

Systematic Analysis of the PMBM, PHD, JPDA and GNN Multi-Target Tracking Filters

Julian Smith, Florian Particke, Markus Hiller and Jörn Thielecke
Department of Electrical, Electronic and Communication Engineering
Information Technology
Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU)

Abstract—Multiple-target tracking has increasingly gained attention over the last 60 years, with the data association task being one of the most challenging aspects in sensor data fusion due to its computational burden. Hence, a plethora of algorithms has been proposed to solve this data association problem. However, most approaches are solely evaluated in comparison to algorithms of the same class. Therefore, this paper tries to give an overview and intends to evaluate systematically the four main classes of multiple-target tracking filters, namely the non-Bayesian data association filters, the Bayesian data association filters, the intensity filters and the multi-Bernoulli filters. These four classes are exemplified by respective filters, namely the Global Nearest Neighbor filter, the Joint Probabilistic Data Association filter, the Probability Hypothesis Density filter and the Poisson multi-Bernoulli mixture filter. These four filters are evaluated on two challenging simulated scenarios comprising situations with false measurements as well as birth and death of targets. The performance is assessed using the well-established GOSPA metric. It is shown that the Poisson multi-Bernoulli mixture filter outperforms the other three filters regarding the GOSPA metric in these scenarios, yielding a smaller mean error and deviation in its position estimates. This accuracy comes at the cost of a higher runtime performance, as the other three filter types require less computational time to produce their state estimates.

Index Terms—PMBM, PHD, JPDA, GNN, GOSPA, multiple-target tracking

I. INTRODUCTION

In a general multiple-target system, the states of the targets vary with time, while the number of targets might also change due to appearing and disappearing targets in the scene. The measurements of the sensor observing the targets are pre-processed into a set of points at each time step. It is important to be noted that existing targets may not be detected and that false measurements may occur. As a result, at each time step, the multi-target observation is a set of detections for which only some are actually generated by real targets. In addition, it is not known which target is responsible for which measurement.

A fundamental problem in multiple-target tracking (MTT) is the unknown association of measurements with appropriate targets. A number of MTT algorithms are used in various tracking applications at present. The Global Nearest Neighbor (GNN) [1], [25], [27], [28] filter is the earliest and simplest MTT algorithm [24]. The traditional GNN method is the multiple-target tracking extension of the Nearest Neighbor (NN) method [25], [39] which assigns the sample point or observation to the nearest target. The GNN method attempts to find and to propagate the single most likely hypothesis at each

scan [1]. One advantage of the GNN filter is that it requires low computation time and storage space [25], as it only uses the most probable hypothesis for data association [1]. A major problem of the GNN filter is that with some probability measurements are assigned to wrong targets, which can not be corrected. Invented in the 1980s, besides the multiple hypothesis tracker (MHT) [24] and probabilistic MHT (PMHT) [40], the Joint Probabilistic Data Association (JPDA) [6], [7], [26], [29], [30], [31] filter uses all possible measurements by weighting the innovation of the different measurements with the respective association probability at each time step [26]. The JPDA filter can yield better performance than the GNN filter in a heavily cluttered multiple-target environment [25], but due to its combinatorial nature, the data association problem makes up the bigger part of the computational load in multiple-target tracking algorithms [2].

To reduce this computational burden, the Random Finite Set (RFS) approach is an emerging and promising alternative to the traditional association-based methods [2], [4]. In the RFS formulation, the collection of individual targets is treated as a set-valued state and the collection of individual observations is treated as a set-valued observation [4]. These set-valued states and set-valued observations are modeled as RFSs. In addition, modeling birth and death processes of potential targets as Poisson processes, intensity based filters like the Probability Hypothesis Density (PHD) [12], [13], [14], [15], [16] filter or the Cardinalized PHD (CPHD) [41] filter avoid explicit association of measurements to targets. As a result, clustering on the intensity densities is needed to extract explicit estimates for the target states.

To overcome the problem that only suboptimal clustering algorithms yielding suboptimal results can be used, the focus was shifted towards the use of probability densities back in target space, whereas birth and death processes are still modeled as Poisson processes. With the development of conjugate multiple-target distributions for multi-Bernoulli distributions, meaning that the posterior multiple-target distribution has the same functional form as the prior, a significant trend in the RFS-based MTT algorithms has become obvious in the recent years [9]. Two increasingly popular MTT conjugate prior filters are the Generalized Labeled multi-Bernoulli (GLMB) filter [14], [17], [18], [19], [20], [21], [22], [23] and the Poisson multi-Bernoulli mixture (PMBM) filter [9], [10], [21], [32], [33]. This paper focuses on the PMBM filter, as it uses

no labels for its targets, similar to the GNN, JPDA and PHD filter. The PMBM filter consists of the union of a Poisson process and a multi-Bernoulli mixture. The multi-Bernoulli mixture considers all the data association hypotheses, whereas the Poisson process represents all undetected targets.

led as Poisson processes. With the development of conjugate multiple-target distributions for multi-Bernoulli distributions, meaning that the posterior multiple-target distribution has the same functional form as the prior, a significant trend in the RFS-based MTT algorithms has become obvious in the recent years [9]. Two increasingly popular MTT conjugate prior filters are the Generalized Labeled multi-Bernoulli (GLMB) filter [14], [17], [18], [19], [20], [21], [22], [23] and the Poisson multi-Bernoulli mixture (PMBM) filter [9], [10], [21], [32], [33]. This paper focuses on the PMBM filter, as it uses no labels for its targets, similar to the GNN, JPDA and PHD filter. The PMBM filter consists of the union of a Poisson process and a multi-Bernoulli mixture. The multi-Bernoulli mixture considers all the data association hypotheses, whereas the Poisson process represents all undetected targets.

The aim of this paper is to compare and evaluate structurally the performance of the different classes of MTT filters. In Section II, the GNN filter, the JPDA filter, the PHD filter and the PMBM filter are introduced and their mathematical background is briefly reviewed. In Section III, the four filters are compared and evaluated on two challenging simulated scenarios regarding their capability to estimate the number of targets and the targets' states. Their individual performance is assessed by using the state of the art metric GOSPA [5]. Additionally, the run-time performance of each filter is briefly discussed. In Section IV, concluding remarks summarizing the accuracy and performance of the four filters are provided.

II. THE MULTIPLE-TARGET FILTERS

This section provides a brief review of the underlying principles of the four MTT filters investigated in this work, starting in chronological order with the GNN filter, up to the most recent PMBM filter.

A. The Global Nearest Neighbor Filter

The simplest multiple-target filter is the GNN filter, an extension of the NN filter to the multiple-target case [1]. The GNN filter searches for the unique joint association of measurements to targets that minimizes a total distance or likelihood. Hence, every measurement is only assigned to one track neglecting less probable candidates. Realizations of the GNN filter use Munkres algorithm [38] for the cost optimal association of targets to measurements and to find and propagate the single most likely hypothesis at each time step [1]. Afterward, the filter performs standard Bayes filtering for each target using these associated measurements directly. Although the GNN scheme is intuitively appealing and simple to implement, it is susceptible to track loss and exhibits consequently poor performance when targets are not well separated. As a consequence, with some probability, measurements are assigned to wrong targets, which can not be corrected. Besides,

the GNN filter does not model birth and death processes. Therefore, it is limited to a fixed and known number of targets.

B. The Joint Probabilistic Data Association Filter

The JPDA filter is an extension of the Probabilistic Data Association (PDA) filter [26] to a fixed and known number of targets. The JPDA filter uses joint association events and joint association probabilities in order to avoid conflicting measurement to track assignments in the presence of multiple targets. The complexity of the calculation for joint association probabilities grows exponentially with the number of targets and the number of measurements. Moreover, since the basic JPDA filter can only accommodate a fixed and known number of targets, several novel extension have been proposed to accommodate an unknown and time varying number of targets [34], [35].

The key feature of the JPDA filter is the evaluation of the conditional probabilities of the following joint association events θ pertaining to the current time step k , which is omitted for simplicity, consisting of θ_{jt_j} , the event that measurement j originated from target t_j , with $j = 1, \dots, m$, $t = 0, 1, \dots, n_T$. The target under consideration, to which measurement j is associated in the event θ_{jt_j} , has the index t_j . For a more detailed explanation, the interested reader is referred to [2], [6], [7].

C. The Probability Hypothesis Density Filter

In scenarios with many targets, the conventional multiple-target Bayes filter tends to reach its computational limitations. Therefore, in many applications it is computationally advantageous to track target groupings instead of individual targets. The conventional approach to such problems is the attempt to track individual targets and assemble the group from them. Such an approach is highly impracticable in dense formations. Thus, in such applications, a standard technique such as the multiple-target Bayes filter may begin to fail. The strategy employed in the PHD filter is opposite to that used in conventional approaches [8]. Initially, it tracks only the overall group behavior, before attempting to detect and track individual targets to the extent which the quantity and quality of the data permits. The PHD filter is founded on the RFS framework and is based on a recursion that propagates the first-order statistical moment, or intensity, of the RFS states in time [36]. The intensity is also known in the tracking literature as the probability hypothesis density. This approximation was developed to alleviate the computational intractability in the multiple-target Bayes filter, which stems from the combinatorial nature of the multiple-target densities and the multiple integrations over the multiple-target state space. The PHD filter operates on the single-target state space and avoids the combinatorial problem that arises from data association. These remarkable features render the PHD filter extremely attractive [4].

The intensity, i.e., the first-order moment of a RFS X on the state space \mathcal{X} with probability distribution f , is a non-negative function $v(\cdot)$ such that for each region $S \subseteq \mathcal{X}$, we have

$$\int_S v(x)dx = \int_S x \cdot f(x)dx. \quad (1)$$

Consequently, the integral of v over any region S results in the expected number \hat{N} of elements of X that are in S . The local maxima of the intensity v are points in \mathcal{X} with the highest local concentration of expected number of elements, and therefore can be used to generate estimates for the elements of X . The simplest approach is to round \hat{N} to the nearest whole number and choose the resulting number of highest peaks from the intensity. A short presentation of the PHD filter recursion, which uses this intensity density, is given in the following, based on [4].

Let v_k and $v_{k|k-1}$ denote the respective intensities associated with the multiple-target posterior density f_k , and the multiple-target predicted density $f_{k|k-1}$ in the recursion of the multiple-target Bayes filter. The posterior intensity can be propagated in time via the PHD recursion with the predicted intensity $v_{k|k-1}$ which consists of three parts. The first part (2a) denotes the intensity of the surviving targets which transition via a transition density $g_{k|k-1}$ from a state x_{k-1} at time step $k-1$ to a state x_k at time step k with a probability of survival p_S . The second part (2b) describes the expected intensity $\beta_{k|k-1}$ of targets spawned from already existing targets with intensity v_{k-1} . The third part (2c) denotes the intensity $\gamma_{k|k-1}$ of new born targets at time step k . The predicted intensity at time step k is given by

$$v_{k|k-1}(x) = \int p_S(x_{k-1})g_{k|k-1}(x|x_{k-1})v_{k-1}(x_{k-1})dx_{k-1} \quad (2a)$$

$$+ \int \beta_{k|k-1}(x|x_{k-1})v_{k-1}(x_{k-1})dx_{k-1} \quad (2b)$$

$$+ \gamma_{k|k-1}(x). \quad (2c)$$

After having calculated the predicted intensity, the intensity must be updated on the basis of the received measurements. The updated intensity v_k consists of two parts. The first part (3a) takes the missed measurements into account, indicating that a part of the predicted intensity exists, but is still undetected. The second part (3b) comprises all received measurements, even those coming from false alarm or clutter, with a likelihood density $l_k(z|x)$ and the probability $p_D(x)$ of detecting the target. The intensity arising from clutter is denoted by $c_k(z)$. The updated intensity at time step k is then

$$v_k(x) = [1 - p_D(x)] \cdot v_{k|k-1}(x) \quad (3a)$$

$$+ \sum_{z \in Z_k} \frac{p_D(x)l_k(z|x)v_{k|k-1}(x)}{c_k(z) + \int p_D(x_{k-1})l_k(z|x_{k-1})v_{k|k-1}(x_{k-1})}. \quad (3b)$$

It has been successfully shown in [4] that the Gaussian mixture PHD (GM-PHD) filter admits a closed form solution to the PHD recursion. This GM-PHD filter is used in the evaluation part in Section IV.

D. The Poisson multi-Bernoulli mixture Filter

The PMBM filter is also based on the RFS framework to model the multiple-target tracking problem in a Bayesian fashion. As the name already indicates, it consists of the union of a Poisson process and a multi-Bernoulli mixture. The important property of the PMBM filter is the conjugacy of its component probability distribution, which has been shown in [33]. The multi-Bernoulli mixture, which considers all the data association hypotheses, can be efficiently implemented using a track-oriented multiple hypotheses tracking (MHT) formulation. The Poisson component considers all targets that have not yet been detected and enables an efficient management of the number of hypotheses covering potential targets.

The Poisson component models the undetected targets, which represent targets that exist at the current time step but have not yet been detected. Each measurement at each time step gives rise to a new potentially detected target. The multi-Bernoulli mixture component models the potentially detected targets. Any new measurement has the possibility of being the first detection of a target, but it can also correspond to a previously detected target. Furthermore, it can correspond to clutter, in which case there is no new target. As this target may or may not exist, its distribution is Bernoulli. The following derivation is based on [3].

The considered multiple-target density uses a Poisson density $f^{\text{Poisson}}(\cdot)$ on the set of undetected targets X_U and a multi-Bernoulli mixture $f^{\text{mbm}}(\cdot)$ on the set of detected targets X_T . Hence, the multiple-target density under consideration is given by

$$f(X) = \sum_{X_U \uplus X_T = X} f^{\text{Poisson}}(X_U) f^{\text{mbm}}(X_T). \quad (4)$$

The Poisson density is

$$f^{\text{Poisson}}(X_U) = e^{-\int \mu(x)dx} [\mu(\cdot)]^{X_U}, \quad (5)$$

where $\mu(\cdot)$ represents its intensity, which is the first-order statistical moment of the set of undetected targets X_U .

The multi-Bernoulli mixture has multiplicative weights such that

$$f^{\text{mbm}}(X_T) \propto \sum_j \sum_{X_1 \uplus \dots \uplus X_n = X_T} \prod_{i=1}^n \omega_{j,i}^b f_{j,i}^b(X_i), \quad (6)$$

where \propto stands for proportionality, j is an index over all global hypotheses (components of the mixtures), n is the number of potentially detected targets and $\omega_{j,i}^b$ and $f_{j,i}^b(\cdot)$ are the weight and the Bernoulli density of a potentially detected target i under the j^{th} global hypothesis. The set X_T of states in (6) is $X_T = X_1 \uplus \dots \uplus X_n$, where X_i is the state of target i . The symbol \uplus represents the disjoint union, indicating that X_T is the union of the mutually disjoint sets X_1, \dots, X_n . The weight of global hypothesis j is proportional to the product of the hypothesis weights $\prod_{i=1}^n \omega_{j,i}^b$ for n potentially detected targets. The Bernoulli densities have the expression

$$f_{j,i}^b(X) = \begin{cases} 1 - r_{j,i} & X = \emptyset \\ r_{j,i} f_{j,i}(x) & X = \{x\}, \\ 0 & \text{otherwise} \end{cases} \quad (7)$$

where $r_{j,i}$ is the probability of existence and $f_{j,i}(\cdot)$ is the state density.

In addition, for each potentially detected target, there are single-target hypotheses, which represent possible histories of target-to-measurement (or false detection) associations. A single-target hypothesis, along with the existence probability of the corresponding Bernoulli RFS, incorporates information about the following events:

- The target never existed.
- The target exists at the current time.
- The target existed, but death occurred subsequent to its latest detection.

Finally, a global hypothesis contains one single-target hypothesis for each potential target, in conjunction with the constraint that each measurement must be contained in exactly one single-target hypothesis.

Global hypotheses can be expressed in terms of single-target hypotheses. A single-target hypothesis corresponds to a sequence of measurements associated to a potentially detected target. Given a single-target hypothesis, this potentially detected target follows a Bernoulli distribution. Therefore, each measurement initiates a new single-target hypothesis. At subsequent time steps, new single-target hypotheses are created by associating single-target hypotheses with current measurements or with a misdetection. Global hypotheses are thus a collection of such single-target hypotheses, with the auxiliary conditions that no measurement remains unassociated, and that a measurement can be assigned only to one single-target hypothesis. Consistent with the GM-PHD, a GM-PMBM implementation is chosen in Section IV.

III. EVALUATION OF THE FOUR FILTERS

In this section, the evaluation results based on different simulated scenarios are presented to allow the direct comparison of the four introduced filter types, namely represented by the GNN, JPDA, PHD and PMBM filter. An area of size $[0, 100] \times [0, 100]$, where all units are in international system, is considered. Target states consist of 2D position and velocity $[p_x, p_y, v_x, v_y]^T$. The GNN and JPDA filters assume that four targets are present in the area of interest, all in the proximity of $[50, 50, 0, 0]^T$. The PHD and PMBM filter consider targets being born according to a Poisson process of intensity 0.005 and Gaussian density with mean $[50, 50, 0, 0]^T$ and covariance $\text{diag}([50^2, 50^2, 1, 1])$, to cover the region of interest. Each target has a probability of survival $p_S = 0.99$ and follows linear Gaussian dynamics with

$$\mathbf{F} = \begin{bmatrix} \mathbf{I}_2 & \Delta \mathbf{I}_2 \\ \mathbf{0}_2 & \mathbf{I}_2 \end{bmatrix}, \quad \mathbf{Q} = \sigma_\nu^2 \begin{bmatrix} \frac{\Delta^4}{4} \mathbf{I}_2 & \frac{\Delta^3}{2} \mathbf{I}_2 \\ \frac{\Delta^3}{2} \mathbf{I}_2 & \Delta^2 \mathbf{I}_2 \end{bmatrix}, \quad (8)$$

where \mathbf{I}_n and $\mathbf{0}_n$ denote, respectively, the $n \times n$ identity and zero matrices, $\Delta = 1$ is the sampling period and $\sigma_\nu = 1$ is the standard deviation of the process noise. Each target is detected with probability $p_D = 0.9$ and the measurement follows a observation model with

$$\mathbf{H} = [\mathbf{I}_2 \quad \mathbf{0}_2], \quad \mathbf{R} = \sigma_\epsilon^2 \mathbf{I}_2, \quad (9)$$

where $\sigma_\epsilon = 1$ is the standard deviation of the measurement noise. We also consider Poisson clutter uniform over the region of interest with λ_c denoting the number of expected false measurements at each time step. The GNN, JPDA and PMBM filters use ellipsoidal gating on measurements in order to reduce computational complexity. The gates' threshold is 20. The PHD filter uses the same ellipsoidal gate to merge components in close proximity. Moreover, the PMBM filter is set to use a maximum of 30 global hypotheses at each time step.

A. Scenario with Clutter

In order to investigate the performance and operational principle of each filter, the behavior of the filters under different scenarios is going to be investigated. The first scenario is visualized in Figure 1. Each of the four targets moves in a linear straight line, following a constant velocity model as defined in (8). The false measurements are Poisson clutter with a uniform distribution in the region of interest, where λ_c denotes the expected number of false alarms per time step in the region of interest. Different values of λ_c are used to distinguish five sub-scenarios, namely $\lambda_c \in \{0, 1, 2, 5, 10\}$. To get a good sense of how well these filters work in the presence of clutter, each of these five sub-scenarios has been simulated a total of 100 times, and the clutter measurements in each case randomly appear in varying positions. The metric, which is used to evaluate the four filters, is the GOSPA metric, with the parameters $\alpha = 2, c = 10, p = 2$, as proposed in [5]. The GOSPA metric consists of two parts: localization errors for the properly detected targets and a cardinality error for missed and false targets. A cardinality error of one corresponds for this parametrization to a localization error of $\sqrt{50}$. The




Fig. 1. Four targets move in straight lines. Clutter measurements, which are not shown here, are added to this scenario.

results for the four filters are summarized in Figure 2. The results are visualized using boxplots [37]. As can be seen, the PMBM filter outperforms the other three filters in every scenario. The PMBM filter is very stable with regard to clutter measurements. This is due to the target-to-measurement association. Only multi-Bernoulli components with a probability of existence higher than a certain threshold (in this case $r = 0.60$) are considered to belong to a true target. Therefore,

the filter recognizes most clutter measurements as false alarms. The PMBM filter only makes a false estimation regarding the cardinality when two clutter measurements fall in close proximity to one another between successive time steps. This can be seen by the increasing GOSPA error with an increasing number of clutter measurements. The PHD filter has difficulty with clutter, as it does not make any target-to-measurement associations. Every clutter measurement adds to the intensity density, leading to a larger GOSPA error. It is to be noted that PHD components with an intensity greater 0.60 are considered to be a true target. The JPDA filter has the greatest difficulty in the reliable estimation of target states. It commits no error in the cardinality, but the presence of clutter tends to deflect its estimates away from their true states, resulting in a large GOSPA error, which becomes more severe with increasing clutter. The GNN filter has a rather good performance in an environment, where the number of targets is known and a low number of clutter measurements can be expected. But as the number of clutter measurements increases, the performance of the GNN filter decreases. This is due to the rising probability of using a false measurement to update the targets' state. The significant differences in the filters' performances become even more obvious by specifically analyzing one iteration with $\lambda_c = 5$ clutter measurements, depicted in Figure 3(a). The PHD filter often misestimates the number of targets, whereas the PMBM filter is very stable, only misestimating the number of targets four times, namely at time steps $k \in \{0, 5, 6, 13\}$. This is depicted in Figure 3(b). In the beginning, the filter assumes that no targets are present, which results in a high GOSPA error. The GNN and JPDA filter have a high average GOSPA error, as they have difficulties in environments with many clutter measurements to estimate the correct target states. The localization error of the GOSPA metric for the PHD and PMBM filters is depicted in Figure 3(c).

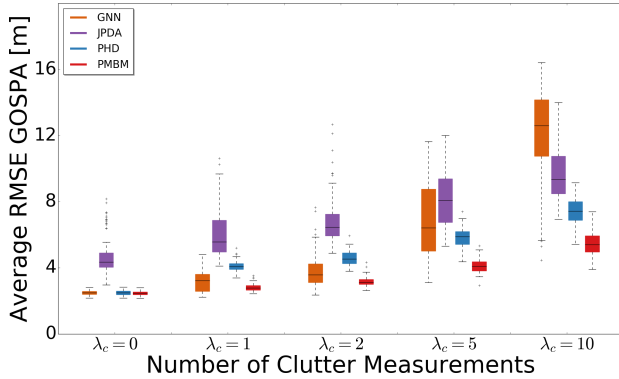




Fig. 2. The root mean square GOSPA error averaged over all time steps and iterations of the four algorithms for different values of λ_c . The results are visualized using boxplots [37].

B. Scenario with Clutter and Birth/Death of Targets


As indicated in Figure 4, this scenario is similar to the previous one. The primary difference now is that the red target dies after 8 time steps, and the green target is born after 16



(a) Accuracy errors relative to GOSPA metric of the GNN, JPDA, PHD and PMBM filters. Shown is the overall GOSPA error for all four filters.



(b) The cardinality error of the GOSPA metric. A cardinality error of one adds to the GOSPA metric as a localization error of $\sqrt{50}$. Visualized are the PHD and PMBM filters, as only these have a cardinality error.



(c) The localization error of the GOSPA metric. In the beginning, the GOSPA metric has no localization error, as the error consists only of the cardinality error.

Fig. 3. The GOSPA metric of the GNN, JPDA, PHD and PMBM filters. Shown is one explicit GOSPA error of each filter during each of the 20 time steps for one iteration with $\lambda_c = 5$ in Figure (a). The cardinality error of the GOSPA metric is depicted in Figure (b) and the localization error in Figure (c).

time steps. At each time step clutter is included, as in the previous scenario. These false measurements are modeled by Poisson clutter with a uniform distribution over the region of interest and with an expected number $\lambda_c = \{1, 2, 5, 10\}$ of clutter measurements per time step. The targets follow a constant velocity model, where process and measurement noise are included as in the scenarios before. As before, each sub-scenario has been simulated a total of 100 times for each clutter number.




Fig. 4. Four targets move in straight lines. At time step $k = 8$ the red target dies. At time step $k = 16$ the green target is born. Clutter is added to this scenario, but not shown in the figure.

In Figure 5, the average GOSPA error over all time steps and iterations of each filter is shown. In most cases, the PMBM filter outperforms the PHD filter. Only in the case with $\lambda_c = 10$ and an decreased probability of detection $p_D = 0.6$, the PHD filter has a slightly smaller GOSPA error than the PMBM filter. The JPDA and GNN filters are excluded from this scenario, because they cannot handle target birth or death events. In Figure 6(a), the GOSPA error for




Fig. 5. The root mean square GOSPA error averaged over all time steps and iterations of the PHD and PMBM filter for different values of λ_c and p_D .

the PMBM and PHD filters for one iteration with birth and death event of targets is displayed, with the parameters $\lambda_c = 5$ and $p_D = 0.9$. As before, the cardinality error in Figure 6(b) and the localization error in Figure 6(c) are depicted as well. The PMBM filter continues to handle clutter well, the PHD filter, however, is having more problems with clutter. The

PMBM filter manages the death of a target at time step $k = 8$ well, detecting it when the measurement of the dead target is no longer made. The PHD filter has difficulties and takes some time to estimate the correct number of targets present. Both filters take some time to recognize the birth event, with the PMBM filter responding more rapidly than the PHD filter. As can be seen at time step $k = 16$, when the green target is born, both filters have problems detecting it immediately. This results in a larger GOSPA error at this time step. The PMBM filter recognizes the newborn target sooner and therefore decreases its GOSPA error faster than the PHD filter.

C. Runtime Performance of the four filters

As the four filters have been evaluated regarding the GOSPA metric, a brief overview is given regarding the filters' runtime performance. All computations were performed in a development and simulation environment with the following specifications:


- Intel(R) Core(TM) i5-4690 CPU @ 3.50 GHz
- 16 GB RAM
- Windows 7 (64 bit operating system)
- Python 2.7.10, using only the numpy library.

The ideal scenario with $\lambda_c = 0$ false measurements has been evaluated in this case. The GNN filter is the fastest filter regarding the runtime, needing a mean time of 0.021 s for one iteration. The PHD filter follows with a mean runtime of 0.044 s for one iteration. The JPDA filter needs an average of 0.251 s for one iteration. Finally, the PMBM filter, due to its global hypotheses, needs an average of 0.646 s to perform one whole iteration. This is because of the filters' approaches to data association. The GNN uses only the single best hypothesis data association. The PHD uses no data association at all. The JPDA and PMBM filters use most of their computational load for the data association.


In the scenario with a total of $\lambda_c = 10$ false measurements, the GNN filter is still the fastest filter regarding the runtime, needing a mean of 0.076 s for one iteration. As before, the PHD filter follows with a mean runtime of 0.529 s for one iteration. Other than before, the PMBM filter needs an average of 2.996 s for one iteration, surpassing the JPDA filter. With an increasing number of measurements, the performance of the JPDA filter decreases significantly. It needs an average of 22.662 s to complete one iteration, due to its use of joint association events and joint association probabilities.

IV. CONCLUSION


Four different types of multiple-target filters have been examined and evaluated according to the GOSPA metric. Scenarios with the inclusion of clutter as well as birth and death events of targets have been simulated, to assess the performance of handling challenging scenarios. As can be seen, the quality and accuracy of multiple-target tracking filters have risen in the course of time. With the upcoming of the RFS formulation, filters do no longer need to know the number of targets to produce accurate state estimates and are therefore



(a) Accuracy errors relative to GOSPA metric of the PHD and PMBM filters for one iteration with $\lambda_c = 5$ and $p_D = 0.9$. Shown is the overall GOSPA error for the two filters.



(b) The cardinality error of the GOSPA metric. A cardinality error of one adds to the GOSPA metric as a localization error of $\sqrt{50}$.



(c) The localization error of the GOSPA metric. In the beginning, the GOSPA metric has no localization error, as the error consists only of the cardinality error.

Fig. 6. Accuracy errors relative to GOSPA metric of the PHD and PMBM filters for one iteration with $\lambda_c = 5$ and $p_D = 0.9$. The cardinality error of the GOSPA metric is depicted in Figure (b) and the localization error in Figure (c).

an improvement compared to classical multiple-target filters as the GNN or the JPDA filters. This development accumulates in the PMBM filter, one of the newest filters based on the RFS formulation.

In the scenario without birth and death events of targets, the PMBM filter outperforms the others, as its data association allows it to handle clutter measurements successfully. The JPDA filter and the GNN filter cannot handle the birth and death of targets. In most scenarios, the PMBM filter outperforms the PHD filter.

Regarding the runtime performance of the filters, the GNN filter performs the fastest, followed by the PHD filter. The JPDA and PMBM filters require the highest computational time due to the intensive data association task.

In case of tracking many targets in close proximity, it is recommended to use the PHD filter, as it is able to track groups instead of separated targets, with a low runtime performance. On the other hand, when it is of importance to clearly separate targets, the PMBM filter is recommended, assumed enough computing power is available.

As an outlook, to enhance the runtime performance of the PMBM filter further, a change from a Python implementation to a C++ implementation is proposed. Furthermore, labeled multi-Bernoulli filters have shown to perform efficiently in severe environments with higher computation burden [22] and can be implemented very efficiently [18] and are therefore a suitable alternative to the PMBM filter.

REFERENCES

- [1] P. Konstantinova, A. Udvarov and T. Semerdjiev, "A Study of a Target Tracking Algorithm Using Global Nearest Neighbor Approach," International Conference on Computer Systems and Technologies - CompSysTech'2003
- [2] B.-N. Vo, M. Mallick, Y. Bar-Shalom, S. Coraluppi, R. Osborne, R. Mahler and B.-T. Vo, "Multitarget Tracking," Preprint: Wiley Encyclopedia of Electrical and Electronics Engineering, Wiley, Sept. 2015.
- [3] Á. F. García-Fernández and J. L. Williams and K. Granström and L. Svensson, "Poisson Multi-Bernoulli Mixture Filter: Direct Derivation and Implementation," IEEE Transactions on Aerospace and Electronic Systems, vol 54, 2018, pp.1883–1901.
- [4] B.-N. Vo and W.-K. Ma, "The Gaussian Mixture Probability Hypothesis Density Filter," IEEE Transactions on Signal Processing, vol 54, 2006, pp. 4091–4104.
- [5] A. S. Rahmathullah and Á. F. García-Fernández and L. Svensson, "Generalized optimal sub-pattern assignment metric," 2017 20th International Conference on Information Fusion (Fusion), 2017, pp. 1–8.
- [6] T. Fortmann and Y. Bar-Shalom and M. Scheffe, "Sonar tracking of multiple targets using joint probabilistic data association," IEEE Journal of Oceanic Engineering, vol 8, 1983, pp.173–184.
- [7] S. H. Rezatofighi and A. Milan and Z. Zhang and Q. Shi and A. Dick and I. Reid, "Joint Probabilistic Data Association Revisited," 2015 IEEE International Conference on Computer Vision (ICCV), 2015, pp. 3047–3055.
- [8] Ronald P.S. Mahler, "Statistical Multisource-multitarget Information Fusion", Artech House Information Warfare Library, Artech House, Inc, 2007.
- [9] Y. Xia and K. Granström and L. Svensson and Á. F. García-Fernández, "An Implementation of the Poisson Multi-Bernoulli Mixture Trajectory Filter via Dual Decomposition," 2018 21st International Conference on Information Fusion (FUSION), 2018, pp. 1–8.
- [10] K. Granström and L. Svensson and Y. Xia and J. Williams and Á. F. García-Fernández, "Poisson Multi-Bernoulli Mixture Trackers: Continuity Through Random Finite Sets of Trajectories," 2018 21st International Conference on Information Fusion (FUSION), 2018, pp. 1–5.

- [11] L. Cament and M. Adams and J. Correa, "A Multi-Sensor, Gibbs Sampled, Implementation of the Multi-Bernoulli Poisson Filter," 2018 21st International Conference on Information Fusion (FUSION), 2018, pp. 2580–2587.
- [12] Á. F. García-Fernández and L. Svensson, "Trajectory probability hypothesis density filter," 2018 21st International Conference on Information Fusion (FUSION), 2018, pp. 1430–1437.
- [13] C. Berry and D. J. Bucci and S. W. Schmidt, "Passive Multi-Target Tracking Using the Adaptive Birth Intensity PHD Filter," 2018 21st International Conference on Information Fusion (FUSION), 2018, pp. 353–360.
- [14] L. Legrand and A. Giremus and É. Grivel and L. Ratton and B. Joseph and C. Magnant, "A Hierarchical LMB/PHD Filter for Multiple Groups of Targets with Coordinated Motions," 2018 21st International Conference on Information Fusion (FUSION), 2018, pp. 1–8.
- [15] Z. Fu and S. M. Naqvi and J. A. Chambers, "Collaborative Detector Fusion of Data-Driven PHD Filter for Online Multiple Human Tracking," 2018 21st International Conference on Information Fusion (FUSION), 2018, pp. 1976–1981.
- [16] Y. Guo and J. Pan and D. Peng and Z. Chen and H. Shentu, "GMPHD Based Multi-Scan Clutter Sparsity Estimation," 2018 21st International Conference on Information Fusion (FUSION), 2018, pp. 1453–1458.
- [17] B. Yang and J. Wang and W. Wang and S. Wei, "Multipath Generalized Labeled Multi-Bernoulli Filter," 2018 21st International Conference on Information Fusion (FUSION), 2018, pp. 1423–1429.
- [18] B. Vo and B. Vo, "Multi-Scan Generalized Labeled Multi-Bernoulli Filter," 2018 21st International Conference on Information Fusion (FUSION), 2018, pp. 195–202.
- [19] S. Panicker and A. K. Gostar and A. Bab-Haidashar and R. Hosein-zhad, "Sensor Control for Selective Object Tracking Using Labeled Multi-Bernoulli Filter," 2018 21st International Conference on Information Fusion (FUSION), 2018, pp. 2218–2224.
- [20] D. Kellner and M. Aeberhard, "Handling of Multiple Measurement Hypotheses in an Efficient Labeled Multi-Bernoulli Filter," 2018 21st International Conference on Information Fusion (FUSION), 2018, pp. 203–210.
- [21] A. Saucan and P. K. Varshney, "Distributed Cross-Entropy δ -GLMB Filter for Multi-Sensor Multi-Target Tracking," 2018 21st International Conference on Information Fusion (FUSION), 2018, pp. 1559–1566.
- [22] J. Wang and Z. Jing and P. Dong and J. Cheng, "A Variational Bayesian Labeled Multi-Bernoulli Filter for Tracking with Inverse Wishart Distribution," 2018 21st International Conference on Information Fusion (FUSION), 2018, pp. 219–225.
- [23] M. I. Hossain and A. K. Gostar and A. Bab-Hadiashar and R. Hosein-zhad, "Visual Mitosis Detection and Cell Tracking Using Labeled Multi-Bernoulli Filter," 2018 21st International Conference on Information Fusion (FUSION), 2018, pp. 1–5.
- [24] S.S. Blackman and R. Popoli, "Design and Analysis of Modern Tracking Systems," Boston: Artech House, 1999.
- [25] M. Shi and Q. Ling and Z. Yu and J. Zhu, "Association using modified Global Nearest Neighbor in the presence of bias," Proceedings of the 32nd Chinese Control Conference, 2013, pp. 4688–4691.
- [26] Y. Bar-Shalom and F. Daum and J. Huang, "The probabilistic data association filter," IEEE Control Systems Magazine, vol 29, 2009, pp. 82–100.
- [27] Y. Ruan and L. Hong and D. Wicker, "Analytic performance prediction of feature-aided global nearest neighbour algorithm in dense target scenarios," IET Radar, Sonar Navigation, 2007, pp. 369–376.
- [28] Lu-Min Zhu, Yong-Liang Zhang, Wen Zhang, Zhu-Cheng Tao and Chao-Fan Liu, "Fish motion tracking based on RGB color space and interframe global nearest neighbour," International Conference on Automatic Control and Artificial Intelligence (ACAI 2012), Xiamen, 2012, pp. 1061–1064.
- [29] M. A. Tinati and T. Y. Rezaii and M. J. Museviniya, "Regularized and simplified Monte Carlo Joint Probabilistic Data Association Filter for multi-target tracking in wireless sensor networks," 2009 IEEE International Symposium on Signal Processing and Information Technology (ISSPIT), 2009, pp. 600–605.
- [30] B. Habtemariam and R. Tharmarasa and T. Thayaparan and M. Mallick and T. Kirubarajan, "A Multiple-Detection Joint Probabilistic Data Association Filter," IEEE Journal of Selected Topics in Signal Processing, 2013, pp. 461–471.
- [31] S. S. Ahmeda and M. Keche and I. Harrison and M. S. Woolfson, "Adaptive joint probabilistic data association algorithm for tracking multiple targets in cluttered environment," IEE Proceedings - Radar, Sonar and Navigation, 1997, pp. 309–314.
- [32] K. Granström and M. Fatemi and L. Svensson, "Gamma Gaussian inverse-Wishart Poisson multi-Bernoulli filter for extended target tracking," 2016 19th International Conference on Information Fusion (FUSION), 2016, pp. 893–900.
- [33] J. L. Williams, "Marginal multi-bernoulli filters: RFS derivation of MHT, JIPDA, and association-based member," IEEE Transactions on Aerospace and Electronic Systems, 2015, pp. 1664–1687.
- [34] Srečko Juric-Kavelj and Ivan Markovic and Ivan Petrovic, "People Tracking with Heterogeneous Sensors using JPDAF with Entropy Based Track Management," ECMR, 2011.
- [35] S. He and H. Shin and A. Tsourdos, "Joint probabilistic data association filter with unknown detection probability and clutter rate," 2017 IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI), 2017, pp. 559–564.
- [36] R. Mahler "Multi-target Bayes filtering via first-order multi-target moments," IEEE Trans. Aerosp. Electron. Syst., vol 39, 2003, pp. 1152–1178.
- [37] John W. Turkey, "Exploratory data analysis," Addison-Wesley, 1977.
- [38] H. W. Kuhn, "The Hungarian Method for the Assignment Problem," Naval Research Logistics Quarterly 2, 1955, pp. 83–97.
- [39] T. M. Cover and P. E. Hart, "Nearest Neighbor Pattern Classification," IEEE Trans. on Information Theory, vol. 13, 1967, pp. 21.
- [40] R. L. Streit and T. E. Luginbuhl, "Maximum likelihood method for probabilistic multi-hypothesis tracking," SPIE International Symposium, Signal and Data Processing of Small Targets, 1994, pp. 394–405.
- [41] R. Mahler, "PHD filters of higher order in target number," IEEE Trans. Aerospace and Electronic Systems, 2007, pp. 1523–1543.