Message Passing for Joint Registration and Tracking in Multistatic Radar

David Cormack
School of Engineering and Physical Sciences
Heriot-Watt University
Edinburgh, UK, EH14 4AS
drc9@hw.ac.uk

James R. Hopgood
School of Engineering
The University of Edinburgh
Edinburgh, UK, EH9 3FG
james.hopgood@ed.ac.uk

***Abstract*—Sensor registration is fundamental in sensor fusion. Inaccuracies in sensor location and rotation can manifest themselves into the measurements used in Multiple Target Tracking (MTT), and dramatically degrade its performance. These registration parameters are often estimated separately to any multi-target estimation, which could lead to increased computational expense, and also to systematic errors. Recent works have shown that MTT algorithms derived from Belief Propagation (BP) are computationally efficient and highly scalable for large tracking scenarios. This work presents a hierarchical Bayesian model inspired by single-cluster methods from the Random Finite Set (RFS) literature, that allow for the registration parameters to be estimated jointly with the multiple target tracking. Simulations are carried out on a multistatic radar network containing two radars with a relative range and azimuth bias between them. Results are presented for a particle-BP MTT algorithm, and it's performance is compared to that of a Sequential Monte Carlo (SMC)-Probability Hypothesis Density (PHD) filter. The results show that the BP algorithm outperforms the PHD implementation in terms of accuracy by around 10%.**

***Index Terms*—Data fusion, sensor registration, belief propagation, radar, target tracking**

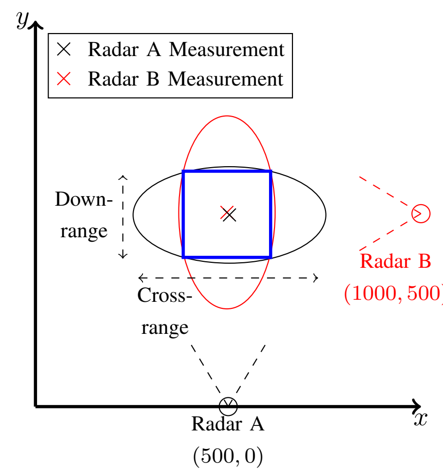

Fig. 1. Uncertainty reduction using two radars at approximately 90° to one another. The down-range uncertainty in Radar A can help limit the cross-range uncertainty of Radar B and vice-versa.

I. Introduction

The necessity for maintaining ground and airborne surveillance continues to grow. Tracking targets from different aspects using multiple radars is an important capability, as large volumes of space can be scanned in a short time period. By using sensor fusion, the output should be better in some way, compared with using each sensor independently [1]. Sensor networks that use commercial off-the-shelf (COTS) sensors that provide asynchronous measurements can make the fusion problem challenging. However, having incorrect sensor registration could vastly decrease fusion performance. This paper presents two different implementations using a hierarchical Bayesian model, inspired by the single-cluster framework from Random Finite Set (RFS) methods [2]. This framework allows for multiple estimation problems to be resolved simultaneously. The parent process resolves the sensor registration, with the offspring process estimating the states of multiple targets.

Recent advances have shown that accurate and scalable Multiple Target Tracking (MTT) methods can be developed through Belief Propagation (BP) and Message Passing (MP) [3], [4]. This paper will use a particle-based BP algorithm for MTT, and compare performance against a Sequential Monte Carlo (SMC) implementation of the Probability Hypothesis Density (PHD) filter [5], [6]. The BP algorithm faciliates a framework that will enable development of an efficient and scalable method for joint registration and fusion in the future. The MP method scales well, and the implementation does not require any high-dimensional or complex operations.

In this work, the 2-D setup shown in Fig. 1 is considered, where the two radars are located approximately 700 metres from one another in fixed and known locations, although there is uncertainty in their orientations. These radars detect the same target with similar accuracy in range and bearing within their respective frames of reference. Modern radar systems will often have measurement uncertainty that is much smaller down-range than it is in cross-range. Down-range uncertainty is determined by how fast the hardware can sample the received signal, and in modern radars this could be in the order of a few metres. Cross-range uncertainty is usually determined by the width of the radar beam. If this beam was projected onto the ground, the cross-range distance may be in the order of hundreds of metres. By using the down-range measurement from one radar to correct for the cross-range measurement in

978-1-7281-0503-1/19/$31.00 ©2019 IEEE

the other radar, a much smaller approximation of uncertainty is obtained as highlighted in blue in Fig. 1 [7].

A number of previous methods have been developed for mitigating the sensor registration problem. Pseudomeasurements are a common method for registration [8]–[10], however these methods treat this problem separately from the target tracking process, and assume much smaller biases than those presented in this work. Other previous works in the literature [11], [12] use BP methods on Markov Random Field (MRF) representations of sensor networks to accurately calibrate sensors, while using RFS tracking methods. The joint approach used here has been used in a varied range of applications previously [13]–[15] and is a special case of group or extended target tracking [2]. This allows for modelling of objects which are conditioned on a single, or potentially multiple, registration parameter(s). This paper presents the ”single-cluster method” as a solution to the sensor registration problem in a defence context. The algorithm is demonstrated on a set of simulations using two radars, where multiple targets are tracked and multiple registration parameters are estimated. Section II will formulate the problem and a present derivations of the parameter likelihoods. MTT and sensor fusion are introduced in Section III, with simulation details and results given in Section IV. Conclusions and future work are provided in Section V.

II. Sensor Registration

A. Overview

For this application in particular, where multiple sensor measurements are being fused, it is important to consider the sensor registration problem alongside that of sensor fusion. All measurements must be projected to a common frame of reference correctly so that they can be fused without any biases. Such biases could include angular orientation, timing, or possibly sensor location itself [10]. If these biases are not accounted for, any attempts at fusing the measurements could lead major tracking inaccuracies. Registration errors could stem from different sources such as incorrect alignment when being installed, sensor drift, or even from vibrations caused by the platform they are attached to.

The main objective of this work is to estimate the relative bias between the two radars, as shown in Fig. 2. When both sets of measurements are projected into a common frame, there appears an offset in both range and azimuth, which is denoteed by the bias vector $b_n = [r_b, \phi_b]^T$. These registration errors must also be estimated recursively along with the multiple target states. The targets that are tracked in the surveillance region make a contribution towards resolving these biases.

The joint estimation strategy presented here follows a similar process to that outlined in [15] and is highlighted in Fig. 3. The parent process in this work will be a 2-dimensional estimation problem to resolve both bias components in b_n. This will be performed using a SMC approach [16], with each particle representing a potential sensor configuration. It is noted that in practice these components may vary over time, however for this work only static parameters are considered.

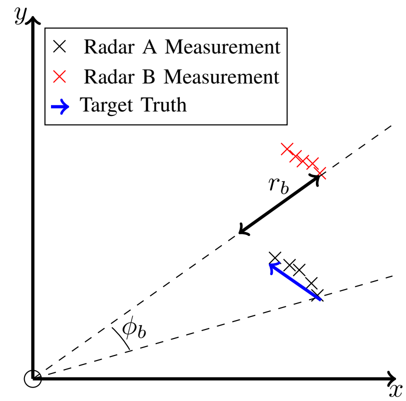

Fig. 2. An example of the sensor registration problem between two radars. Ideally sensor fusion would be performed between the two sets of measurements, however the systematic errors r_b and ϕ_b exist between the two sets of measurements when projected into the same frame of reference.

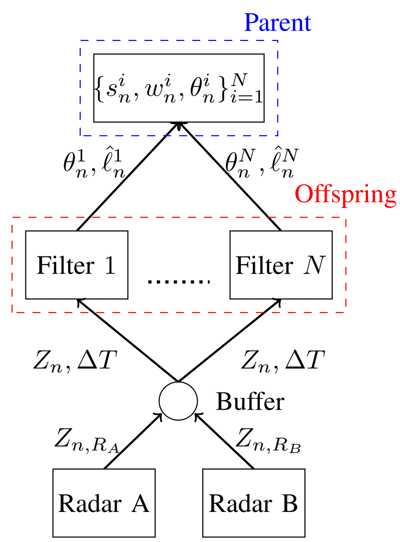

Fig. 3. Flowchart of the joint registration and fusion process. Asynchronous measurements are sent to the buffer when ready, and then used in the offspring layer for tracking. When measurements are received from Radar B, the MOL $\hat{\ell}_n^i = \hat{\ell}_n(s_n^i|Z_n)$ is calculated to update the sensor registration parameter(s).

B. Parameter Likelihoods

The parameter likelihood, $\hat{\ell}_n(s|\mathbf{z}_n)$, where s is the current sensor registration configuration and $\mathbf{z}_n$ are the observations to time-step n, relates the offspring process to the parent process. This is propagated using a Bayes recursion by finding the posterior $\hat{P}_n(s|\mathbf{z}_n)$. This consists of the prediction and update steps:

$$\hat{P}_{n|n-1}(s|\mathbf{z}_{n-1}) = \int_S \hat{f}_{n|n-1}(s|s')\hat{P}_{n-1}(s'|\mathbf{z}_{n-1})\mathrm{d}s', \quad (1a)$$

$$\hat{P}_n(s|\mathbf{z}_n) = \frac{\hat{\ell}_n(s|\mathbf{z}_n)\hat{P}_{n|n-1}(s|\mathbf{z}_{n-1})}{\int_S \hat{\ell}_n(s'|\mathbf{z}_n)\hat{P}_{n|n-1}(s'|\mathbf{z}_{n-1})\mathrm{d}s'} \quad (1b)$$

where s' is the previous sensor registration configuration, and $\hat{f}_{n|n-1}(s|s')$ is its transition probability density function (pdf). These equations are evaluated after the MTT in the offspring process has been performed. The transition function $\hat{f}_{n|n-1}(s|s')$ is applied to the parent likelihood first. As the parameters are assumed to be static, small amounts of random noise are applied to each particle state to account for small perturbations in the parameters over time [16]. The update step in (1b) uses the parameter likelihood value computed in each of MTT filters to determine the new parameter estimates.

1) Likelihood for BP Algorithm: By using MP, also known as BP or the Sum-Product Algorithm (SPA), MTT methods with lower computational cost and better scalability can be developed [4]. In this section, the notation closely follows that of [3], [4]. To use a BP type algorithm for registration, a suitable sensor-parameter likelihood function is needed. An augmented target is defined as $\mathbf{y}_{n,k} = [\mathbf{x}_{n,k}, r_{n,k}]^T$, where $\mathbf{x}_{n,k}$ is a single target state, and $r_{n,k}$ is a binary existence variable. Similar to (1a), the prediction equation is:

$$p(\mathbf{y}_n \mid \mathbf{z}_{n-1}) = \int f(\mathbf{y}_n \mid \mathbf{y}_{n-1})\, p(\mathbf{y}_{n-1} \mid \mathbf{z}_{n-1})\, \mathrm{d}\mathbf{x}_{n-1} \tag{2}$$

where it is assumed the posterior at time-step $n-1$ can be factorised as $p(\mathbf{y}_{n-1} \mid \mathbf{z}_{n-1}) = \prod_{k=1}^{K} p(\mathbf{y}_{n-1,k} \mid \mathbf{z}_{n-1})$, as can the state evolution $f(\mathbf{y}_n \mid \mathbf{y}_{n-1}) = \prod_{k=1}^{K} f(\mathbf{y}_{n,k} \mid \mathbf{y}_{n-1,k})$, i.e. the joint target density is a product over all individual target densities. Following [3, Sect V.A], this is written as

$$p(\mathbf{y}_n \mid \mathbf{z}_{n-1}) = \prod_{k=1}^{K} \alpha(\mathbf{y}_{n,k}) \tag{3}$$

where $\alpha(\mathbf{y}_{n,k}) \equiv \alpha(\mathbf{x}_{n,k},\, r_{n,k})$ is the marginal prediction:

$$\alpha(\mathbf{y}_{n,k}) = \sum_{r_{n,k}} \int f(\mathbf{y}_{n,k} \mid \mathbf{y}_{n-1,k}) \times p(\mathbf{y}_{n-1,k} \mid \mathbf{z}_{n-1})\, \mathrm{d}\mathbf{x}_{n-1,k}$$

Introducing the target-oriented association variables, $\mathbf{a}_n$, the update equation for the BP implementation can be written as,

$$p(\mathbf{y}_n, \mathbf{a}_n \mid \mathbf{z}_n) = \frac{p(\mathbf{z}_n \mid \mathbf{y}_n, \mathbf{a}_n)\, p(\mathbf{y}_n \mid \mathbf{z}_{n-1})}{p(\mathbf{z}_n \mid \mathbf{z}_{n-1})} \tag{4}$$

where $p(\mathbf{z}_n \mid \mathbf{y}_n, \mathbf{a}_n)$ is the single-object association likelihood, $p(\mathbf{y}_n \mid \mathbf{z}_{n-1})$ contains the predicted target states from (2), and $p(\mathbf{z}_n \mid \mathbf{z}_{n-1})$ is the evidence term necessary for deriving the sensor-parameter likelihood function. A stretching process is used to reduce message dimensionality and computational complexity. Here, the random vector $\mathbf{b}_n$ is introduced, which is an alternative measurement-oriented association variable, and can be derived directly from $\mathbf{a}_n$ [3]. Stretching introduces loops into factor graphs, and these loops create instances where the messages or beliefs are no longer exact. However, in [4] it is shown the algorithm will still converge. By using $\mathbf{a}_n$ and $\mathbf{b}_n$, high-dimensional factors in the graph have been replaced with many lower-dimensional factors. Using stretching [3, Eq. (27)] and (4), it is shown:

$$p(\mathbf{y}_n, \mathbf{a}_n|\mathbf{z}_n) = \frac{\prod_{k=1}^{K} v(\mathbf{y}_{n,k}, a_{n,k}|\mathbf{z}_n) \prod_{k=1}^{K} \alpha(\mathbf{y}_{n,k})}{p(\mathbf{z}_n|\mathbf{z}_{n-1})}. \tag{5}$$

where (see [3, Eqs. 11 & 17] for complete details and caveats):

$$v(\mathbf{y}_{n,k}, a_{n,k}|\mathbf{z}_n) = \frac{f(\mathbf{z}_{n,m}|\mathbf{x}_{n,k})}{f_{FA}(\mathbf{z}_{n,m})} \frac{P_d(\mathbf{x}_n)}{\mu_c}, \tag{6}$$

with $f_{FA}(\mathbf{z}_{n,m})$ representing the clutter distribution, μ_c being the mean number of false alarms and M is the number of measurements, $m \in \{1, \ldots, M\}$. Marginalising (5) over $\mathbf{a}_n$ and $\mathbf{y}_n$ gives the evidence term:

$$p(\mathbf{z}_n \mid \mathbf{z}_{n-1}) = \sum_{a_{n,k}} \int \cdots \int \prod_{k=1}^{K} v(\mathbf{y}_{n,k}, a_{n,k}|\mathbf{z}_n) \alpha(\mathbf{y}_{n,k}) \mathrm{d}\mathbf{x}_{n,1} \ldots \mathrm{d}\mathbf{x}_{n,k} \tag{7}$$

which reduces, using the definition in [3, Eq. (31)], to:

$$\hat{\ell}_n(s|\mathbf{z}_n) = p(\mathbf{z}_n|\mathbf{z}_{n-1}) = \sum_{a_k=0}^{\mathbf{a}_k} \prod_{k=1}^{K} \beta(a_{n,k}) \tag{8}$$

where $\beta(a_{n,k})$ is defined by

$$\beta(a_{n,k}) = \sum_{r_{n,k}} \int v(\mathbf{y}_{n,k}, a_{n,k}; \mathbf{z}_n) \alpha(\mathbf{y}_{n,k})\, \mathrm{d}\mathbf{y}_{n,k} \tag{9}$$

Note that $\beta(a_{n,k})$ can be interpreted as an approximation of the single-target association weights commonly found in the Probabilistic Data Association (PDA) and Joint Probabilistic Data Association (JPDA) filters [17].

2) Likelihood for PHD Filter: The parameter likelihood for the PHD filter is derived in [18], [19], and for a given sensor configuration s is,

$$\hat{\ell}_n(s|\mathbf{z}_n) = \frac{\prod_{z \in \mathbf{z}} [\mu_{c,n}(\mathbf{z}) + \int_{\mathcal{X}} p_d(\mathbf{x}_n)\, g_n(\mathbf{z}_m \mid \mathbf{x}_k, s)\, \mu_{n|n-1}(\mathrm{d}\mathbf{x} \mid s)]}{\exp\left[\int_{\mathcal{Z}} \mu_{c,n}(\mathbf{z})\, \mathrm{d}\mathbf{z} + \int_{\mathcal{X}} p_d(\mathbf{x}_n)\, \mu_{n|n-1}(\mathrm{d}\mathbf{x} \mid s)\right]} \tag{10}$$

where $p_d(\mathbf{x}_n)$ is the probability of detection, $g_n(\mathbf{z}_m \mid \mathbf{x}_k, s)$ is the single-object likelihood, $\mu_{c,n}(\mathbf{z})$ is the clutter intensity and $\mu_{n|n-1}(\mathrm{d}\mathbf{x} \mid s)$ is the predicted intensity.

III. Particle Implementations of MTT

A. Particle-BP Algorithm

These methods are a close approximation to Bayesian inference, and provide a flexible trade-off between computation time and accuracy. By using a particle implementation of BP, it is possible to overcome non-linearities between the state space and observation space, along with the potentially unknown and time-varying number of targets [4].

The main attraction of using the SPA is it's efficiency in computing the marginal posterior pdfs that are required for target tracking, compared to computing them using direct marginalisation. To use SPA for marginalisation, it is assumed the joint posterior pdf $f(\mathbf{x} \mid \mathbf{z}_{1:n})$, for arbitrary state $\mathbf{x}$, can be seen as a product of M lower-dimensional factors of subsets of $\mathbf{x}$, denoted $\mathbf{x}^{(m)}$, such that

$$f(\mathbf{x} \mid \mathbf{z}_{1:n}) \propto \prod_{m=1}^{M} \psi_m\left(\mathbf{x}^{(m)}\right). \tag{11}$$

The MP implementation used in this work is given in [3]. This implementation can be seen as an SPA-based reformulation of the Joint Integrated Probabilistic Data Association (JIPDA) filter [20]. Each target is represented by it's own distribution of N particles, which are fitted to a Gaussian distribution in order to perform the data association (DA) step. Measurements are associated to these Gaussians using the Sum-Product Algorithm for Data Association (SPADA) [21], and $\beta(a_{n,k})$, and in turn $\hat{\ell}_n(s|\mathbf{z}_n)$, can be evaluated.

B. PHD Approach

The PHD filter is a recent development in the MTT field [5], [6], [22]. This filter propagates the first-order moment of the target distribution, and assumes both the predicted number of targets and the clutter cardinality are both Poisson distributed. The parameter likelihood in (10) is implemented as part of the filter update step. This function will simplify as the integrals become summations of all components for one measurement, i.e. a sum of the corresponding component weights.

This particle-based implementation of the PHD filter can work explicitly with non-linearities. Direct application of typical SMC methods to propagate the intensity function used in the PHD filter would fail as the intensity function is not strictly a pdf, and the recursion used in the filter is not the exact standard Bayes recursion. Instead for this SMC implementation, the intensity function is represented by a large set of weighted random samples which are propagated over time using a generalised importance sampling and resampling strategy [6]. The number of particles used can be continually adapted, depending on the estimated number of targets in the surveillance region. Note that the main implementation difference between this and the MP approach is that the PHD implementation uses a general particle distribution to represent all targets in the state-space, rather than having N particles specifically for each individual target.

IV. Simulation

In order to compare each of the offspring processes, a simulated scenario has been made, consisting of two crossing targets; a challenging manoeuvre for tracking algorithms to correctly decipher at the crossing point. The parent process in this work is represented by 300 particles. These particles are initialised on a uniform grid between $r_b = [-150m \to +150m]$ and $\phi_b = [-3^\circ \to +3^\circ]$, and are resampled using stratified resampling [23] when the effective sample size is below a given threshold [16]. All results shown have been averaged over 50 Monte-Carlo (MC) runs. Tracking accuracy for each type of offspring process is compared by using the Optimal Sub-Pattern Assignment (OSPA) metric [24]. The parameters for the OSPA metric are $p = 2$ and $c = 100m$.

From Fig. 4, the importance of taking the sensor registration into account when performing fusion can be seen. If fusion is performed with uncalibrated sensors, a large increase in the OSPA distance is observed. Tracking with a single radar is accurate in itself, but performance can be increased by including a second radar. The method presented in this paper is shown by the estimated registration plot, which does not reach the optimal correct registration result, but is a vast increase in accuracy over the incorrect registration result.

TABLE I
Tracking Parameters

Quantity	Symbol	Value
Detection Probability	p_d	0.99
Survival Probability	p_s	0.95
Pruning Threshold	τ_{prune}	0.001
Extraction Threshold	$\tau_{extract}$	0.5
False Alarm Rate	λ_r	2
Birth Intensity	μ_b	0.01
Acceleration Noise	q	$1\ ms^{-2}$
Radar Measurement Noise	σ_{r_r}, σ_{ϕ_r}	$5m, 0.002^\circ$
Particles per Target	N	1000
Maximum Number of Targets	K	5

Fig. 5 shows the parameter estimates taken from the parent process. These are extracted as the Maximum A Posteriori (MAP) of the particle distribution. The estimates for ϕ_b are accurate to within 0.2° for the PHD approach, whereas the MP approach gives consistently more accurate estimates to within 0.1°. The r_b estimates appear more erratic around the true value however, due to the physics of the presented problem. Angular errors could result in biases of tens or hundreds of metres, making the estimation much less accurate. Small changes in range ($< \sigma_{r_r}$) will make little difference to the estimation and to the likelihood value passed to the parent process, hence why variations of approximately $\pm 3m$ around the true value can be observed.

V. Conclusions and Future Work

From the results, it can be seen that the MP approach gives more accurate tracking in terms of the OSPA metric, and also gives more accurate estimates of the registration parameters.

Following on from this work, we wish to improve the performance of the MP approach further. This could be achieved by replacing the single-cluster method with further MP operations, such that the joint estimation is completely evaluated using messages and beliefs.

Acknowledgements

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.

References

[1] W. Koch, *Tracking and Sensor Data Fusion: Methodological Framework and Selected Applications*, 1st ed. Berlin: Springer Verlag, 2014.

[2] A. Swain and D. Clark, "The single-group PHD filter: An analytic solution," in *2011 Proceedings of the 14th International Conference on Information Fusion (FUSION)*. Chicago, IL: IEEE, 2011.

[3] 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.

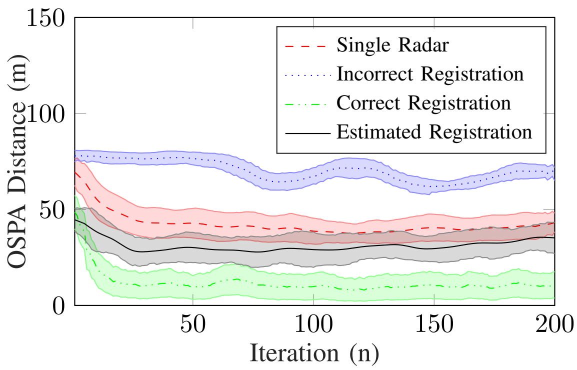

(a) PHD approach

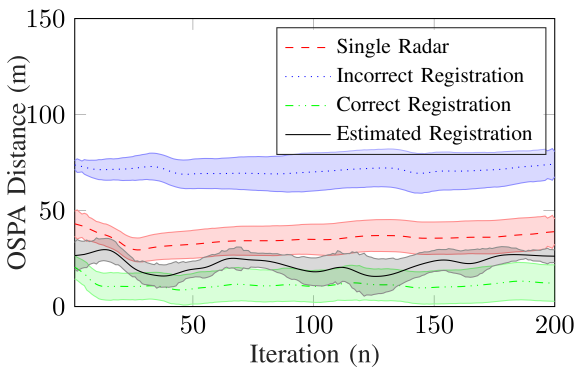

(b) MP approach

Fig. 4. OSPA Distance over Iteration

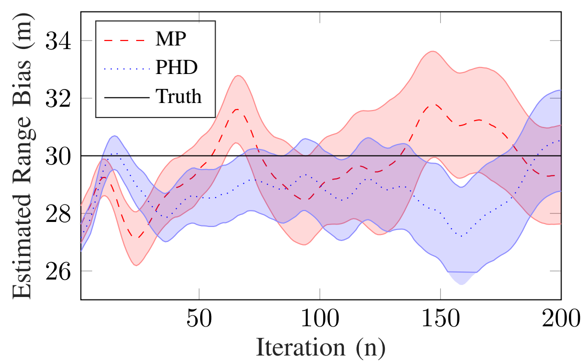

(a) Estimated Range Bias r_b

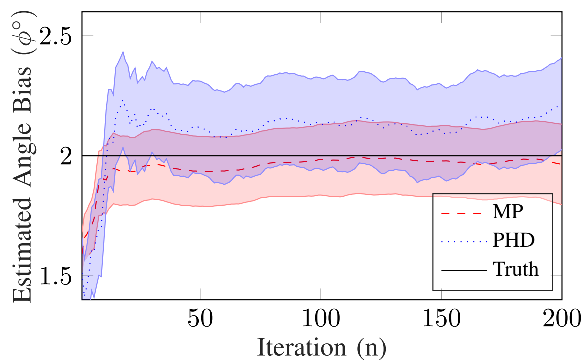

(b) Estimated Angle Bias ϕ_b

Fig. 5. Estimated Parameters over Iteration

[4] 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.

[5] 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.

[6] B.-N. Vo, S. Singh, and A. Doucet, "Sequential Monte Carlo methods for multitarget filtering with random finite sets," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 41, no. 4, pp. 1224 – 1245, 2005.

[7] M. Skolnik, *Introduction to Radar Systems*, 3rd ed. New York: McGraw-Hill, 2001.

[8] 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.

[9] 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.

[10] 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.

[11] M. Uney, B. Mulgrew, and D. E. Clark, "A Cooperative Approach to Sensor Localisation in Distributed Fusion Networks," *IEEE Transactions on Signal Processing*, vol. 64, no. 5, pp. 1187–1199, 2016.

[12] ——, "Latent Parameter Estimation in Fusion Networks Using Separable Likelihoods," *IEEE Transactions on Signal and Information Processing over Networks*, vol. 4, no. 4, pp. 752 – 768, 2018.

[13] 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.

[14] 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.

[15] 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.

[16] A. Doucet, N. de Freitas, and N. Gordon, *Sequential Monte Carlo Methods in Practice*. New York, NY: Springer-Verlag, 2001.

[17] Y. Bar-Shalom, P. K. Willett, and X. Tian, *Tracking and Data Fusion: A Handbook of Algorithms*, Storrs, CT, USA, 2011.

[18] A. Swain, "Group and extended target tracking with the probability hypothesis density filter," Ph.D. dissertation, Heriot-Watt University, 2013.

[19] 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

[20] D. Musicki, R. Evans, and S. Stankovic, "Integrated probabilistic data association," *IEEE Transactions on Automatic Control*, vol. 39, no. 6, pp. 1237 – 1241, 1994.

[21] J. Williams and R. Lau, "Approximate evaluation of marginal association probabilities with belief propagation," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 50, no. 4, pp. 2942 – 2959, 2014.

[22] R. Mahler, *Statistical Multisource-Multitarget Information Fusion*. Norwood, MA: Artech House, 2007.

[23] 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.

[24] 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.