Importance Sampling With Stochastic Particle Flow and Diffusion Optimization

Wenyu Zhang, Mohammad J. Khojasteh, Nikolay A. Atanasov, and Florian Meyer

Abstract—Particle flow (PFl) is an effective method for overcoming particle degeneracy, the main limitation of particle filtering. In PFl, particles are migrated towards regions of high likelihood based on the solution of a partial differential equation. Recently proposed stochastic PFl introduces a diffusion term in the differential equation that describes the trajectory of particles. This diffusion term reduces the stiffness of the differential equation and makes it possible to perform PFl with a lower number of numerical integration steps compared to traditional deterministic PFl. In this work, we introduce a general approach to perform importance sampling (IS) based on stochastic PFl. Our method makes it possible to evaluate a "flow-induced" proposal probability density function (PDF) after the parameters of the prior/predicted PDF represented by Gaussian mixture model (GMM) have been migrated by stochastic PFl. Compared to conventional stochastic PFl, the resulting processing step is asymptotically optimal. Within our method, it is possible to optimize the diffusion matrix that characterizes the diffusion term of the differential equation to improve the accuracy-computational complexity tradeoff. Our simulation results in a highly nonlinear 3-D source localization scenario showcase a reduced stiffness of the resulting stochastic PFl and an improved estimating accuracy compared to state-of-the-art deterministic and stochastic PFl.

I. INTRODUCTION

The particle filter is probably the most widely used method for nonlinear sequential Bayesian estimation [1], [2]. In a strategy known as IS, particles are first sampled from an arbitrary proposal PDF and then weighted based on the predicted/prior PDF and likelihood function. IS is asymptotically optimal, assuming weak conditions are satisfied [3], [4], but known to suffer from particle degeneracy in higher dimensional problems [5]. Here, due to the curse of dimensionality, too few particles have a significant weight after the update step.

To overcome particle degeneracy, PFl [6], [7] migrates particles sampled from a predicted/prior PDF to regions of high likelihood. This transition from the predicted/prior PDF to the posterior PDF is described by a homotopy function. By using the homotopy function to constrain the Fokker-Planck equation and corresponding Langevin stochastic differential equation (SDE), one can derive a drift vector and a diffusion matrix that describe the particle trajectories of the resulting stochastic PFl. By setting the diffusion matrix equal to zero, a deterministic PFl is obtained. In particular, the exact Daum and Huang (EDH) flow [8] and the Gromov's flow [9] are popular deterministic and stochastic PFls with closed-form solutions. Stochastic PFls tend to require a lower number of numerical

integration steps due to improved transient dynamics, i.e., a reduced stiffness of the underlying Langevin dynamics, and thus lead to a reduced overall computational complexity. As proposed and demonstrated recently [10], in stochastic PFl it is possible to optimize the diffusion term to further reduce stiffness and thus computational complexity. Although it has been demonstrated that PFl can overcome particle degeneracy in a variety of nonlinear and high-dimensional problems [6]–[9], it has no asymptotic optimality guarantees if directly applied to the particles [11].

For asymptotically optimal estimation, PFl can be used to develop a proposal PDF for IS. In particular, the particle flow particle filter (PFPF) [11] embeds an EDH PFl into IS by introducing a flow-induced proposal PDF. The evaluation of this proposal PDF at the migrated particles requires an invertible mapping from the predicted/prior PDF at particles before the flow and the PDF represented by the particles after the flow. Invertible deterministic PFl was recently combined with a GMM and used within a belief propagation (BP) framework for the 3-D tracking of an unknown number of sources in the presence of data association uncertainty [12], [13]. This approach has recently been used in the context of marine mammal research [14]. However, this invertible mapping is limited to deterministic flows like the EDH [15]. For stochastic PFl, in [15] particles are drawn at each flow step from a different proposal PDF, which leads to a computationally expensive weights computation for IS. The work in [16], [17] utilizes auxiliary variables and their filtering for IS with embedded PFl. In particular, [17] extended this framework to the stochastic Gromov's flow. Although the method in [17] is computationally less expensive than [15], it has twice computational complexity of PFl due to an additional auxiliary variable for each particle. Furthermore, it relies on heuristic and suboptimal diffusion matrix selections without optimization. Note that IS based on PFl requires explicit forms and derivatives of the measurement model and the predicted/prior PDF.

In this work, we introduce an approach that combines IS with general stochastic flows that include optimizable diffusion terms. The resulting IS framework is combined with a GMM and used within a BP framework for the detection and localization of an unknown number of sources in 3-D. The main contributions of our work are summarized as follows.

- We develop efficient IS where the parameters of a GMM are migrated by stochastic PFls that rely on a SDE with an optimizable diffusion term.
- We evaluate our method in a challenging 3-D multi-source localization problem and demonstrate significant improvements compared to state-of-the-art methods.

This paper advances over the preliminary account of our method provided in the conference publication [18] by (i) extending the approach to general stochastic flows that include optimizable diffusion terms and (ii) introducing additional and more extensive numerical results.

II. STOCHASTIC PFL

Consider a random state to be estimated, $\boldsymbol{x} \in \mathbb{R}^N$, in a Bayesian setting. PFl [6], [7] establishes a continuous mapping w.r.t. pseudo-time $\lambda \in [0, 1]$, to migrate particles sampled from the prior PDF $\boldsymbol{x}_0 \sim f(\boldsymbol{x})$ such that they represent the posterior PDF $\boldsymbol{x}_1 \sim f(\boldsymbol{x}|\boldsymbol{z})$.

Let $f(\boldsymbol{x})$ be the prior PDF and $h(\boldsymbol{x}) = f(\boldsymbol{z}|\boldsymbol{x})$ be the likelihood function, where $\boldsymbol{z}$ is observed and thus fixed. Following Bayes' rule, a log-homotopy function [6], [7] can be introduced as $\phi(\boldsymbol{x}, \lambda) = \log f(\boldsymbol{x}) + \lambda \log h(\boldsymbol{x})$. This log-homotopy function is used to constrain the Fokker-Plank equation and the corresponding Langevin SDE of the form

$$\mathrm{d}\boldsymbol{x} = \boldsymbol{\zeta}(\boldsymbol{x}, \lambda)\mathrm{d}\lambda + \boldsymbol{Q}(\lambda)^{1/2}\mathrm{d}\boldsymbol{w}. \tag{1}$$

Here, $\boldsymbol{\zeta}(\boldsymbol{x}, \lambda) \in \mathbb{R}^N$ is the drift vector, $\boldsymbol{Q}(\lambda) \in \mathbb{R}^{N \times N}$ is the diffusion matrix, and $\boldsymbol{w} \in \mathbb{R}^N$ is Brownian motion. The constrained Fokker-Plank equation can be used to derive the functional forms of $\boldsymbol{\zeta}(\boldsymbol{x}, \lambda)$ and $\boldsymbol{Q}(\lambda)$ [9]. (Note that by setting the diffusion matrix equal to zero, i.e., $\mathrm{d}\boldsymbol{x} = \boldsymbol{\zeta}(\boldsymbol{x}, \lambda)\mathrm{d}\lambda$, a deterministic flow is obtained.)

Since the drift is both time and state-dependent, directly integrating λ from 0 to 1 is analytically infeasible. Thus, the Euler-Maruyama method [19] is commonly used for numerical integration. Here, particle migration is performed by evaluating $\boldsymbol{\zeta}(\boldsymbol{x}, \lambda)$ at N_λ discrete values of λ, i.e., $0 = \lambda_0 < \lambda_1 < ... < \lambda_{N_\lambda} = 1$. First, N_s particles $\{\boldsymbol{x}_0^{(i)}\}_{i=1}^{N_\mathrm{s}} = \{\boldsymbol{x}_{\lambda_0}^{(i)}\}_{i=1}^{N_\mathrm{s}}$ are drawn from $f(\boldsymbol{x})$. Next, each particle $i \in \{1, \ldots, N_\mathrm{s}\}$ is migrated sequentially across the discrete pseudo time steps $l \in \{1, \ldots, N_\lambda\}$, i.e.,

$$\boldsymbol{x}_{\lambda_l}^{(i)} = \boldsymbol{x}_{\lambda_{l-1}}^{(i)} + \boldsymbol{\zeta}_\mathrm{s}(\boldsymbol{x}_{\lambda_{l-1}}^{(i)}, \lambda_l)\Delta_l + \sqrt{\Delta_l \boldsymbol{Q}(\lambda_l)}\,\boldsymbol{w}_l^{(i)} \tag{2}$$

where $\boldsymbol{w}_l^{(i)}$ is Gaussian distributed with unit variance and $\Delta_l = \lambda_l - \lambda_{l-1}$. In this way, particles $\{\boldsymbol{x}_1^{(i)}\}_{i=1}^{N_\mathrm{s}} = \{\boldsymbol{x}_{\lambda_{N_\lambda}}^{(i)}\}_{i=1}^{N_\mathrm{s}}$ representing the posterior PDF $f(\boldsymbol{x}|\boldsymbol{z}) \propto \exp\big(\phi(\boldsymbol{x}, \lambda = 1)\big)$ are finally obtained.

Recent results demonstrate that, based on an appropriate choice of the diffusion term $\boldsymbol{Q}(\lambda)$, stochastic PFl can provide a strongly reduced number of integration steps and computational complexity compared deterministic PFl [10]. A popular stochastic flow is Gromov's flow [9], given by the drift

$$\boldsymbol{\zeta}_\mathrm{g}(\boldsymbol{x}, \lambda) = -\big(\nabla_{\boldsymbol{x}}\nabla_{\boldsymbol{x}}^\mathrm{T}\phi\big)^{-1}\nabla_{\boldsymbol{x}}\log h \tag{3}$$

and the diffusion matrix

$$\boldsymbol{Q}_\mathrm{g}(\lambda) = -\big(\nabla_{\boldsymbol{x}}\nabla_{\boldsymbol{x}}^\mathrm{T}\phi\big)^{-1}\big(\nabla_{\boldsymbol{x}}\nabla_{\boldsymbol{x}}^\mathrm{T}\log h\big)\big(\nabla_{\boldsymbol{x}}\nabla_{\boldsymbol{x}}^\mathrm{T}\phi\big)^{-1}. \tag{4}$$

Here, we used the short notation $\phi \triangleq \phi(\boldsymbol{x}, \lambda)$ and $h \triangleq h(\boldsymbol{x})$.

For a linear measurement model $\boldsymbol{z} = \boldsymbol{H}\boldsymbol{x} + \boldsymbol{v}$ with zero-mean additive Gaussian noise $\boldsymbol{v}$ and covariance matrix $\boldsymbol{R}$, we have $\nabla_{\boldsymbol{x}}\log h = \boldsymbol{H}^\mathrm{T}\boldsymbol{R}^{-1}(\boldsymbol{z} - \boldsymbol{H}\boldsymbol{x})$, $\nabla_{\boldsymbol{x}}\nabla_{\boldsymbol{x}}^\mathrm{T}\log h = -\boldsymbol{H}^\mathrm{T}\boldsymbol{R}^{-1}\boldsymbol{H}$, and $\nabla_{\boldsymbol{x}}\nabla_{\boldsymbol{x}}^\mathrm{T}\phi = -\boldsymbol{P}^{-1} - \lambda\boldsymbol{H}^\mathrm{T}\boldsymbol{R}^{-1}\boldsymbol{H}$.

We can now rewrite drift (3) and diffusion (4) of Gromov's flow as

$$\boldsymbol{\zeta}_\mathrm{g}(\boldsymbol{x}, \lambda) = \big(\boldsymbol{P}^{-1} + \lambda\boldsymbol{H}^\mathrm{T}\boldsymbol{R}^{-1}\boldsymbol{H}\big)^{-1}\boldsymbol{H}^\mathrm{T}\boldsymbol{R}^{-1}(\boldsymbol{z} - \boldsymbol{H}\boldsymbol{x})$$

$$\boldsymbol{Q}_\mathrm{g}(\lambda) = (\boldsymbol{P}^{-1} + \lambda\boldsymbol{H}^\mathrm{T}\boldsymbol{R}^{-1}\boldsymbol{H})^{-1}(\boldsymbol{H}^\mathrm{T}\boldsymbol{R}^{-1}\boldsymbol{H})$$
$$\times (\boldsymbol{P}^{-1} + \lambda\boldsymbol{H}^\mathrm{T}\boldsymbol{R}^{-1}\boldsymbol{H})^{-1}. \tag{5}$$

For a nonlinear measurement model $\boldsymbol{z} = \boldsymbol{h}(\boldsymbol{x}) + \boldsymbol{v}$, one can use the solution of Gromov's flow in (5) based on a local linearization of the measurement function $\boldsymbol{h}(\cdot)$. Gromov's flow applied to linearized measurement models has been demonstrated to outperform deterministic flows [17], [18], [20]. However, as other deterministic and stochastic PFls, due to approximations made for numerical integration, Gromov's flow is not asymptotically optimal, i.e., the PDF represented by the particles after the flow is only an approximation of $f(\boldsymbol{x}|\boldsymbol{z})$.

An alternative approach for asymptotically optimal estimation is to use PFl methods for IS within a particle filtering framework. Due to the lack of an invertible mapping in stochastic PFl, evaluating the PDF after the flow, as required for proposal evaluation [11], is typically infeasible.

III. IS WITH STOCHASTIC PFL

In this work, we propose to use PFl based on a linearized measurement model to develop a "flow-induced" GMM as proposal PDF. For example, let us first transform the drift of Gromov's flow to an affine function, i.e., $\boldsymbol{\zeta}_\mathrm{g}(\boldsymbol{x}, \lambda) = \boldsymbol{A}_\mathrm{g}(\lambda)\boldsymbol{x} + \boldsymbol{b}_\mathrm{g}(\lambda)$, with

$$\boldsymbol{A}_\mathrm{g}(\lambda) = -\big(\boldsymbol{P}^{-1} + \lambda\boldsymbol{H}^\mathrm{T}\boldsymbol{R}^{-1}\boldsymbol{H}\big)^{-1}\boldsymbol{H}^\mathrm{T}\boldsymbol{R}^{-1}\boldsymbol{H}$$

$$\boldsymbol{b}_\mathrm{g}(\lambda) = \big(\boldsymbol{P}^{-1} + \lambda\boldsymbol{H}^\mathrm{T}\boldsymbol{R}^{-1}\boldsymbol{H}\big)^{-1}\boldsymbol{H}^\mathrm{T}\boldsymbol{R}^{-1}\boldsymbol{z}\,.$$

Next, consider a single Gaussian $f(\boldsymbol{x}) = \mathcal{N}(\boldsymbol{x}; \boldsymbol{\mu}_0, \boldsymbol{P}_0)$ with mean $\boldsymbol{\mu}_0$ and covariance matrix $\boldsymbol{P}_0$ as predicted/prior PDF. Based on the affine form introduced above, we can migrate the mean and covariance of this Gaussian predicted/prior PDF based on PFl, i.e.,

$$\boldsymbol{\mu}_l = \boldsymbol{\mu}_{l-1} + \boldsymbol{\zeta}_\mathrm{s}(\boldsymbol{\mu}_{\lambda_{l-1}}, \lambda_l)\Delta_l \tag{6}$$

$$\boldsymbol{P}_l = [\boldsymbol{I} + \Delta_l\boldsymbol{A}_\mathrm{s}(\lambda_l)]\boldsymbol{P}_{l-1}[\boldsymbol{I} + \Delta_l\boldsymbol{A}_\mathrm{s}(\lambda_l)]^\mathrm{T} + \Delta_l\boldsymbol{Q}(\lambda_l) \tag{7}$$

for $l = 1, \ldots, N_\lambda$ and by setting $\boldsymbol{\zeta}_\mathrm{s}(\boldsymbol{\mu}_{\lambda_{l-1}}, \lambda_l) = \boldsymbol{\zeta}_\mathrm{g}(\boldsymbol{\mu}_{\lambda_{l-1}}, \lambda_l)$ and $\boldsymbol{A}_\mathrm{s} = \boldsymbol{A}_\mathrm{g}$. This principle can be extended to GMMs in a straightforward way. Assuming a prior/predicted PDF represented by a GMM, i.e., $f(\boldsymbol{x}) = \sum_{n=1}^{N_\mathrm{g}} w^{(n)}\mathcal{N}\big(\boldsymbol{x}; \boldsymbol{\mu}_0^{(n)}, \boldsymbol{P}_0^{(n)}\big)$, one can use (6) and (7) to perform N_g PFls, one for each GMM component. The resulting GMM $q(\boldsymbol{x}) = \sum_{n=1}^{N_\mathrm{g}} w^{(n)}\mathcal{N}\big(\boldsymbol{x}; \boldsymbol{\mu}_1^{(n)}, \boldsymbol{P}_1^{(n)}\big)$ can be used as a "flow-induced" proposal PDF for asymptotically optimal IS.

For an accurate numerical implementation of the PFl, the step sizes, Δ_l, need to be adapted to the stiffness of the flow. A flow with reduced stiffness can be implemented with larger step sizes, i.e., fewer steps, and thus yields reduced computational complexity. Next, we investigate how to develop stochastic PFls with reduced stiffness. Recently, it has been shown that a solution to (1) can be obtained by choosing $\boldsymbol{Q}(\lambda)$ arbitrarily and computing $\boldsymbol{\zeta}_\mathrm{s}(\boldsymbol{x}, \lambda)$ according to [10]

$$\boldsymbol{\zeta}_{\rm s}(\boldsymbol{x},\lambda) = \boldsymbol{\zeta}_{\rm d}(\boldsymbol{x},\lambda) + \frac{1}{2}\boldsymbol{Q}(\lambda)\,\nabla_{\boldsymbol{x}}\phi(\boldsymbol{x},\lambda) \tag{8}$$

where $\boldsymbol{\zeta}_{\rm d}(\boldsymbol{x},\lambda)$ is the drift of a deterministic flow, i.e., $\mathrm{d}\boldsymbol{x} = \boldsymbol{\zeta}(\boldsymbol{x},\lambda)\mathrm{d}\lambda$. Based on this result, one can choose a deterministic PFl and then design a diffusion matrix to get a stochastic PFl with reduced stiffness.

Consider Gromov's flow with stochastic drift and diffusion as in (3) and (4). Based on (8), the corresponding deterministic drift can be obtained as

$$\boldsymbol{\zeta}_{\rm d\text{-}g}(\boldsymbol{x},\lambda) = \boldsymbol{\zeta}_{\rm g}(\boldsymbol{x},\lambda) - \frac{1}{2}\boldsymbol{Q}_{\rm g}(\lambda)\,\nabla_{\boldsymbol{x}}\phi(\boldsymbol{x},\lambda). \tag{9}$$

By using an arbitrary diffusion matrix $\boldsymbol{Q}(\lambda)$ and by substituting $\boldsymbol{\zeta}_{\rm d}(\boldsymbol{x},\lambda)$ in (8) by $\boldsymbol{\zeta}_{\rm d\text{-}g}(\boldsymbol{x},\lambda)$ in (9), a new stochastic drift can be developed as

$$\boldsymbol{\zeta}_{\rm s}(\boldsymbol{x},\lambda) = \boldsymbol{\zeta}_{\rm g}(\boldsymbol{x},\lambda) + \frac{1}{2}\left(\boldsymbol{Q}(\lambda) - \boldsymbol{Q}_{\rm g}(\lambda)\right)\nabla_{\boldsymbol{x}}\phi(\boldsymbol{x},\lambda). \tag{10}$$

By considering a linearized measurement model, which results in $\nabla_{\boldsymbol{x}}\phi(\boldsymbol{x},\lambda) = \boldsymbol{P}^{-1}(\boldsymbol{\mu}_0-\boldsymbol{x}) + \lambda\boldsymbol{H}^{\rm T}\boldsymbol{R}^{-1}(\boldsymbol{z}-\boldsymbol{H}\boldsymbol{x})$, we can transform (10) to an affine function, i.e.,

$$\boldsymbol{A}_{\rm s}(\lambda) = -\frac{1}{2}\left(\boldsymbol{P}^{-1} + \lambda\boldsymbol{H}^{\rm T}\boldsymbol{R}^{-1}\boldsymbol{H}\right)^{-1}\boldsymbol{H}^{\rm T}\boldsymbol{R}^{-1}\boldsymbol{H}$$
$$-\frac{1}{2}\boldsymbol{Q}(\lambda)\left(\boldsymbol{P}^{-1} + \lambda\boldsymbol{H}^{\rm T}\boldsymbol{R}^{-1}\boldsymbol{H}\right) \tag{11}$$

$$\boldsymbol{b}_{\rm s}(\lambda) = \left(\boldsymbol{P}^{-1} + \lambda\boldsymbol{H}^{\rm T}\boldsymbol{R}^{-1}\boldsymbol{H}\right)^{-1}\boldsymbol{H}^{\rm T}\boldsymbol{R}^{-1}\boldsymbol{z}$$
$$+\frac{1}{2}\left(\boldsymbol{Q}(\lambda) - \boldsymbol{Q}_{\rm g}(\lambda)\right)\left(\boldsymbol{P}^{-1}\boldsymbol{\mu}_0 + \lambda\boldsymbol{H}^{\rm T}\boldsymbol{R}^{-1}\boldsymbol{z}\right).$$

Based on (6) and (7), this affine function can again be used to migrate the means and covariances of a GMM representing a predicted/prior PDF based on PFl.

The arbitrariness in the choice of the diffusion matrix in (10) enables us to design a diffusion matrix such that stiffness is reduced and numerical efficiency is improved. In particular, the transient dynamics of a PFl is measured by the condition number $\kappa(\cdot)$ of the nonsingular matrix $\boldsymbol{A}_{\rm s}(\lambda)$ in (11), i.e., the ratio of the largest singular value to the smallest singular value of nonsingular $\boldsymbol{A}_{\rm s}(\lambda)$. A small condition number close to one usually implies a reduced stiffness of the SDE [19]. Consider the following function form of $\boldsymbol{Q}(\lambda)$ in (11), i.e., $\boldsymbol{Q} = c(\boldsymbol{P}^{-1} + \lambda\boldsymbol{H}^{\rm T}\boldsymbol{R}^{-1}\boldsymbol{H})^{-1}$ where c is a constant. We can get $\lim_{c\to\infty}\kappa(\boldsymbol{A}_{\rm s}) = 1$ for $\lim_{c\to\infty}||\boldsymbol{Q}|| = \infty$. However, we cannot make $||\boldsymbol{Q}||$ too large, since it defines an upper bound of the numerical integration error using the Euler-Maruyama method [10].

To balance the stiffness reduction and the error of numerical evaluation of the SDE, we adopt the objective function [10] $J(\boldsymbol{Q}) = \kappa(\boldsymbol{A}) + \alpha c$ where α is a hyperparameter. The optimal solution $\boldsymbol{Q}^*$ obtained by minimizing $J(\boldsymbol{Q})$ is $\boldsymbol{Q}^* = c^*(\boldsymbol{P}^{-1} + \lambda\boldsymbol{H}^{\rm T}\boldsymbol{R}^{-1}\boldsymbol{H})^{-1}$, with $c^* = \max\left\{\sqrt{k_{\max} - k_{\min}}/\sqrt{\alpha} - k_{\min}, 0\right\}$, where $k_{\max}$ and $k_{\min}$ are the absolute values of the largest and smallest eigenvalues of the negative Jacobian matrix of $\boldsymbol{\zeta}_{\rm d\text{-}g}(\boldsymbol{x},\lambda)$ in (9), i.e., $-\boldsymbol{A}_{\rm d\text{-}g} = \frac{1}{2}(\boldsymbol{P}^{-1} + \lambda\boldsymbol{H}^{\rm T}\boldsymbol{R}^{-1}\boldsymbol{H})^{-1}\boldsymbol{H}^{\rm T}\boldsymbol{R}^{-1}\boldsymbol{H}$. For a detailed proof, see [10]. The resulting flow induced proposal PDF can be used for IS in a traditional particle filter [1], for Monte Carlo integration [2], or for the computation of messages in BP-based estimation framework [21].

IV. NUMERICAL EXPERIMENTS AND RESULTS

We evaluate our flow-induced proposal PDF in a 3-D source localization scenario where a volumetric array of receivers provides time-difference of arrival (TDOA) measurements [22]. This scenario is complicated by (i) the highly nonlinear TDOA measurement model, (ii) measurement-origin uncertainty, and (iii) an unknown number of sources to be localized [12], [13]. To address (ii) and (iii), we make use of the BP-based message passing framework introduced in [23], [24]. To address (i), we use our flow-induced GMM proposal PDF for weight computation in the belief update step and Monte Carlo integration in the message computation step. For more details on how PFl-based proposal PDF can be used within BP-based message passing framework, see [12], [13], [18].

A. Source Localization Scenario and Implementation Aspects

In this work, we consider the localization of an unknown number of static sources in a 3-D region of interest (ROI). There are V receivers. Pairs of receivers provide TDOA measurements obtained by cross-correlation. In particular, the TDOA measurement model related to measurement with index m provided by receivers with indexes a and b, can be expressed as $z_{ab}^{(m)} = h_{ab}(\boldsymbol{x}^{(j)}) + v_{ab}^{(m)}$. Here, $h_{ab}(\boldsymbol{x}^{(j)}) = 1/c(||\boldsymbol{x}^{(j)} - \boldsymbol{p}^{(a)}|| - ||\boldsymbol{x}^{(j)} - \boldsymbol{p}^{(b)}||)$, $\boldsymbol{p}^{(a)}$ and $\boldsymbol{p}^{(a)}$ are the 3-D positions of the receivers, c is the propagation speed in the considered medium, and $v_{ab}^{(m)}$ is the additive white noise with variance σ_v^2. The noise $v_{ab}^{(m)}$ is statistically independent across m and across all receiver pairs (a,b). The dependence of a measurement $z_{ab}^{(m)}$ on the source-location $\boldsymbol{x}^{(j)}$ is described by the likelihood function $f(z_{ab}^{(m)}|\boldsymbol{x}^{(j)})$ that can be directly obtained from the TDOA measurement model. This likelihood function has the shape of a hyperboloid (cf. Fig. 1-a). For unambiguouse source localization, the measurements of multiple receiver pairs have to be used. Our scenario is further complicated by (ii) and (iii) discussed above.

Each receiver pair is considered one of S sensors indexed by $s \in \{1, \ldots, S\}$. The receivers of sensor s are indexed (s_a, s_b) and the number of measurements at sensor s is M_s. We consider a topology with $V = 6$ receivers and $S = 9$ sensors is shown in Fig. 1-b. Furthermore, we set $\sigma_v = 0.5\text{ms}$
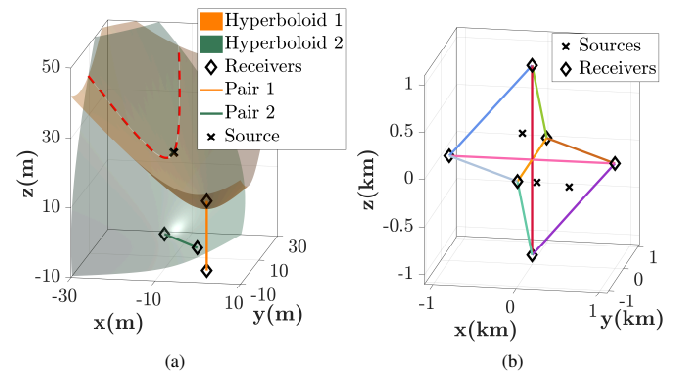

Fig. 1: (a). Source position and hyperboloids resulting from the TDOA measurements of two sensors. Each sensor consists of a receiver pair. A dashed red line indicates the intersection of the two hyperboloids. (b). Receiver and sensor topology used in our simulation. There are six receivers located at the center of each face of the ROI cube. Three sources are randomly placed in the ROI.

and $c = 1500$m/s. The clutter measurements at sensor s follow a uniform PDF on $\|\boldsymbol{q}^{(s_a)} - \boldsymbol{q}^{(s_b)}\|/c$. The mean number of clutter measurements is $\mu_{\mathrm{c}} = 1$ and the probability of source detection, p_{d}, is set to 0.95. The ROI is defined as $[-1000\,\mathrm{m}, 1000\,\mathrm{m}] \times [-1000\,\mathrm{m}, 1000\,\mathrm{m}] \times [-1000\,\mathrm{m}, 1000\,\mathrm{m}]$.

Following [13], [23], TDOA measurements are processed sequentially across sensors. More precisely, let $f(\boldsymbol{x}|\boldsymbol{z}_{1:s-1})$ be the multimodal posterior PDF after sensor update $s-1$. This PDF is represented by N_{g} Gaussian mixture components. For each component n and each TDOA measurement, $z_m^{(s)}$, of sensor s, we perform PFl and update each kernel mean and covariance matrix based on (6) and (7). The resulting $N_{\mathrm{g}} \times M_s$ Gaussian components will be used for IS and Monte Carlo integration within BP-based message passing [23]. The result is an approximate posterior PDF $f(\boldsymbol{x}|\boldsymbol{z}_{1:s})$ represented by N_{g} Gaussian mixture components.

As a reference method for the proposed IS with *optimized stochastic* PFl ("PFl-OS"), we consider bootstrap IS ("BS"), which directly uses the posterior PDF $f(\boldsymbol{x}|\boldsymbol{z}_{1:s-1})$ from previous sensor $s-1$ as proposal PDF. We also use methods with "flow-induced" proposal PDFs using the *deterministic* EDH PFl ("PFl-D") [13]) and the *stochastic* Gromov's PFl ("PFl-S") [18]). Since every method yields different stiffness, for numerical integration, every method requires a different resolution of the pseudo-time. This resolution is defined as the inverse of the time-interval Δ_l. In addition, for every method, a higher resolution is typically needed at the first few steps of the numerical integration. We use an exponentially increasing ratio of the time-interval, i.e., $\Delta_l = \beta \Delta_{l-1}$, where β is a fixed ratio for $l = 2, \ldots, N_\lambda - 1$. Then, we control the interval by two parameters, the initial difference Δ_1 and the increasing ratio β. Note that a larger Δ_1 and β result in fewer discrete steps N_λ and thus to a reduced runtime. For each method, we choose Δ_1 and β to obtain a good runtime-accuracy tradeoff.

B. Results

Table I shows the mean optimal sub-pattern assignment (OSPA) error [25] (with a cutoff threshold at 30) and runtime per run for the different methods and different parametric settings. OSPA and runtime are averaged over 100 Monte Carlo runs. We also list the number of Gaussian components, N_{g}, as well as the number of samples per component, N_{p}.

ID	Method	$(N_{\mathrm{g}}, N_{\mathrm{p}})$	(β, Δ_1)	OSPA	Runtime(s)
1	BS	$(-, 2\mathrm{e}6)$	$(-, -)$	25.89	47.0
2	BS	$(-, 4\mathrm{e}7)$	$(-, -)$	14.73	718.8
3	PFl-D	$(100, 5\mathrm{e}3)$	$(1.5, 1\mathrm{e}-7)$	11.1	202.2
4	PFl-S	$(100, 5\mathrm{e}3)$	$(1.5, 1\mathrm{e}-7)$	8.36	612.7
5	PFl-OS	$(100, 5\mathrm{e}3)$	$(1.5, 1\mathrm{e}-5)$	6.33	490.9
6	PFl-OS	$(100, 5\mathrm{e}3)$	$(2, 1\mathrm{e}-4)$	6.70	295.2

TABLE I: Simulated mean OSPA error and runtime per run for different methods and system parameters.

BS suffers from particle degeneracy and thus yields the highest OSPA. In addition, the large number of $4\mathrm{e}7$ particles results in the largest memory requirements. In contrast, PFl-based methods required much fewer particles. It can be seen that, for the same N_λ value, PFl-S has a smaller OSPA error

compared to PFl-D while its runtime is three times larger. PFl-OS relies on a diffusion matrix $\boldsymbol{Q}$ optimized as discussed in Section III, by setting $\alpha = 0.1$. Notably, PFl-OS yields the lowest OSPA error and, at the same time, has a low runtime.

To better understand the influence of α on the estimation error for different step sizes, we compare the proposed methods for different values of α and different resolutions of pseudo time. Results are shown in Table II. PFl-D and PFl-S results are also listed. When the step size is small, PFl can lead to accurate results despite the stiffness of the underlying differential equation. In this case, traditional methods such as EDH and Gromov's flow also perform well. However, the small step size comes at the cost of a strongly increased runtime. As the step size increases, only PFl-OS with carefully chosen parameter α results in a good estimation accuracy. Notably, our results indicate that PFl-OS has a significantly improved complexity-accuracy tradeoff compared to PFl-D and PFl-S. Further numerical analysis is provided in [26], where the tradeoff between the condition number of the Jacobian matrix $\boldsymbol{A}$ and the norm of the diffusion matrix $\boldsymbol{Q}$ is numerically analyzed for different values of α.

(β, Δ_1) Method	$(1.3, 1\mathrm{e}-13)$	$(1.5, 1\mathrm{e}-5)$	$(2, 1\mathrm{e}-4)$
PFl-D	6.29	30	30
PFl-S	8.70	24.51	28.4
PFl-OS ($\alpha = 0.01$)	4.79	30	30
PFl-OS ($\alpha = 0.1$)	10.28	6.33	6.70
PFl-OS ($\alpha = 0.5$)	10.32	9.36	12.85

TABLE II: Mean OSPA error of different PFl-based IS method and different integration step sizes.

V. CONCLUSION

In this paper, we introduced a general approach to perform IS based on stochastic PFl. Stochastic PFl introduces a diffusion term in the partial differential equation that describes the trajectory of particles. A carefully determined diffusion term reduces the stiffness of this partial differential equation and makes it possible to perform PFl with a lower number of numerical integration steps compared to traditional deterministic PFl. Our method makes it possible to evaluate a "flow-induced" proposal PDF after the parameters of a GMM have been migrated by stochastic PFl. Compared to conventional stochastic PFl, the resulting updating step is asymptotically optimal. Within our method, it is possible to optimize the diffusion matrix that describes the diffusion term to improve the accuracy-computational complexity tradeoff. The presented numerical results in a highly nonlinear 3-D source localization scenario showcased a reduced stiffness and an improved estimating accuracy compared to state-of-the-art deterministic and stochastic PFl. Further research includes flow-induced IS for different types of measurement models [27], application of flow-induced IS to real-world problems [28], [29], and efficient information-seeking control methods based on PFl [30]–[32].

REFERENCES

[1] M. S. Arulampalam, S. Maskell, N. Gordon, and T. Clapp, "A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking," *IEEE Trans. Signal Process.*, vol. 50, no. 2, pp. 174–188, Feb. 2002.

[2] A. Doucet, N. de Freitas, and N. Gordon, *Sequential Monte Carlo Methods in Practice*.   Springer, 2001.

[3] D. Crisan and A. Doucet, "A survey of convergence results on particle filtering methods for practitioners," *IEEE Trans. Signal Process.*, vol. 50, no. 3, pp. 736–746, 2002.

[4] S. Maskell and S. Julier, "Optimised proposals for improved propagation of multi-modal distributions in particle filters," in *Proc. FUSION-13*, 2013, pp. 296–303.

[5] P. Bickel, B. Li, and T. Bengtsson, "Sharp failure rates for the bootstrap particle filter in high dimensions," in *Pushing the Limits of Contemporary Statistics: Contributions in Honor of Jayanta K. Ghosh*, vol. 3. Beachwood, OH, USA: Inst. Math. Statist., 2008, pp. 318–329.

[6] F. Daum and J. Huang, "Nonlinear filters with log-homotopy," in *Proc. SPIE-07*, Aug. 2007, pp. 423–437.

[7] ——, "Nonlinear filters with particle flow induced by log-homotopy," in *Proc. SPIE-09*, May 2009, pp. 76–87.

[8] F. Daum, J. Huang, and A. Noushin, "Exact particle flow for nonlinear filters," in *Proc. SPIE-10*, Apr. 2010, pp. 92–110.

[9] ——, "New theory and numerical results for Gromov's method for stochastic particle flow filters," in *Proc. FUSION-18*, 2018, pp. 108–115.

[10] L. Dai and F. Daum, "On the design of stochastic particle flow filters," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 59, no. 3, pp. 2439–2450, 2023.

[11] Y. Li and M. Coates, "Particle filtering with invertible particle flow," *IEEE Trans. Signal Process.*, vol. 65, no. 15, pp. 4102–4116, Aug. 2017.

[12] W. Zhang and F. Meyer, "Graph-based multiobject tracking with embedded particle flow," in *Proc. IEEE RadarConf-21*, Atlanta, GA, USA, May 2021.

[13] ——, "Multisensor multiobject tracking with improved sampling efficiency," *IEEE Trans. Signal Process.*, vol. 72, pp. 2036–2053, 2024.

[14] J. Jang, F. Meyer, E. R. Snyder, S. M. Wiggins, S. Baumann-Pickering, and J. A. Hildebrand, "Bayesian detection and tracking of odontocetes in 3-D from their echolocation clicks," *J. Acoust. Soc. Am.*, vol. 153, no. 5, p. 2690–2705, May 2023.

[15] P. Bunch and S. Godsill, "Approximations of the optimal importance density using Gaussian particle flow importance sampling," *J. Amer. Statist. Assoc.*, vol. 111, no. 514, pp. 748–762, Aug. 2016.

[16] Y. Li, L. Zhao, and M. Coates, "Particle flow auxiliary particle filter," in *Proc. IEEE CAMSAP-15*, Cancun, Mexico, Dec. 2015, pp. 157–160.

[17] S. Pal and M. Coates, "Particle flow particle filter using Gromov's method," in *Proc. IEEE CAMSAP-19*, Guadeloupe, France, Dec. 2019, pp. 634–638.

[18] W. Zhang, M. J. Khojasteh, and F. Meyer, "Particle flows for source localization in 3-D using TDOA measurements," in *Proc. FUSION-24*, Venice, Italy, 2024.

[19] P. E. Kloeden and E. Platen, *Numerical Solution of Stochastic Differential Equations*.   Berlin, Germany: Springer, 1992.

[20] D. F. Crouse, "Particle flow filters: Biases and bias avoidance," in *Proc. FUSION-22*, Ottawa, Canada, 2019, pp. 1–8.

[21] F. R. Kschischang, B. J. Frey, and H.-A. Loeliger, "Factor graphs and the sum-product algorithm," *IEEE Trans. Inf. Theory*, vol. 47, no. 2, pp. 498–519, Feb. 2001.

[22] F. Gustafsson and F. Gunnarsson, "Positioning using time-difference of arrival measurements," in *Proc. IEEE ICASSP-03*, vol. 6, Hong Kong, China, Apr. 2003, pp. 553–556.

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

[24] F. Meyer, P. Braca, P. Willett, and F. Hlawatsch, "A scalable algorithm for tracking an unknown number of targets using multiple sensors," *IEEE Trans. Signal Process.*, vol. 65, no. 13, pp. 3478–3493, Jul. 2017.

[25] D. Schuhmacher, B.-T. Vo, and B.-N. Vo, "A consistent metric for performance evaluation of multi-object filters," *IEEE Trans. Signal Process.*, vol. 56, no. 8, pp. 3447–3457, 2008.

[26] W. Zhang, M. J. Khojasteh, N. A. Atanasov, and F. Meyer, "Importance sampling with stochastic particle flow and diffusion optimization: Supporting results," 2024, https://fmeyer.ucsd.edu/SPL-2024-SR.pdf.

[27] F. Meyer and J. L. Williams, "Scalable detection and tracking of geometric extended objects," *IEEE Trans. Signal Process.*, vol. 69, pp. 6283–6298, Oct. 2021.

[28] M. Liang and F. Meyer, "Neural enhanced belief propagation for multiobject tracking," *IEEE Trans. Signal Process.*, vol. 72, pp. 15–30, 2024.

[29] L. Watkins, P. Stinco, A. Tesei, and F. Meyer, "A probabilistic focalization approach for single receiver underwater localization," in *Proc. FUSION-24*, Venice, Italy, 2024.

[30] W. Zhang, B. Teague, and F. Meyer, "Active planning for cooperative localization: A Fisher information approach," in *Proc. Asilomar-22*, Pacific Grove, CA, USA, 2022, pp. 795–800.

[31] J. A. Placed, J. Strader, H. Carrillo, N. Atanasov, V. Indelman, L. Carlone, and J. A. Castellanos, "A survey on active simultaneous localization and mapping: State of the art and new frontiers," *IEEE Trans. Robot.*, vol. 39, no. 3, pp. 1686–1705, 2023.

[32] A. Asgharivaskasi and N. Atanasov, "Semantic OcTree mapping and shannon mutual information computation for robot exploration," *IEEE Trans. Robot.*, vol. 39, no. 3, pp. 1910–1928, 2023.