Sensor management for search and track using the Poisson multi-Bernoulli mixture filter

Per Boström-Rost, Daniel Axehill and Gustaf Hendeby

The self-archived postprint version of this journal article is available at Linköping University Institutional Repository (DiVA):

<http://urn.kb.se/resolve?urn=urn:nbn:se:liu:diva-174679>

N.B.: When citing this work, cite the original publication.

Boström-Rost, P., Axehill, D., Hendeby, G., (2021), Sensor management for search and track using the Poisson multi-Bernoulli mixture filter, *IEEE Transactions on Aerospace and Electronic Systems*.

<https://doi.org/10.1109/TAES.2021.3061802>

Original publication available at:

<https://doi.org/10.1109/TAES.2021.3061802>

Copyright: Institute of Electrical and Electronics Engineers

<http://www.ieee.org/index.html>

Sensor management for search and track using the Poisson multi-Bernoulli mixture filter

Per Boström-Rost, Daniel Axehill, *Senior Member, IEEE*, Gustaf Hendeby, *Senior Member, IEEE*
Dept. of Electrical Engineering, Linköping University
 Linköping, Sweden

Abstract—A sensor management method for joint multi-target search and track problems is proposed, where a single user-defined parameter allows for a trade-off between the two objectives. The multi-target density is propagated using the Poisson multi-Bernoulli mixture filter, which eliminates the need for a separate handling of undiscovered targets and provides the theoretical foundation for a unified search and track method. Monte Carlo simulations of two scenarios are used to evaluate the performance of the proposed method.

Index Terms—Sensor management, search and track, Poisson multi-Bernoulli mixture filter, multi-target tracking, informative path planning, receding horizon control, Monte Carlo tree search

I. INTRODUCTION

This paper considers sensor management for multi-target tracking. The problem is to control a sensor such that it maintains accurate estimates of the states of the discovered targets and at the same time searches for targets that are still undiscovered. Complicating factors for the problem are, apart from the sensor's limited field of view and sensing range, that the number of targets is both unknown and time-varying, the measurements are corrupted by noise, there are misdetections, false alarms, and unknown measurement origins.

Classical approaches to multi-target tracking include multi-hypothesis tracking (MHT) [1] and the joint probabilistic data association (JPDA) filter [2]. In recent years, much attention has been focused on methods based on the theory of random finite sets (RFSs) [3, 4], which enables a Bayesian approach to the problem of multi-target tracking. In general, the exact multi-target posterior is intractable to compute, but several tractable algorithms have been developed based on different approximations. Examples of these include the probability hypothesis density (PHD) filter [5], the generalized labeled multi-Bernoulli (GLMB) filter [6, 7], and the Poisson multi-Bernoulli mixture (PMBM) filter [8]. In contrast to the GLMB filter, the PMBM filter operates on unlabeled sets of target states and does not formally provide track continuity. An alternative method is the PMBM tracker [9], which estimates the posterior density over the set of target trajectories and thereby provides track continuity without introducing labels.

Sensor management, *i.e.*, optimization of control inputs for sensors, becomes interesting in situations where the state of

the sensor affects the quality of the measurements. The RFS framework has been used in several sensor management applications with focus on improving the estimation performance of discovered targets [10–13]. The PHD representation has also been used to plan sensor trajectories with the aim of discovering previously undetected targets [14].

Both planning for improving the tracking performance and planning the search for new targets are arguably difficult problems. Even more challenging is the problem of planning for jointly maintaining track of discovered targets and searching for undiscovered targets, as the two objectives are often competing for the same sensor resources. In [15], the PHD filter is used in combination with a Voronoi-based control method to enable a team of mobile sensors to concurrently search for and track an unknown number of targets. The same problem is studied in [16, 17], where information-theoretic objective functions are used to prioritize between searching and tracking. Whereas [15] considers only a single time step, the methods in [16, 17] are non-myopic and optimize sensor control inputs over a horizon of multiple time steps. Both [16, 17] assume known measurement origins, track each detected target separately, and make use of an occupancy grid filter to represent presence of undetected targets.

In this paper, a new non-myopic sensor management approach based on the PMBM filter¹ is proposed for multi-target tracking. The PMBM density, which has been shown to be a conjugate prior for the multi-target Bayes filter [8], allows for an elegant separation of the set of targets into two disjoint subsets: targets that have been detected, and targets that are still undiscovered. This provides a solid theoretical foundation for a unified planning algorithm for search and track. The explicit modeling of undetected targets in the PMBM filter is exploited to trade-off between searching for new targets and tracking detected targets. Furthermore, the underlying PMBM filter eliminates the need of a separate handling of undiscovered targets and known measurement origins as in [16, 17].

A. Notation

Vectors and scalars are described by lower-case letters x , matrices and sets by upper-case letters X , index sets by blackboard bold letters $\mathbb{X}$. The cardinality of a set X is

This work was partially supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation; the Industry Excellence Center LINKSIC funded by The Swedish Governmental Agency for Innovation Systems (VINNOVA); and Saab AB.

¹The proposed method relies on the PMBM filter, but could with minor adaptations instead be based on the PMBM tracker. However, the benefit of this from a practical planning point of view is limited.

denoted $|X|$. The operator $\uplus$ denotes the disjoint set union, i.e., $X^1 \uplus X^2 = X$ means $X^1 \cup X^2 = X$ and $X^1 \cap X^2 = \emptyset$. The inner product of $a(x)$ and $b(x)$ is $\langle a, b \rangle = \int a(x)b(x) dx$.

B. Outline

In Section II, the general sensor management problem for multi-target tracking is presented. Section III briefly reviews RFS-based target tracking and gives an introduction to the PMBM filter. The details of the proposed PMBM-based sensor management method are presented in Section IV, followed by an approach to solve the planning problem using Monte Carlo tree search in Section V. Results from simulation studies are presented in Section VI and Section VII concludes the paper.

II. PROBLEM FORMULATION

The problem of finding the optimal sequence of control inputs for a sensor can be formulated as a stochastic optimal control problem [18]. A problem of this type is often referred to as a partially observable Markov decision process (POMDP) [19, 20], which provides a framework for planning under uncertainty. This section outlines general models of the system dynamics and measurements. These models are used to formulate the problem of sensor management for multi-target tracking as a stochastic optimal control problem.

A. Dynamics and measurement models

Consider a controllable sensor of which the dynamics are governed by the model

$$s_{k+1} = g(s_k, u_k), \quad (1)$$

where s_k is the state of the sensor and $u_k \in \mathcal{U}$ is the control input applied at time k , and $\mathcal{U}$ is the set of admissible control inputs. The state of the sensor is assumed to be perfectly known at all times.

The models of the multi-target dynamics and measurements presented here are inspired by [10]. Let X_k denote the set of targets at time k . The dynamics of the multi-target state X_k is defined as follows. As time progresses from k to $k+1$, each individual target $x_k \in X_k$ generates a new set $S_{k+1|k}(x_k)$, which is empty if the target terminates and contains a single element x_{k+1} if the target survives. By representing the set of newborn targets with B_{k+1} , the multi-target state transition equation is given by

$$X_{k+1} = f(X_k) = \bigcup_{x_k \in X_k} S_{k+1|k}(x_k) \cup B_{k+1}. \quad (2)$$

Note that the multi-target set X_k includes *both* the detected targets and targets that have not yet been detected by the sensor.

The multi-target measurement model is defined as follows. At time k , each target $x_k \in X_k$ generates a measurement set $D_k(x_k, s_k)$. This set contains a single element if the target is detected by the sensor with state s_k , and is empty if the target is undetected. By denoting the set of false alarms by $F_k(s_k)$, the full set of measurements Z_k generated by X_k is given by the multi-target measurement equation

$$Z_k = h(X_k, s_k) = \bigcup_{x_k \in X_k} D_k(x_k, s_k) \cup F_k(s_k). \quad (3)$$

B. Stochastic optimal control problem

The state of the multi-target set X_k is estimated using the measurements obtained by the sensor system. Let $\pi_{k|k} = \pi_{k|k}(X_k | Z_{1:k}, s_{1:k})$ denote the posterior multi-target state density at time k after observing the measurement sets $Z_{1:k}$ obtained at sensor states $s_{1:k}$. Here, $Z_{1:k}$ is the sequence of measurement sets up until and including time k and $s_{1:k}$ is defined analogously for the sensor state. Furthermore, let $f_{k|k-1}(\cdot)$ and $f_{k|k}(\cdot)$ denote the prediction and update steps of the multi-target Bayes filter [3], which will be described in Section III-B.

The considered sensor management problem can then be formulated as follows. Given the current sensor state s_k , the multi-target density $\pi_{k|k}$, and a planning horizon N , an optimal sequence of sensor control inputs $u_{k:k+N-1}$ is obtained by solving the following stochastic optimal control problem:

$$\begin{aligned} \underset{u_{k:k+N-1}}{\text{minimize}} \quad & \mathbf{E}_{Z_{k+1:k+N}} \left[\sum_{t=k+1}^{k+N} \ell(\pi_{t|t}) \right] \\ \text{subject to} \quad & u_t \in \mathcal{U}, \\ & s_{t+1} = g(s_t, u_t), \\ & \pi_{t+1|t+1} = f_{k|k}(f_{k|k-1}(\pi_{t|t}), Z_{t+1}, s_{t+1}), \\ & X_{t+1} = f(X_t), \\ & Z_t = h(X_t, s_t), \end{aligned} \quad (4)$$

where $\ell(\cdot)$ is the stage cost function, which is used to trade-off between tracking discovered targets and searching for new targets. The argument of the stage cost function is the resulting multi-target density at the corresponding time step, i.e., $\pi_{t|t}$. Since this density depends on the measurement sets $Z_{k+1:t}$, which at the time of planning are not yet obtained, the expectation of the objective function is taken with respect to the future measurement sets $Z_{k+1:k+N}$. This approach is commonly used in sensor management applications [10, 12, 21].

Since the targets maneuver and the scenario changes over time, it is necessary to re-plan the sequence of sensor control inputs online as new measurements are obtained. This is done in standard receding horizon fashion [22]:

- 1) At time k , update the multi-target state density to $\pi_{k|k}$ using the set of measurements Z_k of the set of targets X_k , obtained by the sensor system with state s_k .
- 2) Compute a sequence of control inputs $u_{k:k+N-1}$ by solving (4).
- 3) Apply the first element of the computed sequence of control inputs, u_k , to the sensor system, and discard the remaining elements. At time $k+1$, the process is repeated from step 1.

III. BACKGROUND ON MULTI-TARGET FILTERING

The problem formulated in the previous section requires the estimation of a multi-target state density, for which methods based on RFS theory can be applied. This section provides an overview of RFS models, multi-target filtering using RFS, and the PMBM filter.

A. Random finite set models

A *random finite set* is a set with a random number of elements which are themselves random [3]. This means that the cardinality of an RFS is a discrete random variable, and its elements are random variables. This makes the RFS framework convenient for representing sets of multiple targets and multiple sensor measurements. In this work, the following RFSs will be of interest:

1) *Poisson point process*: A Poisson point process (PPP) is an RFS of which the cardinality is Poisson distributed with rate μ and the elements, given the cardinality, are independent and identically distributed (IID) according to $p(x)$. The rate μ and distribution $p(x)$ form the intensity $\lambda(x)$ of the PPP as $\lambda(x) = \mu p(x)$. The density of a PPP X is given by [23]:

$$\pi(X) = e^{-\langle \lambda, 1 \rangle} \prod_{x \in X} \lambda(x). \quad (5)$$

where the notation $\langle a, b \rangle = \int a(x)b(x)dx$ is used for the inner product of $a(x)$ and $b(x)$. In this work, PPPs are used to model clutter measurements, target birth, and undetected targets.

2) *Bernoulli RFS*: The cardinality of a Bernoulli RFS is Bernoulli distributed with parameter $r \in [0, 1]$. It is either empty, with probability $1 - r$, or, with probability r , contains a single element with density $p(x)$. Thus, the density of a Bernoulli RFS X is [23]:

$$\pi(X) = \begin{cases} 1 - r, & X = \emptyset, \\ rp(x), & X = \{x\}, \\ 0, & |X| \geq 2. \end{cases} \quad (6)$$

As the Bernoulli RFS captures both the uncertainty regarding a target's existence (through the parameter r) and state (through the density $p(x)$), the Bernoulli RFS provides a convenient representation for single targets.

3) *Multi-Bernoulli RFS*: The disjoint union of a fixed number of Bernoulli RFSs is a multi-Bernoulli (MB) RFS. Its density is defined by the parameters $\{r^i, p^i\}_{i \in \mathbb{I}}$, where $\mathbb{I}$ is an index set:

$$\pi(X) = \begin{cases} \sum_{\sqcup_{i \in \mathbb{I}} X^i = X} \prod_{i \in \mathbb{I}} \pi^i(X^i), & |X| \leq |\mathbb{I}|, \\ 0, & |X| > |\mathbb{I}|. \end{cases} \quad (7)$$

The notation $X^1 \sqcup X^2 = X$ denotes disjoint union, i.e., $X^1 \cup X^2 = X$ and $X^1 \cap X^2 = \emptyset$.

4) *Multi-Bernoulli mixture RFS*: A normalized, weighted sum of MB RFSs is referred to as a multi-Bernoulli mixture (MBM) RFS. Its density can be stated as [24]:

$$\pi(X) = \sum_{j \in \mathbb{J}} w^j \sum_{\sqcup_{i \in \mathbb{I}^j} X^i = X} \prod_{i \in \mathbb{I}^j} \pi^{j,i}(X^i) \quad (8)$$

and is defined by the set of parameters $\{w^j, \{r^{j,i}, p^{j,i}\}_{i \in \mathbb{I}^j}\}_{j \in \mathbb{J}}$, where $\mathbb{J}$ is an index set for the MB components of the MBM, $\mathbb{I}^j$ is an index set for the Bernoullis of the j th MB RFS, and w^j is the weight of the j th MB.

B. Multi-target filtering using RFS theory

In the RFS-based filtering approach, the multi-target state and set of measurements at time k are modeled as two RFSs denoted X_k and Z_k , respectively. The aim is to estimate the posterior multi-target state density $\pi_{k|k}(X_k | Z_{1:k}, s_{1:k})$.

Similarly as in the standard single-target case, the multi-target posterior density can be computed recursively via prediction and measurement update steps. With the Bayes multi-target filter [3], the posterior multi-target density at time $k-1$ is propagated in time using the prediction step:

$$\pi_{k|k-1}(X_k | Z_{1:k-1}, s_{1:k-1}) = \int \phi_X(X_k | X_{k-1}) \pi_{k-1|k-1}(X_{k-1} | Z_{1:k-1}, s_{1:k-1}) \delta X_{k-1}, \quad (9)$$

where $\phi_X(X_k | X_{k-1})$ is the multi-target transition density. The standard multi-target transition model [3] is used, which means that $\phi_X(X_k | X_{k-1})$ models a Markovian process for individual targets with transition density $p(x_k | x_{k-1})$ and state-dependent probability of survival $p_s(x)$, combined with a PPP birth process with intensity $\lambda^b(x)$.

Given a measurement set Z_k obtained using a sensor with state s_k , the predicted density is updated using the update step:

$$\pi_{k|k}(X_k | Z_{1:k}, s_{1:k}) = \frac{\phi_Z(Z_k | X_k, s_k) \pi_{k|k-1}(X_k | Z_{1:k-1}, s_{1:k-1})}{\int \phi_Z(Z_k | X_k, s_k) \pi_{k|k-1}(X_k | Z_{1:k-1}, s_{1:k-1}) \delta X_k}, \quad (10)$$

where $\phi_Z(Z_k | X_k, s_k)$ is the sensor-dependent multi-target measurement set likelihood function. The standard multi-target measurement model [3] is used, which means that $\phi_Z(Z_k | X_k, s_k)$ models noisy measurements of individual targets with state-dependent probability of detection $p_D(x | s)$, combined with PPP clutter with intensity $\lambda^{fa}(z | s)$. At most one measurement is generated for each target in each time step, and each measurement is the result of at most one target. Each measurement is independent of all other targets and other measurements conditioned on the same target; the single target measurement likelihood is $p(z | x, s)$. The integrals in (9) and (10) are set integrals, defined in [3].

C. Poisson Multi-Bernoulli mixture filter

The RFS-based filter used in this work is the PMBM filter, which is derived in [8, 25]. This section gives an overview of the key concepts in the PMBM filter based on [24, 25].

1) *Interpretation*: In the PMBM filter, the set of targets X is divided into two disjoint subsets $X^u \sqcup X^d = X$, corresponding to the undetected targets X^u and the targets that have been detected at least once X^d . The distribution of the targets is modeled by the union of a PPP component and an MBM component.

The MBM component models the potentially detected targets and considers all possible data association hypotheses. Each measurement at each time step may be either the first detection of a new target, a new measurement of a target that has previously been detected, or a false alarm. Each measurement thus results in a new potentially detected target which distribution is Bernoulli, as it might not exist. For each

Fig. 1. The intensity of the Poisson component in the PMBM density indicates parts of the state space where undetected targets are likely to be found. This can be exploited when planning the search for new targets.

potentially detected target, there are single target hypotheses which represent sequences of target-to-measurement associations. A global hypothesis, represented by an MB component in the MBM, consists of a number of single target hypotheses from different potentially detected targets. The unknown data associations lead to an intractably large number of components in the MBM density, and approximations are necessary for tractability, *e.g.*, using a track-oriented MHT formulation [25].

An interesting unique feature of the PMBM filter is the explicit modeling of targets that are not yet detected. The distribution of these targets are represented by the PPP component. This enables the user to incorporate prior knowledge about where new targets are likely to appear and behave before they are detected by the sensor. An example of how the intensity of undetected targets might be distributed is given in Fig. 1. This explicit modeling of undetected targets can also be used in sensor management applications in order to guide the sensor towards areas where new targets are likely to be discovered.

2) *Mathematical representation*: The PMBM density is a combination of a PPP density π^P on the set of undetected targets X^u and an MBM density π^{MBM} on the set of detected targets X^d . The full multi-target density is given by

$$\pi^{\text{PMBM}}(X) = \sum_{X^u \uplus X^d = X} \pi^P(X^u) \pi^{\text{MBM}}(X^d), \quad (11)$$

where $\pi^P(X^u)$ is given by (5) and $\pi^{\text{MBM}}(X^d)$ is given by (8). The full PMBM density is completely defined by the following set of parameters:

$$\lambda^u, \{w^j, \{r^{j,i}, p^{j,i}\}_{i \in \mathbb{I}^j}\}_{j \in \mathbb{J}}, \quad (12)$$

where λ^u is the intensity of undetected targets, $\mathbb{J}$ is an index set for the global hypotheses, *i.e.*, the components of the MBM. The j th MB component has weight w^j and density π^j and consists of $|\mathbb{I}^j|$ Bernoulli components. The Bernoulli density of the i th potentially detected target under the j th global hypothesis is defined by the probability of existence $r^{j,i}$ and spatial density $p^{j,i}$.

The PMBM density has been shown to be a multi-target tracking conjugate prior [8], which means that the prediction (9) and measurement update (10) steps correspond to

computing new values for the set of parameters in (12). The equations for the update steps are outlined in Appendix A. For derivations and more details, the reader is referred to [8, 25].

There are several ways to extract a set of estimated target states from the PMBM density [25]. The method used in this paper is based on the global hypothesis with the highest weight. Given the set of PMBM parameters (12), the index of the global hypothesis with the highest weight is obtained as

$$j^* = \arg \max_{j \in \mathbb{J}} w^j. \quad (13)$$

From the corresponding MB component with parameters $\{r^{j^*,i}, p^{j^*,i}\}_{i \in \mathbb{I}^{j^*}}$, all Bernoulli components with probability of existence larger than a threshold τ are selected, and the expected values of the target states are included in the set of target estimates $\hat{X}$, *i.e.*,

$$\hat{X} = \{\hat{x}^{j^*,i}\}_{i \in \mathbb{I}^{j^*}: r^{j^*,i} > \tau} \quad (14a)$$

$$\hat{x}^{j^*,i} = \mathbf{E}[p^{j^*,i}]. \quad (14b)$$

IV. PMBM-BASED SENSOR MANAGEMENT

This section formulates a tractable approximation of the original problem (4) based on the PMBM filter that allows for a trade-off between tracking and searching. The planning for tracking of discovered targets is based on the best hypothesis in the MBM component of the PMBM density, whereas the intensity of the Poisson component is used to guide the sensor towards regions where the probability of discovering undetected targets is high.

A. Assumptions

We begin by introducing a number of simplifying assumptions. First of all, it is assumed that the multi-target posterior $\pi_{k|k}$ is described by a PMBM density defined by the parameters in (12). The detection and survival probabilities are approximated as

$$p_D(x | s) \approx p_D(\mathbf{E}[x] | s), \quad (15a)$$

$$p_S(x) \approx p_S(\mathbf{E}[x]), \quad (15b)$$

and the birth intensity λ^b is modeled as a Gaussian mixture,

$$\lambda^b(x) = \sum_{i=1}^{N^b} w^{b,i} \mathcal{N}(x; x^{b,i}, P^{b,i}). \quad (16)$$

Then, by using the prediction and measurement update steps of the extended Kalman filter (EKF) [26] to compute the predictions and measurement updates of the PMBM filter, the intensity of the undetected targets λ^u is a Gaussian mixture and the Bernoulli densities $p^{j,i}$ are Gaussian:

$$\lambda^u(x) = \sum_{i=1}^{N^u} w^{u,i} \mathcal{N}(x; x^{u,i}, P^{u,i}), \quad (17a)$$

$$p^{j,i}(x) = \mathcal{N}(x; \hat{x}^{j,i}, P^{j,i}). \quad (17b)$$

Furthermore, it is assumed that the set of admissible control inputs $\mathcal{U}$ is a finite set. With these assumptions in place, we are ready to discuss planning for tracking and searching.

B. Tracking discovered targets

From the tracking perspective, the objective of the planning problem is to reduce the uncertainty in the state estimates of the discovered targets. Each admissible sequence of control inputs results in a unique sensor state trajectory, along which corresponding measurement sets can be predicted. Using these measurement sets together with the prediction and update steps of the PMBM filter, it is possible to predict the evolution of the state estimates. These predictions can be used to compare future uncertainties in the state estimates after applying different sequences of control inputs.

1) *Approximations to reduce the computational cost:* The predicted ideal measurement set approach [27] is employed to avoid the computationally expensive expectation over the future measurement sets in (4). This means that the planning problem is solved under the assumption that the future measurement sets Z_t , for $t = k + 1, \dots, k + N$, are generated under ideal conditions, *i.e.*, every target within field of view (FoV) generates a noise-free measurement and there are no false alarms. Thus, given sensor state s_t and a set of predicted target states $\hat{X}_t^{\tau_p}$, the predicted ideal measurement set $\hat{Z}_t$ is computed as

$$\hat{Z}_t = \bigcup_{\hat{x} \in \hat{X}_t^{\tau_p} : \hat{x} \in \text{FoV}(s_t)} \left\{ \arg \max_z p(z | \hat{x}, s_t) \right\}. \quad (18)$$

The set of predicted target states $\hat{X}_t^{\tau_p}$ is extracted from the best hypothesis in the MBM component of the PMBM density. The index of the MB corresponding to this hypothesis is denoted j^* and is found using (13). The extraction is similar to the estimation step of the PMBM filter (14), but with a different threshold on the probability of existence $\tau_p < \tau$:

$$\hat{X}_t^{\tau_p} = \{ \hat{x}_{t|k}^{j^*, i} \}_{i \in \mathbb{I}^{j^*} : r_{k|k}^{j^*, i} > \tau_p}. \quad (19)$$

Newly detected targets typically have a low probability of existence and the lower threshold τ_p allows for including measurements of these targets in the predicted ideal measurement set. This is a way to encourage the sensor to confirm or reject new possible tracks.

Note that the predicted ideal measurement sets are generated with known measurement origins. This can be exploited to decrease the computational cost related to solving the planning problem. As the predicted measurements match hypothesis j^* exactly except for targets with probability of existence lower than τ_p , this hypothesis will with high probability remain as the most likely hypothesis after updating the PMBM density using these measurements. Thus, with little risk of introducing errors, no other hypothesis than the one with index j^* need to be updated and no new association hypotheses need to be introduced.

For notational compactness, we define the function $\hat{h}(\cdot)$ and write the predicted ideal measurement set as

$$\hat{Z}_t = \hat{h}(s_t) \quad (20)$$

where the dependency on $\hat{X}_t^{\tau_p}$ is implicit.

2) *Stage cost:* As estimates of individual target states are Gaussian distributed, a measure on the corresponding covariance matrices gives an indication of the tracking performance. Several scalar measures of covariance matrices exist, see *e.g.* [28]. Here, the state estimation accuracy is measured by summing the trace of the covariance matrix of each track:

$$\ell^{\text{track}}(\pi_{t|t}) = \sum_{i \in \mathbb{I}^{j^*}} \text{tr}(W P_{t|t}^{j^*, i} W^\top), \quad (21)$$

where W is a matrix of appropriate dimensions used to extract and weigh the relevant components of the covariance matrix. The variable $P_{t|t}^{j^*, i}$ is the covariance matrix of track i in hypothesis j^* at time t given the actual measurements up until time k and predicted ideal measurements from time $k + 1$ to t .

C. Searching for new targets

The intensity of the Poisson component in the PMBM density provides a convenient representation of where undetected targets are likely to be found. This is illustrated in Fig. 1. Areas with high birth intensity and parts of the surveillance volume that have not been visited by the sensor in a long time will have a higher intensity of undetected targets.

The intensity λ^u is a deterministic function of the sensor state trajectory that can be computed using the prediction and update steps of the PMBM filter. Thus, the resulting intensity of undetected targets after applying different sequences of control inputs can be predicted and evaluated based on a suitable performance measure.

1) *Stage cost:* As the objective of the planning problem from the search perspective is to discover new targets, we propose to use the expected number of undetected targets as stage cost. This is given by the integral of the intensity of undetected targets, which is equal to the sum of the weights in the Gaussian mixture:

$$\ell^{\text{search}}(\pi_{t|t}) = \int \lambda_{t|t}^u(x) dx = \sum_{i=1}^{N_t^u} w_{t|t}^{u, i}. \quad (22)$$

Here, $\lambda_{t|t}^u$ is the resulting intensity of undetected targets at time t given the actual sensor trajectory up until time k and planned trajectory from time $k + 1$ to t , and N_t^u is the number of components in the Gaussian mixture at time t .

D. Resulting optimal control problem

For the full planning problem, the proposed stage cost function is a weighted sum of $\ell^{\text{track}}(\cdot)$ and $\ell^{\text{search}}(\cdot)$, *i.e.*

$$\ell(\pi_{t|t}) = \ell^{\text{track}}(\pi_{t|t}) + \eta \ell^{\text{search}}(\pi_{t|t}), \quad (23)$$

where η is a user-defined weight used to trade-off between tracking and searching. The parameter η can be considered as the cost of not detecting a target, as $\ell^{\text{search}}(\cdot)$ corresponds to the expected number of undetected targets.

Fig. 2. Overview of the Monte Carlo tree search method. Figure adapted from [29].

The resulting deterministic approximation of the original problem (4), which is solved when planning the sensor control input at time k , is formulated as follows:

$$\begin{aligned}
 & \underset{u_{k:k+N-1}}{\text{minimize}} && \sum_{t=k+1}^{k+N} \ell(\pi_t|t) \\
 & \text{subject to} && u_t \in \mathcal{U}, \\
 & && s_{t+1} = g(s_t, u_t), \\
 & && \pi_{t+1|t+1} = \rho_{k|k}(\rho_{k|k-1}(\pi_t|t), \hat{Z}_{t+1}, s_{t+1}), \\
 & && \hat{Z}_t = \hat{h}(s_t)
 \end{aligned} \tag{24}$$

where $\rho_{k+1|k}(\cdot)$ and $\rho_{k|k}(\cdot)$ represent the prediction and measurement update steps of the PMBM filter, respectively.

If $\mathcal{U}$ is a finite set of admissible controls, the solution to (24) can in theory be obtained by explicitly enumerating all possible control sequences $u_{k:k+N-1}$ and selecting the one that results in the smallest objective function value. However, the computational complexity of this method is prohibitive for many problems. For larger problems, we propose to use a Monte Carlo tree search (MCTS) method [29] that incrementally builds a search tree and allocates its search time to the most promising parts of the tree. Such a method is described next.

V. MONTE CARLO TREE SEARCH

Standard Monte Carlo tree search (MCTS) [29] is a stochastic method for finding the optimal actions in search problems. This is done by iteratively building a search tree until some predefined computational budget is reached, after which the best control input at the root node is returned. Each iteration consists of four phases: selection, expansion, simulation, and backpropagation [29]. These phases are often grouped into a *tree policy* (selection and expansion) and a *rollout policy* (simulation). The backpropagation step is used to update node statistics that affect the tree policy decisions in future iterations. Fig. 2 illustrates one iteration of the MCTS method and shows how the search tree is iteratively constructed.

The MCTS method tailored for PMBM-based sensor management is outlined in Algorithm 1. Each node n in the tree holds information about how many times the node has been visited so far during the search $c_v(n)$, the corresponding sensor state and PMBM density, and the average total cost $\hat{J}(n)$ of control input sequences that pass through the node. An edge in the search tree represents a control input applied to the sensor.

In the selection phase (Algorithm 1, line 3), a node n that still has unvisited children is selected for expansion. This is

done by starting at the root node and recursively selecting child nodes until a node with at least one unvisited child is reached. Child nodes are selected using the upper-confidence bound for trees (UCT) policy [30], i.e., a child node n' of node n is selected as

$$\arg \max_{n' \in \text{children of } n} \left\{ -\hat{J}(n') + \varepsilon \sqrt{\frac{\log c_v(n')}{c_v(n)}} \right\}, \tag{25}$$

where the second term encourages selection of control inputs that have rarely been tried, and ε is an exploration parameter to trade-off between exploration and exploitation.

In the expansion phase (Algorithm 1, line 4), a new child node n' is added to the node selected for expansion n . This corresponds to selecting a previously untried control input, updating the state of the sensor, generating the predicted ideal measurement set, and predicting and updating the PMBM density.

A simulation (Algorithm 1, line 5) is then run from the new node n' until the planning horizon, computing the stage costs along the way. This produces an estimate Δ of the total cost of a control input sequence that starts with the control inputs leading up to n' . During the simulation phase, the control inputs are selected according to a rollout policy, which could correspond to selecting control inputs at random or based on heuristic for the specific problem [31]. In the backpropagation phase (Algorithm 1, line 6), the result Δ is used to update the statistics of the nodes selected for this iteration. Each node's visit count is incremented and its average total cost updated according to Δ .

VI. SIMULATION STUDY

In this section, the proposed PMBM-based sensor management method is evaluated using a series of simulated scenarios. The achieved performance is compared for three different stage cost functions: (i) the proposed stage cost function (23) for search and track (ST), (ii) the stage cost function (21) for track only (TO), and (iii) the stage cost function (22) for search only (SO). The PMBM-based method is also compared to the grid-based method proposed in [17]. The method from [17] assumes known measurement origins, which simplifies the problem as the true measurement-to-target associations are given. A Gaussian mixture implementation of the Bernoulli filter [32] is used as underlying tracker for each individual target, and an occupancy grid filter is used to represent undetected targets.

A. Performance evaluation

The generalized optimal subpattern assignment (GOSPA) metric [33] is a unified performance metric for multi-target tracking. It penalizes both localization errors for detected targets and errors due to missed and false targets. According to [33], with $\alpha = 2$ and given $c > 0$, $1 \leq p < \infty$, a metric $d(\cdot)$ in the single target space, the ground truth set

Algorithm 1 Overview of MCTS for PMBM-based sensor management

```

1: Initialize tree  $\mathcal{T}$  with root node  $n_0$  with current sensor state and PMBM density
2: for  $N_{\text{iter}}$  iterations do
3:    $n \leftarrow \text{SELECTNODE}(\mathcal{T})$  ▷ Select node to expand using UCT policy
4:    $n' \leftarrow \text{EXPANDTREE}(n)$  ▷ Add new child to node  $n$ ; simulate measurements and update PMBM density
5:    $\Delta \leftarrow \text{SIMULATE}(n')$  ▷ Perform rollout until planning horizon; compute total cost
6:    $\mathcal{T} \leftarrow \text{BACKPROPAGATE}(\mathcal{T}, n', \Delta)$  ▷ Update node statistics
7: end for
8: return  $\arg \min_{n \in \text{children of } n_0} \hat{J}(n)$ 

```

$X = \{x^1, \dots, x^{|X|}\}$ and its estimate $\hat{X} = \{\hat{x}^1, \dots, \hat{x}^{|\hat{X}|}\}$, the GOSPA metric can be written as

$$d_p^{(c,2)}(X, \hat{X}) = \min_{\gamma \in \Gamma} \left(\sum_{(i,j) \in \gamma} d(x^i, \hat{x}^j)^p + \frac{c^p}{2} (|X| + |\hat{X}| - 2|\gamma|) \right)^{\frac{1}{p}}, \quad (26)$$

where Γ is the set of all possible assignment sets γ between $\{1, \dots, |X|\}$ and $\{1, \dots, |\hat{X}|\}$. Such an assignment set satisfies $\gamma \subseteq \{1, \dots, |X|\} \times \{1, \dots, |\hat{X}|\}$, $(i, j), (i, j') \in \gamma \implies j = j'$ and $(i, j), (i', j) \in \gamma \implies i = i'$, where the last two properties ensure that every i and j gets at most one assignment [33]. Note that $|\gamma|$ is the number of properly detected targets and that the definition implies that $|\gamma| \leq \min\{|X|, |\hat{X}|\}$.

Given the optimal assignment set γ^* , the GOSPA metric can be decomposed as [33]:

$$d_p^{(c,2)}(X_k, \hat{X}_k) = (d^{\text{loc}}(X_k, \hat{X}_k, \gamma^*) + d^{\text{miss}}(X_k, \gamma^*) + d^{\text{false}}(\hat{X}_k, \gamma^*))^{\frac{1}{p}}, \quad (27)$$

where $d^{\text{loc}}(\cdot)$ is the cost related to the localization of properly assigned targets, and $d^{\text{miss}}(\cdot)$ and $d^{\text{false}}(\cdot)$ are the costs related to missed and false targets. The three terms in the decomposed metric are computed as

$$d^{\text{loc}}(X, \hat{X}, \gamma^*) = \sum_{(i,j) \in \gamma^*} d(x^i, \hat{x}^j)^p, \quad (28a)$$

$$d^{\text{miss}}(X, \gamma^*) = \frac{c^p}{2} (|X| - |\gamma^*|), \quad (28b)$$

$$d^{\text{false}}(\hat{X}, \gamma^*) = \frac{c^p}{2} (|\hat{X}| - |\gamma^*|), \quad (28c)$$

where $|X| - |\gamma^*|$ and $|\hat{X}| - |\gamma^*|$ are the number of missed and false targets, respectively. Note that the choice of the maximum allowable localization cost c defines the cost of not detecting a target as $c^p/2$, which can be used to find an appropriate value for the weight parameter η in the objective function (23).

To evaluate the overall performance in the simulation study, GOSPA is used as the main metric. For the parameters, we use $c = 50$, $p = 2$, and define the base distance $d(x, y)$ as the 2-norm of the position components of $x - y$. We also report (i) the normalized localization error

$$d^{\text{nle}}(X, \hat{X}, \gamma^*) = \left(\frac{1}{|\gamma^*|} d^{\text{loc}}(X, \hat{X}, \gamma^*) \right)^{\frac{1}{p}}, \quad (29)$$

Fig. 3. Example of simulated scenario. There are a total of five targets appearing every 80 s. Two targets disappear at time steps 280 and 360, respectively. The targets' initial positions are marked with filled circles. The field of view is orange and the light blue circles illustrate the Gaussian mixture intensity of the undetected targets.

i.e., the average localization error of properly assigned targets as a localization performance measure, (ii) the number of missed targets as a target discovery performance measure, (iii) the number of false targets, and (iv) the number of lost targets.

B. Beam pointing for a stationary sensor

In this scenario, the pointing direction of a stationary range-bearing sensor is controlled to track a varying number of moving targets. The sensor is located at the origin and the surveillance area is defined by the half disc with radius 1 km. A new target with initial position uniformly sampled in the surveillance area appears every 80 s. One target disappears at time 280, and another one at time 360. The full scenario runs for 400 s. One realization is illustrated in Fig. 3.

The state of an individual target $x = [\tilde{x}^T, \omega]^T$ consists of its two-dimensional position and velocity $\tilde{x} = [p_1, v_1, p_2, v_2]^T$ and turn rate ω . Each target follows a coordinated turn model [34], $p(x_k | x_{k-1}) = \mathcal{N}(x_k; F_{\text{CT}}(\omega_{k-1})x_{k-1}, Q)$, where the state transition matrix is given by

$$F_{\text{CT}}(\omega) = \begin{bmatrix} 1 & \frac{\sin(\omega T)}{\omega} & 0 & -\frac{1 - \cos(\omega T)}{\omega} & 0 \\ 0 & \cos(\omega T) & 0 & -\sin(\omega T) & 0 \\ 0 & \frac{1 - \cos(\omega T)}{\omega} & 1 & \frac{\sin(\omega T)}{\omega} & 0 \\ 0 & \sin(\omega T) & 0 & \cos(\omega T) & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix} \quad (30)$$

and the covariance matrix is given by

$$\begin{bmatrix} \sigma_w^2 G G^\top & 0 \\ 0 & \sigma_u^2 \end{bmatrix}, \quad G = I_2 \otimes \begin{bmatrix} T^2 \\ T \end{bmatrix}, \quad (31)$$

where I_2 is the 2×2 identity matrix and $\otimes$ is the Kronecker product. The sampling period is $T = 1$ s, $\sigma_w = 0.01$ m/s² is the standard deviation of the acceleration noise, and $\sigma_u = 0.01\pi/180$ rad/s² is the standard deviation of the turn rate noise. The survival probability is assumed constant $p_s = 0.99$ for each target.

The sensor's FoV is limited, both in angle and range. The beamwidth is $\theta_{BW} = \pi/6$ rad centered around s_k and the maximum sensing range is $r_{\max} = 1$ km. The center of the beam defines the state of the sensor which is controlled as

$$s_{k+1} = g(s_k, u_k) = u_k, \quad (32)$$

i.e., the beam can be moved instantaneously from any pointing direction to any other. The set of admissible control inputs, or pointing directions, is defined as

$$\mathcal{U} = \left\{ \frac{n\pi}{180} \text{ rad}, n \in \{0, 1, \dots, 180\} \right\}. \quad (33)$$

The probability of detection is 0.9 for targets within the sensor's FoV and zero elsewhere. A detected target produces a noisy range-bearing measurement z with likelihood $p(z|x) = \mathcal{N}(z; h(x), R)$, where

$$h(x) = \begin{bmatrix} \sqrt{p_1^2 + p_2^2} \\ \arctan\left(\frac{p_2}{p_1}\right) \end{bmatrix} \text{ and } R = \begin{bmatrix} \sigma_r^2 & 0 \\ 0 & \sigma_\theta^2 \end{bmatrix} \quad (34)$$

with $\sigma_r = 5$ m and $\sigma_\theta = \pi/180$ rad. The Poisson clutter is uniformly distributed in the field of view with 0.5 expected false alarms per time step.

The Poisson birth intensity is a Gaussian mixture tuned to approximate a uniform intensity in the surveillance region. It is given by $\lambda^b(x) = \sum_{i=1}^{N_b} w^{b,i} \mathcal{N}(x; x^{b,i}, P^{b,i})$, with $N_b = 40$, $w^{b,i} = 3.125 \cdot 10^{-4}$ (corresponding to a new target appearing every 80 s), and $P^{b,i} = \text{diag}([100, 2, 100, 2, \pi/180]^2)$. The mean of the Gaussian components are $x^{b,i} = [p_1^{b,i}, 0, p_2^{b,i}, 0, 0]^\top$, where the position components are selected as the grid points of $p \in \{-900, -700, \dots, 900\} \times \{100, 300, \dots, 900\}$ that satisfy $\|p\|_2 \leq r_{\max}$. This gives a non-normalized Gaussian mixture intensity of undetected targets, as illustrated in Fig. 3.

Since the beam can be shifted instantaneously between any pointing directions, the planning algorithm only decides one control input at a time, *i.e.*, the planning horizon is $N = 1$. This means that the planning problem (24) can be solved in reasonable time using an exhaustive search in each iteration of the receding horizon control loop. For the GOSPA calculations we use the maximum allowable localization error $c = 50$, which gives the weight parameter in the ST case $\eta = 50^2/2 = 1250$. The weight matrix W in $\ell^{\text{track}}(\cdot)$ extracts the position components of the covariance matrices.

Fig. 4 shows the performance measures at each time step averaged over 100 Monte Carlo runs. The results confirm that the proposed stage cost function ST achieves better overall performance than SO and TO in terms of GOSPA metric. The TO case gives accurate state estimates of the detected

Fig. 4. Comparison of performance measures in the beam pointing scenario for the cases of planning based on search and track (ST), search only (SO), track only (TO), and the method proposed in [17].

targets, as indicated by the normalized localization error, but at the expense of a higher number of missed targets. The ST case performs better than the SO case in terms of missed targets, which might be surprising. The reason is that in the SO case, the sensor is pointed towards areas with high intensity of undetected targets which makes it more prone to lose track of targets that it has already discovered. There is no significant difference between the three cases in terms of the number of false targets, but TO is marginally faster than the other cases in terms of removing terminated targets. The results also indicate that ST and TO perform equally well in avoiding losing track of targets. The simulations indicate that the proposed method detects new targets faster than [17], while [17] is faster in terms of removing terminated targets. This indicates that [17] focuses more on tracking existing targets than searching for new ones, compared to ST. Both methods yield similar performance in terms of GOSPA metric.

Fig. 5. True target trajectories and initial sensor position with the field of view shown in orange. The targets' initial positions are marked with filled circles. The light blue circles illustrate the Gaussian mixture intensity of the undetected targets.

C. Trajectory planning for a mobile sensor

The second scenario runs for 3000 seconds and includes two targets in a surveillance area defined by a rectangle of $1500\text{ m} \times 750\text{ m}$. The trajectory of a mobile sensor with limited sensing range is controlled to discover and track the targets. As illustrated in Fig. 5, one of the targets is outside the sensor's sensing range at the start of the scenario and the sensor has to move away from the first target in order to discover the second one.

The state of an individual target $x = [p_1, v_1, p_2, v_2]^T$ consists of its two-dimensional position and velocity and its dynamics are modeled by a nearly constant velocity model with discrete white noise acceleration. The transition density is given by $p(x_k | x_{k-1}) = \mathcal{N}(x_k; F_{\text{cv}}x_{k-1}, Q)$, where

$$F_{\text{cv}} = I_2 \otimes \begin{bmatrix} 1 & T \\ 0 & 1 \end{bmatrix}, \quad Q = \sigma_w^2 G G^T, \quad G = I_2 \otimes \begin{bmatrix} \frac{T^2}{2} \\ T \end{bmatrix}, \quad (35)$$

where $T = 10\text{ s}$, $\sigma_w = 0.05\text{ m/s}^2$. The survival probability is assumed constant $p_s = 0.99$ for each target.

The Poisson birth intensity is given by $\lambda^b(x) = \sum_{i=1}^{N_b} w^{b,i} \mathcal{N}(x; x^{b,i}, P^{b,i})$, with $N_b = 128$, $w^{b,i} = 5.2 \cdot 10^{-5}$ (corresponding to a new target appearing every 1500 s), and $P^{b,i} = \text{diag}([50, 2, 50, 2])^2$. The mean of the Gaussian components are $x^{b,i} = [p_1^{b,i}, 0, p_2^{b,i}, 0]^T$, where the position components defines the grid $p \in \{-750, -650, \dots, 750\} \times \{50, 150, \dots, 750\}$. This gives a Gaussian mixture intensity of undetected targets, as illustrated in Fig. 5.

The sensor has a limited sensing range of $r_{\text{max}} = 150\text{ m}$. For targets within sensing range, the probability of detection is 0.9. Each detection results in a linear measurement of the position of the corresponding target. The measurement likelihood has Gaussian density $p(z | x) = \mathcal{N}(z; Hx, R)$ with parameters

$$H = I_2 \otimes [1 \ 0], \quad R = I_2. \quad (36)$$

The Poisson clutter is uniformly distributed in the field of view with five expected false alarms per time step.

The mobile sensor moves according to differential drive dynamics with constant speed of 5 m/s and the control input,

which is updated every ten seconds, determines the heading rate.

For the planning problem, the set of admissible control inputs is given by $\mathcal{U} = \{\omega \text{ rad/s}, \omega \in \{0, \pm \frac{\pi}{8}, \pm \frac{\pi}{4}\}\}$ and the planning horizon is $N = 10$ steps. This gives a total of $|\mathcal{U}|^N \approx 10^7$ possible sequences of control inputs to evaluate. Hence, the problem is too large to be solved using an exhaustive search and instead the MCTS method described in Section V is applied. The implemented MCTS search alternates between two different rollout policies. The first one selects control inputs $u \in \mathcal{U}$ at random. The second one is a heuristic policy that recursively selects the control input $u \in \mathcal{U}$ that minimizes the distance between the mobile sensor and the target with the most uncertain state estimate. A maximum allowable localization error of $c = 50\text{ m}$ is used for the GOSPA calculations, which gives the weight parameter in the ST case $\eta = 50^2/2 = 1250$. The weight matrix W in $\ell^{\text{track}}(\cdot)$ extracts the position components of the covariance matrices. Based on trial and error, the MCTS exploration parameter is $\varepsilon = 2500$ and the number of iterations is $N_{\text{iter}} = 100$.

The objective function used in [17] does not involve any tuning parameters but requires an exhaustive search to be performed. For this reason, a smaller set of admissible control sequences or maneuvers $\mathcal{U}^*$ is constructed. Each maneuver $u \in \mathcal{U}^*$ lasts for 120 s and consists of a course change selected from the set $\{n\pi/9 \text{ rad}, n \in \{-9, -8, \dots, 9\}\}$ performed with turn rate $\pi/40 \text{ rad/s}$ followed by a straight path. For a fair comparison, the proposed method for search and track is evaluated also using the smaller set of maneuvers $\mathcal{U}^*$. This combination of planning method and set of maneuvers is denoted ST*.

Fig. 6 shows the results averaged over 100 Monte Carlo runs. It is clear that ST performs better than both SO and TO in terms of GOSPA. With TO, the sensor remains close to the blue target in Fig. 5 and never discovers the orange target. Thus, the normalized localization error is low for TO and the target is not lost, but one target is never detected. SO results in more false targets than the other cases. As this case focuses on discovering new targets, the time between measurement updates for the existing tracks can be quite large. When the estimation error for a certain target becomes greater than the GOSPA cut-off distance $c = 50\text{ m}$, it is counted both as a false target and a lost target, which is seen between time 400 s and 1000 s. At around time 1000 s, the existence probability of the first target decreases to below the estimation threshold τ , after which it is removed and no longer counted as a false target. The simulation results indicate that in this scenario, the proposed method ST* performs better than [17] in terms of GOSPA and that both methods result in similar normalized localization errors.

Fig. 7 shows snapshots from one realization where the ST case is used to plan the sensor trajectory. The figures indicate that the proposed objective function yields the desired behavior: the sensor searches for new targets in areas with high intensity of undetected targets and maintains accurate state estimates of the discovered targets.

Fig. 6. Comparison of performance measures in the trajectory planning scenario for the cases of planning based on search and track (ST), search only (SO), track only (TO), the method proposed in [17], and search and track with a smaller set of admissible control sequences (ST*).

VII. CONCLUSIONS

This paper has considered the problem of jointly searching and tracking an unknown and varying number of targets. The problem was formulated as a multi-objective stochastic optimal control problem, where the conflicting objectives of search and track are competing for the same sensor resources. A method was designed to minimize both the uncertainty in the state estimates of discovered targets and the number of undiscovered targets, with a user-defined parameter determining the relative importance of the two objectives.

The PMBM filter was proposed as underlying framework when solving the optimization problem, as it provides a theoretical formulation of the multi-target tracking problem that comprises a notion of both known and undiscovered targets. As such, it not only supports the proposed method, but also provides a solid theoretical foundation for previously

proposed similar approaches which treated the two separately.

In two simulation studies, the proposed method has been illustrated and shown to provide reasonable results. The proposed method is promising for future full-scale experiments.

REFERENCES

- [1] S. Blackman and R. Popoli, *Design and Analysis of Modern Tracking Systems*. Norwood, MA, USA: Artech house, 1999.
- [2] Y. Bar-Shalom, P. K. Willett, and X. Tian, *Tracking and data fusion: A handbook of algorithms*. Storrs, CT, USA: YBS publishing, 2011.
- [3] R. Mahler, *Statistical Multisource-Multitarget Information Fusion*. Norwood, MA, USA: Artech House, 2007.
- [4] —, *Advances in statistical Multisource-Multitarget Information Fusion*. Norwood, MA, USA: Artech House, 2014.
- [5] —, “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.-T. Vo and B.-N. Vo, “Labeled random finite sets and multi-object conjugate priors,” *IEEE Transactions on Signal Processing*, vol. 61, no. 13, pp. 3460–3475, 2013.
- [7] B.-N. Vo, B.-T. Vo, and D. Phung, “Labeled random finite sets and the Bayes multi-target tracking filter,” *IEEE Transactions on Signal Processing*, vol. 62, no. 24, pp. 6554–6567, 2014.
- [8] J. L. Williams, “Marginal multi-Bernoulli filters: RFS derivation of MHT, JIPDA, and association-based MeMBer,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 51, no. 3, pp. 1664–1687, 2015.
- [9] K. Granström, L. Svensson, Y. Xia, J. Williams, and Á. F. García-Fernández, “Poisson multi-Bernoulli mixture trackers: Continuity through random finite sets of trajectories,” in *Proceedings of the 21st International Conference on Information Fusion*, Cambridge, UK, 2018.
- [10] M. Beard, B.-T. Vo, B.-N. Vo, and S. Arulampalam, “Void probabilities and Cauchy-Schwarz divergence for generalized labeled multi-Bernoulli models,” *IEEE Transactions on Signal Processing*, vol. 65, no. 19, pp. 5047–5061, 2017.
- [11] H. G. Hoang and B. T. Vo, “Sensor management for multi-target tracking via multi-Bernoulli filtering,” *Automatica*, vol. 50, no. 4, pp. 1135–1142, 2014.
- [12] B. Ristic and B.-N. Vo, “Sensor control for multi-object state-space estimation using random finite sets,” *Automatica*, vol. 46, no. 11, pp. 1812–1818, 2010.
- [13] H. Van Nguyen, H. Rezatofighi, B.-N. Vo, and D. C. Ranasinghe, “Online UAV path planning for joint detection and tracking of multiple radio-tagged objects,” *IEEE Transactions on Signal Processing*, vol. 67, no. 20, pp. 5365–5379, 2019.
- [14] J. Olofsson, G. Hendeby, T. R. Lauknes, and T. A. Johansen, “Multi-agent informed path planning using the probability hypothesis density,” *Autonomous Robots*, 2020.
- [15] P. M. Dames, “Distributed multi-target search and tracking using the PHD filter,” *Autonomous robots*, vol. 44, no. 3, pp. 673–689, 2020.
- [16] B. Charrow, N. Michael, and V. Kumar, “Active control strategies for discovering and localizing devices with range-only sensors,” in *Algorithmic Foundations of Robotics XI*, H. L. Akin, N. M. Amato, V. Isler, and A. F. van der Stappen, Eds. Springer, 2015, pp. 55–71.
- [17] H. Van Nguyen, H. Rezatofighi, B.-N. Vo, and D. C. Ranasinghe, “Multi-objective multi-agent planning for jointly discovering and tracking mobile objects,” in *Proceedings of the 34th AAAI Conference on Artificial Intelligence*, New York, NY, USA, 2020.
- [18] D. P. Bertsekas, *Dynamic programming and optimal control*. Athena Scientific, Belmont, MA, USA, 2005, vol. 1.
- [19] L. P. Kaelbling, M. L. Littman, and A. R. Cassandra, “Planning and acting in partially observable stochastic domains,” *Artificial intelligence*, vol. 101, no. 1-2, pp. 99–134, 1998.
- [20] D. A. Castañón and L. Carin, “Stochastic control theory for sensor management,” in *Foundations and applications of sensor management*, A. O. Hero, D. A. Castañón, D. Cochran, and K. Kastella, Eds. New York, NY, USA: Springer, 2008, pp. 7–32.
- [21] P. Boström-Rost, D. Axehill, and G. Hendeby, “Informative path planning for active tracking of agile targets,” in *Proceedings of IEEE Aerospace Conference*, Big Sky, MT, USA, 2019.
- [22] J. M. Maciejowski, *Predictive control: With constraints*. Englewood Cliffs, NJ, USA: Prentice Hall, 2002.
- [23] B. Ristic, B.-T. Vo, B.-N. Vo, and A. Farina, “A tutorial on Bernoulli filters: theory, implementation and applications,” *IEEE Transactions on Signal Processing*, vol. 61, no. 13, pp. 3406–3430, 2013.

Fig. 7. Snapshots of a simulated scenario where the sensor trajectory (green path) is planned using the proposed method for joint search and track. The shaded orange area is the sensor's field of view, filled gray circles are true target positions, filled orange circles are estimated target positions, dashed orange lines indicate covariances, and light blue circles represent the intensity of undetected targets.

- [24] K. Granström, M. Fatemi, and L. Svensson, "Poisson multi-Bernoulli mixture conjugate prior for multiple extended target filtering," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 56, no. 1, pp. 208–225, 2020.
- [25] Á. F. García-Fernández, J. L. Williams, K. Granström, and L. Svensson, "Poisson multi-Bernoulli mixture filter: direct derivation and implementation," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 54, no. 4, pp. 1883–1901, 2018.
- [26] A. H. Jazwinski, *Stochastic Processes and Filtering Theory*. New York, NY, USA: Academic Press, 1970.
- [27] R. Mahler, "Multitarget sensor management of dispersed mobile sensors," in *Theory and Algorithms for Cooperative Systems*, D. Grundel, R. Murphey, and P. Pardalos, Eds. Singapore: World Scientific Publishing Co, 2004, ch. 12, pp. 239–310.
- [28] C. Yang, L. Kaplan, and E. Blasch, "Performance measures of covariance and information matrices in resource management for target state estimation," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 48, no. 3, pp. 2594–2613, 2012.
- [29] C. B. Browne, E. Powley, D. Whitehouse, S. M. Lucas, P. I. Cowling, P. Rohlfshagen, S. Tavener, D. Perez, S. Samothrakis, and S. Colton, "A survey of Monte Carlo tree search methods," *IEEE Transactions on Computational Intelligence and AI in games*, vol. 4, no. 1, pp. 1–43, 2012.
- [30] L. Kocsis and C. Szepesvári, "Bandit based Monte-Carlo planning," in *European conference on machine learning*, Berlin, Germany, 2006, pp. 282–293.
- [31] S. James, G. Konidaris, and B. Rosman, "An analysis of Monte Carlo tree search," in *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*, San Francisco, CA, USA, 2017, pp. 3576–3582.
- [32] B.-N. Vo, C. M. See, N. Ma, and W. T. Ng, "Multi-sensor joint detection and tracking with the Bernoulli filter," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 48, no. 2, pp. 1385–1402, 2012.
- [33] A. S. Rahmathullah, Á. F. García-Fernández, and L. Svensson, "Generalized optimal sub-pattern assignment metric," in *Proceedings of the 20th International Conference on Information Fusion*, Xi'an, China, 2017.
- [34] X. R. Li and V. P. 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.

Per Boström-Rost is a Ph.D. student in the Division of Automatic Control, Department of Electrical Engineering, Linköping University. He received his M.Sc. degree in Applied Physics and Electrical Engineering in 2013 and his Lic.Eng. in Electrical Engineering in 2019, both from Linköping University. His main research interests are sensor fusion and sensor management with applications to target tracking.

of Automatic Control at Linköping University. His research interests are related to optimization, optimal control, motion planning, hybrid systems, and applications of control.

Gustaf Hendeby is Associate Professor in the Division of Automatic Control, Department of Electrical Engineering, Linköping University. He received his MSc in Applied Physics and Electrical Engineering in 2002 and his PhD in Automatic Control in 2008, both from Linköping University. He worked as Senior Researcher at the German Research Center for Artificial Intelligence (DFKI) 2009–2011, and Senior Scientist at Swedish Defense Research Agency (FOI) and held an adjunct Associate Professor position at Linköping University 2011–2015.

Dr. Hendeby's main research interests are stochastic signal processing and sensor fusion with applications to nonlinear problems, target tracking, and simultaneous localization and mapping (SLAM). He has experience of both theoretical analysis as well as implementation aspects.

APPENDIX A PMBM FILTER RECURSION

The PMBM filter recursion are here given in the style of [24]. The derivation of the filter equations and more details can be found in [8, 25].

A. Prediction

Given a posterior PMBM density with parameters

$$\lambda^u, \{w_+^j, \{r_+^{j,i}, p_+^{j,i}\}_{i \in \mathbb{I}^j}\}_{j \in \mathbb{J}} \quad (37)$$

and the standard dynamic model (Section III-B) with single-target transition density $p_{k+1,k}$, the predicted density is on the same form with parameters

$$\lambda_+^u, \{w_+^j, \{r_+^{j,i}, p_+^{j,i}\}_{i \in \mathbb{I}_+^j}\}_{j \in \mathbb{J}_+}, \quad (38)$$

where

$$\lambda_+^u(x) = \lambda^b(x) + \langle \lambda^u, p_s p_{k+1,k} \rangle, \quad (39a)$$

$$w_+^j = w^j, \quad (39b)$$

$$r_+^{j,i} = \langle p_+^{j,i}, p_s \rangle r^{j,i}, \quad (39c)$$

$$p_+^{j,i}(x) = \frac{\langle p_+^{j,i}, p_s p_{k+1,k} \rangle}{\langle p_+^{j,i}, p_s \rangle}, \quad (39d)$$

and $w_+^j = w^j$, $\mathbb{I}_+^j = \mathbb{I}^j$, $\mathbb{J}_+ = \mathbb{J}$.

B. Measurement update

1) *Data association*: As the true origins of measurements in are unknown, association hypotheses are required. Let $\mathbb{M}$ be an index set for the elements of the measurement set Z , i.e.,

$$Z = \{z^m\}_{m \in \mathbb{M}} \quad (40)$$

and let $\mathcal{A}^j$ be a collection of all possible association hypotheses A for the j th global hypothesis, i.e., the j th MB, of which the targets are indexed by $\mathbb{I}^j$. Then, an association hypothesis $A \in \mathcal{A}^j$ is a partition of $\mathbb{M} \cup \mathbb{I}^j$ into nonempty disjoint subsets $C \in A$, called index cells [24].

The standard assumptions in multi-target tracking that the targets are independent of each other implies that an index cell contains at most one target index and at most one measurement index, i.e., $|C \cap \mathbb{I}^j| \leq 1$ and $|C \cap \mathbb{M}| \leq 1$ for all $C \in A \in \mathcal{A}^j$. In the following, let i_C and m_C denote the target and measurement indices corresponding to index cell C .

2) *Update equations*: Given a prior PMBM density with parameters according to (38), a set of measurements Z , and the standard measurement model (Section III-B) with measurement likelihood and probability of missed detection given by

$$l_z(z|x) = p_D(x)p(z|x), \quad (41a)$$

$$q_D(x) = 1 - p_D(x), \quad (41b)$$

the updated density is a PMBM density given by

$$\pi^{\text{PMBM}}(X|Z) = \sum_{X^u \cup X^d = X} \pi^P(X^u) \pi^{\text{MBM}}(X^d), \quad (42a)$$

$$\pi^P(X^u) = e^{-\langle q_D, \lambda_+^u, 1 \rangle} \prod_{x \in X^u} q_D(x) \lambda_+^u(x), \quad (42b)$$

$$\pi^{\text{MBM}}(X^d) = \sum_{j \in \mathbb{J}_+} \sum_{A \in \mathcal{A}^j} w_A^j \pi_A^j(X^d), \quad (42c)$$

$$\pi_A^j(X^d) = \sum_{\Theta_C \in A} \prod_{C \in A} \pi_C^j(X^{i_C}). \quad (42d)$$

where the weights are given by

$$w_A^j = \frac{w_+^j \prod_{C \in A} \mathcal{L}_C}{\sum_{j \in \mathbb{J}} \sum_{A \in \mathcal{A}^j} w_+^j \prod_{C \in A} \mathcal{L}_C} \quad (43a)$$

$$\mathcal{L}_C = \begin{cases} \lambda^{\text{fa}}(z^{m_C}) + \langle \lambda_+^u, l_z(z^{m_C} | \cdot) \rangle, & \text{if } C \cap \mathbb{I}^j = \emptyset, C \cap \mathbb{M} \neq \emptyset, \\ r_+^{j,i_C} \langle p_+^{j,i_C}, l_z(z^{m_C} | \cdot) \rangle, & \text{if } C \cap \mathbb{I}^j \neq \emptyset, C \cap \mathbb{M} \neq \emptyset, \\ 1 - r_+^{j,i_C} + r_+^{j,i_C} \langle p_+^{j,i_C}, q_D \rangle, & \text{if } C \cap \mathbb{I}^j \neq \emptyset, C \cap \mathbb{M} = \emptyset, \end{cases} \quad (43b)$$

and the densities $\pi_C^j(X)$ are Bernoulli densities with parameters

$$r_C^j = \begin{cases} \frac{\langle \lambda_+^u, l_z(z^{m_C} | \cdot) \rangle}{\lambda^{\text{fa}}(z^{m_C}) + \langle \lambda_+^u, l_z(z^{m_C} | \cdot) \rangle}, & \text{if } C \cap \mathbb{I}^j = \emptyset, C \cap \mathbb{M} \neq \emptyset, \\ 1, & \text{if } C \cap \mathbb{I}^j \neq \emptyset, C \cap \mathbb{M} \neq \emptyset, \\ \frac{r_+^{j,i_C} \langle p_+^{j,i_C}, q_D \rangle}{1 - r_+^{j,i_C} + r_+^{j,i_C} \langle p_+^{j,i_C}, q_D \rangle}, & \text{if } C \cap \mathbb{I}^j \neq \emptyset, C \cap \mathbb{M} = \emptyset, \end{cases} \quad (44a)$$

$$p_C^j(x) = \begin{cases} \frac{l_z(z^{m_C} | x) \lambda_+^u(x)}{\langle \lambda_+^u, l_z(z^{m_C} | \cdot) \rangle}, & \text{if } C \cap \mathbb{I}^j = \emptyset, C \cap \mathbb{M} \neq \emptyset, \\ \frac{l_z(z^{m_C} | x) p_+^{j,i_C}(x)}{\langle p_+^{j,i_C}, l_z(z^{m_C} | \cdot) \rangle}, & \text{if } C \cap \mathbb{I}^j \neq \emptyset, C \cap \mathbb{M} \neq \emptyset, \\ \frac{q_D(x) p_+^{j,i_C}(x)}{\langle p_+^{j,i_C}, q_D \rangle}, & \text{if } C \cap \mathbb{I}^j \neq \emptyset, C \cap \mathbb{M} = \emptyset, \end{cases} \quad (44b)$$