Stochastic Particle Flow for Nonlinear High-Dimensional Filtering Problems

Flávio Eler De Melo[*1], Simon Maskell[†1], Matteo Fasiolo[‡2], and Fred Daum[§3]

[1]Department of Electrical Engineering and Electronics, University of Liverpool
[2]Department of Mathematics, University of Bristol
[3]Raytheon Company

March 24, 2017

Abstract

A series of novel filters for probabilistic inference that propose an alternative way of performing Bayesian updates, called particle flow filters, have been attracting recent interest. These filters provide approximate solutions to nonlinear filtering problems. They do so by defining a continuum of densities between the prior probability density and the posterior, i.e. the filtering density. Building on these methods' successes, we propose a novel filter. The new filter aims to address the shortcomings of sequential Monte Carlo methods when applied to important nonlinear high-dimensional filtering problems. The novel filter uses equally weighted samples, each of which is associated with a local solution of the Fokker-Planck equation. This hybrid of Monte Carlo and local parametric approximation gives rise to a global approximation of the filtering density of interest. We show that, when compared with state-of-the-art methods, the Gaussian-mixture implementation of the new filtering technique, which we call Stochastic Particle Flow, has utility in the context of benchmark nonlinear high-dimensional filtering problems. In addition, we extend the original particle flow filters for tackling multi-target multi-sensor tracking problems to enable a comparison with the new filter.

1 Introduction

Stochastic filtering in high-dimensional spaces is a challenging estimation task because of two fundamental issues:

- *The curse of dimensionality.* In a statistical experiment, as the sample space's dimensionality increases a finite number of realizations can only populate the space to an increasingly sparse extent [1]. This issue makes it challenging to use approximation based on realizations of the state.

- *The infinite number of parameters required to describe a general probability density on a continuous state space.* Such a density, in common with any other real function, can always be exactly described using a power series with infinitely many terms. In all but a very few cases, where the density is known to have a specific parametric form, using a finite set of parameters is necessarily an approximation to this complete description. The fidelity of such approximations falls rapidly as dimension increases. This issue makes it challenging to define a parameterization that uses a number of parameters that scales only gently with dimension.

The development of the vast majority of practical filters focuses on how to accurately represent generic probability densities. However, in the view of the authors, relatively few filters are systematically developed with the explicit intent of efficiently expressing densities in high-dimensional spaces. There does appear to be a consensus that the statistical efficiency associated with expressing high-dimensional filtering densities can be improved by simulating tempering distributions [2–5]. Such approaches involve introducing intermediate distributions such that it is easier to migrate between these intermediate distributions than it is to migrate directly

---

[*]flavio.de-melo@liverpool.ac.uk

[†]s.maskell@liverpool.ac.uk

[‡]matteo.fasiolo@bristol.ac.uk

[§]frederick_e_daum@raytheon.com

arXiv:1511.01448v3 [stat.ME] 23 Mar 2017

from the prior to the posterior. The use of such intermediate distributions stabilizes the sampling procedure and maintain the variance of the Monte Carlo weights at an acceptable level. Bickel *et al.* [6] considered, in the context of a bootstrap particle filter, the number of intermediate distributions needed to use such tempering successfully. They prove that, as the dimensionality increases, the number of intermediate distributions needed to accurately represent a high-dimensional density becomes practically infinite. This implies that considering a continuum, i.e. infinite number, of intermediating distributions [4] might be the basis of a successful approach. This implication is corroborated by the reported success of Markov chain Monte Carlo (MCMC) algorithms that populate high-dimensional state spaces efficiently using approximations of problem-specific continuous-time processes [7–10][1].

Techniques for continuous-time processes stem from the seminal work by Stratonovich [11], Kushner [12] and Zakai [13] on filtering theory. The most popular instance of such filters is the so-called Kalman-Bucy filter [14], the continuous-time counterpart of the Kalman filter. More general filters directly approximate solutions to the Kushner-Stratonovich equation either by a finite-dimensional density parameterization [15, 16] or by Monte Carlo methods [17–19]. Other important finite-order filters that appeal to an unusual formalism of multiple stochastic integrals [20, 21] are worth mentioning as well.

Continuous-time filtering may be seen by some as an idealized problem of limited practical utility. However, recent research [22, 23] has shown that continuous-time filtering can offer key insights into the fundamental principles necessarily associated with successful filtering in high-dimensions: the effects of local, continuous spatial properties of the observation process need to be incorporated in the solution. As identified by Bickel *et al.* [6], the information in the data gives rise to the notion of the *effective dimension* of the space. It is this effective dimension, and not the dimension of the state space itself, that actually affects the statistical efficiency of inference algorithms. This implies that tempering only addresses part of the problem, remaining the local observation properties to be incorporated. Recently, based on a principled approach, Rebeschini & van Handel [23] proposed to decompose the state space into separate blocks. The global solution to the inference problem is then constructed by combining the local solutions for each of the blocks. The paper goes on to demonstrate that, by using the *decay of correlations* property[2], it is possible to develop particle filters based on local solutions in such a way that the approximation error does not depend on the dimension of the state space.

Largely independently, the idea of filtering via a continuum of intermediate distributions seems to have its appeal revigorated as several new methods have been proposed for progressive Bayesian updates, whose continuity is considered in the limit, aiming to gradually introduce the effect of each observation. Those filters have emerged either in a variational, ensemble-based, or sequential Monte Carlo framework. In the variational framework, the new methods presented in [24–26] pose the filtering problem as a multi-step optimization problem for which the cost function is an approximated distance between a parameterised density and the actual filtering density. In the ensemble-based framework, the methods [27–29] are focused on data assimilation problems and apply ideas of optimal transport along with continuous-time filtering to generate multiple independent solutions that are combined to obtain a single solution of an inference problem. The methods in the sequential Monte Carlo framework explore extensions or alternatives to particle filters (e.g., [30–34]), or simply capitalize on techniques for properly choosing a sequence of *bridging* importance densities (e.g., [3–5, 35]), carrying on the intent to overcome the widely known problem of particle filters called *degeneracy* or collapse of weights [36–39].

Among those methods one in particular has recently attracted interest in an Engineering context where it has been described as *particle flow*. The performance that has been reported is remarkable and the literature is extensive with several variants having been developed over recent years (see, for example, [32, 40–43]). The development of particle flow draws on analogies to problems that arise in Fluid Dynamics and Electromagnetism. These filters *flow* probability masses (particles) from a prior probability space to one that is updated according to a set of measurements without the need to perform a Bayesian update explicitly. All particle flow algorithms explore the concept of a homotopy between the prior and posterior probability spaces, implicitly describing a joint measure that couples the prior and posterior probability measures. This idea is in the heart of the Kantorovich's optimal tranportation problem [44] that, by evoking deterministic transport maps for very simple cost functions and dynamic constraints, yields an essential explanation on why original particle flow methods work well based on deterministic rules to flow the particles.

When the sequential filtering problem involves non-compactly supported densities, solving it via deterministic (optimal) transport is not straightforward. A solution would require either a non-trivial approximation of the highly nonlinear Monge-Ampère equation [45] or adapting classical solutions constructed for measures on

---

[1]Note that while practical implementation of these techniques necessarily involves finite time-horizons, the continuous-time processes are typically designed such that, as time tends to infinity, the distribution of the samples from the process tends to the distribution of interest. This is in contrast to the use of tempering distributions, where samples from the posterior are generated after a defined (finite) number of steps between intermediate distributions.

[2]A spatial counterpart of the stability property of nonlinear filters, by which a probability mass is strongly correlated to masses within its neighbourhood but has negligible correlation with respect to the remaining areas of the state space.

bounded sets [46, 47]. In these approaches, severe technical difficulties may arise and not all the effects on the estimation errors are clearly known. A continuously evolving, exact, optimal transport map would require a complete description, at all time instants, of an embedding dynamic field that induces a transference plan to correctly move particles. If the posterior density could be completely characterized beforehand then the optimal transport problem could be numerically solved by the multiple-step augmented-Lagrangian optimization method as proposed by Benamou & Brenier [48]. However, detailed knowledge of the posterior would imply a direct answer to the filtering problem. As an alternative, theoretically speaking, a complete description of the optimal field could be achieved by solving the Monge-Ampère equation for any possible location of particles on the state space. Notwithstanding, the Monge-Ampère equation admits exact solutions only for few particular cases [45] and would also require a thorough description of the posterior density in advance.

In this scenario, one feasible approach is that advocated by particle flow methods, which take simplifying assumptions on the embedding dynamic field in order to avoid both optimization over a parametric class of transport maps and explicit solution of the associated elliptic partial differential equation. However, in our experience, these symplifying assumptions result in approximated filtering densities providing accurate estimates for the first-order moment but estimates for second and higher-order moments whose quality is highly dependent on the problem and algorithm settings (e.g., [49]). In practice, particle flow methods address this latter issue by either relying on a companion filter [50, 51] or using the sample covariance matrix with shrinkage and Tikhonov regularization [52] to be able to estimate the second-order moment.

We conjectured if appealing to stochastic transport could provide a new avenue for solving the filtering problem. Fortunately, a variational formulation of the Fokker-Planck equation as a gradient flow, as exposed by Jordan *et al.* [53], enables the precise interpretation that, if a transport operation is to be understood as a diffusion, then it minimizes the free energy functional of the process with respect to the Wasserstein metric over an admissible class of probability measures. Relying on this formulation, it is straightforward to obtain a transport rule, optimal in terms of minimizing the free energy functional, as a Langevin stochastic process. This rule is based simply on the assumptions of stationarity of the filtering distribution (Gibb's distribution) and on potential conditions, for which an embedding stationary field is exactly derived.

In this article we take into consideration the findings presented by Jordan *et al.* [53], incorporate the description of statistically efficient processes in high-dimensional spaces as proposed by Girolami & Calderhead [10], and incorporate local properties of the observation process to formulate a stochastic variant of particle flow[3]. This new *stochastic particle flow* (SPF) involves defining a Langevin diffusion such that a posterior measure from a previous step, under a known stationary potential field, is diffused onto the current posterior measure, satisfying the Fokker-Planck equation to produce an accurate approximation of the filtered density. This process involves guiding local solutions of the Fokker-Planck equation in such a way that we construct a mixture that approximates the posterior. As we will discuss later on, the SPF method we propose is essentially built as a Gaussian sum filter (SPF-GS), nevertheless, it is possible to use a similar formulation to define an implementation strategy based on a marginal particle filter (SPF-MPF). This variant demonstrates versatility of the SPF to algorithm settings.

It is worth mentioning that our resulting SPF technique is in the same ethos as the method recently developed by Bunch & Godsill [35, 55]. However, in constrast to our approach, their method (i) is based on the homotopy between the prior and posterior spaces, (ii) assumes the particle flow is an Ornstein-Uhlenbeck process whose scaling parameter determines the rate of diffusion of samples' paths, (iii) proposes weights that must be updated iteratively by a partial differential equation (PDE) describing how the unnormalized log-density evolves with a pseudo-time variable; (iv) is articulated as a standard (not marginal) particle filter.

The outline of the article is as follows. We begin by reviewing the stochastic filtering problem in a sequential Monte Carlo framework in Section 2. We abstract the solution in terms of a general map that could adopt any valid method to perform the filtering update. In Section 3, we present a brief overview of the original particle flow methods. We discuss their principles in order to further clarify these methods and motivate the natural step towards stochastic particle flow. In Section 4 we derive the generic SPF algorithm by describing the proposed dynamics of probability masses, describing the associated stationary solution to the Fokker-Planck equation, and constructing the stochastic flow. Algorithmic details are given and relate to how to compute the diffusion matrix, how to integrate the stochastic flow, and how to select the simulation time horizon and integration step size. We present the stochastic particle flow implementation using a Gaussian sum filter (SPF-GS) in Section 5. We achieve this by considering the posterior to be well approximated as a mixture of local solutions to the flow. Similarly, in Section 6 we show the SPF articulated as a marginal particle filter (SPF-MPF) by setting the importance density as a mixture analogous to that generated by the SPF-GS. Section 7 then illustrates the SPF's properties by a series of toy problems, and compares the performance of the SPF and

---

[3]Existing particle flow algorithms (including, perhaps surprisingly, that known as non-zero diffusion particle flow [54]) propagates particles deterministically.

other state-of-the-art methods in the context of three instructive multi-sensor or multi-target tracking problems: multi-sensor bearing-only tracking, convoy tracking and inference on a large network of sensors (as in [56]). In the comparisons for the multi-sensor bearing-only and convoy tracking problems, we included extensions to two of the most effective (original) particle flows, namely, the Gaussian particle flow (GPF) [57] and the scaled-drift particle flow (SDPF) [54]. Finally, Section 8 concludes.

2 Sequential Monte Carlo Filtering

In this section we report the filtering framework within which the particle flows may be formalized. Let $\{x_t \in \mathcal{X} : t \in \mathbb{R}_+\}$ be a sequence of states generated through time by a known continuous-time state process, modelled as a Markov process, and $\{y_{t_k} \in \mathcal{Y} : t_k \in \mathbb{R}_+, k \in \mathbb{N}\}$ be a sequence of discrete-time observations of the process generated by an observation model. In the classical filtering problem, one is required to compute the best estimate of a function of interest φ of the state, given all observations realized up to the time instant t_k, i.e.,

$$\hat{\varphi}_k = \mathbb{E}\left[\varphi(x_{t_k}) | y_{t_1}, y_{t_2}, \ldots, y_{t_k}\right]. \tag{2.1}$$

To simplify notation, we will denote all variables at discretized time instants by the time indexes $k \in \mathbb{N}$, and write $y_{1:k} \triangleq \{y_1, y_2, \ldots, y_k\}$. Now consider a set of particles $\{x_{k-1}^{(i)}, w_{k-1}^{(i)} : i = 1, \ldots, N\}$ constituting samples that can be used to approximate a filtering probability density $p(x_{k-1}|y_{1:k-1})$ by means of a Monte Carlo measure satisfying

$$\sum_{i=1}^{N} w_{k-1}^{(i)} \delta(x_{k-1} - x_{k-1}^{(i)}) \xrightarrow{N \to \infty} p(x_{k-1}|y_{1:k-1}), \tag{2.2}$$

to represent the convergence as follows for any test function $\varphi : \mathcal{X} \to \mathbb{R}$:

$$\sum_{i=1}^{N} w_{k-1}^{(i)} \varphi(x_{k-1}^{(i)}) \xrightarrow{N \to \infty} \int_{\mathcal{X}} \varphi(x_{k-1}) p(x_{k-1}|y_{1:k-1}) dx_{k-1} \quad \text{almost surely.} \tag{2.3}$$

Given a new observation obtained at instant k, one wishes to find a procedure to transform the set of particles $\{x_{k-1}^{(i)}, w_{k-1}^{(i)}\}$ into a new set of particles $\{x_k^{(i)}, w_k^{(i)} : i = 1, \ldots, N\}$ that incorporates the effect of the latest observation in order to estimate the filtered entity as

$$\hat{\varphi}_k \approx \sum_{i=1}^{N} w_k^{(i)} \varphi(x_k^{(i)}). \tag{2.4}$$

In theory, the filtering problem in the sequential Monte Carlo form can be solved by any map $\mathcal{T} : \mathcal{X} \times \mathcal{Y} \to \mathcal{X}'$, $\mathcal{T} \in \mathcal{C}^1(\mathbb{R}^{n_x}) \times \mathcal{C}^0(\mathbb{R}^{n_y})$, where $|\mathcal{X}'| = |\mathcal{X}|$, that implements

$$x_k^{(i)} \quad := \quad \mathcal{T}(x_{k-1}^{(i)}, y_k); \ i = 1, \ldots, N; \tag{2.5}$$

$$w_k^{(i)} \quad := \quad \det \mathcal{J}_{x_{k-1}}[\mathcal{T}]^{-1} w_{k-1}^{(i)}; \tag{2.6}$$

where $\mathcal{J}_{x_{k-1}}[.]$ is the Jacobian matrix with respect to x_{k-1}, and such that

$$\sum_{i=1}^{N} w_k^{(i)} \delta(x_k - x_k^{(i)}) \xrightarrow{N \to \infty} p(x_k|y_{1:k}). \tag{2.7}$$

Although most practical filters implement the mapping (2.5) in terms of discrete Bayesian updates, there should be no objection to the general idea of considering the map $\mathcal{T}$ as a transform continuous in time within $t_{k-1} < t \leq t_k$. This idea establishes the basis for the particle flow filters.

3 Particle Flow

This section aims to present a brief overview on the particle flow methods, to discuss their principles, and to set the background for the introduction of the stochastic particle flow. The key idea of the particle flow is to transfer a set of probability masses by an operation that transports the prior probability measure onto the posterior measure. This operation realizes the measurement update smoothly in order to express a filtering

entity, usually an estimate. The mechanism implied is, therefore, a filtering algorithm that avoids the need to perform a Bayesian measurement update explicitly.

Given a set of particles $\left\{ x^{(i)}(\lambda) \in \mathbb{R}^{n_x} : i = 1, \ldots, N \right\}$ dependent on a continuous pseudo-time variable $\lambda \in [0, 1]$, where n_x is the number of dimensions of the state space, and such that $x^{(i)}(0) = x_{k-1}^{(i)}$ and $x^{(i)}(1) = x_k^{(i)}$, the transformation of the particles is accomplished by solving through $0 < \lambda \leq 1$ an ordinary differential equation (ODE) referred to as the flow equation

$$\frac{dx}{d\lambda} = \mu(x, \lambda), \quad x^{(i)}(0) \sim p_0(x). \tag{3.1}$$

The varieties of particle flow methods rely on how one defines the flow drift $\mu(x, \lambda)$, which in turn depends on the assumptions made to solve the associated continuity equation

$$\frac{\partial p}{\partial \lambda} = -\nabla_x \cdot (\mu \cdot p), \quad p(x, 0) = p_0(x). \tag{3.2}$$

The operator $\nabla_x \cdot (.)$ is the divergence operator and the drift can be understood as a vector field $\mu(x, \lambda) \in \mathbb{R}^{n_x}$ that is not uniquely determined for a given probability density $p(x, \lambda)$. In the optimal transportation literature the vector field is usually determined by the constraint that it minimizes the kinetic energy. In that case, the flow equation (3.1) can be written in terms of a dynamic potential field as $\mu(x, \lambda) = \mathcal{M}^{-1}\nabla_x \psi(x, \lambda)$ [44], where $\mathcal{M}$ is a positive-definite mass matrix, ∇_x is the gradient operator, and $\psi(x, \lambda)$ is a dynamic potential function that satisfies the elliptic PDE

$$\nabla_x \cdot \left(p(x, \lambda) \mathcal{M}^{-1}\nabla_x \psi(x, \lambda) \right) = -\partial_\lambda p(x, \lambda). \tag{3.3}$$

An exact solution to equation (3.3) has been derived by Reich [27] considering Gaussian likelihood functions. In more general settings, if the target posterior density $\pi(x)$ could be thoroughly characterized in advance, the numerical solution to this problem could be achieved by the multiple-step augmented-Lagrangian optimization method as proposed by Benamou & Brenier [48]. However, availability of a detailed description of the posterior density would constitute a direct answer to the filtering problem. Similarly, the well known flow constructed by Dacorogna & Moser [47], appropriate for mapping measures on bounded open sets, could be adapted for problems involving non-compactly supported densities as the solution of the p-Laplacian equation [58]

$$\nabla_x \cdot (a(x, \lambda)\nabla_x \vartheta(x, \lambda)) = \pi(x) - p(x, \lambda), \tag{3.4}$$

where $p(x, \lambda)$ and $\pi(x)$ are the intermediate and target densities respectively. Function $a(x, \lambda) \geq 0$, $a(x, \lambda) \in \mathbb{L}^\infty$ ($\mathbb{L}^\infty$-space[4]), is a Lagrange multiplier that scales the distance of optimal transportation, whereas the term $\nabla_x \vartheta(x, \lambda)$ gives the direction of optimal tranportation. As mentioned before, these transport-based solutions are not straightforwardly applicable to filtering problems as they would require anticipative approximations of the target probability density, and the solution by Dacorogna & Moser [47] would require truncation of the involved densities to bound their support.

Indeed original particle flows do not follow the classical transport-based methodology but rather take simplifying assumptions on the dynamic potential field, avoiding the complexity of solving the elliptic PDEs (3.3) and (3.4). Specifically, the particle flows are derived from a programmed sequence of a dynamic potential field that roughly solves the equation (3.2). As examples we refer the reader to the incompressible particle flow [32], the Gaussian or exact particle flow [57], and the non-zero *"diffusion"* particle flow [54], which is not actually a diffusion, but simply takes into account a diffusion term to scale and/or offset the drift term.

In a closely related problem, as an alternative to the solution of elliptical equations or to original particle flows, it is possible to demonstrate that if the drift solves the continuity equation (3.2), under a stationary potential field (conservative) related to an invariant, locally[5] log-concave density of the form $p(x, T) = \pi(x) \propto \exp(-\psi(x))$, then the flow (3.1) produces the maximum-a-posteriori (MAP) estimate, $\hat{x}_{MAP}$, after an appropriate time horizon $\lambda \geq T$ (see *Theorem 12* in the Appendix A.2). A similar concept is used in optimization algorithms based on gradient descent. An evident problem with this approach is that, regardless of providing a MAP estimate, it is unable to capture higher-order aspects of a target posterior density. Thus, under the assumption of a stationary potential field, a stochastic particle flow seems suitable to describe a filtering density precisely up to an arbitrary moment order, by following the dynamics of a diffusion that minimizes the free energy functional (see [53] for details). Such stochastic flow would propagate a probability density according to the Fokker-Planck equation.

---

[4]The $\mathbb{L}^\infty$-space generalises the $\mathbb{L}^p$-spaces to $p = \infty$. An $\mathbb{L}^p$-space describes the set of all functions f for which the norm $\|f\|_p = \left(\int_{\mathcal{X}} |f|^p \right)^{1/p}$ converges. The concept is analogous for the $\mathbb{L}^\infty$-space although its norm is defined by the essential supremum.
[5]Log-concave in the vicinity of the density maxima.

This observation becomes fundamental when we note that, loosely speaking, obtaining a precise approximation of a stationary potential field requires less effort than obtaining a sequence of accurate approximations of a dynamic potential field. In this context, approximating a dynamic potential field forms the basis for the classical transport methodology (e.g., [27]).

4 Stochastic Particle Flow

This section derives stochastic particle flow based on a stationary solution to the Fokker Planck equation. We capitalize on the fact that, under certain conditions on the drift and diffusion terms of a stochastic process, there is a stationary solution that satisfies a variational principle, minimizing a certain convex free energy functional over an admissible class of probability densities. The Fokker–Planck equation is shown to follow the direction of steepest descent of the associated free energy functional [53] at each instant of time, rendering a process where the entropy is maximized, i.e., a diffusion.

In Section 4.1 we set dynamics for stochastic particle flow. Section 4.2 derives the stationary solution to the Fokker-Planck equation such that the particles follow the Langevin dynamics. In Section 4.3 we show how to specify the Langevin dynamics to solve the specific problem of interest. In Section 4.4 we discuss the interpretation of and possible choices for the diffusion matrix; in Section 4.5 we present the integration methods used to sample from the Langevin dynamics; and in Section 4.6 we discuss criteria for choosing the algorithm's parameters (the step size and time horizon).

4.1 Dynamics of Particles

Assuming that a set of particles $\{x^{(i)}(\lambda) : i = 1, \ldots, N\}$ follows a diffusion process $\{X_\lambda\}_{\lambda \geq 0}$ when subject to a Bayesian measurement update, the dynamics of the particles can, in general, be described by the Îto stochastic differential equation

$$dX_\lambda = \mu(X_\lambda, \lambda)\, d\lambda + \sigma(X_\lambda, \lambda)\, dW_\lambda, \quad X_0 = X(0); \tag{4.1}$$

such that the associated probability distribution, $p(x, \lambda)$, is continuously evolving with respect to the pseudo-time variable $\lambda \in \mathbb{R}_+$, where $\{W_\lambda\}$ is a standard Brownian motion, $\mu(X_\lambda, \lambda)$ is the drift vector and $\sigma(X_\lambda, \lambda)$ is the diffusion coefficient. It is well known [59, 60] that the probability density $p(x, \lambda)$ of an n_x-dimensional random state vector x under the dynamics of (4.1) has a deterministic evolution according to the Fokker-Planck equation

$$\frac{\partial}{\partial \lambda} p(x, \lambda) = -\sum_{i=1}^{n_x} \frac{\partial}{\partial x_i} [\mu_i(x, \lambda)\, p(x, \lambda)]$$
$$+ \frac{1}{2} \sum_{i=1}^{n_x} \sum_{j=1}^{n_x} \frac{\partial}{\partial x_i} \frac{\partial}{\partial x_j} [D_{ij}(x, \lambda)\, p(x, \lambda)], \tag{4.2}$$
$$p(x, 0) = p_0(x), \ \lambda \geq 0;$$

where $x = [x_1, \ldots, x_{n_x}]^T$, $\mu = [\mu_1, \ldots, \mu_{n_x}]^T$, and

$$D_{ij}(x, \lambda) = \sum_{k=1}^{n_x} \sigma_{ik}(x, \lambda)\, \sigma_{jk}(x, \lambda), \tag{4.3}$$

for an n_x-dimensional Wiener process $\{W_\lambda\}$. In its usual form, as described in Physics, the equation reads

$$\frac{\partial}{\partial \lambda} p = -\nabla_x \cdot [\mu p] + \frac{1}{2} \nabla_x \cdot [D \nabla_x p]. \tag{4.4}$$

We assume that the diffusion coefficient σ is locally independent of x, giving rise to a local diffusion matrix $D(\lambda) = \sigma(\lambda)\sigma(\lambda)^T$ that is invariant to the divergence operator in the vicinity of each particle. This means that, at a given time instant, the diffusion term in (4.4) evolves at a rate proportional to the curvature of a (Riemann) manifold that is approximately constant in the neighbourhood of each particle. This assumption does not affect the generality of the concepts applied in our derivation for two reasons: it results in a stochastic particle flow that is missing a simple term, of the form $\sigma(x, \lambda) \cdot \partial_x [\sigma(x, \lambda)]$, that could be incorporated if needed; in practice, any probability density can be well approximated by a mixture of densities whose covariances are locally constant with respect to the state [61] (i.e., $\partial_x [\sigma(x, \lambda)] = 0$ locally). Additionally, as evidenced in [10], keeping the diffusion coefficient fixed for each sampling step does not perturb the target distribution.

4.2 Stationary Solution of the Fokker-Planck Equation

A stationary solution to the equation (4.4) should satisfy

$$\frac{\partial}{\partial \lambda} p\left(\mathrm{x}, \lambda\right) \overset{\lambda \to \infty}{\longrightarrow} 0. \tag{4.5}$$

By writting

$$\nabla_\mathrm{x} \cdot S \triangleq \nabla_\mathrm{x} \cdot \left[\mu\, p\right] - \frac{1}{2} \nabla_\mathrm{x} \cdot \left[D \nabla_\mathrm{x} p\right], \tag{4.6}$$

the definition of the probability current becomes clear:

$$
\begin{aligned}
S\left(\mathrm{x}, \lambda\right) &= \mu\left(\mathrm{x}, \lambda\right) p\left(\mathrm{x}, \lambda\right) - \frac{1}{2} D\left(\lambda\right) \cdot \nabla_\mathrm{x} p\left(\mathrm{x}, \lambda\right) \\
&= p\left(\mathrm{x}, \lambda\right) \left[\mu\left(\mathrm{x}, \lambda\right) - \frac{1}{2} D\left(\lambda\right) \cdot \nabla_\mathrm{x} \log p\left(\mathrm{x}, \lambda\right)\right].
\end{aligned}
\tag{4.7}
$$

Since the stationary condition requires

$$\frac{\partial}{\partial \lambda} p\left(\mathrm{x}, \lambda\right) = -\nabla_\mathrm{x} \cdot S\left(\mathrm{x}, \lambda\right) \overset{\lambda \to \infty}{\longrightarrow} 0, \tag{4.8}$$

the probability current is required to vanish as $\lambda \to \infty$. The probability current can only vanish if the drift $\mu\left(\mathrm{x}, \lambda\right)$ can be expressed as the gradient of a potential function [62], cancelling out the terms within brackets in (4.7). We write the drift as the gradient of a stationary potential function according to

$$\mu\left(\mathrm{x}, \lambda\right) = -\frac{1}{2} D\left(\lambda\right) \cdot \nabla_\mathrm{x} \Phi\left(\mathrm{x}\right). \tag{4.9}$$

The necessary and sufficient conditions for the existence of $\Phi\left(\mathrm{x}\right)$ are the potential conditions [62]

$$\frac{\partial \mu_i}{\partial x_j} = \frac{\partial \mu_j}{\partial x_i}, \quad \forall i \neq j. \tag{4.10}$$

Provided that the probability current vanishes as $\nabla_\mathrm{x} \log p\left(\mathrm{x}, \lambda\right) \to -\nabla_\mathrm{x} \Phi\left(\mathrm{x}\right)$, we obtain the stationary solution, $p_{st}\left(\mathrm{x}\right)$, as

$$p\left(\mathrm{x}, \lambda\right) \overset{\lambda \to \infty}{\longrightarrow} p_{st}\left(\mathrm{x}\right) = \frac{1}{Z} e^{-\Phi\left(\mathrm{x}\right)}, \tag{4.11}$$

where

$$Z = \int_{\mathbb{R}^{n_x}} e^{-\Phi\left(\mathrm{x}\right)} d\mathrm{x} \tag{4.12}$$

must be positive and finite. We promptly recognise (4.11) as analogous to the Gibbs distribution. It is verifiable that (see, for example, [63]) the Gibbs distribution minimizes the free energy functional over all probability densities on $\mathbb{R}^{n_x}$. It can also be shown that the stationary solution is the first eigenfunction of the Fokker-Planck equation, corresponding to the eigenvalue zero [62].

4.3 The Stochastic Flow

The general stochastic particle flow is derived by setting the stationary solution, $p_{st}\left(\mathrm{x}\right)$, to be the target posterior density, $\pi\left(\mathrm{x}\right) = p\left(\mathrm{x}|\mathrm{y}_{1:k}\right)$, to give

$$
\begin{aligned}
p_{st}\left(\mathrm{x}\right) &:= p\left(\mathrm{x}|\mathrm{y}_{1:k}\right), \\
\frac{e^{-\Phi\left(\mathrm{x}\right)}}{Z} &= \frac{p\left(\mathrm{y}_k|\mathrm{x}\right) p\left(\mathrm{x}|\mathrm{y}_{1:k-1}\right)}{p\left(\mathrm{y}_k|\mathrm{y}_{1:k-1}\right)}, \\
\Phi\left(\mathrm{x}\right) &= -\log p\left(\mathrm{y}_k|\mathrm{x}\right) - \log p\left(\mathrm{x}|\mathrm{y}_{1:k-1}\right).
\end{aligned}
\tag{4.13}
$$

Given a valid potential function $\Phi\left(\mathrm{x}\right)$ provides the stationary solution, all potential functions of the form $\Phi\left(\mathrm{x}\right) \pm K$ for any constant $K \in \mathbb{R}$ are also valid. We can therefore choose a valid potential function in (4.13) such that $p\left(\mathrm{y}_k|\mathrm{y}_{1:k-1}\right) = Z$. By using equation (4.9), we obtain

$$
\begin{aligned}
\mu\left(\mathrm{x}, \lambda\right) &= -\frac{1}{2} D\left(\lambda\right) \cdot \nabla_\mathrm{x} \Phi\left(\mathrm{x}\right) = \frac{1}{2} D\left(\lambda\right) \cdot \nabla_\mathrm{x} \log \pi\left(\mathrm{x}\right) \\
&= \frac{1}{2} D\left(\lambda\right) \cdot \left[\nabla_\mathrm{x} \log p\left(\mathrm{y}_k|\mathrm{x}\right) + \nabla_\mathrm{x} \log p\left(\mathrm{x}|\mathrm{y}_{1:k-1}\right)\right].
\end{aligned}
\tag{4.14}
$$

Substituting (4.14) into (4.7), it is easy to see that the probability current vanishes as $\lambda \to \infty$. Additionally, it is important to note that continuous multivariate probability densities commonly used in parametric statistics (e.g., Gaussian, Student's t, Mises-Fisher, Pareto of first kind, Cauchy etc) satisfy the potential conditions (4.10) that suffice for $\Phi\left(\mathrm{x}\right)$ to exist.

Based on equation (4.1) and on the drift obtained from the stationary solution (4.14), the dynamics of a set of particles $\left\{\mathrm{x}^{(i)}\left(\lambda\right) : i = 1, \ldots, N\right\}$ can be described by the stochastic differential equation

$$dx = \mu\left(\mathrm{x}, \lambda\right) d\lambda + \sigma\left(\mathrm{x}, \lambda\right) d\mathrm{w}_\lambda, \quad \mathrm{x}_0^{(i)} = \mathrm{x}_{k-1}^{(i)};$$
$$dx = \frac{1}{2} D\nabla_\mathrm{x} \log \pi\left(\mathrm{x}\right) d\lambda + D^{1/2} d\mathrm{w}_\lambda; \tag{4.15}$$

where $\pi\left(\mathrm{x}\right)$ is the target (posterior) probability density, $\{\mathrm{w}_\lambda\}$ is the standard Wiener process (Brownian motion) and D is the diffusion matrix. The stochastic process described by (4.15) is known in the literature to follow the Langevin dynamics and, except for few special cases, cannot be exactly simulated. Thus, the most common way to ensure the simulation provides samples from the correct target distribution, $\pi\left(\mathrm{x}\right)$, is to set the discretized dynamics as a proposal within a Markov chain Monte Carlo framework, which leads to the Metropolis-adjusted Langevin algorithm (MALA) [9].

By defining the distribution at λ and target distribution as $d\mathcal{P} = p\left(\mathrm{x}, \lambda\right) d\mathrm{x}$ and $d\mathcal{P}_\pi = \pi\left(\mathrm{x}\right) d\mathrm{x}$ respectively, one can articulate the total-variation distance between the probability measures $\mathcal{P}\left(d\mathrm{x}\right)$ and $\mathcal{P}_\pi\left(d\mathrm{x}\right)$ defined on $(\mathbb{R}^{n_x}, \mathscr{B}(\mathbb{R}^{n_x}))^6$ as $\|\mathcal{P} - \mathcal{P}_\pi\|_{\mathrm{TV}}$. It can be shown that (e.g., [64]), if the SDE (4.15) is integrated over a sufficiently long (finite) time horizon $T \in \mathbb{R}_+$, then[7]

$$\|\mathcal{P} - \mathcal{P}_\pi\|_{\mathrm{TV}} = \frac{1}{2} \int_{\mathbb{R}^{n_x}} |p\left(\mathrm{x}, \lambda\right) - \pi\left(\mathrm{x}\right)| \, d\mathrm{x} \leq \varepsilon \tag{4.16}$$

for any $\lambda > T$, under a desired precision ε. The implication is that stochastic particle flow implements the filtering mapping (2.5) with increasing accuracy as λ progresses.

Stochastic particle flow can be interpreted as a continuous-time filtering method in the classical sense. Under the abstraction of a continuously interpolated observation process, the method has a direct correspondence to the Kallianpur-Striebel formula and satisfies the Zakai equation as we demonstrate by the *Theorem 17* and *Corollary 19* in the Appendix A.2. Practically speaking, the major difference between stochastic particle flow and other Langevin-based algorithms is the way that the target density is sequentially approximated via local representations that compose a mixture. This will be discussed later in Section 5.

4.4 The Diffusion Matrix

As explained by Girolami & Calderhead [10], the space of parameterized probability density functions is endowed with a natural Riemann geometry, where the diffusion matrix arises as the inverse of a position-specific metric tensor, $G\left(\mathrm{x}\left(\lambda\right)\right)$. This metric tensor maps the distances inscribed in a Riemann manifold to distances in the Euclidean space and, therefore, constitutes a means to constrain the dynamics of any stochastic process to the geometric structure of the parametric probability space. Rao [66] showed the tensor $G\left(\mathrm{x}\left(\lambda\right)\right)$ to be the expected Fisher information matrix

$$G\left(\mathrm{x}\left(\lambda\right)\right) = -\mathbb{E}_{\mathrm{y}|\mathrm{x}} \left[\mathcal{H}_\mathrm{x} \left[\log p\left(\mathrm{y}|\mathrm{x}\right)\right]\right], \tag{4.17}$$

where $\mathcal{H}_\mathrm{x}\left[.\right]$ is the Hessian matrix with respect to x. In a Bayesian context, Girolami & Calderhead [10] suggested a metric tensor that includes the prior information as

$$G\left(\mathrm{x}\left(\lambda\right)\right) = -\mathbb{E}_{\mathrm{y}|\mathrm{x}} \left[\mathcal{H}_\mathrm{x} \left[\log p\left(\mathrm{y}|\mathrm{x}\right)\right]\right] - \mathcal{H}_\mathrm{x} \left[\log p_x\left(\mathrm{x}\right)\right], \tag{4.18}$$

although many possible choices of metric for a specific manifold could be advocated. Because we are interested in local (curvature) properties of stochastic particle flow, a sensible choice for the metric tensor $G\left(\mathrm{x}\left(\lambda\right)\right)$ is the observed Fisher information matrix incorporating the prior information. In this case, the diffusion matrix becomes

$$D = G\left(\mathrm{x}\left(\lambda\right)\right)^{-1} = \left[-\mathcal{H}_\mathrm{x} \left[\log \pi\left(\mathrm{x}\right)\right]\right]_{\mathrm{x}=\mathrm{x}_\lambda}^{-1}, \tag{4.19}$$

where the Hessian matrix is locally evaluated at $\mathrm{x} = \mathrm{x}_\lambda^{(i)}$ for the ith sample, and the resulting diffusion matrix is kept constant for each integration step to obtain the subsequent sample. A problem with this choice is that

---

[6]$\mathscr{B}(\mathbb{R}^{n_x})$ is the σ-field of Borel sets of $\mathbb{R}^{n_x}$.

[7]The total variation norm for probability measures have an equivalence to the $\mathbb{L}^1$-norm as presented in (4.16). A simple argument for this equivalence is given in [65], chapter 4, proposition 4.2.

the expression (4.19) may not be strictly positive definite at specific points of the state space for some types of probability distributions (e.g., mixtures). In order to solve that problem, one could appeal to methods for regularizing the diffusion matrix such as the Tikhonov regularization, the technique to find the nearest (in terms of minimum Fröbenius norm) positive definite matrix [67], or *SoftAbs* [68], a technique that implements a smooth absolute transformation of the eigenvalues to map the negative-Hessian metric into a positive-definite matrix. Another possibility is adopting an empirical estimate to (4.18).

4.5 Integration Method

Among the discretization methods that could be used to integrate the SDE (4.15), we advocate the use of Ozaki's discretization [69] of the Langevin diffusion. This is more accurate than methods based on the Euler discretization. Ozaki's discretization is only possible for target densities that are continuously differentiable and have a smooth Hessian matrix. These requirements may be fulfilled by a solution that constitutes a superposition of conveniently parameterized local approximations to a density.

The algorithm that enables simulation from the SDE (4.15) using Ozaki's discretization is generally called *Langevin Monte Carlo with Ozaki discretization* (LMCO) in the MCMC community (see [64]). Provided an appropriate time horizon, T, by discretizing the interval $0 \leq \lambda \leq T$ into L sub-intervals $\{\lambda_0 = 0, \lambda_1, \ldots \lambda_l, \ldots, \lambda_L = T\}$, the discretized particle flow equation using Ozaki's method is given by

$$
\begin{aligned}
\mathrm{x}(\lambda_{l+1}) = \mathrm{x}(\lambda_l) &+ \left(\mathbb{I}_{n_x} - e^{-\frac{\Delta\lambda}{2}D(\lambda_l)^{-1}} \right) D(\lambda_l)^2 \nabla_{\mathrm{x}} \log \pi\left(\mathrm{x}(\lambda_l)\right) \\
&+ \left[\left(\mathbb{I}_{n_x} - e^{-\Delta\lambda D(\lambda_l)^{-1}} \right) D(\lambda_l)^2 \right]^{1/2} \mathrm{w}_{l+1},
\end{aligned}
\tag{4.20}
$$

where $\{\mathrm{w}_l : l = 1, \ldots, L\}$ is a sequence of independent random vectors distributed according to $\mathrm{w}_l \sim \mathcal{N}(\mathrm{w}; 0_{n_x}, \mathbb{I}_{n_x})$. The need to compute the exponential matrices in (4.20) implies an increment in complexity typically bounded by $\mathcal{O}\left(NLn_x^3\right)$ computations, which may not be justifiable for some applications. A cheaper alternative is achieved by linearizing (4.15) in the neighbourhood of the current state, assuming $D(\lambda)$ piecewise constant in pseudo-time, transforming the linearized equation by the Laplace transform, solving it in the Laplace domain, and transforming it back. The result is

$$
\begin{aligned}
\mathrm{x}(\lambda_{l+1}) = \mathrm{x}(\lambda_l) &+ \left(1 - e^{-\frac{1}{2}\Delta\lambda} \right) D(\lambda_l) \nabla_{\mathrm{x}} \log \pi\left(\mathrm{x}(\lambda_l)\right) \\
&+ \left(1 - e^{-\Delta\lambda} \right)^{1/2} D(\lambda_l)^{1/2} \mathrm{w}_{l+1},
\end{aligned}
\tag{4.21}
$$

where the exponential matrices are avoided but the exponential effect on the integration variable (time step) is kept. See Appendix B for the derivation of this latter integration rule.

It is important to note that, upon integration of the SDE (4.15) by a numerical method, convergence to the invariant distribution is no longer guaranteed for any finite step size. This is due to the first-order integration error that is introduced. When tackling difficult nonlinear filtering problems where the integration error becomes significant, a correction can be carried out by employing a Metropolis acceptance step after each integration step to ensure convergence to the invariant measure.

4.6 Selection of Time Horizon and Integration Step Size

To successfully implement stochastic particle flow, it is necessary to select an adequate time horizon, T, and an integration step size, $\Delta\lambda$. These parameters need to be chosen such that stationarity is reached and convergence to the invariant measure is achieved. There are several routes one could take to solve this problem with each making different assumptions about the probability measures involved and about the regularity properties of the stationary distribution. One could also pose related questions in the context of specific implementations. Answering such questions might, for example, involve selecting the time horizon and integration step size that minimizes the variance of the samples' weights.

The view adopted here is that, since computational effort is a fundamental issue for implementing stochastic particle flow, we should choose these parameters to minimize computational effort. More specifically, we want to minimize the number of integration steps that need to be performed to achieve

$$
\|\mathcal{P}_{\tilde{\mathcal{L}}[\Delta\lambda],T} - \mathcal{P}_{\pi}\|_{\mathrm{TV}} \leq \varepsilon
\tag{4.22}
$$

for an acceptable precision level ε, where $\mathcal{P}_{\tilde{\mathcal{L}}[\Delta\lambda],T}(d\mathrm{x})$ is the approximating probability measure achieved by sampling from the discretized Langevin stochastic process over $L = \lceil T/\Delta\lambda \rceil$ steps.

Defining near-optimal choices of these parameters for general target measures, including mixtures and highly skewed distributions, would require a more thorough study that extrapolates the scope of this work. Instead, in this paper, we propose two pragmatic approaches to choosing both the time horizon and integration step size.

Approach 1

Our first approach builds on results concerning the scaling of Langevin-based MCMC algorithms: the interested reader is referred to Roberts & Rosenthal [70] and a recent extension by Pillai *et al.* [71] that treat high-dimensional target measures that are not of the product form. In summary, these analyses show that the number of steps required to sample the target measure by the Metropolis-Adjusted Langevin algorithm (MALA) grows as $L \sim \mathcal{O}(n_x^{1/3})$. In addition, these papers work out the optimal step size by maximizing the "speed function" or, equivalently, producing the optimal average acceptance rate. Although this optimal criterion is only applicable to algorithms that employ a Metropolis acceptance step, tuning the step size used in stochastic particle flow for an "emulated" (hypothetical) acceptance rate of interest can guide the rate of convergence (even if the accept-reject step is suppressed in practice). Abusing the methodology presented by Pillai *et al.* [71] and using *Proposition 2.4* from Roberts *et al.* [72], let us denote the asymptotic acceptance probability $\alpha(l)$ as a function of a scaling parameter $l \in \mathbb{R}$, such that the speed function $h(l)$ for high-dimensional MALA can be approximated as [71]

$$
\begin{aligned}
h\left(l\right) &= l \cdot \alpha\left(l\right) \approx l \cdot \mathbb{E}_\pi \left[1 \wedge e^{\mathcal{N}\left(-l^3/4, l^3/2\right)}\right] \\
&= l \cdot \left[\mathcal{N}_{\mathrm{cdf}}\left(\frac{-l^3/4}{\sqrt{l^3/2}}\right) + \exp\left(-\frac{l^3}{4} + \frac{\left(\sqrt{l^3/2}\right)^2}{2}\right)\mathcal{N}_{\mathrm{cdf}}\left(-\sqrt{l^3/2} - \frac{-l^3/4}{\sqrt{l^3/2}}\right)\right] \\
&= l \cdot 2\mathcal{N}_{\mathrm{cdf}}\left(-\sqrt{l^3/8}\right),
\end{aligned}
\tag{4.23}
$$

where $\mathcal{N}_{\mathrm{cdf}}(.)$ is a standard normal cumulative distribution function (cdf). One can observe the maximum occurs at $l_{opt} \approx 1.3620$ and corresponds to the theoretical optimal acceptance rate, $\alpha(l_{opt}) \approx 0.5741$.

Our approach is then based on the notion that, for a conveniently selected acceptance rate, the step size should scale as $\Delta\lambda \propto l \cdot n_x^{-\xi}$ such that the total number of steps is scaled as $L \propto n_x^\xi$, and the exponent ξ depends on whether the Metropolis adjustment is used or not. Theoretically, if the accept-reject step is not adopted then[8] $\xi = 1$, otherwise the optimal choice is $\xi = 1/3$ [71].

To exemplify the use of this approach, to produce an emulated (asymptotic) acceptance rate of $\alpha = 0.80$, a stochastic particle flow should be scaled as $l \approx 2\left[-\mathcal{N}_{\mathrm{cdf}}^{-1}\left(\alpha/2\right)\right]^{2/3} = 0.8008$. Using the Langevin dynamics considered herein, if $n_x = 10$ and $\xi = 1$ then $\Delta\lambda = 2l \cdot n_x^{-\xi} \approx 0.1602$.

If the accept-reject step is present, a criterion to stop the simulation could be established online. Various MCMC convergence diagnostic methods are applicable to this task. However, in our experience with stochastic particle flow, such approaches to online determination of convergence may indicate more steps are needed than are actually necessary to obtain good results, and so give a pessimistic view of the amount of computation required. To set the time horizon when accept-reject step is not used, we note that $T = T_0 + L_s \cdot \Delta\lambda$, where T_0 is the time required to take the chain to the region of high acceptance probability ("warm-up") and L_s is the number of steps to explore the invariant measure. We determine both L_s and T_0 either by presetting reasonable values, or based on the second approach to be explained next.

Approach 2

Our second approach is an extension of the method proposed by Dalalyan [64]. It is useful due to its ease of application and suitability to "nicely" measurable filtering quantities although, strictly speaking, the method is only applicable to target densities that are log-concave. The criteria are presented as follows.

Theorem 1. *Let* $\Phi : \mathbb{R}^{n_x} \to \mathbb{R}$ *be a measurable convex function satisfying*

$$
\int_{\mathbb{R}^{n_x}} \exp\{-\Phi\left(\mathrm{x}\right)\} < \infty,
\tag{4.24}
$$

$$
\Phi\left(\mathrm{x}\right) - \Phi\left(\bar{\mathrm{x}}\right) - \nabla_\mathrm{x}\Phi\left(\bar{\mathrm{x}}\right)^T\left(\mathrm{x} - \bar{\mathrm{x}}\right) \geq \frac{1}{2}m\left\|\mathrm{x} - \bar{\mathrm{x}}\right\|_2^2,
\tag{4.25}
$$

$$
\left\|\nabla_\mathrm{x}\Phi\left(\mathrm{x}\right) - \nabla_\mathrm{x}\Phi\left(\bar{\mathrm{x}}\right)\right\|_2 \leq M\left\|\mathrm{x} - \bar{\mathrm{x}}\right\|_2, \quad \forall\mathrm{x}, \bar{\mathrm{x}} \in \mathbb{R}^{n_x},
\tag{4.26}
$$

---

[8]In view of (4.28) and (4.29), $T/\Delta\lambda \sim \mathcal{O}(n_x)$.

for two existing positive constants m and M. Let $\bar{\mathrm{x}} \in \mathbb{R}^{n_x}$ be the global minimum of $\Phi(\mathrm{x})$. Suppose a discrete-time Langevin Monte Carlo algorithm integrates (4.15), targeting the invariant density $\pi(\mathrm{x}) \propto \exp\{-\Phi(\mathrm{x})\}$ with measure $\mathcal{P}_\pi(d\mathrm{x})$, and with the initial density $\nu(\mathrm{x}) = \delta(\mathrm{x} - \mathrm{x}_\nu)$ (a probability mass initially located at $\mathrm{x} = \mathrm{x}_\nu$). In addition, assume that for some $\gamma \geq 1$ we have $\Delta\lambda \leq (\gamma M)^{-1}$, and $K = \sup_{\mathrm{x}} \|D(\mathrm{x})\|_2$ where $D_\lambda = D(\mathrm{x}_\lambda)$ is the diffusion matrix. Then, for a time horizon, T, and step size, $\Delta\lambda$, the total-variation distance between the target measure $\mathcal{P}_\pi$ and the approximated measure $\mathcal{P}_{\tilde{\mathcal{L}}(\Delta\lambda),T}$ furnished by the discrete-time Langevin Monte Carlo algorithm satisfies

$$\|\mathcal{P}_{\tilde{\mathcal{L}}[\Delta\lambda],T} - \mathcal{P}_\pi\|_{TV} \leq \frac{1}{2} \exp\left\{-\frac{1}{2}mT + \frac{n_x}{2}\log\left(\frac{M}{m}\right) - \log\left[\Gamma_u\left(\frac{n_x}{2}, \frac{M\|\bar{\mathrm{x}} - \mathrm{x}_\nu\|_2^2}{2}\right)\right]\right\}$$
$$+ \frac{1}{2} - \frac{1}{2}\exp\left\{-\frac{n_x}{2}\frac{M^3 K^4 \gamma}{48(2\gamma - 1)}\left(\frac{1}{n_x}\|\bar{\mathrm{x}} - \mathrm{x}_\nu\|_2^2 + 2T\right)\Delta\lambda^2 - \frac{n_x M^2 K^3 T}{16}\Delta\lambda\right\}, \quad (4.27)$$

where $\Gamma_u(s, x) \triangleq \Gamma(s)^{-1} \int_x^\infty t^{s-1} e^{-t} dt$ is the upper incomplete gamma function.

Corollary 2. *Let $n_x \geq 2$, Φ satisfy (4.24), (4.25) and (4.26), and $\varepsilon \in (0, 1/2)$ be a desired precision level. Let the time horizon T and the step size $\Delta\lambda$ be defined by*

$$T \geq \frac{2\log(1/\varepsilon) + n_x \log\left(\frac{M}{m}\right) - 2\log\left[\Gamma_u\left(\frac{n_x}{2}, \frac{M\|\bar{\mathrm{x}} - \mathrm{x}_\nu\|_2^2}{2}\right)\right]}{m}, \quad (4.28)$$

$$\Delta\lambda \leq \frac{-\frac{T}{16} + \sqrt{\left(\frac{T}{16}\right)^2 + \frac{\gamma}{48(2\gamma - 1)}\left(\frac{1}{n_x}\|\bar{\mathrm{x}} - \mathrm{x}_\nu\|_2^2 + 2T\right)M^{-1}K^{-2}\left[\frac{2}{n_x}\log\left(\frac{1}{1-\varepsilon}\right)\right]}}{\frac{\gamma}{48(2\gamma - 1)}\left(\frac{1}{n_x}\|\bar{\mathrm{x}} - \mathrm{x}_\nu\|_2^2 + 2T\right)MK}, \quad (4.29)$$

where $\gamma \geq 1$. Then the resulting probability distribution of a Langevin Monte Carlo algorithm that integrates (4.15) after $L = \lceil T/\Delta\lambda \rceil$ steps, satisfies $\|\mathcal{P}_{\tilde{\mathcal{L}}[\Delta\lambda],T} - \mathcal{P}_\pi\|_{TV} \leq \varepsilon$.

Theorem 1 is thoroughly underpinned by the findings of Dalalyan [73], with specific settings changed to match the Langevin algorithm proposed in this paper, and both a more general bound for the time horizon and a tightened bound for the step size (to reduce computational effort). We recommend the reader interested in the proof to first refer to [73] and then follow the missing arguments for its proof in the Appendix A.1.

Corollary 2 is a direct criterion for selecting the time horizon and step size, arising from the right-hand side of the inequality (4.27) being set to be a desired precision level. It is essential to clarify that some practical issues arise here: in this form, the method holds for log-concave densities only; the positive constants, m and M, are assumed known a priori; and the approximation of the filtering density is not taken into account in the error budget. Rigorously speaking, the method does not apply to more general cases. However, the method has utility as the basis of an approximate (and pragmatic) mechanism for obtaining the required parameters. In making this approximation, we explicitly acknowledge that we are assuming that:

1. The target density can be well characterized by a central tendency statistic, $\bar{\mathrm{x}}_c$, that replaces and roughly represents $\bar{\mathrm{x}}$ in all aspects of the analysis.

2. The initial measure is composed of a superposition of N probability masses described by

$$N^{-1} \sum_{i=1}^N \delta(\mathrm{x} - \mathrm{x}^{(i)})$$

   or, equivalently, an empirical distribution with mean μ_ν and covariance matrix V, spatially encompassing all initial samples (from the previous filtering iteration), which is assumed to constrain the constants M and m by

$$M < \|V\|_2/2, \quad (4.30)$$

$$(\mathrm{x}^{(i)} - \bar{\mathrm{x}}_c)^T \frac{m}{2}(\mathrm{x}^{(i)} - \bar{\mathrm{x}}_c) \geq \chi_{\mathrm{cdf}}^{-2}(0.99, n_x), \quad \forall \mathrm{x}^{(i)} : i = 1, \ldots, N, \quad (4.31)$$

   where $\chi_{\mathrm{cdf}}^{-2}(P, \kappa)$ is the inverse of chi-square cdf for probability P and κ degrees of freedom.

3. In accordance with *Lemma 4* in [64], the constant M is also constrained by

$$\Phi(\mathrm{x}^{(i)}) - \Phi(\bar{\mathrm{x}}_c) - \nabla_\mathrm{x}\Phi(\bar{\mathrm{x}}_c)^T(\mathrm{x}^{(i)} - \bar{\mathrm{x}}_c) \leq \frac{M}{2}\|\mathrm{x} - \bar{\mathrm{x}}_c\|_2^2, \quad \forall \mathrm{x}^{(i)} : i = 1, \ldots, N. \quad (4.32)$$

4. The positive constants M and m can be roughly estimated by

   (a) approximating the statistic $\bar{\mathrm{x}}_c$ of the target density (e.g., obtaining a maximum-a-posteriori estimate by optimization or an approximated mean by the EKF),

   (b) inverting conditions (4.25) and (4.26), and incorporating the constraint (4.32), to give

$$\tilde{M} = 2 \sup_{i \in [1,N]} \max \left[\frac{\left\| \nabla_{\mathrm{x}} \Phi(\mathrm{x}^{(i)}) - \nabla_{\mathrm{x}} \Phi(\bar{\mathrm{x}}_c) \right\|_2}{2 \left\| \mathrm{x}^{(i)} - \bar{\mathrm{x}}_c \right\|_2}, \frac{\Phi(\mathrm{x}^{(i)}) - \Phi(\bar{\mathrm{x}}_c) - \nabla_{\mathrm{x}} \Phi(\bar{\mathrm{x}}_c)^T (\mathrm{x}^{(i)} - \bar{\mathrm{x}}_c)}{\left\| \mathrm{x}^{(i)} - \bar{\mathrm{x}}_c \right\|_2^2} \right], \quad (4.33)$$

$$\tilde{m} = 2 \inf_{i \in [1,N]} \frac{\Phi(\mathrm{x}^{(i)}) - \Phi(\bar{\mathrm{x}}_c) - \nabla_{\mathrm{x}} \Phi(\bar{\mathrm{x}}_c)^T (\mathrm{x}^{(i)} - \bar{\mathrm{x}}_c)}{\left\| \mathrm{x}^{(i)} - \bar{\mathrm{x}}_c \right\|_2^2}, \quad (4.34)$$

where all quantities can be computed from definition $\Phi(\mathrm{x}) \triangleq -\log p(\mathrm{y}_k|\mathrm{x}) - \log p(\mathrm{x}|\mathrm{y}_{1:k-1})$ given an approximation to the prior pdf; the resulting values of $\tilde{M}$ and $\tilde{m}$ must also satisfy (4.30) and (4.31).

Once the positive constants M and m have been estimated, obtaining T and $\Delta\lambda$ follows from (4.28) and (4.29) respectively. In our experience, for very simple problems, (4.28) may produce overestimated time horizons and, as a consequence, may cause (4.29) to produce underestimated step sizes for stochastic particle flow. This happens because the bound for the time horizon becomes loose, in view of *Lemma 7*, for initial distributions that are far from the target distribution. In simple cases, a closer approximation can by achieved by assuming 1-uniform ergodicity of the Markov chain to give

$$T \geq \frac{2\log(1/\varepsilon) + n_x \log R}{\tilde{m}}, \quad (4.35)$$

for a finite $R \in \mathbb{R}_+$, at the cost of having to determine R empirically. Similarly, for simple cases, expression (4.29) can be replaced with an empirical rule of the form

$$\Delta\lambda \leq 2\frac{\sqrt{\tilde{m}}}{\tilde{M}}, \quad (4.36)$$

which satisfies $\Delta\lambda \leq \left(\gamma\tilde{M}\right)^{-1}$ for $\gamma \leq \left(2\sqrt{\tilde{m}}\right)^{-1}$ as required by *Theorem 1*, but is not guaranteed to satisfy *Corollary 2*.

5 Stochastic Particle Flow as a Gaussian Sum Filter

In this section we use stochastic particle flow to derive a filter that approximates the posterior probability density as a Gaussian mixture. We refer to the resulting filter as the stochastic particle flow Gaussian sum filter (SPF-GS).

5.1 The Mixture-Based Approximating Measure

Given a set of samples drawn from an importance distribution, $\left\{ \mathrm{x}^{(i)} \in \mathcal{X} : i \in 1, \ldots, N \right\}$, if one is required to solve the filtering problem by a standard Monte Carlo method, then the stochastic filter adopts the following approximation

$$\begin{aligned}
\hat{\varphi} &= \int_{\mathcal{X}} \varphi(\mathrm{x}) \pi(\mathrm{x}) \, dx \\
&\approx \int_{\mathcal{X}} \varphi(\mathrm{x}) \sum_{i=1}^{N} w(\mathrm{x}^{(i)}) \delta(\mathrm{x} - \mathrm{x}^{(i)}) dx \\
&= \sum_{i=1}^{N} w^{(i)} \varphi(\mathrm{x}^{(i)}),
\end{aligned} \quad (5.1)$$

where $w(\mathrm{x}^{(i)}) = w^{(i)}$ are the importance weights. Now suppose that we have access to an approximating measure $\tilde{\mathcal{P}}_{\pi}(dx)$ on $(\mathcal{X}, \mathscr{B}(\mathcal{X}))$ with an associated density such that $d\tilde{\mathcal{P}}_{\pi} = \tilde{\pi} dx$. If the density $\tilde{\pi}$ involves a mixture of

N Gaussians according to

$$\tilde{\pi}(\mathrm{x}) = \sum_{i=1}^{N} w_m^{(i)} \mathcal{N}(\mathrm{x}; \mu_m(\mathrm{x}^{(i)}), \Sigma_m(\mathrm{x}^{(i)}))$$

$$= \sum_{i=1}^{N} w_m^{(i)} \mathcal{N}(\mathrm{x}; \mu_m^{(i)}, \Sigma_m^{(i)}), \tag{5.2}$$

where $\{w_m^{(i)}, \mu_m^{(i)}, \Sigma_m^{(i)}\}$ are computed based on the samples $\{\mathrm{x}^{(i)}\}$, then the solution is given by

$$\hat{\varphi} = \int_{\mathcal{X}} \varphi(\mathrm{x}) \pi(\mathrm{x}) \, d\mathrm{x}$$

$$\approx \int_{\mathcal{X}} \varphi(\mathrm{x}) \sum_{i=1}^{N} w_m^{(i)} \mathcal{N}(\mathrm{x}; \mu_m^{(i)}, \Sigma_m^{(i)}) d\mathrm{x}$$

$$= \sum_{i=1}^{N} w_m^{(i)} \int_{\mathcal{X}} \varphi(\mathrm{x}) \mathcal{N}(\mathrm{x}; \mu_m^{(i)}, \Sigma_m^{(i)}) d\mathrm{x}$$

$$= \sum_{i=1}^{N} w_m^{(i)} \mathbb{E}_{\mathcal{N}} \left[\varphi(\mathrm{x}) \mid \mu_m^{(i)}, \Sigma_m^{(i)} \right]. \tag{5.3}$$

In this setting, it is possible to prove that $\tilde{\pi}(\mathrm{x}) \to \pi(\mathrm{x})$ as $N \to \infty$ almost surely if $\mu_m^{(i)} \to \mathrm{x}^{(i)}$ and $\Sigma_m^{(i)} \to 0$, by appealing to convergence proofs for mixture-based estimators (see [74], pages 197–199). Also, it is worth noting that this procedure is quite general in the sense that (5.2) could be replaced by a mixture of any convenient parametric distribution.

5.2   The Stochastic-Particle-Flow Gaussian Sum Filter

Building upon the results presented in the previous section, SPF-GS uses samples to propagate local Gaussian components, which can together provide an accurate approximation to the posterior probability density of the form (5.2). More specifically, given a new measurement and a set of samples and moments $\{\mathrm{x}_{k-1}^{(i)}, \mu_{m,k-1}^{(i)}, \Sigma_{m,k-1}^{(i)} : i = 1, \ldots, N\}$, the filtering procedure consists of integrating the SDE (4.15) for each particle $\mathrm{x}^{(i)}(\lambda)$ and propagating its associated moments $(\mu_m^{(i)}(\lambda), \Sigma_m^{(i)}(\lambda))$ through the interval $0 < \lambda \leq T$, which corresponds to the interval $t_{k-1} < t \leq t_k$. The integration process is performed until one achieves the posterior set of samples and parameters $\{\mathrm{x}_k^{(i)}, \mu_{m,k}^{(i)}, \Sigma_{m,k}^{(i)}\} := \{\mathrm{x}^{(i)}(T), \mu_m^{(i)}(T), \Sigma_m^{(i)}(T)\}$, where $\mathrm{x}^{(i)}(0) = \mathrm{x}_{k-1}^{(i)}$, so that

$$\frac{1}{2} \int_{\mathcal{X}} \left| \sum_{i=1}^{N} \frac{1}{N} \mathcal{N}(\mathrm{x}; \mu_m^{(i)}(T), \Sigma_m^{(i)}(T)) - \pi(\mathrm{x}) \right| d\mathrm{x} \leq \varepsilon',$$

under a desired precision level ε' as $N \to \infty$. In practice, the integration of stochastic particle flow (4.15) over $0 < \lambda \leq T$ involves multiple intermediate sampling steps that evolve the samples $\mathrm{x}^{(i)}(\lambda)$ to populate the local regions of the state space where the target distribution will be described. At the same time, the mixture components are propagated to define the local approximations and thereby the global approximation to the filtering density. A fundamental aspect of the procedure proposed herein is that, for each filtering step, the intermediate moments $(\mu_m^{(i)}(\lambda), \Sigma_m^{(i)}(\lambda))$ are initialized as departing from the corresponding samples obtained from the previous step, i.e., $\mu_m^{(i)}(0) := \mathrm{x}_{k-1}^{(i)}$ and $\Sigma_m^{(i)}(0) := 0_{n_x \times n_x}$, and evolved onto the local posterior moments, $\mu_m^{(i)}(T)$ and $\Sigma_m^{(i)}(T)$. In this setting, each component of the filtering mixture is associated, via Fokker-Planck equation (Langevin dynamics), with the mapping

$$\frac{1}{N} \delta(\mathrm{x}_{k-1} - \mathrm{x}_{k-1}^{(i)}) \mapsto \frac{1}{N} \mathcal{N}(\mathrm{x}_k; \mu_{m,k}^{(i)}, \Sigma_{m,k}^{(i)}). \tag{5.4}$$

It is also important to highlight that we perceive the most informative approximation of the posterior from the previous iteration as provided by the set of moments from the previous iteration, $\{\mu_{m,k-1}^{(i)}, \Sigma_{m,k-1}^{(i)}\}$, and not by the samples. We therefore use the previous iteration's mixture to define the invariant target measure but samples from this mixture to integrate the SDE. In our experience, this setting is beneficial because: it avoids

the bias that would result from propagating a mixture-only approximation from one filtering iteration to the next; it dismisses the need to explicitly compute mixture weights since the ergodic Markov chain that carries out (5.4) is known to converge to the invariant measure no matter where it starts [75].

More specifically, in the Langevin diffusion setting presented in Section 4, the mixture measure forms the basis of approximating the target log-density. Each mixture component from the previous filtering step enables a local approximation of the prior density as

$$\tilde{p}^{(i)}\left(\mathrm{x}_k'|\mathrm{y}_{1:k-1}\right) = \int_{\mathbb{R}^{n_x}} p_t(\mathrm{x}_k'|\mathrm{x}_{k-1})\mathcal{N}(\mathrm{x}_{k-1}; \mu_{m,k-1}^{(i)}, \Sigma_{m,k-1}^{(i)})d\mathrm{x}_{k-1}, \tag{5.5}$$

where $p_t\left(\mathrm{x}_k'|\mathrm{x}_{k-1}\right)$ is the state-process transition kernel, and the resulting prior density is locally approximated as a Gaussian (e.g., via the Unscented Transform). Thus, provided a known likelihood function, $p\left(\mathrm{y}_k|\mathrm{x}\right)$, one Langevin transition kernel is computed per sample based on $\nabla_\mathrm{x} \log \tilde{\pi}^{(i)}(\mathrm{x}) = \nabla_\mathrm{x} \log p\left(\mathrm{y}_k|\mathrm{x}\right) + \nabla_\mathrm{x} \log \tilde{p}^{(i)}\left(\mathrm{x}|\mathrm{y}_{1:k-1}\right)$.

We propose an approximate method to propagate the moments, $\mu_m(\lambda)$ and $\Sigma_m(\lambda)$, by linearizing the flow locally, in the neighbourhood of a probability mass located at x_l. In the Appendix C we provide an argument (which we regard as useful, if not rigorous) to justify why this local flow approximation should produce acceptable errors on the propagated moments. The procedure produces a negligible error for a small state displacement given a small increment of pseudo-time, $\Delta\lambda$, so that stochastic particle flow (4.15) can be approximated within the region $\|\mathrm{x} - \mathrm{x}_l\| < \zeta$, for a sufficiently small $\zeta \in \mathbb{R}_+$, as

$$d\mathrm{x} = \frac{1}{2}D(\lambda)\nabla_\mathrm{x}\log\pi\left(\mathrm{x}\right)d\lambda + D(\lambda)^{1/2}d\mathrm{w}_\lambda, \quad \lambda \in (\lambda_l, \lambda_l + \Delta\lambda], \mathrm{x}(\lambda_l) = \mathrm{x}_l;$$
$$d\mathrm{x} \approx \left[C(\mathrm{x}_l, \lambda) \cdot \mathrm{x} + c(\mathrm{x}_l, \lambda)\right]d\lambda + D(\lambda)^{1/2}d\mathrm{w}_\lambda. \tag{5.6}$$

As a consequence of integrating the flow, the corresponding component moments are evolved according to the locally approximated ordinary differential equations (Appendix C)

$$\frac{d\mu_m^{(i)}(\lambda)}{d\lambda} = C(\mathrm{x}_l^{(i)})\mu_m^{(i)}(\lambda) + c(\mathrm{x}_l^{(i)}), \tag{5.7}$$

$$\frac{d\Sigma_m^{(i)}(\lambda)}{d\lambda} = C(\mathrm{x}_l^{(i)})\Sigma_m^{(i)}(\lambda) + \Sigma_m^{(i)}(\lambda)C^T(\mathrm{x}_l^{(i)}) + D^{(i)}. \tag{5.8}$$

For nonlinear Gaussian problems, the locally approximated flow implies that

$$C\left(\mathrm{x}_l, \lambda\right) = -\frac{1}{2}D(\lambda)P_{k|k-1}^{-1}$$
$$-\frac{1}{2}D(\lambda)\mathcal{J}_\mathrm{x}\left[h(\mathrm{x}_l)\right]^T R_k^{-1}\mathcal{J}_\mathrm{x}\left[h(\mathrm{x}_l)\right], \tag{5.9}$$

$$c\left(\mathrm{x}_l, \lambda\right) = \frac{1}{2}D(\lambda)P_{k|k-1}^{-1}f(\mu_{m,k-1})$$
$$+\frac{1}{2}D(\lambda)\mathcal{J}_\mathrm{x}\left[h(\mathrm{x}_l)\right]^T R_k^{-1}\mathcal{J}_\mathrm{x}\left[h(\mathrm{x}_l)\right]\cdot\mathrm{x}_l$$
$$+\frac{1}{2}D(\lambda)\mathcal{J}_\mathrm{x}\left[h(\mathrm{x}_l)\right]^T R_k^{-1}\left(\mathrm{y}_k - h(\mathrm{x}_l)\right), \tag{5.10}$$

where $\mathcal{J}_\mathrm{x}\left[\cdot\right]$ is the Jacobian matrix with respect to the state, $f(\cdot)$ is the state process function, $h(\cdot)$ is the observation function, and where

$$P_{k|k-1} = \mathbb{E}\left[(\mathrm{x}_{k|k-1} - f(\mu_{m,k-1}))(\mathrm{x}_{k|k-1} - f(\mu_{m,k-1}))^T\right], \tag{5.11}$$

$$R_k = \mathbb{E}\left[(\mathrm{y}_k - h(\mathrm{x}_k))(\mathrm{y}_k - h(\mathrm{x}_k))^T\right], \tag{5.12}$$

are, respectively, the covariance matrix of the prior probability density and the covariance matrix of the observation noise.

Notice that the resulting algorithm is somewhat similar to the Kalman-Bucy filter, in that it avoids an explicit discrete-time measurement update. One could also interpret the SPF-GS as a Monte-Carlo and continuous-time version of the original Gaussian sum filter [76, 77], albeit modified to explore the Riemannian geometric structure of the probability space. However, in contrast to the Gaussian sum filter (and the particle filter), the structure of the SPF-GS removes the need to explicitly compute mixture weights: it relies on multiple independent (ergodic) Markov chains that start at the samples to produce equally weighted mixture components, each describing a local version of the underlying geometric structure of the posterior measure.

The SPF-GS has features that are apparently similar to those of the Gaussian sum particle filter [78], however, in reality, these filters rely on distinct fundamental principles that make them very different. The principle of the Gaussian sum particle filter is using importance sampling to estimate the moments of a mixture's components for approximating a target density. In contrast, the SPF-GS evolves a mixture through multiple intermediate steps by exploring the local properties of a stochastic flow in order to translate probability masses from the previous iteration to a mixture on the posterior probability space. The SPF-GS is also very different from that proposed by Terejanu *et al.* [79], which is a Gaussian sum filter analogous to an extended Kalman-Bucy filter, but providing an estimate of the predicted mixture weights based on an optimization procedure.

The stochastic particle flow Gaussian-sum filter (SPF-GS) is summarized in *Algorithm 1*. The algorithm is expressed in a quite general form, including an accept-reject step that enforces theoretical convergence to the invariant target measure. Our empirical experience indicates that having this step was often not necessary (at least in the cases we have considered). As has already been touched on, if computational efficiency is a primary goal, the SPF-GS should avoid the accept-reject step: calculating the acceptance probability requires double[9] the number of computations of the gradients and inverted Hessian matrices. In addition, using this step requires the explicit evaluation of the kernels and target densities themselves. When the step is removed, there may be the need to deal with outliers, which can significantly affect the estimates because all components in the propagated mixture are treated as equally important. In all numerical examples studied in this paper, none adopted the Metropolis adjustment. For the high-dimensional examples, the (occasional) outliers are removed online by an empirical test to identify the samples that fall outside a credible inference region.

6   Stochastic Particle Flow as a Marginal Particle Filter

In this section we derive a marginal particle filter whose proposal density is built upon a Gaussian mixture obtained via stochastic particle flow. The resulting filter is referred to as the stochastic-particle-flow marginal particle filter (SPF-MPF).

6.1   Marginal Particle Filtering

In the standard setting, particle filters don't target the marginal filtering distribution $p\left(\mathrm{x}_k|\mathrm{y}_{1:k}\right)$, a characteristic inherited from the first particle filters, which were designed to be relatively simple to implement. The main problem with the standard particle filters arises because they construct importance densities that target a joint filtering density $p\left(\mathrm{x}_{0:k}|\mathrm{y}_{1:k}\right)$. A typical particle filter incrementally draws *path samples,* $\{\mathrm{x}_{0:k}^{(i)} \in \mathcal{X}^{k+1} : i = 1, \ldots, N\}$, from a joint importance density $q\left(\mathrm{x}_{0:k}|\mathrm{y}_{1:k}\right)$, and ignores the past of the sampled paths ($\{\mathrm{x}_{0:k-1}^{(i)} \in \mathcal{X}^k\}$) when computing (filtered) expectations of interest. Thus, although these algorithms provide a simple way to perform measurement update, they perform importance sampling in the joint space along all time steps, i.e., in $\mathcal{X}^{k+1} = \mathcal{X}\left(0\right) \times \mathcal{X}\left(1\right) \times \cdots \times \mathcal{X}\left(k\right)$. The result is precipitation of the degeneracy phenomenon: the set of paths become increasingly sparse on the joint space $\mathcal{X}^{k+1}$, leading to a quick increase in the weights' variance while most paths have vanishingly small probability. In high-dimensional applications this problem becomes even more pronounced, rendering the standard particle filters to be practically infeasible.

With the mindset of improving this shortcoming in particle filters, Klaas *et al.* [80] proposed the marginal particle filter. The marginal particle filter targets the marginal posterior distribution $p\left(\mathrm{x}_k|\mathrm{y}_{1:k}\right)$, performing importance sampling on the marginal state space, $\mathcal{X}\left(k\right)$, to produce samples with commensurate sparsity over time. The samples are drawn from an importance density of the form

$$q\left(\mathrm{x}_k|\mathrm{y}_{1:k}\right) \propto \int_{\mathcal{X}} q\left(\mathrm{x}_k|\mathrm{x}_{k-1}, \mathrm{y}_k\right) q\left(\mathrm{x}_{k-1}|\mathrm{y}_{1:k-1}\right) d\mathrm{x}_{k-1}, \tag{6.1}$$

to target the posterior density

$$p\left(\mathrm{x}_k|\mathrm{y}_{1:k}\right) \propto p\left(\mathrm{y}_k|\mathrm{x}_k\right) \int_{\mathcal{X}} p_t\left(\mathrm{x}_k|\mathrm{x}_{k-1}\right) p\left(\mathrm{x}_{k-1}|\mathrm{y}_{1:k-1}\right) d\mathrm{x}_{k-1}, \tag{6.2}$$

with the importance weights

$$w\left(\mathrm{x}_k\right) \propto \frac{p\left(\mathrm{x}_k|\mathrm{y}_{1:k}\right)}{q\left(\mathrm{x}_k|\mathrm{y}_{1:k}\right)}. \tag{6.3}$$

---

[9]Because of the need to construct both the forward and backward transition kernels.

Algorithm 1: Stochastic particle flow - Gaussian sum filter

---

1 **Initialization:**

2 **if** *time $k = 0$* **then**

3 $\quad$ Sample $x_0^{(i)} \sim p(x_0)$, $\forall i = 1, \ldots, N$

4 $\quad$ Set $w_{m,0}^{(i)} := N^{-1}$, $\mu_{m,0}^{(i)} := \mathbb{E}_{p_0}[x_0]$, $\Sigma_{m,0}^{(i)} := \mathbb{E}_{p_0}\left[(x_0 - \bar{x}_0)(x_0 - \bar{x}_0)^T\right]$, $\forall i = 1, \ldots, N$

5 **end**

6 **Steps:**

7 **for** *time $k \geq 1$* **do**

8 $\quad$ Compute the time horizon T and step size $\Delta\lambda$ (section 4.6)

9 $\quad$ Discretize the interval $0 \leq \lambda \leq T$ into L sub-intervals $\{\lambda_0 = 0, \ldots \lambda_l, \ldots, \lambda_L = T\}$

10 $\quad$ Set $x_{l=0}^{(i)} := x_{k-1}^{(i)}$, $\mu_{l=0}^{(i)} := x_{k-1}^{(i)}$, $\Sigma_{l=0}^{(i)} := 0_{n_x \times n_x}$, $\forall i = 1, \ldots, N$

11 $\quad$ **for** $l = 1$ **to** L **do**

12 $\quad\quad$ **for** $i = 1, \ldots, N$ **do**

13 $\quad\quad\quad$ Simulate

$$x_l^{\star(i)} \leftarrow x_{l-1}^{(i)} + \frac{1}{2} \int_{\lambda_{l-1}}^{\lambda_l} D(x_{l-1}^{(i)}) \nabla_x \log \tilde{\pi}(x_{l-1}^{(i)}) d\lambda + \int_{\lambda_{l-1}}^{\lambda_l} D(x_{l-1}^{(i)})^{1/2} dw_\lambda \qquad (5.13)$$

14 $\quad\quad\quad$ Compute the MH acceptance probability $\rho^{(i)} = \min\left[1, \frac{\tilde{\pi}(x_l^{\star(i)})}{q(x_l^{\star(i)}|x_{l-1}^{(i)})} \frac{q(x_{l-1}^{(i)}|x_l^{\star(i)})}{\tilde{\pi}(x_{l-1}^{(i)})}\right]$

$\quad\quad\quad$ `/* We advocate using the Metropolis-adjustment step as practically optional, i.e., only for`
$\quad\quad\quad\quad$ `difficult problems. For many Engineering problems the approximation achieved by suppressing`
$\quad\quad\quad\quad$ `the MH step may be enough and will be more computationally efficient.         */`

15 $\quad\quad\quad$ Simulate $z^{(i)} \sim \mathcal{U}(0,1)$

16 $\quad\quad\quad$ **if** $z^{(i)} \leq \rho^{(i)}$ **then**

17 $\quad\quad\quad\quad$ Set $x_l^{(i)} \leftarrow x_l^{\star(i)}$

18 $\quad\quad\quad\quad$ Propagate

$$\mu_l^{(i)} \leftarrow \mu_{l-1}^{(i)} + \int_{\lambda_{l-1}}^{\lambda_l} \left[C(x_{l-1}^{(i)})\mu_{l-1}^{(i)} + c(x_{l-1}^{(i)})\right] d\lambda, \qquad (5.14)$$

$$\Sigma_l^{(i)} \leftarrow \Sigma_{l-1}^{(i)} + \int_{\lambda_{l-1}}^{\lambda_l} \left[C(x_{l-1}^{(i)})\Sigma_{l-1}^{(i)} + \Sigma_{l-1}^{(i)}C^T(x_{l-1}^{(i)}) + D(x_{l-1}^{(i)})\right] d\lambda \qquad (5.15)$$

19 $\quad\quad\quad$ **else**

20 $\quad\quad\quad\quad$ Set $x_l^{(i)} \leftarrow x_{l-1}^{(i)}$, $\mu_l^{(i)} \leftarrow \mu_{l-1}^{(i)}$, $\Sigma_l^{(i)} \leftarrow \Sigma_{l-1}^{(i)}$

21 $\quad\quad\quad$ **end**

22 $\quad\quad$ **end**

23 $\quad$ **end**

24 $\quad$ Set $x_k^{(i)} := x_{l=L}^{(i)}$, $\mu_{m,k}^{(i)} := \mu_{l=L}^{(i)}$, $\Sigma_{m,k}^{(i)} := \Sigma_{l=L}^{(i)}$, $\forall i = 1, \ldots, N$

$\quad$ **Output:** Approximation of the filtering density as

25

$$\tilde{p}(x_k|y_{1:k}) = \sum_{i=1}^{N} \frac{1}{N} \mathcal{N}(x_k; \mu_{m,k}^{(i)}, \Sigma_{m,k}^{(i)})$$

26 **end**

In practical terms, particles and weights from the previous iteration are used to compose both an approximation of the target density (6.2) and the importance density (6.1), in order to obtain particles and weights for the current iteration. Even though the marginal particle filter is more robust than the standard particle filter against degeneracy, and thereby more suitable to high-dimensional problems in principle, its success is highly dependent on the validity of sequential representations of the target density. Problems may arise in situations where the usual approximation

$$\tilde{p}\left(\mathrm{x}_k | \mathrm{y}_{1:k}\right) \propto p(\mathrm{y}_k | \mathrm{x}_k) \sum_{i=1}^{N} w_{k-1}^{(i)} p_t(\mathrm{x}_k | \mathrm{x}_{k-1}^{(i)}), \tag{6.4}$$

is prone to relevant statistical or numerical errors, e.g., when the transition density $p_t\left(\mathrm{x}_k | \mathrm{x}_{k-1}\right)$ describes a Markov process with small variance and the observation y_k lies relatively far from the current set of particles $\{\mathrm{x}_{k-1}^{(i)}, \mathrm{w}_{k-1}^{(i)}\}$ on the state space (see the linear, univariate example in Section 7). Moreover, owing to the curse of dimensionality, the usual approximation (6.4) is corrupted by a Monte Carlo error that increases geometrically with the number of state dimensions. This may cripple the marginal particle filter in very high-dimensional problems. Because of this limitation, marginal particle filters are likely to perform well only in moderately high-dimensional problems. We illustrate this limitation of marginal particle filters by numerical examples in Section 7.

As well covered in [80], there exist several possibilities to choose the marginal importance density (6.1), among which the auxiliary marginal proposal density is particularly interesting because it emulates an optimal importance density in the sense of minimizing the weights' variance. The marginal optimal (auxiliary) proposal density is usually approximated as

$$\tilde{q}\left(\mathrm{x}_k | \mathrm{y}_{1:k}\right) = \sum_{i=1}^{N} w_{q,k-1}^{(i)} p(\mathrm{x}_k | \mathrm{x}_{k-1}^{(i)}, \mathrm{y}_k), \tag{6.5}$$

$$w_{q,k-1}^{(i)} \propto w_{k-1}^{(i)} p(\mathrm{y}_k | \mathrm{x}_{k-1}^{(i)}).$$

It is straightforward to verify that, in the usual setting, the marginal optimal proposal implies that weights never change:

$$w_k \propto \frac{\tilde{p}\left(\mathrm{x}_k | \mathrm{y}_{1:k}\right)}{\tilde{q}\left(\mathrm{x}_k | \mathrm{y}_{1:k}\right)} \propto \frac{p(\mathrm{y}_k | \mathrm{x}_k) \sum_{i=1}^{N} w_{k-1}^{(i)} p_t(\mathrm{x}_k | \mathrm{x}_{k-1}^{(i)})}{\sum_{i=1}^{N} w_{q,k-1}^{(i)} p(\mathrm{x}_k | \mathrm{x}_{k-1}^{(i)}, \mathrm{y}_k)}$$

$$\propto \frac{p(\mathrm{y}_k | \mathrm{x}_k) \sum_{i=1}^{N} w_{k-1}^{(i)} p_t(\mathrm{x}_k | \mathrm{x}_{k-1}^{(i)})}{\sum_{i=1}^{N} w_{k-1}^{(i)} p(\mathrm{y}_k | \mathrm{x}_{k-1}^{(i)}) \cdot \frac{p(\mathrm{y}_k | \mathrm{x}_k) p_t(\mathrm{x}_k | \mathrm{x}_{k-1}^{(i)})}{p(\mathrm{y}_k | \mathrm{x}_{k-1}^{(i)})}}$$

$$= \text{constant.}$$

This feature is crucial because it endows a particle filter with low variance of weights, which essentially turns into statistical efficiency. This finding motivates the marginal optimal proposal density as the foundation for a marginal particle filter based on the stochastic particle flow. The resulting filter is expected to work well for moderately high-dimensional problems.

6.2 Difficulties from a Usual Marginal Importance Density

This section discusses the problems that naturally arise when considering a standard Monte Carlo setting as (5.1) to build a marginal importance density based on the stochastic particle flow. If one regards the proposal distribution as the result of a sequence of L Markov transitions through a discretization of the interval $0 < \lambda \leq T$ onto the sub-intervals $\{\lambda_0 = 0, \lambda_1, \ldots \lambda_l, \ldots, \lambda_L = T\}$, where $\mathrm{x}_k \triangleq \mathrm{x}_L$ and $\mathrm{x}_{k-1} \triangleq \mathrm{x}_0$, then the sequence of transitions would provide the importance density

$$q\left(\mathrm{x}_k | \mathrm{y}_{1:k}\right) = \int_{\mathcal{X}} \int_{\mathcal{X}} \cdots \int_{\mathcal{X}} q(\mathrm{x}_L | \mathrm{x}_{L-1}, \mathrm{y}_k) q(\mathrm{x}_{L-1} | \mathrm{x}_{L-2}, \mathrm{y}_k) \ldots$$

$$q(\mathrm{x}_1 | \mathrm{x}_0, \mathrm{y}_k) q(\mathrm{x}_0 | \mathrm{y}_{1:k-1}) d\mathrm{x}_{L-1} d\mathrm{x}_{L-2} \ldots d\mathrm{x}_0. \tag{6.6}$$

In order to evaluate this importance density over a set of N particles, incorporating the previous set of

samples and importance weights, one would be required to compute

$$\tilde{q}(\mathbf{x}_k^{(i)}|\mathbf{y}_{1:k}) = \sum_{j=1}^{N} w_{k-1}^{(j)} \tilde{q}(\mathbf{x}_k^{(i)}|\mathbf{x}_{k-1}^{(j)}, \mathbf{y}_k), \tag{6.7}$$
$$i = 1, \dots, N.$$

This implementation depends on the set of conditional kernels $\tilde{q}(\mathbf{x}_k|\mathbf{x}_{k-1}, \mathbf{y}_k) = \tilde{q}(\mathbf{x}_L|\mathbf{x}_0, \mathbf{y}_k)$ that could be achieved in terms of a recursion of the form

$$w_1^{(i|j)} \triangleq q(\mathbf{x}_1^{(i)}|\mathbf{x}_0^{(j)}, \mathbf{y}_k), \quad i, j = 1, \dots, N;$$

$$w_2^{(i|j)} \triangleq \tilde{q}(\mathbf{x}_2^{(i)}|\mathbf{x}_0^{(j)}, \mathbf{y}_k) = \sum_{n=1}^{N} w_1^{(n|j)} q(\mathbf{x}_2^{(i)}|\mathbf{x}_1^{(n)}, \mathbf{x}_0^{(j)}, \mathbf{y}_k);$$

$$\vdots$$

$$w_l^{(i|j)} \triangleq \tilde{q}(\mathbf{x}_l^{(i)}|\mathbf{x}_0^{(j)}, \mathbf{y}_k) = \sum_{n=1}^{N} w_{l-1}^{(n|j)} q(\mathbf{x}_l^{(i)}|\mathbf{x}_{l-1}^{(n)}, \mathbf{x}_0^{(j)}, \mathbf{y}_k);$$

$$\vdots$$

$$\tilde{q}(\mathbf{x}_L^{(i)}|\mathbf{x}_0^{(j)}, \mathbf{y}_k) = \sum_{n=1}^{N} w_{L-1}^{(n|j)} q(\mathbf{x}_L^{(i)}|\mathbf{x}_{L-1}^{(n)}, \mathbf{x}_0^{(j)}, \mathbf{y}_k); \tag{6.8}$$

where $q(\mathbf{x}_l|\mathbf{x}_{l-1}, \mathbf{x}_0, \mathbf{y}_k)$ are the one-step proposal kernels conditioned on the initial state (prior samples), which are directly available from the discretized version of (4.15). Computing the conditional proposal components (6.8) and the marginal proposal (6.7) involves high computational effort, bounded by $\mathcal{O}\left((L-1)N^3 + N^2\right)$ evaluations. In addition, the main complication of this realization is due to the mixing properties of (6.7), leading to significant errors built up through the sequence of finite-sample approximations in (6.8) along with the prohibitively high variance of the resulting importance weights (6.3).

While these problems could be tentatively worked around by a judicious choice of a variance reduction method, it is worth looking how the implementation difficulties would turn out to be by evoking a hypothetical "continuity" between sampling steps. It is well known that in the limit $\Delta\lambda \to 0$, the proposal density (6.6) defines a path integral. Based on the concept of path probability density [81] of a Markov process

$$W_\infty \left[\mathbf{x}\left(\lambda\right)\right] \left[dx\right] \propto e^{-\int_0^T \left[\frac{1}{2}(\dot{\mathbf{x}} - \mu(\mathbf{x}))^T D^{-1}(\dot{\mathbf{x}} - \mu(\mathbf{x})) + \frac{1}{2}\nabla_\mathbf{x}\cdot\mu(\mathbf{x})\right]d\lambda}, \tag{6.9}$$

for samples describing continuous paths, the proposal could be written as a functional integral [82] of the form

$$q_c\left(\mathbf{x}_k|\mathbf{y}_{1:k}\right) \propto \int e^{-\int_0^T \left[\frac{1}{2}(\dot{\mathbf{x}} - \mu(\mathbf{x}))^T D^{-1}(\dot{\mathbf{x}} - \mu(\mathbf{x})) + \frac{1}{2}\nabla_\mathbf{x}\cdot\mu(\mathbf{x})\right]d\lambda} \left[dx\right], \tag{6.10}$$

where $[dx] = dx_{L-1} \dots dx_0$ as $\Delta\lambda \to 0$. Solving path integrals in general is a daunting task, nevertheless, a density of interest could be approximately obtained in terms of a Gaussian mixture, under the assumption of local Gaussianity of probability paths. Within this framework, an ensemble of independently selected Gaussian densities can be analytically integrated to achieve local solutions to (6.10). This fundamental idea is equivalent to what stochastic particle flow proposes when the filtering solution is formulated as the mixture (5.2).

6.3   The Stochastic-Particle-Flow Marginal Particle Filter

In marginal particle filtering, the best importance density one could achieve is the proposal density $q\left(\mathbf{x}_k|\mathbf{y}_{1:k}\right)$ when computed exactly. This density enables inference of the actual posterior pdf, $p\left(\mathbf{x}_k|\mathbf{y}_{1:k}\right)$. Composing the marginal optimal proposal requires computing $p(\mathbf{y}_k|\mathbf{x}_{k-1}^{(i)})$ exactly, which is not possible in general. In addition, the same scenarios that produce considerable errors in computing the empirical target, $\tilde{\pi}\left(\mathbf{x}_k\right) = \tilde{p}\left(\mathbf{x}_k|\mathbf{y}_{1:k}\right)$, according to (6.4), will also affect evaluation of the proposal $\tilde{q}\left(\mathbf{x}_k|\mathbf{y}_{1:k}\right)$, computed by (6.5), as illustrated by the first example in Section 7. In these cases, one can benefit from the inherent characteristics of stochastic

particle flow to construct a proposal density with better regularity properties by doing

$$
\begin{aligned}
\tilde{q}\left(\mathrm{x}_k|\mathrm{y}_{1:k}\right) &= \sum_{i=1}^{N} w_{k-1}^{(i)} p(\mathrm{y}_k|\mathrm{x}_k) p_t(\mathrm{x}_k|\mathrm{x}_{k-1}^{(i)}) \\
&= \sum_{i=1}^{N} w_{k-1}^{(i)} \frac{p_t(\mathrm{x}_k|\mathrm{x}_{k-1}^{(i)})}{\tilde{p}^{(i)}(\mathrm{x}_k|\mathrm{y}_{1:k-1})} p(\mathrm{y}_k|\mathrm{x}_k) \tilde{p}^{(i)}(\mathrm{x}_k|\mathrm{y}_{1:k-1}) \\
&\propto \sum_{i=1}^{N} w_{k-1}^{(i)} \frac{p_t(\mathrm{x}_k|\mathrm{x}_{k-1}^{(i)})}{\tilde{p}^{(i)}(\mathrm{x}_k|\mathrm{y}_{1:k-1})} w_{m,k}^{(i)} \mathcal{N}(\mathrm{x}_k; \mu_{m,k}^{(i)}, \Sigma_{m,k}^{(i)}), \\
\tilde{q}\left(\mathrm{x}_k|\mathrm{y}_{1:k}\right) &= \sum_{i=1}^{N} w_q^{(i)}(\mathrm{x}_k) \mathcal{N}(\mathrm{x}_k; \mu_{m,k}^{(i)}, \Sigma_{m,k}^{(i)});
\end{aligned}
\tag{6.11}
$$

where

$$
\bar{w}_q^{(i)}(\mathrm{x}_k) = w_{k-1}^{(i)} w_{m,k}^{(i)} \frac{p_t(\mathrm{x}_k|\mathrm{x}_{k-1}^{(i)})}{\tilde{p}^{(i)}(\mathrm{x}_k|\mathrm{y}_{1:k-1})},
\tag{6.12}
$$

$$
w_q^{(i)}(\mathrm{x}_k) = \frac{\bar{w}_q^{(i)}(\mathrm{x}_k)}{\sum_{i=1}^{N} \bar{w}_q^{(i)}(\mathrm{x}_k)};
\tag{6.13}
$$

and $\tilde{p}^{(i)}(\mathrm{x}_k|\mathrm{y}_{1:k-1})$ is a per-sample, local prior density. For a known Markov transition density $p_t\left(\mathrm{x}_k'|\mathrm{x}_{k-1}\right)$, we recall the local prior density as given by

$$
\tilde{p}^{(i)}\left(\mathrm{x}_k'|\mathrm{y}_{1:k-1}\right) = \int_{\mathcal{X}} p_t(\mathrm{x}_k'|\mathrm{x}_{k-1}) \mathcal{N}(\mathrm{x}_{k-1}; \mu_{m,k-1}^{(i)}, \Sigma_{m,k-1}^{(i)}) d\mathrm{x}_{k-1}.
\tag{6.14}
$$

As in the classical Gaussian-sum setting, the mixture weights $\{w_{m,k}^{(i)}\}$ are given by (see [74], pages 214 and 215)

$$
w_{m,k}^{(i)} \propto \frac{1}{N} \int_{\mathcal{X}} p\left(\mathrm{y}_k|\mathrm{x}_k'\right) \tilde{p}^{(i)}\left(\mathrm{x}_k'|\mathrm{y}_{1:k-1}\right) d\mathrm{x}_k',
\tag{6.15}
$$

where the proportionality to N^{-1} holds because stochastic particle flow generates equally weighted mixture components. The mixture weights $\{w_{m,k}^{(i)}\}$ generated by this method are only applicable in the context of the proposal (6.11), and shall be re-evaluated in the same way whenever a new instance of the marginal proposal is constructed. It is relevant to make clear the distinction $\mathrm{x}_k' \neq \mathrm{x}_k$ in the expressions (6.15) and (6.14), bearing in mind that x_k' corresponds to the state that the flow would reach when considering only the prior density as the target $\pi_{\mathrm{prior}}(\mathrm{x}') = p_x(\mathrm{x}') \triangleq p(\mathrm{x}'|\mathrm{y}_{1:k-1})$. We note that the involved integrals may not be tractable in general and may require approximation either by a Gaussian representation of the likelihood, or adequate quadrature rules (e.g., Gauss-Hermite).

This formulation evokes stochastic particle flow to promote an accurate approximation to the marginal optimal proposal density. Given a set of samples and parameters $\{\mathrm{x}_{k-1}^{(i)}, w_{k-1}^{(i)}, \mu_{m,k-1}^{(i)}, \Sigma_{m,k-1}^{(i)}\}$ from a previous filtering iteration, where $w_{k-1}^{(i)}$ are importance weights, the algorithm integrates the SDE (4.15) for each sample and propagates the associated parameters through the interval $0 < \lambda \leq T$. As result, the procedure acquires posterior samples and parameters, $\{\mathrm{x}_k^{(i)}, w_k^{(i)}, \mu_{m,k}^{(i)}, \Sigma_{m,k}^{(i)}\}$, which are used to evaluate the marginal proposal (6.11) and enable filtering as by a marginal particle filter. The moments of the mixture's components are evolved in accordance with (5.7) and (5.8), and the importance weights are updated by

$$
w_k(\mathrm{x}_k) \propto \frac{\sum_{j=1}^{N} w_{k-1}^{(j)} p(\mathrm{y}_k|\mathrm{x}_k) p_t(\mathrm{x}_k|\mathrm{x}_{k-1}^{(j)})}{\sum_{j=1}^{N} w_q^{(j)}(\mathrm{x}_k) \mathcal{N}(\mathrm{x}_k; \mu_{m,k}^{(j)}, \Sigma_{m,k}^{(j)})}.
\tag{6.16}
$$

The resulting filter, called stochastic-particle-flow marginal particle filter (SPF-MPF), is summarized in *Algorithm 2*. It is worth noting that a simpler alternative to (6.11) could be chosen by considering

$$
\tilde{q}\left(\mathrm{x}_k|\mathrm{y}_{1:k}\right) = \sum_{i=1}^{N} \frac{1}{N} \mathcal{N}(\mathrm{x}_k; \mu_{m,k}^{(i)}, \Sigma_{m,k}^{(i)}),
\tag{6.17}
$$

however, in that case, the importance density would not be affected by the same errors as the empirical target, $\tilde{\pi}(\mathrm{x}_k) = \tilde{p}(\mathrm{x}_k|\mathrm{y}_{1:k})$, as computed by (6.4), because each component in (6.17) targets a local instance of the posterior density itself. As a result, even though the importance density could approximate the true posterior density accurately, it would not directly approach the target density. In situations where the empirical target density cannot represent the true posterior density as well as a mixture of the form (6.17), the SPF-MPF with such a proposal would fail because of the mismatch originated from distinctions in the approximation methods. As a consequence, the importance weights would have infeasibly high variance. The described issue is equivalent to treat errors in the standard Monte Carlo measure (5.1) as comparable to errors in the mixture measure (5.3), which is not true except for rare cases. This scenario is well illustrated by two examples in Section 7.

7 Examples

In this section we present some illustrative toy examples and experimental results for three instructive applications in the multi-sensor multi-target tracking context: a multi-sensor bearing-only problem, a convoy tracking problem, and inference on a large spatial sensor network as presented by Septier & Peters [56].

In the experimental results for the bearing-only and convoy tracking examples, we compared the SPF-GS against standard target trackers and extensions of two of the most effective particle flows, namely, the Gaussian particle flow (GPF) and the scaled-drift particle flow (SDPF). The GPF was first called *exact particle flow* in [57] and the SDPF was first called *non-zero diffusion particle flow* in [54]. Actually, this latter is a particle flow with the drift scaled by a diffusion coefficient, but the filter itself is not a diffusion.

It is important to mention that, in order to work properly, both the Gaussian particle flow and the scaled-drift particle flow are implemented with the aid of a companion filter such that the state covariance matrix can be correctly estimated. Implementation details have been presented by Choi *et al.* [50] and Ding & Coates [51], who advocate using the EKF (or UKF) as a companion filter to estimate the associated covariance matrices. Another option is to shrink the empirical covariance and apply Tikhonov regularization [52]. In contrast, the stochastic particle flow does not require any auxiliary technique to estimate the second order moment, relying solely on its mixture measure. In the toy examples a companion filter was not necessary for the original particle flows since a single filtering cycle has been analyzed. In the multi-sensor and multi-target examples we adopted baseline filters, which are the most structurally similar to the EKF, as companion filters for the particle flows (GPF, SDPF).

In the example of the large spatial sensor network, we compared the SPF-GS, a particle filter (Sequential Importance Resampling - SIR), a block particle filter (block SIR), and two of the best sequential MCMC filters [10, 56]: the Sequential manifold Metropolis-Adjusted Algorithm (SmMALA) and the Sequential manifold Hamiltonian Monte Carlo (SmHMC). The block particle filter partitions the state space into separate subspaces of smaller dimensions and run a particle filter on each subspace [23].

7.1 Toy Examples

The toy examples are Gaussian processes chosen to demonstrate the properties of the stochastic particle flow, summarized as

- Univariate
  - linear,
  - quadratic,
  - cubic;
- Bivariate
  - multimodal, linear,
  - nonlinear (banana-shaped pdf).

In all cases, we analyze the filters for a single filtering cycle. Generally, we describe the state process, the observation process and the initial distribution for these examples as

$$\mathrm{x}_k = f(\mathrm{x}_{k-1}) + \mathrm{u}_k, \ \mathrm{u}_k \sim \mathcal{N}(\mathrm{u}_k; 0, Q_k), \tag{7.1}$$

$$\mathrm{y}_k = h(\mathrm{x}_k \quad) + \mathrm{v}_k, \ \mathrm{v}_k \sim \mathcal{N}(\mathrm{v}_k; 0, R_k), \tag{7.2}$$

$$p_0(\mathrm{x}_{k-1}) = \mathcal{N}(\mathrm{x}_{k-1}; \bar{\mathrm{x}}_{k-1}, P_{k-1}). \tag{7.3}$$

Algorithm 2: Stochastic particle flow - marginal particle filter

---

1 Initialization:
2 if *time $k = 0$* **then**
3   Sample $x_0^{(i)} \sim p(x_0)$ and set $w_0^{(i)} := N^{-1}, \forall i = 1, \ldots, N$
4   Set $w_{m,0}^{(i)} := N^{-1}, \mu_{m,0}^{(i)} := \mathbb{E}_{p_0}[x_0], \Sigma_{m,0}^{(i)} := \mathbb{E}_{p_0}\left[(x_0 - \bar{x}_0)(x_0 - \bar{x}_0)^T\right], \forall i = 1, \ldots, N$
5 end

6 Steps:
7 for *time $k \geq 1$* **do**
8   Compute the time horizon T and step size $\Delta\lambda$ (section 4.6)
9   Discretize the interval $0 \leq \lambda \leq T$ into L sub-intervals $\{\lambda_0 = 0, \ldots \lambda_l, \ldots, \lambda_L = T\}$
10   Set $x_{l=0}^{(i)} := x_{k-1}^{(i)}, \mu_{l=0}^{(i)} := x_{k-1}^{(i)}, \Sigma_{l=0}^{(i)} := 0_{n_x \times n_x}, \forall i = 1, \ldots, N$
11   **for** $l = 1$ **to** L **do**
12   **for** $i = 1, \ldots, N$ **do**
13   Simulate

$$x_l^{(i)} \leftarrow x_{l-1}^{(i)} + \frac{1}{2}\int_{\lambda_{l-1}}^{\lambda_l} D(x_{l-1}^{(i)})\nabla_x \log \tilde{\pi}(x_{l-1}^{(i)})d\lambda + \int_{\lambda_{l-1}}^{\lambda_l} D(x_{l-1}^{(i)})^{1/2}dw_\lambda$$

14   Propagate

$$\mu_l^{(i)} \leftarrow \mu_{l-1}^{(i)} + \int_{\lambda_{l-1}}^{\lambda_l} \left[C(x_{l-1}^{(i)})\mu_{l-1}^{(i)} + c(x_{l-1}^{(i)})\right] d\lambda,$$

$$\Sigma_l^{(i)} \leftarrow \Sigma_{l-1}^{(i)} + \int_{\lambda_{l-1}}^{\lambda_l} \left[C(x_{l-1}^{(i)})\Sigma_{l-1}^{(i)} + \Sigma_{l-1}^{(i)}C^T(x_{l-1}^{(i)}) + D(x_{l-1}^{(i)})\right] d\lambda$$

15   **end**
16   **end**
17   Set $x_k^{(i)} := x_{l=L}^{(i)}, \mu_{m,k}^{(i)} := \mu_{l=L}^{(i)}, \Sigma_{m,k}^{(i)} := \Sigma_{l=L}^{(i)}, \forall i = 1, \ldots, N$
18   Compute the normalized proposal weights, $\forall i = 1, \ldots, N$, by

$$w_{m,k}^{(i)} \propto \frac{1}{N}\int_{\mathcal{X}} p\left(y_k|x_k'\right)\tilde{p}^{(i)}\left(x_k'|y_{1:k-1}\right)dx_k',$$

$$w_q^{(i)} \propto w_{k-1}^{(i)}w_{m,k}^{(i)}\frac{p_t(x_k^{(i)}|x_{k-1}^{(i)})}{\tilde{p}^{(i)}(x_k^{(i)}|y_{1:k-1})}$$

19   Compute the normalized importance weights, $\forall i = 1, \ldots, N$, by

$$w_k^{(i)} \propto \frac{\sum_{j=1}^N w_{k-1}^{(j)}p(y_k|x_k^{(i)})p_t(x_k^{(i)}|x_{k-1}^{(j)})}{\sum_{j=1}^N w_q^{(j)}\mathcal{N}(x_k^{(i)};\mu_{m,k}^{(j)},\Sigma_{m,k}^{(j)})}$$

20   **if** $ESS_k < 0.5N$ **then** resample: $\{x_k^{(i)}, N^{-1}\} \leftarrow \{x_k^{(i)}, w_k^{(i)}\}$
       Output: Approximation of the filtering distribution by the empirical measure
21

$$\tilde{p}(x_k|y_{1:k}) = \sum_{i=1}^N w_k^{(i)}\delta(x_k - x_k^{(i)})$$

22 end

We consider four different types of particle filters based on the marginal importance density

$$\tilde{q}\left(\mathrm{x}_k|\mathrm{y}_{1:k}\right) = \sum_{i=1}^{N} w_{k-1}^{(i)} q(\mathrm{x}_k|\mathrm{x}_{k-1}^{(i)}, \mathrm{y}_k),$$

where

- for the marginal bootstrap particle filter (MBPF), the proposal's components are set as the Markov transition kernel: $q(\mathrm{x}_k|\mathrm{x}_{k-1}^{(i)}, \mathrm{y}_k) = p_t(\mathrm{x}_k|\mathrm{x}_{k-1}^{(i)})$;

- for the marginal EKF-based particle filter (MEPF), the proposal's components are computed by the EKF: $q(\mathrm{x}_k|\mathrm{x}_{k-1}^{(i)}, \mathrm{y}_k) = p_{\mathrm{EKF}}(\mathrm{x}_k|\mathrm{x}_{k-1}^{(i)}, \mathrm{y}_k)$;

- for the marginal UKF-based particle filter (MUPF), the proposal's components are computed by the UKF: $q(\mathrm{x}_k|\mathrm{x}_{k-1}^{(i)}, \mathrm{y}_k) = p_{\mathrm{UKF}}(\mathrm{x}_k|\mathrm{x}_{k-1}^{(i)}, \mathrm{y}_k)$; and

- for the marginal auxiliary particle filter (MAPF) [80], the importance density is given by (6.5).

When comparing probability densities furnished by different filters, we include the empirical marginal target, $\tilde{\pi}\left(\mathrm{x}_k\right) = \tilde{p}\left(\mathrm{x}_k|\mathrm{y}_{1:k}\right)$, evaluated according to (6.4) for samples obtained by stochastic particle flow. For all filters, when applicable, we calculate the average of the effective sample size

$$\mathrm{ESS} = \left(\sum_{i=1}^{N} w_k^{(i)\,2} \right)^{-1} \tag{7.4}$$

over 100 Monte Carlo runs, for 1000 particles. For all marginal proposal densities, we analyze their similiarity to the true posterior probability density by averaging their empirical Jensen-Shannon divergence (JSD) with respect to the true posterior, which is obtained to high numerical precision. The Jensen-Shannon divergence is defined as

$$\begin{aligned} \mathrm{JSD}\left(P \parallel Q\right) &= \frac{1}{2}\mathrm{D}_{\mathrm{KL}}(P \parallel (P+Q)/2) \\ &\quad + \frac{1}{2}\mathrm{D}_{\mathrm{KL}}(Q \parallel (P+Q)/2), \end{aligned} \tag{7.5}$$

where the Kullback–Leibler divergence, $\mathrm{D}_{\mathrm{KL}}(\cdot \parallel \cdot)$, is computed using the base-2 logarithm such that the Jensen-Shannon divergence is bounded as $0 \leq \mathrm{JSD}\left(P \parallel Q\right) \leq 1$. The Jensen-Shannon divergence is symmetric and equals zero when the compared densities are equal. In the bivariate examples we also consider the original particle flow methods, the Gaussian particle flow (GPF) and scaled-drift particle flow (SDPF), for which the Jensen-Shannon divergence with respect to the true posterior is evaluated based on empirical densities constructed by (bidimensional) histograms of samples.

7.1.1 Linear, Univariate Model

The simplest example is a linear, univariate model, with parameters set as in the table below.

Parameters for the linear, univariate model	
Initial distribution	$\bar{\mathrm{x}}_{k-1} = 0,\ P_{k-1} = 20$
Markov transition pdf	$f(\mathrm{x}_{k-1}) = \mathrm{x}_{k-1},\ Q_k = 5$
Likelihood function	$h(\mathrm{x}_k) = \mathrm{x}_k,\ R_k = 10$
Observation	$\mathrm{y}_k = 30$

Although very simple, this example was proposed to demonstrate a scenario where the empirical marginal target, $\tilde{\pi}\left(\mathrm{x}_k\right) = \tilde{p}\left(\mathrm{x}_k|\mathrm{y}_{1:k}\right)$, is prone to relevant statistical and numerical errors. This is done by setting a situation where the transition kernel describes a Markov process with small variance and the observation lies relatively far from the initial distribution. In this scenario, statistical inefficiency emerges because the observation provides little information in the space region where probability masses are more densely distributed by the state process. Not incidentally, this is also the main source of degeneracy in standard particle filters. Additionally, there may exist round-off errors when evaluating the empirical marginal target owing to samples

being located relatively far from the posterior mean, several standard deviations apart, in the tail of each proposal component.

As depicted in Figure 7.1, the importance density proposed by the SPF-MPF (red x's) is successful at aiming the empirical marginal target (blue circles), generating a high effective sample size. However, since the empirical target constitutes a poor approximation to the true posterior pdf (black line), importance sampling clearly fails and the SPF-MPF leads to a solution excessively biased. In contrast, the direct filtering density generated by the SPF-GS approximates the true posterior pdf accurately, generating a satisfactory solution. These findings are quantified by the Jensen-Shannon divergences averaged over 100 Monte Carlo runs and presented in Table 1. Table 1 shows a neglible divergence between the density filtered by the SPF-GS and the true posterior whereas the divergences computed for the target density and for the proposal density constructed by the SPF-MPF are significant.
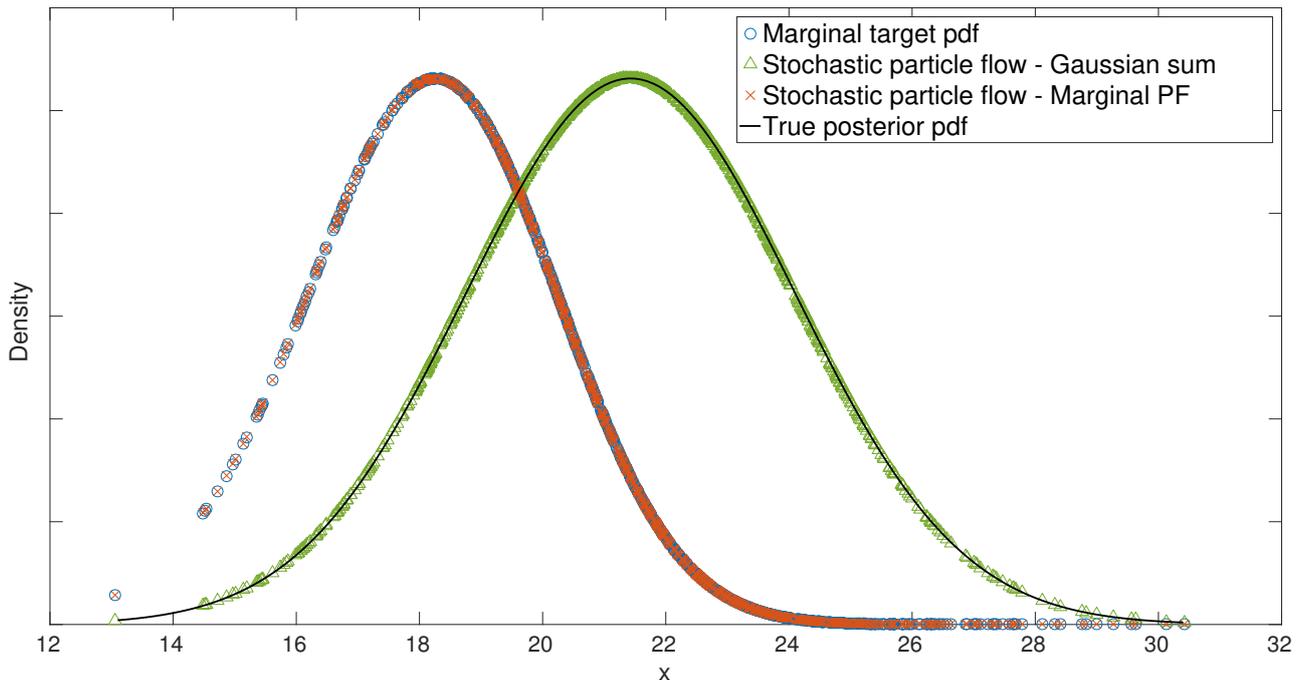
Figure 7.1: Densities generated by the SPF-GS and SPF-MPF for the linear, univariate example

7.1.2 Quadratic, Univariate Model

The quadratic, univariate model was tested with parameters set as shown in the following table. This model is interesting because nonlinearity of the observation process leads to bimodality of the filtered density.

Parameters for the quadratic, univariate model	
Initial distribution	$\bar{x}_{k-1} = 0$, $P_{k-1} = 20$
Markov transition pdf	$f(x_{k-1}) = x_{k-1}$, $Q_k = 20$
Likelihood function	$h(x_k) = x_k^2/20$, $R_k = 50$
Observation	$y_k = 30$

This nonlinear example was set to be favourable for marginal importance sampling such that it would be possible to compare different marginal particle filters against the SPF-MPF. The original particle flows, GPF and SDPF, are compared to the SPF-MPF as well. The quantified performances for this quadratic univariate model are shown in Table 1.

Firstly, we compare the sequence of histograms achieved when propagating samples by the GPF, by the SDPF and by the SPF-GS. As it can be seen in Figure 7.2, for this example, stochastic particle flow provides the best distribution of particles to approximate the posterior density, denoting a higher level of accuracy and regularity of the flow formulated as a diffusion.

Regarding the marginal importance densities illustrated in Figure 7.3, we observe a high degree of similarity of the SPF-MPF proposal density to the marginal target density. In the same manner, the filtering density
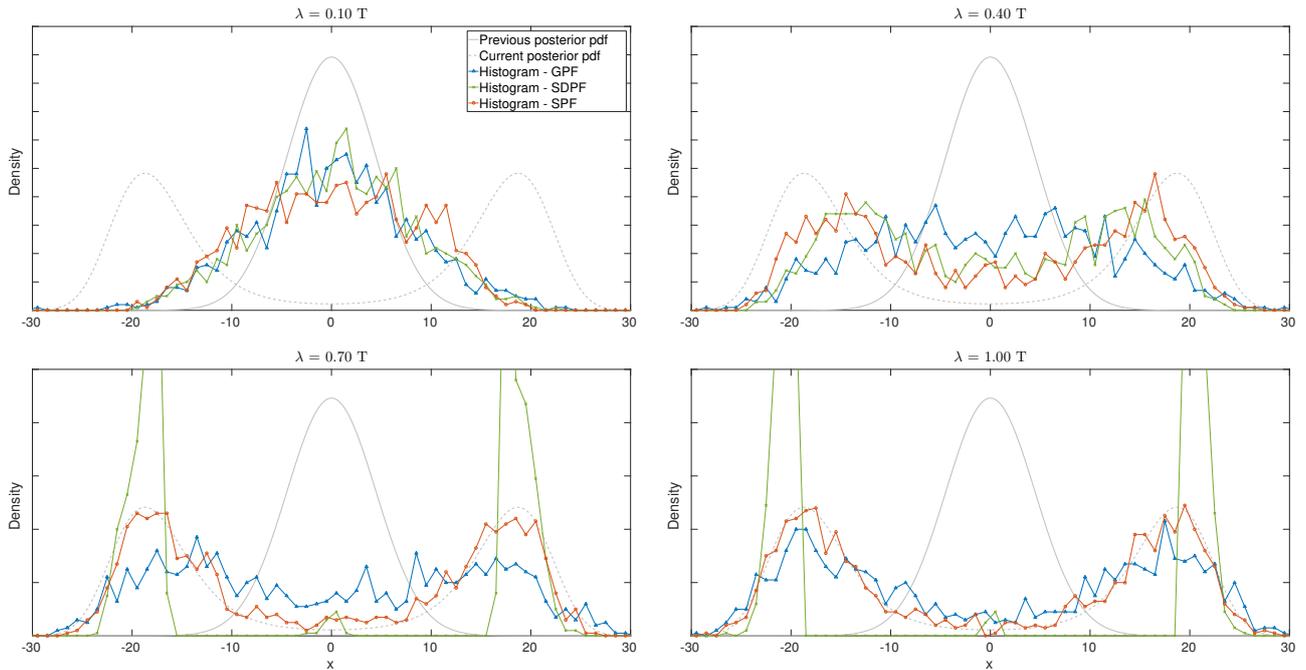
Figure 7.2: Sequence of histograms achieved by propagating particles for the quadratic, univariate example

achieved by the SPF-GS accurately approximates the true posterior density, as evidenced in Table 1. In Figure 7.4 we can see in detail the proximity of the SPF-MPF proposal density to both the marginal target density and the true posterior density, along with some of the proposal mixture components. The density proposed by the marginal (optimal) auxiliary particle filter (MAPF) is also very similar to the marginal target, providing an accurate solution, whereas all other filters propose densities less effective for this example. These observations are quantitatively captured by the performance data summarized in Table 1.
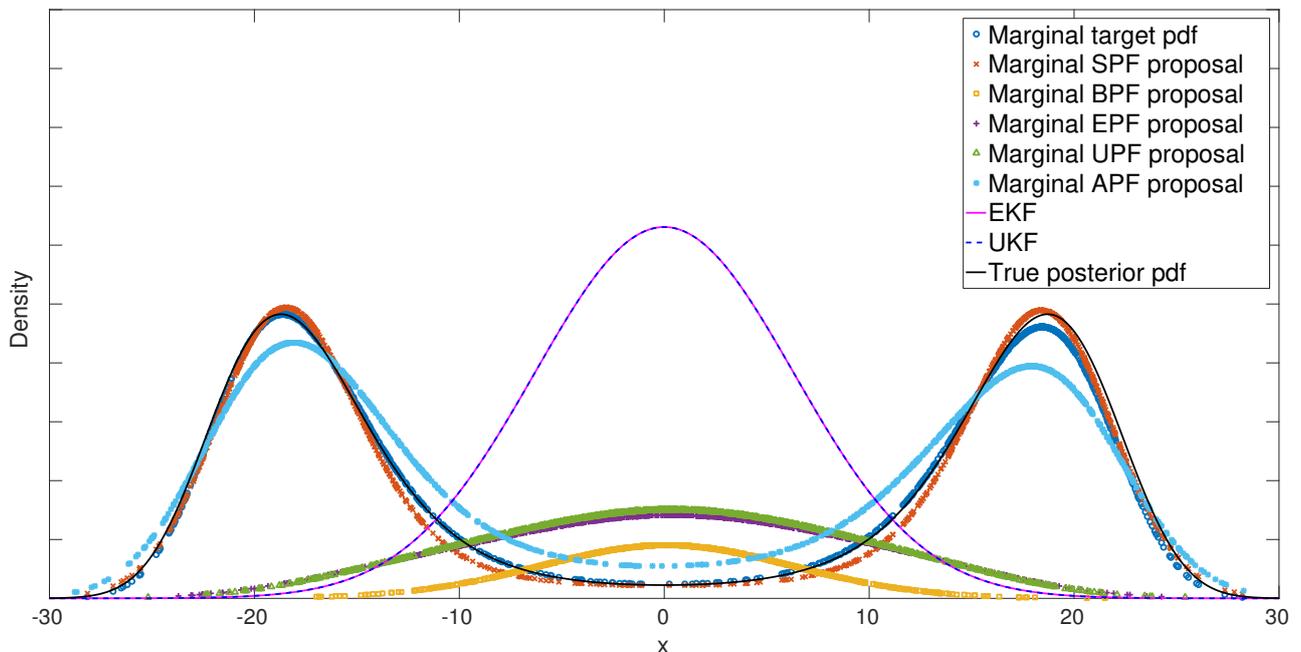
Figure 7.3: Comparison of proposal densities for the quadratic, univariate example
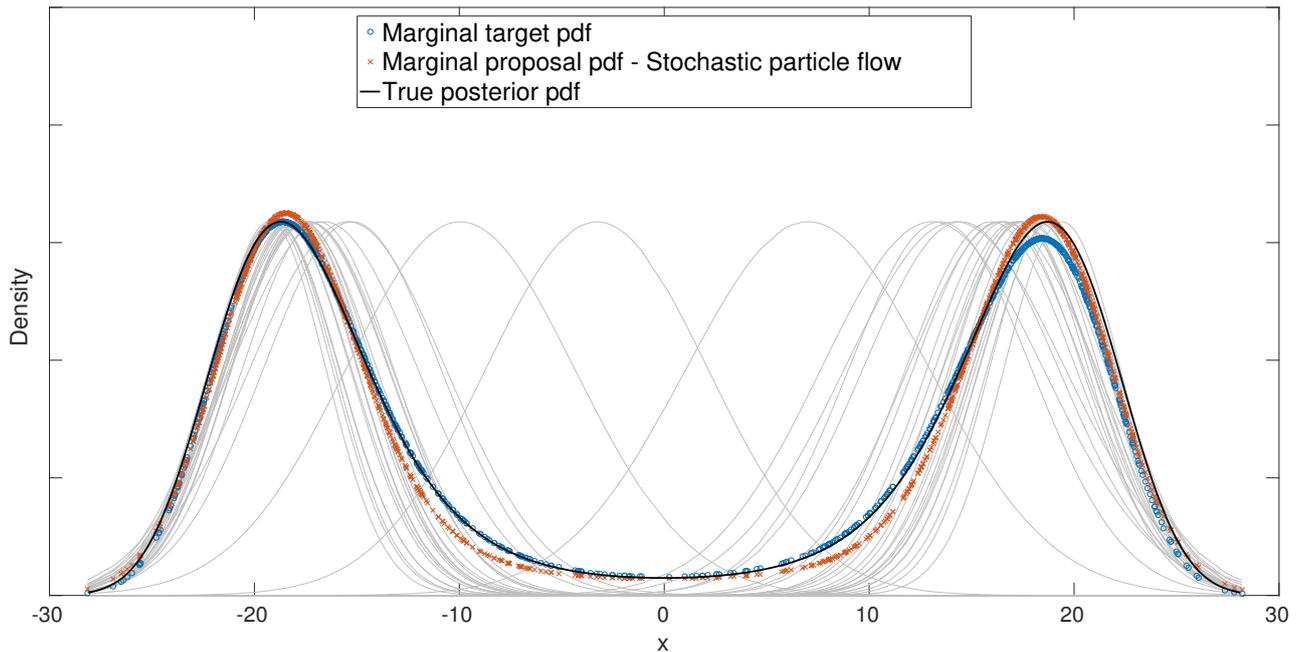
Figure 7.4: Marginal proposal density based on the stochastic particle flow

7.1.3 Cubic, Univariate Model

The cubic, univariate model was tested with parameters set as shown in the following table.

Parameters for the cubic, univariate model	
Initial distribution	$\bar{x}_{k-1} = 0,\ P_{k-1} = 20$
Markov transition pdf	$f(x_{k-1}) = x_{k-1},\ Q_k = 20$
Likelihood function	$h(x_k) = x_k^3/120,\ R_k = 50$
Observation	$y_k = 20$

This nonlinear example was also set to be favourable for marginal importance sampling, i.e., avoiding the scenario described in the first toy example where importance sampling fails. By comparing the resulting histograms achieved when propagating samples by the GPF, by the SDPF and by the stochastic particle flow, it is remarkable in Figure 7.5 that the stochastic particle flow provides a fairly superior distribution of particles to approximate the posterior density. This superiority is incorporated in the importance density proposed by the SPF-MPF as can be seen in Figure 7.6. The importance density proposed by the marginal auxiliary particle filter (MAPF) also provides an accurate solution to the filtering problem, but it is slightly less effective than the SPF-MPF. All the other marginal particle filters present less effective solutions. The comparison of all filters for this example is quantified in Table 1.

$$\lambda = 1.00 \text{ T}$$
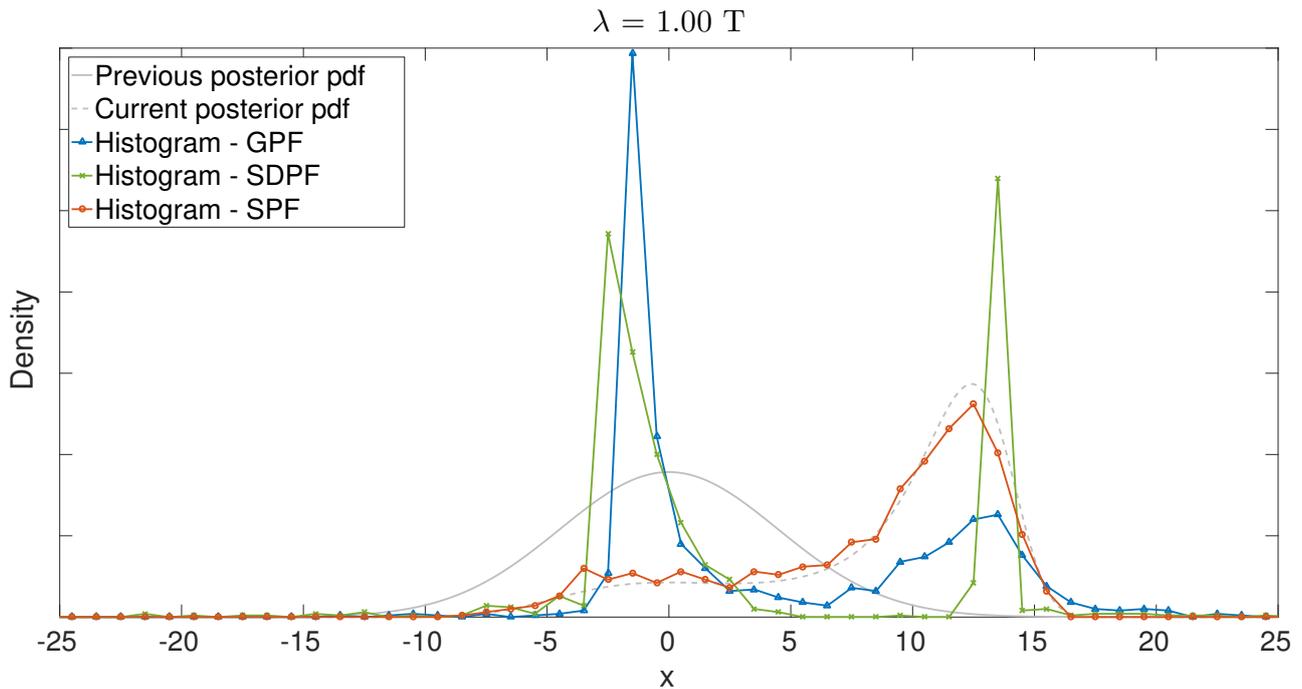
Figure 7.5: Resulting histograms of particles for the cubic, univariate example
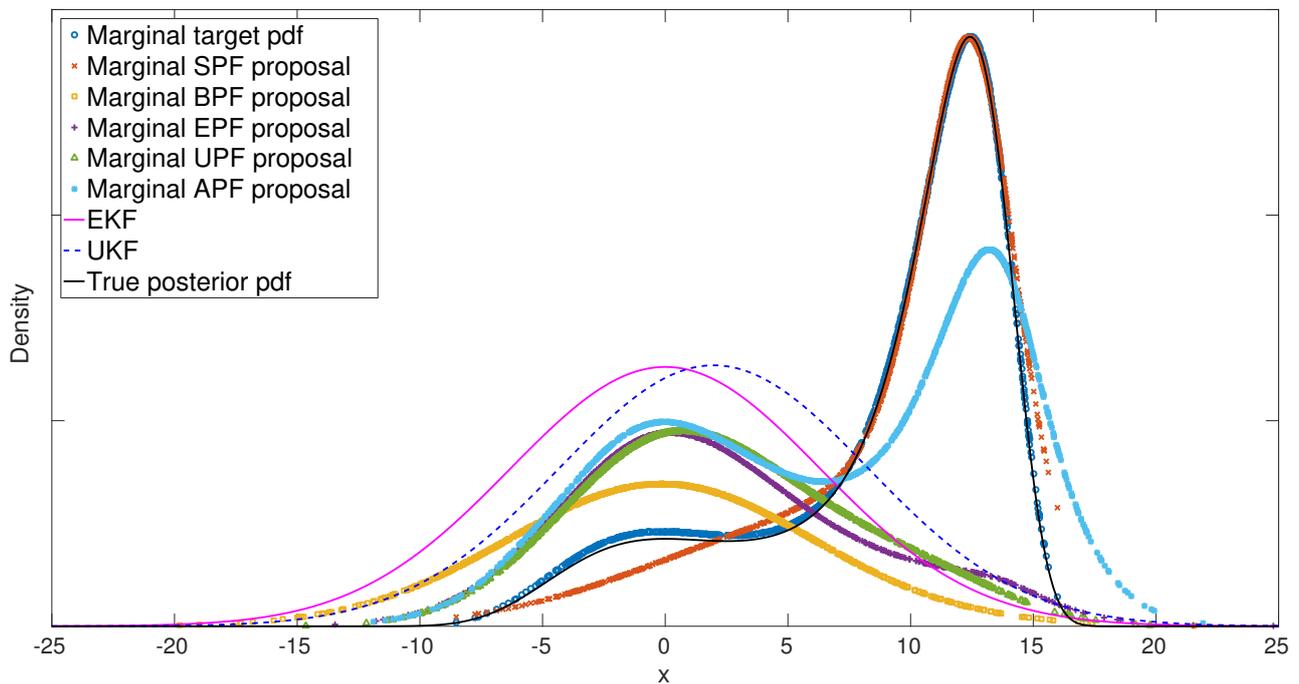
Figure 7.6: Comparison of proposal densities for the cubic, univariate example

7.1.4 Linear, Bimodal, Bivariate Model

This example poses a bimodal model where the modes arise from two different observations with a joint likelihood explicitly known. In the algorithm for propagating particles, we implemented a scheme that preselects samples to be filtered for either observation. This is done according to a set of indexes that are sampled from a binomial distribution $B\left(u_1, u_2; 1, w_{l,(1)}, w_{l,(2)}\right) \propto w_{l,(1)}^{u_1} w_{l,(2)}^{u_2}$ where $u_1, u_2 \in [0, 1]$, $u_1 + u_2 = 1$, such that indexes are uniquely associated to either event u_1 or u_2, with probability of either observation, $w_{l,(1)}$ or $w_{l,(2)}$ respectively.

Density	Linear		Quadratic		Cubic	
	JSD_{avg}	ESS_{avg}	JSD_{avg}	ESS_{avg}	JSD_{avg}	ESS_{avg}
Marginal target	0.1572	-	0.0028	-	0.0001	-
SPF-GS	0.0000	-	0.0013	-	0.0165	-
SPF-MPF	0.1574	100.00%	0.0052	97.12%	0.0071	96.74%
Marginal BPF	0.9876	0.21%	0.2641	1.79%	0.1723	12.60%
Marginal EPF	0.7857	1.90%	0.3097	16.69%	0.1820	28.63%
Marginal UPF	0.7870	2.04%	0.3112	14.46%	0.1674	25.32%
Marginal APF	0.0670	100.00%	0.0153	92.92%	0.0596	72.00%

Table 1: Comparison of densities for the univariate examples

The linear, bimodal, bivariate model was tested with parameters set as shown in Table 2. These parameters were chosen to result in quite distinct local properties of the two modes.

Parameters for the linear, bimodal, bivariate model	
Initial distribution	$\bar{\mathrm{x}}_{k-1} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}$, $P_{k-1} = \begin{pmatrix} 9 & 0 \\ 0 & 9 \end{pmatrix}$
Markov transition pdf	$f(\mathrm{x}_{k-1}) = \mathrm{x}_{k-1}$, $Q_k = \begin{pmatrix} 16 & 0 \\ 0 & 16 \end{pmatrix}$
Likelihood function:	$h(\mathrm{x}_k) = \mathrm{x}_k$
Mode 1	$R_{k,(1)} = \begin{pmatrix} 0.8 & 0 \\ 0 & 0.2 \end{pmatrix}$, $w_{l,(1)} = 0.2$
Mode 2	$R_{k,(2)} = \begin{pmatrix} 4.0 & 0 \\ 0 & 1.0 \end{pmatrix}$, $w_{l,(2)} = 0.8$
Observations	$\mathrm{y}_{k,(1)} = \begin{pmatrix} +10 \\ +20 \end{pmatrix}$, $\mathrm{y}_{k,(2)} = \begin{pmatrix} +10 \\ -20 \end{pmatrix}$

Table 2: Parameters for the bimodal bivariate model

For this example, we analyze stochastic particle flow methods, SPF-GS and SPF-MPF, against original particle flow methods only. We exemplify the sequence of particles' distributions acquired by the GPF, by the SDPF and by the stochastic particle flow in Figure 7.7. It becomes clear that the final distribution generated by the stochastic particle flow is closely similar to the true posterior density, precisely describing the local moments of the two modes. In opposition, the GPF generates a distribution that is excessively biased for the most peaky mode whereas the SDPF generates a distribution that does not describe correctly the covariances of each mode.

These findings are quantified by the average Jensen-Shannon divergences presented in Table 5. Table 5 shows a small divergence of the density filtered by the SPF-GS with respect to the true posterior, a small divergence of the SPF-MPF proposal density as well as of the target density, whereas the divergences of the original particle flows are fairly big. The SPF-MPF provides a high effective sample size.

7.1.5 Nonlinear, Unimodal, Bivariate Model

The nonlinear bivariate model was tested in two cases:

1. favourable for marginal particle filters, and

2. unfavourable, i.e., emulating a scenario similar to that presented in the first toy example where importance sampling fails.

The parameters used for cases 1 and 2 are presented in Table 3 and Table 4, respectively.

In either cases the sequence of distributions generated by the original particle flows and by the stochastic particle flow are as illustrated in Figure 7.8. Once more it becomes evident that the stochastic particle flow provides a superior distribution of samples to approximate the posterior density, which demonstrates its higher level of accuracy and regularity. Similarly to results presented for previous examples, the GPF seems to generate substantially biased distributions whereas the SDPF seems highly prone to regularity problems. These aspects are well corroborated by the average Jensen-Shannon divergences presented in Table 5.
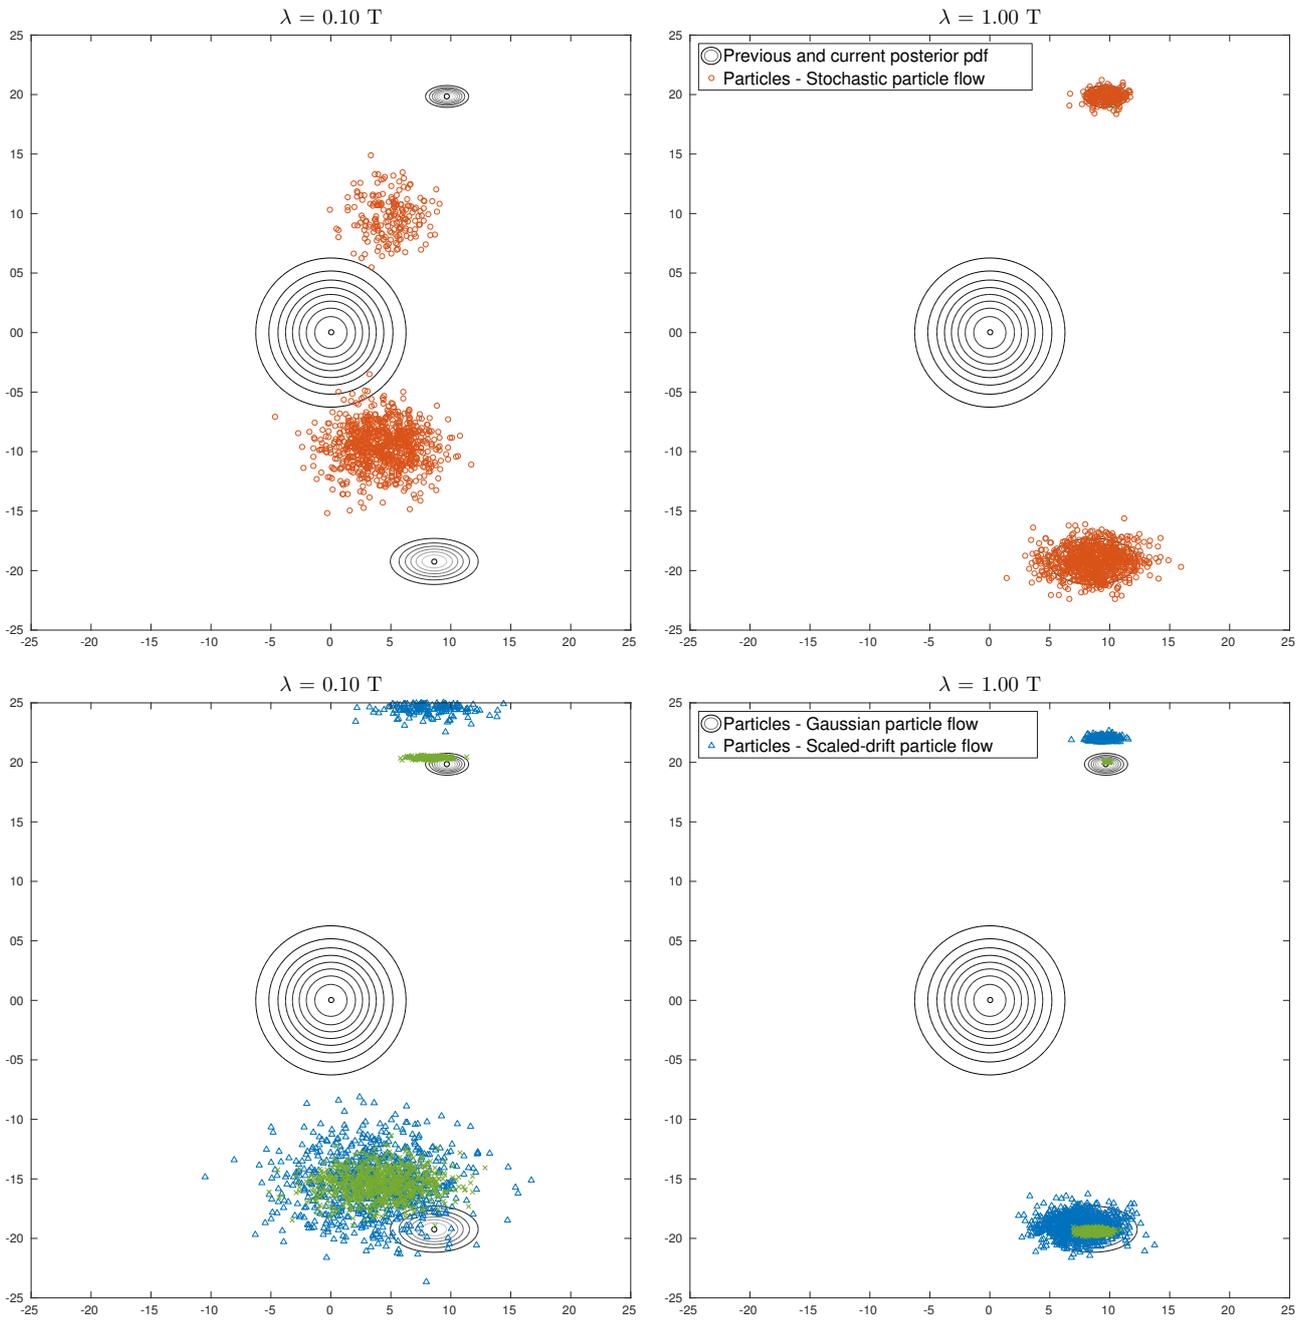
Figure 7.7: Sequence of distributions achieved by propagating particles for the bimodal, bivariate example

Parameters for the nonlinear bivariate model, case 1	
Initial distribution	$\bar{\mathrm{x}}_{k-1} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}, P_{k-1} = \begin{pmatrix} 20 & 0 \\ 0 & 20 \end{pmatrix}$
Markov transition pdf	$f(\mathrm{x}_{k-1}) = \mathrm{x}_{k-1}, Q_k = \begin{pmatrix} 20 & 0 \\ 0 & 20 \end{pmatrix}$
Likelihood function	$h(\mathrm{x}_k) = \begin{pmatrix} \sqrt{\mathrm{x}_k(1)^2 + \mathrm{x}_k(2)^2} \\ \mathrm{atan}\left(\mathrm{x}_k(2)/\mathrm{x}_k(1)\right) \end{pmatrix},$ $R_k = \begin{pmatrix} 1.00 & 0 \\ 0 & 0.16 \end{pmatrix}$
Observation	$\mathrm{y}_k = \begin{pmatrix} 20 \\ 0° \end{pmatrix}$

Table 3: Parameters for the nonlinear bivariate model, case 1

Parameters for the nonlinear bivariate model, case 2	
Initial distribution	$\bar{\mathrm{x}}_{k-1} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}, P_{k-1} = \begin{pmatrix} 10 & 0 \\ 0 & 10 \end{pmatrix}$
Markov transition pdf	$f(\mathrm{x}_{k-1}) = \mathrm{x}_{k-1}, Q_k = \begin{pmatrix} 5 & 0 \\ 0 & 5 \end{pmatrix}$
Likelihood function	$h(\mathrm{x}_k) = \begin{pmatrix} \sqrt{\mathrm{x}_k(1)^2 + \mathrm{x}_k(2)^2} \\ \mathrm{atan}\left(\mathrm{x}_k(2)/\mathrm{x}_k(1)\right) \end{pmatrix},$ $R_k = \begin{pmatrix} 1.00 & 0 \\ 0 & 0.16 \end{pmatrix}$
Observation	$\mathrm{y}_k = \begin{pmatrix} 20 \\ 0° \end{pmatrix}$

Table 4: Parameters for the nonlinear bivariate model, case 2

In the comparison we also included other marginal particle filters. For *case 1* (favourable), we illustrate in Figures 7.9 and 7.10 how the marginal importance densities, projected (marginalized) onto the horizontal and vertical planes, would look like as proposed by the marginal auxiliary particle filter (MAPF) and by the SPF-MPF. It is clear that in this case both MAPF and SPF-MPF generate proposal densities quite proximate of the empirical marginal target, which in turn approximates well the true posterior. Additionally, it is possible to visualize that the SPF-MPF provides a slightly better proposal density in terms of similarity to the target density, which is corroborated by a greater average effective sample size as presented in Table 5. All other marginal particle filters don't generate effective importance densities in terms of approximating either the true posterior or the target density.

For *case 2* (unfavourable), importance sampling fails as exemplified by the projections of the importance density proposed by the MAPF depicted in Figure 7.11. By the same reason explained before, the importance sampling procedure fails to provide a satisfactory filtering measure owing to the errors that affect evaluations of both the marginal target density and the marginal importance density. As a consequence, in this case, any marginal particle filter generates a poor solution, although the MAPF provides a high effective sample size. The SPF-MPF generates a remarkably poor solution for *case 2* because it distributes particles to approximate the true posterior density by design, but must constrain the proposal mixture components to match a very inaccurate empirical target density.

In contrast, in both cases *1* and *2* the SPF-GS proposes a direct filtering density that accurately approximates the true posterior density. The SPF-GS is demonstrated to be insensitive to the issues caused by an observation located relatively far from the initial distribution. These features are quantitatively captured by the performance indexes summarized in Table 5.
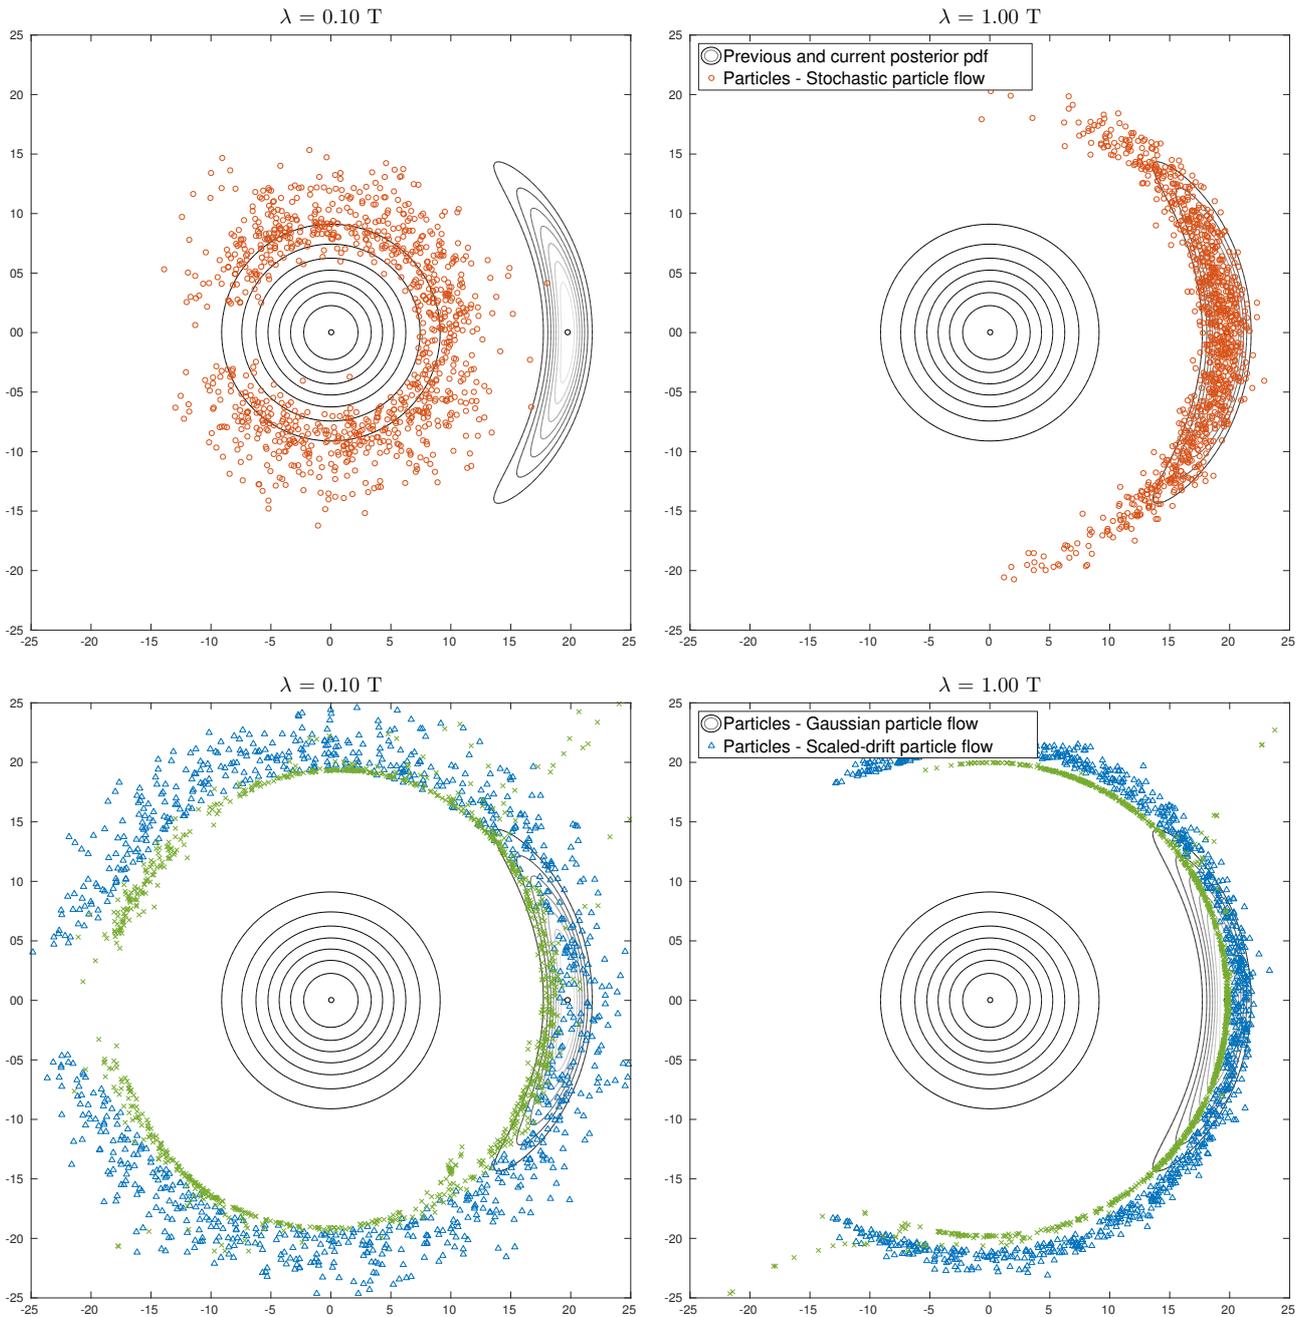
Figure 7.8: Sequence of distributions achieved by propagating particles for the nonlinear, bivariate example

Density	Multimodal, linear		Nonlinear - case 1		Nonlinear - case 2	
	JSD$_{avg}$	ESS$_{avg}$	JSD$_{avg}$	ESS$_{avg}$	JSD$_{avg}$	ESS$_{avg}$
Marginal target	0.0118	-	0.0074	-	0.2444	-
SPF-GS	0.0003	-	0.0133	-	0.0755	-
SPF-MPF	0.0217	93.00%	0.0112	84.01%	0.2746	10.79%
Gaussian particle flow	0.2647	-	0.6563	-	0.5279	-
Scaled-drift particle flow	0.3866	-	0.4962	-	0.5804	-
Marginal BPF	-	-	0.9969	0.37%	0.9998	0.13%
Marginal EPF	-	-	0.3131	27.57%	0.5714	8.83%
Marginal UPF	-	-	0.7753	4.44%	0.8136	2.31%
Marginal APF	-	-	0.0119	81.32%	0.1467	85.11%

Table 5: Comparison of densities for the bivariate examples
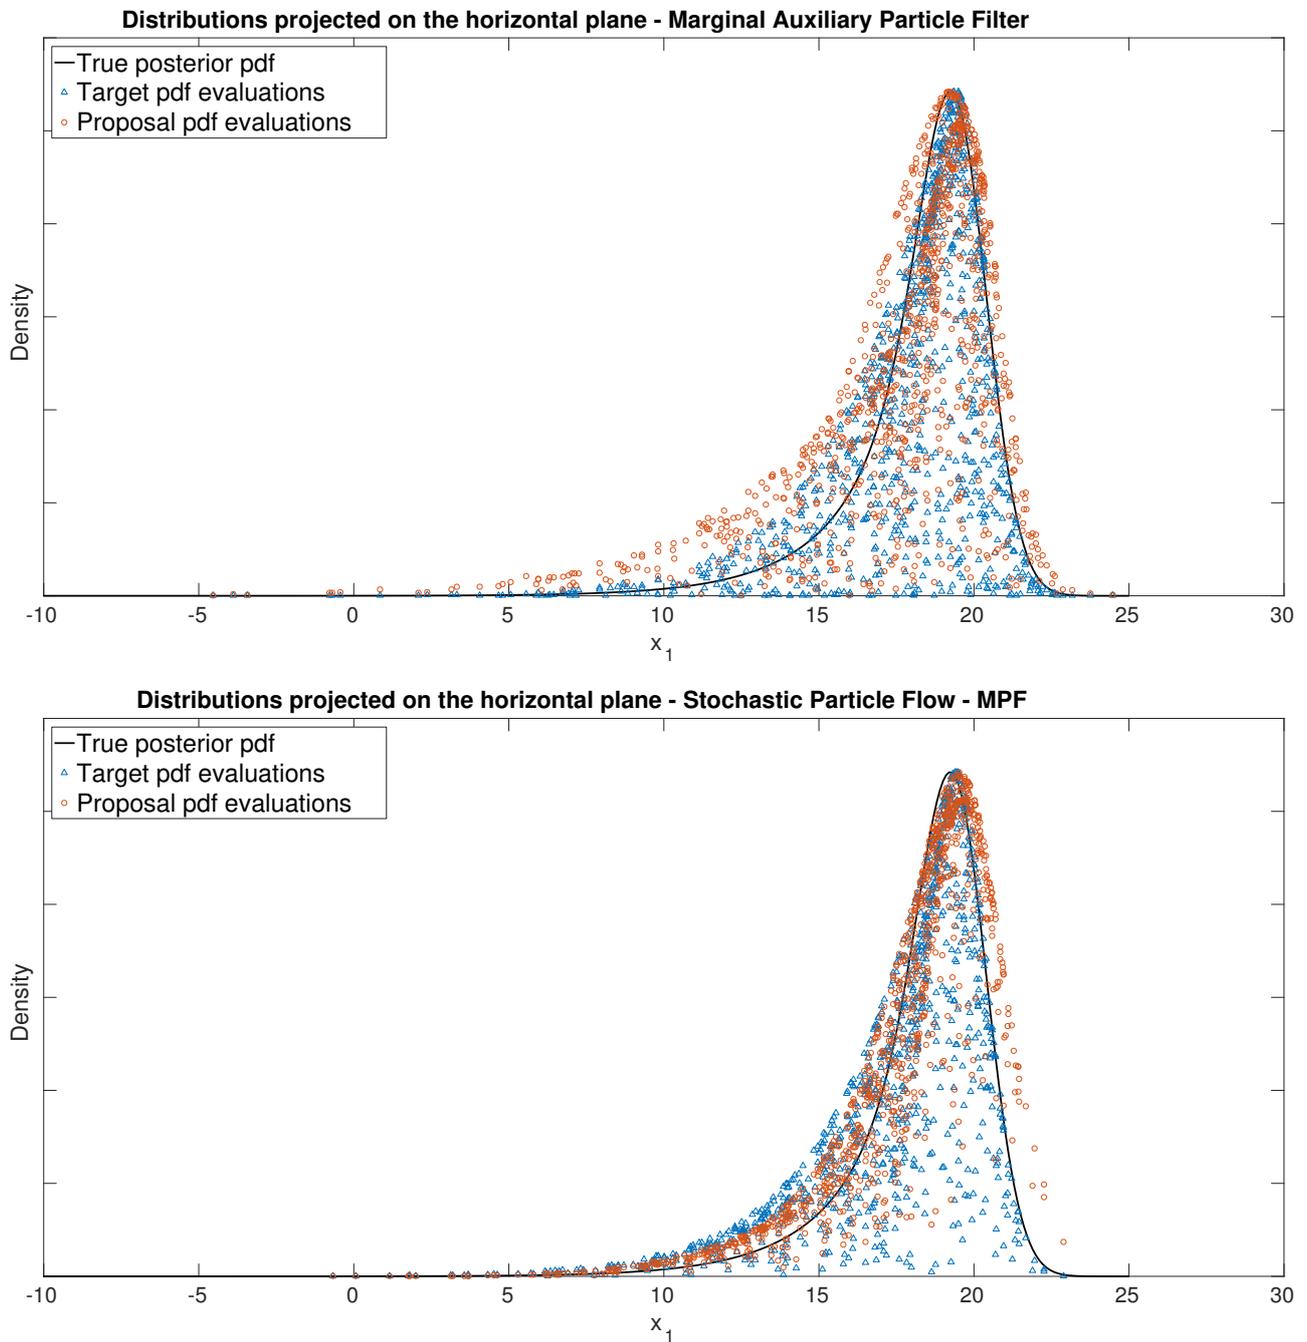
Figure 7.9: Horizontal-plane projection of densities for the MAPF and SPF-MPF (nonlinear, bivariate example)
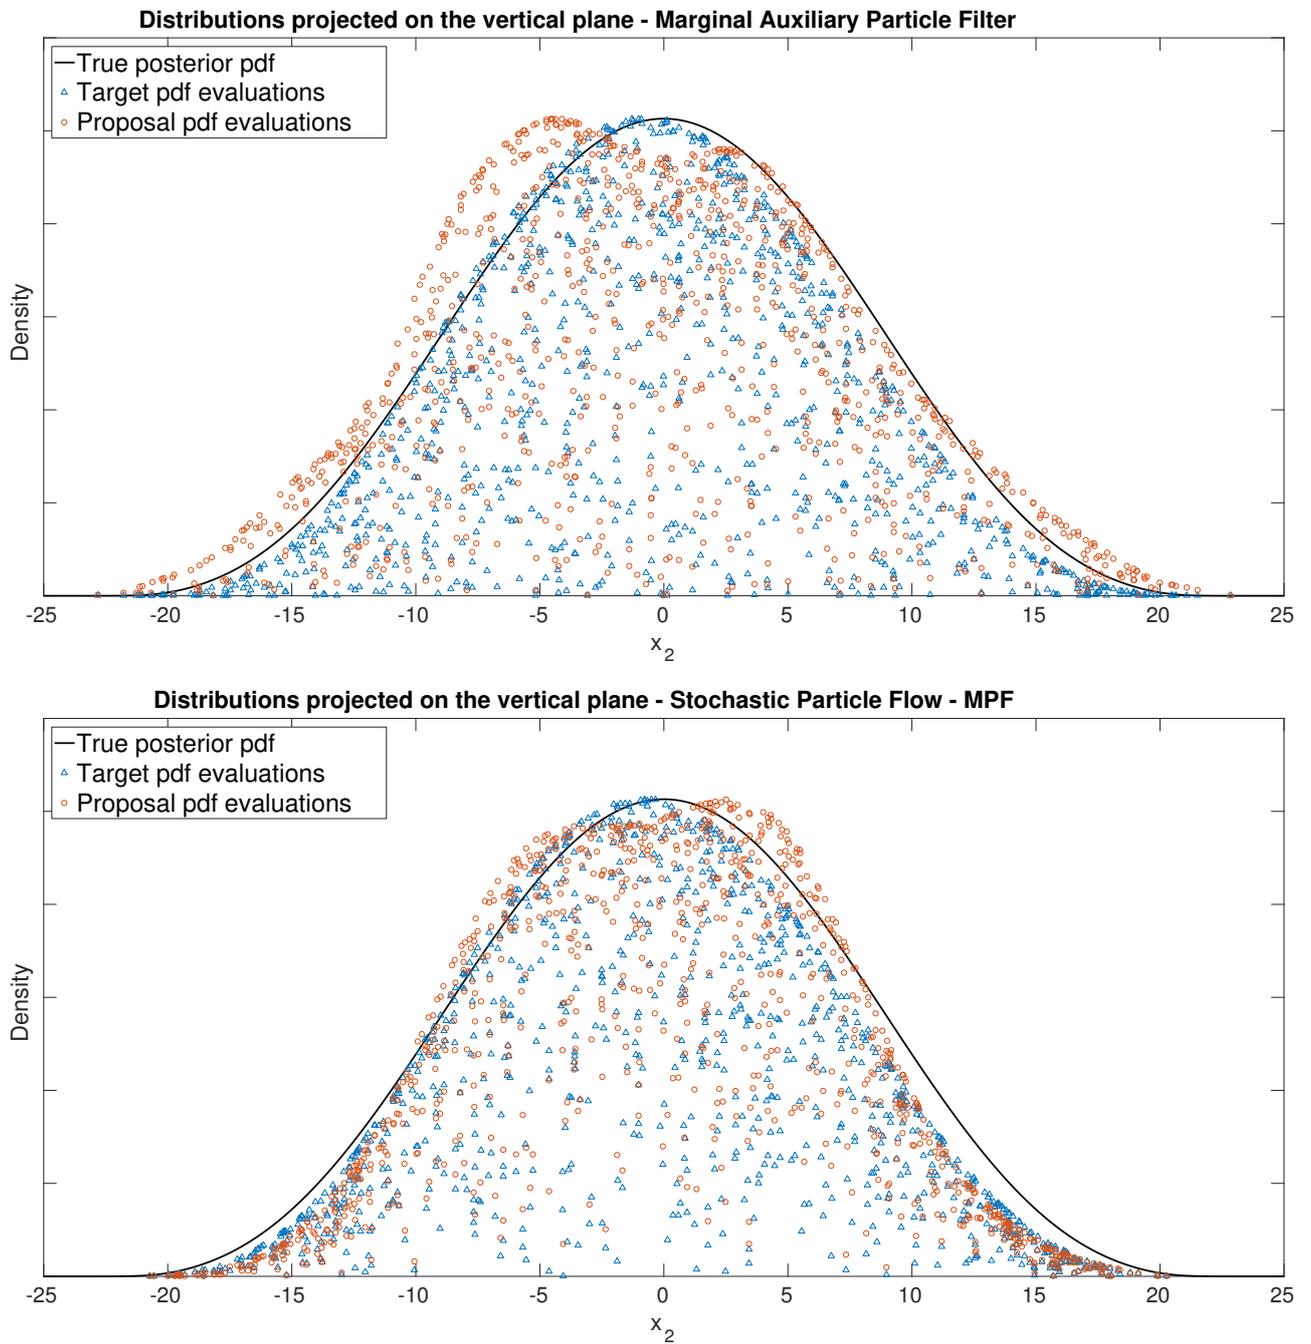
Figure 7.10: Vertical-plane projection of densities for the MAPF and SPF-MPF (nonlinear, bivariate example)
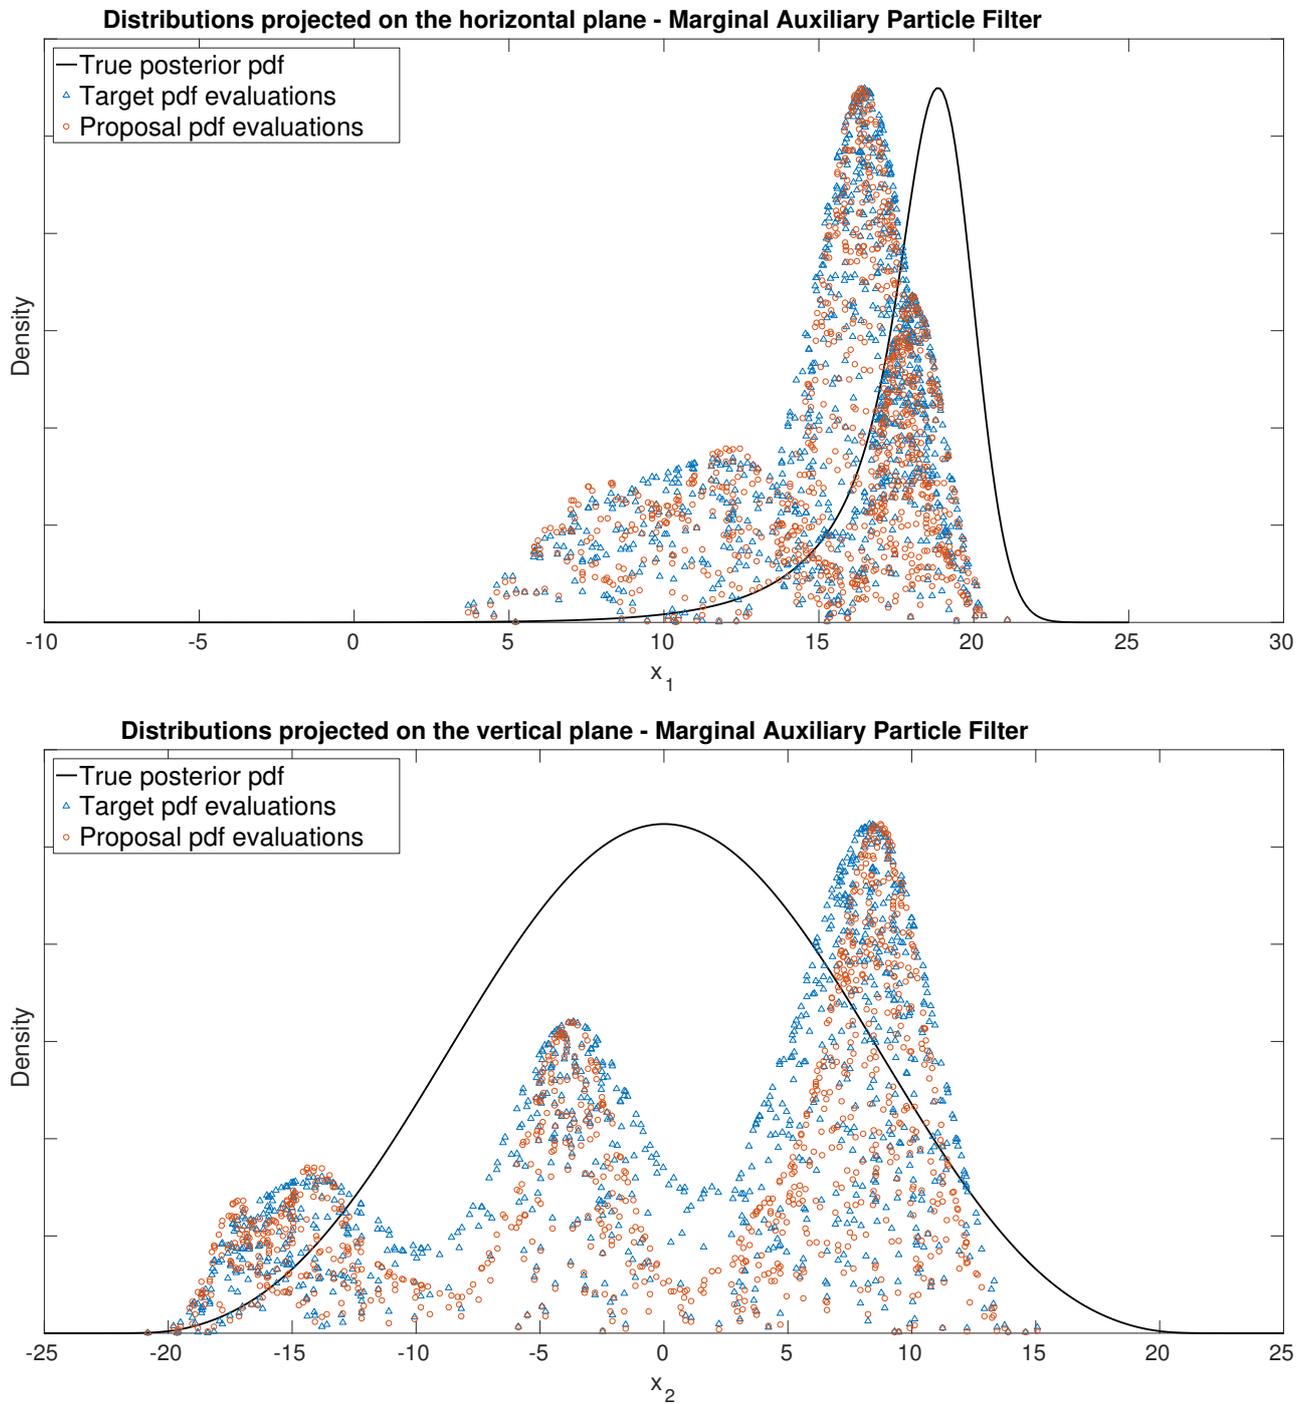
Figure 7.11: Failure of marginal importance sampling for the nonlinear, bivariate example

7.2 Multi-Sensor Bearings-Only Tracking

Estimation in clutter of a target's position and velocity based solely on angular measurements is a relevant problem that finds direct application in airborne radar and sonar in passive listening mode. We propose an example where a single target is observed by a circumferential array of sensors. Each sensor measures the target's bearing with respect to its own position.

In this example, we compare performances of the following filters:

- multi-sensor EKF that performs a series of centralized measurement updates, considering each sensor in sequence;

- Information Matrix Fusion filter (IMF-EKF) [83] that fuses distributed estimates (in parallel) into a global estimate using the Information Matrix form of the EKF;

- bootstrap particle filter (SIR);

- Gaussian particle flow (GPF);

- scaled-drift particle flow (SDPF); and

- stochastic-particle-flow Gaussian sum (SPF-GS).

As mentioned before, in order to work properly, both the GPF and SDPF are implemented based on a companion filter that estimates the state covariance matrix. This is in accordance with implementations suggested by Choi *et al.* [50] and Ding & Coates [51]. In this example we used the multi-sensor EKF as companion filter for both the GPF and SDPF. In contrast, the stochastic particle flow does not require a companion filter.

The bootstrap particle filter (SIR), the GPF and SDPF, and the SPF-GS consider all measurements jointly according to a joint likelihood function described in the next section. Performance is analyzed by computing the root-mean-square error (RMSE) of estimates and the normalized-estimation error squared (NEES) over 100 Monte Carlo runs. All particle-based filters use 200 samples.

7.2.1 A Multi-Sensor Bearings-Only Model

When tracking in clutter based on multiple measurements, the usual treatment rests on the probabilistic data association (PDA) [84]. In the PDA model, a set of m_k valid measurements is received at each time step k and assumed to be generated according to the possibilities: (i) all measurements are false alarms (clutter), (ii) one of the measurements is originated from the target and the remaining are false alarms. Let $\theta_{k,i}$ be the association event that the ith measurement is target-originated. The PDA filter computes the association probabilities $p(\theta_{k,i}|y_{1:k})$ conditional on the set of all received measurements up to time instant k, and calculates the target state posterior density, $p(x_k|y_{1:k})$, by marginalizing the joint density $p(x_k, \theta_{k,1:m_k}|y_{1:k})$ over all possible associations.

In our example a single target is tracked by a set of N_s sensors located along a circumference that encloses the surveillance region, at equally-spaced angular positions. As per the PDA model, one target is known to exist a priori, detected with probability $P_{d,j}$ by the jth sensor, and the number of clutter detections per sensor is Poisson-distributed with mean $\lambda_c \cdot V$, where λ_c is the clutter spatial density and V is the surveillance region's volume. For any given set of N_s sensors, the expected likelihood can be easily obtained by extending the procedure established by Marrs *et al.* [85] to multiple sensors, to give

$$p\left(y_k|x_k, y_{1:k-1}\right) = \prod_{j=1}^{N_s} V^{-m_{k,j}} \frac{(\lambda_c V)^{m_{k,j}} e^{-\lambda_c V}}{m_{k,j}!} \left[\lambda_c \left(1 - P_{d,j}\right) + \sum_{i=1}^{m_{k,j}} P_{d,j} \mathcal{N}\left(y_{k,i(j)}; h_j(x_k), R_{k,j}\right)\right], \quad (7.6)$$

where $m_{k,j}$ is the total number of validated measurements for the jth sensor, $y_{k,i(j)}$ is the ith measurement received by the jth sensor, $h_j(\cdot)$ and $R_{k,j}$ are the observation function and the observation noise variance for the jth sensor, respectively.

On a bidimensional state space, the bearing observations are modeled by

$$h_j(x_k) = \operatorname{atan}\left(\frac{x_k(2) - p_{2,j}}{x_k(1) - p_{1,j}}\right), \quad (7.7)$$

where $p_j = (p_{1,j}, p_{2,j})^T$ are the position coordinates for the jth sensor. We assume a target moving according to the nearly-constant velocity model,

$$x_k = F x_{k-1} + u_k, \ u_k \sim \mathcal{N}(u_k; 0_{n_x}, Q_k), \quad (7.8)$$

where $x_k = (p_{x_1}, p_{x_2}, v_{x_1}, v_{x_2})_k^T$ is the state vector composed of position and velocity in Cartesian coordinates (x_1, x_2), and

$$F = \begin{pmatrix} 1 & 0 & \Delta t & 0 \\ 0 & 1 & 0 & \Delta t \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \tag{7.9}$$

$$Q_k = \begin{pmatrix} \Delta t^3/3 & 0 & \Delta t^2/2 & 0 \\ 0 & \Delta t^3/3 & 0 & \Delta t^2/2 \\ \Delta t^2/2 & 0 & \Delta t & 0 \\ 0 & \Delta t^2/2 & 0 & \Delta t \end{pmatrix} \sigma_q^2. \tag{7.10}$$

The multi-sensor joint likelihood (7.6) is incorporated in the bootstrap (SIR) filter, the GPF and SDPF, and the SPF-GS by considering their filtering densities to target a posterior density involving such joint likelihood. Regarding the implementation of particle flows, specifically for this problem, the GPF and SDPF reinterpret the filtered density empirically as a Gaussian pdf at the end of each iteration in order to avoid exponential growth of the number of mixture components over time. This practical aspect does not affect the SPF-GS, whose filtered density is a mixture composed invariably of N local solutions to the actual posterior pdf, where N is the number of samples (and mixture components).

Generally speaking, multi-sensor bearings-only tracking is a difficult problem to solve when the observation noise has high variance, the probability of detection is relatively low and the probability of having clutter in the surveillance region is not negligible. In this scenario, the difficulty stems from the fact that the joint multi-sensor likelihood (7.6) is a product of mixtures composed of several nonlinear and non-informative likelihood terms: when nonlinearity is pronounced by a high-variance observation noise, the resulting posterior density may not be well expressed by simple parametric densities. In addition, this difficulty is modulated by the amount of information available: the fewer the number of sensors the more difficult to solve the problem. Another aspect that poses additional concern is the system's observability. It is highly dependent on the relative position of a sensor with respect to the target's trajectory, i.e., trajectories radially aligned with a sensor's position provide less information on the target's velocity.

7.2.2 Results

A challenging scenario was set for comparing the filters in order to exacerbate differences of their performances to a noticeable level. In this very difficult scenario, state process noise is assumed with variance scaled by $\sigma_q^2 = 25\,m^2$, observation noise variance is $R_k = 25\,deg^2$, $P_d = 0.50$ and $\lambda_c \times V = 1.00$ false alarm/sensor/scan, for identical sensors. Even though estimation errors generated for such a scenario might not be feasible as an Engineering solution, it is certainly of practical interest to examine how the estimates' errors scale to such extreme scenarios, which might happen in real applications. This problem is particularly interesting because the smaller the number of sensors the more difficult to achieve reasonable estimates since, in this case, the low signal-to-noise ratio would deteriorate inference. For this example, the SPF-GS has been set with time horizon $T = 10\,s$ and integration step size $\Delta\lambda = 1\,s$.

The resulting track of an exemplar run is shown in Figure 7.12. The track initiation is based on an overdetermined triangulation of measurements for the first two steps. No gating has been performed, i.e., no preprocessing to discard measurements that fall outside a high-confidence region of each sensor. Figure 7.12 depicts a successful tracking of the target despite the difficult scenario.
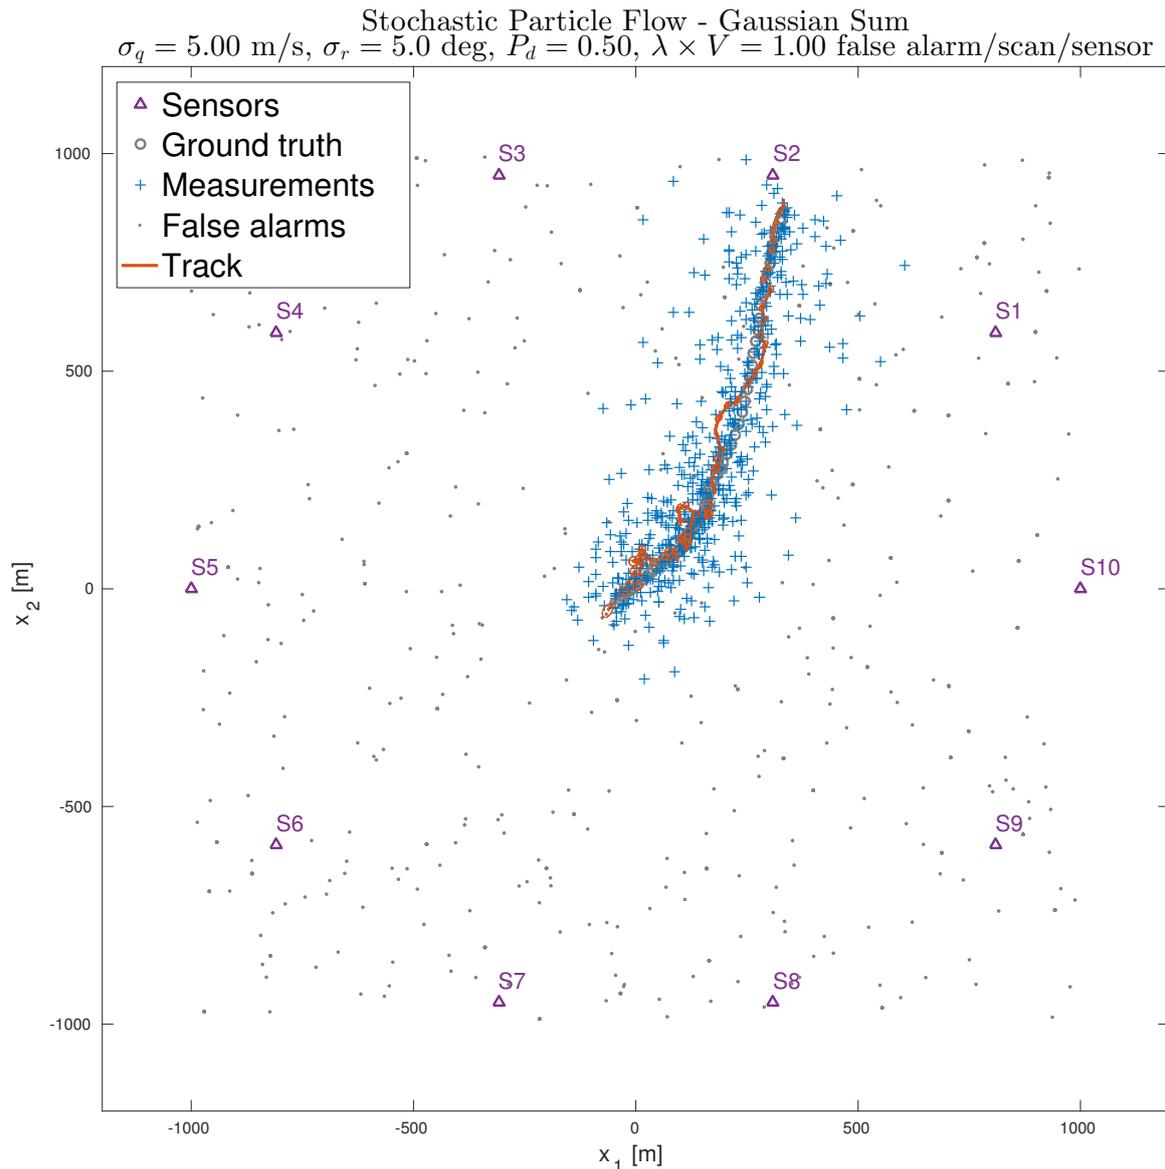
Figure 7.12: Illustration of bearing-only multi-sensor tracking

The resulting root-mean-square error (RMSE), normalized-estimation error squared (NEES), and average computation time (per time step) of all filters for different numbers of sensors are shown in Figures 7.13, 7.14 and 7.15 respectively. The following important aspects can be observed from Figure 7.13:

- Somewhat counterintuitively, the multi-sensor serial EKF provides better estimates than that of Information Matrix Fusion EKF (IMF-EKF), both in terms of precision (RMSE) and "consistency" (or credibility[10] as indicated by NEES).

- The scaled-drift particle flow (SDPF) shows remarkably poor performance.

- The bootstrap particle filter (SIR) provides mediocre performance, eventually becoming better than the IMF-EKF as the number of sensors increases.

- Accuracies shown by the Gaussian particle flow (GPF) and the serial EKF are commensurate and similar to that of stochastic particle flow (SPF-GS) when the number of sensors is high.

- As expected, the overall estimation accuracy is improved as the number of sensors is increased.

---

[10]As presented in [86], the normalized-estimation error squared (NEES) is the simplest metric that indicates an estimator's credibility.

- The SPF-GS provides the most accurate estimates in difficult scenarios, i.e., when the number of sensors is small.

- Estimation by the SPF-GS is more consistent (or credible), which is denoted by an NEES closer to one ($\log_{10} NEES \to 0$) from above.

It is worth commenting on the results comparing GPF and SPF-GS. Specifically for this problem, when the number of sensors is sufficiently high, the GPF provides estimates as accurate as those of SPF-GS at a slightly lower computational cost. It is also remakable the successful synergy between the GPF and its companion filter, a multi-sensor serial EKF that provides the covariance estimates. However, it is difficult to justify the calculated NEES for original particle flows since their first and second moment estimates are underpinned by distinct filtering methods. On the other hand, stochastic particle flow (SPF-GS) provides fairly accurate state estimates and securely constitutes the most credible estimator among all evaluated filters.
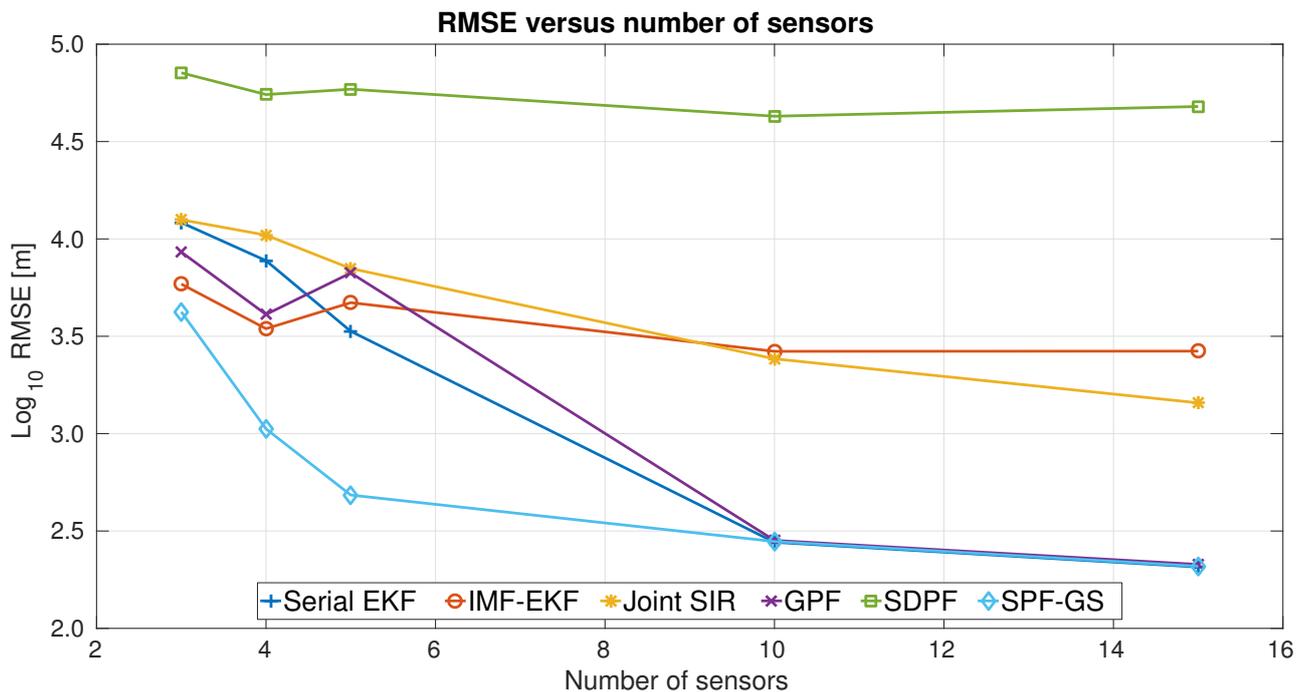
Figure 7.13: RMSE for the multi-sensor bearing-only tracking example
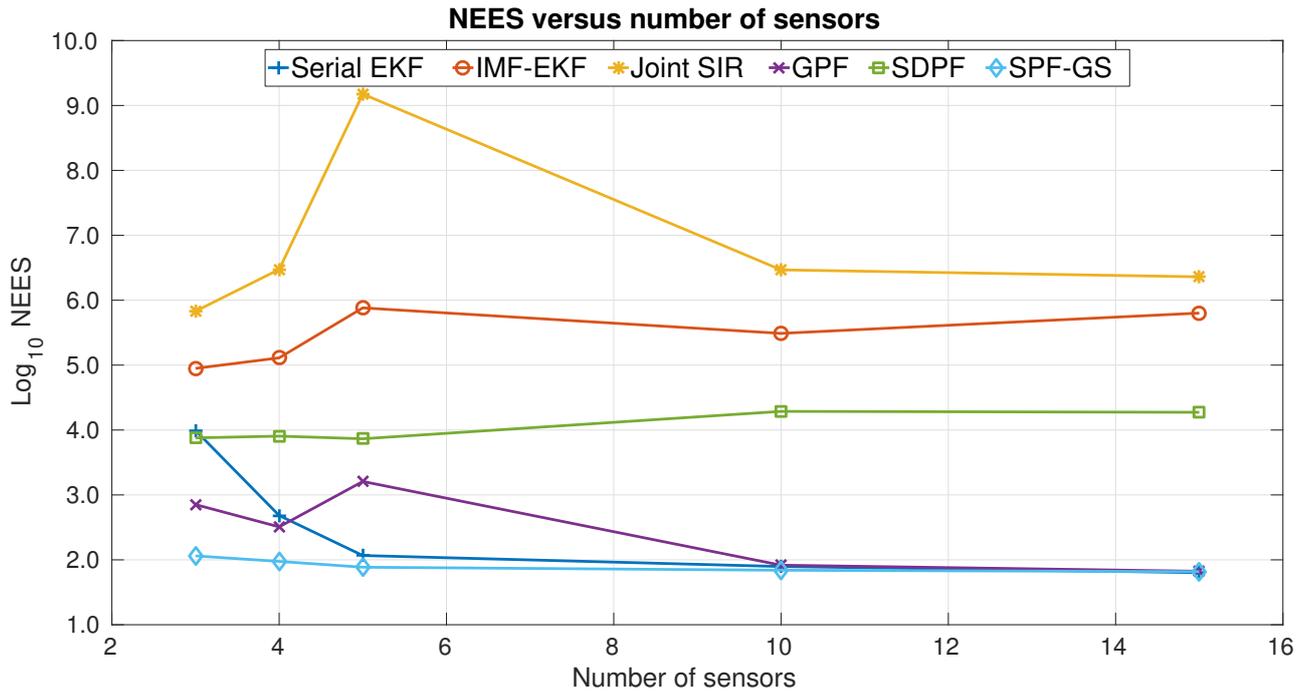
Figure 7.14: NEES for the multi-sensor bearing-only tracking example
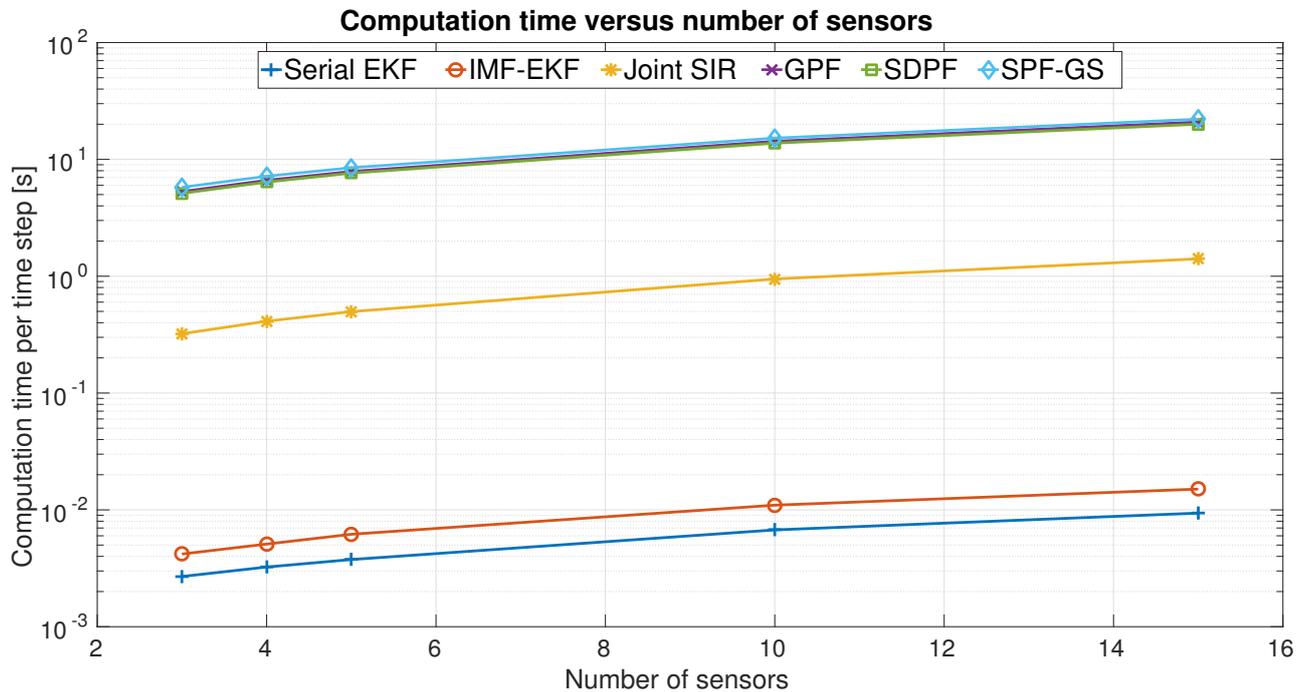
Figure 7.15: Average computation time for the multi-sensor bearing-only tracking example

7.3 Convoy Tracking

Tracking multiple objects in clutter is as challenging as important for real applications. In the multi-target tracking standard methods, the most common treatment assumes the targets' states to be independent so that the joint probability density is the product of their marginal densities. While this assumption is fairly reasonable for applications where objects are far apart most of the time, the same cannot be stated for cases where objects are in proximity for a considerable part of time. This latter cases elicit tracking all targets jointly in the hope of implicitly capturing dependencies between targets. When targets are tracked jointly the problem's

dimensionality scales with the number of targets.

To illustrate this situation, we propose an example of a convoy of vehicles that are forced into mutual proximity when trafficking on a road . The vehicles have explicit interactions as each driver aims driving at the maximum allowed speed unless there is another vehicle immediately in front at a slower speed. This scenario demands care for a safety distance. The intent of the tracker is then to provide the best estimate of each vehicle on a convoy, given a set of non-identified measurements corrupted by noise and possible false alarms (clutter) reported by a position sensor.

We compare performances of the following filters:

- Joint Probabilistic Data Association (JPDA) filter [84];

- Global Nearest Neighbor Data Association (GNN) filter;

- JPDA with a Gaussian mixture per target (JPDA-GM), applying mixture reduction [87];

- Coupled Probabilistic Data Association (CPDA) filter [88];

- multi-target bootstrap particle filter (joint SIR) based on the description by Blom & Bloem [89];

- Gaussian particle flow (GPF);

- scaled-drift particle flow (SDPF); and

- stochastic-particle-flow Gaussian sum (SPF-GS).

The GPF and the SDPF rely on a companion filter to estimate state covariance matrices correctly, according to implementation guidelines by Choi *et al.* [50] and Ding & Coates [51]. In this example, we used the CPDA [88] as companion filter for the original particle flows (GPF, SDPF). In contrast, the SPF-GS does not require a companion filter.

The CPDA, the joint bootstrap particle filter (SIR), the GPF and SDPF, and the SPF-GS consider all targets' states jointly, as a single high-dimensional state. In contrast, the classical multi-target filters track targets separately, where each target's state is described by the nearly-constant velocity model. Performance is analyzed by computing the root-mean-square error (RMSE) of estimates and the normalized-estimation error squared (NEES) over 100 Monte Carlo runs. The particle filter (SIR), the original particle flows and stochastic particle flow use 200 samples.

7.3.1 The Intelligent Driver Model

The Intelligent Driver Model (IDM) [90] is a model[11] used in Traffic Engineering to simulate phenomena such as congestion and to analyze the traffic behaviour as a response to changes in the transport system. Because the interaction between vehicles is explicitly taken into account by the IDM, tracking based on it involves consideration of the joint state of multiple targets. Even though the IDM establishes an empirical description of traffic for multiple vehicles, it has not been previously used in the context of multi-target tracking. We propose a stochastic version of the IDM and discretize it in order to make it compatible with multi-target trackers formulated on the joint state space.

The IDM describes the dynamics of vehicles in traffic, in terms of positions and velocities, incorporating the interaction between each vehicle and the vehicle directly in front. Provided a vehicle indexed as α with length l_α, the dynamics of its position x_α and velocity v_α are given by the following (continuous-time) stochastic differential equations:

$$dx_\alpha = v_\alpha dt, \tag{7.11}$$

$$dv_\alpha = \underbrace{a \left[1 - \left(\frac{v_\alpha}{v_0} \right)^\delta \right] dt}_{\dot{v}_\alpha^{\text{free road}}} - \underbrace{a \left[\frac{\bar{s}}{s_\alpha} \right]^2 dt}_{\dot{v}_\alpha^{\text{interaction}}} + dw_t, \tag{7.12}$$

where $\{w_t\}_{t \geq 0}$ is a Wiener process, $s_\alpha = x_{\alpha-1} - x_\alpha - l_\alpha$ is the net distance between vehicles, the approaching rate is given by $\Delta v_\alpha = v_{\alpha-1} - v_\alpha$, and $\bar{s} = \bar{s}(v_\alpha, \Delta v_\alpha)$ is the expected distance defined as

$$\bar{s}(v_\alpha, \Delta v_\alpha) = s_0 + v_\alpha T_h + \frac{v_\alpha \Delta v_\alpha}{2\sqrt{a \cdot b}}. \tag{7.13}$$

---

[11]In its simplest form, the IDM is focused on the interaction of vehicles moving along a single-carriageway road. More complex variants exist to model overtaking, for example, and consider factors such as the politeness of the driver.

The model dynamics is such that when a vehicle is travelling on a free road it will predominantly accelerate according to $\dot{v}_\alpha^{\text{free road}}$ up to the maximum allowed speed v_0, whereas when it approaches another vehicle immediately in front, the decrement in acceleration according to $\dot{v}_\alpha^{\text{interaction}}$ becomes relevant to maintain a safe-time headway T_h and to avoid approaching closer than the minimum safe distance s_0. The IDM parameters are summarized in the following table.

Parameter	Description
a	nominal maximum acceleration
b	comfortable braking decceleration
δ	acceleration exponent (driver dynamics)
v_0	free-road desired velocity
s_0	minimum allowed distance between vehicles
T_h	safe-time headway
$\alpha - 1$	index of the vehicle direcly in front

In order to use the stochastic IDM as the state process for a multi-target tracker, its continuous-time equations are discretized by a first-order approximation (Markov random field). This assumes that the state's derivative with respect to time is linear in time between two subsequent measurements, but the interactions between non-adjacent vehicles are negligible when compared to the interactions between adjacent vehicles. The discretized version of the stochastic IDM is presented in the Appendix D.

7.3.2   The Multi-Target Joint Likelihood Function

The joint multi-target filters extend the joint probabilistic data association (JPDA) [84] framework for situations where the targets' states are not mutually independent conditioned on the past observations. This formulation has been first proposed as the JPDA Coupled filter (JPDAC) [84] and further generalized by Blom & Bloem [88, 89] who consider the measurement-to-target associations implicitly.

In the JPDA model, a set of N_m valid measurements is received at each time step k and assumed to be generated according to the possibilities: (i) each of the measurements may be originated from each target, considering all possible associations, (ii) a measurement not originated from any target is due to a false alarm (clutter). These possibilities are exhaustive such that a measurement can have only one source, and at most one of the validated measurements can originate from a target.

Let $\phi_{k,i} \in \{0, 1, \ldots, N_m\}$ be an association event that maps each target $i \in \{1, \ldots, N_t\}$ to the measurement indexed as $\phi_{k,i}$, where $\phi_{k,i} = 0$ means that no measurement is associated to the ith target. The Coupled JPDA filter computes the joint association probabilities $p(\phi_{k,1:N_t}|y_{1:k})$ conditional on the set of all received measurements up to time instant k, and calculates the joint state posterior density, $p(x_{k,1:N_t}|y_{1:k})$, by marginalizing $p(x_{k,1:N_t}, \phi_{k,1:N_t}|y_{1:k})$ over all possible joint associations.

In the JPDA framework, N_t targets are known to exist a priori, detected with probability P_d by a single sensor; the number of clutter detections is Poisson-distributed with mean $\lambda_c \cdot V$, where λ_c is the clutter spatial density and V is the surveillance region's volume; the location of each clutter detection is independently distributed according to a spatial density $\eta_c(y)$; and the likelihood function of the jth measurement being originated from the ith detected target is $p(y_{k,j}|x_{k,i})$. Denoting the joint multi-target state as $x_{k,1:N_t}$ and the joint observation as $y_{k,1:N_m}$, the joint likelihood can be either obtained as in [89] or by a formulation equivalent to the Coupled JPDA as

$$p\left(y_{k,1:N_m}|x_{k,1:N_t}\right)$$
$$= \frac{\left[\prod_{j=1}^{N_m} \eta_c\left(y_{k,j}\right)\right]}{N_t!} \sum_{N_d=0}^{N_t} \frac{(\lambda_c V)^{N_m-N_d} e^{-\lambda_c V}}{(N_m-N_d)!} P_d^{N_d} \left(1-P_d\right)^{N_t-N_d} \sum_{\phi_{k,1:N_t}|N_d} \prod_{i=1}^{N_t} \frac{p\left(y_{k,\phi_{k,i}}|x_{k,i}\right)}{\eta_c\left(y_{k,\phi_{k,i}}\right)}. \quad (7.14)$$

The joint state vector $x_{k,1:N_t} = (p_1, \ldots, p_{N_t}, v_1, \ldots, v_{N_t})_k^T$ is composed of position and velocity of all vehicles in the convoy, and the joint observation $y_{k,1:N_m} = (y_1, \ldots, y_{N_m})_k^T$ contains position measurements of all targets and possible false alarms obtained at a given time instant k.

The joint bootstrap particle filter, the GPF and SDPF, and the SPF-GS consider their filtering densities to target a joint posterior density incorporating the joint multi-target likelihood function (7.14). Regarding the implementation of original particle flows, the GPF and SDPF reinterpret the filtered density empirically as a Gaussian pdf at the end of each iteration in order to avoid exponential growth of the number of mixture components over time. In contrast, this practical aspect does not affect the stochastic particle flow, whose filtered density is a mixture composed of a fixed number of local solutions to the actual posterior pdf.

7.3.3   Results

We simulated the trajectories of vehicles on a ring road by integrating the continuous-time stochastic IDM over 60 seconds with the parameters presented as follows. The convoy was set to start from rest with the vehicles initially positioned apart, led by a truck so that the queue of cars is slowed down and forced into mutual proximity. The minimum allowed distance between vehicles was set to be exaggeratedly small ($s_0 = 0.5\,m$) to induce the model to control the distance between cars mainly based on the safe-time headway T_h. In this case, the safe-time headway indirectly determines the desired distance between vehicles, which is denoted as *target distance* in the table below. At the final steady state, the net speed of the convoy is dominated by the free-road speed of the truck, which motivates the safe-time headway being computed based on $v_{0,\text{truck}}$.

Parameter	car	truck
a	$0.5\,m/s^2$	$0.4\,m/s^2$
b	$1.5\,m/s^2$	$1.2\,m/s^2$
δ	4	4
v_0	$15\,m/s$	$10\,m/s$
s_0	$0.5\,m$	$0.5\,m$
T_h	$\frac{\{\text{target distance (m)}\}}{v_{0,\text{truck}}}$	-
l_α	$5\,m$	$20\,m$

The joint state process covariance matrix is assumed as scaled by $\sigma_q^2 = 0.0625\ (m/s)^2$, each position observation has variance $\sigma_r^2 = 4\,m^2$, $P_d = 0.80$ and $\lambda_c \times V = 0.01$ false alarm/scan, and the surveillance region's "volume", V, is in fact the length covered by a confidence region ($\approx 99.73\%$) that contains all the vehicles. Proposing a method to effectively initiate tracks was out of the example's scope, thus track initiation was considered to be ideal, i.e., the initial position and velocity of the targets are known with initial uncertainty scaled by the observation noise. The stochastic particle flow has been set with time horizon $T = 15\,s$ and integration step size $\Delta\lambda = 1\,s$.

Figure 7.20 shows two frames of an exemplar run, demonstrating the situation where a queue of cars is slowed down by a truck, forcing them into proximity. The non-filled rectangles depicted in Figure 7.20 denote the position estimates provided by the filter applied for that run. Interactions between vehicles in the convoy, due to proximity, can be well perceived by position estimates of the exemplar run as shown in Figure 7.19.

The resulting root-mean-square error (RMSE), normalized-estimation error squared (NEES), and average computation time (per time step) of all filters, for different numbers of vehicles, and target distance between vehicles $d = 10\,m$, are shown in Figures 7.16, 7.17 and 7.18 respectively. RMSE and NEES were computed over 100 Monte Carlo runs, with the particle-based filters using 200 samples. The following important aspects can be noted from Figure 7.16:

- In general, trackers that estimate on the joint $2N_t$-dimensional state space clearly outperform the classical multi-target trackers (JPDA, GNN and JPDA-GM), both in terms of precision (RMSE) and credibility (NEES).

- Among the classical multi-target trackers, the Global Nearest Neighbor (GNN) association filter is the one that provides the most accurate estimates. This can be explained by an increasingly detrimental effect of the association uncertainty on estimation, which is more prominent in the JPDA and less prominent on the GNN filter.

- Estimation errors committed by the multi-target particle filter (joint SIR) grow exponentially with the number of state-space dimensions ($2N_t$), as expected, due to the curse of dimensionality.

- The Coupled PDA (CPDA) and the SPF-GS present commensurate root-mean-square errors, suggesting that most of their accuracy gain originates from tracking on the joint $2N_t$-dimensional state space and accounting for inherent dependencies between targets.

- The original particle flows (GPF, SDPF) show significant values of NEES. Most likely this is because the evaluated implementation of these filters cannot provide reliable estimates for state covariance matrices and depend on a dissimilar companion filter to work around it, which affects consistency of their estimates.

- The SPF-GS provides overall higher estimation accuracy (RMSE) and consistency (NEES), with low sensitivity to increasing the problem's number of dimensions.

Based on the results for this example, two important remarks are worth making. Firstly, the results show a notable performance improvement with adoption of filtering on the joint $2N_t$-dimensional state space: about 30-fold improvement in estimation precision (RMSE) and nearly 10-fold in credibility (NEES). The fact that the Coupled PDA performs as well as the SPF-GS suggests that modeling inherent dependencies between targets and filtering on the joint space provides most of the performance gain. Secondly, the example not only illustrates well the curse of dimensionality for the multi-target particle filter (joint SIR), but also corroborates the success of principled choices made in the SPF's formulation in order to avoid degeneracy in high-dimensional problems. This latter observation becomes clear when we realize that the performance indexes for stochastic particle flow scale gently with the number of dimensions.

Additionally, it is also worth noting that the original particle flows provide relatively accurate estimates, scaling well with the number of dimensions. For the evaluated implementation, both particle flow filters (GPF, SDPF) rely on covariance matrices estimated by the CPDA as a companion filter. Due to this fact, calculated NEES for these filters is not reliable since their first and second moment estimates are underpinned by distinct filtering methods. This does not disqualify the original particle filters per se since observed characteristics are probably due to the implementation settings. Under these circumstances, actual consistency (credibility) of their estimates cannot be quantified and, ultimately, evokes the question about the extent to which the success of the adopted implementation is due to the companion filter.
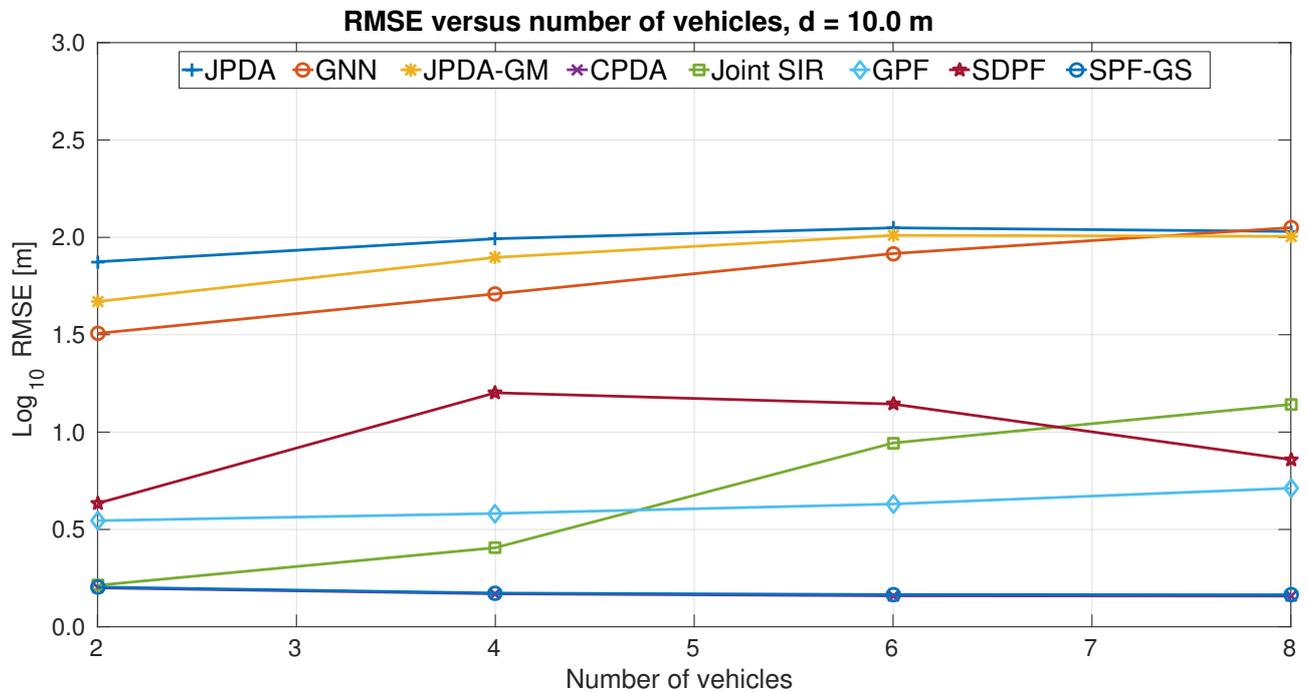
Figure 7.16: RMSE for the convoy tracking example
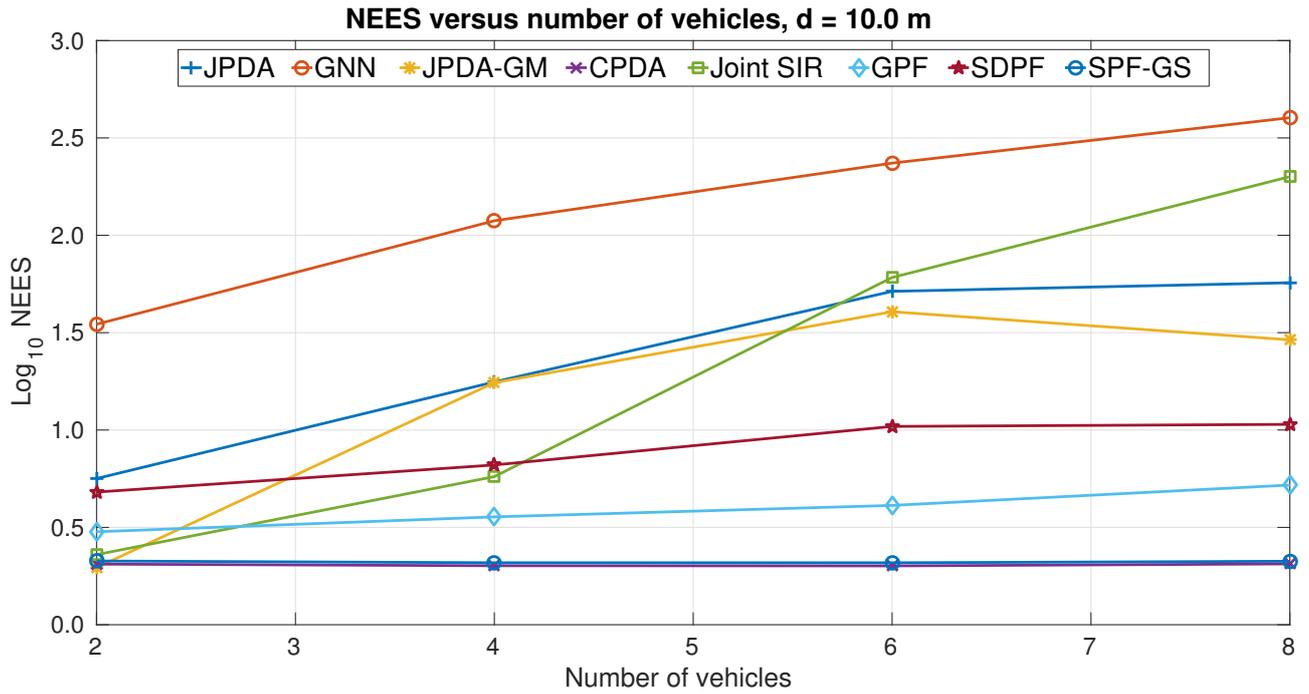
Figure 7.17: NEES for the convoy tracking example
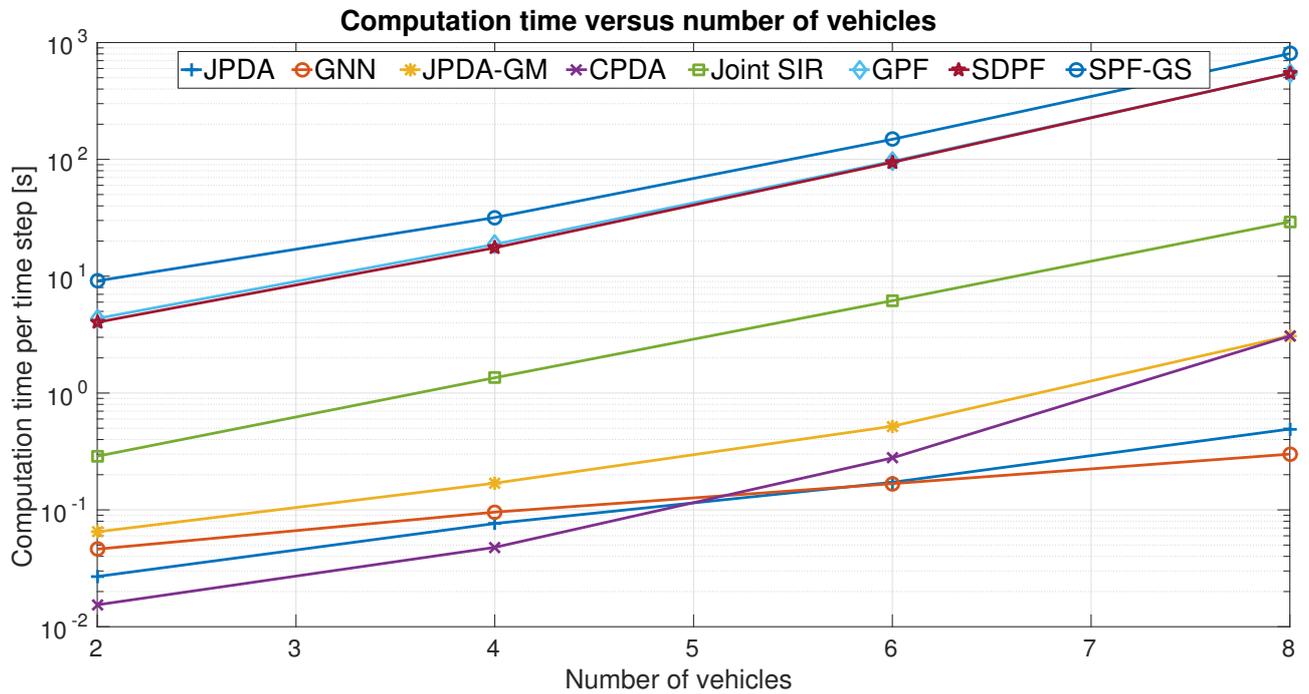
Figure 7.18: Average computation time for the convoy tracking example
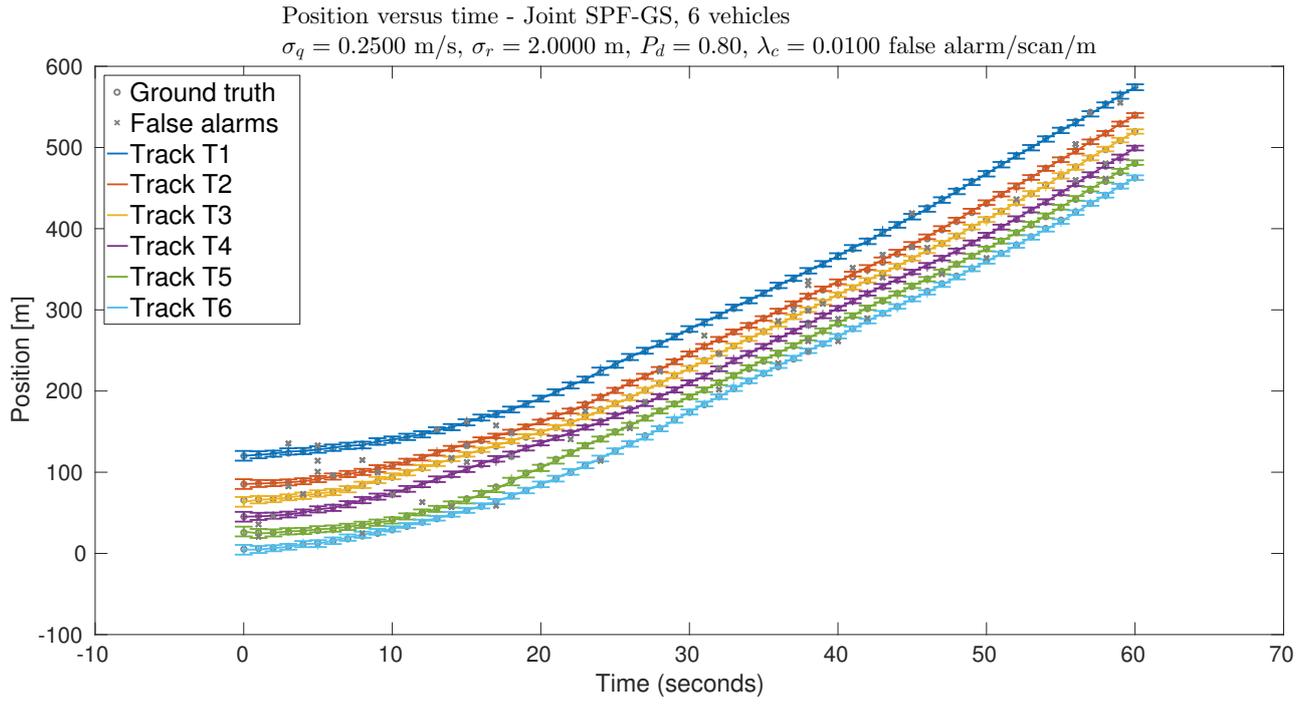
Figure 7.19: Position estimates for an exemplar run of the convoy tracking

Convoy tracking, Joint SPF-GS, step $k = 15$
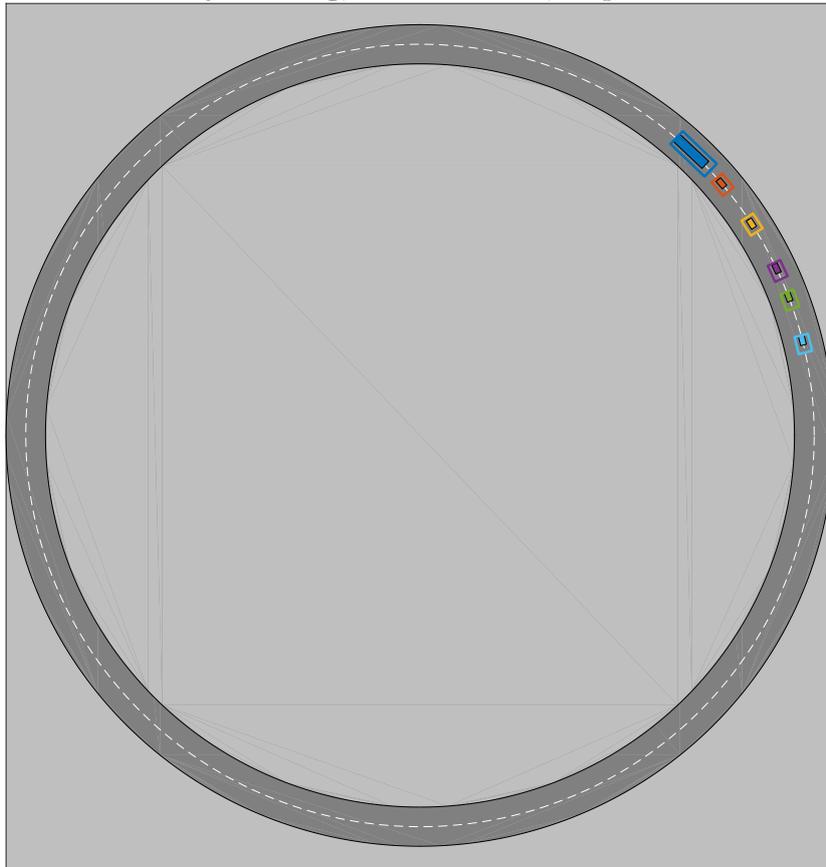
Convoy tracking, Joint SPF-GS, step $k = 45$
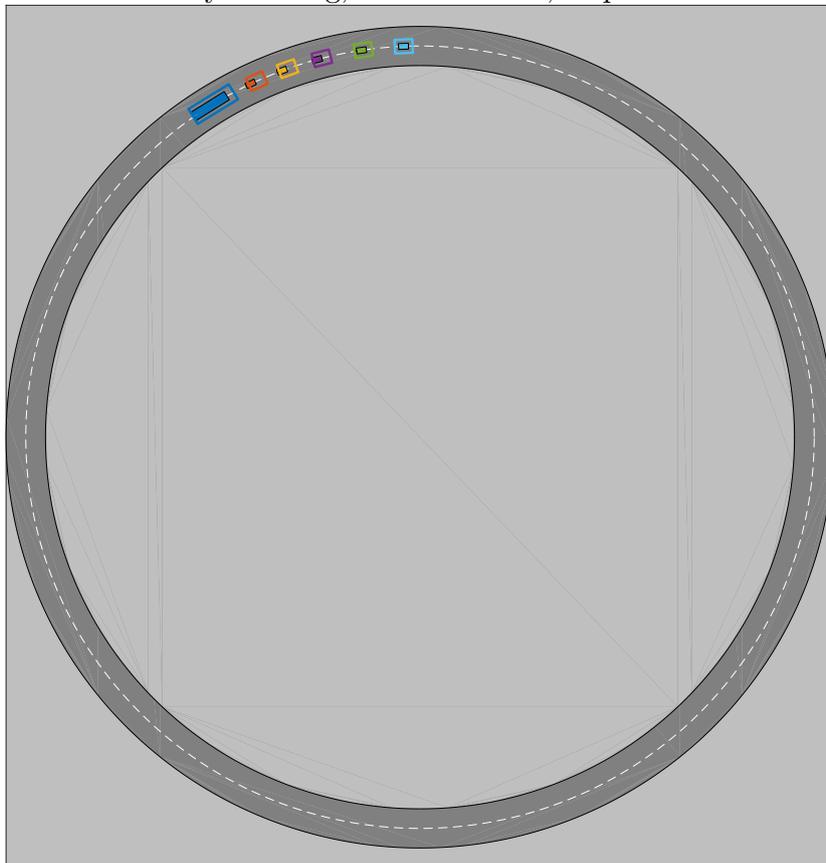
Figure 7.20: Illustration of convoy tracking on a ring road

7.4 Inference on Large Spatial Sensor Networks

In this section we consider the problem proposed by Septier & Peters [56], posed to address inference of physical quantities of complex phenomena from a collection of noisy measurements obtained by a large network of spatially distributed sensors. According to Septier & Peters [56], a large number of applications could adopt such sensor networks to make inferences related to complex phenomena. Applications include environmental monitoring, weather forecasting etc. In this framework, a fusion center would regularly receive observations from sensors set up as a grid, which monitor a time-varying physical phenomenon presenting spatially diverse attributes such as pressure, temperature, concentrations of substance, radiation levels, seismic activity etc. Upon fusing the observations, the solution to the problem consists of estimating the phenomenon state at the current time instant at each of the sensor's positions. The problem becomes particularly challenging as the number of sensors in the grid increases, since solving the problem then demands efficient algorithms for inference in high-dimensions.

The physical phenomenon is modeled as a time-varying spatially-dependent continuous process defined over a two-dimensional space which is observed sequentially in time by a 2D spatial grid of $n_x = N_s$ sensors, where n_x is the state dimension. At time instant k, each sensor independently produces a noisy measurement of an attribute of interest about the phenomenon at its specific location, giving $y_{k,j} | x_k \sim p(y_k^{(j)} | x_k)$, $\forall j = 1, \ldots, N_s$. Based on the historic set of observations $Y_{1:k} := \{Y_{k'} : k' = 1, \ldots, k\}$, where $Y_k := \{y_{k,j} : j = 1, \ldots, N_s\}$, one is required to estimate, at time k, the state of the physical phenomenon $x_k \in \mathbb{R}^{n_x}$ across the locations of all sensors in the grid. The state process that models the time-varying physical phenomenon is considered to follow a transition multivariate Generalized Hyperbolic (GH) density as

$$p_t(x_k | x_{k-1}) \propto \frac{\mathbb{K}_{c_1 - n_x/2}\left[\sqrt{(c_2 + Q(x_k, x_{k-1}))(c_3 + \gamma^T \Sigma^{-1} \gamma)}\right]}{\left(\sqrt{(c_2 + Q(x_k, x_{k-1}))(c_3 + \gamma^T \Sigma^{-1} \gamma)}\right)^{n_x/2 - c_1}} \cdot e^{(x_k - \alpha x_{k-1})\Sigma^{-1}\gamma} \tag{7.15}$$

where $Q(x_k, x_{k-1}) = (x_k - \alpha x_{k-1})^T \Sigma^{-1}(x_k - \alpha x_{k-1})$, $\alpha \in \mathbb{R}$ is the location constant, and $\mathbb{K}_{c_1}[\cdot]$ denotes the modified Bessel function of the second kind, of order c_1. The parameters c_1, c_2, and c_3 are scalar values that determine the shape of the distribution, $\Sigma \in \mathbb{R}^{n_x \times n_x}$ is the dispersion matrix, and the vector $\gamma \in \mathbb{R}^{n_x}$ is the skewness parameter. The choice of transition density in 7.15 can account for heavy-tailed and asymmetric data [56], which is beneficial when modeling physical process with extremal behavior. In special cases, the transition density becomes the normal, normal inverse Gaussian, skewed-t, and other densities. To generate the prior distribution at the first time step, we take $p_x(x_0) = p_t(x_0 | x_{-1} = 0)$.

The dispersion matrix is positive definite and is defined such that the degree of spatial correlation across sites of a physical phenomenon is given in terms of the separation between locations as

$$[\Sigma]_{ij} = \alpha_0 \exp\left[-\beta^{-1} \|\mathcal{S}_i - \mathcal{S}_j\|_2^2\right] + \alpha_1 \delta_{ij}, \tag{7.16}$$

where $\|\cdot\|_2$ is the L2-norm, δ_{ij} the Kronecker symbol, $\alpha_0, \alpha_1 \in \mathbb{R}$, and $\mathcal{S}_m \in \mathbb{R}^2$ are the physical locations of the sensors for $m = 1, \ldots, N_s$.

For this example, we compare the performance of following filters:

- the Sequential Importance Resampling filter (SIR);

- the block SIR filter, which partitions the state space into separate subspaces of smaller dimensions (blocks of 4 sensors each) and run a particle filter on each subspace [23];

- the Sequential manifold Metropolis-Adjusted Algorithm (SmMALA) filter [10, 56];

- the Sequential manifold Hamiltonian Monte Carlo (SmHMC) filter [10, 56];

- the Stochastic Particle Flow, Gaussian sum (SPF-GS).

These filters are compared for two cases:

- Gaussian state process and Gaussian likelihood;

- Skewed-t state process and Poisson-distributed observations.

Note that SmMALA and SmHMC are chosen because we perceive they constitute two of the best sequential MCMC filters that exist. It is essential to justify why we have not included annealed importance sampling (AIS) [4] and SMC samplers [5] in our comparisons. Although these techniques are built on fast mixing Markov

chains, their filtering procedures operate on the joint space along the complete path of samples. This makes them highly prone to the curse of dimensionality for long-time horizons. A careful explanation of this issue can be found in [80] (and Section 6.1), but the key point is that as dimension increases, it becomes increasingly important to avoid consideration of the path. SMC samplers could be adapted to filter in the marginal space (at the cost of up to $\mathcal{O}(N^3)$ evaluations) but this would require an ad-hoc approximation to the target pdf. Developing the approximations that would be needed to enable SMC samplers to be applied to the problem we consider is not the focus of our paper. We have compared performance of SPF-GS with state-of-the-art techniques (SmMALA and SmHMC) that consider the marginal distribution as well as techniques (SIR and block SIR) that we perceive to be good examples of the class of algorithms, which also includes AIS and SMC samplers, that consider the joint space of the complete sample path and can be applied, without modifications, to the problem we are focused on solving.

These filters are compared for two cases:

- Gaussian state process and Gaussian likelihood;

- Skewed-t state process and Poisson-distributed observations.

The implementations of SmMALA and SmHMC used are exactly as made available by Septier & Peters [56] on their web page[12]. These algorithms make use of a refinement step of the state [56] at the current time, performed with a random partitioning of size 4, by using the empirical approximation of the previous posterior distribution as proposal distribution.

7.4.1 Results

Gaussian State Process and Gaussian Likelihood

We first consider a trivial special case of the GH family as the transition density, namely the multivariate normal distribution. In this setting, each sensor measures the attribute of a physical process with some Gaussian noise. The resulting model is given by

$$
\begin{aligned}
p_t(\mathrm{x}_k|\mathrm{x}_{k-1}) &= \mathcal{N}(\mathrm{x}_k;\ \alpha\mathrm{x}_{k-1}, \Sigma), \\
p(\mathrm{y}_k|\mathrm{x}_k) &= \mathcal{N}(\mathrm{y}_k;\quad \mathrm{x}_k,\quad R),
\end{aligned}
\tag{7.17}
$$

where $R = \sigma_y^2 \mathbb{I}_{n_x}$, and the following model parameters are used: $\alpha = 0.9$, $\sigma_y^2 = 2$, and with the dispersion matrix constructed using $\alpha_0 = 3$, $\alpha_1 = 0.01$, $\beta = 20$. When performing the comparison between the filters, we use as a reference the estimates provided by the Kalman filter (which is optimal in this special case). For the SmMALA and SmHMC, the proposed metric tensor is given by $G(\mathrm{x}_k) = R^{-1} + \Sigma^{-1}$.

The methodology presented in [56] for performance evaluation was reproduced. Instead of presenting the baseline (provided by the Kalman filter) performance explicitly, the accuracy of each filter is evaluated with respect to the baseline performance, i.e., we observe the difference between the log mean square error for each filter against the log mean square error for the Kalman filter. The relative log mean square error, log normalized-estimation error squared (NEES) and average computation time (per time step) of all filters, for different numbers of sensors in the grid, are shown in Figures 7.21, 7.22, 7.23 respectively. The mean square error of the estimates were computed over 100 Monte Carlo runs, with the particle-based filters using 200 samples. The step size adopted for the sequential MCMC filters is $\Delta\lambda = 0.5$ and the number of steps obtained as $L = L_0 + N$, where $L_0 = 0.2N$ is the number of steps for the burn-in phase, whereas for the SPF-GS we applied the empirical rules for time horizon and step step size as presented in subsection 4.6. From Figures 7.21, 7.22, 7.23, we can note that for this example:

- The performances of all sequential MCMC filters are in accordance with the results shown in [56].

- The SPF-GS outperforms all other filters in terms of mean square error and normalized-estimation error squared.

- The SPF-GS demands the highest computational effort when the number of sensors is small, and its computing time scales better than that of SmMALA and worse than that of SmHMC for higher dimensions.

---

[12]Code available at http://pagesperso.telecom-lille.fr/septier/software.html.
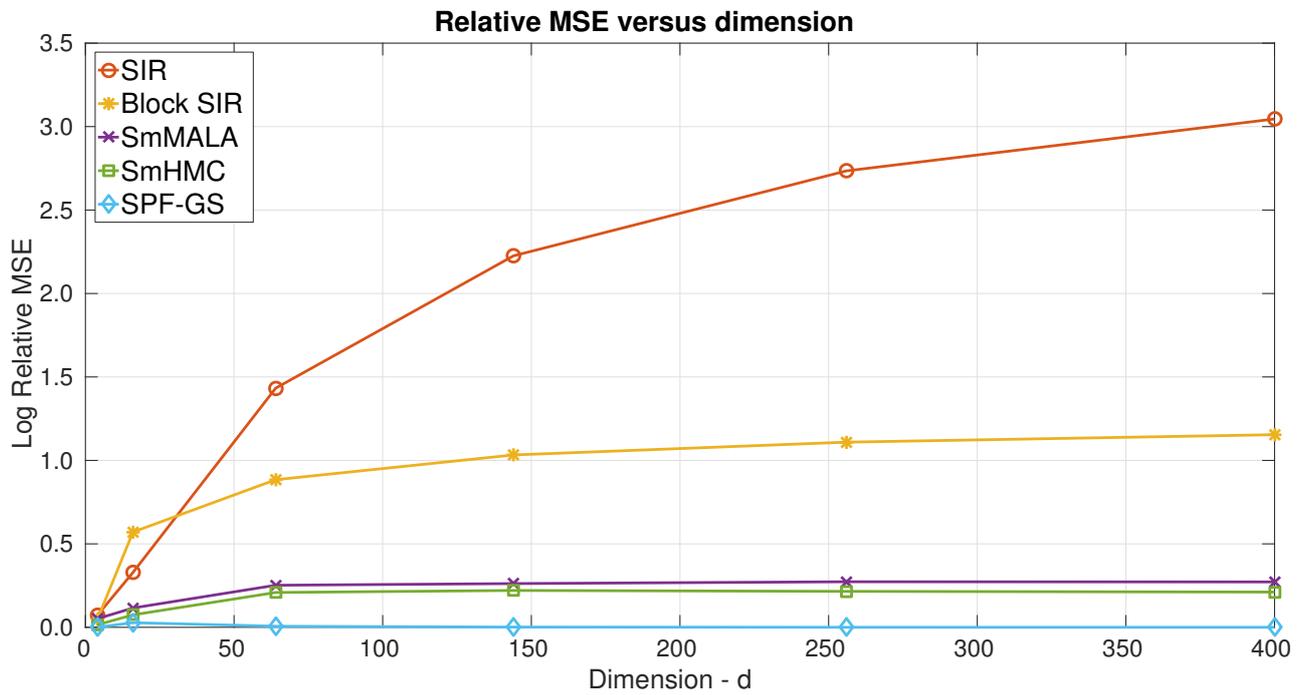
Figure 7.21: Relative MSE for the linear, Gaussian sensor network example
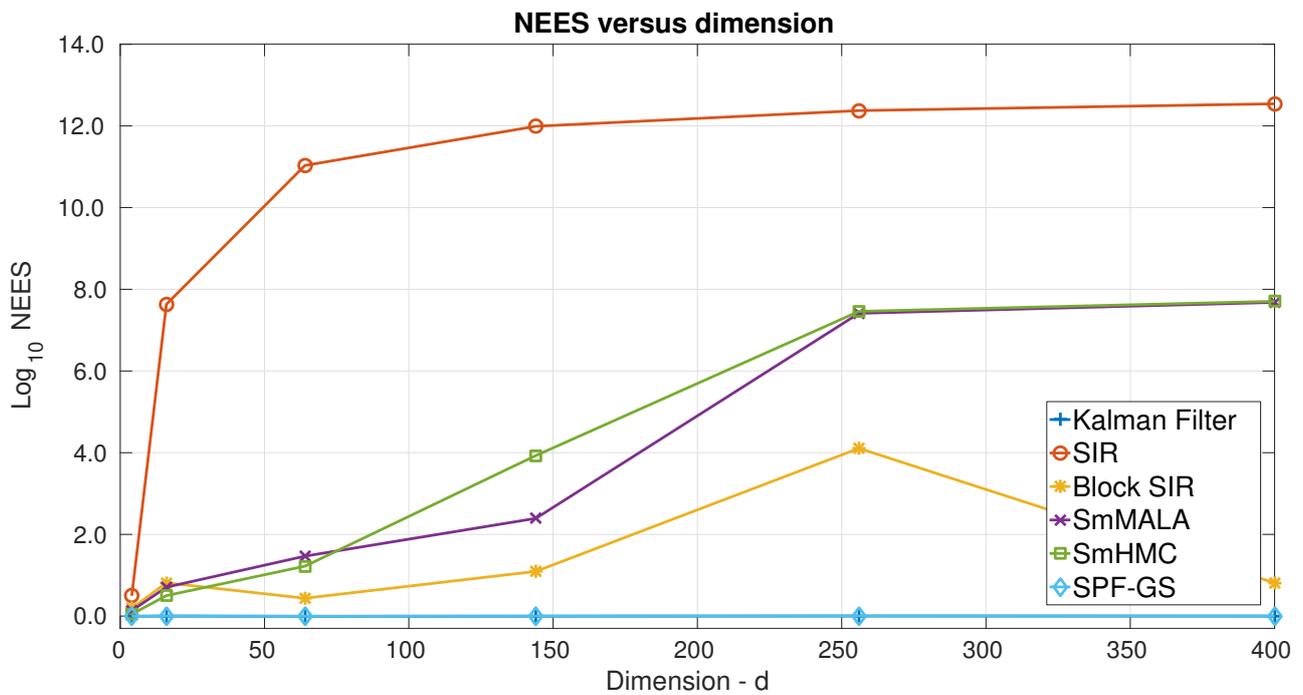
Figure 7.22: NEES for the linear, Gaussian sensor network example
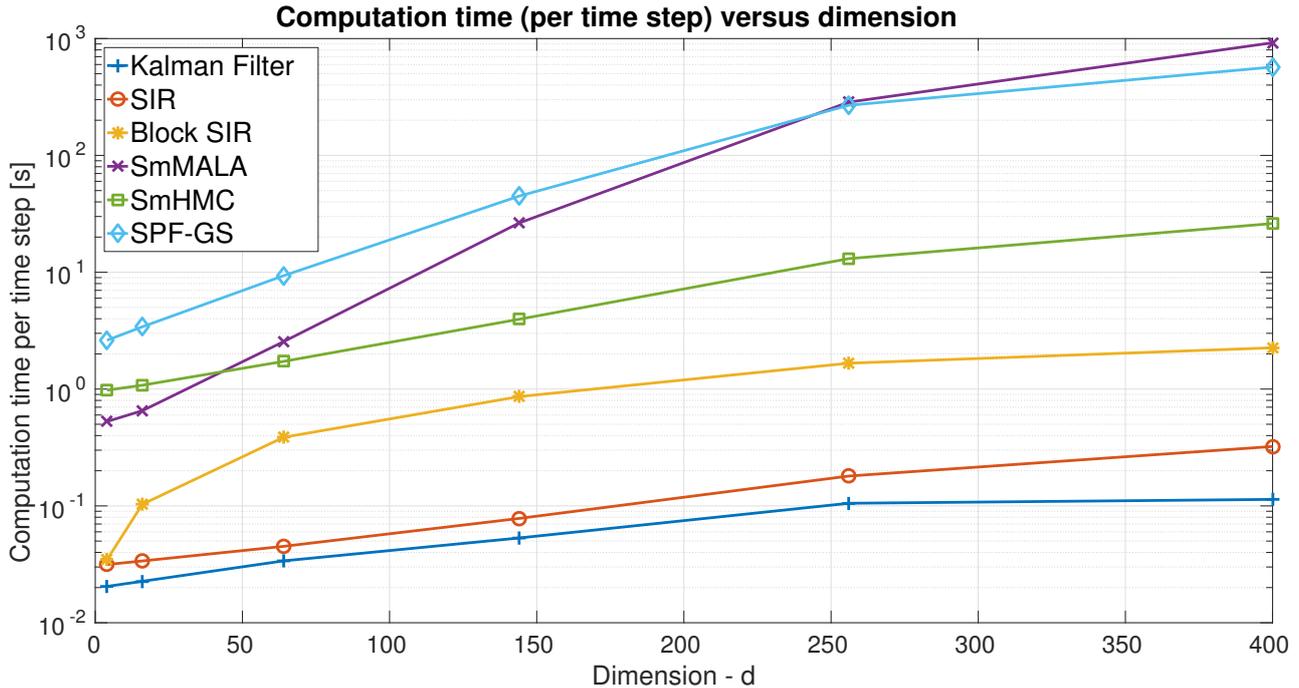
Figure 7.23: Average computation time for the linear, Gaussian sensor network example

The results suggest that SPF-GS is the most accurate among the compared methods. It is unambiguous that the consistency of the estimates produced by the SPF-GS is better than that by all other filters: the NEES for SPF-GS is very close to one (from above) for all evaluated dimensions. It may be well possible that better results could be achieved for the sequential MCMC filters by carefully choosing the step size and number of steps. This would slightly change the performance indexes. However, we strongly believe that such changes would not be enough to modify the conclusions. Additional tests with the SPF-GS demonstrate the computational cost can be directly traded with estimation accuracy. By allowing more steps, via the criteria of Section 4.6, SPF-GS is very close to optimal as depicted in Figure 7.21. On the other hand, if one fixes the step size to $\Delta\lambda = 0.5$ and number of steps to $L = 20$, the results are as presented in Figures 7.21 and 7.23, where it is becomes clear computational cost alleviation at the expense of slightly degrading mean square error.

Skewed-t State Process with Poisson-Distributed Observations

A high-dimensional non-linear and non-Gaussian state-space model is now studied. The transition kernel is proposed to be a multivariate GH skewed-t density described by (7.15) with $c_1 = -\nu/2$, $c_2 = \nu$ and $c_3 = 0$. The likelihood function is assumed to be a Poisson distribution, highly non-linear on the state x_k, given by

$$p(\mathrm{y}_k|\mathrm{x}_k) = \prod_{j=1}^{n_x} \frac{\lambda_j(\mathrm{x}_{j,k})^{\mathrm{y}_{j,k}}}{\mathrm{y}_{j,k}!} e^{-\lambda_j(\mathrm{x}_{j,k})}, \quad \lambda_j(\mathrm{x}_{j,k}) = m_1 e^{m_2 \mathrm{x}_{j,k}}, \tag{7.18}$$

such that $\mathrm{x}_k = (\mathrm{x}_{1,k}, \ldots, \mathrm{x}_{n_x,k})^T$ and $\mathrm{y}_k = (\mathrm{y}_{1,k}, \ldots, \mathrm{y}_{n_x,k})^T$. The model parameters are fixed as $m_1 = 1$, $m_2 = 1/3$, $\alpha = 0.9$, $\nu = 7$, $\gamma = 0.3\mathbb{1}_{n_x \times 1}$ with the dispersion matrix constructed using $\alpha_0 = 3$, $\alpha_1 = 0.01$ and $\beta = 20$. The implied prior density is not log concave, and thus the tensor metric that defines the diffusion coefficient for the sequential MCMC algorithms is modified according to

$$G(\mathrm{x}_k) = \Lambda(\mathrm{x}_k) + \tilde{\Sigma}^{-1}, \tag{7.19}$$

where

$$\Lambda(\mathrm{x}_k) = m_1 m_2^2 \begin{pmatrix} e^{m_2 \mathrm{x}_{1,k}} & & 0 \\ & \ddots & \\ 0 & & e^{m_2 \mathrm{x}_{n_x,k}} \end{pmatrix}, \tag{7.20}$$

$$\tilde{\Sigma} = \frac{\nu}{\nu-2}\Sigma + \frac{\nu^2}{(2\nu-8)(1/2-1)^2}\gamma\gamma^{\mathrm{T}}, \quad \nu > 4. \tag{7.21}$$
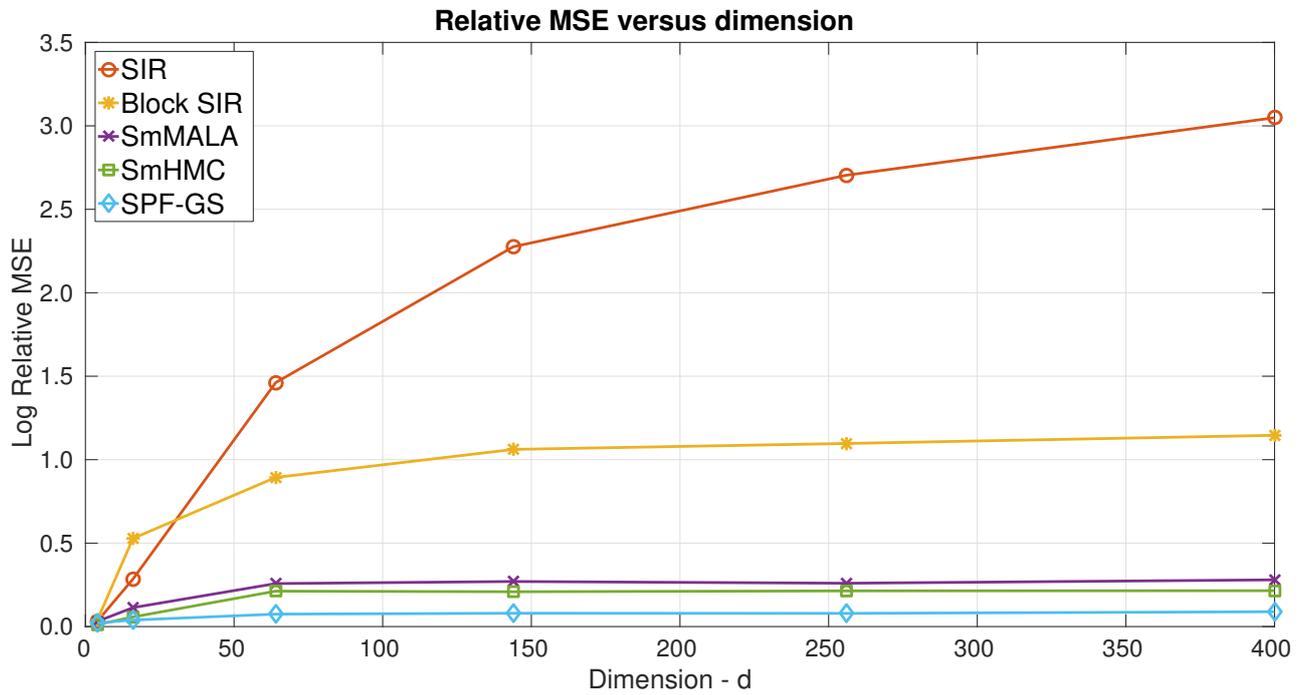
Figure 7.24: Relative MSE for the linear, Gaussian example (SPF-GS with $\Delta\lambda = 0.5$, $L = 20$)
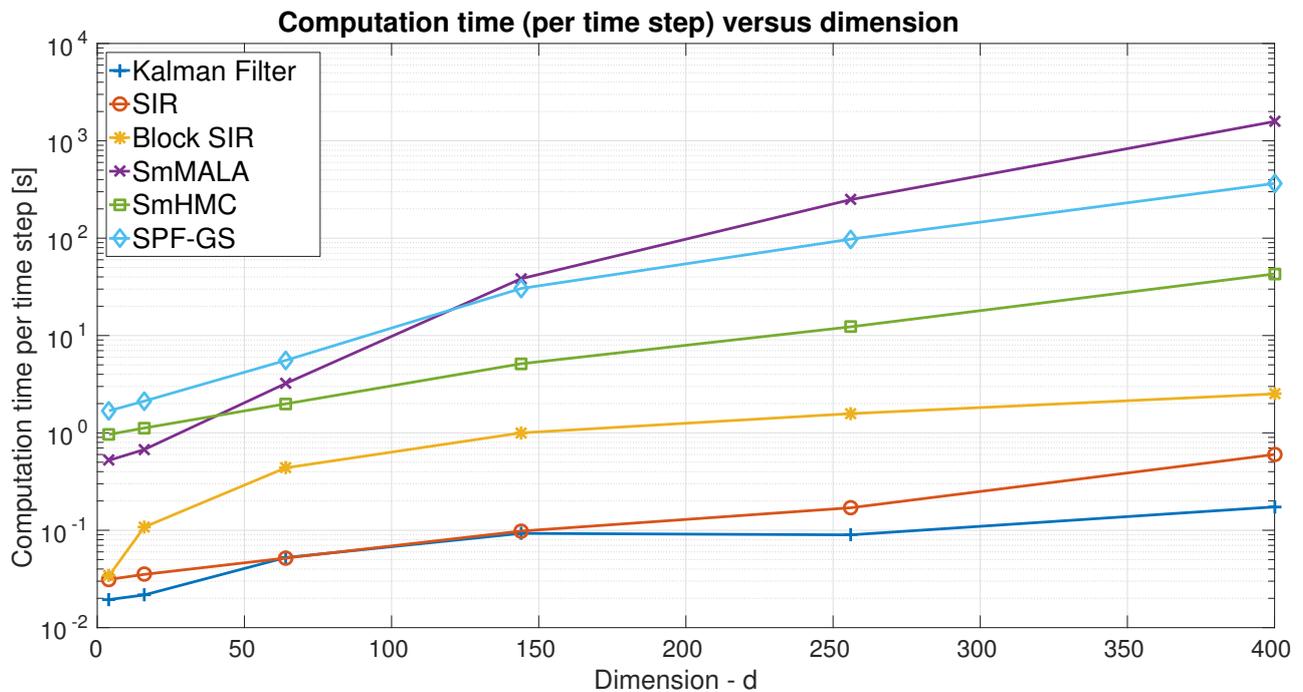
Figure 7.25: Computation time for the linear, Gaussian example (SPF-GS with $\Delta\lambda = 0.5$, $L = 20$)

For this problem, the local linearization of stochastic particle flow around a probability mass x_l, analogous to (5.9) and (5.10), is given by

$$\frac{1}{2}D\nabla_{\mathrm{x}}\log\tilde{\pi}(\mathrm{x}) \approx C(\mathrm{x}_l)\cdot\mathrm{x} + c(\mathrm{x}_l),$$

$$C(\mathrm{x}_l) = \frac{1}{2}D\left(-\mathrm{V}(\mathrm{x}_l)^{-1} - P_{k|k-1}^{-1}\right), \tag{7.22}$$

$$c(\mathrm{x}_l) = \frac{1}{2}D\left(m_2\mathrm{y}_k - \mathrm{v}(\mathrm{x}_l) + \mathrm{V}(\mathrm{x}_l)^{-1}\mathrm{x}_l + P_{k|k-1}^{-1}\alpha\mu_{m,k-1}\right), \tag{7.23}$$

where $\mathrm{V}(\mathrm{x}_l)^{-1} = \Lambda(\mathrm{x}_l)$, $\mathrm{v}(\mathrm{x}_l) = m_1m_2e^{m_2\mathrm{x}_l}$, and

$$P_{k|k-1} = \mathbb{E}\left[(\mathrm{x}_{k|k-1} - \alpha\mu_{m,k-1})(\mathrm{x}_{k|k-1} - \alpha\mu_{m,k-1})^T\right]. \tag{7.24}$$

Once again we follow the methodology presented in [56] for performance evaluation. The log root-mean-square error, log normalized-estimation error squared (NEES) and average computation time (per time step) of all filters, for different numbers of sensors in the grid, are shown in Figures 7.26, 7.27, and 7.28 respectively. The mean square error of the estimates were computed over 100 Monte Carlo runs, with the particle-based filters using 200 samples. The step size adopted for the sequential MCMC filters is $\Delta\lambda = 0.5$ and the number of steps computed as $L = L_0 + N$, where $L_0 = 0.2N$ is the number of steps for the burn-in phase, whereas for the SPF-GS we applied the empirical rules for time horizon and step-step size as presented in Subsection 4.6. In Figure 7.29 we illustrate the posterior means and variances of the state across the sensors grid ($n_x = 400$) at different time steps, for all evaluated filters. From Figures 7.26, 7.27, and 7.28, we can note that

- The performances of all sequential MCMC filters are in accordance with the results shown in [56].

- The SPF-GS presents performance commensurate to that of the SmHMC filter in terms of root-mean-square error, outperforming all other filters.

- The SPF-GS outperforms all other filters in terms normalized-estimation error squared.

- The SPF-GS demands a computational effort higher than that of SmMALA when the number of sensors is small, but its computing time scales better than that of SmMALA and similarly to that of SmHMC for higher dimensions.

For this example, the results indicate that the SPF-GS is as accurate as the SmHMC filter in general, and that SPF-GS is the most accurate method in high dimensional problems. Once again it is clear that the consistency (credibility) of estimates by the SPF-GS is higher than that by all other filters: the NEES is very close to one (from above) for all evaluated dimensions. It is worth noting that, to keep the computational cost for the SPF-GS competitive, the total number of steps was limited. This did result in some loss of accuracy but the implication is perhaps that improved performance is possible if sufficient computational resources are available. The associated trade-off is sufficiently complex to form a hard obstacle against systematic solutions. However, the techniques presented in Section 4.6 proved sufficient to generate the results presented here.

8 Conclusions and Future Work

This paper builds on concepts, such as continuous-time filtering and sequential Monte Carlo methods, being intensely studied by the research community. The paper provides a description of these concepts that aims to draw out some key insights from the theoretical research into sequential Monte Carlo filtering while also responding to the significant empirical challenges encountered when applying existing and emerging tools to difficult real-world problems. More specifically, the paper is part of a growing body of research which aims to apply concepts from the sequential Monte Carlo community to problems involving high-dimensional spaces. This broader body of research and this paper in particular is motivated by the increasing demand for more statistically efficient methods to solve difficult inference problems exemplified by those involving large numbers of dimensions and relevant to a vast range of applications.

The paper capitalizes on some important findings that have been reported recently [22, 23, 56] regarding how the local properties of sequential Monte Carlo filtering measures impact algorithms' abilities to solve high-dimensional problems. We exploit the observation made by Rebeschini & van Handel [23], that, by using the decay of correlations property, it is possible to develop particle filters based on local solutions whose approximation error becomes less sensitive to augmenting the number of state dimensions.
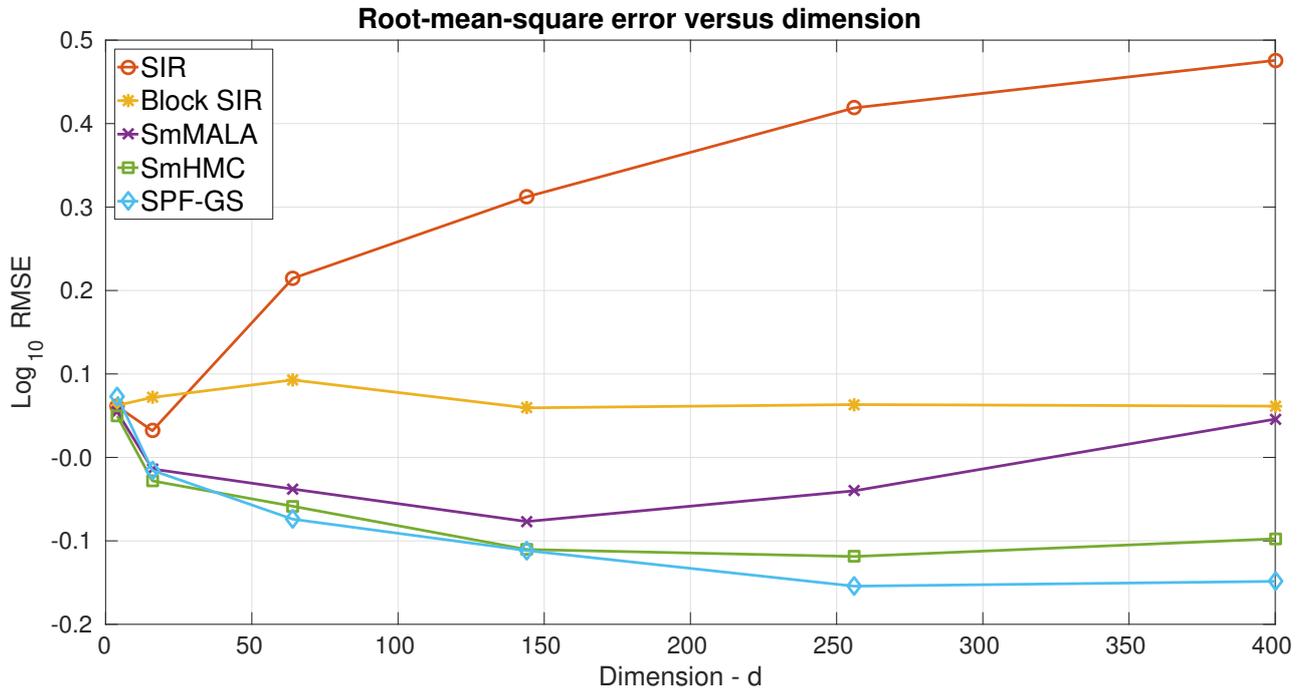
Figure 7.26: RMSE for the nonlinear, non-Gaussian sensor network example
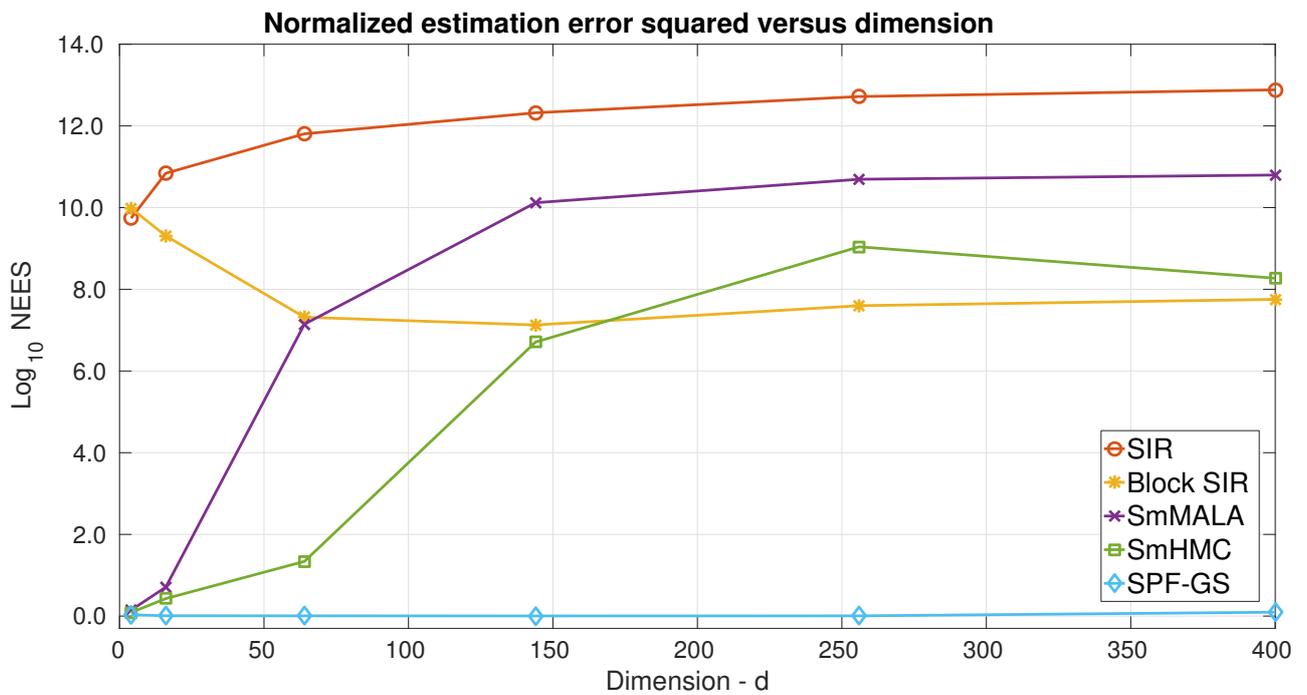
Figure 7.27: NEES for the nonlinear, non-Gaussian sensor network example
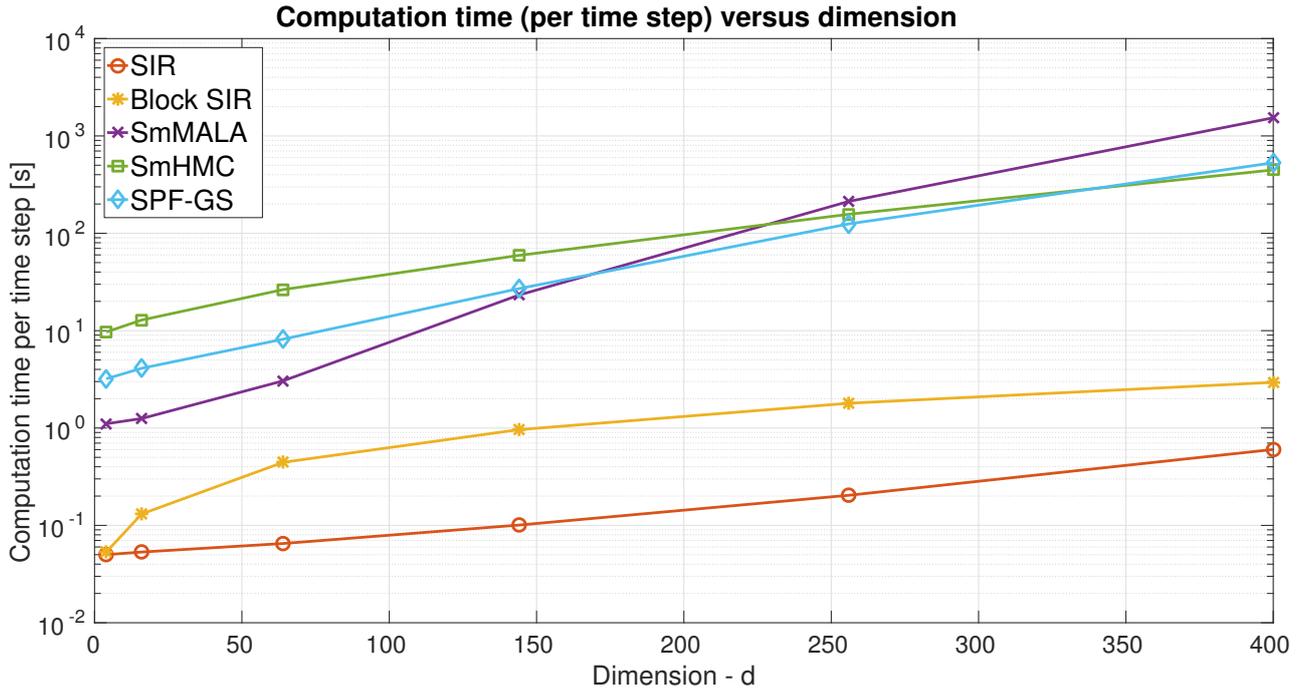
Figure 7.28: Average computation time for the nonlinear, non-Gaussian sensor network example

Within this context, we proposed a novel filter which aims to address the well-known shortcomings of sequential Monte Carlo methods when applied to nonlinear high-dimensional filtering problems. The novel method uses a Monte Carlo procedure to generate a sequence of equally-weighted samples that each guide a local solution of the Fokker-Planck equation. Using these local approximations, a mixture is produced that approximates the filtering density. The result is a statistically-sound general-purpose class of algorithms. In the context of a simple, though not trivial, high-dimensional inference problem and in comparison with state-of-the-art algorithms, the proposed approach has been shown to offer significant improvement in statistical consistency with commensurate computational expense.

In its most computationally efficient form, SPF-GS, the proposed filter has a complexity bounded by only $\mathcal{O}\left(NLn_x^3\right)$ evaluations[13] and it has the appealing property that its operations per sample (and the associated mixture component) can be parallelized. When articulated as a marginal particle filter, SPF-MPF, the complexity is bounded by $\mathcal{O}\left(NLn_x^3 + N + N^2\right)$ evaluations. That said, our investigations indicate that further consideration is required in order to explore fully the potential of stochastic particle flow. Future work will focus on the computational cost of the algorithms and guaranteeing certain properties of the diffusion matrix, (e.g., ensuring the matrix is positive definite and not singular). It should be possible, at least in the context of some statistical models, to reduce the computational complexity by exploiting or imposing sparsity in the diffusion matrix. Girolami & Calderhead [10] suggest that the use of guiding Hamiltonians [7] could provide a way of reducing such computational cost, but it is currently unclear how such a solution would be adopted in the context of stochastic particle flow. Another potentially promising future direction would be adopting the same approach as the well-known Broyden-Fletcher-Goldfarb-Shanno (BFGS) algorithm and thereby work around the need to explicitly evaluate the Hessian matrix at all simulation steps. We also strongly believe that it would be possible for future work to result in further improvements on the bounds used to compute the step size and number of simulation steps. Such advances will likely improve the computational efficiency of the algorithm and are the subject of ongoing development.

---

[13]It is important to remark that the stochastic particle flow is much more computationally complex than the original particle flows, which are generally bounded by $\mathcal{O}(N)$ computations both in theory and practice.
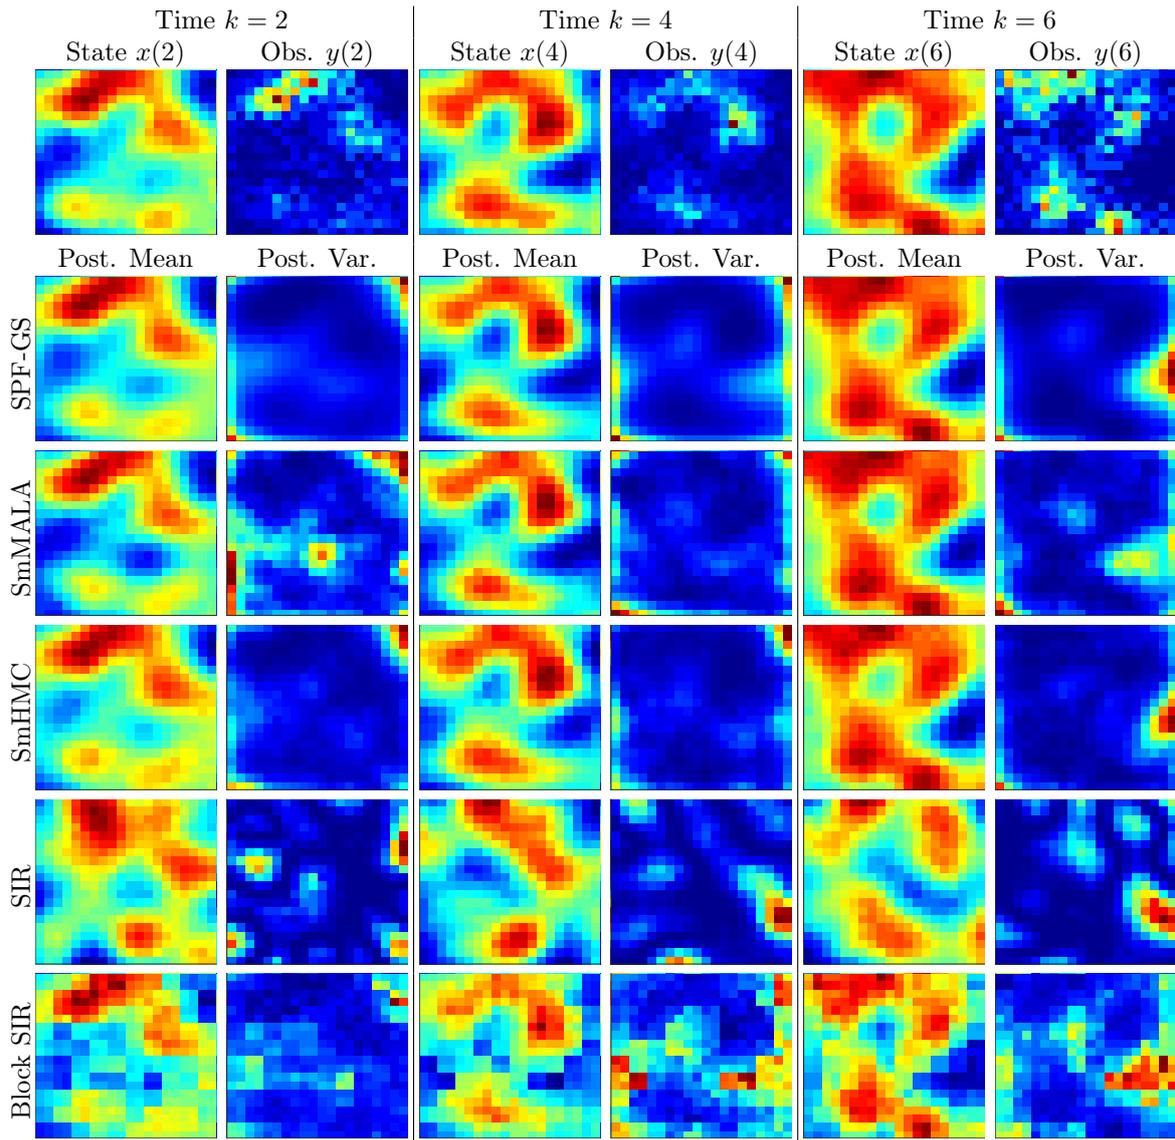
Figure 7.29: Posterior statistics for the nonlinear, non-Gaussian sensor network example ($n_x = 400$)

A   Proofs

A.1   Bounds for the Time Horizon and Step Size

General Assumptions

Let $\Phi : \mathbb{R}^{n_x} \to \mathbb{R}$ be a measurable convex function, satisfying

$$\int_{\mathbb{R}^{n_x}} \exp\{-\Phi(x)\} < \infty, \tag{A.1}$$

$$\Phi(x) - \Phi(\bar{x}) - \nabla_x \Phi(\bar{x})^T (x - \bar{x}) \geq \frac{1}{2} m \|x - \bar{x}\|_2^2, \tag{A.2}$$

$$\|\nabla_x \Phi(x) - \nabla_x \Phi(\bar{x})\|_2 \leq M \|x - \bar{x}\|_2, \quad \forall x, \bar{x} \in \mathbb{R}^{n_x}, \tag{A.3}$$

for two existing positive constants m and M. Let $\bar{x} \in \mathbb{R}^{n_x}$ be the global minimum of $\Phi(x)$. We define the log-concave target density for a Langevin algorithm as $\pi(x) = e^{-\Phi(x)} \left(\int_{\mathbb{R}^{n_x}} e^{-\Phi(x)} dx \right)^{-1}$.

Lemma 3. *Suppose a probability measure $\mathcal{P}_{\mathcal{L},T}$ produced by the exact integration, up to time horizon T, of the Langevin diffusion SDE*

$$d\mathcal{L}_\lambda = -\frac{1}{2} D_\lambda \nabla \Phi(\mathcal{L}_\lambda) d\lambda + D_\lambda^{1/2} dW_\lambda, \ \lambda \geq 0, \ \mathcal{L}_0 = 0, \tag{A.4}$$

departing from the initial density $\nu(x)$ and targeting the invariant density $\pi(x) \propto \exp\{-\Phi(x)\}$ with measure $\mathcal{P}_\pi(dx)$. Process $\{W_\lambda\}_{\lambda \geq 0}$ is the standard Wiener process and D_λ is the diffusion matrix. Under assumptions (A.1), (A.2), and (A.3),

$$\|\mathcal{P}_{\mathcal{L},T} - \mathcal{P}_\pi\|_{TV} \leq e^{-\frac{1}{2}Tm} \mathbb{E}^{\mathcal{P}_\pi} \left[\nu^2/\pi^2 \right]^{1/2}, \ T \geq 0. \tag{A.5}$$

Proof. By denoting the Markov transition kernel as $P_t(x, \cdot)$, an argument similar to that given by Dalalyan [64], invoking the Cauchy-Schwarz inequality and the spectral gap bound for the transition operator, gives

$$\|\mathcal{P}_{\mathcal{L},T} - \mathcal{P}_\pi\|_{TV} = \sup_{\mathcal{A} \in \mathscr{B}(\mathbb{R}^{n_x})} \left| \int_{\mathbb{R}^{n_x}} P_t(x, \mathcal{A}) \nu(x) dx - \mathcal{P}_\pi(\mathcal{A}) \right|$$

$$= \sup_{\mathcal{A} \in \mathscr{B}(\mathbb{R}^{n_x})} \left| \int_{\mathbb{R}^{n_x}} (P_t(x, \mathcal{A}) - \mathcal{P}_\pi(\mathcal{A})) \nu(x) dx \right|$$

$$= \sup_{\mathcal{A} \in \mathscr{B}(\mathbb{R}^{n_x})} \left| \int_{\mathbb{R}^{n_x}} (P_t(x, \mathcal{A}) - \mathcal{P}_\pi(\mathcal{A})) \frac{\nu(x)}{\pi(x)} \pi(x) dx \right|$$

$$\leq \sup_{\mathcal{A} \in \mathscr{B}(\mathbb{R}^{n_x})} \int_{\mathbb{R}^{n_x}} |P_t(x, \mathcal{A}) - \mathcal{P}_\pi(\mathcal{A})| \left| \frac{\nu(x)}{\pi(x)} \right| \pi(x) dx$$

$$\leq \sup_{\mathcal{A} \in \mathscr{B}(\mathbb{R}^{n_x})} \left[\int_{\mathbb{R}^{n_x}} |P_t(x, \mathcal{A}) - \mathcal{P}_\pi(\mathcal{A})|^2 \pi(x) dx \right]^{1/2} \left(\int_{\mathbb{R}^{n_x}} |\nu(x)/\pi(x)|^2 \pi(x) dx \right)^{1/2}$$

$$\text{(Dalalyan, 2014)} \leq \frac{1}{2} e^{-\frac{1}{2} mT} \mathbb{E}^{\mathcal{P}_\pi} \left[\nu^2/\pi^2 \right]^{1/2}, \tag{A.6}$$

where $\mathscr{B}(\mathbb{R}^{n_x})$ is the σ-algebra of Borel sets of $\mathbb{R}^{n_x}$. $\qquad\square$

Lemma 4. *Under conditions (A.1), (A.2) and (A.3), and assumptions of Lemma 3, given the initial probability density $\nu(x) = \mathcal{N}(x; \mu_\nu, \sigma_\nu^2 \mathbb{I}_{n_x})$, for $\sigma_\nu^2 < 2M^{-1}$, then*

$$\|\mathcal{P}_{\mathcal{L},T} - \mathcal{P}_\pi\|_{TV} \leq \frac{1}{2} \exp \left\{ -\frac{1}{2} mT + \frac{n_x}{4} \log \left(\frac{1}{m\sigma_\nu^2(2 - M\sigma_\nu^2)} \right) + \frac{M}{2(2 - M\sigma_\nu^2)} \|\bar{x} - \mu_\nu\|_2^2 \right\}. \tag{A.7}$$

Proof. According to *Lemma 4* in [64], if (A.3) holds then

$$\Phi(x) - \Phi(\bar{x}) - \nabla_x \Phi(\bar{x})^T (x - \bar{x}) \leq \frac{M}{2} \|x - \bar{x}\|_2^2. \tag{A.8}$$

Thus,

$$\pi(x) \geq e^{-\frac{M}{2}\|x-\bar{x}\|_2^2 - \Phi(\bar{x})} \left(\int_{\mathbb{R}^{n_x}} e^{-\frac{m}{2}\|x-\bar{x}\|_2^2 - \Phi(\bar{x})} dx \right)^{-1}$$

$$= e^{-\frac{M}{2}\|x-\bar{x}\|_2^2 - \Phi(\bar{x}) + \Phi(\bar{x})} \left(\int_{\mathbb{R}^{n_x}} e^{-\frac{m}{2}\|x-\bar{x}\|_2^2} dx \right)^{-1},$$

$$\pi(x) \geq \left(2\pi m^{-1} \right)^{-n_x/2} e^{-\frac{M}{2}\|x-\bar{x}\|_2^2}. \tag{A.9}$$

We use (A.9), $\sigma_\nu^{-2} > M/2$, and define $W = 2\sigma_\nu^{-2}$ to compute

$$\mathbb{E}^{\mathcal{P}_\pi}\left[\nu^2/\pi^2\right] = \int_{\mathbb{R}^{n_x}}\left(\frac{\nu}{\pi}\right)^2 \pi(\mathrm{x})d\mathrm{x} = \int_{\mathbb{R}^{n_x}}\frac{\nu^2}{\pi}d\mathrm{x}$$

$$\leq \int_{\mathbb{R}^{n_x}}\frac{\left[\left(2\pi\sigma_\nu^2\right)^{-n_x/2}e^{-\frac{\sigma_\nu^{-2}}{2}\|\mathrm{x}-\mu_\nu\|_2^2}\right]^2}{\left(2\pi m^{-1}\right)^{-n_x/2}e^{-\frac{M}{2}\|\mathrm{x}-\bar{\mathrm{x}}\|_2^2}}d\mathrm{x}$$

$$= \left(2\pi\sigma_\nu^2\right)^{-n_x}\left(2\pi m^{-1}\right)^{n_x/2}\int_{\mathbb{R}^{n_x}}\frac{e^{-\sigma_\nu^{-2}\|\mathrm{x}-\mu_\nu\|_2^2}}{e^{-\frac{M}{2}\|\mathrm{x}-\bar{\mathrm{x}}\|_2^2}}d\mathrm{x}$$

$$= \left(2\pi\cdot 2W^{-1}\right)^{-n_x}\left(2\pi m^{-1}\right)^{n_x/2}\int_{\mathbb{R}^{n_x}}e^{-\frac{W}{2}\|\mathrm{x}-\mu_\nu\|_2^2+\frac{M}{2}\|\mathrm{x}-\bar{\mathrm{x}}\|_2^2}d\mathrm{x}$$

$$= \frac{\left(2\pi m^{-1}\right)^{n_x/2}}{\left(2\pi\cdot 2W^{-1}\right)^{n_x}}e^{\frac{1}{2}(W-M)^{-1}WM\|\bar{\mathrm{x}}-\mu_\nu\|_2^2}\int_{\mathbb{R}^{n_x}}e^{-\frac{(W-M)}{2}\|\mathrm{x}-(W-M)^{-1}(W\mu_\nu-M\bar{\mathrm{x}})\|_2^2}d\mathrm{x}$$

$$= \frac{\left(2\pi m^{-1}\right)^{n_x/2}}{\left(2\pi\cdot 2W^{-1}\right)^{n_x}}e^{\frac{1}{2}\left(M^{-1}-W^{-1}\right)^{-1}\|\bar{\mathrm{x}}-\mu_\nu\|_2^2}\left(2\pi(W-M)^{-1}\right)^{n_x/2}$$

$$= \left(\frac{W^2}{4m(W-M)}\right)^{n_x/2}e^{\frac{1}{2}\left(M^{-1}-W^{-1}\right)^{-1}\|\bar{\mathrm{x}}-\mu_\nu\|_2^2}$$

$$= \left(\frac{\sigma_\nu^{-4}}{m(2\sigma_\nu^{-2}-M)}\right)^{n_x/2}e^{\frac{1}{2}\left(M^{-1}-\frac{1}{2}\sigma_\nu^2\right)^{-1}\|\bar{\mathrm{x}}-\mu_\nu\|_2^2}$$

$$= \left(\frac{1}{m\sigma_\nu^2(2-M\sigma_\nu^2)}\right)^{n_x/2}e^{M\left(2-M\sigma_\nu^2\right)^{-1}\|\bar{\mathrm{x}}-\mu_\nu\|_2^2}$$

$$\therefore \mathbb{E}^{\mathcal{P}_\pi}\left[\nu^2/\pi^2\right]^{1/2} \leq \exp\left\{\frac{n_x}{4}\log\left(\frac{1}{m\sigma_\nu^2(2-M\sigma_\nu^2)}\right)+\frac{M}{2(2-M\sigma_\nu^2)}\|\bar{\mathrm{x}}-\mu_\nu\|_2^2\right\}. \tag{A.10}$$

The result is complete by incorporating (A.10) into the result (A.6). $\qquad\square$

Lemma 5. *For $d\nu(\mathrm{x}) \sim \delta(\mathrm{x}-\mathrm{x}_\nu)d\mathrm{x}$, the Radon-Nikodym derivative*

$$\frac{d\nu}{d\mathcal{P}_\pi} = \frac{1}{\int_{\Omega(\mathrm{x}_\nu)}\pi(\mathrm{x}')d\mathrm{x}'} \tag{A.11}$$

is valid, where $d\mathcal{P}_\pi(\mathrm{x}) = \pi(\mathrm{x})d\mathrm{x}$ and the integral is taken over $\Omega(\mathrm{x}_\nu) = \{\mathrm{x}' \in \mathbb{R}^{n_x} : \|\mathrm{x}'-\bar{\mathrm{x}}\|_2 \geq \|\mathrm{x}_\nu-\bar{\mathrm{x}}\|_2\}$.

Proof. The Dirac-delta measure is singular, i.e., not absolutely continuous with respect to the Lebesgue measure, and hence its Radon-Nikodym derivative may not be formally defined in general. However, by abusing notation one can obtain a functional expression for the enunciated derivative. Observe the definition of the Radon-Nikodym derivative as

$$\frac{d\mathcal{P}_\pi}{d\nu} = f \implies \mathcal{P}_\pi(\mathcal{D}) = \int_{\mathcal{D}} f\, d\nu, \tag{A.12}$$

for any measurable domain $\mathcal{D} \subseteq \mathbb{R}^{n_x}$. Given $\nu(\mathrm{x}_\nu) = \delta(\mathrm{x}_\nu-\mathrm{x}) \equiv \delta(\mathrm{x}-\mathrm{x}_\nu) = \nu(\mathrm{x})$ and setting $f(\mathrm{x}) := \int_{\Omega(\mathrm{x})}\pi(\mathrm{x}')d\mathrm{x}'$, we write

$$\mathcal{P}_\pi(\mathrm{x}) = \int_{\mathbb{R}^{n_x}}f(\mathrm{x}_\nu)\,d\nu(\mathrm{x}_\nu) = \int_{\mathbb{R}^{n_x}}f(\mathrm{x}_\nu)\,\delta(\mathrm{x}_\nu-\mathrm{x})d\mathrm{x}_\nu = f(\mathrm{x}) = \int_{\Omega(\mathrm{x})}\pi(\mathrm{x}')d\mathrm{x}', \tag{A.13}$$

where $\frac{d\mathcal{P}_\pi}{d\mathrm{x}} = \pi(\mathrm{x})$ accordingly, and $f(\mathrm{x}_\nu) = \int_{\Omega(\mathrm{x}_\nu)}\pi(\mathrm{x}')d\mathrm{x}'$. Notice that $f(\mathrm{x})$ plays the role of a cumulative distribution function that is integrated over the "tail" of $\pi(\mathrm{x}')$, in the region defined by $\Omega(\mathrm{x})$.

Thus, applying definition (A.12), we get

$$\frac{d\mathcal{P}_\pi}{d\nu} := f(\mathrm{x}_\nu) = \int_{\Omega(\mathrm{x}_\nu)}\pi(\mathrm{x}')d\mathrm{x}',$$

$$\therefore \frac{d\nu}{d\mathcal{P}_\pi} = \frac{1}{\int_{\Omega(\mathrm{x}_\nu)}\pi(\mathrm{x}')d\mathrm{x}'}. \tag{A.14}$$

$\qquad\square$

Remark 6. Notice that if $\pi(\mathrm{x}') = (2\pi P^{-1})^{-\frac{n_x}{2}} e^{-\frac{P}{2}\|\mathrm{x}'-\bar{\mathrm{x}}\|_2^2}$, for some $P \in \mathbb{R}_+$, we can compute the resulting integral in *Lemma 5* using the spherical symmetry of the Gaussian function to give

$$
\begin{aligned}
\int_{\Omega(\mathrm{x}_\nu)} \pi(\mathrm{x}')d\mathrm{x}' &= (2\pi P^{-1})^{-\frac{n_x}{2}} \int_{\Omega(\mathrm{x}_\nu)} e^{-\frac{P}{2}\|\mathrm{x}'-\bar{\mathrm{x}}\|_2^2} d\mathrm{x}' \\
&= (2\pi)^{-\frac{n_x}{2}} \int_{\{\|u\|_2 \geq \sqrt{P}\|\mathrm{x}_\nu-\bar{\mathrm{x}}\|_2\}} e^{-\frac{1}{2}\|u\|_2^2} du \\
&= (2\pi)^{-\frac{n_x}{2}} \int_{r_\nu}^{\infty} \int_{\mathcal{S}^{n_x-1}(r)} e^{-\frac{1}{2}r^2} d\mathcal{A}\, dr \\
&= (2\pi)^{-\frac{n_x}{2}} \int_{r_\nu}^{\infty} e^{-\frac{1}{2}r^2} \mathcal{A}_{n_x-1}(r)\, dr \\
&= (2\pi)^{-\frac{n_x}{2}} \mathcal{A}_{n_x-1}(1) \int_{r_\nu}^{\infty} e^{-\frac{1}{2}r^2} r^{n_x-1}\, dr \\
&= (2\pi)^{-\frac{n_x}{2}} \frac{2\pi^{\frac{n_x}{2}}}{\Gamma\left(\frac{n_x}{2}\right)} \int_{r_\nu}^{\infty} e^{-\frac{1}{2}r^2} r^{n_x-1}\, dr \\
&= (2\pi)^{-\frac{n_x}{2}} \frac{2\pi^{\frac{n_x}{2}}}{\Gamma\left(\frac{n_x}{2}\right)} 2^{\frac{n_x}{2}-1} \int_{r_\nu^2/2}^{\infty} e^{-t} \cdot t^{\frac{n_x}{2}-1}\, dt \\
&= \frac{1}{\Gamma\left(\frac{n_x}{2}\right)} \int_{r_\nu^2/2}^{\infty} t^{\frac{n_x}{2}-1} e^{-t}\, dt \\
&= \Gamma_u\left(\frac{n_x}{2}, \frac{r_\nu^2}{2}\right) = \Gamma_u\left(\frac{n_x}{2}, \frac{P\|\mathrm{x}_\nu-\bar{\mathrm{x}}\|_2^2}{2}\right),
\end{aligned}
$$

where $du = d\mathcal{A}\, dr$ for a volume element $du \in \mathbb{R}^{n_x}$, area element $d\mathcal{A} \in \mathbb{R}^{n_x-1}$ and radius element $dr \in \mathbb{R}_+$. In addition, $\mathcal{S}^{n_x-1}(r)$ denotes the $(n_x - 1)$-sphere of radius r, with a total surface area of $\mathcal{A}_{n_x-1}(r)$, and the recursion $\mathcal{A}_{n_x-1}(r) = \mathcal{A}_{n_x-1}(1)r^{n_x-1}$ has been applied with $\mathcal{A}_{n_x-1}(1) = 2\pi^{\frac{n_x}{2}}/\Gamma(n_x/2)$. The lower integration extreme is taken as $r_\nu := \sqrt{P}\|\mathrm{x}_\nu - \bar{\mathrm{x}}\|_2$, and $\Gamma_u(s, x) = \Gamma(s)^{-1} \int_x^{\infty} t^{s-1} e^{-t} dt$ is the upper incomplete gamma function. Also, we recall that $\int_{\Omega(\mathrm{x}_\nu)} \pi(\mathrm{x}')d\mathrm{x}' = 1 - \int_{\mathbb{R}^{n_x}\setminus\Omega(\mathrm{x}_\nu)} \pi(\mathrm{x}')d\mathrm{x}'$.

Lemma 7. *Under conditions (A.1), (A.2) and (A.3), and assumptions of Lemma 3, given a initial probability mass located at* $\mathrm{x} = \mathrm{x}_\nu$, *i.e.,* $\nu(\mathrm{x}) = \delta(\mathrm{x} - \mathrm{x}_\nu)$ *(Dirac delta), then*

$$
\|\mathcal{P}_{\mathcal{L},T} - \mathcal{P}_\pi\|_{TV} \leq \frac{1}{2} \exp\left\{ -\frac{1}{2}mT + \frac{n_x}{2}\log\left(\frac{M}{m}\right) - \log\left[\Gamma_u\left(\frac{n_x}{2}, \frac{M\|\bar{\mathrm{x}} - \mathrm{x}_\nu\|_2^2}{2}\right)\right] \right\}. \tag{A.15}
$$

Proof. Considering $\nu(\mathrm{x}) = \delta(\mathrm{x} - \mathrm{x}_\nu)$ on the definition of total variation distance, we use *Lemma 5* and *Remark 6*.

The procedure follows as

$$
\begin{aligned}
\|\mathcal{P}_{\mathcal{L},T} - \mathcal{P}_\pi\|_{\text{TV}} &= \sup_{\mathcal{A}\in\mathscr{B}(\mathbb{R}^{n_x})} \left| \int_{\mathbb{R}^{n_x}} P_t(\mathrm{x},\mathcal{A}) d\nu(\mathrm{x}) - \mathcal{P}_\pi(\mathcal{A}) \right| \\
&= \sup_{A\in\mathscr{B}(\mathbb{R}^{n_x})} \left| \int_{\mathbb{R}^{n_x}} \left(P_t(\mathrm{x},\mathcal{A}) - \mathcal{P}_\pi(\mathcal{A}) \right) d\nu(\mathrm{x}) \right| \\
&= \sup_{A\in\mathscr{B}(\mathbb{R}^{n_x})} \left| \int_{\mathbb{R}^{n_x}} \left(P_t(\mathrm{x},\mathcal{A}) - \mathcal{P}_\pi(\mathcal{A}) \right) \frac{d\nu}{d\mathcal{P}_\pi} d\mathcal{P}_\pi(\mathrm{x}) \right| \\
(\text{Lemma 5}) &= \sup_{A\in\mathscr{B}(\mathbb{R}^{n_x})} \left| \int_{\mathbb{R}^{n_x}} \left(P_t(\mathrm{x},\mathcal{A}) - \mathcal{P}_\pi(\mathcal{A}) \right) \frac{1}{\int_{\Omega(\mathrm{x}_\nu)} \pi(\mathrm{x}')d\mathrm{x}'} d\mathcal{P}_\pi(\mathrm{x}) \right| \\
&\leq \frac{1}{\int_{\Omega(\mathrm{x}_\nu)} \pi(\mathrm{x}')d\mathrm{x}'} \cdot \sup_{A\in\mathscr{B}(\mathbb{R}^{n_x})} \int_{\mathbb{R}^{n_x}} \left| P_t(\mathrm{x},\mathcal{A}) - \mathcal{P}_\pi(\mathcal{A}) \right| \pi(\mathrm{x})d\mathrm{x} \\
&\leq \frac{1}{\int_{\Omega(\mathrm{x}_\nu)} \pi(\mathrm{x}')d\mathrm{x}'} \cdot \sup_{A\in\mathscr{B}(\mathbb{R}^{n_x})} \left[\int_{\mathbb{R}^{n_x}} \left| P_t(\mathrm{x},\mathcal{A}) - \mathcal{P}_\pi(\mathcal{A}) \right|^2 \pi(\mathrm{x})d\mathrm{x} \right]^{1/2} \\
&\leq \frac{(2\pi m^{-1})^{n_x/2}}{\int_{\Omega(\mathrm{x}_\nu)} e^{-\frac{M}{2}\|\mathrm{x}'-\bar{\mathrm{x}}\|_2^2} d\mathrm{x}'} \sup_{A\in\mathscr{B}(\mathbb{R}^{n_x})} \left[\int_{\mathbb{R}^{n_x}} \left| P_t(\mathrm{x},\mathcal{A}) - \mathcal{P}_\pi(\mathcal{A}) \right|^2 \pi(\mathrm{x})d\mathrm{x} \right]^{1/2} \\
&= \frac{(2\pi m^{-1})^{n_x/2} (2\pi M^{-1})^{-n_x/2}}{\int_{\Omega(\mathrm{x}_\nu)} \mathcal{N}(\mathrm{x}';\bar{\mathrm{x}}, M^{-1}\mathbb{I}_{n_x})d\mathrm{x}'} \sup_{A\in\mathscr{B}(\mathbb{R}^{n_x})} \left[\int_{\mathbb{R}^{n_x}} \left| P_t(\mathrm{x},\mathcal{A}) - \mathcal{P}_\pi(\mathcal{A}) \right|^2 \pi(\mathrm{x})d\mathrm{x} \right]^{1/2} \\
(\text{Remark 6}) &= \frac{(Mm^{-1})^{n_x/2}}{\Gamma_u\left(\frac{n_x}{2}, \frac{M\|\bar{\mathrm{x}}-\mathrm{x}_\nu\|_2^2}{2}\right)} \sup_{A\in\mathscr{B}(\mathbb{R}^{n_x})} \left[\int_{\mathbb{R}^{n_x}} \left| P_t(\mathrm{x},\mathcal{A}) - \mathcal{P}_\pi(\mathcal{A}) \right|^2 \pi(\mathrm{x})d\mathrm{x} \right]^{1/2} \\
(\text{Dalalyan, 2014}) &\leq \left(\frac{M}{m}\right)^{n_x/2} \cdot \Gamma_u\left(\frac{n_x}{2}, \frac{M\|\mathrm{x}_\nu-\bar{\mathrm{x}}\|_2^2}{2}\right)^{-1} \cdot \frac{1}{2} e^{-\frac{1}{2}mT} \\
&= \frac{1}{2}\exp\left\{ -\frac{1}{2}mT + \frac{n_x}{2}\log\left(\frac{M}{m}\right) - \log\left[\Gamma_u\left(\frac{n_x}{2}, \frac{M\|\bar{\mathrm{x}}-\mathrm{x}_\nu\|_2^2}{2}\right)\right] \right\}. \quad (\text{A.16})
\end{aligned}
$$

$\square$

Lemma 8. *Suppose a probability measure $\mathcal{P}_{\tilde{\mathcal{L}}[\Delta\lambda],T}$ produced by numerical integration with step size, $\Delta\lambda$, up to time horizon $T = L\Delta\lambda$, of the Langevin diffusion SDE according to*

$$
d\tilde{\mathcal{L}}_\lambda = -\frac{1}{2}\sum_{l=0}^{L-1} D_{l\Delta\lambda}\nabla\tilde{\Phi}(\tilde{\mathcal{L}}_{l\Delta\lambda})\mathbb{1}_{[l\Delta\lambda,(l+1)\Delta\lambda)}d\lambda + \sum_{l=0}^{L-1} D_{l\Delta\lambda}^{1/2}dW_\lambda, \, 0 \leq \lambda \leq L\Delta\lambda, \, \tilde{\mathcal{L}}_0 = 0. \quad (\text{A.17})
$$

The resulting Markov chain is assumed to depart from the initial density $\nu(\mathrm{x}_0) = \mathcal{N}(\mathrm{x}_0; \mu_\nu, \sigma_\nu^2\mathbb{I}_{n_x})$ and targets the invariant density $\pi(\mathrm{x}) \propto \exp\{-\Phi(\mathrm{x})\}$ with measure $\mathcal{P}_\pi(d\mathrm{x})$. The process $\{W_\lambda\}_{\lambda\geq 0}$ is the standard Wiener process and the diffusion matrix $D_\lambda = D(\mathrm{x}_\lambda)$ is bounded by $K = \sup_{\mathrm{x}}\|D(\mathrm{x})\|_2$. Under assumptions (A.1), (A.2), and (A.3), then

$$
\begin{aligned}
\|\mathcal{P}_{\mathcal{L},T} - \mathcal{P}_{\tilde{\mathcal{L}}[\Delta\lambda],T}\|_{TV} \leq \frac{1}{2} \\
-\frac{1}{2}\exp\left\{ -\frac{n_x}{2}\frac{M^3K^4\gamma}{48(2\gamma-1)}\left(\sigma_\nu^2 + \frac{1}{n_x}\|\bar{\mathrm{x}}-\mu_\nu\|_2^2 + 2T\right)\Delta\lambda^2 - \frac{n_xM^2K^3T}{16}\Delta\lambda \right\}, \, T \geq 0; \quad (\text{A.18})
\end{aligned}
$$

where $\mathcal{P}_{\mathcal{L},T}$ is the probability measure produced by the exact integration of the Langevin diffusion SDE.

Proof. We take a different approach as that proposed by Dalalyan [64] for this proof. Instead of bounding the

total variance distance $\|\mathcal{P}_{\mathcal{L},T} - \mathcal{P}_{\tilde{\mathcal{L}}[\Delta\lambda],T}\|_{\mathrm{TV}}$ by the Pinsker inequality, we treat it directly. First, identify that

$$
\begin{aligned}
\|\mathcal{P}_{\mathcal{L},T} - \mathcal{P}_{\tilde{\mathcal{L}}[\Delta\lambda],T}\|_{\mathrm{TV}} &= \frac{1}{2} \int_{\mathbb{R}^{n_x}} \left| d\mathcal{P}_{\mathcal{L},T} - d\mathcal{P}_{\tilde{\mathcal{L}}[\Delta\lambda],T} \right| \\
&= \frac{1}{2} \int_{\mathbb{R}^{n_x}} \left| 1 - \frac{d\mathcal{P}_{\tilde{\mathcal{L}}[\Delta\lambda],T}}{d\mathcal{P}_{\mathcal{L},T}} \right| d\mathcal{P}_{\mathcal{L},T} \\
&= \frac{1}{2} \int_{\mathbb{R}^{n_x}} \left| 1 - \frac{\nu(\mathrm{x}_0) dP_t^{\tilde{\mathcal{L}}}(\mathrm{x}_0,\mathrm{x})}{\nu(\mathrm{x}_0) dP_t^{\mathcal{L}}(\mathrm{x}_0,\mathrm{x})} \right| d\mathcal{P}_{\mathcal{L},T} \\
&= \frac{1}{2} \int_{\mathbb{R}^{n_x}} \left| 1 - \frac{dP_t^{\tilde{\mathcal{L}}}(\mathrm{x}_0,\mathrm{x})}{dP_t^{\mathcal{L}}(\mathrm{x}_0,\mathrm{x})} \right| d\mathcal{P}_{\mathcal{L},T}(\mathrm{x}) \\
&= \frac{1}{2} \int_{\mathbb{R}^{n_x}} \mathbb{E}^{P_t^{\mathcal{L}}} \left[\left| 1 - \frac{dP_t^{\tilde{\mathcal{L}}}(\mathrm{x}_0,\mathrm{x})}{dP_t^{\mathcal{L}}(\mathrm{x}_0,\mathrm{x})} \right| \Big| \mathrm{x}_0 \right] d\nu(\mathrm{x}_0), \quad \text{(A.19)}
\end{aligned}
$$

where $\mathcal{P}_{\mathcal{L},T}(dx) = \int_{\mathbb{R}^{n_x}} P_t^{\mathcal{L}}(\mathrm{x}_0, dx)\nu(\mathrm{x}_0)dx_0$, and $P_t^{\tilde{\mathcal{L}}}(\mathrm{x}_0,\cdot)$ and $P_t^{\mathcal{L}}(\mathrm{x}_0,\cdot)$ are the Markov transition kernels for the discrete-time and continuous-time Langevin dynamics respectively. By applying Girsanov's theorem to change the measure from $\mathcal{P}_{\mathcal{L},T}$ to $\mathcal{P}_{\tilde{\mathcal{L}}[\Delta\lambda],T}$, one obtains (see step 3 of the proof of *Proposition 2* in [91])

$$
\frac{dP_t^{\tilde{\mathcal{L}}}}{dP_t^{\mathcal{L}}}(\mathrm{x})
$$
$$
= \exp\left\{ \int_0^T \frac{1}{2} \left(-\sum_{l=0}^{L-1} D_{l\Delta\lambda} \nabla\tilde{\Phi}(\mathrm{x}_{l\Delta\lambda})\mathbb{1}_{[l\Delta\lambda,(l+1)\Delta\lambda)} + D_\lambda \nabla\Phi(\mathrm{x}_\lambda) \right)^T D_\lambda^{-1} \left(d\mathcal{L}_\lambda + \frac{1}{2}D_\lambda \nabla\Phi(\mathcal{L}_\lambda)d\lambda \right) \right\}
$$
$$
\times \exp\left\{ -\frac{1}{2} \int_0^T \left\| -\frac{1}{2}\sum_{l=0}^{L-1} D_{l\Delta\lambda} \nabla\tilde{\Phi}(\mathrm{x}_{l\Delta\lambda})\mathbb{1}_{[l\Delta\lambda,(l+1)\Delta\lambda)} + \frac{1}{2}D_\lambda \nabla\Phi(\mathrm{x}_\lambda) \right\|_2^2 d\lambda \right\}. \quad \text{(A.20)}
$$

Thus, since $\mathbb{E}[e^z] \geq e^{\mathbb{E}[z]}$ by Jensen's inequality, $D_\lambda = D(\mathrm{x}_\lambda) \leq K$, and $\nabla\Phi(\mathrm{x})$ is Lipschitz-continuous with constant M (A.3), we obtain

$$
\begin{aligned}
\mathbb{E}^{P_t^{\mathcal{L}}} \left[\frac{dP_t^{\tilde{\mathcal{L}}}}{dP_t^{\mathcal{L}}}(\mathrm{x})|\mathrm{x}_0 \right] &\geq \exp\left\{ -\mathbb{E}^{P_t^{\mathcal{L}}}\left[\frac{1}{2}\sum_{l=0}^{L-1} \int_{l\Delta\lambda}^{(l+1)\Delta\lambda} \left\| \frac{1}{2}D_\lambda \nabla\Phi(\mathrm{x}_\lambda) - \frac{1}{2}D_{l\Delta\lambda} \nabla\tilde{\Phi}(\mathrm{x}_{l\Delta\lambda}) \right\|_2^2 d\lambda |\mathrm{x}_0 \right] \right\} \\
&= \exp\left\{ -\frac{1}{2}\sum_{l=0}^{L-1} \int_{l\Delta\lambda}^{(l+1)\Delta\lambda} \mathbb{E}^{P_t^{\mathcal{L}}}\left[\left\| \frac{1}{2}D_\lambda \nabla\Phi(\mathrm{x}_\lambda) - \frac{1}{2}D_{l\Delta\lambda} \nabla\tilde{\Phi}(\mathrm{x}_{l\Delta\lambda}) \right\|_2^2 |\mathrm{x}_0 \right] d\lambda \right\} \\
&\geq \exp\left\{ -\frac{K^2}{8}\sum_{l=0}^{L-1} \int_{l\Delta\lambda}^{(l+1)\Delta\lambda} \mathbb{E}^{P_t^{\mathcal{L}}}\left[\left\| \nabla\Phi(\mathrm{x}_\lambda) - \nabla\tilde{\Phi}(\mathrm{x}_{l\Delta\lambda}) \right\|_2^2 |\mathrm{x}_0 \right] d\lambda \right\} \\
&\geq \exp\left\{ -\frac{K^2 M^2}{8}\sum_{l=0}^{L-1} \int_{l\Delta\lambda}^{(l+1)\Delta\lambda} \mathbb{E}^{P_t^{\mathcal{L}}}\left[\|\mathrm{x}_\lambda - \mathrm{x}_{l\Delta\lambda}\|_2^2 |\mathrm{x}_0 \right] d\lambda \right\}. \quad \text{(A.21)}
\end{aligned}
$$

For each step $\mathrm{x}_\lambda - \mathrm{x}_{l\Delta\lambda} = -\frac{1}{2}D_{l\Delta\lambda}\nabla\Phi(\mathrm{x}_{l\Delta\lambda})(\lambda - l\Delta\lambda) + D_{l\Delta\lambda}^{1/2}(W_\lambda - W_{l\Delta\lambda})$, hence

$$
\mathbb{E}^{P_t^{\mathcal{L}}}\left[\frac{dP_t^{\tilde{\mathcal{L}}}}{dP_t^{\mathcal{L}}}(\mathrm{x})|\mathrm{x}_0\right] \geq \exp\left\{-\frac{K^2M^2}{8}\sum_{l=0}^{L-1}\int_{l\Delta\lambda}^{(l+1)\Delta\lambda}\mathbb{E}^{P_t^{\mathcal{L}}}\left[\left\|\frac{1}{2}D(\mathrm{x}_{l\Delta\lambda})\nabla\Phi(\mathrm{x}_{l\Delta\lambda})\right\|_2^2(\lambda - l\Delta\lambda)^2|\mathrm{x}_0\right]d\lambda\right\}
$$

$$
\times \exp\left\{-\frac{K^2M^2}{8}\sum_{l=0}^{L-1}\int_{l\Delta\lambda}^{(l+1)\Delta\lambda}\mathbb{E}^{P_t^{\mathcal{L}}}\left[\left\|D(\mathrm{x}_{l\Delta\lambda})^{1/2}(W_\lambda - W_{l\Delta\lambda})\right\|_2^2|\mathrm{x}_0\right]d\lambda\right\}
$$

$$
\geq \exp\left\{-\frac{K^2M^2}{8}\sum_{l=0}^{L-1}\int_{l\Delta\lambda}^{(l+1)\Delta\lambda}\frac{K^2}{4}\mathbb{E}^{P_t^{\mathcal{L}}}\left[\|\nabla\Phi(\mathrm{x}_{l\Delta\lambda})\|_2^2|\mathrm{x}_0\right](\lambda - l\Delta\lambda)^2 d\lambda\right\}
$$

$$
\times \exp\left\{-\frac{K^2M^2}{8}\sum_{l=0}^{L-1}\int_{l\Delta\lambda}^{(l+1)\Delta\lambda}K(\lambda - l\Delta\lambda)n_x d\lambda\right\}
$$

$$
= \exp\left\{-\frac{M^2K^4\Delta\lambda^3}{96}\sum_{l=0}^{L-1}\mathbb{E}^{P_t^{\mathcal{L}}}\left[\|\nabla\Phi(\mathrm{x}_{l\Delta\lambda})\|_2^2|\mathrm{x}_0\right] - \frac{n_x M^2 K^3 L\Delta\lambda^2}{16}\right\}. \tag{A.22}
$$

Now we invoke a result from *Corollary 4* in [64], that gives

$$
\Delta\lambda\sum_{l=0}^{L-1}\mathbb{E}\left[\|\nabla\Phi(\mathrm{x}_{l\Delta\lambda})\|_2^2\right] \leq \frac{M\gamma}{2\gamma-1}\mathbb{E}\left[\|\mathrm{x}_0 - \bar{\mathrm{x}}\|_2^2\right] + \frac{2\gamma M T n_x}{2\gamma-1}, \tag{A.23}
$$

for some $\gamma > 1$, $\Delta\lambda \leq (\gamma M)^{-1}$ and $L > 1$. Substituting $L\Delta\lambda = T$ and incorporating (A.23):

$$
\mathbb{E}^{P_t^{\mathcal{L}}}\left[\frac{dP_t^{\tilde{\mathcal{L}}}}{dP_t^{\mathcal{L}}}(\mathrm{x})|\mathrm{x}_0\right] \geq \exp\left\{-\frac{M^2K^4\Delta\lambda^2}{48}\frac{\gamma}{2\gamma-1}\left(\frac{M}{2}\mathbb{E}^{P_t^{\mathcal{L}}}\left[\|\mathrm{x}_0 - \bar{\mathrm{x}}\|_2^2|\mathrm{x}_0\right] + n_x M T\right) - \frac{n_x M^2 K^3 T\Delta\lambda}{16}\right\}
$$

$$
= \exp\left\{-\frac{M^2K^4\Delta\lambda^2}{48}\frac{\gamma}{2\gamma-1}\left(\frac{M}{2}\|\mathrm{x}_0 - \bar{\mathrm{x}}\|_2^2 + n_x M T\right) - \frac{n_x M^2 K^3 T\Delta\lambda}{16}\right\}. \tag{A.24}
$$

Since

$$
0 \leq \mathbb{E}^{P_t^{\mathcal{L}}}\left[\frac{dP_t^{\tilde{\mathcal{L}}}}{dP_t^{\mathcal{L}}}(\mathrm{x})|\mathrm{x}_0\right] \leq 1,
$$

and applying expression (A.24) in (A.19), we have

$$
\|\mathcal{P}_{\mathcal{L},T} - \mathcal{P}_{\tilde{\mathcal{L}}[\Delta\lambda],T}\|_{\mathrm{TV}} = \frac{1}{2}\int_{\mathbb{R}^{n_x}}\left|1 - \frac{dP_t^{\tilde{\mathcal{L}}}(\mathrm{x}_0,\mathrm{x})}{dP_t^{\mathcal{L}}(\mathrm{x}_0,\mathrm{x})}\right|d\mathcal{P}_{\mathcal{L},T}(\mathrm{x})
$$

$$
= \frac{1}{2}\int_{\mathbb{R}^{n_x}}d\mathcal{P}_{\mathcal{L},T}(\mathrm{x}) - \frac{1}{2}\int_{\mathbb{R}^{n_x}}\mathbb{E}^{P_t^{\mathcal{L}}}\left[\frac{dP_t^{\tilde{\mathcal{L}}}(\mathrm{x}_0,\mathrm{x})}{dP_t^{\mathcal{L}}(\mathrm{x}_0,\mathrm{x})}|\mathrm{x}_0\right]d\nu(\mathrm{x}_0)
$$

$$
\leq \frac{1}{2} - \frac{1}{2}\int_{\mathbb{R}^{n_x}}e^{-\frac{M^2K^4\Delta\lambda^2}{48}\frac{\gamma}{2\gamma-1}\left(\frac{M}{2}\|\mathrm{x}_0-\bar{\mathrm{x}}\|_2^2 + n_x M T\right) - \frac{n_x M^2 K^3 T\Delta\lambda}{16}}\nu(\mathrm{x}_0)d\mathrm{x}_0. \tag{A.25}
$$

Finally, using the substitutions $A_1 = \frac{M^3K^4}{96}\frac{\gamma}{2\gamma-1}$, $A_2 = \frac{n_x M^3 T K^4}{48}\frac{\gamma}{2\gamma-1}$, $B = \frac{n_x M^2 K^3 T}{16}$, $\mathbb{E}[e^{-z}] \geq e^{-\mathbb{E}[z]}$ by Jensen's inequality, and $\nu(\mathrm{x}_0) = \mathcal{N}(\mathrm{x}_0; \mu_\nu, \sigma_\nu^2\mathbb{I}_{n_x})$, we get

$$
\|\mathcal{P}_{\mathcal{L},T} - \mathcal{P}_{\tilde{\mathcal{L}}[\Delta\lambda],T}\|_{\mathrm{TV}} \leq \frac{1}{2} - \frac{1}{2}e^{-A_2\Delta\lambda^2 - B\Delta\lambda}\int_{\mathbb{R}^{n_x}}e^{-A_1\Delta\lambda^2\|\mathrm{x}_0-\bar{\mathrm{x}}\|_2^2}\nu(\mathrm{x}_0)d\mathrm{x}_0
$$

$$
\leq \frac{1}{2} - \frac{1}{2}e^{-A_2\Delta\lambda^2 - B\Delta\lambda}e^{-A_1\Delta\lambda^2\int_{\mathbb{R}^{n_x}}\|\mathrm{x}_0-\bar{\mathrm{x}}\|_2^2\nu(\mathrm{x}_0)d\mathrm{x}_0}
$$

$$
= \frac{1}{2} - \frac{1}{2}e^{-A_2\Delta\lambda^2 - B\Delta\lambda - A_1\Delta\lambda^2\left(n_x\sigma_\nu^2 + \|\bar{\mathrm{x}}-\mu_\nu\|_2^2\right)}
$$

$$
= \frac{1}{2} - \frac{1}{2}e^{-\left(A_1\left(n_x\sigma_\nu^2 + \|\bar{\mathrm{x}}-\mu_\nu\|_2^2\right) + A_2\right)\Delta\lambda^2 - B\Delta\lambda}
$$

$$
= \frac{1}{2} - \frac{1}{2}e^{-\frac{n_x}{2}\frac{M^3K^4\gamma}{48(2\gamma-1)}\left(\sigma_\nu^2 + \frac{1}{n_x}\|\bar{\mathrm{x}}-\mu_\nu\|_2^2 + 2T\right)\Delta\lambda^2 - \frac{n_x M^2 K^3 T}{16}\Delta\lambda}. \tag{A.26}
$$

$\square$

Lemma 9. *Under the same assumptions as of Lemma 8, except for a initial density $\nu(\mathrm{x}_0) = \delta(\mathrm{x}_0 - \mathrm{x}_\nu)$ (Dirac delta), i.e., a probability mass initially located at $\mathrm{x}_0 = \mathrm{x}_\nu$, then*

$$\|\mathcal{P}_{\mathcal{L},T} - \mathcal{P}_{\tilde{\mathcal{L}}[\Delta\lambda],T}\|_{TV} \leq \frac{1}{2}$$
$$- \frac{1}{2}\exp\left\{-\frac{n_x}{2}\frac{M^3 K^4 \gamma}{48(2\gamma - 1)}\left(\frac{1}{n_x}\|\bar{\mathrm{x}} - \mathrm{x}_\nu\|_2^2 + 2T\right)\Delta\lambda^2 - \frac{n_x M^2 K^3 T}{16}\Delta\lambda\right\}, \ T \geq 0; \quad \text{(A.27)}$$

Proof. The proof follows straightforwardly by substituting $\nu(\mathrm{x}_0) = \delta(\mathrm{x}_0 - \mathrm{x}_\nu)$ into (A.25). $\square$

Theorem 10. *Let a convex function Φ satisfy the general assumptions (A.1), (A.2) and (A.3). Suppose a discrete-time Langevin Monte Carlo algorithm integrates (4.15), targeting the invariant density $\pi(\mathrm{x}) \propto \exp\{-\Phi(\mathrm{x})\}$ with measure $\mathcal{P}_\pi(d\mathrm{x})$. In addition, assume that for some $\gamma \geq 1$ we have $\Delta\lambda \leq (\gamma M)^{-1}$, and $K = \sup_{\mathrm{x}}\|D(\mathrm{x})\|_2$ where $D_\lambda = D(\mathrm{x}_\lambda)$ is the diffusion matrix. Then, for a time horizon, T, and step size, $\Delta\lambda$, the total-variation distance between the target measure $\mathcal{P}_\pi$ and the approximated measure $\mathcal{P}_{\tilde{\mathcal{L}}(\Delta\lambda),T}$ furnished by a discrete-time Langevin Monte Carlo algorithm with initial density $\nu(\mathrm{x}) = \mathcal{N}(\mathrm{x}; \mu_\nu, \sigma_\nu^2 \mathbb{I}_{n_x})$, for $\sigma_\nu^2 < 2M^{-1}$, satisfies*

$$\|\mathcal{P}_{\tilde{\mathcal{L}}[\Delta\lambda],T} - \mathcal{P}_\pi\|_{TV} \leq \frac{1}{2}\exp\left\{-\frac{1}{2}mT + \frac{n_x}{4}\log\left(\frac{1}{m\sigma_\nu^2(2 - M\sigma_\nu^2)}\right) + \frac{M}{2(2 - M\sigma_\nu^2)}\|\bar{\mathrm{x}} - \mu_\nu\|_2^2\right\}$$
$$+ \frac{1}{2} - \frac{1}{2}\exp\left\{-\frac{n_x}{2}\frac{M^3 K^4 \gamma}{48(2\gamma - 1)}\left(\sigma_\nu^2 + \frac{1}{n_x}\|\bar{\mathrm{x}} - \mu_\nu\|_2^2 + 2T\right)\Delta\lambda^2 - \frac{n_x M^2 K^3 T}{16}\Delta\lambda\right\}. \quad \text{(A.28)}$$

Proof. The proof follows from the triangle inequality

$$\|\mathcal{P}_{\tilde{\mathcal{L}}[\Delta\lambda],T} - \mathcal{P}_\pi\|_{\text{TV}} \leq \|\mathcal{P}_{\mathcal{L},T} - \mathcal{P}_\pi\|_{\text{TV}} + \|\mathcal{P}_{\mathcal{L},T} - \mathcal{P}_{\tilde{\mathcal{L}}[\Delta\lambda],T}\|_{\text{TV}}, \quad \text{(A.29)}$$

on which we subsitute the results of *Lemmas 4* and *8* to give the final result (A.28). $\square$

Corollary 11. *Let $n_x \geq 2$, Φ satisfy (A.1), (A.2) and (A.3), and $\varepsilon \in (0, 1/2)$ be a desired precision level. Let the time horizon, T, and the step size, $\Delta\lambda$, be defined by*

$$T \geq \frac{2\log(1/\varepsilon) + \frac{n_x}{2}\log\left(\frac{1}{m\sigma_\nu^2(2 - M\sigma_\nu^2)}\right) + \frac{M}{(2 - M\sigma_\nu^2)}\|\bar{\mathrm{x}} - \mu_\nu\|_2^2}{m}, \quad \text{(A.30)}$$

$$\Delta\lambda \leq \frac{-\frac{T}{16} + \sqrt{\left(\frac{T}{16}\right)^2 + \frac{\gamma}{48(2\gamma - 1)}\left(\sigma_\nu^2 + \frac{1}{n_x}\|\bar{\mathrm{x}} - \mu_\nu\|_2^2 + 2T\right)M^{-1}K^{-2}\left[\frac{2}{n_x}\log\left(\frac{1}{1 - \varepsilon}\right)\right]}}{\frac{\gamma}{48(2\gamma - 1)}\left(\sigma_\nu^2 + \frac{1}{n_x}\|\bar{\mathrm{x}} - \mu_\nu\|_2^2 + 2T\right)MK}, \quad \text{(A.31)}$$

where $\gamma \geq 1$. Then the resulting probability distribution of a Langevin Monte Carlo algorithm that integrates (4.15) after $L = \lceil T/\Delta\lambda \rceil$ steps, satisfies $\|\mathcal{P}_{\tilde{\mathcal{L}}[\Delta\lambda],T} - \mathcal{P}_\pi\|_{TV} \leq \varepsilon$.

Proof. Bound each term on the right-hand side of (A.29) by half of the required precision ε, i.e.,

$$\|\mathcal{P}_{\mathcal{L},T} - \mathcal{P}_\pi\|_{\text{TV}} \leq \varepsilon/2, \ \|\mathcal{P}_{\mathcal{L},T} - \mathcal{P}_{\tilde{\mathcal{L}}[\Delta\lambda],T}\|_{\text{TV}} \leq \varepsilon/2.$$

For the first term:

$$\frac{1}{2}\exp\left\{-\frac{1}{2}mT + \frac{n_x}{4}\log\left(\frac{1}{m\sigma_\nu^2(2 - M\sigma_\nu^2)}\right) + \frac{M}{2(2 - M\sigma_\nu^2)}\|\bar{\mathrm{x}} - \mu_\nu\|_2^2\right\} \leq \frac{\varepsilon}{2},$$
$$-\frac{1}{2}mT + \frac{n_x}{4}\log\left(\frac{1}{m\sigma_\nu^2(2 - M\sigma_\nu^2)}\right) + \frac{M}{2(2 - M\sigma_\nu^2)}\|\bar{\mathrm{x}} - \mu_\nu\|_2^2 \leq \log\varepsilon,$$
$$\frac{1}{2}mT \geq -\log\varepsilon + \frac{n_x}{4}\log\left(\frac{1}{m\sigma_\nu^2(2 - M\sigma_\nu^2)}\right) + \frac{M}{2(2 - M\sigma_\nu^2)}\|\bar{\mathrm{x}} - \mu_\nu\|_2^2,$$
$$T \geq \frac{2\log(1/\varepsilon) + \frac{n_x}{2}\log\left(\frac{1}{m\sigma_\nu^2(2 - M\sigma_\nu^2)}\right) + \frac{M}{(2 - M\sigma_\nu^2)}\|\bar{\mathrm{x}} - \mu_\nu\|_2^2}{m}. \quad \text{(A.32)}$$

And for the second term:

$$\frac{1}{2} - \frac{1}{2}\exp\left\{-\overbrace{\frac{n_x}{2}\frac{M^3K^4\gamma}{48(2\gamma-1)}\left(\sigma_\nu^2 + \frac{1}{n_x}\|\bar{\mathrm{x}} - \mu_\nu\|_2^2 + 2T\right)}^{a\,(>0)}\Delta\lambda^2 - \overbrace{\frac{n_xM^2K^3T}{16}}^{b\,(>0)}\Delta\lambda\right\} \leq \frac{\varepsilon}{2},$$

$$\exp\left\{-a\Delta\lambda^2 - b\Delta\lambda\right\} \geq 1 - \varepsilon,$$

$$-a\Delta\lambda^2 - b\Delta\lambda \geq \log(1-\varepsilon),$$

$$a\Delta\lambda^2 + b\Delta\lambda \leq -\log(1-\varepsilon),$$

$$a\Delta\lambda^2 + b\Delta\lambda + c \leq 0,\ c = \log(1-\varepsilon) = -\log\left(\frac{1}{1-\varepsilon}\right),$$

$$0 < \Delta\lambda \leq \frac{-b + \sqrt{b^2 - 4ac}}{2a},$$

$$\Delta\lambda \leq \frac{-\frac{n_xM^2K^3T}{16} + \sqrt{\left(\frac{n_xM^2K^3T}{16}\right)^2 + 4\frac{n_x}{2}\frac{M^3K^4\gamma}{48(2\gamma-1)}\left(\sigma_\nu^2 + \frac{1}{n_x}\|\bar{\mathrm{x}} - \mu_\nu\|_2^2 + 2T\right)\log\left(\frac{1}{1-\varepsilon}\right)}}{2\frac{n_x}{2}\frac{M^3K^4\gamma}{48(2\gamma-1)}\left(\sigma_\nu^2 + \frac{1}{n_x}\|\bar{\mathrm{x}} - \mu_\nu\|_2^2 + 2T\right)},$$

$$\Delta\lambda \leq \frac{-\frac{T}{16} + \sqrt{\left(\frac{T}{16}\right)^2 + \frac{\gamma}{48(2\gamma-1)}\left(\sigma_\nu^2 + \frac{1}{n_x}\|\bar{\mathrm{x}} - \mu_\nu\|_2^2 + 2T\right)M^{-1}K^{-2}\left[\frac{2}{n_x}\log\left(\frac{1}{1-\varepsilon}\right)\right]}}{\frac{\gamma}{48(2\gamma-1)}\left(\sigma_\nu^2 + \frac{1}{n_x}\|\bar{\mathrm{x}} - \mu_\nu\|_2^2 + 2T\right)MK}. \tag{A.33}$$

$\square$

Proof of Theorem 1

Proof. We follow the same procedure established for proving *Theorem 10*, subsituting the results of *Lemmas 7* and *9* into (A.29) to give

$$\|\mathcal{P}_{\tilde{\mathcal{L}}[\Delta\lambda],T} - \mathcal{P}_\pi\|_{\mathrm{TV}} \leq \frac{1}{2}\exp\left\{-\frac{1}{2}mT + \frac{n_x}{2}\log\left(\frac{M}{m}\right) - \log\left[\Gamma_u\left(\frac{n_x}{2}, \frac{M\|\bar{\mathrm{x}} - \mathrm{x}_\nu\|_2^2}{2}\right)\right]\right\}$$
$$+ \frac{1}{2} - \frac{1}{2}\exp\left\{-\frac{n_x}{2}\frac{M^3K^4\gamma}{48(2\gamma-1)}\left(\frac{1}{n_x}\|\bar{\mathrm{x}} - \mathrm{x}_\nu\|_2^2 + 2T\right)\Delta\lambda^2 - \frac{n_xM^2K^3T}{16}\Delta\lambda\right\}. \tag{A.34}$$

$\square$

Proof of Corollary 2

Proof. We follow the same procedure established for proving *Corollary 11*, bounding each term on the right-hand side of (A.29) in view of (A.34) to obtain

$$T \geq \frac{2\log(1/\varepsilon) + n_x\log\left(\frac{M}{m}\right) - 2\log\left[\Gamma_u\left(\frac{n_x}{2}, \frac{M\|\bar{\mathrm{x}} - \mathrm{x}_\nu\|_2^2}{2}\right)\right]}{m}, \tag{A.35}$$

$$\Delta\lambda \leq \frac{-\frac{T}{16} + \sqrt{\left(\frac{T}{16}\right)^2 + \frac{\gamma}{48(2\gamma-1)}\left(\frac{1}{n_x}\|\bar{\mathrm{x}} - \mathrm{x}_\nu\|_2^2 + 2T\right)M^{-1}K^{-2}\left[\frac{2}{n_x}\log\left(\frac{1}{1-\varepsilon}\right)\right]}}{\frac{\gamma}{48(2\gamma-1)}\left(\frac{1}{n_x}\|\bar{\mathrm{x}} - \mathrm{x}_\nu\|_2^2 + 2T\right)MK}. \tag{A.36}$$

$\square$

A.2 On the Filtering Properties of the Stochastic Particle Flow

Theorem 12. *Define* $\mathrm{x} \in \mathbb{R}^{n_x}$ *to describe an* n_x*-dimensional vector state. Let the vector field* $\mu : \mathbb{R}^{n_x} \to \mathbb{R}^{n_x}$*,* $\mu(\mathrm{x}) \in \mathcal{C}^1(\mathbb{R}^{n_x})$*, be a conservative field, i.e., there exists a scalar potential function* $\psi : \mathbb{R}^{n_x} \to \mathbb{R}$*,* $\psi(\mathrm{x}) \in \mathcal{C}^2(\mathbb{R}^{n_x})$*, such that*

$$\mu(\mathrm{x}) = -\nabla_\mathrm{x}\psi(\mathrm{x}). \tag{A.37}$$

Let $p(\mathrm{x}, \lambda)$ *be the density of an ensemble of particles and, without loss of generality, can be assumed to be a continuous probability density function on* $\mathbb{R}^{n_x}$ *that depends on the pseudo-time variable* $\lambda \in \mathbb{R}$*,* $\lambda \geq 0$*. Set*

$\pi\left(\mathrm{x}\right) \propto e^{-\psi(\mathrm{x})}$ *to be an invariant, locally log-concave probability density to which the density $p\left(\mathrm{x}, \lambda\right)$ is expected to converge weakly at a stationary state in a finite time horizon $\lambda \geq T$, $T \in \mathbb{R}_{+}$, i.e.,*

$$\mathbb{E}_p\left[\varphi\left(\mathrm{x}\right)\right] \to \mathbb{E}_\pi\left[\varphi\left(\mathrm{x}\right)\right], \quad \lambda \to T; \tag{A.38}$$

for all bounded, continuous functions φ, and where $\mathbb{E}_p\left[.\right]$ is the expectation with respect to the probability density $p\left(\mathrm{x}, \lambda\right)$. If the probability density $p\left(\mathrm{x}, \lambda\right)$ satisfies the continuity equation (Liouville's equation)

$$\frac{\partial p}{\partial \lambda} = -\nabla_\mathrm{x} \cdot \left(p\,\mu\right), \quad \lambda \geq 0; \tag{A.39}$$

with the initial condition

$$p\left(\mathrm{x}, \lambda\right) = p_0\left(\mathrm{x}\right), \quad \lambda = 0; \tag{A.40}$$

then any probability mass (particle) $\mathrm{x}_m\left(0\right) \sim p_0\left(\mathrm{x}\right)$, when evolved according to the associated state equation

$$d\mathrm{x}_m(\lambda) = \mu\left(\mathrm{x}_m(\lambda)\right)d\lambda, \quad \lambda \geq 0; \tag{A.41}$$

converges to

$$\mathrm{x}_m\left(T\right) = argmax\left[\pi\left(\mathrm{x}\right)\right], \quad \lambda \geq T, \tag{A.42}$$

at a stable equilibrium.

Proof. The general solution of the continuity equation without sources (A.39) assumes the form (see for example [92])

$$\begin{aligned} p\left(\mathrm{x}, \lambda\right) &= p_0\left(\mathrm{x}_m(\mathrm{x}, \lambda)\right)\left|\frac{\partial \mathrm{x}_m}{\partial \mathrm{x}}\right| \\ &= p_0\left(\mathrm{x}_m(\mathrm{x}, \lambda)\right)\left|\mathcal{J}_\mathrm{x}\left[\mathrm{x}_m(\mathrm{x}, \lambda)\right]\right|, \end{aligned} \tag{A.43}$$

where $\mathrm{x}_m(\mathrm{x}, \lambda)$ is an arbitrary element of mass that is regarded as a function of the pseudo-time λ and of the state x that it can possibly reach. The matrix $\mathcal{J}_\mathrm{x}\left[\mathrm{x}_m(\mathrm{x}, \lambda)\right]$ is the Jacobian matrix of $\mathrm{x}_m(\mathrm{x}, \lambda)$ with respect to x. Conceptually, at the stationary state $\mathrm{x}_m\left(\mathrm{x}_T, T\right) = \mathrm{x}_T$ the continuity equation (A.39) reads

$$\frac{\partial p}{\partial \lambda} = 0, \quad \lambda \geq T. \tag{A.44}$$

Using the general solution (A.43) to verify the stationary condition (A.44), we conclude that

$$p_0\left(\mathrm{x}_m(\mathrm{x}_T, T)\right)\left|\mathcal{J}_\mathrm{x}\left[\mathrm{x}_m(\mathrm{x}_T, T)\right]\right|$$

must be constant with respect to the pseudo-time, thus

$$\frac{d\mathrm{x}_m(\lambda)}{d\lambda} = \mu\left(\mathrm{x}_m(\lambda)\right) = 0, \quad \lambda \geq T. \tag{A.45}$$

Following the assumption of conservative field, $\mu\left(\mathrm{x}_T\right) = -\nabla_\mathrm{x}\psi\left(\mathrm{x}_T\right) = 0$ implies that the stationary state x_T is an equilibrium point, i.e., an extreme of the potential function ψ. In addition, since the potential function is assumed to be related to the stationary distribution as $\psi\left(\mathrm{x}\right) \propto -\log \pi\left(\mathrm{x}\right)$, the stationary state x_T is an extreme of the stationary density.

A valid Lyapunov function of the flow is $V\left(\mathrm{x}\right) = \psi\left(\mathrm{x}\right)$, which is positive semi-definite ($\psi\left(\mathrm{x}\right) \geq 0$) in the neighbourhood of the equilibrium point due to the local log-concavity of the invariant density $\pi\left(\mathrm{x}\right)$. Analysing the (Lie) time derivative of the Lyapunov function in the neighbourhood of the equilibrium point, $\|\mathrm{x} - \mathrm{x}_T\| < \varepsilon$ for a sufficiently small $\varepsilon \in \mathbb{R}_{+}$, we have

$$\begin{aligned} \frac{dV\left(\mathrm{x}\right)}{d\lambda} &= \nabla_\mathrm{x}V\left(\mathrm{x}\right)^T \cdot \frac{d\mathrm{x}}{d\lambda} = \nabla_\mathrm{x}V\left(\mathrm{x}\right)^T \cdot \mu\left(\mathrm{x}\right), \\ \dot{V}\left(\mathrm{x}\right) &= \nabla_\mathrm{x}\psi\left(\mathrm{x}\right)^T \cdot \left(-\nabla_\mathrm{x}\psi\left(\mathrm{x}\right)\right), \\ \dot{V}\left(\mathrm{x}\right) &= -\left\|\nabla_\mathrm{x}\psi\left(\mathrm{x}\right)\right\|^2 \leq 0, \|\mathrm{x} - \mathrm{x}_T\| < \varepsilon; \end{aligned} \tag{A.46}$$

from which we conclude that x_T is a point of (uniformily) stable equilibrium. Therefore, under the established hypotheses, any arbitrary probability mass $\mathrm{x}_m\left(\lambda\right)$ evolved according to (A.41) converges to

$$\begin{aligned} \mathrm{x}_m\left(T\right) &= \operatorname{argmin}\left[\psi\left(\mathrm{x}\right)\right] = \operatorname{argmin}\left[-\log \pi\left(\mathrm{x}\right)\right], \\ \mathrm{x}_m\left(T\right) &= \operatorname{argmax}\left[\pi\left(\mathrm{x}\right)\right], \quad \lambda \geq T; \end{aligned}$$

at a stable equilibrium. $\qquad\square$

Lemma 13. *Let $\{X_\lambda : t \leq \lambda \leq T\}$ be a diffusion process in $\mathbb{R}^{n_x}$ (hence a Markov process), solution of*

$$dX_\lambda = \mu_f(X_\lambda, \lambda)\, d\lambda + D_f(\lambda)^{1/2}\, dW_\lambda, \tag{A.47}$$
$$X_\lambda(t) = x_t, t \leq \lambda \leq T;$$

where $\{W_\lambda : t \leq \lambda \leq T\}$ is a standard Wiener process in $\mathbb{R}^{n_x}$ under the probability measure $\mathcal{P}$, $\mu_f : \mathbb{R}^{n_x} \times [t, T] \to \mathbb{R}^{n_x}$ is the drift and $D_f : [t, T] \to \mathbb{R}^{n_x \times n_x}$ is a diffusion coefficient invariant over the space at any time instant. There exists an equivalent process $\{\bar{X}_\tau, \mathcal{V}_\tau : t \leq \tau \leq T\}$, which is probabilistically the same as the original process, called reverse process *on the interval $[t, T]$ (see [93]), that provides the solution of the stochastic system*

$$d\bar{X}_\tau = \mu_r(\bar{X}_\tau, \tau)\, d\tau + D_r(\tau)^{1/2}\, d\bar{W}_\tau, \quad \bar{X}_\tau(t) = \bar{x}_t; \tag{A.48}$$
$$d\mathcal{V}_\tau = v_r(\bar{X}_\tau, \tau)\, \mathcal{V}_\tau d\tau, \qquad\qquad \mathcal{V}_\tau(t) = 1; \tag{A.49}$$

for a standard Wiener process $\{\bar{W}_\tau : t \leq \tau \leq T\}$ in $\mathbb{R}^{n_x}$ under the measure $\mathcal{P}$, with the reverse drift and diffusion coefficients given, respectively, by

$$\mu_r(\bar{X}_\tau, \tau) = -\mu_f(\bar{X}_\tau, T + t - \lambda), \tag{A.50}$$
$$D_r(\tau) = D_f(T + t - \lambda). \tag{A.51}$$

Proof. The Markov process $\{X_\lambda\}$, as an existing solution to the SDE (A.47), has an associated probability density $p(x_\lambda, \lambda)$ that must satisfy the Kolmogorov forward equation (Fokker-Planck equation):

$$\frac{\partial}{\partial\lambda} p = -\nabla_x \cdot (\mu_f p) + \frac{1}{2}\nabla_x \cdot (D_f \nabla_x p), \; \lambda \geq t,$$
$$p(x_\lambda, t) = p_t(x_t), x_\lambda \in \mathbb{R}^{n_x}.$$

The Fokker-Planck equation can be written in the non-divergence form as

$$\frac{\partial}{\partial\lambda} p = \hat{\mu}^T \nabla_x p + \frac{1}{2}\nabla_x \cdot (D_f \nabla_x p) + \hat{v} \cdot p, \tag{A.52}$$

where

$$\hat{\mu}(x_\lambda) = -\mu_f(x_\lambda),$$
$$\hat{v}(x_\lambda) = -\nabla_x \cdot \mu_f(x_\lambda).$$

We introduce the reverse time variable $\tau = T + t - \lambda$, so that

$$p(x_\lambda(T + t - \lambda), T + t - \lambda) \equiv \hat{p}(\bar{x}_\tau, \tau),$$

and hence $-\partial_\lambda p = \partial_\tau \hat{p}$. Thus, rewritting (A.52) with respect to $\hat{p}(\bar{x}_\tau, \tau)$ for $\tau \leq T$, $\bar{x}_\tau \in \mathbb{R}^{n_x}$, and performing the substitutions

$$\mu_r(\bar{x}_\tau, \tau) = -\mu_f(x_\lambda(T + t - \lambda), T + t - \lambda),$$
$$D_r(\tau) = D_f(T + t - \lambda),$$

we obtain

$$\frac{\partial}{\partial\lambda} p = -\mu_f^T \nabla_x p + \frac{1}{2}\nabla_x \cdot (D_f \nabla_x p) + (-\nabla_x \cdot \mu_f) \cdot p,$$
$$-\frac{\partial}{\partial\tau} \hat{p} = \mu_r^T \nabla_{\bar{x}} \hat{p} + \frac{1}{2}\nabla_{\bar{x}} \cdot (D_r \nabla_{\bar{x}} \hat{p}) - (-\nabla_{\bar{x}} \cdot \mu_r) \cdot \hat{p},$$
$$-\frac{\partial}{\partial\tau} \hat{p} = \mu_r^T \nabla_{\bar{x}} \hat{p} + \frac{1}{2}\nabla_{\bar{x}} \cdot (D_r \nabla_{\bar{x}} \hat{p}) - v_r \cdot \hat{p}, \; \tau \leq T, \tag{A.53}$$
$$\hat{p}(\bar{x}_\tau, T) = \hat{p}_T(\bar{x}_T) = p_t(x_t), \bar{x}_\tau \in \mathbb{R}^{n_x};$$

where

$$v_r(\bar{x}_\tau, \tau) = -\nabla_{\bar{x}} \mu_r(\bar{x}_\tau, \tau). \tag{A.54}$$

Solving (A.53) corresponds to the Cauchy problem in reverse time $\tau \leq T$, which is equivalent to solve the stochastic system stated by (A.48) and (A.49). Therefore, because the solution to the SDE (A.47) is assumed to exist and corresponds to the solution of (A.53) for $\tau \leq T$, then there exists the equivalent reverse process $\{\bar{X}_\tau, \mathcal{V}_\tau\}_{\tau \in [t, T]}$ that solves the stochastic system (A.48) and (A.49). $\qquad\square$

Remark 14. Despite its name, inherited from [93], it is worth stressing that $\left\{\bar{\mathrm{X}}_\tau, \mathcal{V}_\tau\right\}$ is the solution of a stochastic system forward in time on the interval $[t, T]$, which may be properly understood as a smoothing process.

Remark 15. It is clear that the solution of (A.49) at $\tau = T$ is

$$\mathcal{V}_T \equiv \mathcal{V}_\tau (T) = e^{- \int_t^T v_r(\bar{\mathrm{x}}_\tau, \tau) d\tau},$$

which evokes the solution of (A.53) by the Feynman-Kac formula

$$\hat{p}(\mathrm{x}, \tau) = \mathbb{E}^{\mathcal{P}} \left[e^{- \int_\tau^T v_r\left(\bar{\mathrm{x}}_{\tau'}, \tau'\right) d\tau'} \hat{p}_T (\bar{\mathrm{x}}_T) \, | \bar{\mathrm{x}}_\tau = \mathrm{x} \right]. \tag{A.55}$$

A more general form of the *Lemma 13* can be found in [93].

Lemma 16. *The reverse process $\left\{\bar{\mathrm{X}}_\tau, \mathcal{V}_\tau : t \leq \tau \leq T\right\}$ described by (A.48) and (A.49), has an associated smooth probability density $\hat{p}(\bar{\mathrm{x}}_\tau, \tau)$ that satisfies, for the initial value problem, the Kolmogorov forward equation*

$$\frac{\partial}{\partial \tau} \hat{p} = -\nabla_{\bar{\mathrm{x}}} \cdot (\mu_r \hat{p}) + \frac{1}{2} \nabla_{\bar{\mathrm{x}}} \cdot (D_r \nabla_{\bar{\mathrm{x}}} \hat{p}) - v_r \cdot \hat{p}, \ t \leq \tau \leq T, \tag{A.56}$$

$$\hat{p}(\bar{\mathrm{x}}_\tau(t), \tau = t) = \hat{p}_t (\bar{\mathrm{x}}_t), \bar{\mathrm{x}}_\tau \in \mathbb{R}^{n_x}.$$

Proof. We consider a continuous function of the process $\left\{\bar{\mathrm{X}}_\tau, \mathcal{V}_\tau\right\}_{\tau \in [t, T]}$, declared as $\varphi : \mathbb{R}^{n_x} \times \mathbb{R} \to \mathbb{R}$, which is assumed to be $\varphi(\bar{\mathrm{X}}_\tau, \mathcal{V}_\tau) \in \mathcal{C}^2(\mathbb{R}^{n_x}, \mathbb{R})$, bounded and integrable on the product space $\mathbb{R}^{n_x} \times \mathbb{R}$. Applying Îto's lemma to φ and substituting (A.48) and (A.49) one obtains

$$d\varphi = (\nabla_{\bar{\mathrm{X}}_\tau} \varphi)^T d\bar{\mathrm{X}}_\tau + \frac{1}{2} d\bar{\mathrm{X}}_\tau^T \mathcal{H}_{\bar{\mathrm{X}}_\tau} [\varphi] d\bar{\mathrm{X}}_\tau + \partial_{\mathcal{V}_\tau} \varphi \, d\mathcal{V}_\tau$$

$$= (\nabla_{\bar{\mathrm{X}}_\tau} \varphi)^T \cdot \left(\mu_r d\tau + D_r^{1/2} d\bar{W}_\tau \right) + \frac{1}{2} \mathrm{tr} \left\{ D_r \mathcal{H}_{\bar{\mathrm{X}}_\tau} [\varphi] \right\} d\tau + \partial_{\mathcal{V}_\tau} \varphi \cdot (v_r \mathcal{V}_\tau d\tau),$$

$$= \left[(\nabla_{\bar{\mathrm{X}}_\tau} \varphi)^T \mu_r + \frac{1}{2} \mathrm{tr} \left\{ D_r \mathcal{H}_{\bar{\mathrm{X}}_\tau} [\varphi] \right\} + \partial_{\mathcal{V}_\tau} \varphi \cdot v_r \mathcal{V}_\tau \right] d\tau + (\nabla_{\bar{\mathrm{X}}_\tau} \varphi)^T D_r^{1/2} d\bar{W}_\tau. \tag{A.57}$$

Consider the expected (average) rate of change of the projections of φ defined by:

$$\langle \dot{\varphi}_\mathcal{V} \rangle (\tau) \triangleq \left\langle \frac{d}{d\tau} \int_1^{\mathcal{V}_T} \varphi(\bar{\mathrm{X}}_\tau, \mathcal{V}_\tau) d\mathcal{V}_\tau \right\rangle$$

$$= \left\langle \int_1^{\mathcal{V}_T} \partial_\tau \varphi(\bar{\mathrm{X}}_\tau, \mathcal{V}_\tau) d\mathcal{V}_\tau \right\rangle$$

$$\equiv \int_{\mathbb{R}^{n_x}} \int_1^{v_T} \varphi(\bar{\mathrm{x}}_\tau, v_\tau) \cdot \partial_\tau \hat{p} \, dv_\tau d\bar{\mathrm{x}}_\tau. \tag{A.58}$$

Substituting (A.57) into (A.58), we have

$$\langle \dot{\varphi}_\mathcal{V} \rangle (\tau) = \left\langle \int_1^{\mathcal{V}_T} \partial_\tau \varphi(\bar{\mathrm{X}}_\tau, \mathcal{V}_\tau) d\mathcal{V}_\tau \right\rangle$$

$$= \left\langle \int_1^{\mathcal{V}_T} \left[(\nabla_{\bar{\mathrm{X}}_\tau} \varphi)^T \mu_r + \frac{1}{2} \mathrm{tr} \left\{ D_r \mathcal{H}_{\bar{\mathrm{X}}_\tau} [\varphi] \right\} + \partial_{\mathcal{V}_\tau} \varphi \cdot v_r \mathcal{V}_\tau \right] d\mathcal{V}_\tau \right\rangle$$

$$= \int_1^{v_T} \int_{\mathbb{R}^{n_x}} \left[(\nabla_{\bar{\mathrm{x}}_\tau} \varphi)^T \mu_r + \frac{1}{2} \mathrm{tr} \left\{ D_r \mathcal{H}_{\bar{\mathrm{x}}_\tau} [\varphi] \right\} \right] \hat{p} \, d\bar{\mathrm{x}}_\tau dv_\tau$$

$$+ \int_{\mathbb{R}^{n_x}} \int_1^{v_T} \left[\partial_{v_\tau} \varphi \cdot v_r \, v_\tau \right] \hat{p} \, dv_\tau d\bar{\mathrm{x}}_\tau. \tag{A.59}$$

Integrating (A.59) by parts and offsetting the integration constant to cancel out the surface terms, we have

$$
\begin{aligned}
\langle \dot{\varphi}_{\mathcal{V}} \rangle (\tau) &= \int_1^{\upsilon_T} \int_{\mathbb{R}^{n_x}} \varphi \cdot \left[-\nabla_{\bar{x}_\tau} \cdot (\mu_r \hat{p}) + \frac{1}{2} \nabla_{\bar{x}_\tau} \cdot (D_r \nabla_{\bar{x}_\tau} \hat{p}) \right] d\bar{x}_\tau d\upsilon_\tau \\
&\quad + \int_{\mathbb{R}^{n_x}} \int_1^{\upsilon_T} \varphi \cdot \left[-\upsilon_r \partial_{\upsilon_\tau} (\upsilon_\tau \hat{p}) \right] d\upsilon_\tau d\bar{x}_\tau \\
&= \int_{\mathbb{R}^{n_x}} \int_1^{\upsilon_T} \varphi \cdot \left[-\nabla_{\bar{x}_\tau} \cdot (\mu_r \hat{p}) + \frac{1}{2} \nabla_{\bar{x}_\tau} \cdot (D_r \nabla_{\bar{x}_\tau} \hat{p}) \right] d\upsilon_\tau d\bar{x}_\tau \\
&\quad + \int_{\mathbb{R}^{n_x}} \int_1^{\upsilon_T} \varphi \cdot \left[-\upsilon_r \hat{p} \right] d\upsilon_\tau d\bar{x}_\tau \\
&= \int_{\mathbb{R}^{n_x}} \int_1^{\upsilon_T} \varphi \cdot \left[-\nabla_{\bar{x}_\tau} \cdot (\mu_r \hat{p}) + \frac{1}{2} \nabla_{\bar{x}_\tau} \cdot (D_r \nabla_{\bar{x}_\tau} \hat{p}) - \upsilon_r \hat{p} \right] d\upsilon_\tau d\bar{x}_\tau.
\end{aligned}
\tag{A.60}
$$

The proof of the lemma is complete by comparing (A.60) to the definition (A.58) and noting that their integrands must be equal. $\qquad\square$

Theorem 17. *Let $(\Omega, \mathcal{F}, \mathcal{P})$ to be a complete probability space and let $\{\mathcal{F}_\lambda\}_{\lambda \geq 0}$, $\lambda \in [0, T]$, be an increasing family of sub σ-fields of $\mathcal{F}$. Let $\{X_\lambda : 0 < \lambda \leq T\}$ be an $\mathcal{F}_\lambda$-adapted process, considered to be the signal process with state equation*

$$
\begin{aligned}
dX_\lambda &= \mu_f (X_\lambda) \, d\lambda + D_f (\lambda)^{1/2} \, dW_\lambda, \\
X_\lambda (0) &= x_0, \, 0 \leq \lambda \leq T;
\end{aligned}
\tag{A.61}
$$

for a Wiener process $\{W_\lambda\}_{\lambda \in [0, T]}$ under the probability measure $\mathcal{P}$. Assume $p(x_\lambda, \lambda)$ defined on $(\Omega, \mathcal{F})$ to be the probability density of the measure $\mathcal{P}$, which

(a) is the probabilistic representation of the process $\{X_\lambda\}_{\lambda \in [0,T]}$,

(b) is absolutely continuous with respect to the Lesbesque measure,

(c) approaches a stationary measure as $p(x_\lambda) \xrightarrow{\lambda \to T} \pi (x_\lambda) \propto e^{-\Phi(x_\lambda)}$, for a sufficiently long horizon T.

Let $\{\bar{X}_\tau, \mathcal{V}_\tau : \lambda < \tau \leq T\}$ be the reverse process of $\{X_\lambda\}_{\lambda \in [0,T]}$, as established in Lemma 13 by the stochastic system (A.48) and (A.49), so that the reverse drift and diffusion coefficients are, respectively,

$$
\begin{aligned}
\mu_r (\bar{x}_\tau) &= -\frac{1}{2} D_r (\lambda) \nabla_{\bar{x}} \log \pi (\bar{x}_\tau), \\
D_r (\tau) &= \left[-\mathcal{H}_{\bar{x}} \left[\log \pi (\bar{x}_\tau) \right] \right]^{-1}.
\end{aligned}
$$

Assume $\hat{p}(\bar{x}_\tau, \tau)$ to be the probability density on $(\Omega, \mathcal{F})$ describing the reverse process $\{\bar{X}_\tau, \mathcal{V}_\tau\}_{\tau \in [\lambda, T]}$, under the same measure $\mathcal{P}$, which must satisfy the Kolmogorov forward equation (provided a known initial condition) according to Lemma 16:

$$
\begin{aligned}
\frac{\partial}{\partial \tau} \hat{p} &= -\nabla_{\bar{x}} \cdot (\mu_r \hat{p}) + \frac{1}{2} \nabla_{\bar{x}} \cdot (D_r \nabla_{\bar{x}} \hat{p}) - \upsilon_r \cdot \hat{p}, \, \lambda \leq \tau \leq T, \\
\hat{p} (\bar{x}_\tau (\lambda), \tau = \lambda) &= \hat{p}_\lambda (\bar{x}_\tau (\lambda)) = \pi (x_T), \bar{x}_\tau \in \Omega;
\end{aligned}
\tag{A.62}
$$

If the stationary density is set to be

$$
\pi (x) := \frac{p (y_k | x) \, p (x | y_{1:k-1})}{Z_1} = \frac{p (y_k | x) \, p_x (x)}{Z_1},
\tag{A.63}
$$

where the prior density $p_x (x)$ and the likelihood $p (y_k | x)$ are integrable functions with respect to x, $Z_1 = p (y_k | y_{1:k-1})$ is a normalization constant, and the discrete-time observation process $\{y_k \in \mathbb{R}^{n_y} : k \in \mathbb{N}\}$ is described as

$$
y_k = h (x_k) + R^{1/2} \xi_k, \quad \xi_k \sim \mathcal{N}(\xi_k; 0_{n_y}, I_{n_y});
\tag{A.64}
$$

then the probability density corresponding to the signal process (A.61) is equivalent to the following filtering entity

$$
p(x, \lambda | \mathcal{F}_\lambda) = \frac{\mathbb{E}^P \left[e^{h(x_T)^T R^{-1} y_k - \frac{1}{2} h(x_T)^T R^{-1} h(x_T)} | x \right] p (x | y_{1:k-1})}{Z}.
\tag{A.65}
$$

In addition, the expression (A.65) can be interpreted as the analogous to the well known result (see [94])

$$p(\mathrm{x}, \lambda | \mathcal{F}_\lambda) = \frac{\mathbb{E}^P\left[e^{\int_0^T h_T(\mathrm{x}_\lambda)^T R_T^{-1} dy_\lambda - \frac{1}{2}\int_0^T h_T(\mathrm{x}_\lambda)^T R_T^{-1} h_T(\mathrm{x}_\lambda) d\lambda} | \mathrm{x}\right]}{Z} p(\mathrm{x}|y_{1:k-1}), \qquad (A.66)$$

for a discrete-time observation process whose analog continuous-time (interpolated) version has the observation function $h_T(.)$ and covariance matrix R_T.

Proof. By definition of the stochastic system described by (A.48) and (A.49), and *Lemma 16*, the reverse process $\{\bar{X}_\tau, \mathcal{V}_\tau\}$ is known to satisfy the Kolmogorov forward equation (A.62) in reverse time $\lambda \leq \tau \leq T$, for which the stationary density π is an initial condition (initial value problem). Using the reverse time variable $\tau = T - \lambda$, so that $p(\mathrm{x}_\lambda, \lambda) \equiv \hat{p}(\bar{\mathrm{x}}_\tau(T-\tau), T-\tau)$ and $\partial_\lambda p = -\partial_\tau \hat{p}$, and applying the relations (A.50), (A.51) and (A.54) from *Lemma 13*, we rewrite the equation (A.62) for $0 \leq \lambda \leq T$ as

$$\frac{\partial}{\partial \tau}\hat{p} = -\nabla_{\bar{\mathrm{x}}} \cdot (\mu_r \hat{p}) + \frac{1}{2}\nabla_{\bar{\mathrm{x}}} \cdot (D_r \nabla_{\bar{\mathrm{x}}} \hat{p}) - v_r \cdot \hat{p}, \ \lambda \leq \tau \leq T,$$

$$\frac{\partial}{\partial \tau}\hat{p} = -(\nabla_{\bar{\mathrm{x}}} \cdot \mu_r)\hat{p} - \mu_r^T \nabla_{\bar{\mathrm{x}}} \hat{p} + \frac{1}{2}\nabla_{\bar{\mathrm{x}}} \cdot (D_r \nabla_{\bar{\mathrm{x}}} \hat{p}) + (\nabla_{\bar{\mathrm{x}}} \cdot \mu_r)\hat{p},$$

$$\frac{\partial}{\partial \tau}\hat{p} = -\mu_r^T \nabla_{\bar{\mathrm{x}}} \hat{p} + \frac{1}{2}\nabla_{\bar{\mathrm{x}}} \cdot (D_r \nabla_{\bar{\mathrm{x}}} \hat{p}),$$

$$-\frac{\partial}{\partial \lambda}p = +\mu_f^T \nabla_{\mathrm{x}} p + \frac{1}{2}\nabla_{\mathrm{x}} \cdot (D_f \nabla_{\mathrm{x}} p), \ 0 \leq \lambda \leq T, \qquad (A.67)$$

$$p(\mathrm{x}_\lambda, T) = \pi(\mathrm{x}_T), \mathrm{x}_\lambda \in \Omega.$$

Now we have a Kolmogorov backward equation in $p(\mathrm{x}_\lambda, \lambda)$ with a terminal value problem for the forward process $\{X_\lambda\}$. Hence, we can apply the Feynman-Kac formula for the terminal condition $p(\mathrm{x}, T) = \pi(\mathrm{x})$, with $\mathrm{x}_T = \mathrm{x}_\lambda(T)$, to give

$$p(\mathrm{x}, \lambda) \triangleq \mathbb{E}^{\mathcal{P}}[\pi(\mathrm{x}_T)|\mathrm{x}_\lambda = \mathrm{x}]$$

$$= \frac{\mathbb{E}^{\mathcal{P}}[p(y_k|\mathrm{x}_T)p(\mathrm{x}_T|y_{1:k-1})|\mathrm{x}_\lambda = \mathrm{x}]}{Z_1}$$

$$= \frac{\mathbb{E}^{\mathcal{P}}\left[e^{-\frac{1}{2}(y_k - h(\mathrm{x}_T))^T R^{-1}(y_k - h(\mathrm{x}_T))}p_x(\mathrm{x}_T)|\mathrm{x}\right]}{Z_1(2\pi R)^{n_y/2}}$$

$$= \frac{\mathbb{E}^{\mathcal{P}}\left[e^{h(\mathrm{x}_T)^T R^{-1} y_k - \frac{1}{2}h(\mathrm{x}_T)^T R^{-1} h(\mathrm{x}_T)}p_x(\mathrm{x}_T)|\mathrm{x}\right]}{Z_1(2\pi R)^{n_y/2}e^{+\frac{1}{2}y_k^T R^{-1} y_k}}$$

$$= \frac{\mathbb{E}^{\mathcal{P}}\left[e^{h(\mathrm{x}_T)^T R^{-1} y_k - \frac{1}{2}h(\mathrm{x}_T)^T R^{-1} h(\mathrm{x}_T)}p_x(\mathrm{x}_T)|\mathcal{F}_\lambda, \mathrm{x}\right]}{Z}. \qquad (A.68)$$

Let us reinterpret the discrete-time observation process as a continuous-time process for which we only obtain a realization at $\lambda = T$, by linearly interpolating it along the interval $0 < \lambda \leq T$ to write

$$dy_\lambda = \frac{1}{T}h(\mathrm{x}_\lambda)d\lambda + \left(\frac{R}{T}\right)^{1/2}\xi_k d\lambda^{1/2},$$

$$dy_\lambda^{\mathcal{Q}} = h_T(\mathrm{x}_\lambda)d\lambda + R_T^{1/2}d\bar{\xi}_\lambda^{\mathcal{P}}, \quad 0 < \lambda \leq T. \qquad (A.69)$$

where $\{\bar{\xi}_\lambda\}$ is an interpolated Wiener process, under the probability measure $\mathcal{P}$, that produces the observation noise $\int_0^T R_T^{1/2}d\bar{\xi}_k \equiv R^{1/2}\xi_k$ at $\lambda = T$. The probability measure $\mathcal{P}$ is induced in the space of paths jointly described by the state process and observation noise $(\{X_\lambda\}, \{\bar{\Xi}_\lambda\})$. The interpolation is established such that $y_\lambda(T) = y_k$ is the realization of the observation process under the probability measure $\mathcal{Q}$, which is induced in the space of paths jointly described by the state and observation processes $(\{X_\lambda\}, \{Y_\lambda\})$. By applying the Girsanov theorem, we know that the Radon-Nykodym derivative to change the measure from $\mathcal{P}$ to $\mathcal{Q}$ assumes the form (see [13] for example)

$$\left.\frac{d\mathcal{Q}}{d\mathcal{P}}\right|_{\mathcal{F}_T} = e^{\int_0^T h_T(\mathrm{x}_\lambda)^T R_T^{-1} dy_\lambda - \frac{1}{2}\int_0^T h_T(\mathrm{x}_\lambda)^T R_T^{-1} h_T(\mathrm{x}_\lambda)d\lambda}$$

$$\propto e^{h(\mathrm{x}_T)^T R^{-1} y_k - \frac{1}{2}h(\mathrm{x}_T)^T R^{-1} h(\mathrm{x}_T)}. \qquad (A.70)$$

Rewritting (A.68) in terms of (A.70) and manipulating it further, we obtain

$$
p(\mathrm{x}, \lambda) \propto \mathbb{E}^{\mathcal{P}} \left[e^{h(\mathrm{x}_T)^T R^{-1} \mathrm{y}_k - \frac{1}{2} h(\mathrm{x}_T)^T R^{-1} h(\mathrm{x}_T)} p_x \left(\mathrm{x}_T \right) | \mathcal{F}_\lambda, \mathrm{x} \right]
$$

$$
\equiv \mathbb{E}^{\mathcal{P}} \left[\frac{d\mathcal{Q}}{d\mathcal{P}} \left(T \right) \cdot p_x \left(\mathrm{x}_T \right) | \mathcal{F}_\lambda, \mathrm{x} \right]
$$

$$
= \mathbb{E}^{\mathcal{P}} \left[\frac{d\mathcal{Q}}{d\mathcal{P}} \left(T \right) \cdot \mathbb{E}^{\mathcal{Q}} \left[p_x \left(\mathrm{x}_\lambda \right) | \mathcal{F}_T, \mathrm{x} \right] | \mathcal{F}_\lambda, \mathrm{x} \right]
$$

$$
= \mathbb{E}^{\mathcal{P}} \left[\frac{d\mathcal{Q}}{d\mathcal{P}} \left(T \right) | \mathcal{F}_\lambda, \mathrm{x} \right] \cdot \mathbb{E}^{\mathcal{Q}} \left[p_x \left(\mathrm{x}_\lambda \right) | \mathcal{F}_T, \mathrm{x} \right]
$$

$$
\equiv \mathbb{E}^{\mathcal{P}} \left[e^{h(\mathrm{x}_T)^T R^{-1} \mathrm{y}_k - \frac{1}{2} h(\mathrm{x}_T)^T R^{-1} h(\mathrm{x}_T)} | \mathcal{F}_\lambda, \mathrm{x} \right] p_x \left(\mathrm{x} \right), \tag{A.71}
$$

where we take into account the smoothing property for conditional expectations as

$$
p(\mathrm{x}, \lambda) \propto \mathbb{E}^{\mathcal{P}} \left[\frac{d\mathcal{Q}}{d\mathcal{P}} \left(T \right) \cdot p_x \left(\mathrm{x}_T \right) | \mathcal{F}_\lambda, \mathrm{x}_\lambda = \mathrm{x} \right]
$$

$$
= \mathbb{E}^{\mathcal{Q}} \left[p_x \left(\mathrm{x}_T \right) | \mathcal{F}_\lambda, \mathrm{x} \right]
$$

$$
= \mathbb{E}^{\mathcal{Q}} \left[\mathbb{E}^{\mathcal{Q}} \left[p_x \left(\mathrm{x}_\lambda \right) | \mathcal{F}_T, \mathrm{x} \right] | \mathcal{F}_\lambda, \mathrm{x} \right]
$$

$$
= \mathbb{E}^{\mathcal{P}} \left[\frac{d\mathcal{Q}}{d\mathcal{P}} \left(T \right) \cdot \mathbb{E}^{\mathcal{Q}} \left[p_x \left(\mathrm{x}_\lambda \right) | \mathcal{F}_T, \mathrm{x} \right] | \mathcal{F}_\lambda, \mathrm{x} \right],
$$

the $\mathcal{F}_\lambda$-measurability of $\mathbb{E}^{\mathcal{Q}} \left[p_x \left(\mathrm{x}_\lambda \right) | \mathcal{F}_T, \mathrm{x}_\lambda = \mathrm{x} \right]$, and

$$
\mathbb{E}^{\mathcal{Q}} \left[p_x \left(\mathrm{x}_\lambda \right) | \mathcal{F}_T, \mathrm{x}_\lambda = \mathrm{x} \right] = \int p_x \left(\mathrm{x} \right) d\mathcal{Q}
$$

$$
= p_x \left(\mathrm{x} \right) = p \left(\mathrm{x} | \mathrm{y}_{1:k-1} \right).
$$

As a result, the expression (A.71) can be written in the normalized form (A.65), proving the theorem statement. The proof is complete by inserting the continuous-time (interpolated) version of (A.70) into (A.71) to verify the analogy with (A.66). □

Remark 18. A result more general than the one presented by *Theorem 17*, in terms of McKean-Vlasov diffusions, can be found in [19].

Corollary 19. *The signal process with state equation (A.61), under the hypotheses of Theorem 17, filters its associated (unnormalized) probability density in accordance with the Zakai equation*

$$
dp_u = \mathcal{L} \left[p_u \right] d\lambda + p_u \cdot h_T \left(\mathrm{x}_\lambda \right)^T R_T^{-1} d\mathrm{y}_\lambda, \quad 0 < \lambda \le T; \tag{A.72}
$$

where $\mathcal{L} [.] = -\nabla_{\mathrm{x}} \cdot (\mu \cdot) + 1/2 \nabla_{\mathrm{x}} \cdot (D \nabla_{\mathrm{x}}(\cdot))$ is the forward Kolmogorov operator, and $\{\mathrm{y}_\lambda : 0 < \lambda \le T\}$ is the continuous, linearly interpolated observation process defined by (A.69) for which the realization is only taken at $\lambda = T$.

Proof. Define

$$
d\zeta_\lambda = h_T \left(\mathrm{x}_\lambda \right)^T R_T^{-1} d\mathrm{y}_\lambda - \frac{1}{2} h_T \left(\mathrm{x}_\lambda \right)^T R_T^{-1} h_T \left(\mathrm{x}_\lambda \right) d\lambda
$$

$$
= \frac{1}{2} h_T \left(\mathrm{x}_\lambda \right)^T R_T^{-1} h_T \left(\mathrm{x}_\lambda \right) d\lambda + h_T \left(\mathrm{x}_\lambda \right)^T R_T^{-1/2} d\bar{\xi}_\lambda, \tag{A.73}
$$

and recognize the unnormalized probability density to be the numerator of (A.66):

$$
p_u = \mathbb{E}^{\mathcal{P}} \left[e^{\zeta_T} | \mathrm{x} \right] p_x \left(\mathrm{x} \right). \tag{A.74}
$$

Applying Îto's Lemma to p_u we get

$$
dp_u = \partial_\lambda p_u d\lambda + \partial_\zeta p_u d\zeta_\lambda
$$

$$
+ \frac{1}{2} \left[h_T \left(\mathrm{x}_\lambda \right)^T R_T^{-1/2} \right] \left[R_T^{-1/2} h_T \left(\mathrm{x}_\lambda \right) \right] \partial_{\zeta\zeta}^2 p_u d\lambda
$$

$$
= \partial_\lambda p_u d\lambda + \partial_\zeta p_u d\zeta_\lambda
$$

$$
+ \frac{1}{2} h_T \left(\mathrm{x}_\lambda \right)^T R_T^{-1} h_T \left(\mathrm{x}_\lambda \right) \partial_{\zeta\zeta}^2 p_u d\lambda. \tag{A.75}
$$

Because

$$\partial_\zeta p_u = \partial_{\zeta\zeta}^2 p_u = p_u,$$
$$\partial_\lambda p_u = \mathbb{E}^{\mathcal{P}}\left[e^{-\int_0^T d\zeta_\lambda}|\mathrm{x}\right]\partial_\lambda p_x(\mathrm{x}) = \mathcal{L}[p_u];$$

the expression (A.75) becomes the Zakai equation as

$$dp_u = \mathcal{L}[p_u]\, d\lambda$$
$$+ p_u \cdot \left[\frac{1}{2}h_T(\mathrm{x}_\lambda)^T R_T^{-1}h_T(\mathrm{x}_\lambda)\, d\lambda + h_T(\mathrm{x}_\lambda)^T R_T^{-1/2}d\bar\xi_\lambda\right]$$
$$+ p_u \cdot \frac{1}{2}h_T(\mathrm{x}_\lambda)^T R_T^{-1}h_T(\mathrm{x}_\lambda)\, d\lambda$$
$$= \mathcal{L}[p_u]\, d\lambda$$
$$+ p_u \cdot \left[h_T(\mathrm{x}_\lambda)^T R_T^{-1}h_T(\mathrm{x}_\lambda)\, d\lambda + h_T(\mathrm{x}_\lambda)^T R_T^{-1/2}d\bar\xi_\lambda\right]$$
$$= \mathcal{L}[p_u]\, d\lambda + p_u \cdot h_T(\mathrm{x}_\lambda)^T R_T^{-1}\left[h_T(\mathrm{x}_\lambda)\, d\lambda + R_T^{1/2}d\bar\xi_\lambda\right]$$
$$= \mathcal{L}[p_u]\, d\lambda + p_u \cdot h_T(\mathrm{x}_\lambda)^T R_T^{-1}d\mathrm{y}_\lambda.$$

$\square$

B Derivation of the integration rule

We intend to approximate the integration of the following equation with respect to λ:

$$dx = \frac{1}{2}D(\lambda) \cdot \nabla_{\mathrm{x}}\log\pi(\mathrm{x})\, d\lambda + D(\lambda)^{1/2}\, d\mathrm{w}_\lambda. \tag{B.1}$$

Linearising equation (B.1) w.r.t. x around the current state x_{n-1}, we have

$$dx = A \cdot \mathrm{x}\, d\lambda + B\, d\lambda + D^{1/2}d\mathrm{w}_\lambda, \tag{B.2}$$

where

$$A(\mathrm{x}_{n-1}) = \text{\textonehalf} D(\lambda_{n-1}) \cdot \mathcal{H}_{\mathrm{x}}[\log\pi(\mathrm{x})]_{\mathrm{x}_{n-1}},$$
$$B(\mathrm{x}_{n-1}) = a(\mathrm{x}_{n-1}) - A \cdot \mathrm{x}_{n-1}, \tag{B.3}$$
$$a(\mathrm{x}_{n-1}) = \text{\textonehalf} D(\lambda_{n-1}) \cdot \nabla_{\mathrm{x}}\log\pi(\mathrm{x})|_{\mathrm{x}_{n-1}}.$$

If we apply the definition $D(\lambda_{n-1}) = -\mathcal{H}_{\mathrm{x}}[\log\pi(\mathrm{x})]_{\mathrm{x}_{n-1}}^{-1}$, where $\mathcal{H}_{\mathrm{x}}[\cdot]$ is the Hessian w.r.t. x, we have

$$A = -\frac{1}{2}\mathbb{I}_{n_x}, \tag{B.4}$$

where $\mathbb{I}_{n_x}$ is the identity matrix with dimension $n_x \times n_x$. Based on the Laplace transform, we can obtain the solution for a homogeneous version of the equation (B.2) in discrete time by

$$\mathrm{x}(\lambda) = \mathcal{L}^{-1}\left\{(s \cdot \mathbb{I}_{n_x} - A)^{-1}\mathrm{x}(\lambda_{n-1})\right\}$$
$$= \int_{\lambda_{n-1}}^{\lambda}(s \cdot \mathbb{I}_{n_x} - A)^{-1}e^{s \cdot \mathbb{I}_{n_x} \cdot \tau}\mathrm{x}(\lambda_{n-1})\, ds$$
$$= e^{A \cdot (\lambda - \lambda_{n-1})}\mathrm{x}(\lambda_{n-1}),$$
$$\mathrm{x}_n = e^{A \cdot \Delta\lambda}\mathrm{x}_{n-1}. \tag{B.5}$$

By a similar procedure, and considering the definition of a Wiener integral for the stochastic term, we can

obtain the solution of the complete inhomogeneous equation (B.2) as

$$\mathrm{x}\,(\lambda_n) = e^{A\cdot\Delta\lambda}\mathrm{x}\,(\lambda_{n-1}) + \int_{\lambda_{n-1}}^{\lambda_n} e^{-A\cdot(\tau-\Delta\lambda)}B\,d\tau$$

$$+ \sqrt{\int_{\lambda_{n-1}}^{\lambda_n} e^{-A\cdot(\tau-\Delta\lambda)}D\,(\lambda_{n-1})\,e^{-A^T\cdot(\tau-\Delta\lambda)}d\tau}\cdot\mathrm{w}_n,$$

$$\mathrm{x}_n = e^{A\cdot\Delta\lambda}\mathrm{x}_{n-1} + A^{-1}\left[e^{A\cdot\Delta\lambda}-\mathbb{I}_{n_x}\right]\left[a - A\cdot\mathrm{x}_{n-1}\right]$$

$$+ \sqrt{\int_0^{\Delta\lambda} e^{A\cdot\nu}D\,(\lambda_{n-1})\,e^{A^T\cdot\nu}d\nu}\cdot\mathrm{w}_n$$

$$= \mathrm{x}_{n-1} + A^{-1}\left[e^{A\cdot\Delta\lambda}-\mathbb{I}_{n_x}\right]a\,(\mathrm{x}_{n-1})$$

$$+ \sqrt{\int_0^{\Delta\lambda} e^{A\cdot\nu}D\,(\lambda_{n-1})\,e^{A^T\cdot\nu}d\nu}\cdot\mathrm{w}_n, \tag{B.6}$$

where $\mathrm{w}_n \sim \mathcal{N}(\mathrm{w};0_{n_x},\mathbb{I}_{n_x})$. Substituting (B.4) into (B.6), we have

$$\mathrm{x}_n = \mathrm{x}_{n-1} + \left[-\frac{1}{2}\mathbb{I}_{n_x}\right]^{-1}\left[e^{-\frac{1}{2}\mathbb{I}_{n_x}\cdot\Delta\lambda}-\mathbb{I}_{n_x}\right]a\,(\mathrm{x}_{n-1})$$

$$+ \sqrt{\int_0^{\Delta\lambda} e^{-\frac{1}{2}\mathbb{I}_{n_x}\cdot\nu}D\,(\lambda_{n-1})\,e^{-\frac{1}{2}\mathbb{I}_{n_x}^T\cdot\nu}d\nu}\cdot\mathrm{w}_n. \tag{B.7}$$

By noticing that

$$e^{-\frac{1}{2}\mathbb{I}_{n_x}\Delta\lambda} = e^{-\frac{1}{2}\Delta\lambda}\mathbb{I}_{n_x}, \tag{B.8}$$

the equation (B.7) can be simplified as

$$\mathrm{x}_n = \mathrm{x}_{n-1} - 2\left(e^{-\frac{\Delta\lambda}{2}}-1\right)a\,(\mathrm{x}_{n-1})$$

$$+ \sqrt{\int_0^{\Delta\lambda} e^{-\nu/2}e^{-\nu/2}d\nu}\cdot D\,(\lambda_{n-1})^{1/2}\cdot\mathrm{w}_n$$

$$= \mathrm{x}_{n-1} + 2\left(1-e^{-\frac{\Delta\lambda}{2}}\right)a\,(\mathrm{x}_{n-1})$$

$$+ \sqrt{\int_0^{\Delta\lambda} e^{-\nu}d\nu}\cdot D\,(\lambda_{n-1})^{1/2}\cdot\mathrm{w}_n,$$

$$\mathrm{x}_n = \mathrm{x}_{n-1} + \left(1-e^{-\frac{\Delta\lambda}{2}}\right)D\,(\lambda_{n-1})\cdot\nabla_{\mathrm{x}}\log\pi\,(\mathrm{x}_{n-1})$$

$$+ \left(1-e^{-\Delta\lambda}\right)^{1/2}D\,(\lambda_{n-1})^{1/2}\cdot\mathrm{w}_n. \tag{B.9}$$

C Justification for the local flow linearization

When approximating the stochastic particle flow (4.15) as locally linear in the neighbourhood of a probability mass located at x_l, we expect to produce a negligible error in the propagated moments. Given a small increment of pseudo-time $\Delta\lambda > 0$, the SDE is approximated within the region $\|\mathrm{x}-\mathrm{x}_l\| < \zeta$, for a sufficiently small $\zeta \in \mathbb{R}_+$, as

$$d\mathrm{x} = \frac{1}{2}D(\lambda)\nabla_{\mathrm{x}}\log\pi\,(\mathrm{x})\,d\lambda + D(\lambda)^{1/2}d\mathrm{w}_\lambda, \quad \lambda \in (\lambda_l,\lambda_l+\Delta\lambda], \mathrm{x}(\lambda_l)=\mathrm{x}_l; \tag{C.1}$$

$$d\mathrm{x} \approx \left[C(\mathrm{x}_l,\lambda)\cdot\mathrm{x} + c(\mathrm{x}_l,\lambda)\right]d\lambda + D(\lambda)^{1/2}d\mathrm{w}_\lambda. \tag{C.2}$$

In this section we provide a nonrigorous argument to explain why this local flow approximation produces admissible errors on the propagated moments without major concern. We will look at the expected error with respect to the intermediate marginal measures that follow from the Langevin dynamics for $\lambda \geq \lambda_l$ as

$$q(\mathrm{x}|\mathrm{y}_k) = \int_{\mathcal{X}} p_t(\mathrm{x}|\mathrm{x}_l)p(\mathrm{x}_l|\mathrm{y}_k)d\mathrm{x}_l = \mathbb{E}_{p(l)}\left[p_t(\mathrm{x}|\mathrm{x}_l)\right]$$

$$= \mathbb{E}_{p(l)}\left[\mathcal{N}\left(\mathrm{x};\mathrm{x}_l+\int_{\lambda_l}^{\lambda_l+\Delta\lambda}\mu(\mathrm{x}_l,\lambda)d\lambda, \left(\int_{\lambda_l}^{\lambda_l+\Delta\lambda}D(\lambda)^{1/2}d\mathrm{w}_\lambda\right)^2\right)\right].$$

where $\mu(\mathrm{x}, \lambda) = \frac{1}{2} D(\lambda) \nabla_{\mathrm{x}} \log \pi(\mathrm{x})$.

In the following analysis we will assume that the state is one-dimensional, i.e., $\mathrm{x} \in \mathbb{R}$, just to present a short argument that can be easily extended to the multidimensional case. If one follows a procedure that (i) applies Îto's lemma to a continuous, measurable and nicely behaved function $\mathcal{M}(\mathrm{x})$, (ii) substitutes in the exact stochastic differential (C.1), (iii) takes expectation of the resulting equation, (iv) derivates it with respect to λ, and then (v) integrates its right-hand side by parts, one obtains the so-called moment equation:

$$\frac{d}{d\lambda} \mathbb{E}_q \left[\mathcal{M}(\mathrm{x}) \right] = \mathbb{E}_q \left[\frac{\partial \mathcal{M}(\mathrm{x})}{\partial \mathrm{x}} \mu(\mathrm{x}, \lambda) + \frac{1}{2} D(\lambda) \frac{\partial^2 \mathcal{M}(\mathrm{x})}{\partial \mathrm{x}^2} \right]. \tag{C.3}$$

Note that when applied to $\mathcal{M}(\mathrm{x}) = \mathrm{x}$ and $\mathcal{M}(\mathrm{x}) = (\mathrm{x} - \mu_m)^2$, for the linear approximation $\tilde{\mu}(\mathrm{x}, \lambda) = C(\lambda) \cdot \mathrm{x} + c(\lambda)$, the referred equation gives the ODEs of the approximated mean and variance respectively

$$\frac{d}{d\lambda} \tilde{\mathbb{E}}_q [\mathrm{x}] = \ C(\lambda) \cdot \tilde{\mathbb{E}}_q [\mathrm{x}] + c(\lambda),$$

$$\frac{d}{d\lambda} \tilde{\mathbb{E}}_q \left[(\mathrm{x} - \mu_m)^2 \right] = 2C(\lambda) \cdot \tilde{\mathbb{E}}_q \left[(\mathrm{x} - \mu_m)^2 \right] + D(\lambda).$$

Denote the deviation $\delta_{\mathbb{E}}[\mathrm{x}] = \mathbb{E}_q[\mathrm{x}] - \tilde{\mathbb{E}}_q[\mathrm{x}]$, where $\mathbb{E}_q[\mathrm{x}] = \mu_m$ is the exact mean propagated by the process (C.1) and $\tilde{\mathbb{E}}_q[\mathrm{x}] = \tilde{\mu}_m$ is the approximated mean propagated by the locally linearized process (C.2). First we note that the reverse triangle inequality allows us to state

$$\lim_{\Delta\lambda \to 0} \left| \frac{\|\delta_{\mathbb{E}}[\mathrm{x}(\lambda + \Delta\lambda)]\| - \|\delta_{\mathbb{E}}[\mathrm{x}(\lambda)]\|}{\Delta\lambda} \right| \leq \lim_{\Delta\lambda \to 0} \frac{\|\delta_{\mathbb{E}}[\mathrm{x}(\lambda + \Delta\lambda)] - \delta_{\mathbb{E}}[\mathrm{x}(\lambda)]\|}{\Delta\lambda},$$

$$\left| \lim_{\Delta\lambda \to 0} \frac{\|\delta_{\mathbb{E}}[\mathrm{x}(\lambda + \Delta\lambda)]\| - \|\delta_{\mathbb{E}}[\mathrm{x}(\lambda)]\|}{\Delta\lambda} \right| \leq \lim_{\Delta\lambda \to 0} \left\| \frac{\delta_{\mathbb{E}}[\mathrm{x}(\lambda + \Delta\lambda)] - \delta_{\mathbb{E}}[\mathrm{x}(\lambda)]}{\Delta\lambda} \right\|,$$

$$\therefore \left| \frac{d}{d\lambda} \|\delta_{\mathbb{E}}[\mathrm{x}(\lambda)]\| \right| \leq \left\| \frac{d}{d\lambda} \delta_{\mathbb{E}}[\mathrm{x}(\lambda)] \right\|, \tag{C.4}$$

where $|\cdot|$ is the absolute value, $\|\cdot\|$ the Euclidean norm, and all limits are assumed to exist. We use inequality (C.4) and the moment equation (C.3) to work out

$$\left| \frac{d}{d\lambda} \left\| \mathbb{E}_q[\mathrm{x}] - \tilde{\mathbb{E}}_q[\mathrm{x}] \right\| \right| \leq \left\| \frac{d}{d\lambda} \mathbb{E}_q[\mathrm{x}] - \frac{d}{d\lambda} \tilde{\mathbb{E}}_q[\mathrm{x}] \right\|,$$

$$\left| \frac{d}{d\lambda} \|\mu_m - \tilde{\mu}_m\| \right| \leq \|\mathbb{E}_q[\mu(\mathrm{x}, \lambda)] - \mathbb{E}_q[\tilde{\mu}(\mathrm{x}, \lambda)]\|, \quad \text{(moment equation for } \mathcal{M}(\mathrm{x}) = \mathrm{x})$$

$$\left| \frac{d}{d\lambda} \|\mu_m - \tilde{\mu}_m\| \right| \leq \|\mathbb{E}_q[\mu(\mathrm{x}, \lambda) - \tilde{\mu}(\mathrm{x}, \lambda)]\| \leq \mathbb{E}_q[\|\mu(\mathrm{x}, \lambda) - \tilde{\mu}(\mathrm{x}, \lambda)\|],$$

$$\left| \frac{d}{d\lambda} \|\mu_m - \tilde{\mu}_m\| \right| \leq \mathbb{E}_q[\|\mu(\mathrm{x}, \lambda) - (C(\mathrm{x}_l, \lambda) \cdot \mathrm{x} + c(\mathrm{x}_l, \lambda))\|] \leq K_0 \mathbb{E}_q\left[\left\|(\mathrm{x} - \mathrm{x}_l)^2\right\|\right],$$

$$\frac{d}{d\lambda} \|\mu_m - \tilde{\mu}_m\| \leq K_1 \zeta^2$$

$$\|\mu_m - \tilde{\mu}_m\| \leq K_1 \zeta^2 \Delta\lambda, \quad \text{for } \lambda \in (\lambda_l, \lambda_l + \Delta\lambda], \ \mathrm{x}(\lambda_l) = \mathrm{x}_l, \ K_0, K_1 \in \mathbb{R}_+, \tag{C.5}$$

where the modulus is dismissed because $\|\mu_m - \tilde{\mu}_m\|$ increases monotonically with $\Delta\lambda$. Similarly for the error

on the second moment

$$\left| \frac{d}{d\lambda} \|\Sigma_m - \tilde{\Sigma}_m\| \right| = \left\| \frac{d}{d\lambda} \mathbb{E}_q \left[(\mathrm{x} - \mu_m)^2 \right] - \frac{d}{d\lambda} \tilde{\mathbb{E}}_q \left[(\mathrm{x} - \mu_m)^2 \right] \right\|,$$

$$\left| \frac{d}{d\lambda} \|\Sigma_m - \tilde{\Sigma}_m\| \right| \leq \left\| \mathbb{E}_q \left[2(\mathrm{x} - \mu_m)\mu(\mathrm{x}, \lambda) + D(\lambda) \right] - \mathbb{E}_q \left[2(\mathrm{x} - \mu_m)\tilde{\mu}(\mathrm{x}, \lambda) + D(\lambda) \right] \right\|,$$

$$\left| \frac{d}{d\lambda} \|\Sigma_m - \tilde{\Sigma}_m\| \right| \leq \left\| 2\mathbb{E}_q \left[(\mathrm{x} - \mu_m)(\mu(\mathrm{x}, \lambda) - \tilde{\mu}(\mathrm{x}, \lambda)) \right] \right\| \leq 2\mathbb{E}_q \left[\left\| (\mathrm{x} - \mu_m)(\mu(\mathrm{x}, \lambda) - \tilde{\mu}(\mathrm{x}, \lambda)) \right\| \right],$$

$$\left| \frac{d}{d\lambda} \|\Sigma_m - \tilde{\Sigma}_m\| \right| \leq 2K_0' \mathbb{E}_q \left[\left\| (\mathrm{x} - \mu_m)(\mathrm{x} - \mathrm{x}_l)^2 \right\| \right] = 2K_0' \mathbb{E}_q \left[\left\| (\mathrm{x}_l - \mu_m)(\mathrm{x} - \mathrm{x}_l)^2 + (\mathrm{x} - \mathrm{x}_l)^3 \right\| \right],$$

$$\left| \frac{d}{d\lambda} \|\Sigma_m - \tilde{\Sigma}_m\| \right| \leq 2K_0' \mathbb{E}_q \left[\left(\|\mathrm{x}_l - \mu_m\| \cdot \left\| (\mathrm{x} - \mathrm{x}_l)^2 \right\| + \left\| (\mathrm{x} - \mathrm{x}_l)^3 \right\| \right) \right], \quad \text{(triangle, Cauchy-Schwarz)}$$

$$\left| \frac{d}{d\lambda} \|\Sigma_m - \tilde{\Sigma}_m\| \right| \leq K_{1a}' \zeta^2 + K_{1b}' \zeta^3,$$

$$\frac{d}{d\lambda} \|\Sigma_m - \tilde{\Sigma}_m\| \leq K_1' \zeta^2,$$

$$\|\Sigma_m - \tilde{\Sigma}_m\| \leq K_1' \zeta^2 \Delta\lambda, \quad \text{for } \lambda \in (\lambda_l, \lambda_l + \Delta\lambda], \, \mathrm{x}(\lambda_l) = \mathrm{x}_l, \, , \, K_0', K_{1a}', K_{1b}', K_1' \in \mathbb{R}^+, \tag{C.6}$$

where the modulus is suppressed because $\|\Sigma_m - \tilde{\Sigma}_m\|$ increases monotonically with $\Delta\lambda$.

It is very important to mention that the collection of factors K_0, K_1, K_0', K_{1a}', K_{1b}', K_1', can be different for each possible interval $(\lambda_l, \lambda_l + \Delta\lambda]$. Rigorously speaking those coefficients can depend on pseudo-time λ because

$$K_0, \, K_0' \propto \frac{1}{2} \left\| \frac{\partial^2}{\partial \mathrm{x}^2} \mu(\mathrm{x}, \lambda) \right\|_{\mathrm{x} = \mathrm{x}_l}.$$

However, in accordance with the methodology of the stochastic particle flow, we select these factors so that inequalities (C.5) and (C.6) hold for a specific interval within which the diffusion coefficient is kept fixed as $D(\lambda_l) \equiv D(\mathrm{x}(\lambda_l)) = D(\mathrm{x}_l)$. Given that the diffusion coefficient is piecewise constant in λ, the referred factors are also piecewise constant in λ and directly dependent on the diffusion coefficient. If the target density is Gaussian, then $K_0, K_0' = 0$ and the error commited due to the local linearization is null. For a fixed diffusion coefficient, $K_0, K_0' > 0$ if and only if the target log-density has third or higher-order non-zero derivatives at x_l.

Notice that integrating (C.1) by the Euler-Maruyama scheme would produce $\mathbb{E}_q \left[\|\tilde{\mathrm{x}}_{l+1} - \mathrm{x}_{l+1}\| \right] \leq K \Delta\lambda^{1/2}$ and, therefore, it would be reasonable to expect $\|\mu_m - \tilde{\mu}_m\| \leq K_2 \Delta\lambda^2$ and $\|\Sigma_m - \tilde{\Sigma}_m\| \leq K_2' \Delta\lambda^2$ for some $K_2, K_2' \in \mathbb{R}_+$. In the multivariate case, by applying the method used in this section, the curious reader should learn that the bounds are multiplied by the dimension n_x, to give $\|\mu_m - \tilde{\mu}_m\|_2 \leq K_1 \zeta^2 \Delta\lambda \cdot n_x$ and $\|\Sigma_m - \tilde{\Sigma}_m\|_2 \leq K_1' \zeta^2 \Delta\lambda \cdot n_x$.

D Discrete-time stochastic IDM

This section of the appendix presents the resulting discrete-time approximation of the stochastic Intelligent Driver model. Define the state equation for the discrete-time IDM to be

$$\mathrm{x}_k = A \cdot \mathrm{x}_{k-1} + B + \mathrm{w}_k, \tag{D.1}$$

where the state vector is represented for α vehicles by

$$\mathrm{x}_k = \begin{bmatrix} p_1 \\ p_2 \\ \vdots \\ p_{\alpha-1} \\ p_\alpha \\ v_1 \\ v_2 \\ \vdots \\ v_{\alpha-1} \\ v_\alpha \end{bmatrix}_k . \tag{D.2}$$

The variables p_i and v_i are the position and velocity of the ith vehicle respectively. The state-transition matrix can be written as

$$A = \begin{bmatrix} A_{1,1} & A_{1,2} & \ldots & A_{1,2\alpha} \\ A_{2,1} & A_{2,2} & & \vdots \\ \vdots & & \ddots & \\ A_{2\alpha,1} & \ldots & & A_{2\alpha,2\alpha} \end{bmatrix}_{k-1}. \tag{D.3}$$

For $i, j \in \mathbb{N}$, the diagonal elements of the state-transition matrix are given by

$$A_{i,i} = \begin{cases} 1, & i \in [1, \alpha]; \\ 1 + \frac{\partial \dot{v}_{i-\alpha}}{\partial v_{i-\alpha}} \cdot \Delta t, & i \in (\alpha + 1, 2\alpha]; \end{cases} \tag{D.4}$$

and the off-diagonal elements given by

$$A_{i,j} = \begin{cases} \Delta t, & i \in [1, \alpha], \ j = i + \alpha; \\ \frac{\partial \dot{v}_{i-\alpha}}{\partial p_\alpha} \cdot \Delta t, & i = \alpha + 1, \ j = \alpha; \\ \frac{\partial \dot{v}_{i-\alpha}}{\partial p_{i-\alpha-1}} \cdot \Delta t, & i \in (\alpha + 1, 2\alpha], \ j = i - 1 - \alpha; \\ \frac{\partial \dot{v}_{i-\alpha}}{\partial v_\alpha} \cdot \Delta t, & i = \alpha + 1, \ j = 2\alpha; \\ \frac{\partial \dot{v}_{i-\alpha}}{\partial v_{i-\alpha-1}} \cdot \Delta t, & i \in (\alpha + 1, 2\alpha], \ j = i - 1; \\ 0, & \text{otherwise}; \end{cases} \tag{D.5}$$

where

$$\frac{\partial \dot{v}_n}{\partial p_{n-1}} = +2a \left(\frac{s(v_n, \Delta v_n)^2}{s_n^3} \right), \tag{D.6}$$

$$\frac{\partial \dot{v}_n}{\partial p_n} = -2a \left(\frac{s(v_n, \Delta v_n)^2}{s_n^3} \right), \tag{D.7}$$

$$\frac{\partial \dot{v}_n}{\partial v_{n-1}} = +2a \left(\frac{s(v_n, \Delta v_n)}{s_n^2} \right) \left(\frac{v_n}{2\sqrt{a \cdot b}} \right), \tag{D.8}$$

$$\frac{\partial \dot{v}_n}{\partial v_n} = -a \left(\frac{\delta}{v_0} \right) \left(\frac{v_n}{v_0} \right)^{\delta-1} \\ - 2a \left(\frac{s(v_n, \Delta v_n)}{s_n^2} \right) \left(T_h + \frac{2v_n - v_{n-1}}{2\sqrt{a \cdot b}} \right). \tag{D.9}$$

The model takes into account the fact that, on a ring road, the last vehicle in the convoy can be regarded the one potentially in front of the vehicle leading the queue, assuming that the first vehicle can complete the circuit faster and approach the last one from behind. This is represented by the terms $\partial_{p_\alpha} \dot{v}_1 \cdot \Delta t$ and $\partial_{v_\alpha} \dot{v}_1 \cdot \Delta t$ that appear in (D.5) when $i = \alpha + 1$, which shall be calculated respectively according to expressions analogous to (D.6) and (D.8). The constant term is defined as

$$B = \begin{bmatrix} B_1 \\ B_2 \\ \vdots \\ B_{2\alpha} \end{bmatrix}_k, \tag{D.10}$$

where

$$B_i = \begin{cases} 0, & i \in [1, \alpha]; \\ \langle \dot{v}_i \rangle - \frac{\partial \dot{v}_i}{\partial p_{i-1}} \cdot p_{i-1} - \frac{\partial \dot{v}_i}{\partial p_i} \cdot p_i \\ \quad - \frac{\partial \dot{v}_i}{\partial v_{i-1}} \cdot v_{i-1} - \frac{\partial \dot{v}_i}{\partial v_i} \cdot v_i, & i \in [\alpha + 1, 2\alpha]; \end{cases} \tag{D.11}$$

and

$$\langle \dot{v}_i \rangle = a \left[1 - \left(\frac{v_i}{v_0} \right)^\delta - \left(\frac{s(v_i, \Delta v_i)}{s_i} \right)^2 \right]. \tag{D.12}$$

The covariance matrix $Q_k = \mathbb{E}\left[\mathrm{w}_k \mathrm{w}_k^T\right]$ is defined as

$$Q_k = \begin{bmatrix} Q_{1,1} & Q_{1,2} & \dots & Q_{1,2\alpha} \\ Q_{2,1} & Q_{2,2} & & \vdots \\ \vdots & & \ddots & \\ Q_{2\alpha,1} & \dots & & Q_{2\alpha,2\alpha} \end{bmatrix}_k, \tag{D.13}$$

with diagonal elements

$$Q_{i,i} = \sigma_q^2 \times \begin{cases} Q_{i,i}^{(1)}, & i \in [1,\alpha]; \\ Q_{i,i}^{(2)}, & i = \alpha+1; \\ Q_{i,i}^{(3)}, & i \in (\alpha+1, 2\alpha]; \end{cases} \tag{D.14}$$

where

$$Q_{i,i}^{(1)} = \frac{\Delta t^3}{3} + \Delta t,$$

$$Q_{i,i}^{(2)} = \left(\frac{\partial \dot{v}_{i-\alpha}}{\partial p_\alpha} + \frac{\partial \dot{v}_{i-\alpha}}{\partial p_{i-\alpha}} + \frac{\partial \dot{v}_{i-\alpha}}{\partial v_\alpha} + \frac{\partial \dot{v}_{i-\alpha}}{\partial v_{i-\alpha}} \right) \cdot \frac{\Delta t^3}{3}$$

$$+ \frac{\partial \dot{v}_{i-\alpha}}{\partial v_{i-\alpha}} \cdot \Delta t^2 + \Delta t,$$

$$Q_{i,i}^{(3)} = \left(\frac{\partial \dot{v}_{i-\alpha}}{\partial p_{i-\alpha-1}} + \frac{\partial \dot{v}_{i-\alpha}}{\partial p_{i-\alpha}} + \frac{\partial \dot{v}_{i-\alpha}}{\partial v_{i-\alpha-1}} + \frac{\partial \dot{v}_{i-\alpha}}{\partial v_{i-\alpha}} \right) \cdot \frac{\Delta t^3}{3}$$

$$+ \frac{\partial \dot{v}_{i-\alpha}}{\partial v_{i-\alpha}} \cdot \Delta t^2 + \Delta t; \tag{D.15}$$

and off-diagonal elements

$$Q_{j,i} = Q_{i,j} = \sigma_q^2 \times \begin{cases} Q_{i,j}^{(4)}, & i \in [1,\alpha], \; j = i+\alpha; \\ Q_{i,j}^{(5)}, & i = \alpha+1, \; j = \alpha; \\ Q_{i,j}^{(6)}, & i \in (\alpha+1, 2\alpha], \; j = i-1-\alpha; \\ Q_{i,j}^{(7)}, & i = \alpha+1, \; j = 2\alpha; \\ Q_{i,j}^{(8)}, & i \in (\alpha+1, 2\alpha], \; j = i-1; \\ 0, & \text{otherwise}; \end{cases} \tag{D.16}$$

where

$$Q_{i,j}^{(4)} = \frac{\partial \dot{v}_i}{\partial v_i} \cdot \frac{\Delta t^3}{3} + \left(\frac{\partial \dot{v}_i}{\partial p_i} + 1 \right) \cdot \frac{\Delta t^2}{2},$$

$$Q_{i,j}^{(5)} = \frac{\partial \dot{v}_{i-\alpha}}{\partial v_\alpha} \cdot \frac{\Delta t^3}{3} + \frac{\partial \dot{v}_{i-\alpha}}{\partial p_\alpha} \cdot \frac{\Delta t^2}{2},$$

$$Q_{i,j}^{(6)} = \frac{\partial \dot{v}_{i-\alpha}}{\partial v_{i-\alpha-1}} \cdot \frac{\Delta t^3}{3} + \frac{\partial \dot{v}_{i-\alpha}}{\partial p_{i-\alpha-1}} \cdot \frac{\Delta t^2}{2},$$

$$Q_{i,j}^{(7)} = \frac{\partial \dot{v}_{i-\alpha}}{\partial v_\alpha} \cdot \frac{\Delta t^2}{3},$$

$$Q_{i,j}^{(8)} = \frac{\partial \dot{v}_{i-\alpha}}{\partial v_{i-\alpha-1}} \cdot \frac{\Delta t^2}{3}. \tag{D.17}$$

References

[1] E. Novak, "Some Results on the Complexity of Numerical Integration," in *Monte Carlo and Quasi-Monte Carlo Methods: MCQMC, Leuven, Belgium, April 2014* (R. Cools and D. Nuyens, eds.), pp. 161–183, Springer International Publishing, 2016.

[2] A. Gelman and X.-L. Meng, "Simulating normalizing constants: from importance sampling to bridge sampling to path sampling," *Statistical Science*, vol. 13, pp. 163–185, 05 1998.

[3] J. Deutscher, A. Blake, and I. Reid, "Articulated body motion capture by annealed particle filtering," in *Computer Vision and Pattern Recognition, 2000, IEEE Conference on*, vol. 2, pp. 126–133 vol.2, 2000.

[4] R. M. Neal, "Annealed importance sampling," *Statistics and Computing*, vol. 11, no. 2, pp. 125–139, 2001.

[5] P. Del Moral, A. Doucet, and A. Jasra, "Sequential Monte Carlo samplers," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 68, no. 3, pp. 411–436, 2006.

[6] P. Bickel, B. Li, and T. Bengtsson, *Sharp failure rates for the bootstrap particle filter in high dimensions*, vol. 3 of *Collections*, pp. 318–329. Beachwood, Ohio, USA: Institute of Mathematical Statistics, 2008.

[7] S. Duane, A. D. Kennedy, B. J. Pendleton, and D. Roweth, "Hybrid Monte Carlo," *Physics Letters B*, vol. 195, no. 2, pp. 216 – 222, 1987.

[8] R. M. Neal, *Bayesian Learning for Neural Networks*. Secaucus, NJ, USA: Springer-Verlag New York, Inc., 1996.

[9] G. O. Roberts and O. Stramer, "Langevin Diffusions and Metropolis-Hastings Algorithms," *Methodology And Computing In Applied Probability*, vol. 4, no. 4, pp. 337–357, 2002.

[10] M. Girolami and B. Calderhead, "Riemann manifold Langevin and Hamiltonian Monte Carlo methods," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 73, no. 2, pp. 123–214, 2011.

[11] R. L. Stratonovich, "Optimum nonlinear systems which bring about a separation of a signal with constant parameters from noise," *Radiofizika*, vol. 2, no. 6, pp. 892–901, 1959.

[12] H. J. Kushner, "On the differential equations satisfied by conditional probability densities of Markov processes, with applications," *Journal of the SIAM, Series A (Control)*, vol. 2, no. 1, pp. 106–119, 1964.

[13] M. Zakai, "On the optimal filtering of diffusion processes," *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, vol. 11, no. 3, pp. 230–243, 1969.

[14] R. S. Bucy and P. D. Joseph, *Filtering for Stochastic Processes with Applications to Guidance*. John Wiley & Sons, 2nd ed., 1968.

[15] V. E. Beneš, "Exact finite-dimensional filters for certain diffusions with nonlinear drift," *Stochastics*, vol. 5, no. 1-2, pp. 65–92, 1981.

[16] F. E. Daum, "Exact finite-dimensional nonlinear filters," *Automatic Control, IEEE Transactions on*, vol. 31, pp. 616–622, Jul 1986.

[17] D. Crisan and T. Lyons, "A particle approximation of the solution of the Kushner-Stratonovitch equation," *Probability Theory and Related Fields*, vol. 115, no. 4, pp. 549–578, 1999.

[18] D. Crisan, P. D. Moral, and T. J. Lyons, "Interacting Particle Systems Approximations of the Kushner-Stratonovitch Equation," *Advances in Applied Probability*, vol. 31, no. 3, pp. pp. 819–838, 1999.

[19] D. Crisan and J. Xiong, "Approximate McKean-Vlasov representations for a class of SPDEs," *Stochastics*, vol. 82, pp. 53–68, 2010.

[20] A. Budhiraja and G. Kallianpur, "Approximations to the solution of the Zakai equation using multiple Wiener and Stratonovich integral expansions," *Stochastics and Stochastic Reports*, vol. 56, no. 3-4, pp. 271–315, 1996.

[21] A. Budhiraja and G. Kallianpur, "The Feynman-Stratonovich semigroup and Stratonovich integral expansions in nonlinear filtering," *Applied Mathematics and Optimization*, vol. 35, no. 1, pp. 91–116, 1997.

[22] C. Vergé, C. Dubarry, P. Del Moral, and E. Moulines, "On parallel implementation of sequential Monte Carlo methods: the island particle model," *Statistics and Computing*, vol. 25, no. 2, pp. 243–260, 2015.

[23] P. Rebeschini and R. van Handel, "Can local particle filters beat the curse of dimensionality?," *Annals of Applied Probability*, vol. 25, pp. 2809–2866, 10 2015.

[24] U. D. Hanebeck, K. Briechle, and A. Rauh, "Progressive Bayes: a new framework for nonlinear state estimation," in *Proc. SPIE. 5099, Multisensor, Multisource Information Fusion: Architectures, Algorithms, and Applications 2003*, vol. 5099, pp. 256–267, 2003.

[25] J. Hagmar, M. Jirstrand, L. Svensson, and M. Morelande, "Optimal parameterization of posterior densities using homotopy," in *Information Fusion (FUSION), 2011, Proceedings of the 14th International Conference on*, pp. 1–8, July 2011.

[26] U. D. Hanebeck and J. Steinbring, "Progressive Gaussian filtering based on Dirac Mixture approximations," in *Information Fusion (FUSION), 2012, Proceedings of the 15th International Conference on*, pp. 1697–1704, July 2012.

[27] S. Reich, "A dynamical systems framework for intermittent data assimilation," *BIT Numerical Mathematics*, vol. 51, no. 1, pp. 235–249, 2011.

[28] S. Reich, "A Guided Sequential Monte Carlo Method for the Assimilation of Data into Stochastic Dynamical Systems," in *Recent Trends in Dynamical Systems* (A. Johann, H.-P. Kruse, F. Rupp, and S. Schmitz, eds.), vol. 35 of *Springer Proceedings in Mathematics & Statistics*, pp. 205–220, Springer Basel, 2013.

[29] M. Venugopal, R. M. Vasu, and D. Roy, "An Ensemble Kushner-Stratonovich-Poisson Filter for Recursive Estimation in Nonlinear Dynamical Systems," *Automatic Control, IEEE Transactions on*, vol. PP, no. 99, pp. 1–1, 2015.

[30] N. Oudjane and C. Musso, "Progressive correction for regularized particle filters," in *Information Fusion (FUSION), 2000, Proceedings of the Third International Conference on*, vol. 2, pp. THB2/10–THB2/17 vol.2, July 2000.

[31] D. Crisan, "Particle Filters in a Continuous Time Framework," in *Nonlinear Statistical Signal Processing Workshop, 2006, IEEE*, pp. 73–78, Sept 2006.

[32] F. Daum and J. Huang, "Nonlinear filters with log-homotopy," in *Proc. SPIE. 6699, Signal and Data Processing of Small Targets 2007*, vol. 6699 of *Signal and Data Processing of Small Targets, 2007*, (USA), p. 18, SPIE, August 26, 2007 2007. id: 1; issn: print.

[33] S. Sarkar, S. R. Chowdhury, M. Venugopal, R. M. Vasu, and D. Roy, "A Kushner-Stratonovich Monte Carlo filter applied to nonlinear dynamical system identification," *Physica D: Nonlinear Phenomena*, vol. 270, pp. 46 – 59, 2014.

[34] D. Crisan and K. Li, "Generalised particle filters with Gaussian mixtures," *Stochastic Processes and their Applications*, vol. 125, no. 7, pp. 2643 – 2673, 2015.

[35] P. Bunch and S. Godsill, "Particle filtering with progressive Gaussian approximations to the optimal importance density," in *Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP), 2013, IEEE 5th International Workshop on*, pp. 360–363, 2013. ID: 1.

[36] M. S. Arulampalam, S. Maskell, N. Gordon, and T. Clapp, "A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking," *Signal Processing, IEEE Transactions on*, vol. 50, pp. 174–188, Feb 2002.

[37] F. Daum and J. Huang, "Curse of dimensionality and particle filters," in *Aerospace Conference, 2003, IEEE*, vol. 4, pp. 1979–1993, March 2003.

[38] C. Snyder, T. Bengtsson, P. Bickel, and J. Anderson, "Obstacles to High-Dimensional Particle Filtering," *Monthly Weather Review*, vol. 136, pp. 4629–4640, dec 2008.

[39] A. Doucet and A. M. Johansen, *A Tutorial on Particle Filtering and Smoothing: Fifteen Years Later*. Oxford University Press, 2009.

[40] F. Daum and J. Huang, "Particle flow for nonlinear filters with log-homotopy," in *Proc. SPIE. 6969, Signal and Data Processing of Small Targets 2008*, vol. 6969, pp. 696918–696918–12, 2008.

[41] F. Daum and J. Huang, "Nonlinear filters with particle flow," in *Proc. SPIE. 7445, Signal and Data Processing of Small Targets 2009*, vol. 7445, pp. 74450R–74450R–9, 2009.

[42] F. Daum and J. Huang, "Generalized particle flow for nonlinear filters," in *Proc. SPIE. 7698, Signal and Data Processing of Small Targets 2010*, vol. 7698, pp. 76980I–76980I–12, 2010.

[43] F. Daum and J. Huang, "Particle flow for nonlinear filters, Bayesian decisions and transport," in *Information Fusion (FUSION), 2013, 16th International Conference on*, pp. 1072–1079, 2013.

[44] C. Villani, *Topics in Optimal Transportation.* Graduate studies in mathematics, American Mathematical Society, 2003.

[45] C. E. Gutiérrez, *The Monge-Ampère Equation.* Progress in Nonlinear Differential Equations and Their Applications, Birkhäuser Boston, 2001.

[46] J. Moser, "On the Volume Elements on a Manifold," *Transactions of the American Mathematical Society*, vol. 120, no. 2, pp. 286–294, 1965.

[47] B. Dacorogna and J. Moser, "On a partial differential equation involving the jacobian determinant," *Annales de l'I.H.P. Analyse non linéaire*, vol. 7, no. 1, pp. 1–26, 1990.

[48] J.-D. Benamou and Y. Brenier, "A computational fluid mechanics solution to the Monge-Kantorovich mass transfer problem," *Numerische Mathematik*, vol. 84, no. 3, pp. 375–393, 2000.

[49] J. Heng, A. Doucet, and Y. Pokern, "Gibbs Flow for Approximate Transport with Applications to Bayesian Computation," *ArXiv e-prints*, Sept. 2015.

[50] S. Choi, P. Willett, F. Daum, and J. Huang, "Discussion and application of the homotopy filter," in *Proc. SPIE. 8050, Signal Processing, Sensor Fusion, and Target Recognition XX*, vol. 8050, pp. 805021–805021–12, 2011. journal: Proc. SPIE.

[51] T. Ding and M. J. Coates, "Implementation of the Daum-Huang exact-flow particle filter," in *Statistical Signal Processing Workshop (SSP), 2012, IEEE*, pp. 257–260, 2012. ID: 1.

[52] M. Khan and M. Ulmke, "Improvements in the implementation of log-homotopy based particle flow filters," in *Information Fusion (FUSION), 2015, Proceedings of the 18th International Conference on*, pp. 74–81, July 2015.

[53] R. Jordan, D. Kinderlehrer, and F. Otto, "The Variational Formulation of the Fokker–Planck Equation," *SIAM Journal on Mathematical Analysis*, vol. 29, no. 1, pp. 1–17, 1998.

[54] F. Daum and J. Huang, "Particle flow with non-zero diffusion for nonlinear filters, Bayesian decisions and transport," in *Proc. SPIE. 8857, Signal and Data Processing of Small Targets 2013*, vol. 8745, pp. 87450P–87450P–13, 2013.

[55] P. Bunch and S. Godsill, "Approximations of the Optimal Importance Density using Gaussian Particle Flow Importance Sampling," *Journal of the American Statistical Association*, vol. 0, no. ja, pp. 0–0, 2015.

[56] F. Septier and G. W. Peters, "Langevin and Hamiltonian Based Sequential MCMC for Efficient Bayesian Filtering in High-Dimensional Spaces," *Signal Processing, IEEE Journal of Selected Topics in*, vol. 10, pp. 312–327, March 2016.

[57] F. Daum, J. Huang, and A. Noushin, "Exact particle flow for nonlinear filters," in *Proc. SPIE 7697, Signal Processing, Sensor Fusion, and Target Recognition XIX*, vol. 7697, pp. 769704–769704–19, 2010.

[58] L. C. Evans and W. Gangbo, "Differential equations methods for the Monge-Kantorovich mass transfer problem," *Memoirs of the American Mathematical Society*, vol. 137, p. 653, 1999.

[59] A. H. Jazwinski, *Stochastic Processes and Filtering Theory.* Academic Press, 1970.

[60] C. W. Gardiner, *Handbook of Stochastic Methods.* Berlin Heidelberg New York: Springer, 2nd enlarged ed., 1985.

[61] B. W. Silverman, *Density Estimation for Statistics and Data Analysis.* London: Chapman and Hall, April 1986.

[62] H. Risken, *The Fokker-Planck equation: methods of solution and applications.* 2nd, Springer-Verlag, 1989.

[63] R. Jordan and D. Kinderlehrer, "An Extended Variational Principle," in *Partial Differential Equations and Applications: Collected Papers in Honor of Carlo Pucci* (G. Talenti, E. Vesentini, and P. Marcellini, eds.), Lecture Notes in Pure and Applied Mathematics, ch. 18, Taylor & Francis, 1996.

[64] A. S. Dalalyan, "Theoretical guarantees for approximate sampling from smooth and log-concave densities," *ArXiv e-prints*, Dec 2014.

[65] D. A. Levin, Y. Peres, and E. L. Wilmer, *Markov chains and mixing times*, ch. 4. Providence, R.I. American Mathematical Society, 2009. With a chapter on coupling from the past by James G. Propp and David B. Wilson.

[66] R. C. Rao, "Information and the accuracy attainable in the estimation of statistical parameters," *Bulletin of Calcutta Mathematical Society*, vol. 37, pp. 81–91, 1945.

[67] N. Higham, "Computing the nearest correlation matrix—a problem from finance," *IMA Journal of Numerical Analysis*, vol. 22, pp. 329–343, 2002.

[68] M. Betancourt, *A General Metric for Riemannian Manifold Hamiltonian Monte Carlo*, pp. 327–334. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013.

[69] T. Ozaki, "A bridge between nonlinear time series models and nonlinear stochastic dynamical systems: a local linearization approach," *Statistica Sinica*, vol. 2, no. 1, pp. 113–135, 1992.

[70] G. O. Roberts and J. S. Rosenthal, "Optimal Scaling of Discrete Approximations to Langevin Diffusions," *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, vol. 60, no. 1, pp. 255–268, 1998.

[71] N. S. Pillai, A. M. Stuart, and A. H. Thiéry, "Optimal Scaling and Diffusion Limits for the Langevin Algorithm in High Dimensions," *Annals of Applied Probability*, vol. 22, pp. 2320–2356, 12 2012.

[72] G. O. Roberts, A. Gelman, and W. R. Gilks, "Weak convergence and optimal scaling of random walk Metropolis algorithms," *Annals of Applied Probability*, vol. 7, pp. 110–120, 02 1997.

[73] A. S. Dalalyan, "Theoretical guarantees for approximate sampling from smooth and log-concave densities," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 2016.

[74] B. D. O. Anderson and J. B. Moore, *Optimal Filtering*. Englewood Cliffs, NJ: Prentice-Hall, 1979.

[75] S. Meyn and R. L. Tweedie, *Markov Chains and Stochastic Stability*. New York, NY, USA: Cambridge University Press, 2nd ed., 2009.

[76] H. W. Sorenson and D. L. Alspach, "Recursive Bayesian Estimation Using Gaussian Sums," *Automatica*, vol. 7, pp. 465–479, July 1971.

[77] D. L. Alspach and H. W. Sorenson, "Nonlinear Bayesian estimation using Gaussian sum approximations," *Automatic Control, IEEE Transactions on*, vol. 17, pp. 439–448, Aug 1972.

[78] J. H. Kotecha and P. M. Djurić, "Gaussian sum particle filtering," *Signal Processing, IEEE Transactions on*, vol. 51, pp. 2602–2612, Oct 2003.

[79] G. Terejanu, P. Singla, T. Singh, and P. D. Scott, "A novel Gaussian Sum Filter Method for accurate solution to the nonlinear filtering problem," in *Information Fusion (FUSION), 2008, Proceedings of the 11th International Conference on*, pp. 1–8, June 2008.

[80] M. Klaas, N. de Freitas, and A. Doucet, "Toward Practical N^2 Monte Carlo: the Marginal Particle Filter," *ArXiv e-prints*, July 2012.

[81] R. Graham, "Statistical Theory of Instabilities in Stationary Nonequilibrium Systems with Applications to Lasers and Nonlinear Optics," in *Springer Tracts in Modern Physics* (G. Höhler, ed.), pp. 1–97, Springer Berlin Heidelberg, 1973.

[82] H.-K. Janssen, "On a Lagrangean for classical field dynamics and renormalization group calculations of dynamical critical properties," *Zeitschrift für Physik B Condensed Matter*, vol. 23, no. 4, pp. 377–380, 1976.

[83] Y. Bar-Shalom, P. K. Willett, and X. Tian, *Tracking and Data Fusion: A Handbook of Algorithms*. YBS Publishing, April 2011.

[84] Y. Bar-Shalom, F. Daum, and J. Huang, "The probabilistic data association filter," *Control Systems, IEEE*, vol. 29, pp. 82–100, Dec 2009.

[85] A. Marrs, S. Maskell, and Y. Bar-Shalom, "Expected likelihood for tracking in clutter with particle filters," in *Proc. SPIE. 4728, Signal and Data Processing of Small Targets 2002*, vol. 4728, pp. 230–239, 2002.

[86] X. R. Li and Z. Zhao, "Measuring Estimator's Credibility: Noncredibility Index," in *Information Fusion (FUSION), 2006, Proceedings of the 9th International Conference on*, pp. 1–8, July 2006.

[87] L. Y. Pao, "Multisensor multitarget mixture reduction algorithms for tracking," *Journal of Guidance, Control, and Dynamics*, vol. 17, pp. 1205–1211, 2014/04/27 1994.

[88] H. A. P. Blom and E. A. Bloem, "Probabilistic data association avoiding track coalescence," *Automatic Control, IEEE Transactions on*, vol. 45, pp. 247–259, Feb 2000.

[89] H. A. P. Blom and E. A. Bloem, "Joint Particle Filtering of Multiple Maneuvering Targets From Unassociated Measurements," *Journal of Advances in Information Fusion*, vol. 1, no. 1, pp. 15–34, 2006.

[90] M. Treiber and A. Kesting, *Traffic flow dynamics: Data, Models and Simulation.* Springer-Verlag Berlin Heidelberg, 2013.

[91] A. S. Dalalyan and A. B. Tsybakov, "Sparse Regression Learning by Aggregation and Langevin Monte-Carlo," *Journal of Computing and System Sciences*, vol. 78, pp. 1423–1443, Sep 2012.

[92] H. Lamb, *Hydrodynamics.* University Press, 1895.

[93] G. N. Milstein, J. G. M. Schoenmakers, and V. Spokoiny, "Transition Density Estimation for Stochastic Differential Equations via Forward-Reverse Representations," *Bernoulli*, vol. 10, no. 2, pp. pp. 281–312, 2004.

[94] R. Bucy, "Nonlinear filtering theory," *Automatic Control, IEEE Transactions on*, vol. 10, pp. 198–198, Apr 1965.