An Adaptive and Scalable Multi-object Tracker based on the Non-homogeneous Poisson Process

Qing Li, *Student Member, IEEE*, Runze Gan, *Student Member, IEEE*, Jiaming Liang, and Simon J. Godsill, *Fellow, IEEE*

***Abstract*—This paper proposes a new adaptive framework for tracking multiple objects in the presence of data association uncertainty and heavy clutter, either with or without knowledge of the measurement rates and/or target shapes. Built upon an online Gibbs sequential Markov chain Monte Carlo sampling scheme, the adaptive tracker is Bayesian optimal and robust, requiring no additional approximations or measurement partition steps. With a non-homogeneous Poisson process measurement model, our tracker can tackle the data association task with linear computational complexity. Meanwhile, we study generalised inverse Gaussian and inverse Wishart distributions for modelling Poisson rates and object shapes, respectively; these prior models ensure closed-form full conditionals in our online Gibbs sampling steps, under which object states and shapes can be jointly estimated with associations and Poisson rates in a parallel fashion. Furthermore, a fast Rao-Blackwellisation scheme for linear Gaussian dynamics is designed and demonstrated to significantly improve both tracking efficiency and accuracy. We validate the efficacy of our method on real and simulated data.**

***Index Terms*—sequential Markov chain Monte Carlo, extended target tracking, data association, Rao-Blackwellisation**

I. Introduction

Recent advances in sensor technologies have led to an increase in sensor resolution, making the traditional point target assumption less conducive to many modern applications, including surveillance, unmanned aerial vehicles (UAVs), and autonomous driving, where multiple detections are captured for each object per time step. The non-homogeneous Poisson process (NHPP) measurement model [1], also referred to as the Poisson point process (PPP) model, is a common extended object measurement model adopted in the existing extended target trackers, e.g., [2]–[7]. It assumes a Poisson-distributed number of measurements from targets and the clutter process, and the measurement process is modelled as an NHPP from the superposition of conditionally independent NHPPs from each target as well as the clutter.

In comparison to conventional point target tracking, extended or group target tracking poses more challenges to maintaining tracking efficiency and accuracy. To begin with, the complexity of data association in extended target tracking dramatically increases with the number of measurements, making traditional methods such as the joint probabilistic data association (JPDA) [8] and the multiple hypothesis tracker (MHT) [9] fail due to the requirement of exhaustive enumeration of all possible association hypotheses. Another challenge is that, in practice, the Poisson rates of the NHPP measurement model may be unknown to users and/or time-varying. Therefore, efficient learning of these time-varying parameters would be essential to maintaining the robustness of the tracker. In addition, practical applications may require the estimation of the target shape, which also demands the robustness of the tracker, since the estimation accuracy of the target position, measurement rates, and target shapes are highly correlated. Therefore, this paper proposes an adaptive and robust extended target tracker that can provide a fast and reliable estimation of target position and shape under varying and unknown detection environments.

A. Related Work

The data association problem has been intensively studied for tracking multiple point targets in clutter. Besides the classical JPDA and MHT algorithms, other popular point target trackers include the probabilistic MHT (PMHT) [10]–[12], the graphical model approaches [13], [14], the sampling-based methods [15]–[17], the random finite set (RFS) approaches [18], and references therein. To tackle the data association uncertainty in extended target tracking, many point multi-target trackers have been developed to accommodate extended targets that can generate more than one measurement. However, a direct extension of the point target tracker to the extended target version, such as the extended target JPDA filter in [19], can experience a large growth in combinatorial complexity as the number of targets and measurements increases.

One way to alleviate the computational burden of data association is to implement a heuristic measurement partition before the data association step. Based on this measurement clustering technique, a two-stage MHT algorithm has been proposed in [20] by using a generalisation of track-oriented MHT recursion to handle repeated measurements. Similarly, a JPDA-based tracker has been devised in [21] that can estimate both the targets' positions and their sizes with a limited computational burden. This measurement partition strategy has also been combined with the recently popular RFS-based approaches to deal with extended target tracking, e.g., the extended target probability hypothesis density (PHD) filter

This research is sponsored by the US Army Research Laboratory and the UK MOD University Defence Research Collaboration (UDRC) in Signal Processing under the SIGNeTS project. It is accomplished under Cooperative Agreement Number W911NF-20-2-0225. The views and conclusions contained in this document are of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory, the MOD, the U.S. Government or the U.K. Government. The U.S. Government and U.K. Government are authorised to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein.

Q. Li, R. Gan, J. Liang, and S. J. Godsill are with the Engineering Department, University of Cambridge, Cambridge CB2 1PZ, U.K. e-mail: {ql289, rg605, jl809, sjg30}@cam.ac.uk

[2] and the extended target Poisson multi-Bernoulli mixture (PMBM) filter [4]. Nevertheless, the performance of these clustering-based methods is closely dependent on the clustering result, and this preprocessing step could deteriorate tracking accuracy, especially in challenging cases of heavy clutter and/or closely-spaced targets [6], [22].

To mitigate the measurement partition error, several sampling methods are employed in the PMBM filter to sample the candidate associations directly without conducting measurement partitioning [22]. In spite of improved tracking results, there are several issues with this method. Firstly, its stochastic optimization sampling scheme is not a rigorous MCMC method and cannot guarantee sampling from the desired posterior distribution. In addition, the MCMC sampling schemes in [22] are not scalable and require significant computation cost because they cannot sample all association components in parallel, and the relative likelihoods need to be recalculated each time an association component is sampled. A further issue is that it heuristically truncates hypotheses to avoid the exhaustive enumeration of possible data association hypotheses, which may lead to undesirable tracking performance in demanding scenarios with high association uncertainty [6].

Another extended target tracker that requires no measurement partition to solve data association is built upon the scalable sum-product algorithm (SPA) data association framework for point targets [13]. For point targets generating no more than one measurement, closed-form SPA calculations are available for linear Gaussian models [13]. However, for extended targets that produce multiple measurements, closed-form SPA operations are not available under a linear Gaussian system. Hence, a particle filter implementation of the SPA algorithm was proposed in [23] for tracking a fixed number of extended targets, which was later developed in [6] to include target detection, termination, and shape estimation. The SPA-based tracker has demonstrated performance advantages compared to the PMBM filter [6]. Nonetheless, this SPA method is suboptimal in a Bayesian sense, e.g., it approximates the joint posterior distribution by a product of marginal distributions at each time step. Furthermore, the cycles of the constructed factor graph raise concerns regarding, e.g., the convergence of the loopy SPA algorithm and the order of message computation, since these factors may affect the estimation result.

The non-homogeneous Poisson process, besides serving as an extended target measurement model, also provides an alternative perspective to solve data association efficiently, free from measurement partitioning. The idea of exploiting the NHPP model to simplify the data association was first presented in [1], [24]. Based on the NHPP model, the likelihood function can be evaluated without constructing the association hypotheses, resulting in a compact particle filter implementation of a Bayesian extended target tracker that circumvents the data association step. A similar strategy for untangling the NP hard data association appears in the probabilistic multiple hypothesis tracker (PMHT), which relaxes the restriction of one measurement per target for point targets so that the PMHT requires neither exhaustive enumeration of associations nor pruning [10]. Although their relationship is rarely mentioned in literature, we will later prove that the NHPP measurement model is equivalent to the point target measurement model of the general PMHT when conditional on the measurement number, where the association probability vector of the PMHT can be calculated by a vector of normalised Poisson rates in the NHPP model (i.e., the Π defined with π_i in (23)).

A recent application of the NHPP measurement model is the extended target JPDA filter [5], where the marginal association probabilities can be obtained with linear complexity in the number of measurements and targets. An improved version of [5] that incorporates detection and existence probabilities can be found in [25]. Nevertheless, the independent marginal association posterior is a rough approximation due to the assumption of an independent predictive likelihood for each measurement conditional on associations. An additional approximation is the sequential PDA state update strategy, which is a recursive moment matching procedure that is sensitive to the order of the measurements [5], [25]. Therefore, a scalable data association framework has been devised in our previous paper [26], which can achieve parallel computing of target states as well as association variables under a sequential Markov chain Monte Carlo (SMCMC) inference framework. Compared to the methods in [5], [6], it theoretically converges to an optimal Bayesian filter with a sufficiently large sample size [17], [27]. Compared to the sampling-based PMBM filter in [22] that cannot independently sample associations and requires suboptimal truncations to avoid enumeration of all association hypotheses, our method is scalable, and both associations and states can be sampled in parallel.

Besides data association uncertainty, another important issue in extended target tracking is the robustness of the algorithm. For one thing, the Poisson rates, which are parameters of the NHPP measurement model, may be unknown and/or time-varying. Previous studies that adopt NHPP measurement models assume Poisson rates either to be known constants [1], [5] or predefined functions of the extended target state [2]. This prior information, if incorrect, could lead to false association results, which further influences the tracking accuracy. Therefore, measurement rate estimation has been considered in several extended target RFS-based trackers with a heuristic design of the predicted density by using a Gamma distribution prior [2], [4]. The same heuristic prediction step has also been adopted in the PMHT framework to estimate the target measurement rates [11], implemented by a batch expectation–maximization (EM) method. An online PMHT method has also been developed for cases of unknown Poisson rates [12]. However, this online PMHT method cannot return the original rate estimates as it lacks the estimation of the normalisation constant. Another problem is the false deduction of the clutter's association probability (i.e., π_0) in [12], which should have been estimated along with the association probabilities of the targets. The particle filter implementation also limits its ability to track a large number of targets due to the degeneration problem of the particle filter [27].

Contrary to the heuristic design of the transition density in [4], [11], we develop a Bayesian optimal multi-target tracker that can capture the temporal characteristics of rates by exploring the generalised inverse Gaussian (GIG) family [26]. Specifically, we explore the GIG distribution as the prior for

Poisson rates, which is conjugate to the Poisson distribution despite rarely being applied in these contexts. Compared to the Gamma prior adopted in [2], [4], the three-parameter GIG distribution is more flexible and known to be a better fit for highly dispersed detection data [28].

B. Contributions

This paper presents a novel adaptive framework that can jointly estimate measurement rates, target kinematic states, and/or shapes in the presence of data association uncertainty and clutter. Here, a scalable solution to the data association problem is derived by exploiting the NHPP measurement model, under which the associations are conditionally independent and can be estimated in parallel without an exhaustive enumeration of all associations. Another key innovation is a fast Rao-Blackwellised SMCMC implementation that is distinct from the well-known Rao-Blackwellisation scheme for linear Gaussian systems [29]–[31]. This fast Rao-Blackwellisation scheme enables the tracker to maintain a parallel processing structure, making it advantageous in large-scale tracking scenarios. Further, we enhance its adaptability to handle realistic situations with unknown and time-varying Poisson rates/target shapes. Throughout this paper, the number of targets is assumed to be known for simplicity, which is applicable to cases where the target number has been acquired manually or through a track-by-detection paradigm. However, the framework may be extended to accommodate an unknown and time-varying number of targets, e.g., by introducing an existence vector as in [17]. Extensions to the joint detection and tracking framework will be presented in future work.

An earlier conference version of this paper has been presented in [26]. In comparison, this paper makes significant improvements and proposes a number of novel developments as follows. First, this paper provides much more detailed derivations and implementation of the algorithm, as well as a full performance analysis by evaluating it in comparison with other popular algorithms including the LT-JPDA filter [5], online PMHT [12], different versions of the SPA-based trackers [6], [23] and the GGIW-PMBM filters [4]; a real animal behavioural dataset has also been studied to demonstrate the efficacy of the proposed method in group target tracking scenarios. Moreover, we analyse the conditions for equivalence between the NHPP trackers and the PMHT algorithm, and provide a direct mathematical relationship between Poisson rates and association probabilities defined in PMHT algorithms.

One of the major advances is that we devised two Rao-Blackwellisation SMCMC schemes for linear Gaussian systems, based on the proposed association based NHPP (AbNHPP) tracker in [26]. The Rao-Blackwellised particle filter has been empirically demonstrated to be capable of alleviating the degeneration of particle filters and thus improving tracking performance, e.g., see [15], [32] for a point target tracking case. Nonetheless, these particle filter based methods would still suffer from severe degeneration in high dimensional problems. In contrast, our proposed Rao-Blackwellisation schemes are based on the SMCMC sampling method, the superiority of which over regular particle filters has recently been demonstrated both empirically and theoretically [17], [27], [33]–[35]. Most importantly, we propose a fast Rao-Blackwellisation scheme that is comparable to the standard Rao-Blackwellisation scheme in estimation accuracy while maintaining an efficient parallel sampling structure and a better mixing property. To our knowledge, this is the first time that this type of Rao-Blackwellisation has been developed for the SMCMC methods in this paper. Typically, the standard Rao-Blackwellisation scheme (e.g. in Section IV-A, or other general Rao-Blackwellisation schemes [15], [29]) operate by sampling from a reduced state space. In contrast, the proposed fast Rao-Blackwellisation scheme in Section IV-B first samples from the original state space to maintain desired sampling and computational features; thereafter, only the samples in the reduced state space are used to approximate the intractable expectation, whereby the estimates enjoy the benefit of Rao-Blackwellisation. We can later see from the results that the advantage brought by this fast Rao-Blackwellised SMCMC scheme is significant in both tracking accuracy and efficiency, compared to the standard joint SMCMC sampler for the considered multi-object tracking problem. The underlying reason is that the stationary distribution of the fast Rao-Blackwellised scheme is built on a mixture Gaussian approximation, which is more accurate than that of the standard joint SMCMC built on the weighted Dirac masses.

Another improvement is that we develop an adaptive extended target tracker that can deal with challenging scenarios where both the rates and target shapes are time-varying by utilising the GIG family and the inverse Wishart distribution [36]. Other target extent models (e.g., [3]) can be easily accommodated in our scheme if necessary. Unlike the popular Gamma Gaussian inverse-Wishart (GGIW) model [4], which uses a heuristic predicted density for both rates and target shapes, our method is mathematically more rigorous, with a closed-form transition density that leads to an optimal Bayesian filter; the conjugacy of the GIG distribution and inverse Wishart to the likelihood function also leads to simple and closed-form full conditionals for Gibbs sampling steps, allowing efficient parallelisation. This novel adaptive extended target tracker can also be employed in group tracking scenarios, where we can acquire the shape, the location, as well as other information (e.g., the approximate target number in the group) for tracking a group of objects. An example result of a group tracking case can be seen in Section. VI-F.

C. Paper Outline

The remainder of this paper is structured as follows. Section II presents the NHPP measurement model and data association framework. Section III presents the AbNHPP tracker under known measurement rates; Rao-Blackwellised AbNHPP trackers for linear Gaussian systems are detailed in Section IV. Section V introduces the proposed adaptive AbNHPP tracker that can handle multi-target tracking tasks with unknown/time-varying Poisson rates and/or object extents. Results and conclusions are given in Sections VI and VII.

II. Model

Assuming that there are K objects moving independently, the overall target state at time step n is $X_n =$

$[X_{n,1}^\top, ..., X_{n,K}^\top]^\top$, where each vector $X_{n,i}, i \in \{1, ..., K\}$ denotes the i-th target's kinematic state, and K is known as a prior. The M_n measurements received at time step n are denoted by $Z_n = [Z_{n,1}, ..., Z_{n,M_n}]$.

A. Dynamical Model

Assume that targets move in a D-dimensional surveillance area, and the target state for each object i in the d-th dimension is $X_{n,i}^d = [x_{n,i}^d, \dot{x}_{n,i}^d]^T$, which contains the target's position and velocity. We assume a continuous time linear Gaussian dynamical model for the kinematic state $X_i^d(t)$, expressed in a stochastic differential equation (SDE) [37]:

$$dX_i^d(t) = A_i^d X_i^d(t)dt + C_i dB_i^d(t), \tag{1}$$

where $B_i^d(t)$ is a one-dimensional Brownian motion with variance Q_i. Subsequently, the state at time $t+\tau$ can be derived by integrating (1) from t to $t+\tau$:

$$X_i^d(t+\tau) = F_i X_i^d(t) + \boldsymbol{w}_{i,t}^d, \;\; \boldsymbol{w}_{i,t}^d \sim \mathcal{N}(0, P_i) \tag{2}$$

with the transition matrix F_i and the noise covariance P_i being

$$F_i^d = e^{\tau A_i}, \;\; P_i^d = \int_0^\tau e^{tA_i} C_i Q_i C_i^\top e^{tA_i^\top} dt. \tag{3}$$

By direct discretisation of (2), the transition density at time step n is $p(X_{n,i}^d | X_{n-1,i}^d) = \mathcal{N}(F_i^d X_{n-1,i}^d, P_i^d)$. For independent moving targets, the joint transition density can be computed as $p(X_n|X_{n-1}) = \prod_{d=1}^D \prod_{i=1}^K p(X_{n,i}^d | X_{n-1,i}^d)$.

For a constant velocity (CV) model, we have $A_i = [0\ 1; 0\ 0]$, $C_i = [0; 1]$, and its transition matrix and noise covariance are

$$F_i^d = \begin{bmatrix} 1 & \tau \\ 0 & 1 \end{bmatrix}, P_i^d = Q_i \begin{bmatrix} \tau^3/3 & \tau^2/2 \\ \tau^2/2 & \tau \end{bmatrix}. \tag{4}$$

B. NHPP Measurement Model

This paper considers an NHPP measurement model as in [1]. Denote the set of Poisson rates by $\Lambda = \{\Lambda_i; i = 0, 1, ..., K\}$, where Λ_0 is the clutter rate, and Λ_i is the i-th target rate, $i = 1, ..., K$. We assume the measurement process of each target i is an NHPP with intensity $\lambda_i(Z_n|X_{n,i})$, where the measurement number $m_{n,i}$ is Poisson distributed with rate Λ_i defined over the observation area V, and $\Lambda_i = \int_V \lambda_i(Z_n|X_{n,i}) dZ_n$. The clutter process is a homogeneous Poisson process (HPP) with intensity $\lambda_0(Z_n|X_{n,0})$ and Poisson rate Λ_0, where $X_{n,0}$ denotes the parameter/information of the clutter. By superposition of the conditional independent NHPP/HPP measurement process from K targets and the clutter, the total measurement process remains an NHPP with intensity $\lambda(Z_n|X_n) = \sum_{i=0}^K \lambda_i(Z_n|X_{n,i})$ and the total number of measurements follows a Poisson distribution with rate $\Lambda_s = \sum_{i=0}^K \Lambda_i$ and $M_n = \sum_{i=0}^K m_{n,i}$:

$$p(M_n) = \frac{e^{-\Lambda_s}(\Lambda_s)^{M_n}}{M_n!}. \tag{5}$$

The measurements are conditionally independent when conditional on the measurement number M_n and target state X_n,

$$p(Z_n|X_n, M_n) = \prod_{j=1}^{M_n} p(Z_{n,j}|X_n), \tag{6}$$

where each measurement $Z_{n,j}$ is an i.i.d. sample from the probability density function (PDF) $p(Z_{n,j}|X_n)$ given by

$$p(Z_{n,j}|X_n) = \frac{\lambda(Z_{n,j}|X_n)}{\Lambda_s} \tag{7}$$

where the intensity function $\lambda(Z_{n,j}|X_n)$ is $\lambda(Z_{n,j}|X_n) = \sum_{i=0}^K \lambda_i(Z_{n,j}|X_{n,i})$ and each intensity can be computed by $\lambda_i(Z_{n,j}|X_{n,i}) = \Lambda_i p(Z_{n,j}|X_{n,i})$. Subsequently, by multiplying (5)-(6) using Bayes' Theorem, the likelihood function of the measurement process is

$$p(Z_n, M_n|X_n) = \frac{e^{-\Lambda_s}}{M_n!} \prod_{j=1}^{M_n} \sum_{i=0}^K \Lambda_i p(Z_{n,j}|X_{n,i}). \tag{8}$$

For the measurement model, we assume the target originated measurement follows a linear and Gaussian model while the clutter measurement is uniformly distributed in the observation area of volume V:

$$p(Z_{n,j}|X_{n,i}) = \begin{cases} \mathcal{N}(HX_{n,i}, R_i), & i \neq 0; \quad \text{(object)} \\ \frac{1}{V}, & i = 0; \quad \text{(clutter)} \end{cases} \tag{9}$$

where H is the observation matrix. For point target i, R_i indicates the sensor noise covariance; for extended target, R_i represents the target extent where the target shape is modelled as a Gaussian distribution,. For time-varying target extent or measurement noise covariance, R_i will be replaced by $R_{n,i}$. We will discuss this time-varying $R_{n,i}$ in Section V-B.

C. Data Association

Under the NHPP measurement process in Section II-B, the total measurements at each time step are generated by $K+1$ conditionally independent NHPP measurement processes from K targets and clutter. However, we still do not know to which NHPP process each measurement belongs, and the associations have to be integrated out by summation over all possible associations in the likelihood function (8).

To retrieve the measurement-target correspondence, we reintroduce the data association variables to avoid the summation involved in the evaluation of the likelihood function. We define data association as a random variable $\theta_n = [\theta_{n,1}, ..., \theta_{n,M_n}]$, with each component $\theta_{n,j}, j \in \{1, ..., M_n\}$ indicating the origin of each measurement $Z_{n,j}$; $\theta_{n,j} = 0$ indicates that $Z_{n,j}$ is a clutter measurement from the clutter's HPP process with $\lambda_0(Z_n|X_{n,0})$, and $\theta_{n,j} = i, i \in \{1, \ldots, K\}$ means that $Z_{n,j}$ is generated from the NHPP of target i with $\lambda_i(Z_n|X_{n,i})$.

Under the assumption of the NHPP measurement model, M_n measurements $\{Z_{n,j}\}_{j=1}^{M_n}$ along with associations $\{\theta_{n,j}\}_{j=1}^{M_n}$ are conditionally independent. Therefore we have

$$p(Z_n, \theta_n|X_n, M_n) = \prod_{j=1}^{M_n} p(Z_{n,j}, \theta_{n,j}|X_n) \tag{10}$$

For each measurement $Z_{n,j}$, the likelihood function is

$$p(Z_{n,j}, \theta_{n,j}|X_n) = \frac{\lambda_{\theta_{n,j}}(Z_{n,j}|X_{n,\theta_{n,j}})}{\Lambda_s}, \tag{11}$$

where $\lambda_{\theta_{n,j}}(Z_{n,j}|X_{n,\theta_{n,j}}) = \Lambda_{\theta_{n,j}} p(Z_{n,j}|X_{n,\theta_{n,j}})$. It can be easily verified that integrating $\theta_{n,j}$ out of $p(Z_{n,j}, \theta_{n,j}|X_n)$ in (11) leads to $p(Z_{n,j}|X_n)$ in (7).

By using Bayes' theorem and (11), we can deduce that the prior for each association $p(\theta_{n,j})$ is a categorical distribution with support $\theta_{n,j} \in \{0, .., K\}$, $[\cdot]$ is the Iverson bracket and $[\theta_{n,j} = i]$ evaluates to 1 if $\theta_{n,j} = i$, and 0 otherwise.

$$p(\theta_{n,j}) = \frac{p(Z_{n,j}, \theta_{n,j}|X_n)}{p(Z_{n,j}|X_{n,\theta_{n,j}})} = \prod_{i=0}^K \left(\frac{\Lambda_i}{\Lambda_s}\right)^{[\theta_{n,j}=i]}. \tag{12}$$

The conditional $p(\theta_{n,j}|Z_{n,j}, X_n)$, which is proportional to (11), is also a categorical distribution as follows:

$$p(\theta_{n,j}|Z_{n,j}, X_n) = \left(\frac{\Lambda_0}{V\tilde{l}}\right)^{[\theta_{n,j}=0]} \prod_{i=1}^{K} \left(\frac{\Lambda_i l_{ij}}{\tilde{l}}\right)^{[\theta_{n,j}=i]}, \quad (13)$$

where $l_{ij} = \mathcal{N}(Z_{n,j}; HX_{n,i}, R_i)$ under the measurement model in Section II, and $\tilde{l}$ is a normalisation constant to ensure the sum of all categories' probabilities equals one. We can see that the full conditional for each $\theta_{n,j}$ has an explicit form, and each association $\theta_{n,j}$ can be directly sampled in parallel. Hence, once the associations are known, we can perform the update step in parallel on the condition that the targets move independently. Therefore, it motivates us to formulate a multi-target tracker based on an online Gibbs sampling scheme such that both association and target state sampling steps are parallelisable. We will present this method in section III.

III. Association based NHPP tracker

This section presents the formulation of the AbNHPP tracker. We assume as a starting point that both Poisson rates and object shapes are known as prior information, meaning that all deduced distributions are implicitly conditional on them. The joint filtering distribution of interest is given by

$$p(X_n, \theta_n|Z_{1:n}) \propto p(M_n)p(Z_n, \theta_n|X_n, M_n) \times \int p(X_n|X_{n-1})p(X_{n-1}|Z_{1:n-1})dX_{n-1}, \quad (14)$$

where $p(M_n)$ and $p(Z_n, \theta_n|X_n, M_n)$ have been defined in (5) and (10), respectively, and $p(X_n|X_{n-1})$ is the state transition density given in Section II-A. As this joint posterior is intractable, here we present an online Gibbs sequential MCMC scheme to jointly estimate the target states and associations. To avoid the integration in (14), we target $p(X_{n-1:n}, \theta_n|Z_{1:n})$ instead of $p(X_n, \theta_n|Z_{1:n})$ at each time step n, where $X_{n-1:n}$ denotes states $\{X_{n-1}, X_n\}$. The AbNHPP tracker for the joint inference of $p(X_n, \theta_n|Z_{1:n})$ is summarised in Algorithm 1.

The implementation of Gibbs sampling at each time step n requires full conditionals of all variables in the joint target distribution. Assume at time step $n-1$, we have a set of N_p unweighted samples from the converged chain $\{X_{n-1}^{(p)}, \theta_{n-1}^{(p)}\}_{p=1}^{N_p}$ to approximate the stationary distribution $p(X_{n-1}, \theta_{n-1}|Z_{1:n-1})$. First, we deduce the full conditionals of the Gibbs sampling block of θ_n, in which we sample each $\theta_{n,j}$ for $j = 1, ..., M_n$ from conditional $p(\theta_{n,j}|\theta_{n,-j}, X_{n-1:n}, Z_{1:n})$, $\theta_{n,-j}$ denotes all $\{\theta_{n,k}\}_{k=1,k\neq j}^{M_n}$:

$$\begin{aligned} p(\theta_{n,j}|\theta_{n,-j}, X_{n-1:n}, Z_{1:n}) &= \frac{p(\theta_n, Z_n|X_{n-1:n}, Z_{1:n-1})}{p(\theta_{n,-j}, Z_n|X_{n-1:n}, Z_{1:n-1})} \\ &= \frac{\prod_{k\neq j} p(\theta_{n,k}, Z_{n,k}|X_n)p(Z_{n,j}, \theta_{n,j}|X_n)}{\prod_{k\neq j} p(\theta_{n,k}, Z_{n,k}|X_n)p(Z_{n,j}|X_n)} \\ &= p(\theta_{n,j}|Z_{n,j}, X_n) \end{aligned} \quad (15)$$

where $p(\theta_{n,j}|Z_{n,j}, X_n)$ has an explicit form given in (13), and the second line is derived from the conditional independence assumption of the NHPP measurement model. Therefore, each $\theta_{n,j}$, $j = 1, ..., M_n$ can be sampled independently in parallel.

Next, we deduce the full conditional of X_{n-1}. For K independently moving targets, the transition density has the property that $p(X_{n-1}|X_n) = \prod_{i=1}^{K} p(X_{n-1,i}|X_{n,i})$. The conditional of X_{n-1} can be deduced by

$$\begin{aligned} p(X_{n-1}|X_n, \theta_n, Z_{1:n}) &\propto p(X_n|X_{n-1})p(X_{n-1}|Z_{1:n-1}) \\ &\approx \frac{1}{N_p}\sum_{p=1}^{N_p}\prod_{i=1}^{K} p(X_{n,i}|X_{n-1,i}^{(p)})\delta_{X_{n-1}^{(p)}}(X_{n-1}) \\ &= \frac{1}{N_p}\sum_{p=1}^{N_p}\prod_{i=1}^{K} \mathcal{N}(X_{n,i}; F_i X_{n-1,i}^{(p)}, P_i)\delta_{X_{n-1}^{(p)}}(X_{n-1}). \end{aligned} \quad (16)$$

Finally, we deduce the full conditional of $X_{n,i}$ for each object $i \in \{1, ..., K\}$ as follows:

$$\begin{aligned} p(X_{n,i}|Z_{1:n}, X_{n,-i}, X_{n-1}, \theta_n) &\propto p(Z_n|X_n, \theta_n)p(X_n|X_{n-1}) \\ &\propto \prod_{j\in\Theta_n^i} p(Z_{n,j}|X_{n,i})p(X_{n,i}|X_{n-1,i}) \\ &\propto \mathcal{N}(\tilde{Z}_n^i; HX_{n,i}, \tilde{R}_i)\mathcal{N}(X_{n,i}; F_i X_{n-1,i}, P_i) \\ &\propto \mathcal{N}(X_{n,i}; \mu_{n|n,i}, \Sigma_{n|n,i}), \end{aligned} \quad (17)$$

where Θ_n^i includes the indexes of all measurements generated from target i, i.e., $\Theta_n^i = \{j|j \in \{1, ..., M_n\}, \theta_{n,j} = i\}$, $X_{n,\sim i}$ denotes all $\{X_{n,k}\}_{k=1,k\neq i}^{K}$. The likelihood $\mathcal{N}(\tilde{Z}_n^i; HX_{n,i}, \tilde{R}_i)$ is obtained by the product formula for multiple multivariate Gaussians [38], and

$$\tilde{Z}_n^i = \frac{1}{|\Theta_n^i|}\sum_{j\in\Theta_n^i} Z_{n,j}, \quad \tilde{R}_i = \frac{1}{|\Theta_n^i|}R_i, \quad (18)$$

where $|\cdot|$ denotes the cardinality of the set.

Thereafter, $\mu_{n|n,i}, \Sigma_{n|n,i}$ in (17) can be calculated by Kalman filtering as follows:

$$\begin{aligned} &\mu_{n|n-1,i} = F_i X_{n-1,i}, \ \Sigma_{n|n-1,i} = P_i \\ &k_n = \Sigma_{n|n-1,i}H^T\left(H\Sigma_{n|n-1,i}H^T + \tilde{R}_i\right)^{-1} \\ &\mu_{n|n,i} = \mu_{n|n-1,i} + k_n(\tilde{Z}_n^i - H\mu_{n|n-1,i}) \\ &\Sigma_{n|n,i} = (I - k_n H)\Sigma_{n|n-1,i} \end{aligned} \quad (19)$$

Therefore, conditioned on the association θ_n, we can update the target states in parallel.

Algorithm 1: Association-based NHPP tracker

```
for time step n = 1 to T do
    Initialization
    for iteration m = 1 to M_iter do
        Sample θ^m_{n,1}, ..., θ^m_{n,M_n} from (13) in parallel;
        Sample X^m_{n-1} from (16);
        For every object i = 1, ..., K,
        Sample X^m_{n,i} from (17) in parallel.
    end
    After a burn-in time: keep the subsequent N_p
      samples {X^(p)_n, θ^(p)_n}_{p=1}^{N_p} ~ p(X_n, θ_n|Z_{1:n}).
end
```

A. Discussion

1) Relationship to probabilistic multi-hypothesis tracker: In the basic PMHT [39], a fundamental assumption is that each target can generate more than one measurement. Accordingly, it defines an association probability vector (also known as

measurement probability vector [10]) of K targets and clutter $\Pi = [\pi_0, ..., \pi_K]$, where $\sum_{i=0}^{K} \pi_i = 1$. By relaxing the one measurement per target restriction, a probabilistic structure for the association is defined as $\Pr(\theta_{n,j} = i) = \pi_i$ (see eq.(6) in [39]), which implies that the association prior has a form of

$$p(\theta_{n,j}) = \prod_{i=0}^{K} (\pi_i)^{[\theta_{n,j}=i]}. \tag{20}$$

Subsequently, the likelihood of a general PMHT algorithm can be written as follows according to [12], [39]:

$$p(Z_n|X_n, M_n) = \prod_{j=1}^{M_n} \sum_{i=0}^{K} \pi_i p(Z_{n,j}|X_{n,i}). \tag{21}$$

Recall that in the NHPP measurement model, the likelihood function in (8), if conditional on the measurement number M_n, can be deduced by dividing (8) by the prior of M_n in (5):

$$p(Z_n|X_n, M_n) = \prod_{j=1}^{M_n} \sum_{i=0}^{K} \frac{\Lambda_i}{\Lambda_s} p(Z_{n,j}|X_{n,i}) \tag{22}$$

Therefore, we can observe that the the relationship between association probabilities and measurement rates is:

$$\pi_i = \frac{\Lambda_i}{\Lambda_s}, \quad i = 0, ..., K. \tag{23}$$

where $\Lambda_s = \sum_{i=0}^{K} \Lambda_i$ is the normalisation constant. The underlying reason for this relationship is that the measurement number generated from targets and clutter in the PMHT is multinomial distributed [11], and there is a mathematical relationship between the Poisson and multinomial distributions as in (23). Subsequently, we can see that the likelihood function and association prior of the NHPP trackers and PMHT algorithms are actually equivalent, when conditional on a known measurment number M_n.

2) AbNHPP tracker for point target tracking: With this relationship, the association prior of PMHT in (20) and that of the AbNHPP tracker in (12) equal; thereafter, the conditional of the AbNHPP tracker in (13) can be rewritten in the context of the PMHT by substituting each Λ_i with π_i (with a different normalization constant). Therefore, we can see that the AbHHPP tracker in Algorithm 1 can also be implemented in point target tracking as a counterpart to the PMHT algorithms. As PMHT algorithms has been validated in literature to be capable of tracking point targets, the AbHHPP tracker, due to this equivalence, can also be implemented for point target tracking. For unknown or time-varying association probabilities, likewise, we can first infer the measurement rates, and then calculate the association probabilities using (23). On the contrary, the PMHT cannot infer the Poisson rates in a reverse manner from the estimated association probabilities if the normalisation constant is not given.

In point target tracking, it usually defines a detection probability $P_{d,i}$ for each target i, by which the measurement number (0 or 1 in this case) generated by each target is Bernoulli distributed with mean $P_{d,i}$. Note that under the above-mentioned point target AbNHPP tracker, the detection probability $P_{d,i}$ is not specifically defined. Rather, it is implicitly calculated by the Poisson rate λ_i based on the first moment approximation of the Bernoulli and Poisson distributions as in [1], [11]. Therefore, the approximated estimation of detection probabilities may be inaccurate when the detection probabilities are close to one due to the divergence between Poisson and Bernoulli.

IV. Rao-Blackwellised AbNHPP tracker

In the case of a linear and Gaussian dynamic system defined in Section II-A, a Rao-Blackwellisation scheme can be applied to replace the joint Monte Carlo estimation in Section III; such a Rao-Blackwellised estimator is theoretically proved to have smaller or at least equal variance compared to the joint estimator, and, in practice, can greatly enhance tracking performance. To see this, we first split the joint filtering probability density $p(\theta_{1:n}, X_n|Z_{1:n})$ with regard to the marginal-conditional decomposition:

$$p(\theta_{1:n}, X_n|Z_{1:n}) = p(\theta_{1:n}|Z_{1:n}) p(X_n|\theta_{1:n}, Z_{1:n}), \tag{24}$$

where $p(X_n|\theta_{1:n}, Z_{1:n})$ is linear Gaussian conditional on $\theta_{1:n}$.

For this partially tractable state space model, we only need to sample the association $\theta_{1:n}$. Specifically, we adopt a Monte Carlo approximation, and the posterior distribution $p(\theta_{1:n}|Z_{1:n})$ at each time step n is approximated by a set of N_p unweighted samples $\{\theta_{1:n}^{(p)}\}_{p=1}^{N_p}$:

$$\hat{p}(\theta_{1:n}|Z_{1:n}) = \frac{1}{N_p} \sum_{p=1}^{N_p} \delta_{\theta_{1:n}^{(p)}}(\theta_{1:n}). \tag{25}$$

Subsequently, the marginal posterior of X_n can be approximated by a Gaussian mixture:

$$p(X_n|Z_{1:n}) \approx \frac{1}{N_p} \sum_{p=1}^{N_p} p(X_n|\theta_{1:n}^{(p)}, Z_{1:n}), \tag{26}$$

and each $p(X_n|\theta_{1:n}^{(p)}, Z_{1:n})$ is a Gaussian deduced as follows:

$$\begin{aligned}
p(X_n|\theta_{1:n}^{(p)}, Z_{1:n}) &\propto p(Z_n|X_n, \theta_n^{(p)}) p(X_n|\theta_{1:n-1}^{(p)}, Z_{1:n-1}) \\
&\propto \prod_{i=1}^{K} \prod_{j \in \Theta_n^i} p(Z_{n,j}|X_{n,i}) p(X_{n,i}|\theta_{1:n-1}^{(p)}, Z_{1:n-1}) \\
&\propto \prod_{i=1}^{K} \mathcal{N}(\tilde{Z}_n^i; HX_{n,i}, \tilde{R}_i) \mathcal{N}(X_{n,i}; \mu_{n|n-1,i}^{(p)}, \Sigma_{n|n-1,i}^{(p)}) \\
&\propto \prod_{i=1}^{K} \mathcal{N}(X_{n,i}|\mu_{n|n,i}^{(p)}, \Sigma_{n|n,i}^{(p)})
\end{aligned} \tag{27}$$

where $\Theta_n^i = \{j | j \in \{1, ..., M_n\}, \theta_{n,j}^{(p)} = i\}$, $\tilde{Z}_n^i$ and $\tilde{R}_i$ is computed by (18), and

$$\begin{aligned}
\mu_{n|n-1,i}^{(p)} &= F_i \mu_{n-1|n-1,i}^{(p)} \\
\Sigma_{n|n-1,i}^{(p)} &= F_i \Sigma_{n-1|n-1,i}^{(p)} F_i^\top + P_i.
\end{aligned} \tag{28}$$

The $\mu_{n|n,i}^{(p)}, \Sigma_{n|n,i}^{(p)}$ in (27) are then updated by Kalman filtering. From the final form of the conditional $p(X_n|\theta_{1:n}^{(p)}, Z_{1:n})$, we can see that each object $X_{n,i}$ can be updated independently.

For the intractable part $p(\theta_{1:n}|Z_{1:n})$, we adopt a SMCMC inference scheme by considering the posterior $p(\theta_{1:n}|Z_{1:n})$ as the stationary distribution of the Markov chain, where the posterior is derived as:

$$\begin{aligned}
p(\theta_{1:n}|Z_{1:n}) &\propto p(Z_n|Z_{1:n-1}, \theta_{1:n}) p(\theta_n|M_n) p(\theta_{1:n-1}|Z_{1:n-1}) \\
&\propto p(\theta_n|M_n) \sum_{p=1}^{N_p} p(Z_n|Z_{1:n-1}, \theta_{1:n-1}^{(p)}, \theta_n) \delta_{\theta_{1:n-1}^{(p)}}(\theta_{1:n-1})
\end{aligned} \tag{29}$$

In (29), the joint prior $p(\theta_n|M_n)$ can be calculated from the product of M_n independent association priors by using (12)

$$p(\theta_n|M_n) = \prod_{j=1}^{M_n} p(\theta_{n,j}) = \frac{1}{\Lambda_s^{M_n}} \prod_{i=0}^{K} (\Lambda_i)^{m_{n,i}}, \tag{30}$$

$$m_{n,i} = \sum_{j=1}^{M_n} [\theta_{n,j} = i], \tag{31}$$

where $m_{n,i}$ is the number of measurements generated by object i according to θ_n. Note that by definition of Θ_n^i in (17), we have $m_{n,i} = |\Theta_n^i|$ for all $i = 0, 1, ..., K$.

Each marginal likelihood $p(Z_n|Z_{1:n-1}, \theta_{1:n-1}^{(p)}, \theta_n)$ in (29) can be decomposed as

$$\begin{aligned}
&p(Z_n|Z_{1:n-1}, \theta_{1:n-1}^{(p)}, \theta_n) \\
&= \int p(Z_n|X_n, \theta_n) p(X_n|Z_{1:n-1}, \theta_{1:n-1}^{(p)}, \theta_n) dX_n \\
&= \prod_{j \in \Theta_n^0} p(Z_{n,j}|X_{n,0}) \prod_{i=1}^{K} \int \prod_{j \in \Theta_n^i} p(Z_{n,j}|X_{n,i}) \\
&\qquad \times p(X_{n,i}|\theta_{1:n-1}^{(p)}, Z_{1:n-1}) dX_{n,i} \\
&= \left(\frac{1}{V}\right)^{m_{n,0}} \prod_{i=1}^{K} \int \mathcal{N}(\tilde{Z}_n^i; HX_{n,i}, \tilde{R}_i) \\
&\qquad \times \mathcal{N}(X_{n,i}; \mu_{n|n-1,i}^{(p)}, \Sigma_{n|n-1,i}^{(p)}) dX_{n,i} \\
&= \left(\frac{1}{V}\right)^{m_{n,0}} \prod_{i=1}^{K} \mathcal{N}(\tilde{Z}_n^i; \mu_{\tilde{Z}_n^i}, \Sigma_{\tilde{Z}_n^i}),
\end{aligned} \tag{32}$$

where $\tilde{Z}_n^i$ and $\tilde{R}_i$ are defined in (18), and

$$\mu_{\tilde{Z}_n^i} = H\mu_{n|n-1,i}^{(p)}, \tag{33}$$

$$\Sigma_{\tilde{Z}_n^i} = H\Sigma_{n|n-1,i}^{(p)} H^\top + \tilde{R}_i. \tag{34}$$

Now we can see that $\theta_{1:n}$ cannot be directly sampled from the $p(\theta_{1:n}|Z_{1:n})$ in (29), as this requires the evaluation of (29) to (32) for each possible realization of $\theta_{1:n}$. Here, we first introduce a standard Rao-Blackwellisation algorithm, and then we develop a fast Rao-Blackwellisation version that can maintain the parallel computing efficiency.

A. Standard Rao-Blackwellisation Scheme

In the standard Rao-Blackwellisation scheme, we target the posterior $p(\theta_{1:n}|Z_{1:n})$ as the stationary distribution. Therefore, in the online Gibbs sampling steps, we sequentially infer the association $\theta_{1:n-1}$ and $\theta_{n,1}, ..., \theta_{n,M_n}$ with the full conditionals $p(\theta_{1:n-1}|\theta_n, Z_{1:n})$ and $p(\theta_{n,j}|\theta_{n,-j}, \theta_{1:n-1}, Z_{1:n})$, $j \in \{1, ..., M_n\}$.

By using (29), we can deduce that each conditional has the following expressions:

$$\begin{aligned}
&p(\theta_{1:n-1}|\theta_n, Z_{1:n}) \propto p(Z_n|Z_{1:n-1}, \theta_{1:n}) p(\theta_{1:n-1}|Z_{1:n-1}) \\
&\approx \frac{1}{N_p} \sum_{p=1}^{N_p} p(Z_n|Z_{1:n-1}, \theta_{1:n-1}^{(p)}, \theta_n) \delta_{\theta_{1:n-1}^{(p)}}(\theta_{1:n-1}),
\end{aligned} \tag{35}$$

$$\begin{aligned}
&p(\theta_{n,j}|\theta_{n,-j}, \theta_{1:n-1}, Z_{1:n}) \propto p(\theta_{n,j}, Z_n|\theta_{n,-j}, \theta_{1:n-1}, Z_{1:n-1}) \\
&= p(Z_n|Z_{1:n-1}, \theta_{1:n-1}, \theta_n) p(\theta_{n,j}) \\
&\propto \sum_{i=0}^{K} \Lambda_i p(Z_n|Z_{1:n-1}, \theta_{1:n-1}, \theta_n) \delta_i(\theta_{n,j}).
\end{aligned} \tag{36}$$

where the marginal likelihood $p(Z_n|Z_{1:n-1}, \theta_{1:n-1}, \theta_n)$ can be calculated by (32) and both conditionals are in closed forms.

However, it is noted that the associations in (36) are no longer mutually independent and thus cannot be sampled in parallel. This dependence also makes it computationally inefficient especially for a large data set, as the likelihood needs to be evaluated repeatedly in calculating probabilities of each support $\{\theta_{n,j} = i\}_{i=0}^K$. Therefore, a more efficient Rao-Blackwellisation framework is proposed in Section IV-B.

B. Fast Rao-Blackwellised AbNHPP Tracker

The objective distribution here is $p(\theta_{1:n}|Z_{1:n})$, which is the same as the standard Rao-Blackwellisation scheme. In contrast with the standard approach, the sampling steps are modified to include an auxiliary sampling step for X_n; that is, to obtain $p(\theta_{1:n}|Z_{1:n})$, we instead target $p(\theta_{1:n}, X_n|Z_{1:n})$ at each time step n, while only keeping the converged sample set of $\theta_{1:n}$ to approximate the posterior $p(\theta_{1:n}|Z_{1:n})$. Afterwards, the posterior $p(X_n|Z_{1:n})$ can be approximated by a Gaussian mixture distribution as in (26). This fast Rao-Blackwellised AbNHPP (RB-AbNHPP) tracker is summarised in Algorithm 2. Here we deduce the full conditionals of $\theta_{1:n-1}$, θ_n, and X_n that are required in the implementation of the online Gibbs sequential MCMC sampling scheme.

First, we sample each association $\theta_{n,j}$ from the conditional $p(\theta_{n,j}|\theta_{n,-j}, \theta_{1:n-1}, X_n, Z_{1:n})$, $j = 1, ..., M_n$. The association is independent of previous $\theta_{1:n-1}$ given X_n, and the full conditional equals to $p(\theta_{n,j}|Z_{n,j}, X_n)$ as in (15). Therefore, we can still independently sample each association $\theta_{n,j}$ in this fast RB-AbNHPP tracker from (13). Next, the conditional of $\theta_{1:n-1}$ can be deduced as

$$\begin{aligned}
&p(\theta_{1:n-1}|\theta_n, X_n, Z_{1:n}) \\
&\propto p(X_n|\theta_{1:n-1}, Z_{1:n-1}) p(\theta_{1:n-1}|Z_{1:n-1}) \\
&\approx \frac{1}{N_p} \sum_{p=1}^{N_p} \prod_{i=1}^{K} p(X_{n,i}|\theta_{1:n-1}^{(p)}, Z_{1:n-1}) \delta_{\theta_{1:n-1}^{(p)}}(\theta_{1:n-1}),
\end{aligned} \tag{37}$$

where $p(X_{n,i}|\theta_{1:n-1}^{(p)}, Z_{1:n-1})$ is given in (28). We will see later that sampling $\theta_{1:n-1}$ from (37) is to provide the mean and covariance of the sampled r-th component $p(X_n|\theta_{1:n-1}^{(r)}, Z_{1:n-1})$ in (38), and thus it is not necessary to keep the whole history of $\theta_{1:n-1}$. Hence, sampling $\theta_{1:n-1}$ can be understood as sampling an index r from $\{1, ..., N_p\}$, and the corresponding Gaussian component $p(X_n|\theta_{1:n-1}^{(r)}, Z_{1:n-1})$ will be used in the following Gibbs step of sampling X_n.

Finally, we sample the block X_n from the conditional $p(X_n|\theta_{1:n}, Z_{1:n})$, which is the same as (27) with $\theta_{1:n}$ being the sample from this iteration instead of a component $\theta_{1:n}^{(p)}$. Thus, each target state can be sampled in parallel with the conditional distribution as follows:

$$\begin{aligned}
&p(X_{n,i}|X_{n,-i}, \theta_{1:n}, Z_{1:n}) \propto p(X_n|\theta_{1:n}, Z_{1:n}) \\
&\propto \mathcal{N}(\tilde{Z}_n^i; HX_{n,i}, \tilde{R}_i) \mathcal{N}(X_{n,i}; \mu_{n|n-1}^{(r)}, \Sigma_{n|n-1}^{(r)}) \\
&\propto \mathcal{N}(X_{n,i}; \mu_{n|n,i}, \Sigma_{n|n,i}),
\end{aligned} \tag{38}$$

where r is the index of the sampled Gaussian component, and $\theta_{1:n-1}^{(r)} = \theta_{1:n-1}$. The mean $\mu_{n|n-1}^{(r)}$ and variance $\Sigma_{n|n-1}^{(r)}$ of the r-th component are given in (28), and $\mu_{n|n,i}, \Sigma_{n|n,i}$ in (38) can then be computed by Kalman filtering.

If the Markov chain in this fast RB-AbNHPP tracker at each time step n converges to the stationary distribution $p(\theta_{1:n}, X_n|Z_{1:n})$, the output samples $\{\theta_{1:n}^{(p)}\}_{p=1}^{N_p}$ from Algorithm 2 are implicitly drawn exactly from $p(\theta_{1:n}|Z_{1:n})$, which is the stationary distribution in the standard Rao-Blackwellisation scheme. Therefore, we can expect this fast Rao-Blackwellisation scheme to have comparable performance to the standard Rao-Blackwellisation scheme in Section IV-A while maintaining an efficient parallel sampling structure. Therefore, all the RB-AbNHPP trackers deduced in the following sections are based on this fast RB-AbNHPP tracker.

Compared to the joint inference scheme in Section III, the fast RB-AbNHPP tracker enhances the estimation performance without introducing a large extra computational burden, and we can see from the results in Section VI-C that the advantage brought by the Rao-Blackwellisation is significant. Besides theoretically having a smaller estimation variance, another explanation for the performance boost may be due to the fact that the posterior of object states in the RB-AbNHPP tracker is approximated by a Gaussian mixture, which can normally preserve more information and lead to a more accurate stationary distribution at future time steps than approximated by a group of particles in the joint inference scheme. This improvement in accuracy due to the Rao-Blackwellisation scheme is especially crucial for tracking multiple objects in clutter, where an inaccurate stationary distribution at any time step may cause a large estimation error and further result in track loss, which will make it difficult to restore the track in future time steps. Comparatively, other estimation tasks may have less obvious performance improvements by adopting a Rao-Blackwellisation scheme, e.g. in [29], [30], [32].

Algorithm 2: Fast Rao-Blackwellised AbNHPP tracker

for ***time step*** $n = 1$ *to* T **do**
  Initialization $X_n^{(0)}$
  for ***iteration*** $m = 1$ *to* M_{iter} **do**
    Sample $\theta_{n,j}^m, j = 1, ..., M_n$ in parallel by (13);
    Sample $\theta_{1:n-1}^m$ in parallel by (37);
    Sample $X_{n,i}^m, i = 1, ..., K$ in parallel by (38).
  end
  After a burn-in time: keep the subsequent N_p samples $\{\theta_{1:n-1}^{(p)}, \theta_n^{(p)}\}_{p=1}^{N_p} \sim p(\theta_{1:n}|\Lambda_n, Z_{1:n})$.
  Each object state is updated by Kalman filtering and $p(X_n|Z_{1:n}) \approx \frac{1}{N_p}\sum_{p=1}^{N_p} p(X_n|\theta_{1:n}^{(p)}, Z_{1:n})$
end

V. Adaptive AbNHPP tracker for point or extended targets

The AbNHPP trackers developed in Section III and IV is on the premise that Poisson rates and object extents (i.e., sensor noises for point targets) are known as constants. However, this assumption can be restrictive in real-world applications where such information may be unknown and/or time-varying. Therefore, in this section we extend the AbNHPP tracker to a more general framework which can handle multi-target tracking tasks with unknown/time-varying Poisson rates and/or object extents. Specifically, we will first introduce the measurement rate model and the object shape model. Then, we will devise the adaptive AbNHPP tracker based on the proposed fast RB-AbNHPP tracker in Section IV-B, since it outperforms the joint sequential MCMC estimator in Section III.

A. Measurement Rate Model

Recall that under the NHPP measurement model, the number of measurements $m_{n,i}$ generated by object i ($i = 0, .., K$) is Poisson distributed, that is,

$$p(m_{n,i}|\Lambda_{n,i}) = \frac{e^{-\Lambda_{n,i}}(\Lambda_{n,i})^{m_{n,i}}}{m_{n,i}!}, i \in \{0, 1, ..., K\}. \quad (39)$$

Here we assume the Poisson rate $\Lambda_{n,i}$ for each object i to be a time-varying random variable with a prior $g(\Lambda_{n,i})$ implicitly conditional on known parameters (e.g., the hyperparameters of the prior distribution or Poisson rates $\Lambda_{1:n-1,i}$ at previous time steps). To better capture the dynamic detection profile while keeping the efficient inference structure intact, we explore in this section a more flexible conjugate family of distributions, namely, the GIG distributions, for the Poisson rate, in contrast to a common choice such as the Gamma distribution. The GIG distribution is a three-parameter distribution

$$\mathcal{GIG}(\Lambda; a, b, p) = \frac{(a/b)^{p/2}}{2\mathcal{K}_p(\sqrt{ab})}\Lambda^{p-1}\exp(-\frac{a}{2}\Lambda - \frac{b}{2\Lambda}) \quad (40)$$

with the expectation being

$$\mathbb{E}[\Lambda] = \frac{\sqrt{b}\mathcal{K}_{p+1}(\sqrt{ab})}{\sqrt{a}\mathcal{K}_p(\sqrt{ab})}, \quad (41)$$

where $\mathcal{K}_\nu(z)$ denotes the modified Bessel function of the second kind for order ν and argument z. Note that Gamma distribution $\mathcal{GIG}(\Lambda; \sqrt{2a}, 0, p)$, inverse Gaussian $\mathcal{GIG}(\Lambda; a, b, -\frac{1}{2})$, and inverse gamma $\mathcal{GIG}(\Lambda; 0, \sqrt{2a}, -p)$ may be obtained as special cases of the GIG (see [28]).

The transition densities of Poisson rates for every object and the clutter are assumed mutually independent $p(\Lambda_n|\Lambda_{1:n-1}) = \prod_{i=0}^{K} p(\Lambda_{n,i}|\Lambda_{1:n-1,i})$. To satisfy the conjugacy of Bayesian inference, both the prior $p(\Lambda_{0,i})$ and its transition density $p(\Lambda_{n,i}|\Lambda_{1:n-1,i})$ have a form of a GIG distribution. It leads to a closed-form full conditional distribution of $\Lambda_{n,i}$ in online Gibbs sampling steps, which will be deduced in Section V-C. Specifically, we discuss two settings: Poisson rates are time independent, and Poisson rates follow a first-order Markov chain. Inference methods will be described for each setting.

1) Time independent Poisson rates: Assume that the Poisson rates $\Lambda_{n,i}$ ($i = 0, .., K$) at current time step n are not dependent on the previous $\Lambda_{1:n-1,i}$, that is, the transition density follows a zero-order Markov chain $p(\Lambda_{n,i}|\Lambda_{1:n-1,i}) = p(\Lambda_{n,i})$. In this case, we let the prior of each object i at every time step n be a GIG distribution

$$p(\Lambda_{n,i}) = \mathcal{GIG}(\Lambda_{n,i}; a_i, b_i, p_i). \quad (42)$$

The joint distribution $p(m_{n,i}, \Lambda_{n,i})$ equals $p(m_{n,i}|\Lambda_{n,i})$ in (39) times $p(\Lambda_{n,i})$ in (42). By marginalising $\Lambda_{n,i}$ out of the joint distribution $p(m_{n,i}, \Lambda_{n,i})$, the number of measurements from target i turns out to be a Sichel distribution:

$$p(m_{n,i}) = \mathrm{Si}(m_{n,i}; \sqrt{b_i(a_i + 2)}, \frac{2}{a_i + 2}, p_i). \quad (43)$$

The Sichel distribution is a three-parameter distribution:

$$\mathrm{Si}(m;\alpha,\beta,\gamma)=\frac{((1-\beta)^{\frac{1}{2}})^{\gamma}}{\mathcal{K}_{\gamma}(\alpha(1-\beta)^{\frac{1}{2}})}\frac{(\alpha\beta/2)^{m}}{m!}\mathcal{K}_{m+\gamma}(\alpha), \quad (44)$$

where shape parameters $\alpha>0$, $0<\beta<1$, $-\infty<\gamma<\infty$. Compared to the Poisson distribution, the Poisson mixture distributions possess an over-dispersion property such that the variance is always larger than the mean, regardless of the choice of the mixing density. The Sichel distribution is one example of a Poisson mixture distribution when the GIG distribution is applied as the mixing density; a negative binomial distribution is deduced if the Gamma distribution serves as the mixing density, while the Sichel distribution can have a longer tail than the negative binomial distribution. The property of a Sichel distributed measurement number may seem indirect to the current rate estimation framework; nonetheless, the over-dispersion property of these mixed Poisson distributions may lead to a more stable estimation of the target number, e.g., in [40], and this part will be analysed in future work.

2) Time dependent Poisson rates: For cases where the Poisson rates change more gradually over time, we introduce a novel GIG first-order Markov chain as the prior for the Poisson rate, where the Poisson rate is dependent on Poisson rate of the last time step $p(\Lambda_{n,i}|\Lambda_{1:n-1,i})=p(\Lambda_{n,i}|\Lambda_{n-1,i})$, $(i=0,..,K)$. We construct the chain such that the expectation of $\Lambda_{n,i}$ w.r.t. $p(\Lambda_{n,i}|\Lambda_{n-1,i})$ equals the previous time step's Poisson rate $\Lambda_{n-1,i}$. Under this constraint, the transition density of $\Lambda_{n,i}$ for each object i is designed as:

$$p(\Lambda_{n,i}|\Lambda_{n-1,i})=\mathcal{GIG}(\Lambda_{n,i};\frac{r_c r_B}{\Lambda_{n-1,i}},\frac{r_c\Lambda_{n-1,i}}{r_B},p_i), \quad (45)$$

where $r_B=\frac{\mathcal{K}_{p_i+1}(r_c)}{\mathcal{K}_{p_i}(r_c)}, r_c>0$. We can verify by using (41) that the expectation of Λ_n w.r.t. $p(\Lambda_n|\Lambda_{n-1})$ equals to $\Lambda_{n-1,i}$ under such GIG chain prior constructed in (45). Under this setting, Poisson rates are unlikely to change greatly from the last time step, making it suitable for tracking scenarios with a more stable detection environment.

B. Object Shape Model

Here a simple Gaussian distribution is adopted to represent target's extent with an ellipsoidal shape. When the measurement error is negligible compared to target extent, the measurement likelihood for the j-th measurement generated from target i can be defined by $p(Z_{n,j}|X_{n,i})=\mathcal{N}(Z_{n,j};HX_{n,i},R_{n,i})$. The measurement error, if required, can be included as an additive sensor noise as in [3], and our framework can be easily extended to this case.

We adopt a common assumption that transition densities of all object shapes are mutually independent $p(R_n|R_{n-1})=\prod_{i=1}^{K}p(R_{n,i}|R_{n-1,i})$. The evolving model for each object shape is assumed to be a first-order Markov chain, that is, $p(R_{n,i}|R_{1:n-1,i})=p(R_{n,i}|R_{n-1,i})$. In order to satisfy conjugacy to a Gaussian likelihood function, we assume an inverse Wishart distribution for each object's shape $R_{n,i}$. The inverse Wishart distribution [41] defined on a $d\times d$ positive-definite matrices R is expressed as

$$\mathcal{IW}_d(R;v,\psi)=\frac{\det(\psi)^{\frac{v}{2}}\det(R)^{-\frac{(v+d+1)}{2}}}{2^{\frac{vd}{2}}\Gamma_d(\frac{v}{2})}e^{-\frac{\mathrm{Tr}(\psi R^{-1})}{2}} \quad (46)$$

where v denotes degrees of freedom, $v>d-1$, ψ is the $d\times d$ positive-definite scale matrix, and $\Gamma_d(\cdot)$ is the multivariate Gamma function. The expectation is given by

$$\mathbb{E}[R]=\psi/(v-d-1). \quad (47)$$

Specifically, we adopt the similar assumption that object shapes change smoothly over time. Hence, we construct a Markov chain that the expectation of $R_{n,i}$ w.r.t. $p(R_{n,i}|R_{n-1,i})$ equals the previous time step's object shape $R_{n-1,i}$, and its the transition density is devised as

$$p(R_{n,i}|R_{n-1,i})=\mathcal{IW}_d(R_{n,i};\xi_i,R_{n-1,i}(\xi_i-d-1)), \quad (48)$$

where ξ_i is a known constant. It can be easily verified that the expectation of $R_{n,i}$ w.r.t. $p(R_{n,i}|R_{n-1,i})$ is $R_{n-1,i}$.

C. Inference Algorithms

Our task is to infer Poisson rates $\Lambda_{1:n}$ and object shapes $R_{1:n}$ in conjunction with the object states X_n and associations $\theta_{1:n}$. For a linear Gaussian dynamical model, we adopt the fast Rao-Blackwellisation scheme in Section IV-B, where we utilise a SMCMC strategy to approximate the posterior $p(\theta_{1:n},\Lambda_{1:n},R_{1:n}|Z_{1:n})$, based on which the marginal distribution of X_n can be directly estimated as a Gaussian mixture Monte Carlo approximation using Kalman filtering. In particular, we construct a MCMC chain with the stationary distribution being $p(X_n,\theta_{1:n},\Lambda_{1:n},R_{1:n}|Z_{1:n})$; when the chain converges, we keep the samples of $\theta_{1:n},\Lambda_{1:n},R_{1:n}$ to approximate the target distribution $p(\theta_{1:n},\Lambda_{1:n},R_{1:n}|Z_{1:n})$. The stationary distribution is written as:

$$\begin{aligned}&p(X_n,\theta_{1:n},\Lambda_{1:n},R_{1:n}|Z_{1:n})\\&\propto p(M_n|\Lambda_n)p(Z_n|X_n,\theta_n,R_n)p(\theta_n|M_n,\Lambda_n)p(\Lambda_n|\Lambda_{1:n-1})\\&\quad\times p(R_n|R_{1:n-1})p(X_n|\theta_{1:n-1},R_{1:n-1},Z_{1:n-1})\\&\quad\times p(\theta_{1:n-1},\Lambda_{1:n-1},R_{1:n-1}|Z_{1:n-1}),\end{aligned} \quad (49)$$

where the transitions $p(\Lambda_n|\Lambda_{1:n-1})$ are defined according to different measurement rate models in Section V-A, and transition $p(R_n|R_{1:n-1})$ are defined in Section V-B.

To sample from the stationary distribution in (49), once again, a SMCMC with Gibbs sampling steps is adopted. The general procedure is presented in Algorithm 3. Specifically, we sample $\{\theta_{1:n-1},\Lambda_{1:n-1},R_{1:n-1}\}$, θ_n, Λ_n, R_n, and X_n in turns from the conditional distribution of each block. The explicit form of each component's conditional distribution is deduced as follows.

First, the conditional of joint state at last time step can be derived from (49):

$$\begin{aligned}&p(\theta_{1:n-1},\Lambda_{1:n-1},R_{1:n-1}|X_n,\theta_n,\Lambda_n,R_n,Z_{1:n})\\&\propto p(\Lambda_n|\Lambda_{1:n-1})p(R_n|R_{1:n-1})p(X_n|\theta_{1:n-1},R_{1:n-1},Z_{1:n-1})\\&\quad\times p(\theta_{1:n-1},\Lambda_{1:n-1},R_{1:n-1}|Z_{1:n-1})\\&\approx\sum_{p=1}^{N_p}p(X_n|\theta_{1:n-1}^{(p)},R_{1:n-1}^{(p)},Z_{1:n-1})p(\Lambda_n|\Lambda_{1:n-1}^{(p)})\\&\quad\times p(R_n|R_{1:n-1}^{(p)})\delta_{\{\theta,\Lambda,R\}_{1:n-1}^{(p)}}(\theta_{1:n-1},\Lambda_{1:n-1},R_{1:n-1})\end{aligned} \quad (50)$$

Note that this step is equivalent to sample one index from $1,...,N_p$ according to the weight of each sample.

Algorithm 3: Adaptive RB-AbNHPP tracker

for ***time step*** $n = 1$ *to* T **do**
  Initialization $\theta_{1:n}^{(0)}$ $\Lambda_{1:n}^{(0)}$, $R_{1:n}^{(0)}$, $X_n^{(0)}$
  for ***iteration*** $m = 1$ *to* M_{iter} **do**
    Step 1: sample joint state at time step $n-1$
    Sample $\{\theta_{1:n-1}^m, \Lambda_{1:n-1}^m, R_{1:n-1}^m\}$ by (50);
    Step 2: sample measurement rates (optional)
    1) For time independent Poisson rates:
    Sample $\Lambda_{n,i}^m, i = 0, ..., K$ in parallel by (51).
    2) For time dependent Poisson rates:
    Sample $\Lambda_{n,i}^m, i = 0, ..., K$ in parallel by (52).
    Step 3: sample object shapes (optional)
    Sample $R_{n,i}^m, i = 1, ..., K$ in parallel by (57).
    Step 4: sample object states
    Sample $X_{n,i}^m, i = 1, ..., K$ in parallel by (38).
    Step 5: sample associations
    Sample $\theta_{n,j}^m, j = 1, ..., M_n$ in parallel by (13);
  end
  After a burn-in time:
  keep the subsequent N_p samples
  $\{\theta_{1:n}^{(p)}, \Lambda_{1:n}^{(p)}, R_{1:n}^{(p)}\}_{p=1}^{N_p} \sim p(\theta_{1:n}, \Lambda_{1:n}, R_{1:n}|Z_{1:n})$.
end

1) Conditionals of measurement rates: We will derive the conditionals of measurement rates under two measurement rate models, respectively.

a) Time independent Poisson rates: First, we consider a time independent Poisson rate prior for each $\Lambda_{n,i}(i = 0, ..., K)$. By using (49) and the time independence of the rates, the condition $\Lambda_{1:n-1}$ in the full conditional $p(\Lambda_n|X_n, \theta_{1:n}, \Lambda_{1:n-1}, R_{1:n}, Z_{1:n})$ can be omitted, and

$$\begin{aligned} &p(\Lambda_n|X_n, \theta_{1:n}, R_{1:n}, Z_{1:n}) \propto p(M_n|\Lambda_n)p(\theta_n|M_n, \Lambda_n)p(\Lambda_n) \\ &= \frac{e^{-\Lambda_{s,n}}(\Lambda_{s,n})^{M_n}}{M_n!}\frac{1}{\Lambda_{s,n}^{M_n}}\prod_{i=0}^{K}(\Lambda_{n,i})^{m_{n,i}}\mathcal{GIG}(\Lambda_{n,i}; a_i, b_i, p_i) \\ &\propto \prod_{i=0}^{K} e^{-\Lambda_{n,i}}\Lambda_{n,i}^{m_{n,i}}\Lambda_{n,i}^{p_i-1}e^{-\frac{a_i}{2}\Lambda_{n,i}-\frac{b_i}{2\Lambda_{n,i}}} \\ &\propto \prod_{i=0}^{K}\mathcal{GIG}(\Lambda_{n,i}; a_i + 2, b_i, m_{n,i} + p_i), \end{aligned} \tag{51}$$

where $m_{n,i}$ is the number of measurements generated by object i given θ_n. Clearly, the conditional for each Poisson rate has an explicit form of a GIG distribution, which verifies the conjugacy of the GIG prior. In this way, we can sample each $\Lambda_{n,i}(i = 0, 1, ..., K)$ independently from its conditional in a parallel fashion.

b) Time dependent Poisson rates: Under a time dependent Poisson rate model, we can deduce that the conditional of Λ_n is a GIG distribution in a form similar to (51) with different parameters:

$$\begin{aligned} &p(\Lambda_n|X_n, \theta_{1:n}, \Lambda_{1:n-1}, R_{1:n}, Z_{1:n}) \\ &\propto p(M_n|\Lambda_n)p(\theta_n|M_n, \Lambda_n)p(\Lambda_n|\Lambda_{n-1}) \\ &\propto \prod_{i=0}^{K}\mathcal{GIG}(\Lambda_{n,i}; \frac{r_c r_B}{\Lambda_{n-1,i}} + 2, \frac{r_c\Lambda_{n-1,i}}{r_B}, m_{n,i} + p_i). \end{aligned} \tag{52}$$

where we can see that the conditional of each $\Lambda_{n,i}$ is also a GIG distribution and each $\Lambda_{n,i}(i = 0, 1, ..., K)$ can be sampled in parallel.

2) Conditionals of extended target shapes: For tracking a group or a large object that the shape is non-negligible, we can jointly estimate the shape $R_{1:n}$, along with $\theta_n, \Lambda_{1:n}, X_{1:n}$. Under the object shape model in Section V-B, the conditional of shape state R_n, by using (49), is

$$\begin{aligned} &p(R_n|X_n, \theta_{1:n}, R_{1:n-1}, \Lambda_{1:n}, Z_{1:n}) \\ &\propto p(Z_n|X_n, \theta_n, R_n)p(R_n|R_{n-1}) \\ &= \prod_{i=1}^{K}\prod_{j\in\Theta_n^i} p(Z_{n,j}|X_{n,i}, R_{n,i})p(R_{n,i}|R_{n-1,i}) \end{aligned} \tag{53}$$

Note that the expression of $\prod_{j\in\Theta_n^i} p(Z_{n,j}|X_{n,i}, R_{n,i})$ is different from (18) in which case we discard the constant factor as object shape $R_{n,i}$ in (18) is a known constant. Here, the likelihood function is derived as

$$\begin{aligned} &\prod_{j\in\Theta_n^i} p(Z_n|X_{n,i}, R_{n,i}) \\ &\propto \mathcal{N}(\tilde{Z}_n^i; HX_{n,i}, \tilde{R}_{n,i})\mathcal{W}_d(S_n^i; m_{n,i} - 1, R_{n,i}) \qquad (54) \\ &S_n^i = \sum_{j\in\Theta_n^i}(Z_{n,j} - \tilde{Z}_n^i)(Z_{n,j} - \tilde{Z}_n^i)^\top \qquad (55) \end{aligned}$$

where $\mathcal{N}(\tilde{Z}_n^i; HX_{n,i}, \tilde{R}_{n,i})$ has the same form of (18), $\Theta_n^i = \{j|j \in \{1, ..., M_n\}, \theta_{n,j} = i\}$, and $m_{n,i}$ is defined in (31). The Wishart distribution defined on $S^{d\times d}$ is written as

$$\mathcal{W}_d(S; m, R) = \frac{\det(S)^{\frac{1}{2}(m-d-1)}e^{-\frac{1}{2}\operatorname{Tr}(SR^{-1})}}{2^{\frac{md}{2}}\Gamma_d(\frac{m}{2})\det(R)^{\frac{m}{2}}} \tag{56}$$

where m denotes degrees of freedom, and $m \geq d$.

Applying (54) to (53), the full conditional of R_n is

$$\begin{aligned} &p(R_n|X_n, \theta_{1:n}, R_{1:n-1}, \Lambda_{1:n}, Z_{1:n}) \qquad (57) \\ &\propto \prod_{i=1}^{K}\mathcal{N}(\tilde{Z}_n^i; HX_{n,i}, \tilde{R}_{n,i})\mathcal{W}_d(S_n^i; m_{n,i} - 1, R_{n,i}) \\ &\qquad \times \mathcal{IW}_d(R_{n,i}; \xi_i, R_{n-1,i}(\xi_i - d - 1)) \\ &\propto \prod_{i=1}^{K}\mathcal{IW}_d(R_{n,i}; \xi_i + m_{n,i}, R_{n-1,i}(\xi_i - d - 1) + S_n^i). \end{aligned}$$

Hence, we verify that the full conditional of each target shape is also an inverse Wishart distribution, and we can sample the object shapes in parallel for object $i = 1, ..., K$.

3) Conditionals of associations and object states: Given the sampled $\{\theta_{1:n-1}, \Lambda_{1:n-1}, R_{1:n-1}\}$, Λ_n, and R_n, the full conditionals of θ_n and X_n are equivalent to (13) and (38), respectively. Therefore we will not repeat here.

VI. RESULTS

A. Performance Metrics

We use the following metrics to evaluate the performance of proposed algorithms and to compare them with other methods.

1) Optimal sub pattern assignment (OSPA) [42]: The OSPA distance is used to evaluate the tracking accuracy, and we note that this can be regarded as a version of the general OSPA (GOSPA), normalised over the target number and with parameter $\alpha = 1$ [43]. The order is set to $p = 2$

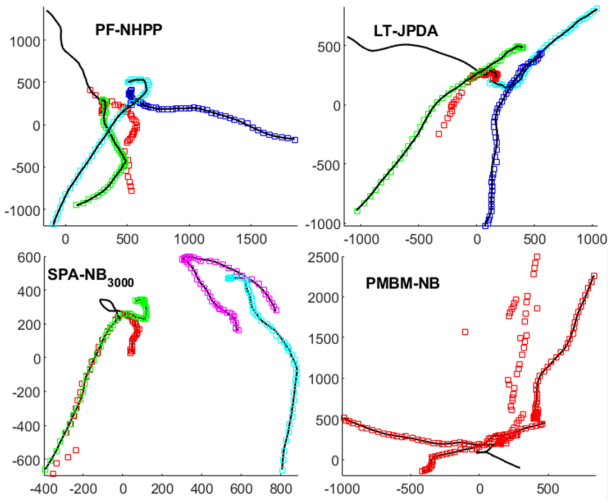

Fig. 1: Example tracks and estimates of scenario 1; black lines are ground truth, and squares are estimates

and the distance cut-off value is $c = 50$. We additionally calculate its mean and standard deviation over Monte Carlo runs for cases 1 and 2 as the simulated trajectories are all different in these datasets.

2) Track loss percentage: At each time step, a target is considered tracked if the current OSPA value is lower than $c = 50$. Then, a successful track is determined if the target is tracked for at least 80% of its lifespan. In this paper, the track number equals the target number K, and the track loss percentage is the ratio of the unsuccessful track number to the total track number K.
3) CPU time: We record the CPU time required at a single time step and average it over all time steps (System: Intel(R) Core(TM) i7-8550 CPU@1.80 GHz, 8 GB RAM).
4) Gaussian Wasserstein distance (GWD)-OSPA metric: The GWD metric in [44] is designed for evaluating a single extended target tracking performance, which considers both the position and target shape estimation error. For evaluating multiple extended targets. The GWD-OSPA metric is computed by using the optimum assignments of the OSPA metric, similar to the GWD-GOSPA in [4].

B. Design and Settings of Simulations and Comparisons

We design various cases of tracking multiple targets in heavy clutter with or without knowing the measurement rates and target shapes in order to demonstrate the adaptiveness of the proposed method. An extensive comparison analysis of existing tracking algorithms is provided, including the linear time JPDA (LT-JPDA) filter [5], the PF-NHPP tracker [1], the online PMHT algorithm [12], and detailed comparisons with the popular PMBM filters [4] and SPA-based trackers [6]. Other RFS-based methods such as the PHD filter are not included here since the PMBM filters have been verified to outperform them in [4].

In Case 1, we validate the efficacy of the proposed Bayesian optimal SMCMC inference schemes over other suboptimal methods in tracking scenarios with known measurement rates and target shapes. Under this fixed rates and target shapes setting, the differences of the underlying modeling assumptions over compared methods in both measurement rates and target shapes are minimised, meaning that a fair comparison can be made since our proposed GIG rate transitions and the mathematically well-defined inverse Wishart shape transition may be advantageous but are not utilised in all of the compared methods. In Case 2, we analyse the advantages and generality of the proposed GIG prior models in modeling both highly dispersed and slowly changing Poisson rates with a small variance. Case 3 verifies the robustness of the proposed method in tracking multiple extended targets with time-varying measurement rates and target shapes. In Section VI-F, we demonstrate the proposed AbNHPP trackers in group target tracking scenarios with a real fish dataset.

In the simulations, we assume the target number is fixed and known. Therefore, to present a fair comparison, this target number information is incorporated into the PMBM filters [4] and SPA-based trackers [6] that are originally devised for tracking a varying number of targets. To be specific, for PMBM filters and SPA-based trackers, we consider both the standard versions with the target birth process and death process termed 'PMBM-B' and 'SPA-B', and the fixed target number versions termed 'PMBM-NB' and 'SPA-NB', respectively. The reasons for this altered versions of PMBM-NB and SPA-NB algorithms are: 1) it matches the modeling setting in our experiment; 2) it is faster because the birth and death operations are deleted in each of the algorithms; however, note that the original recycling step is kept in the PMBM-NB method, therefore a birth process will still be carried out in a small region due to the recycling process. 3) it achieves a much lower OSPA than the standard version.

To fit the setting of a known target number K, both PMBM-B and PMBM-NB filters are initialised with a single multi-Bernoulli hypothesis with K components, each associated with an existence probability of 1, and a ground truth target state prior. Similarly, the SPA-B and SPA-NB trackers are initialised with the ground truth target number K and target state prior, where the SPA-NB tracker corresponds to the fixed target version in [23]. The target number K is also employed to cap the number of output state estimates by K when evaluating the OSPA metric for PMBM-B, PMBM-NB, and SPA-B methods. For both PMBM-B and SPA-B methods, the survival probability is set to 0.99, and targets are born according to a Poisson process of intensity 0.1 and Gaussian density covering the surveillance region of interest. For both the PMBM-NB and SPA-NB methods, the survival probability is set to 1.

Other general parameter settings are as follows. For both the PMBM-B and PMBM-NB filters, the detection probability is set to 1 to match our model. The ellipsoidal gate size in probability is 0.999. The global hypotheses and Bernoulli components pruning thresholds are set to 10^{-2} and 10^{-3}, respectively. The maximum global hypothesis number is 100. The DBSCAN algorithm is run with 50 different distance values equally spaced between 1 and 50, and the maximum assignments number is 20 for the Murty's algorithm. For SPA-B, the threshold for object declaration is 0.5, and the pruning threshold is 10^{-3}. Note that the parameters in the PMBM and SPA methods are manually tuned considering both the tracking

performance and implementation efficiency; for instance, the grid for the DBSCAN clustering is carefully selected to a value such that a finer grid will no longer improve the tracking accuracy but only introduce extra computational time.

C. Case 1: Tracking Multiple Targets with Known Measurement Rates and Target Shapes

To evaluate the robustness of the algorithms in tracking efficiency and accuracy, we design two tracking scenarios. In the scenario 1, there are four targets moving in a 2D surveillance area; the target rates are all set to 5, and the clutter rate is 50. In contrast, the scenario 2 is more challenging with 8 targets and a heavier clutter. The target rates remain 5 while the clutter rate becomes 300. Besides target number and clutter rate, other parameters are all the same for the two scenarios. Specifically, for each scenario, we generated 50 different synthetic simulations. In each simulation, the trajectories are simulated using the CV model in (2), (4) with $Q_i = 25, i = 1, ..., K$. The total time steps are 50, and the time interval between observations is $\tau = 1$s. The target is simulated as an extended target with an elliptical extent, and its covariance in (9) is set to $R_i = 100\mathrm{I}, (i = 1, ..., K)$ where I is a 2-D identity matrix.

The tracking results are presented in Table I and II. Note that the target number K is inherently known in all methods in the table except for the PMBM-B, the PMBM-NB, and the SPA-B methods. Nevertheless, the cardinality errors of the PMBM filters and SPA-B method are almost always 0 in all time steps owing to our modifications. Therefore, the OSPA metric is only used for comparing the localization error in most cases in this tracking scenario. To visualise the reliability of the trackers, we monitor their track loss percentages over all 50 simulations and present their average value in the tables. The number of particles used in the PF-NHPP, the AbNHPP and RB-AbNHPP trackers is denoted in the subscript in the tables. The burn-in time for the AbNHPP$_{2000}$, the AbNHPP$_{10000}$, the RB-AbNHPP$_{100}$ and the RB-AbNHPP$_{500}$ are respectively 500, 500, 50, and 100. For SPA-B and SPA-NB methods, we set the iteration number to 3 as suggested [6]. The particle number is set to $300, 3000$, and the resulting parameter combinations are denoted as SPA-B$_{3000-3}$, SPA-NB$_{300-3}$ and SPA-NB$_{3000-3}$.

TABLE I: Tracking performance of scenario 1 with 4 targets

method	OSPA (mean $\pm 1\sigma$)	track loss(%)	CPU time (s)
LT-JPDA	12.5 $\pm$ 6.43	10.50	0.002
PF-NHPP$_{10000}$	15.29 $\pm$ 5.20	14.00	4.03
PMBM-B	7.00 $\pm$ 0.80	0.00	0.77
PMBM-NB	6.37 $\pm$ 2.34	0.50	0.29
SPA-B$_{3000-3}$	6.44 $\pm$ 1.36	0.00	1.33
SPA-NB$_{300-3}$	6.58 $\pm$ 1.61	0.50	0.10
SPA-NB$_{3000-3}$	6.01 $\pm$ 1.30	0.50	1.23
AbNHPP$_{2000}$	6.61 $\pm$ 0.76	0.00	1.43
AbNHPP$_{10000}$	5.95 $\pm$ 0.40	0.00	16.22
RB-AbNHPP$_{100}$	5.78 $\pm$ 0.31	0.00	0.04

In both Table I and Table II, our proposed RB-AbNHPP evidently outperforms all other competing methods in tracking accuracy, i.e. it has lower mean and variance of OSPA, and a lower track loss percentage. Besides this superior tracking performance, it is the second fastest tracker, only slower than the LT-JPDA tracker which, however, has significantly inferior tracking performance. The advantage of the RB-AbNHPP is more significant for the challenging tracking task in Table II, where our RB-AbNHPP$_{100}$ achieves a clearly lower OSPA value than its next-ranked method, SPA-NB$_{3000-3}$, with a much faster implementation. This advantageous tracking accuracy of the proposed RB-AbNHPP is owing to the asymptotic properties of MCMC and the efficacy of Rao-Blackwellisation. As stated in Section IV-B, the Rao-Blackwellisation in our RB-AbNHPP can construct a more accurate stationary distribution than our standard AbNHPP, hence requiring fewer samples to achieve the same level of accuracy. This can be seen from the Table I, where the standard AbNHPP with 10000 samples is outperformed by the RB-AbNHPP with only 100 samples, which brings a huge saving in CPU time.

TABLE II: Tracking performance of scenario 2 with 8 targets

method	OSPA (mean $\pm 1\sigma$)	track loss(%)	CPU time (s)
LT-JPDA	22.96 $\pm$ 5.69	35.8	0.015
PMBM-B	8.82 $\pm$ 1.47	7.00	6.35
PMBM-NB	8.50 $\pm$ 2.59	1.50	1.61
SPA-NB$_{300-3}$	7.61 $\pm$ 2.19	0.75	1.45
SPA-NB$_{3000-3}$	7.05 $\pm$ 2.04	1.00	23.55
SPA-B$_{3000-3}$	8.08 $\pm$ 2.86	0.25	98.99
RB-AbNHPP$_{100}$	6.49 $\pm$ 1.12	0.25	0.12
RB-AbNHPP$_{500}$	6.35 $\pm$ 0.35	0.00	0.49

Four examples of tracking results from Scenario 1 can be seen in Figure 1, where we show the presence of tracking loss in the compared methods. Here we omit the plot of our proposed RB-AbNHPP tracker because its estimated results are similar to the ground truth and there is no track loss. The tracking performance of other competing methods can be seen in both Table I and Table II. We can see that the fastest method is the LT-JPDA, but its crude approximations lead to a poor tracking accuracy. The PMBM-NB and SPA-NB algorithms achieve a lower OSPA than the standard PMBM-B and SPA-B methods due to the more accurate modeling assumptions; however we notice that the birth process may be able to retrieve some lost track targets, and hence PMBM-B and SPA-B have lower track loss percentages as shown in Table I. In general, the SPA-based methods outperforms PMBM methods given a large enough sample size. The poor performance of PF-NHPP in Table I, even with a very large sample size, is due to its severe particle degeneration caused by high-dimensional sample space. This degeneration can typically be mitigated by the SMCMC methods, as demonstrated by the results of our tracker in the Table I. We can see from both tables that more samples can typically lead to a better performance for SPA-NB, AbNHPP and RB-AbNHPP, and among them the proposed RB-AbNHPP requires the smallest number of samples to achieve the best performance. This is again owing to the nature of MCMC and Rao-Blackwellisation we analysed before. This significantly small sample size in the proposed RB-AbNHPP makes it the fastest method among all tested sampling-based algorithms. Finally, we emphasis that its efficiency can be further improved by exploiting the scalable features as described in Algorithm 2.

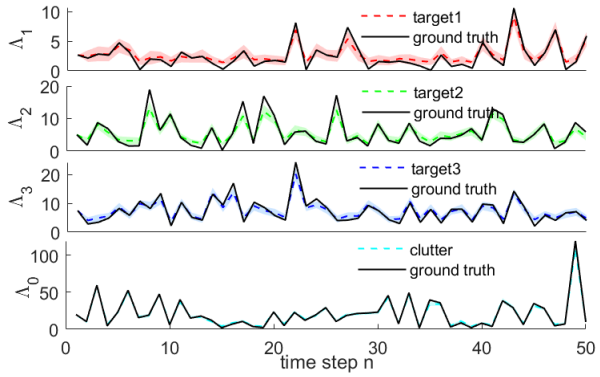

(a) Measurement rates with GIG independent priors

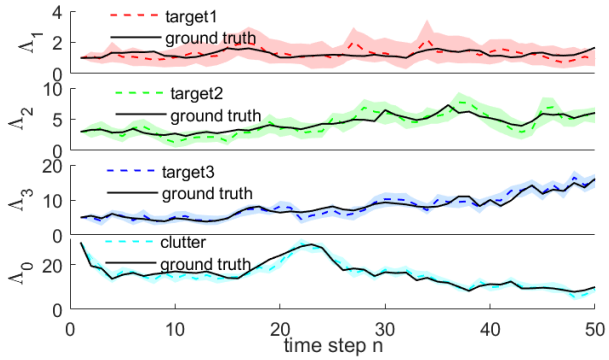

(b) Measurement rates with GIG Markov chain priors

Fig. 2: Estimates and ground truths of the measurement rates for both targets and clutter; the black lines are ground truth; the dotted lines are means of estimated rates and shaded areas represent means $\pm 1\sigma$ (σ-standard deviation).

D. Case 2: Unknown and Varying Measurement Rates

In this simulation, we consider tracking scenarios where measurement rates are unknown and time-varying, and target extents are either negligible (e.g., targets are far away from sensors), or known in advance. Specifically, we analyse the tracking performance and Poisson rate estimation results under two proposed GIG measurement rate models in the following two simulations: (1) multi-target tracking under a time-independent GIG prior rate model for both targets and clutter; (2) multi-target tracking under a GIG Markov chain prior rate model for both targets and clutter. In addition, we compare it with the online PMHT method proposed in [12] to validate the effectiveness of the proposed method.

TABLE III: Results under an independent GIG rate model

method	OSPA (mean $\pm 1\sigma$)	track loss(%)
online PMHT	6.19±3.09	30.0
RB-AbNHPP	1.30±0.08	0.00

1) Poisson rate with a time independent GIG prior: In the first simulation, three extended objects move independently in 2-D over 50 time steps. The trajectories are simulated using the CV model with $Q_i = 1$, and the covariance in (9) is set to $R_i = 4\mathrm{I}, (i = 1, 2, 3)$. The time interval between observations is $\tau = 1$s. We generated 50 datasets to evaluate the performance. For each dataset, the Poisson rates are generated with a GIG independent prior with parameters $a_{1:3} = 0.8$, $a_0 = 0.1$, $b_{1:3} = 0.1$, $b_0 = 10$, $\{p_i = i\}_{i=1}^{3}$ and $p_0 = 0.5$, under which measurements are generated according to Section II-B. For target rates, the GIG priors' means are around 2, 5, and 8; the clutter rate is around 20. The simulated Poisson rates for three targets and clutter are shown in Fig. 2a. We can see that both target rates and clutter rate show an over-dispersed and a long-tailed property.

We use the adaptive RB-AbNHPP tracker introduced in Section V-C with object shape being known constants. In detail, we use a total of 70 particles to approximate $p(\theta_{1:n}, \Lambda_{1:n}|Z_{1:n})$, after a 30-iteration burn-in time. For the online PMHT method, 3000 particles are used in the particle filtering algorithm. Recall that the online PMHT can handle the tracking task under time-varying Poisson rates (see Section III-A1), whereas it is unable to provide valid Poisson rate estimation results. Therefore, we only present the Poisson rate estimation results of our method, shown in Fig. 2a. We can see that the proposed method can correctly estimate rates with high accuracy. From Table III, our method excels in tracking accuracy, while the online PMHT method suffers from track loss even though a large sample size has been used.

TABLE IV: Results under a GIG Markov chain rate model

method	OSPA (mean $\pm 1\sigma$)	track loss(%)
online PMHT	5.53±3.50	26.67
RB-AbNHPP	1.48±0.06	0.00

2) Poisson rate with a GIG Markov chain prior: In the second simulation, trajectories of three extended targets are generated with $Q_i = 1$, and the covariance in (9) is set to $R_i = 25\mathrm{I}, (i = 1, 2, 3)$. The Poisson rates are simulated by a GIG Markov chain prior with parameters $r_c = 10$, $p_{0:3} = 50$. The other settings are the same as in the first simulation. From Fig. 2b, we can observe that the rates change more gradually over time due to the time dependence of Poisson rates. The estimated rates by using the proposed method are shown in Fig. 2b. It can be seen that measurement rates are well estimated in both low and high measurement rate scenarios. The comparison of tracking performance is shown in Table IV. Again, our method has a lower OSPA distance and therefore outperforms the online PMHT in tracking accuracy.

E. Case 3: Tracking Multiple Extended Targets with Varying Measurement Rates and Target Shapes

This section analyses the adaptive RB-AbNHPP tracker's performance for extended target tracking scenarios. We compare our proposed method with the popular GGIW-PMBM filter in [4]. To present a fair comparison, the target number is known to the GGIW-PMBM filter and the modified PMBM-NB implementation is chosen for this GGIW-PMBM filter since it outperforms PMBM-B as described in section VI-C. Since the GGIW-PMBM filter assumes a heuristic predicted density that possesses a time dependency property of Poisson rates, we correspondingly select the time dependent GIG Markov chain rate model in Section V-A2 for a fair

This article has been accepted for publication in IEEE Transactions on Signal Processing. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TSP.2023.3240498

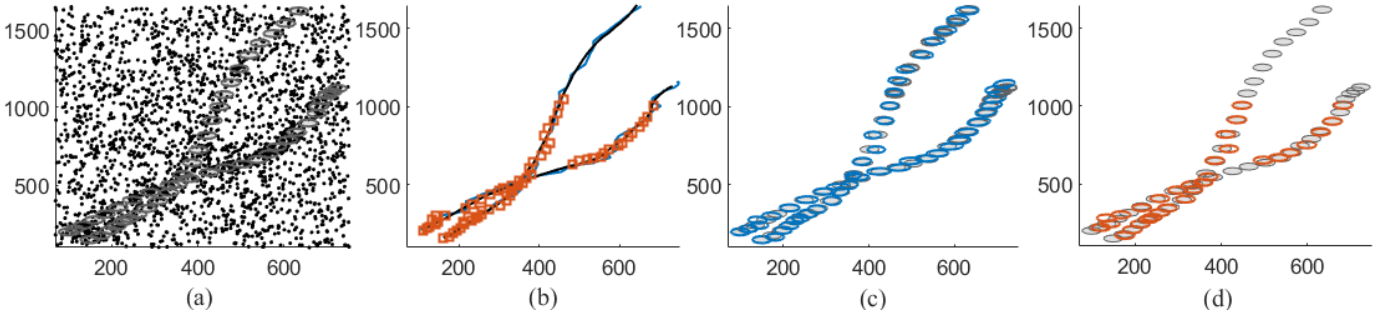

Fig. 3: (a) Measurements and ground truth shapes; grey ellipses denote ground truth target extent at every second time steps; black dots denote measurements; (b) position estimates of two trackers; black lines are ground truth, blue lines are estimates from RB-AbNHPP tracker and tangerine squares are estimates from GGIW-PMBM filter; (c) and (d) are shape estimates of two trackers; blue and tangerine ellipses are estimates from RB-AbNHPP tracker and GGIW-PMBM filter, respectively.

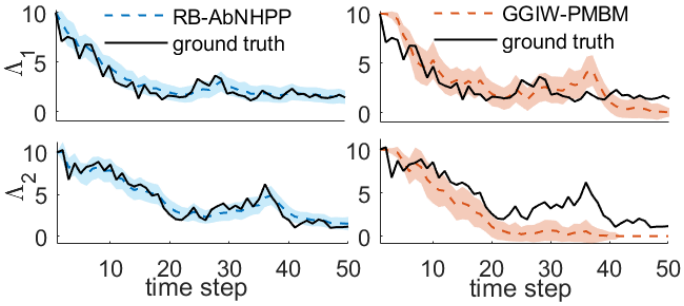

Fig. 4: Measurement rates over 50 time steps; black lines are ground truth; dotted lines and shaded areas are estimated mean and standard derivation of Poisson rates over 50 datasets.

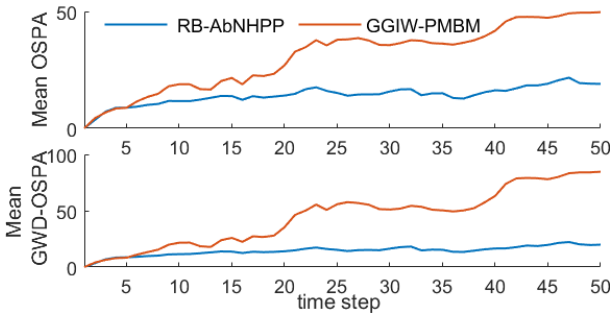

Fig. 5: Mean OSPA and GWD-OSPA over 50 time steps.

comparison. Note that the GGIW-PMBM filters in [4] build on heuristic rate and shape transitions where their heuristic transitions depend on the estimation result at the previous time step, and hence are unable to generate synthetic datasets.

In this simulation, trajectories of two extended targets are generated with a CV model, and $Q_i = 10, i = 1, 2$. The target shape is simulated as defined in Section V-B with $\xi_i = 5000, i = 1, 2$. The Poisson rates follow the GIG Markov chain prior model with parameters $r_c = 5$ and $p_{1:2} = 10$. We assume a known clutter rate $\Lambda_0 = 50$ as the GGIW-PMBM tracker [4] cannot estimate the clutter rate. The measurement data is shown in Fig. 3, where ground truth target extents are presented at every second time steps. The trajectories show that it is a demanding tracking scenario where two targets first move closely in parallel, and then their tracks cross and depart. The ground truth Poisson rates for the two extended targets are shown in Fig. 4. With such a heavy clutter setting in this experiment, the measurements from two targets are buried in clutter. Hence, it is challenging to correctly estimate the measurement rates and track all the targets.

To evaluate the robustness of the algorithm, we generate 50 Monte Carlo runs under the same settings, and the average OSPA distance, the average GWD-OSPA, and estimated measurement rates are calculated over 50 datasets. We can see from Fig. 4 that our method can successfully estimate measurement rates of two targets under this challenging scenario. In comparison, the GGIW-PMBM filter overestimates the Λ_1 from time step 30 to 40 due to the interference of clutter, and it fails to estimate the measurement rate of target 1 and 2 when the target rates are low because of the malfunction of the partition algorithm. From the mean OSPA distance and mean GWD-OSPA metric in Fig. 5, our method has a much lower OSPA and GWD-OSPA value than the GGIW-PMBM filter over all time steps, meaning that it has a more steady and accurate estimation in both position and shape of the targets. One example target extent estimation result is shown in Fig. 3, where the GGIW-PMBM filter frequently loses track, while the proposed method can track all targets and their extent even when target rates are low. In accordance to the implementation simplicity, the GGIW-PMBM filter requires manual tuning of parameters such as measurement partition parameters. Our method, on the other hand, require no manual tuning other than choosing the iteration number, which makes it more adaptive and practically applicable.

F. Real Fish Dataset

We verify the proposed method in a real dataset of a school of 61 golden shiner fish collected in [45]. The video is recorded by a Sony EX1 video camera (1280 x 720 pixels) mounted above the fish tank at a frame rate of 60 Hz over 60 time steps. Based on the preprocessed 2D position measurement data, we additionally add heavy clutter to make this scenario more challenging. The clutter rate is set to 50. The measurement data with clutter is shown in Fig. 6. The figure applies different colors to measurements received at different time steps in order to show the transition of group shape. From Fig. 6, we can observe that the shape of the group changes greatly from its initial shape. In this task, we aim to track the fish school and its shape in an online fashion; the Poisson rate of the fish school will also be estimated to

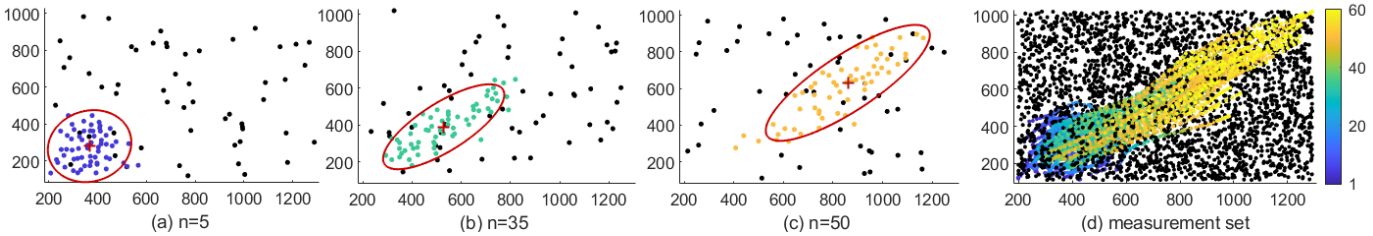

Fig. 6: Total measurement set and estimates at time step 5, 35 and 50; black dots are clutters; the colorbar indicates the time step these measurements are received at and colored dots are measurements of the fish school with the color representing the time step corresponding to the colorbar; red crosses and circles are estimated group center and shape of fish school, respectively.

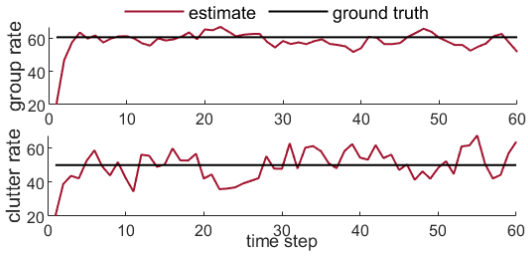

Fig. 7: Estimated Poisson rates of the fish school and clutters.

obtain the density of the fish school. Both the target rate and the clutter rate are unknown to users. Here we implement the adaptive RB-AbNHPP tracker with a measurement rate and target shape model in Section V, where specifically we adopt the GIG Markov chain prior model for Poisson rates.

The estimation of fish school shape, its group center position, and its Poisson rate are shown in Fig. 6 and Fig. 7. Note that as the camera gives one measurement per fish, we can assume the Poisson rate of each fish being one, and thus the ground truth Poisson rate of the whole fish school is 61. From the results, the estimated rate of fish school has a mean around 61, which corresponds to the true number of the fish school. Therefore, we can see that our method can not only track the group center and shape online but also provide us with useful information on the approximated number of fishes in the group based on the estimated Poisson rate.

VII. Conclusion

This paper develops a robust Bayesian multi-object tracker that can estimate online the target states and shapes in conjunction with measurement rates and association variables. A fast Rao-Blackwellisation scheme for linear Gaussian models has been devised which shows great improvement in both the speed and accuracy of the algorithm. We have introduced the powerful GIG family of distributions to model time-varying measurement rates, offering a general solution to the modelling of the time evolution of measurement rates. Future work will include cases with a time-varying number of targets and detection environments with a non-uniform clutter map.

References

[1] K. Gilholm, S. Godsill, S. Maskell, and D. Salmond, "Poisson models for extended target and group tracking," in *Signal and Data Processing of Small Targets 2005*, vol. 5913. SPIE, 2005, pp. 230–241.

[2] K. Granström and U. Orguner, "Estimation and maintenance of measurement rates for multiple extended target tracking," in *2012 15th International Conference on Information Fusion*. IEEE, 2012, pp. 2170–2176.

[3] M. Feldmann, D. Fränken, and W. Koch, "Tracking of extended objects and group targets using random matrices," *IEEE Transactions on Signal Processing*, vol. 59, no. 4, pp. 1409–1420, 2010.

[4] 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, 2019.

[5] S. Yang, K. Thormann, and M. Baum, "Linear-time joint probabilistic data association for multiple extended object tracking," in *2018 IEEE 10th Sensor Array and Multichannel Signal Processing Workshop (SAM)*. IEEE, 2018, pp. 6–10.

[6] F. Meyer and J. L. Williams, "Scalable detection and tracking of geometric extended objects," *IEEE Transactions on Signal Processing*, vol. 69, pp. 6283–6298, 2021.

[7] R. Gan, Q. Li, and S. Godsill, "A variational Bayes association-based multi-object tracker under the non-homogeneous Poisson measurement process," in *2022 25th International Conference on Information Fusion (FUSION)*. IEEE, 2022, pp. 1–8.

[8] Y. Bar-Shalom, T. E. Fortmann, and P. G. Cable, "Tracking and data association," 1990.

[9] D. Reid, "An algorithm for tracking multiple targets," *IEEE transactions on Automatic Control*, vol. 24, no. 6, pp. 843–854, 1979.

[10] R. L. Streit and T. E. Luginbuhl, "Probabilistic multi-hypothesis tracking," Naval Underwater Systems Center Newport RI, Tech. Rep., 1995.

[11] S. Davey, "Probabilistic multi-hypothesis tracker with an evolving Poisson prior," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 51, no. 1, pp. 747–759, 2015.

[12] C. Hue, J.-P. Le Cadre, and P. Pérez, "Tracking multiple objects with particle filtering," *IEEE transactions on aerospace and electronic systems*, vol. 38, no. 3, pp. 791–812, 2002.

[13] F. Meyer, T. Kropfreiter, J. L. Williams, R. Lau, F. Hlawatsch, P. Braca, and M. Z. Win, "Message passing algorithms for scalable multitarget tracking," *Proceedings of the IEEE*, vol. 106, no. 2, pp. 221–259, 2018.

[14] Q. Li, J. Sun, and W. Sun, "An efficient multiple hypothesis tracker using max product belief propagation," in *2017 20th International Conference on Information Fusion (Fusion)*. IEEE, 2017, pp. 1–6.

[15] S. Särkkä, A. Vehtari, and J. Lampinen, "Rao-Blackwellized particle filter for multiple target tracking," *Information Fusion*, vol. 8, no. 1, pp. 2–15, 2007.

[16] W. Ng, J. Li, S. Godsill, and S. K. Pang, "Multitarget initiation, tracking and termination using Bayesian Monte Carlo methods," *The Computer Journal*, vol. 50, no. 6, pp. 674–693, 2007.

[17] F. Septier, S. K. Pang, A. Carmi, and S. Godsill, "On MCMC-based particle methods for Bayesian filtering: Application to multitarget tracking," in *2009 3rd IEEE International Workshop on Computational Advances in Multi-Sensor Adaptive Processing*. IEEE, 2009, pp. 360–363.

[18] B.-n. Vo, M. Mallick, Y. Bar-Shalom, S. Coraluppi, R. Osborne, R. Mahler, and B.-t. Vo, "Multitarget tracking," *Wiley encyclopedia of electrical and electronics engineering*, no. 2015, 2015.

[19] B. Habtemariam, R. Tharmarasa, T. Thayaparan, M. Mallick, and T. Kirubarajan, "A multiple-detection joint probabilistic data association filter," *IEEE Journal of Selected Topics in Signal Processing*, vol. 7, no. 3, pp. 461–471, 2013.

[20] S. P. Coraluppi and C. A. Carthel, "Multiple-hypothesis tracking for targets producing multiple measurements," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 54, no. 3, pp. 1485–1498, 2018.

[21] G. Vivone and P. Braca, "Joint probabilistic data association tracker for extended target tracking applied to x-band marine radar data," *IEEE Journal of Oceanic Engineering*, vol. 41, no. 4, pp. 1007–1019, 2016.

[22] K. Granström, L. Svensson, S. Reuter, Y. Xia, and M. Fatemi, "Likelihood-based data association for extended object tracking using sampling methods," *IEEE Transactions on intelligent vehicles*, vol. 3, no. 1, pp. 30–45, 2017.
[23] F. Meyer and M. Z. Win, "Scalable data association for extended object tracking," *IEEE Transactions on Signal and Information Processing over Networks*, vol. 6, pp. 491–507, 2020.
[24] K. Gilholm and D. Salmond, "Spatial distribution model for tracking extended objects," *IEE Proceedings-Radar, Sonar and Navigation*, vol. 152, no. 5, pp. 364–371, 2005.
[25] S. Yang, L. M. Wolf, and M. Baum, "Marginal association probabilities for multiple extended objects without enumeration of measurement partitions," in *23rd International Conference on Information Fusion*. IEEE, 2020, pp. 1–8.
[26] Q. Li, J. Liang, and S. Godsill, "Scalable data association and multi-target tracking under a Poisson mixture measurement process," in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 5503–5507.
[27] S. Godsill, "Particle filtering: the first 25 years and beyond," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 7760–7764.
[28] M. P. S. Gander and D. A. Stephens, "Stochastic volatility modelling in continuous time with general marginal distributions: Inference, prediction and model selection," *Journal of Statistical Planning and Inference*, vol. 137, no. 10, pp. 3068–3081, 2007.
[29] A. Doucet, S. Godsill, and C. Andrieu, "On sequential Monte Carlo sampling methods for Bayesian filtering," *Statistics and computing*, vol. 10, no. 3, pp. 197–208, 2000.
[30] C. Andrieu and A. Doucet, "Particle filtering for partially observed Gaussian state space models," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 64, no. 4, pp. 827–836, 2002.
[31] R. Gan, J. Liang, B. I. Ahmad, and S. Godsill, "Bayesian intent prediction for fast maneuvering objects using variable rate particle filters," in *2019 IEEE 29th International Workshop on Machine Learning for Signal Processing (MLSP)*. IEEE, 2019, pp. 1–6.
[32] R. Gan, B. I. Ahmad, and S. J. Godsill, "Lévy state-space models for tracking and intent prediction of highly maneuverable objects," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 57, no. 4, 2021.
[33] A. Finke, A. Doucet, and A. M. Johansen, "Limit theorems for sequential MCMC methods," *Advances in Applied Probability*, vol. 52, no. 2, pp. 377–403, 2020.
[34] Q. Li, B. I. Ahmad, and S. J. Godsill, "Sequential dynamic leadership inference using Bayesian Monte Carlo methods," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 57, no. 4, pp. 2039–2052, 2021.
[35] Q. Li, S. J. Godsill, J. Liang, and B. I. Ahmad, "Inferring dynamic group leadership using sequential Bayesian methods," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 8911–8915.
[36] J. Koch, "Bayesian approach to extended object and cluster tracking using random matrices," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 44, 2008.
[37] S. Särkkä and A.Solin, *Applied Stochastic Differential Equations*, ser. Institute of Mathematical Statistics Textbooks. Cambridge University Press, 2019.
[38] P. Bromiley, "Products and convolutions of Gaussian probability density functions," *Tina-Vision Memo*, vol. 3, no. 4, p. 1, 2003.
[39] P. Willett, Y. Ruan, and R. Streit, "PMHT: Problems and some solutions," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 38, no. 3, pp. 738–754, 2002.
[40] I. Schlangen, E. Delande, J. Houssineau, and D. E. Clark, "A PHD filter with negative binomial clutter," in *2016 19th International Conference on Information Fusion (FUSION)*. IEEE, 2016, pp. 658–665.
[41] A. Gelman, J. B. Carlin, H. S. Stern, and D. B. Rubin, *Bayesian data analysis*. Chapman and Hall/CRC, 1995.
[42] D. Schuhmacher, B.-T. Vo, and B.-N. Vo, "A consistent metric for performance evaluation of multi-object filters," *IEEE transactions on signal processing*, vol. 56, no. 8, pp. 3447–3457, 2008.
[43] A. S. Rahmathullah, Á. F. García-Fernández, and L. Svensson, "Generalized optimal sub-pattern assignment metric," in *2017 20th International Conference on Information Fusion*. IEEE, 2017, pp. 1–8.
[44] S. Yang, M. Baum, and K. Granström, "Metrics for performance evaluation of elliptic extended object tracking methods," in *2016 IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI)*. IEEE, 2016, pp. 523–528.
[45] A. Strandburg-Peshkin, C. R. Twomey, N. W. Bode, A. B. Kao, Y. Katz, C. C. Ioannou, S. B. Rosenthal, C. J. Torney, H. S. Wu, S. A. Levin *et al.*, "Visual sensory networks and effective information transfer in animal groups," *Current Biology*, vol. 23, no. 17, pp. R709–R711, 2013.

Qing Li is a PhD student in the Signal Processing and Communications Laboratory at the Department of Engineering, Cambridge University, U.K. since 2018 and a research assistant at the Engineering Department, Cambridge University. She received the M.E. degree in information and communication engineering from Beihang University, China, in 2018. Her research interests include Monte Carlo computational methods, Bayesian statistical inference, sensor fusion, and tracking.

Runze Gan is currently a PhD student and research assistant at the Engineering Department, Cambridge University. He received B.Eng. degree in Electronic and Electrical Engineering from University of Birmingham, U.K., and Huazhong University of Science and Technology, China, in 2018. His research focuses on stochastic modelling for time series and approximate inference methods (particularly Monte Carlo methods and variational inference), with applications in intent inference and object tracking.

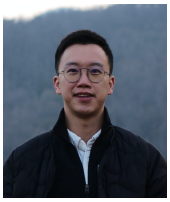

Jiaming Liang received the M.Phil. degree in advanced computer science in 2015, the Ph.D. degree in information engineering in 2020, both from Cambridge University, U.K., where he has worked as a postdoctoral research associate since graduation. His research interests lie primarily in Bayesian statistical methods and probabilistic modelling, with applications in target(s) tracking, intentionality estimation and intelligent vehicles.

Simon J. Godsill is Professor of Statistical Signal Processing in the Engineering Department at Cambridge University. He is also a Professorial Fellow and tutor at Corpus Christi College Cambridge. He coordinates an active research group in Signal Inference and its Applications and is Head of the Information Engineering Division at Cambridge. His group specialises in Bayesian computational methodology, including multiple object tracking, audio and music processing, and financial time series modeling. A particular methodological theme over recent years has been the development of novel techniques for optimal Bayesian filtering and smoothing, using Sequential Monte Carlo or Particle Filtering methods. Prof. Godsill has published extensively in journals, books and international conference proceedings, and has given a number of high profile invited and plenary addresses at conferences such as the Valencia conference on Bayesian Statistics, the IEEE Statistical Signal Processing Workshop, the Conference on Bayesian Inference for Stochastic Processes (BISP), the IEEE Workshop on Machine Learning in Signal Processing (2013) and FUSION (2016). He co-authored a Springer text Digital Audio Restoration with Prof. Peter Rayner in 1998. He was technical chair of the IEEE NSSPW workshop in 2006 on sequential and nonlinear filtering methods, and has been on the conference panel for numerous other conferences/workshops. Prof. Godsill has served as Associate Editor for IEEE Tr. Signal Processing and the journal Bayesian Analysis. He was Theme Leader in Tracking and Reasoning over Time for the UK's Data and Information Fusion Defence Technology Centre (DIF-DTC) and Principal Investigator on many grants funded by the EU, EPSRC, QinetiQ, General Dynamics, MOD, Microsoft UK, Citibank, Mastercard, Google, DSO Singapore, Huawei and Jaguar Landrover. In 2009-10 he was co-organiser of an 18 month research program in Sequential Monte Carlo Methods at the SAMSI Institute in North Carolina and in 2014 he co-organised a research programme at the Isaac Newton Institute on Sequential Monte Carlo methods. In 2018 he was General Chair of the FUSION Conference in Cambridge. Two of his journal papers have recently received Best Paper awards from the IEEE and IET. He continues to be a Director of CEDAR Audio Ltd. (which has received numerous accolades over the years, including a technical Oscar), and for which he was a founding staff member in 1988. The company has commercialised many of the ideas from Professor Godsill's research over the years.