EasyChair Preprint

№ 10464

A Scalable Rao-Blackwellised Sequential MCMC Sampler for Joint Detection and Tracking in Clutter

Qing Li, Runze Gan and Simon Godsill

June 28, 2023

A Scalable Rao-Blackwellised Sequential MCMC Sampler for Joint Detection and Tracking in Clutter

Qing Li, Runze Gan, Simon Godsill
Engineering Department
University of Cambridge
Cambridge, UK
{ql289, rg605, sjg30}@cam.ac.uk

***Abstract*—This paper addresses the joint detection and tracking of an unknown and time-varying number of targets in clutter. Here we formulate the tracking task in a variable-dimension state space, under which the reversible jump sequential Markov chain Monte Carlo sampling methods can be utilised to online estimate the target number, their kinematic states, and the association variables. In particular, a fast Rao-Blackwellisation scheme is devised to improve the tracking accuracy and sampling efficiency for linear Gaussian models. Based on the non-homogeneous Poisson process measurement model, the developed tracker enjoys a partially parallel sampling structure, thereby being able to efficiently tackle the data association under massive measurements and clutter. The simulation results demonstrate that the developed tracker exhibits superior tracking performance in comparison to existing trackers in both accuracy and computational efficiency when tracking multiple targets under heavy clutter.**

***Index Terms*—data association, target tracking, detection, Rao-Blackwellisation, reversible jump sequential MCMC**

I. Introduction

In practical tracking applications, targets may frequently appear or disappear from the field of view due to the limited view of the sensor, making existing trackers that assume a fixed number of targets inapplicable in this case [1], [2]. One traditional strategy to address track formation and termination is to employ statistical test methods such as the sequential probability ratio test (SPRT), in which a log-likelihood ratio is calculated to assess the validity of the track [3]. Likewise, the estimates of target rates are employed to perform statistical tests to confirm the target appearance and disappearance in the probabilistic multiple hypothesis tracker (PMHT) [4].

Another way to handle joint detection and tracking is with random finite set (RFS) methods, which include the recently popular probabilistic hypothesis density (PHD) [5] and the Poisson multiple Bernoulli mixture (PMBM) filters [6]. In contrast to hypothesis testing methods [3], [4] that have no explicit target birth or death probabilistic model, RFS-based approaches require assigning a specific target birth and death prior in a Bayesian recursion, where the two most commonly used birth priors are the Poisson and Multi-Bernoulli birth models, combined with a target death process modelled by survival probabilities. Nonetheless, RFS-based filters often require approximations and heuristic designs in exchange for a faster implementation; typical examples are the Gaussian Mixture implementation of the PHD filter [7] and the PMBM filter with heuristic pruning, gating, and hypothesis management to limit the exponential increase in the global hypotheses [6].

Besides the RFS-based methods, a more straightforward solution is to directly infer the joint posterior distribution of the target number/existence and states under the Bayesian inference framework. To model the target birth and death, the notion of existence was introduced in [8], [9], which is a binary vector to indicate the existence of each target at each time step. Several implementations have been developed under the existence vector representation, including Monte Carlo sampling methods [10], [11] and graphical model approaches [12]. For instance, in [10], the birth and death of each target were modelled by a Bernoulli process, under which a sequential Markov chain Monte Carlo (SMCMC) sampling algorithm was developed to detect and track multiple point targets. Specifically, it avoided the data association problem by applying a non-homogeneous Poisson process (NHPP) measurement model so that the likelihood function can be evaluated in closed form in the SMCMC implementation. Despite being association-free, it could be inefficient since the sampling structure cannot be parallelised as in [13]; in addition, it is also inapplicable to a Rao-Blackwellisation scheme as in [1] for linear Gaussian models. In [12], a Poisson birth process is assumed based on the existence vector representation, and target detection is performed by calculating the marginal posterior existence probabilities and comparing them with a threshold. An approximate particle-based implementation of the sum-product algorithm (SPA) was developed for tracking an unknown number of extended targets. However, using the existence variable may be computationally inefficient due to the need to specify a maximum target number that exceeds the actual target number, meaning that part of the calculation is ineffectual.

Alternatively, the evolution of the target number, the kinematic states, and the associations can be modelled by a Markov switching state space model, also known as a jump Markov system. In [14], a general particle filter framework was outlined for point target tracking, while no details were provided for design proposals. Another particle filter-based method was devised in [15] for tracking an unknown number of point targets, where it assumed a simpler and perhaps more realistic birth/death model that restricted changes in the number of targets to no more than one at a given time step.

In particular, it adopted a M-best 2-D assignment algorithm to approximate the marginal likelihood required in each particle weight calculation. The same target birth and death model was used in [11], and a Rao-Blackwellisation particle filter was devised for linear Gaussian models to improve the sampling efficiency. However, due to the degradation of particle filters on high-dimensional problems, implementations based on particle filters, e.g., [11], [15], would suffer from degraded tracking accuracy, thus limiting their capability to track large numbers of targets. For these variable dimension tasks, the reversible jump MCMC (RJ-SMCMC) method was first presented by Green in [16], and a more general framework was subsequently presented in [17], which is an advantageous alternative to the particle filter in high-dimensional problems. In the aspect of tracking applications, this reversible jump MCMC approach has previously been utilised to jointly estimate the number and state of targets in [18]; nevertheless, it does not take into account the problem of data association that arises when tracking multiple targets in clutter.

Therefore, this paper presents a fast Rao-Blackwellised RJ-SMCMC (RB-RJ-SMCMC) sampler that can jointly estimate the target number and target kinematic states under data association uncertainty and clutter for linear Gaussian models. The developed fast Rao-Blackwellisation scheme is different from standard Rao-Blackwellisation schemes that are implemented by sampling from a reduced sampling space, e.g., [19], [20]; in contrast, it holds the same purpose as [1], which samples from the original sampling space that includes the target states to keep the parallel sampling features, while reconstructing the estimates of the target state using the obtained samples of non-linear states to enjoy the advantage of Rao-Blackwellisation. Based on the previous work [1], here we consider the target number uncertainty, and a RB-RJ-SMCMC sampler is devised to handle the joint detection and tracking task. Specifically, we employ a target birth and death model in which at most one target can be added or removed by each time step, the same as the birth model in [15]. Other birth and death models can be easily incorporated into our framework due to the flexibility of our sampling scheme. Compared to the standard SMCMC approach in [10], our RB-RJ-SMCMC framework does not require assigning a maximum number of targets and sampling the whole vector of the target states. In addition, we develop a more efficient Rao-Blackwellisation scheme and tackle the data association problem by leveraging the scalable association framework [1], which guarantees a partially parallel execution of the RB-RJ-SMCMC sampling scheme and avoids the expensive summations in the likelihood function in [10], thus making it advantageous in adverse conditions with heavy clutter. In comparison to approximated methods [6], [12], it can theoretically converge to the Bayesian optimal filter when the sample size is large. We verify the effectiveness of the proposed tracker in two simulation scenarios.

II. Problem Formulation and Modelling

This paper considers tracking an unknown and time-varying number of targets in clutter. Here we formulate the problem in a Bayesian setting, where we seek to recursively estimate the joint posterior of the target number, the target state, and the association uncertainty. Specifically, we define a variable dimension target state $X_n = [X_{n,1}, X_{n,2}, ..., X_{n,K_n}]^T$, in which K_n denotes the number of objects. Given the sequence of observations from time step 1 to n, $Z_{1:n}$, the associations at time step n are denoted by $\theta_n = [\theta_{n,1}, ..., \theta_{n,M_n}]$, where M_n is the total measurement number. Each component $\theta_{n,j}, j \in \{1, ..., M_n\}$ gives the origin of the measurement $Z_{n,j}$; $\theta_{n,j} = 0$ indicates that $Z_{n,j}$ is a clutter, and $\theta_{n,j} = i, i \in \{1, \ldots, K_n\}$ means that $Z_{n,j}$ is generated from target i. In the subsequent sections, we will present the state space model, the transition densities and the association prior.

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 $X^d_{n,i} = [x^d_{n,i}, \dot{x}^d_{n,i}]^T$, which contains the target's position and velocity. We assume a linear Gaussian dynamical model for the kinematic state $X^d_{n,i}$, represented as

$$X^d_{n,i} = F^d_i X^d_{n-1,i} + w^d_{n,i} \tag{1}$$

where $w^d_{n,i} \sim \mathcal{N}(0, P^d_i)$. Accordingly, the transition density at time step n is $p(X^d_{n,i}|X^d_{n-1,i}) = \mathcal{N}(F^d_i X^d_{n-1,i}, P^d_i)$. The joint transition density for all dimensions can be computed as $p(X_{n,i}|X_{n-1,i}) = \prod_{d=1}^{D} p(X^d_{n,i}|X^d_{n-1,i})$.

For a constant velocity (CV) model, its transition matrix and noise covariance are

$$F^d_i = \begin{bmatrix} 1 & \tau \\ 0 & 1 \end{bmatrix}, P^d_i = Q_i \begin{bmatrix} \tau^3/3 & \tau^2/2 \\ \tau^2/2 & \tau \end{bmatrix}. \tag{2}$$

where τ is the time interval between time steps.

B. Target Birth and Death Model

In this part, we describe the target birth and death under a variable dimensional state space model. For simplicity, the newly-born target state is always added to the end of the target state vector of the previous time step. In addition, we define a target death identifier I_n to indicate the target death event, and I_n records the disappearing target index from the previous target state X_{n-1} whenever a target death happens, where $I_n \in \{0, 1, ..., K_{n-1}\}$ and $I_n = 0$ denotes that no target dies.

We model the evolution as the first-order Markov chain and the joint transition density $p(X_n, K_n, I_n|X_{n-1}, K_{n-1}, I_{n-1})$ can be factored as:

$$\begin{aligned} p(X_n, K_n, I_n|X_{n-1}, K_{n-1}, I_{n-1}) &= p(K_n|K_{n-1}) \\ &\times p(I_n|K_n, K_{n-1}) p(X_n|X_{n-1}, K_n, K_{n-1}, I_n) \end{aligned} \tag{3}$$

where $p(K_n|K_{n-1})$ is the transition density of the target cardinality that implicitly models the birth, death and survival events of the targets, and $p(I_n|K_n, K_{n-1})$ is the prior distribution of the target death identifier which gives the probability of choosing an index I_n from the existing target indices to delete when a target death event happens. $p(X_n|X_{n-1}, K_n, K_{n-1}, I_n)$ defines the motion model under the target birth, death and survival events.

For the consideration of practicality, the dimension change at time step n is limited to a single move to the adjacent dimension from the dimension of the last time step $n-1$, and this simplified assumption is widely adopted in literature, e.g., [15], [21]. To be specific, we assume that at each time step, only one of the three events can happen, i.e., one new target appearing, one existing target disappearing, and all targets surviving. Under such an assumption, at most one target can appear or disappear at a given time step. Therefore, the transition density $p(K_n|K_{n-1})$ is defined as:

$$p(K_n|K_{n-1}) = \begin{cases} P_b & K_n = K_{n-1}+1; \\ 1-P_b-P_d & K_n = K_{n-1}; \\ P_d & K_n = K_{n-1}-1; \\ 0 & \text{otherwise} \end{cases} \tag{4}$$

where P_b and P_d denote the probability of a target appearing and disappearing, respectively.

The transition of the target death identifier $p(I_n|K_n, K_{n-1})$ is time independent of the previous I_{n-1}. A simple choice is a discrete uniform distribution with its probability mass function being $\frac{1}{K_{n-1}}$ and the support being $\{1, ..., K_{n-1}\}$ when target death happens; otherwise $I_n = 0$ with probability one which indicates no target death at time step n:

$$\begin{aligned} &p(I_n|K_n, K_{n-1}) \\ &= \begin{cases} \frac{1}{K_{n-1}} \sum_{i_D=1}^{K_{n-1}} [I_n = i_D] & K_n = K_{n-1}-1; \\ [I_n = 0] & \text{otherwise} \end{cases} \end{aligned} \tag{5}$$

where $[I_n = i_D]$ is 1 if $I_n = i_D$, and 0 otherwise. Alternatively, a non-uniform prior can be designed to assign a larger possibility of deleting one or several targets by monitoring, e.g., the covariance of the estimates.

The transition density $p(X_n|X_{n-1}, K_n, K_{n-1}, I_n)$ is defined based on different events:

$$\begin{aligned} &p(X_n|X_{n-1}, K_n, K_{n-1}, I_n) \\ &= \begin{cases} p_b(X_n|X_{n-1}, K_n) & K_n = K_{n-1}+1; \\ p_s(X_n|X_{n-1}, K_n) & K_n = K_{n-1}; \\ p_d(X_n|X_{n-1}, K_n, I_n) & K_n = K_{n-1}-1; \\ 0 & \text{otherwise} \end{cases} \end{aligned} \tag{6}$$

Here we assume that targets move independently such that the transition density can be further partitioned into each target's transition density as follows:

- **Target birth event**: When the target birth event happens with a probability of P_b at time step n, the target number $K_n = K_{n-1}+1$, and $I_n = 0$. The transition density can be partitioned into surviving targets and new-born targets. Under the assumption that the birth and survival events are mutually independent, we have

$$p_b(X_n|X_{n-1}, K_n) = p_1(X_{n,K_n}) \prod_{i=1}^{K_n-1} p(X_{n,i}|X_{n-1,i}) \tag{7}$$

  where X_{n,K_n} denotes the newborn target, stacked in the K_n-th of X_n with the first $X_{n,1:K_{n-1}}$ being the surviving targets. The transition density $p(X_{n,i}|X_{n-1,i})$ for each surviving target is defined by the motion model, and $p_1(X_{n,K_n})$ is the state prior distribution specified according to various tracking scenarios. In this paper, we assume a Gaussian distribution prior $p_1(X_{n,K_n}) = \mathcal{N}(\mu_b, C_b)$, where μ_b and C_b are defined such that the Gaussian distribution can cover the interested area, e.g., the target birth zone or the whole surveillance region.
- **Target survival event**: Under this event, the dimension is invariant to the last time step $n-1$. Therefore, the target transition density can be directly deduced by:

$$p_s(X_n|X_{n-1}, K_n) = \prod_{i=1}^{K_n} p(X_{n,i}|X_{n-1,i}) \tag{8}$$

- **Target death event**: The target death process involves terminating a target according to I_n. Hence, the transition density under the target death event is

$$\begin{aligned} &p_d(X_n|X_{n-1}, K_n, I_n) \\ &= \prod_{i=1}^{I_n-1} p(X_{n,i}|X_{n-1,i}) \prod_{i^*=I_n}^{K_n} p(X_{n,i^*}|X_{n-1,i^*+1}) \end{aligned} \tag{9}$$

C. Association prior and measurement model

Here, we adopt the NHPP measurement model presented in [1]. Denote the set of Poisson rates by $\Lambda = \{\Lambda_i; i = 0, 1, ..., K_n\}$, where Λ_0 is the clutter rate, and Λ_i is the i-th target rate, $i = 1, ..., K_n$. We assume the measurement process of each target i is an NHPP with a Poisson distributed measurement number with rate Λ_i. The clutter process is a homogeneous Poisson process (HPP) with Poisson rate Λ_0. The total measurement process is the superposition of the conditional independent NHPP/HPP measurement processes from K_n targets and the clutter, and the total measurement number follows a Poisson distribution with rate $\Lambda_s = \sum_{i=0}^{K_n} \Lambda_i$. Under the assumption of the NHPP measurement model, the measurements are conditionally independent when conditional on the measurement number M_n and target state X_n,

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

The M_n measurements $\{Z_{n,j}\}_{j=1}^{M_n}$ along with associations $\{\theta_{n,j}\}_{j=1}^{M_n}$ are also conditionally independent:

$$p(Z_n|X_n, \theta_n, M_n) = \prod_{j=1}^{M_n} p(Z_{n,j}|X_{n,\theta_{n,j}}) \tag{11}$$

Here 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:

$$p(Z_{n,j}|X_{n,i}) = \begin{cases} \mathcal{N}(HX_{n,i}, R_i), & i \neq 0; \quad \text{(object)} \\ \frac{1}{V}, & i = 0; \quad \text{(clutter)} \end{cases} \tag{12}$$

where H is the observation matrix. R_i indicates the noise covariance or the Gaussian target extent. Non-uniformed clutter and non-Gaussian models can be readily included if required.

The joint prior $p(\theta_n|K_n, M_n)$ can be calculated from the product of M_n independent association priors:

$$p(\theta_n|K_n, M_n) = \prod_{j=1}^{M_n} p(\theta_{n,j}|K_n) \tag{13}$$

where the prior for each association $p(\theta_{n,j}|K_n)$ is a categorical distribution with support $\theta_{n,j} \in \{0, .., K_n\}$

$$p(\theta_{n,j}|K_n) = \prod_{i=0}^{K_n} \left(\frac{\Lambda_i}{\Lambda_s}\right)^{[\theta_{n,j}=i]}. \tag{14}$$

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

$$p(\theta_{n,j}|Z_{n,j}, X_n) = \left(\frac{\Lambda_0}{V\tilde{l}}\right)^{[\theta_{n,j}=0]} \prod_{i=1}^{K_n} \left(\frac{\Lambda_i l_{ij}}{\tilde{l}}\right)^{[\theta_{n,j}=i]}, \tag{15}$$

where $l_{ij} = \mathcal{N}(Z_{n,j}; HX_{n,i}, R_i)$ under the linear Gaussian measurement model, and $\tilde{l}$ is a normalisation constant to ensure the sum of all categories' probabilities equals one. Note that since K_n is the cardinality of X_n, that is, $K_n = |X_n|$, the condition on K_n is omitted whenever it is conditional on X_n, and the same applies to M_n, where $M_n = |Z_n|$.

III. Rao-Blackwellised Reversible Jump Sequential MCMC Sampler

This section presents a Rao-Blackwellised reversible jump SMCMC (RB-RJ-SMCMC) sampler for linear Gaussian models where only nonlinear or non-Gaussian states are sampled by the RJ-SMCMC. We adopt a Monte Carlo approximation of the posterior distribution $p(\theta_{1:n}, K_{1:n}, I_{1:n}|Z_{1:n})$, which, at each time step n, is approximated by a set of N_p unweighted samples $\{\theta_{1:n}^{(p)}, K_{1:n}^{(p)}, I_{1:n}^{(p)}\}_{p=1}^{N_p}$. According to the marginal-conditional decomposition, we have

$$\begin{aligned} &p(X_n, \theta_{1:n}, K_{1:n}, I_{1:n}|Z_{1:n}) \\ &= p(\theta_{1:n}, K_{1:n}, I_{1:n}|Z_{1:n}) p(X_n|\theta_{1:n}, K_{1:n}, I_{1:n}, Z_{1:n}) \end{aligned} \tag{16}$$

Subsequently, the marginal posterior of target state X_n can be calculated as a Gaussian mixture by using the empirical distribution of $p(\theta_{1:n}, K_{1:n}, I_{1:n}|Z_{1:n})$:

$$p(X_n|Z_{1:n}) \approx \frac{1}{N_p} \sum_{p=1}^{N_p} p(X_n|\theta_{1:n}^{(p)}, K_{1:n}^{(p)}, I_{1:n}^{(p)}, Z_{1:n}), \tag{17}$$

where each $p(X_n|\theta_{1:n}^{(p)}, K_{1:n}^{(p)}, I_{1:n}^{(p)}, Z_{1:n})$ is a Gaussian that can be updated by Kalman filtering. In the following, we presented the detailed derivations. First, we have

$$\begin{aligned} &p(X_n|\theta_{1:n}^{(p)}, K_{1:n}^{(p)}, I_{1:n}^{(p)}, Z_{1:n}) \\ &\propto p(Z_n|X_n, \theta_n^{(p)}, M_n) p(X_n|K_{1:n}^{(p)}, I_{1:n}^{(p)}, \theta_{1:n-1}^{(p)}, Z_{1:n-1}) \end{aligned} \tag{18}$$

The predicted density $p(X_n|K_{1:n}^{(p)}, I_{1:n}^{(p)}, \theta_{1:n-1}^{(p)}, Z_{1:n-1})$ is in the form of a Gaussian distribution, when assuming a Gaussian initial prior distribution and transition densities in (6)

$$\begin{aligned} &p(X_n|K_{1:n}^{(p)}, I_{1:n}^{(p)}, \theta_{1:n-1}^{(p)}, Z_{1:n-1}) \\ &= \int p(X_n, X_{n-1}|K_{1:n}^{(p)}, I_{1:n}^{(p)}, \theta_{1:n-1}^{(p)}, Z_{1:n-1}) dX_{n-1} \\ &= \int p(X_n|X_{n-1}, K_n^{(p)}, K_{n-1}^{(p)}, I_n^{(p)}) \\ &\quad \times p(X_{n-1}|\theta_{1:n-1}^{(p)}, K_{1:n-1}^{(p)}, I_{1:n-1}^{(p)}, Z_{1:n-1}) dX_{n-1} \\ &= \prod_{i=1}^{K_n^{(p)}} \mathcal{N}(X_{n,i}; \mu_{n|n-1,i}^{(p)}, \Sigma_{n|n-1,i}^{(p)}) \\ &= \begin{cases} \mathcal{N}(X_{n,K_n^{(p)}}; \mu_b, P_b) & K_n^{(p)} = K_{n-1}^{(p)} + 1; \\ \times \prod_{i=1}^{K_{n-1}^{(p)}} \mathcal{N}(X_{n,i}; \mu_{n|n-1,i}^{(p)}, \Sigma_{n|n-1,i}^{(p)}) & \\ \prod_{i=1}^{K_n^{(p)}} \mathcal{N}(X_{n,i}; \mu_{n|n-1,i}^{(p)}, \Sigma_{n|n-1,i}^{(p)}) & K_n^{(p)} = K_{n-1}^{(p)}; \\ \prod_{i=1}^{I_n^{(p)}-1} \mathcal{N}(X_{n,i}; \mu_{n|n-1,i}^{(p)}, \Sigma_{n|n-1,i}^{(p)}) & K_n^{(p)} = K_{n-1}^{(p)} - 1; \\ \times \prod_{i^*=I_n^{(p)}}^{K_n^{(p)}} \mathcal{N}(X_{n,i^*}; \mu_{n|n-1,i^*+1}^{(p)}, \Sigma_{n|n-1,i^*+1}^{(p)}) & \end{cases} \end{aligned} \tag{19}$$

where

$$\mu_{n|n-1,i}^{(p)} = F_i \mu_{n-1|n-1,i}^{(p)} \tag{20}$$

$$\Sigma_{n|n-1,i}^{(p)} = F_i \Sigma_{n-1|n-1,i}^{(p)} F_i^\top + P_i. \tag{21}$$

Subsequently, the marginal posterior in (18) has a form of:

$$\begin{aligned} &p(X_n|\theta_{1:n}^{(p)}, K_{1:n}^{(p)}, I_{1:n}^{(p)}, Z_{1:n}) \\ &\propto \prod_{i=1}^{K_n^{(p)}} \mathcal{N}(\tilde{Z}_n^i; HX_{n,i}, \tilde{R}_i) \mathcal{N}(X_{n,i}; \mu_{n|n-1,i}^{(p)}, \Sigma_{n|n-1,i}^{(p)}) \\ &\propto \prod_{i=1}^{K_n^{(p)}} \mathcal{N}(X_{n,i}; \mu_{n|n,i}^{(p)}, \Sigma_{n|n,i}^{(p)}) \end{aligned} \tag{22}$$

where $\tilde{Z}_n^i$ and $\tilde{R}_i$ are computed as follows, with $\Theta_n^i = \{j | j \in \{1, ..., M_n\}, \theta_{n,j}^{(p)} = i\}$, and $|\cdot|$ denotes the set cardinality:

$$\tilde{Z}_n^i = \frac{1}{|\Theta_n^i|} \sum_{j \in \Theta_n^i} Z_{n,j}, \quad \tilde{R}_i = \frac{1}{|\Theta_n^i|} R_i, \tag{23}$$

The $\mu_{n|n,i}^{(p)}, \Sigma_{n|n,i}^{(p)}$ in (22) are then updated by Kalman filtering. From the final form of the conditional $p(X_n|\theta_{1:n}^{(p)}, K_{1:n}^{(p)}, I_{1:n}^{(p)}, Z_{1:n})$, we can see that each object $X_{n,i}$ can be updated independently.

A. A Parallel Implementation of the RB-RJ-SMCMC Sampler

To keep the parallel sampling efficiency, here we adopt a novel Rao-Blackwellisation scheme similar to [1], where we first perform the joint RJ-SMCMC to sample from the posterior $p(X_n, \theta_{1:n}, K_{1:n}, I_{1:n}|Z_{1:n})$; after it converges, we only keep the samples of $\theta_{1:n}, K_{1:n}, I_{1:n}$ to form the empirical distribution $p(\theta_{1:n}, K_{1:n}, I_{1:n}|Z_{1:n})$, and the marginal posterior of the target state X_n can be calculated as a Gaussian mixture distribution by (17).

To begin with, we present the derivation of the joint RJ-SMCMC sampler with its target distribution being the joint distribution $p(X_n, \theta_{1:n}, K_{1:n}, I_{1:n}|Z_{1:n})$. Assume that the tracking task has been solved at the previous time step $n-1$, where the posterior distribution $p(X_{n-1}, \theta_{1:n-1}, K_{1:n-1}, I_{1:n-1}|Z_{1:n-1})$ is approximated by an unweighted sequence of N_p samples from the convergent Markov chain $\{X_{n-1}^{(p)}, \theta_{1:n-1}^{(p)}, K_{1:n-1}^{(p)}, I_{1:n-1}^{(p)}\}_{p=1}^{N_p}$. By discarding the samples of X_{n-1}, the empirical distribution of $p(\theta_{1:n-1}, K_{1:n-1}, I_{1:n-1}|Z_{1:n-1})$ is expressed as:

$$\begin{aligned} &p(\theta_{1:n-1}, K_{1:n-1}, I_{1:n-1}|Z_{1:n-1}) \\ &\approx \frac{1}{N_p}\sum_{p=1}^{N_p}\delta_{\{\theta,K,L\}_{1:n-1}^{(p)}}(\theta_{1:n-1}, K_{1:n-1}, I_{1:n-1}) \end{aligned} \tag{24}$$

Thus, the Bayesian recursion can be expressed as follows, where the joint posterior $p(X_n, \theta_{1:n}, K_{1:n}, I_{1:n}|Z_{1:n})$ is computed sequentially based on $p(\theta_{1:n-1}, K_{1:n-1}, I_{1:n-1}|Z_{1:n-1})$ at the previous time step $n-1$:

$$\begin{aligned} &p(X_n, \theta_{1:n}, K_{1:n}, I_{1:n}|Z_{1:n}) \\ &= \frac{p(Z_n|\theta_n, X_n, M_n)}{p(Z_n|Z_{1:n-1}, M_n)}p(\theta_n|K_n, M_n)p(K_n|K_{n-1}) \\ &\quad \times p(I_n|K_{n-1}, K_n)p(X_n|\theta_{1:n-1}, K_{1:n}, I_{1:n}, Z_{1:n-1}) \\ &\quad \times p(\theta_{1:n-1}, K_{1:n-1}, I_{1:n-1}|Z_{1:n-1}) \end{aligned} \tag{25}$$

where $p(\theta_{1:n-1}, K_{1:n-1}, I_{1:n-1}|Z_{1:n-1})$ is the empirical distribution in (24), $p(Z_n|\theta_n, X_n, M_n)$ is the measurement likelihood function in (11), and $p(\theta_n|K_n, M_n)$ denotes the association prior with an expression in (13). $p(K_n|K_{n-1})$ and $p(I_n|K_{n-1}, K_n)$ are given in (4) and (5). The predicted density $p(X_n|\theta_{1:n-1}, K_{1:n}, I_{1:n}, Z_{1:n-1})$ is in the form of a Gaussian distribution as in (19).

In this paper, we adopt an independent Metropolis-Hastings formulation of the RJ-SMCMC method, under which the proposal function is independent of the previous iteration's samples. In this part, we present an adaptive proposal that can realise a partially parallel sampling structure and may enjoy a better mixing property than using predictive prior proposals. In accordance to the principle of choosing the proposal density that resembles the target distribution, we use the optimal proposal $p(\theta_{n,j}|Z_{n,j}, X_n)$ for each of the association variables $\theta_{n,j}$ and predictive prior proposals for the target number K_n and the death identifier I_n. Both the predictive prior and measurements are considered when proposing the target state X_n. Accordingly, the proposal density $q(X_n, \theta_{1:n}, K_{1:n}, I_{1:n}|Z_{1:n})$ is designed in the following form:

$$\begin{aligned} q(X_n, \theta_{1:n}, K_{1:n}, I_{1:n}|Z_{1:n}) &= \prod_{j=1}^{M_n} p(\theta_{n,j}|Z_{n,j}, X_n) \\ &\times p(K_n|K_{n-1})p(I_n|K_{n-1}, K_n) \\ &\times q_a(X_n|\theta_{1:n-1}, K_{1:n}, I_{1:n}, Z_{1:n}) \\ &\times p(\theta_{1:n-1}, K_{1:n-1}, I_{1:n-1}|Z_{1:n-1}) \end{aligned} \tag{26}$$

where we can see that all the associations $\theta_{n,j}, j = 1, ..., M_n$ can be sampled directly from (15) in parallel. The proposal $q_a(X_n|\theta_{1:n-1}, K_{1:n}, I_{1:n}, Z_{1:n})$ for X_n considers both the form of the target distribution and the newly-received measurements Z_n at each time step n, with an expression of:

$$\begin{aligned} &q_a(X_n|\theta_{1:n-1}, K_{1:n}, I_{1:n}, Z_{1:n}) \\ &= \begin{cases} q_1(X_{n,K_n}|Z_n) & \\ \times \prod_{i=1}^{K_{n-1}} \mathcal{N}(X_{n,i}; \mu_{n|n-1,i}, \Sigma_{n|n-1,i}) & K_n = K_{n-1}+1; \\ p(X_n|K_{1:n}, I_{1:n}, \theta_{1:n-1}, Z_{1:n-1}) & K_n \neq K_{n-1}+1; \end{cases} \end{aligned} \tag{27}$$

where each $\mathcal{N}(\mu_{n|n-1,i}, \Sigma_{n|n-1,i})$ has the same form in (19) without the index $(\cdot)^{(p)}$. Specifically, instead of the Gaussian birth prior $p_1(X_{n,K_n}) = \mathcal{N}(\mu_b, C_b)$ defined over the surveillance/birth region, here we design a Gaussian mixture proposal where we first conduct a measurement clustering step to find a group of N_b possible birth locations $\{m_r\}_{r=1}^{N_b}$, and thus $q_1(X_{n,K_n}|Z_n) = \frac{1}{N_b}\sum_{r=1}^{N_b}\mathcal{N}(m_r, C_r)$. Other than this target birth proposal, all the remaining proposals are set equal to the target prior transition densities $p(X_n|K_{1:n}, I_{1:n}, \theta_{1:n-1}, Z_{1:n-1})$.

Assume we have samples $\theta_{1:n}^{m-1}, X_n^{m-1}, K_{1:n}^{m-1}, I_{1:n}^{m-1}$ from the last iteration $m-1$, the procedure of the independent Metropolis-Hastings at iteration m is to first draw new samples $\theta_{1:n}^{'}, X_n^{'}, K_{1:n}^{'}, I_{1:n}^{'}$ from the proposal function $q(X_n, \theta_{1:n}, K_{1:n}, I_{1:n}|Z_{1:n})$, and then to accept or reject these new samples based on the acceptance probability, written as:

$$\alpha(X_n^{m-1}, \{\theta, K, I\}_{1:n}^{m-1}; X_n^{'}, \{\theta, K, I\}_{1:n}^{'}) = \min(1, \rho) \tag{28}$$

where ρ is the acceptance ratio, computed as:

$$\begin{aligned} \rho = &\frac{p(X_n^{'}, \theta_{1:n}^{'}, K_{1:n}^{'}, I_{1:n}^{'}|Z_{1:n})}{q(X_n^{'}, \theta_{1:n}^{'}, K_{1:n}^{'}, I_{1:n}^{'}|Z_{1:n})} \\ &\times \frac{q(X_n^{m-1}, \theta_{1:n}^{m-1}, K_{1:n}^{m-1}, I_{1:n}^{m-1}|Z_{1:n})}{p(X_n^{m-1}, \theta_{1:n}^{m-1}, K_{1:n}^{m-1}, I_{1:n}^{m-1}|Z_{1:n})}. \end{aligned} \tag{29}$$

Note that in (29), the Jacobian term for meeting dimension matching condition is omitted as it is equals to one [17].

By using (10)-(11), (13) and the Bayes' formula we have the following relationship:

$$\begin{aligned} \prod_{j=1}^{M_n} p(\theta_{n,j}|Z_{n,j}, X_n) &= \frac{\prod_{j=1}^{M_n} p(\theta_{n,j}, Z_{n,j}|X_n)}{\prod_{j=1}^{M_n} p(Z_{n,j}|X_n)} \\ &= \frac{p(Z_n|\theta_n, X_n, M_n)}{p(Z_n|X_n, M_n)}p(\theta_n|K_n, M_n) \end{aligned} \tag{30}$$

Thus, the acceptance ratio ρ in (29) can be computed by using (25), (26) and (30):

$$\begin{aligned} \rho = &\frac{p(Z_n|X_n^{'}, M_n)}{p(Z_n|X_n^{(m-1)}, M_n)} \\ &\times \frac{p(X_n^{'}|\theta_{1:n-1}^{'}, K_{1:n}^{'}, I_{1:n}^{'}, Z_{1:n-1})}{p(X_n^{(m-1)}|\theta_{1:n-1}^{(m-1)}, K_{1:n}^{(m-1)}, I_{1:n}^{(m-1)}, Z_{1:n-1})} \\ &\times \frac{q_a(X_n^{(m-1)}|\theta_{1:n-1}^{(m-1)}, K_{1:n}^{(m-1)}, I_{1:n}^{(m-1)}, Z_{1:n})}{q_a(X_n^{'}|\theta_{1:n-1}^{'}, K_{1:n}^{'}, I_{1:n}^{'}, Z_{1:n})} \\ = &\rho_0 \times \rho_1 \end{aligned} \tag{31}$$

where the first part ρ_0 is deduced as

$$\begin{aligned}\rho_0 &= \frac{p(Z_n|X_n^{'}, M_n)}{p(Z_n|X_n^{(m-1)}, M_n)} \\ &= \frac{\sum_{\theta_n} p(Z_n, \theta_n|X_n^{'}, K_n^{'}, M_n)}{\sum_{\theta_n} p(Z_n, \theta_n|X_n^{(m-1)}, K_n^{(m-1)}, M_n)} \\ &= \frac{\sum_{\theta_{n,1}} \cdots \sum_{\theta_{n,M_n}} \prod_{j=1}^{M_n} p(Z_{n,j}, \theta_{n,j}|X_n^{'})}{\sum_{\theta_{n,1}} \cdots \sum_{\theta_{n,M_n}} \prod_{j=1}^{M_n} p(Z_{n,j}, \theta_{n,j}|X_n^{(m-1)})} \\ &= \prod_{j=1}^{M_n} \frac{\sum_{i'=0}^{K_n^{'}} p(Z_{n,j}, \theta_{n,j} = i'|X_{n,i'}^{'})}{\sum_{i=0}^{K_n^{(m-1)}} p(Z_{n,j}, \theta_{n,j} = i|X_{n,i}^{(m-1)})} \end{aligned} \tag{32}$$

Subsequently, the acceptance ratio ρ is computed as:

$$\rho = \rho_0 \times \rho_1 = \begin{cases} \rho_0 \times \rho_1^1 & \text{Case 1;} \\ \rho_0 \times \rho_1^2 & \text{Case 2;} \\ \rho_0 \times \rho_1^3 & \text{Case 3;} \\ \rho_0 \times \rho_1^4 & \text{Case 4.} \end{cases} \tag{33}$$

The specific forms of ρ_1 for different events are:

- Case 1: $K_n^{'} = K_{n-1}^{'} + 1$ and $K_n^{(m-1)} = K_{n-1}^{(m-1)} + 1$;

$$\begin{aligned}\rho_1^1 &= \frac{p_1(X_{n,K_n}^{'})}{p_1(X_{n,K_n}^{(m-1)})} \frac{q_1(X_{n,K_n}^{(m-1)}|Z_n)}{q_1(X_{n,K_n}^{'}|Z_n)} \\ &= \frac{\mathcal{N}(X_{n,K_n}^{'}; \mu_b, P_b)}{\mathcal{N}(X_{n,K_n}^{(m-1)}; \mu_b, P_b)} \frac{\sum_{r=1}^{N_b} \mathcal{N}(X_{n,K_n}^{(m-1)}; m_r, C_r)}{\sum_{r=1}^{N_b} \mathcal{N}(X_{n,K_n}^{'}; m_r, C_r)} \end{aligned} \tag{34}$$

- Case 2: $K_n^{'} = K_{n-1}^{'} + 1$ and $K_n^{(m-1)} \neq K_{n-1}^{(m-1)} + 1$;

$$\rho_1^2 = \frac{p_1(X_{n,K_n}^{'})}{q_1(X_{n,K_n}^{'}|Z_n)} = N_b \times \frac{\mathcal{N}(X_{n,K_n}^{'}; \mu_b, P_b)}{\sum_{r=1}^{N_b} \mathcal{N}(X_{n,K_n}^{'}; m_r, C_r)} \tag{35}$$

- Case 3: $K_n^{'} \neq K_{n-1}^{'} + 1$ and $K_n^{(m-1)} = K_{n-1}^{(m-1)} + 1$;

$$\rho_1^3 = \frac{q_1(X_{n,K_n}^{(m-1)}|Z_n)}{p_1(X_{n,K_n}^{(m-1)})} = \frac{\sum_{r=1}^{N_b} \mathcal{N}(X_{n,K_n}^{(m-1)}; m_r, C_r)}{N_b \times \mathcal{N}(X_{n,K_n}^{(m-1)}; \mu_b, P_b)} \tag{36}$$

- Case 4: $K_n^{'} \neq K_{n-1}^{'} + 1$ and $K_n^{(m-1)} \neq K_{n-1}^{(m-1)} + 1$;

$$\rho_1^4 = 1 \tag{37}$$

Therefore, we can observe that the implementation is greatly simplified as most of the terms are cancelled out in the calculation of the acceptance ratio.

In sum, the procedure of the independent Metropolis-Hastings at iteration m with the adaptive proposal includes the following steps, in which it first jointly draws $\theta_{1:n}^{'}, X_n^{'}, K_{1:n}^{'}, I_{1:n}^{'}$, and accept it with a probability of $\alpha(X_n^{m-1}, \{\theta, K, I\}_{1:n}^{m-1}; X_n^{'}, \{\theta, K, I\}_{1:n}^{'})$.

1) draw $\{\theta_{1:n-1}^{'}, K_{1:n-1}^{'}, I_{1:n-1}^{'}\}$ from (24);
2) draw $K_n^{'}$ and $I_n^{'}$ from (4) and (5);
3) draw $\{X_{n,i}^{'}\}_{i=1}^{K_n}$ from (27) in parallel;
4) draw $\{\theta_{j,n}^{'}\}_{j=1}^{M_n}$ from (15) in parallel;
5) accept $\theta_{1:n}^{'}, X_n^{'}, K_{1:n}^{'}, I_{1:n}^{'}$ with a probability of $\alpha(X_n^{m-1}, \{\theta, K, I\}_{1:n}^{m-1}; X_n^{'}, \{\theta, K, I\}_{1:n}^{'})$

IV. Results

This section analyses the performance of the proposed RB-RJ-SMCMC sampler for tracking a time-varying number of targets in clutter, and compares it with existing trackers including the PMBM filter in [6] and the SPA-based tracker in [12]. Here we assume the target shapes and measurement rates are known and all the compared methods are adjusted to this fixed parameter setting. We simulated two datasets to compare the tracking performance of the methods. The first dataset is a simpler scenario with a lower clutter rate and fewer target births and deaths, and the second dataset is more challenging with heavy clutter and frequent target appearing and disappearing events.

We use the optimal sub pattern assignment (OSPA) [22] and general OSPA (GOSPA) [23] metrics to evaluate the tracking performance of all methods. For both the OSPA and GOSPA metrics, the order is set to $p = 1$ and the distance cut-off value is $c = 20$. For the first dataset, we plot its mean OSPA value over time since the track is fixed for all simulations. For both datasets, we calculate the mean OSPA and the mean GOSPA metrics over Monte Carlo runs. To compare computational efficiency, we measure the CPU time required at each time step and average it across all time steps and simulations (System: Intel(R) Core(TM) i7-8550 CPU at 1.80 GHz, 8 GB RAM).

The general parameter settings are as follows. For all datasets, the total time steps are 100, and the time interval between observations is $\tau = 1$s. In the RB-RJ-SMCMC sampler, we use the DBSCAN algorithm to construct the adaptive proposal with the maximum distance being 20, and the minimum cluster member number being 2, and the birth Gaussian prior is set to cover the whole surveillance region. For the PMBM filter, the detection probability is 0.99. The ellipsoidal gate size in probability is 0.999. The global hypotheses and Bernoulli component pruning thresholds are 10^{-2} and 10^{-3}, respectively. The maximum global hypothesis number is 100. To speed up the PMBM filter, the birth area is set to the ground-truth birth region in the simulations. The DBSCAN algorithm in the PMBM filter is run with 50 different distance values equally spaced between 1 and 50, and the maximum assignment number is 20 for Murty's algorithm. For the SPA tracker, the object declaration threshold is 0.5, and the pruning threshold is 10^{-3}. The target birth zone is set to the whole region, and a distance-based clustering is used to select a number of the most probable measurements to initiate the target birth as in [12], instead of using all measurements.

1) Dataset 1 with a lower clutter rate and fewer target births and deaths: In the first simulation, we have a maximum of four targets and the changes in the target number over time are illustrated in Figure 3. The target Poisson rates are set to 5, and the clutter rate is 50. The parameters in the CV model are $Q_i = 8$ and $R_i = 100$I where I is a 2-D identity matrix. To evaluate the robustness of the algorithm, we generate

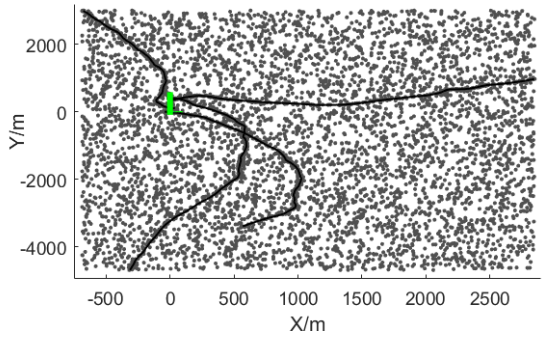

(a) Example measurement set from dataset 1

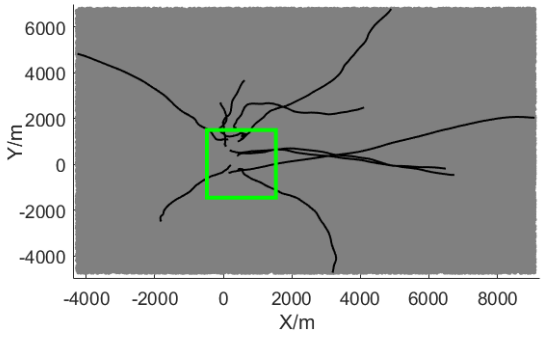

(b) Example measurement set from dataset 2

Fig. 1: Example measurements for each datasets over 100 time steps from 20 Monte Carlo runs; grey dots are measurements; black lines are ground truth tracks; green rectangle denotes the target birth zone.

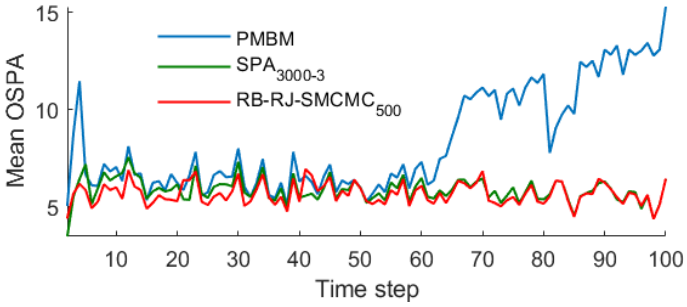

Fig. 2: Mean OSPA over 100 time steps.

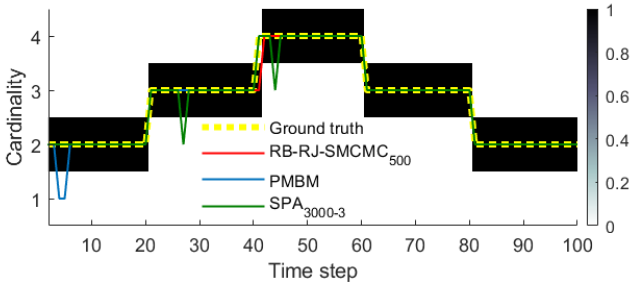

Fig. 3: Ground truth and estimated target number of one example simulation in dataset 1; the colorbar ranging from white to black indicates the probability of the cardinality value at each time step.

20 different measurement sets from the same ground-truth trajectories in Figure 1a, under the same parameter settings. One sample measurement set is shown in Figure 1a. For the RB-RJ-SMCMC sampler, we examine two sample sizes: 200 particles and 500 particles with a burn-in time of 50 and 100 iterations, respectively. For the SPA tracker, 3000 particles with 3 iterations are used in this simulation.

The mean OSPA and mean GOSPA averaged over 20 datasets are shown in Table I. It shows that the proposed RB-RJ-SMCMC samplers have the lowest mean OSPA and mean GOSPA values, and the tracking accuracy increases with the sample size. In addition, the computational time of RB-RJ-SMCMC$_{200}$ is also comparable to the fastest method. To present the tracking performance over time, we plot the mean OSPA of these methods in Figure 2 over 100 time steps. It is observed that the proposed method has the smallest mean OSPA most of the time. In Figure 3, we provide cardinality estimates of all methods for one example simulation, including both the cardinality mean point estimates and colormap plot of the RB-RJ-SMCMC sampler. It verifies that our method outperforms others with fewer errors in cardinality estimation.

2) Dataset 2 with heavy clutter and frequent target births and deaths: We set $Q_i = 9$ and $R_i = 100\mathrm{I}$. The target rates are 10; the clutter rate is 1500. Figure 4 shows the ground truth target number, where targets appear and disappear frequently over the 100 time steps. To evaluate the robustness of the

TABLE I: Tracking performance of dataset 1

method	mean OSPA	mean GOSPA	CPU time (s)
PMBM	8.19	21.33	0.05
SPA$_{3000-3}$	5.77	15.93	0.32
RB-RJ-SMCMC$_{200}$	5.65	15.71	0.08
RB-RJ-SMCMC$_{500}$	5.58	15.56	0.19

algorithm, we generate 20 simulations with different tracks and measurements under the same settings, where one sample measurement set and trajectories are shown in Figure 1b. Here we evaluate the proposed sampler using 500 particles and 1500 particles with a burn-in time of 100 and 200 iterations, respectively. For the SPA tracker, we use up to 200 particles with 3 iterations due to the large computational time it requires as shown in Table II.

Table II gives the mean OSPA and mean GOSPA averaged over 20 datasets. To show the effectiveness of the proposed method more clearly, we analyse one randomly-chosen example from 20 simulations, and the target number estimates are shown in Figure 4. In this heavy clutter case, our trackers are much faster than all other methods with higher accuracy, demonstrating the scalability of the proposed methods with the size of the measurement data.

TABLE II: Tracking performance of dataset 2

method	mean OSPA	mean GOSPA	CPU time (s)
PMBM	10.22	70.39	4.19
SPA_{50-3}	7.13	51.27	25.35
SPA_{200-3}	5.03	36.17	132.3
RB-RJ-SMCMC_{500}	5.01	33.72	0.39
$\text{RB-RJ-SMCMC}_{1500}$	4.84	33.05	1.22

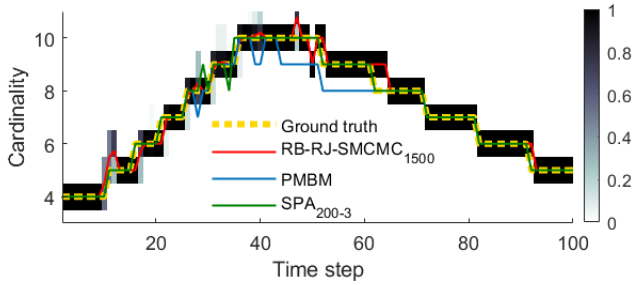

Fig. 4: Ground truth and estimated target number of one example simulation in dataset 2; the colorbar ranging from white to black indicates the probability of the cardinality value at each time step.

V. Conclusion

In this paper, we devise a scalable Rao-Blackwellised reversible jump SMCMC sampler for tracking an unknown number of objects in clutter. Specifically, we develop an independent Metropolis-Hastings formulation with an adaptive proposal function that is independent of the previous iteration's samples. The deduced sampler enjoys a partially parallel sampling structure and the merits of Rao-Blackwellisation, thus being advantageous in both tracking accuracy and efficiency. Future work includes the design of more sophisticated sampling schemes and proposal functions to enhance sampling efficiency. Other modelling assumptions, e.g., [24]–[26] and birth/death priors will also be investigated for a wider application of the current tracker.

Acknowledgements: 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.

References

[1] Q. Li, R. Gan, J. Liang, and S. J. Godsill, "An adaptive and scalable multi-object tracker based on the non-homogeneous Poisson process," *IEEE Transactions on Signal Processing*, 2023.

[2] 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.

[3] Y. Bar-Shalom and X. R. Li, *Multitarget-multisensor tracking: principles and techniques*. YBs Storrs, CT, 1995, vol. 19.

[4] 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.

[5] R. P. Mahler, "Multitarget Bayes filtering via first-order multitarget moments," *IEEE Transactions on Aerospace and Electronic systems*, vol. 39, no. 4, pp. 1152–1178, 2003.

[6] 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.

[7] B.-N. Vo and W.-K. Ma, "The Gaussian mixture probability hypothesis density filter," *IEEE Transactions on signal processing*, vol. 54, no. 11, pp. 4091–4104, 2006.

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

[9] J. Vermaak, S. Maskell, and M. Briers, "A unifying framework for multi-target tracking and existence," in *2005 7th International Conference on Information Fusion*, vol. 1. IEEE, 2005, pp. 9–pp.

[10] F. Septier, S. K. Pang, A. Carmi, and S. Godsill, "On MCMC-based particle methods for Bayesian filtering: Application to multitarget tracking," in *3rd ieee international workshop on computational advances in multi-sensor adaptive processing (camsap)*. Citeseer, 2009, pp. 360–363.

[11] 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.

[12] 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.

[13] 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.

[14] A. Doucet, B.-N. Vo, C. Andrieu, and M. Davy, "Particle filtering for multi-target tracking and sensor management," in *Proceedings of the Fifth International Conference on Information Fusion. FUSION 2002.(IEEE Cat. No. 02EX5997)*, vol. 1. IEEE, 2002, pp. 474–481.

[15] 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.

[16] P. J. Green, "Reversible jump Markov chain Monte Carlo computation and Bayesian model determination," *Biometrika*, vol. 82, no. 4, pp. 711–732, 1995.

[17] S. J. Godsill, "On the relationship between Markov chain Monte Carlo methods for model uncertainty," *Journal of computational and graphical statistics*, vol. 10, no. 2, pp. 230–248, 2001.

[18] Z. Khan, T. Balch, and F. Dellaert, "An MCMC-based particle filter for tracking multiple interacting targets," in *European Conference on Computer Vision*. Springer, 2004, pp. 279–290.

[19] 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.

[20] 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.

[21] W. Ng, J. Li, S. Godsill, and J. Vermaak, "A hybrid approach for online joint detection and tracking for multiple targets," in *2005 IEEE Aerospace Conference*. IEEE, 2005, pp. 2126–2141.

[22] 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.

[23] 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.

[24] 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.

[25] 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.

[26] Q. Li and S. Godsill, "A new leader-follower model for Bayesian tracking," in *2020 IEEE 23rd International Conference on Information Fusion (FUSION)*. IEEE, 2020, pp. 1–7.