in \Pi_n} \sum_{i=1}^m d^{(c)}(x_i, y_{\pi(i)})^p + c^p(n - m) \right]^{\frac{1}{p}}, \quad (24)$$

using an order parameter p and a cut-off distance c . Here, the distance function $d^{(c)}(x, y) = \min(c, d(x, y))$ is an appropriate distance measure, e.g. the Euclidean distance, cut off at c , and Π_n denotes the set of all possible permutations of the numbers $1, \dots, n$. As the order parameter increases, the metric penalises estimates that lie further away from the ground truth more harshly. From [43, Sec. III-D], $p = 2$ is a good practical choice for the order parameter, as it usually gives smoother distance curves, and is consistent with other

²Without loss of generality, it is assumed that X has at most as many elements as Y .

Fig. 5. OSPA Distance over Time, PHD Filter, Poisson Distribution

metrics that use a p -th order average construction. The cut-off distance c determines the trade-off between penalising cardinality errors as opposed to localisation errors. For all OSPA results shown, a cut-off parameter of $c = 100$ m and order parameter $p = 2$ will be used.

In Figs. 5 through to 9, where the OSPA distance and IRST pointing angle are plotted, these are the averages taken across all Monte-Carlo trials. We use a consistent target ground truth in all trials and use this to generate a different set of sensor measurements for each individual trial. The legend shown in Fig. 5a is consistent across Figs. 5, 6, 7, and 10a. The red dashed plots represent the case where only the radar is used and no data fusion is performed. The blue dotted plots show the case where the radar and IRST are incorrectly registered and no correction attempt is made, whereas the dashed-dotted black plots show the results for the joint estimation method presented in this work. Finally, as a benchmark to compare our result to, we have also simulated the case where the radar and IRST are perfectly registered, giving an optimal result. This is shown as the solid green plot on each figure.

A. Poisson Distribution

We start by considering the Poisson distribution as the underlying false alarm model. Here, the mean number of false alarms is equal to the variance, giving a reasonably consistent number per scan. The average false alarm rates for each sensor are shown in Table I. From Figs. 5 and 6, the importance of having the correct registration between sensors is clear. Performing tracking with only radar measurements

(a) $p_d = 0.99$ (b) $p_d = 0.85$

Fig. 6. OSPA Distance over Time, Panjer Filter, Poisson Distribution

(a) PHD Filter, $p_d = 0.85$ (b) Panjer Filter, $p_d = 0.85$

Fig. 7. OSPA Distance over Time, Negative Binomial Distribution

is accurate in itself, however the tracking accuracy can be improved by performing sensor fusion between a radar and a calibrated IRST sensor. Without taking the registration into account, and performing fusion with an IRST sensor that is uncalibrated, tracking performance is decreased and the OSPA distance increases significantly. Comparing Fig. 5a to Fig. 6a, and Fig. 5b to Fig. 6b there is an improvement in tracking performance by using the Panjer PHD filter. In the $p_d = 0.99$ case, the $k = 95$ OSPA distances for the PHD filter and Panjer PHD filters with estimated registration are 21.72 m and 20.04 m respectively, and in the $p_d = 0.85$ case, these are 32.56 m and 30.72 m. This is due to propagation of the variance of the cardinality distribution, as well as the mean.

As with many tracking scenarios, the probability of detection p_d is an important factor and as shown in Figs. 8a and 9a, performance improves in both the target tracking and the registration estimation as p_d increases from 0.7 to 0.99.

B. Negative Binomial Distribution

For this set of simulations, the variance in the number of false alarms is no longer constrained to be equal to the mean. The included variance parameters are shown in Table I. The PHD filter assumes the false alarm distribution is Poisson, and does not take account of the variance information. The Panjer PHD filter however allows for variance information to be included. It can be seen in Fig. 7 that the Panjer PHD filter does outperform the PHD filter in all cases. At $k = 95$ in the estimated registration case, the OSPA distance for

the PHD filter is 35.71 m, whereas for the Panjer PHD filter, this is 30.61 m. Again these results show it is important to take account of the registration between the sensors. A larger performance gain can be seen in Fig. 9. We see that even with lower values of p_d (e.g. 0.7), the Panjer PHD filter is able to give more consistent and accurate estimates of the IRST pointing angle. We require a higher p_d value (> 0.9) to get accurate registration estimates from the PHD filter.

C. Alternative Frame of Reference

We now consider the alternative representation of the sensors, where instead, the radar measurements are assumed to contain some angular bias and the IRST measurements contain no bias. Firstly in Fig. 10a, it can be seen that the tracking accuracy using a single radar has now decreased as expected, and the result where fusion is performed between the unregistered sensors is still poor. The estimated result is close to the unbiased and correctly registered result, with an OSPA distance at $k = 95$ of 27.04 m. In Figs. 10b and 10c, as the probability of detection p_d increases, there is a decrease in the OSPA distance, and the angle estimate tends towards the true registration configuration, following the same trends as the results given in Figs. 8a and 9a.

D. Real Scenario

Fig. 11 shows that after the initialisation to 10° and a transition period, a 1.5° to 2° registration error is estimated

(a) Poisson Distributed

(a) Radar bias, $p_d = 0.99$, PHD filter

(b) Negative Binomial Distributed

(b) Average OSPA Distance

Fig. 8. OSPA Distance over p_d . The average OSPA Distance is taken from between Iteration 80 and Iteration 100.

(a) Poisson Distributed

(b) Negative Binomial Distributed

Fig. 9. IRST Pointing Angle over p_d . The average pointing angle is taken from between Iteration 80 and Iteration 100.

(c) Estimated Alignment Angle w.r.t. IRST frame

Fig. 10. Results for alternative FoR scenario.

Fig. 11. Estimated correction angle for real scenario

Fig. 12. Output track from PHD filter for real scenario.

between the radar and the IRST system. After the transition, it can be seen that the angle only deviates from this a small amount, which could be due to some platform vibration, or gusts of wind for example. This error could account for the fact that the installation was done only by eye and therefore subject to human error. The advantage of performing sensor fusion is apparent in the comparison in Fig. 12. The radar track shown in Fig. 12a is accurate, however there are considerable gaps between consecutive estimates. These gaps are filled using fusion of the IRST information, see Fig. 12b. With range information only available in the radar measurements, we see the two larger "jumps" in the track at $(-2000, 400)$ and $(-1850, 450)$. These are points where radar measurements become available again and range information can be updated.

V. CONCLUSIONS

We have successfully demonstrated a suitable method for performing joint sensor registration and fusion using asynchronous and heterogeneous sensors. The registration parameter is estimated on the fly, based on the performance of the underlying target tracker, which has never been performed in this context before. The simulation results on a challenging scenario clearly highlight the importance of calibrating the network before performing fusion across the sensors, even when using an alternative false alarm model. When applying this method to a real data set, a plausible offset angle was found. For this sensor setup of radar and IRST, the well-documented bearings-only tracking problem was evident when using IRST measurements. By not having a range measurement available, tracking accuracy is reduced. This was compensated for by using only a subset of the IRST images available, giving more weight to the sparse range-bearing measurements of the radar.

Taking into consideration the simulated scenario outlined in this work, it is of vital importance that sensor fusion is used. Using only the slow update rate of the radar system will leave longer periods where no track estimates are given, and therefore track resolution is lost. Exploiting the higher update rate of the IRST system majorly improves the performance where the calibration is accurate, especially at the point where the targets cross. Having more track estimates gives better situational awareness, and helps to resolve potential issues with estimate-to-track association.

REFERENCES

- [1] W. Koch, *Tracking and Sensor Data Fusion: Methodological Framework and Selected Applications*, 1st ed. Berlin: Springer Verlag, 2014.
- [2] W. Elmenreich, "An Introduction to Sensor Fusion," Vienna Institute of Technology, Austria, Vienna, Austria, Tech. Rep., 2002.
- [3] Y. Bar-Shalom, *Multitarget-Multisensor Tracking: Applications and Advances Volume III*, 1st ed., W. D. Blair, Ed. Norwood, MA: Artech House, 2000.
- [4] B. Ristic, D. Clark, and N. Gordon, "Calibration of Multi-Target Tracking Algorithms Using Non-Cooperative Targets," *IEEE Journal of Selected Topics in Signal Processing*, vol. 7, no. 3, pp. 390–398, 2013.
- [5] C. Huazhi and R. Jian, "A Multitarget Tracking Algorithm based on Radar and Infrared Sensor Data Fusion," in *2011 IEEE 3rd International Conference on Communication Software and Networks (ICCSN)*. Xi'an, China: IEEE, 2011, pp. 367–371.
- [6] G. Mirzaci, M. M. Jamali, J. Ross, P. V. Gorsevski, and V. P. Bingman, "Data Fusion of Acoustics, Infrared and Marine Radar for Avian Study," *IEEE Sensors Journal*, vol. 15, no. 11, pp. 6625–6632, 2015.
- [7] E. Taghavi, R. Tharmarasa, T. Kirubarajan, and M. McDonald, "Multisensor-multitarget bearing-only sensor registration," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 52, no. 4, pp. 1654–1666, 2016.
- [8] E. Taghavi, R. Tharmarasa, T. Kirubarajan, Y. Bar-Shalom, and M. McDonald, "A Practical Bias Estimation Algorithm for Multisensor-Multitarget Tracking," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 52, no. 1, pp. 2–19, 2016.
- [9] X. Lin, Y. Bar-Shalom, and T. Kirubarajan, "Multisensor Multitarget Bias Estimation for General Asynchronous Sensors," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 41, no. 3, pp. 899–921, 2005.
- [10] D. Huang, H. Leung, and E. Bosse, "A Pseudo-Measurement Approach to Simultaneous Registration and Track Fusion," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 48, no. 3, pp. 2315–2331, 2012.
- [11] N. Schneider, F. Piewak, C. Stiller, and U. Franke, "RegNet: Multimodal sensor registration using deep neural networks," in *2017 IEEE Intelligent Vehicles Symposium*. Los Angeles, CA: IEEE, 2017, pp. 1803–1810.
- [12] T. Fortmann, Y. Bar-Shalom, and M. Scheffe, "Multi-Target Tracking using Joint Probabilistic Data Association," in *19th IEEE Conference on Decision and Control including the Symposium on Adaptive Processes*. Albuquerque, NM, USA: IEEE, 1980, pp. 807–812.
- [13] —, "Sonar Tracking of Multiple Targets using Joint Probabilistic Data Association," *IEEE Journal of Oceanic Engineering*, vol. 8, no. 3, pp. 173–184, 1983.
- [14] D. B. Reid, "An Algorithm for Tracking Multiple Targets," *IEEE Transactions on Automatic Control*, vol. 24, no. 6, pp. 843–854, 1979.
- [15] B.-N. Vo, M. Mallick, Y. Bar-Shalom, S. Coraluppi, R. Osborne, R. Mahler, and B.-T. Vo, "Multitarget Tracking," in *Wiley Encyclopedia of Electrical and Electronics Engineering*. American Cancer Society, 2015, pp. 1–15.
- [16] F. Meyer, P. Braca, P. Willett, and F. Hlawatsch, "A Scalable Algorithm for Tracking an Unknown Number of Targets Using Multiple Sensors," *IEEE Transactions on Signal Processing*, vol. 65, no. 13, pp. 3478–3493, 2017.
- [17] F. Meyer, T. Kropfreiter, J. L. Williams, R. A. Lau, F. Hlawatsch, P. Braca, and M. Z. Win, "Message Passing Algorithms for Scalable Multitarget Tracking," *Proceedings of the IEEE*, vol. 106, no. 2, pp. 221–259, 2018.
- [18] R. Mahler, "Multitarget Bayes filtering via first-order multitarget moments," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 39, no. 4, pp. 1152–1178, 2003.
- [19] —, *Statistical Multisource-Multitarget Information Fusion*. Norwood, MA: Artech House, 2007.

- [20] B.-N. Vo and W.-K. Ma, "The Gaussian Mixture Probability Hypothesis Density Filter," *IEEE Transactions on Signal Processing*, vol. 54, no. 11, pp. 4091–4104, 2006.
- [21] B.-N. Vo, B.-T. Vo, and H. G. Hoang, "An Efficient Implementation of the Generalized Labeled Multi-Bernoulli Filter," *IEEE Transactions on Signal Processing*, vol. 65, no. 8, pp. 1975 – 1987, 2017.
- [22] A. Swain, "Group and extended target tracking with the probability hypothesis density filter," Ph.D. dissertation, Heriot-Watt University, 2013.
- [23] K. D. Ward, R. J. Tough, and S. Watts, *Sea Clutter: Scattering, the K Distribution and Radar Performance*, 1st ed. London, U.K.: The Institution of Engineering and Technology, 2013.
- [24] I. Schlangen, E. D. Delande, J. Houssineau, and D. E. Clark, "A Second-Order PHD Filter With Mean and Variance in Target Number," *IEEE Transactions on Signal Processing*, vol. 66, no. 1, pp. 48–63, 1 2018.
- [25] I. Schlangen, D. E. Clark, and E. D. Delande, "Single-cluster PHD filter methods for joint multi-object filtering and parameter estimation," *ArXiv e-prints*, *ArXiv:1705.05312*, 5 2017. [Online]. Available: <http://arxiv.org/abs/1705.05312>
- [26] C. S. Lee, D. Clark, and J. Salvi, "SLAM with Dynamic Targets via Single-Cluster PHD Filtering," *IEEE Journal of Selected Topics in Signal Processing*, vol. 7, no. 3, pp. 543–552, 2013.
- [27] J. Houssineau, D. Clark, S. Ivekovic, C. S. Lee, and J. Franco, "A Unified Approach for Multi-Object Triangulation, Tracking and Camera Calibration," *IEEE Transactions on Signal Processing*, vol. 64, no. 11, pp. 2934–2948, 2016.
- [28] I. Schlangen, J. Franco, J. Houssineau, W. Pitkeathly, D. Clark, I. Smal, and C. Rickman, "Marker-less Stage Drift Correction in Super-Resolution Microscopy Using the Single-Cluster PHD Filter," *IEEE Journal of Selected Topics in Signal Processing*, vol. 10, no. 1, pp. 193–202, 2016.
- [29] O. Hagen, J. Houssineau, I. Schlangen, E. Delande, J. Franco, and D. E. Clark, "Joint Estimation of Telescope Drift and Space Object Tracking," in *2016 IEEE Aerospace Conference*. Big Sky, MT: IEEE, 2016, pp. 1–10.
- [30] Z. Li, S. Chen, H. Leung, and E. Bosse, "Joint Data Association, Registration, and Fusion using EM-KF," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 46, no. 2, pp. 496 – 507, 2010.
- [31] X. Lin, Y. Bar-Shalom, and T. Kirubarajan, "Exact Multisensor Dynamic Bias Estimation with Local Tracks," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 40, no. 2, pp. 576 – 590, 2004.
- [32] B. Ristic, S. Arulampalam, and N. Gordon, *Beyond the Kalman Filter: Particle Filters for Tracking Applications*, 1st ed. Boston, MA: Artech House, 2004.
- [33] A. Doucet, N. de Freitas, and N. Gordon, *Sequential Monte Carlo Methods in Practice*. New York, NY: Springer-Verlag, 2001.
- [34] I. Schlangen, "Multi-object filtering with second-order moment statistics," Ph.D. dissertation, Heriot-Watt University, 2017.
- [35] I. Schlangen, V. Bharti, E. Delande, and D. E. Clark, "Joint multi-object and clutter rate estimation with the single-cluster PHD filter," in *IEEE 14th International Symposium on Biomedical Imaging (ISBI 2017)*. Melbourne, Australia: IEEE, 2017, pp. 1087 – 1091.
- [36] T. Li, M. Bolic, and P. M. Djuric, "Resampling Methods for Particle Filtering: Classification, Implementation and Strategies," *IEEE Signal Processing Magazine*, vol. 32, no. 3, pp. 70–86, 2015.
- [37] J. Houssineau and D. Laneuville, "PHD filter with diffuse spatial prior on the birth process with applications to GM-PHD filter," in *13th International Conference on Information Fusion*. Edinburgh, UK: IEEE, 2010, pp. 1 – 8.
- [38] S. Nagappa, D. E. Clark, and R. Mahler, "Incorporating Track Uncertainty into the OSPA Metric," in *2011 Proceedings of the 14th International Conference on Information Fusion (FUSION)*. Chicago, IL: IEEE, 2011, p. 8.
- [39] Y. Bar-Shalom, P. K. Willett, and X. Tian, *Tracking and Data Fusion: A Handbook of Algorithms*, Storrs, CT, USA, 2011.
- [40] X. Rong-Li and V. Jilkov, "Survey of maneuvering target tracking. Part I. Dynamic models," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 39, no. 4, pp. 1333 – 1364, 2003.
- [41] Kelvin Hughes, "Navigation Radar from Kelvin Hughes," 2017. [Online]. Available: <https://www.kelvinhughes.com/maritime/commercial-ships/navigation-radar>
- [42] I. Schlangen, E. Delande, J. Houssineau, and D. Clark, "A PHD Filter with Negative Binomial Clutter," in *19th International Conference on Information Fusion (FUSION)*. Heidelberg, Germany: IEEE, 2016, pp. 658–665.
- [43] D. Schuhmacher, B.-T. Vo, and B.-N. Vo, "A Consistent Metric for Performance Evaluation of Multi-Object Filters," *IEEE Transactions on Signal Processing*, vol. 56, no. 8, pp. 3447–3457, 2008.
- [44] D. Daley and D. Vere-Jones, *An introduction to the theory of point processes. vol. I. , Elementary theory and methods*, 2nd ed. New York: Springer, 2003.
- [45] P. Bernhard, "Chain differentials with an application to the mathematical fear operator," *Nonlinear Analysis: Theory, Methods & Applications*, vol. 62, no. 7, pp. 1225 – 1233, 2005.
- [46] D. Clark and J. Houssineau, "Faa di Brunos formula for chain differentials," *ArXiv e-prints*, *ArXiv:1310.2833*, pp. 1 – 7, 2013. [Online]. Available: <https://arxiv.org/abs/1310.2833>
- [47] S. N. Chiu, D. Stoyan, W. S. Kendall, and J. Mecke, *Stochastic geometry and its applications*. John Wiley & Sons, 2013.

PLACE
PHOTO
HERE

David Cormack received the M.Eng. degree in electrical and electronic engineering with distinction from Heriot-Watt University, Edinburgh, UK in 2016. He is currently working towards the Ph.D. degree in signal processing, awarded jointly by Heriot-Watt University and The University of Edinburgh. He has completed a number of internships at Leonardo MW Ltd, Edinburgh, and they now sponsor his Ph.D. degree. His research interests include radar signal processing, target detection, multiple target tracking, sensor fusion and sensor calibration.

PLACE
PHOTO
HERE

Isabel Schlangen received a German diploma in mathematics from the University of Bonn (Germany) in 2012 and a joint MSc degree in vision and robotics from the Universities of Burgundy (France), Girona (Spain), and Heriot-Watt (Edinburgh, UK) in 2014. In 2017, she was awarded a Ph.D. at Heriot-Watt University for her research on multi-object filtering with second-order moment statistics. She is a research associate at Fraunhofer FKIE (Wachtberg, Germany), working on multi-object estimation and resource management for radar applications.

PLACE
PHOTO
HERE

James R. Hopgood (M'02) received the Ph.D. degree from the University of Cambridge, U.K., in 2001. He was a Research Fellow at Queens' College and the Signal Processing Laboratory in Cambridge until 2004. He has since been at the University of Edinburgh, Scotland, where he is currently a Senior Lecturer with the Institute for Digital Communications, School of Engineering. His research projects span a number of applications in statistical signal and image processing. He has been Editor-in-Chief for the IET Journal of Signal Processing since 2011.

PLACE
PHOTO
HERE

Daniel E. Clark received the Ph.D. degree from Heriot-Watt University, Edinburgh, Scotland, in 2006. He is with the Department CITI, Telecom SudParis, France. From 2007 to 2017, he was with the School of Engineering and Physical Sciences, Heriot-Watt University. His research interests include the development of the theory and applications of multiobject estimation algorithms for sensor fusion problems. He has led a range of projects spanning theoretical algorithm development to practical deployment.

The Multi-Object Likelihood for the PHD and Panjer Filters

A supplement supporting the publication:

Joint Registration and Fusion of an Infra-Red Camera and Scanning Radar in a Maritime Context

David Cormack, Isabel Schlangen, James R. Hopgood and Daniel E. Clark

I. MOTIVATION

The MOLs of the PHD and Panjer filter (Eq. (9) and (13)) can be proved with the mathematical framework of point processes, which is a general form of the Finite Set Statistics (FISST) developed by Mahler [19]. In Sec. II, we will introduce some mathematical tools that will be necessary to conduct the proofs in Sec. III.

II. A SHORT INTRODUCTION TO POINT PROCESSES

A. Probability Generating Functionals

A useful stochastic concept to describe populations of objects is the notion of a *point process* Φ on $\mathcal{X}$, which is a random variable on the space $\mathfrak{X} = \bigcup_{n \geq 0} \mathcal{X}^n$ of finite sequences in $\mathcal{X}$. A realisation of Φ is a sequence¹ $\varphi = (x_1, \dots, x_n) \in \mathcal{X}^n$, where both the number $n \in \mathbb{N}$ and the states $x_i \in \mathcal{X}$ of the objects are random. Each point process Φ can be characterised by its Probability Generating Functional (PGFL)

$$G_\Phi(h) = \sum_{n \geq 0} \int \left[\prod_{i=1}^n h(x_i) \right] P_\Phi^{(n)}(x_1 \dots x_n) dx_1 \dots dx_n, \quad (15)$$

where the $P_\Phi^{(n)}$ are permutation-invariant probability measures on $\mathcal{X}^n$ for all $n \geq 0$, and h is a real-valued function. This type of functional is a useful mathematical representation of a point process since it makes it possible to model the number of objects on the one hand (via considering all cardinalities n in the sum) and their stochastic properties on the other (via the probabilities $P_\Phi^{(n)}$ for every fixed n). Let us look at a few common examples of PGFLs which will also be used later on.

The most simple PGFL is that of a Bernoulli process with parameter $p_\star$ and spatial distribution $s_\star$ given by

$$G_\star(h) = (1 - p_\star) + p_\star \int h(x) s_\star(x) dx. \quad (16)$$

This process is a suitable choice for modelling survival and detection processes since those can have only two outcomes; either, the object dies (or gets miss-detected) with probability $(1 - p_\star)$, or it survives (or is detected) with probability $p_\star$, in which case it evolves according to the spatial distribution $s_\star(\cdot)$.

The PGFL of a Poisson point process with parameter $\lambda_\bullet$ and intensity $\mu_\bullet(\cdot) = \lambda_\bullet s_\bullet(\cdot)$ is of the form [44]

$$G_\bullet(h) = \exp \left(\int [h(x) - 1] \mu_\bullet(x) dx \right). \quad (17)$$

¹If instead of sequences, one assumes that realisations of Φ are finite *sets*, we arrive at the Random Finite Set (RFS) theory derived by Mahler [19]. Sets do not allow state repetitions by definition, whereas sequences generally do.

This compact form is found by inserting the probability mass function of a Poisson distribution in Eq. (15) and using the series expansion of the exponential function to remove the infinite sum. The Poisson point process gained its popularity from the fact that while describing common phenomena in many applications, it only depends on one parameter and its exponential form makes it easy to work with.

A less known concept is that of the Panjer distribution, which is common in actuarial mathematics but only made its way to the engineering community very recently. The PGFL of a Panjer point process with spatial distribution s_o and parameters α_o, β_o is found to be [44], [34]

$$G_o(h) = \left(1 + \frac{1}{\beta_o} \int [1 - h(x)] s_o(x) dx \right)^{-\alpha_o}. \quad (18)$$

One can show that by taking the limit $\alpha \rightarrow \infty$ while keeping the ratio $\frac{\alpha}{\beta}$ constant, the PGFL (18) reduces to that of the Poisson process (17). Furthermore, a binomial (or negative binomial) process is obtained by choosing negative (or positive) values for both α and β , which makes this formulation versatile and useful in cases where the Poisson assumption is too restrictive.

B. PGFLs and functional differentiation

Many properties of a point process are encoded in its statistical moments, which are obtained via differentiation of its PGFL. Two particularly important and useful properties are:

- Calculating the expected value [47]:

$$\mathbb{E}[G_\Phi] = \delta G_\Phi(h; \eta)|_{h=1} \quad (19)$$

- Extracting the probability measure of having exactly n objects:

$$P_\Phi^{(n)}(x_1, \dots, x_n) = \delta^n G_\Phi(h; \delta_{x_1}, \dots, \delta_{x_n})|_{h=0}, \quad (20)$$

where δ_x denotes the Dirac delta function, being nonzero at x only.

The derivation of the MOL heavily relies on Eq. (20) since we want to express the likelihood of having exactly m measurements $z_1, \dots, z_m$, given the current state of the system. The filter equations of the Probability Hypothesis Density (PHD) filter and its variations propagate the expected value which is found via the rule (19). For the purpose of computing explicit formulae for the Multi-Object Likelihood (MOL) functions, let us first write down the so-called chain differential which admits a higher order product and chain rule [46]; other common differential operators are the Gâteaux and Fréchet differential [45].

Definition II.1 (Chain differential [45], [46]). Let $(\eta_n : \mathcal{X} \rightarrow \mathbb{R}^+)_{n \in \mathbb{N}}$ be a sequence of positive, bounded functions converging pointwise to a function $\eta : \mathcal{X} \rightarrow \mathbb{R}^+$, and let $(\varepsilon_n)_{n \in \mathbb{N}}$ be a sequence of positive real values converging to 0. The chain differential of a functional G with respect to its functional argument $h : \mathcal{X} \rightarrow \mathbb{R}^+$ in the direction of η is defined as

$$\delta G(h; \eta) = \lim_{n \rightarrow \infty} \frac{G(h + \varepsilon_n \eta_n) - G(h)}{\varepsilon_n}. \quad (21)$$

If the limit exists, it is unique for any sequence $(\varepsilon_n)_{n \in \mathbb{N}}$ and $(\eta_n : \mathcal{X} \rightarrow \mathbb{R}^+)_{n \in \mathbb{N}}$ with the above properties.

The chain differential yields an n -fold product rule [44]

$$\begin{aligned} \delta^n(F \cdot G)(h; \eta_1, \dots, \eta_n) \\ = \sum_{\omega \subseteq \{1, \dots, n\}} \delta^{|\omega|} F(h; (\eta_i)_{i \in \omega}) \delta^{|\bar{\omega}|} G(h; (\eta_j)_{j \in \bar{\omega}}), \end{aligned} \quad (22)$$

with $\bar{\omega}$ being the set complement $\{1, \dots, n\} \setminus \omega$, and an n -fold chain rule² [46]

$$\begin{aligned} \delta^n(F \circ G)(h; \eta_1, \dots, \eta_n) \\ = \sum_{\pi \in \Pi_n} \delta^{|\pi|} F\left(G(h; (\delta^{|\omega|} G(h; (\eta_i)_{i \in \omega}))_{\omega \in \pi}\right). \end{aligned} \quad (23)$$

C. Concatenation of point processes

Of course, most problems in Bayesian filtering cannot be described by one single point process only since several processes are involved in both the prediction and the update. Joint PGFLs can be written down in analogy to (15) but involving joint probabilities of the form $P_{\Phi}^{(n,m)}(x_1 \dots x_n, z_1, \dots, z_m)$. In general, these do not simplify, however we will be interested in two special cases, namely superposition and branching.

- 1) *Superposition*: If two processes are independent, their joint PGFL decomposes into a product of the form

$$G_{\Phi, \Psi}(h, g) = G_{\Phi}(h) G_{\Psi}(g). \quad (24)$$

This is used, for example, for superimposing spontaneous clutter which is assumed to be independent of the target process.

- 2) *Branching*: If each point in process Φ creates a new process Ψ , its resulting PGFL is the concatenation of the individual functionals:

$$G_{\Phi, \Psi}(h, g) = G_{\Phi}(h G_{\Psi}(g|\cdot)). \quad (25)$$

Such a structure is necessary to describe the detection process of the current targets: each target is either detected or missed, hence one would choose a Bernoulli process for the inner and the predicted process for the outer functional.

With the above considerations in mind, the general form of the PGFL describing the joint target and measurement processes, dependent on sensor state ϕ , is of the form

$$G_J(g, h|\phi) = G_{\text{pr}}(h G_d(g|\cdot, \phi)) G_c(g|\phi). \quad (26)$$

Here, $G_d(g|\cdot, \phi)$ is the Bernoulli detection process for a given target state x and sensor state ϕ with detection

probability $p_*(x) = p_d(x|\phi)$ and single-object single-measurement association likelihood $s_*(\cdot) = l(x|\cdot, \phi)$, i.e.

$$G_d(g|x, \phi) = 1 - p_d(x|\phi) + p_d(x|\phi) \int g(z) l(x|z, \phi) dz. \quad (27)$$

III. DERIVATION OF THE MOLs

The MOL of each filter is obtained by first setting $h = 1$ and then taking the $|Z|$ -fold derivative of (26) w.r.t. g (in the manner of (20)), i.e. differentiating once for each measurement z in the set Z . The only difference between the two forms is the model assumption regarding the predicted and clutter processes.

A. Proof of Theorem II.1

In case of the PHD filter, both G_{pr} and G_c are Poisson PGFLs with predicted intensity $\mu_{\text{pr}}^\bullet$ and clutter intensity $\mu_c^\bullet$, such that (26) takes the form

$$\begin{aligned} G_J^\bullet(g, 1|\phi) = \exp \left[\int (g(z) - 1) \mu_c^\bullet(z|\phi) \right. \\ \left. + \int (G_d(g|x, \phi) - 1) \mu_{\text{pr}}^\bullet(x|\phi) dx \right]. \end{aligned} \quad (28)$$

Differentiation by g requires repeated applications of the chain rule (23) which pulls out one multiplicative term $\mu_c^\bullet(z|\phi) + \int p_d(x|\phi) l(x|z, \phi) \mu_{\text{pr}}^\bullet(x|\phi) dx$ for each $z \in Z_k$. The final result is obtained by setting $g = 0$.

An alternative derivation of the MOL of the PHD filter is found in [22] where the FISST framework is used.

B. Proof of Theorem II.2

The Panjer filter assumes that G_{pr} and G_c are Panjer processes with spatial distributions s_{pr}° and s_c° , respectively, and corresponding parameters $\alpha_{\text{pr}}, \beta_{\text{pr}}$ and α_c, β_c . The explicit form of (26) becomes in this case

$$\begin{aligned} G_J^\circ(g, 1|\phi) = \left[1 + \frac{1}{\beta_c} \int (1 - g(z)) s_c^\circ(z|\phi) dz \right]^{-\alpha_c} \\ \cdot \left[1 + \frac{1}{\beta_{\text{pr}}} \int [1 - G_d(g|x, \phi)] s_{\text{pr}}^\circ(x|\phi) dx \right]^{-\alpha_{\text{pr}}} \end{aligned} \quad (29)$$

which needs to be differentiated using the general product rule (22). Using the notations $G_c^\circ(g|\phi) := F_c(g|\phi)^{-\alpha_c}$ and $G_{\text{pr}}^\circ(g|\phi) := F_d(g|\phi)^{-\alpha_{\text{pr}}}$ for the two right-hand-side terms of (29), the corresponding derivatives are

$$\delta^{|\bar{Z}|} G_c^\circ(g; (\delta_z)_{z \in \bar{Z}}|\phi) = F_c(g|\phi)^{-\alpha_c - |\bar{Z}|} \frac{(\alpha_c)_{|\bar{Z}|} \uparrow}{\beta_c^{|\bar{Z}|}} \prod_{z \in \bar{Z}} s_c^\circ(z) \quad (30)$$

and

$$\begin{aligned} \delta^{|Z|} G_{\text{pr}}^\circ(g; (\delta_z)_{z \in Z}|\phi) = F_d(g|\phi)^{-\alpha_{\text{pr}} - |Z|} \frac{(\alpha_{\text{pr}})_{|Z|} \uparrow}{\beta_{\text{pr}}^{|Z|}} \\ \cdot \prod_{z \in Z} \int p_d(x|\phi) l(x|z, \phi) s_{\text{pr}}^\circ(x|\phi) dx. \end{aligned} \quad (31)$$

Including (30) and (31) into (22), switching the summations and rearranging the terms leads to the desired result.

²This formula is also called Faà di Bruno's formula for chain differentials.