

Particle Filtering for Stochastic Navier-Stokes Signal Observed with Linear Additive Noise

Francesc Pons Llopis*, Nikolas Kantas*, Alexandros Beskos†, Ajay Jasra‡

Abstract

We consider a non-linear filtering problem, whereby the signal obeys the stochastic Navier-Stokes equations and is observed through a linear mapping with additive noise. The setup is relevant to data assimilation for numerical weather prediction and climate modelling, where similar models are used for unknown ocean or wind velocities. We present a particle filtering methodology that uses likelihood informed importance proposals, adaptive tempering, and a small number of appropriate Markov Chain Monte Carlo steps. We provide a detailed design for each of these steps and show in our numerical examples that they are all crucial in terms of achieving good performance and efficiency.

1 Introduction

We focus on a stochastic filtering problem where a space and time varying hidden signal is observed at discrete times with noise. The non-linear filtering problem consists of computing the conditional probability law of the hidden stochastic process (the so-called signal) given observations of it collected in a sequential manner. In particular, we model the signal with a particular dissipative stochastic partial differential equation (SPDE), which is the stochastic Navier-Stokes Equation (NSE). This model, or a variant thereof, is often used in applications to model unknown quantities such as atmosphere or ocean velocity. In the spirit of data assimilation and uncertainty quantification, we wish to extract information for the trajectory of the hidden signal from noisy observations using a Bayesian approach. Typical applications include numerical weather forecasting in meteorology, oceanography and atmospheric sciences, geophysics, hydrology and petroleum engineering; see [2, 29, 5] for an overview.

We restrict to the setting where the state of interest is the time varying velocity field, $V(x, t)$, in some 2D bounded set Ω . The unknown state is modelled using the stochastic NSE

$$dV(x, t) - \nu \Delta V(x, t) dt + B(V, V)(x, t) dt = f(x, t) dt + Q^{\frac{1}{2}} dW(x, t), \quad (1)$$

where Δ is the Laplacian, ν a viscosity constant, B a non-linear operator due to convection, Q a positive, self adjoint, trace class operator, f a deterministic forcing and $W(x, t)$ a space-time white noise as in [9]. This might appear as a restrictive choice for the dynamics, but the subsequent methodology is generic and could be potentially applied to other similar dissipative SPDEs, such as the stochastic Burger's or Kuramoto–Sivashinski equations [20, 6].

The evolution of the unknown state of the SPDE is observed at discrete times and generates a sequence of noisy observations $\mathcal{Y}_n = (Y_{t_1}, \dots, Y_{t_n})$. In order to perform accurate estimation and uncertainty quantification, we are interested not just in approximating a single trajectory estimate of the hidden state, but in the complete filtering distribution,

$$\pi_n(\cdot) = \mathbb{P}[V(x, t_n) \in \cdot | \mathcal{Y}_n], \quad (2)$$

*Department of Mathematics, Imperial College London, UK.

†Department of Statistical Science, University College London, UK.

‡Department of Statistics and Applied Probability, National University of Singapore

that is, the conditional distribution of the state given all the observations obtained up to current time t_n . The main objective is to compute the filtering distribution as it evolves with time, which is an instance of the stochastic filtering problem [1]. The solution of the problem can be formulated rigorously as a recursive Bayesian inference problem posed on an appropriate function space [27]. In contrast to standard filtering problems, the problem setup here is particularly challenging: the prior consists of a complicated probability law generated by the SPDE [9] and the observation likelihoods tend to be very informative about the signal due to the high dimension of the latter.

The aim of this paper is to propose Sequential Monte Carlo (SMC) methods (also known as Particle Filters (PF)) that can approximate effectively these conditional distributions. Computing the evolution of the filtering distribution π_n is not analytically tractable, except in linear Gaussian settings. SMC is a generic Monte Carlo method that approximates the sequence of π_n -s and their normalising constant $\mathbb{P}[\mathcal{Y}_n]$ (known in Statistics as marginal likelihood or evidence). This is achieved by obtaining samples known as particles and combining Importance Sampling (IS), resampling and parallel Markov Chain Monte Carlo (MCMC) steps. The main advantages of the methodology are: i) it is sequential and on-line in nature; ii) it does not require restrictive model assumptions such as Gaussian noise or linear dynamics and observations; iii) it is parallelisable, so one could gain significant speed-up using appropriate hardware (e.g. GPUs, computing clusters) [28]; iv) it is a well-studied principled method with an extensive literature justifying its validity and theoretical properties, see e.g. [11, 10]. So far SMC has been extremely successful in typically low to moderate dimensions [13], but its application in high dimensional settings has been very challenging mainly due to the difficulty to perform IS efficiently in high dimensions [35]. Despite this challenge a few successful high dimensional SMC implementations have appeared recently for applications with discrete time signal dynamics [30, 40, 39, 41, 5, 7, 3].

We will formulate the filtering problem with discrete time observations and continuous time dynamics. This setup has appeared previously in [34, 33] for signals corresponding to low dimensional stochastic differential equations (SDEs). The aim of this paper is to provide a novel, accurate and more efficient SMC design when the hidden signal is modelled by a SPDE with linear Gaussian observation. To achieve this challenging task, the particle filter will use computational tools that have been previously successful in similar high dimensional problems, such as tempering [25] and pre-conditioned Crank Nicholson MCMC steps [19, 8]. Using such tools, we propose a particle algorithm that can be used to approximate π_n when the signal obeys the stochastic NSE and the observations are linear with additive noise. On a general level, the proposed algorithm has a similar structure to [23], but here we additionally adopt the use of IS. We will provide a detailed design of the necessary likelihood informed importance proposals and the MCMC moves used. We extend known IS techniques for SDEs ([18, 42]) and MCMC moves for high dimensional problems ([37, 19, 8]) to make them applicable for filtering problems involving the stochastic NSE or other dissipative SPDEs. In the context of particle filtering, our developments leads to an SMC algorithm that performs effectively for the high dimensional problem at hand using a moderate amount of particles.

The organisation of this paper is as follows: in Section 2 we present some background on the stochastic NSE and in Section 3 we formulate the filtering problem of interest. In Section 4 we present the SMC algorithm and in Section 5 we present a numerical case study that illustrates the performance and efficiency of our method. Finally, in Section 6 we provide some concluding remarks.

2 Background on the Stochastic Navier-Stokes Equations

We present some background on the 2D stochastic NSE defined on an appropriate separable Hilbert space. We restrict the presentation to the case of periodic boundary conditions following the treatment in [14]. This choice is motivated mainly for convenience in exposition and for performing numerical approximations using Fast Fourier transforms. The formulation and properties of the stochastic NSE can allow for the more technically demanding Dirichlet conditions on a smooth boundary [15, 19]. We stress that the subsequent particle filtering methodology is generic and does not rely on the choice of boundary conditions.

2.1 Preliminaries

Let the region of interest be the torus $\Omega := [0, 2\pi]^2$ with $x = (x_1, x_2) \in \Omega$ being a point on the space. The quantity of interest is a time-space varying velocity field $v : \Omega \times [0, \infty) \rightarrow \mathbb{R}^2$, $v(x, t) = (v_1(x, t), v_2(x, t))'$, where \cdot' denotes vector/matrix transpose. Conservation of mass and momentum for Newtonian, incompressible, homogeneous flow gives rise to the traditional NSE expressed as follows:

$$\begin{aligned} \partial_t v - \nu \Delta v + (v \cdot \nabla) v &= f_0 - \nabla p, \quad v(x, 0) = v_0(x), \\ \nabla \cdot v &= 0, \quad \int_{\Omega} v(x, \cdot) dx = 0, \end{aligned} \quad (3)$$

where $\nu > 0$ is the viscosity, $\Delta = (\partial_{x_1}^2 + \partial_{x_2}^2, \partial_{x_1}^2 + \partial_{x_2}^2)'$ the Laplacian operator, $\nabla = (\partial_{x_1}, \partial_{x_2})'$ the gradient, $p : \Omega \times [0, \infty) \rightarrow \mathbb{R}$ the pressure function, and $f_0 : \Omega \times [0, T] \rightarrow \mathbb{R}^2$ an exogenous forcing. $T > 0$ denotes the length of the time period under consideration. The periodic boundary conditions give $v(\cdot, 0, t) = v(\cdot, 2\pi, t)$. Following for instance [16], one can show existence and uniqueness of a weak solution of (3)-(4), over the time interval $[0, T]$, under the assumptions that $v_0 \in H$, $f \in L^2(0, T; H)$ where

$$H = \left\{ v \in L_{\text{per}}^2(\Omega)^2 \mid \nabla \cdot v = 0, \int_{\Omega} v(x) dx = 0 \right\}.$$

Here, $L_{\text{per}}^2(\Omega)$ denotes the Hilbert space of squared-integrable functions with periodic boundary conditions; the solution is such that $v_i, \partial_{x_j} v_i \in L^2(\Omega \times [0, T])$ and $v : [0, T] \mapsto H$ is continuous. The Leray-Helmholtz decomposition expresses in a unique way (given the boundary conditions) vector fields into a gradient plus a curl vector (i.e. vector of zero divergence). The Leray projector $P : (L_{\text{per}}^2(\Omega))^2 \rightarrow H$ maps a vector onto the curl vector in the decomposition. Applying P on both sides of (3) (considering also the boundary conditions and (4)), provides the functional evolution equation

$$dv + \nu Av dt + B(v, v) dt = f(t) dt, \quad v(0) = v_0, \quad (5)$$

where following standard notation, we have defined

$$A := -\Delta, \quad B(u, v) = P((v \cdot \nabla)u), \quad f := Pf_0,$$

with A called the Stokes operator in the literature. It is convenient to work with the Fourier characterisation of the relevant spaces, that is we use the equivalent representation of H ,

$$H = \left\{ u = \sum_{k \in \mathbb{Z}^2 \setminus \{0\}} u_k \psi_k(x) \mid u_{-k} = -\overline{u_k}, \sum_{k \in \mathbb{Z}^2 \setminus \{0\}} |u_k|^2 < \infty \right\} \quad (6)$$

where we have defined the orthonormal basis functions for H ,

$$\psi_k(x) = \frac{1}{2\pi} \frac{k^\perp}{|k|} e^{i k \cdot x}, \quad k \in \mathbb{Z}^2 \setminus \{0\}, \quad k^\perp := (-k_2, k_1)'$$

Given, the re-formulation of NSE in (5) and the basis in (6), one can introduce additive noise in the dynamics in a standard manner. First, we define the upper half-plane of wavenumbers

$$\begin{aligned} \mathbb{Z}_\uparrow^2 &= \{k = (k_1, k_2) \in \mathbb{Z}^2 \setminus \{0\} : k_1 + k_2 > 0\} \\ &\cup \{k = (k_1, k_2) \in \mathbb{Z}^2 \setminus \{0\} : k_1 + k_2 = 0, k_1 > 0\}. \end{aligned}$$

Let

$$Z_k(t) = Z_k^{re}(t) + i Z_k^{im}(t), \quad k \in \mathbb{Z}_\uparrow^2,$$

where $\{Z_k^{re}, Z_k^{im}\}$ are (independent) standard Brownian motions on $[0, T]$. In the spirit of [9, Section 4.1], consider a covariance operator Q such that $Q\psi_k = \sigma_k^2 \psi_k$, for $\sigma_k^2 > 0$, $\sigma_{-k} = \sigma_k$. Then, we can define the Q -Wiener process as

$$Q^{\frac{1}{2}} W(t) := \sum_{k \in \mathbb{Z}^2 \setminus \{0\}} \sigma_k Z_k(t) \psi_k(x), \quad (7)$$

under the requirement

$$Z_{-k} \equiv -\overline{Z_k}, \quad k \in \mathbb{Z}_\uparrow^2. \quad (8)$$

Thus, we are working with a diagonal covariance matrix (w.r.t. the relevant basis of interest), though other choices could easily be considered. We will also work under the scenario that $\sigma_k^2 = O(|k|^{-2(1+\epsilon)})$, for some $\epsilon > 0$, so that $\sum_{k \in \mathbb{Z}^2 \setminus \{0\}} \sigma_k^2 < \infty$, i.e. Q is trace-class operator. We use $\mathbb{W}(\cdot)$ to denote the Q -Wiener measure on $[0, T]$.

Having introduced the random component, we are now interested in weak solutions $V = V(t)$ of the functional SDE,

$$dV(t) + \nu AV(t) dt + B(V(t), V(t))dt = f(t) dt + Q^{\frac{1}{2}} dW(t), \quad V(0) = v_0, \quad (9)$$

with the solution understood pathwise on the probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \geq 0}, \mathbb{P})$. More formally, following [14], we define the spaces

$$\mathcal{V}_s := \left\{ u = \sum_{k \in \mathbb{Z}^2 \setminus \{0\}} u_k \psi_k(x) \mid u_{-k} = -\overline{u_k}, \sum_{k \in \mathbb{Z}^2 \setminus \{0\}} |k|^{2s} |u_k|^2 < \infty \right\}, \quad s \in \mathbb{R}.$$

Since the operator $Q^{1/2}$ is linear, bounded in H and $\text{Im}(Q^{1/2}) \equiv \mathcal{V}_{1+\epsilon}$, [14, Theorem 6.1] implies that for $v_0 \in \mathcal{V}_1$ and $f \in C([0, T]; \mathcal{V}_1)$, there exists a unique solution for (9) such that $V \in C([0, T]; \mathcal{V}_1)$. In [14, 15] one may also find more details on the existence of an invariant distribution, together with irreducibility and Feller properties of the corresponding Markov transition kernel.

2.2 Galerkin Projections and Computational Considerations

Using the Fourier basis (6), we can write the solution as

$$\begin{aligned} V(t) &= \sum_{k \in \mathbb{Z}^2 \setminus \{0\}} u_k(t) \psi_k(x), \quad u_{-k}(t) \equiv -\overline{u_k(t)}, \\ u_k(t) &= \langle V(t), \psi_k \rangle = \int_{\Omega} V(t) \cdot \overline{\psi_k(x)} dx. \end{aligned}$$

Hence, it is equivalent to consider the parameterisation of V via $\{u_k(t)\}_{k \in \mathbb{Z}^2 \setminus \{0\}}$. By taking the inner product with ψ_k on both sides of (9), it is straightforward to obtain that the u_k 's obey the following infinite-dimensional SDE

$$\begin{aligned} du_k(t) &= -\nu |k|^2 u_k(t) dt \\ &\quad - \sum_{m, p \in \mathbb{Z}^2 \setminus \{0\}} b_{k, m, p} u_m(t) u_p(t) dt + f_k(t) dt + \sigma_k dZ_k(t), \quad k \in \mathbb{Z}_\uparrow^2, \end{aligned} \quad (10)$$

with

$$b_{k, m, p} = \langle B(\psi_m, \psi_p), \psi_k \rangle, \quad f_k(t) = \langle f(t), \psi_k \rangle.$$

Recall that due to $V(t)$ being a real field $u_{-k}(t) \equiv -\overline{u_k(t)}$, $k \in \mathbb{Z}_\uparrow^2$. This parameterisation of V is more convenient as it allows performing inference on a vector (even if infinitely long), with coordinates evolving according to an SDE. For numerical purposes one is forced to use Galerkin discretisations, using projections of V onto a finite Hilbert space instead. Consider the set of wavenumbers in

$$\mathbb{L} = \{k \in \mathbb{Z}_\uparrow^2 : (k_1 \vee k_2) \leq L\},$$

for some integer $L > 0$, and define the finite dimensional-subspace H_L via the projection $P_L : H \rightarrow H_L$ so that

$$P_L v = \sum_{k \in \mathbb{L}} \langle v, \psi_k \rangle \psi_k.$$

Then, inferring the Galerkin projection for V corresponds to inferring the vector $\{u_k(t)\}_{k \in \mathbb{L}}$ that obeys the following finite-dimensional SDE

$$du_k(t) = -\nu|k|^2 u_k(t) - \sum_{m,p \in \mathbb{L}} b_{k,m,p} u_m(t) u_p(t) dt + f_k(t) dt + \sigma_k dZ_k(t), \quad k \in \mathbb{L}. \quad (11)$$

This high dimensional SDE will provide an approximation for the infinite dimensional SPDE. Such an inference problem is more standard, but is still challenging due to the high dimensionality of \mathbb{L} and the non-linearities involved in the summation term of the drift function in (11). Since (11) is only an approximation of (10), it will induce a bias in the inferential procedure. In our paper, we do not study the size of this bias. Instead we concentrate our efforts on designing an algorithm to approximate π_n (in (2)) that is robust to *mesh refinement*. This means our method should perform well numerically when one increases L (and, indeed, reducing the bias in the numerical approximation of (10)). Naturally this would be at the expense of adding computational effort at a moderate amount, but this will depend on the particular numerical scheme used to approximate the solution of (11). For instance, for the FFT based numerical schemes used in Section 5 the computational cost is $\mathcal{O}(L^2 \log L)$.

2.3 The Distribution of v_0

We assume that the initial condition of V is random and distributed according to the following Gaussian process prior:

$$\pi_0 = \mathcal{N}(\mu, \beta^2 A^{-\alpha}), \quad \alpha > 2, \beta > 0, \mu \in \mathcal{V}_1, \quad (12)$$

with hyper-parameters α, β affecting the roughness and magnitude of the initial vector field. This is a convenient but still flexible enough choice of a prior; see [9, Sections 2.3 and 4.1] for more details on Gaussian distributions on Hilbert spaces. Notice that π_0 admits the Karhunen-Loève expansion

$$\pi_0 = \mathcal{L}aw \left(\sum_{k \in \mathbb{Z}^2 \setminus \{0\}} \left(\mu_k + \frac{\beta}{\sqrt{2}} |k|^{-\alpha} \xi_k \psi_k \right) \right),$$

with $\mu_k = \langle \mu, \psi_k \rangle$, $k \in \mathbb{Z}_\uparrow^2$, (so, necessarily $\mu_{-k} = -\overline{\mu_k}$, $k \in \mathbb{Z}_\uparrow^2$) and

$$\operatorname{Re}(\xi_k), \operatorname{Im}(\xi_k) \stackrel{iid}{\sim} \mathcal{N}(0, 1), \quad k \in \mathbb{Z}_\uparrow^2; \quad \xi_{-k} = -\overline{\xi_k}, \quad k \in \mathbb{Z}_\uparrow^2.$$

Since the covariance operator is determined via the Stokes operator A , one can easily check that the choice $\alpha > 2$ implies that for $v_0 \in \mathcal{V}_1$, π_0 -a.s., thus the conditions for existence of weak solution of (9) in [14, Theorem 6.1] are satisfied a.s. in the initial condition. Notice that sampling from π_0 is straightforward.

3 The Stochastic Filtering Problem

In Section 2 we defined the SPDE providing the unknown signal, i.e. the object we are interested in performing Bayesian inference upon. In this section we present the non-linear filtering problem in detail. We begin by discussing the observations. We assume that the vector field V is unknown, but generates a sequence of noisy observations $\mathcal{Y}_n = (Y_{t_1}, \dots, Y_{t_n})$ at ordered discrete time instances $(t_p)_{p=1, \dots, n}$ with $t_n < t_{n+1} < T$ for all n , with $Y_{t_i} \in \mathbb{R}^{2d_y}$, for $d_y \geq 1$. Each observation vector Y_{t_i} is further assumed to originate from the following observation equation

$$Y_{t_n} = FV(t_n) + \Xi_n, \quad \Xi_n \sim \mathcal{N}(0, \Sigma), \quad (13)$$

where F is a bounded linear operator $F : H \rightarrow \mathbb{R}^{d_y}$ and $\Sigma \in \mathbb{R}^{d_y \times d_y}$ is symmetric positive-definite. One can then write the observation likelihood at instance t_n as

$$p(Y_{t_n} | V(t_n)) = \frac{\exp \left(-\frac{1}{2} \left| \Sigma^{-\frac{1}{2}} (Y_{t_n} - FV(t_n)) \right|^2 \right)}{(2\pi)^{d_y/2} |\Sigma|^{1/2}}.$$

Using a linear observation model is restrictive but it does include typical observation schemes used in practice. We focus our attention at the case when Y_{t_n} is a noisy measurement of the velocity field at different fixed stationary points $x_l \in \Omega$, $l = 1, \dots, p$. This setting is often referred to as *Eulerian data assimilation*. In particular we have that

$$F = (F'_1, \dots, F'_{d_y})',$$

with F_l denoting a spatial average over a (typically small) region around x_l , $l = 1, \dots, d_y$, say $B_{x_l}(r) = \{x \in \Omega : |x - x_l| \leq r\}$ for some radius $r > 0$; that is, F_l is the following integral operator

$$F_l V(t) = \frac{1}{|B_{x_l}(r)|} \int_{B_{x_l}(r)} V(t, x) dx, \quad (14)$$

with $|B_{x_l}(r)|$ denoting the area of $B_{x_l}(r)$. In what follows, other integral operators could also be similarly used, such as $F_l V(t) = (\int_{\Omega} V(t, x) w_{x_l}(x) dx) / (\int_{\Omega} w_{x_l}(x) dx)$, with $w_{x_l} \in L^2(\Omega)$ being appropriate weighting functions that decay as $|x - x_l|$ grows.

Earlier in the introduction, the filtering problem was defined as the task of computing the conditional distribution $\pi_n(\cdot) = \mathbb{P}[V(t_n) \in \cdot | \mathcal{Y}_n]$. Due to the nature of the observations, it is clear we are dealing with a discrete time filtering problem. A particular challenge here (in common with other typical non-linear SPDEs) is that the distribution of the associated Markov transition kernel, $\mathbb{P}[V(t_n) \in \cdot | V(t_{n-1}) = v]$, is intractable. Still, it is possible to simulate from the unconditional dynamics of $V(t)$ given $V(t_{n-1}) = v$ using standard time discretization techniques. (The simulated path introduces a time discretisation bias, but its effect is ignored in this paper.)

We aim to infer the following posterior distribution, based on the *continuous time signal*

$$\Pi_n(\cdot) = \mathbb{P}[V^n \in \cdot | \mathcal{Y}_n], \quad V^n := (V(t))_{t \in [0, t_n]};$$

we also denote

$$V_{n-1}^n = (V(t))_{t \in (t_{n-1}, t_n]}.$$

This data augmentation approach - when one applying importance sampling on continuous time - has appeared in [19] for a related problem and in [34] for filtering problems involving certain multivariate SDEs. We proceed by writing the filtering recursion for Π_n . We denote the law of V in (9) for the time interval between t_{n-1} and t_n as

$$\mathbb{V}_{n-1}^n(\cdot | v) := \mathbb{P}[(V(t))_{t \in (t_{n-1}, t_n]} \in \cdot | V(t_{n-1}) = v].$$

Then, one may use Bayes rule to write Π_n recursively as

$$\frac{d\Pi_n}{d(\Pi_{n-1} \otimes \mathbb{V}_{n-1}^n)}(V^n) = \frac{p(Y_{t_n} | V(t_n))}{p(Y_{t_n} | \mathcal{Y}_{n-1})}, \quad (15)$$

where $p(Y_{t_n} | \mathcal{Y}_{n-1}) = \int p(Y_{t_n} | V(t_n)) [\Pi_{n-1} \otimes \mathbb{V}_{n-1}^n](dV^n)$.

In addition, one can attempt to propose paths from an appropriate SPDE different from (9), say

$$\begin{aligned} d\tilde{V}(t) + \nu A\tilde{V}(t)dt + B(\tilde{V}(t), \tilde{V}(t))dt \\ = Q^{\frac{1}{2}}g(t, \tilde{V}(t))dt + f(t)dt + Q^{\frac{1}{2}}dW(t), \quad t \in (t_{n-1}, t_n], \end{aligned} \quad (16)$$

where $g : [0, T] \times H \mapsto H$ and $Q^{\frac{1}{2}}W_t$ is a Q -Wiener process on $(t_{n-1}, t_n]$. We define

$$\mathbb{Q}_{n-1}^n(\cdot | v) := \mathbb{P}[(\tilde{V}(t))_{t \in (t_{n-1}, t_n]} \in \cdot | \tilde{V}(t_{n-1}) = v].$$

One needs to ensure that the change of drift g is appropriately chosen so that a Girsanov theorem holds and $\mathbb{V}_{n-1}^n(\cdot | v)$ is absolutely continuous with respect to $\mathbb{Q}_{n-1}^n(\cdot | v)$ for all relevant v , with the recursion in (15) becoming

$$\frac{d\Pi_n}{d(\Pi_{n-1} \otimes \mathbb{Q}_{n-1}^n)}(V^{n-1}, \tilde{V}_{n-1}^n) \propto p(Y_{t_n} | \tilde{V}(t_n)) \cdot \frac{d\mathbb{V}_{n-1}^n}{d\mathbb{Q}_{n-1}^n}(\tilde{V}_{n-1}^n | V^{n-1}(t_{n-1})). \quad (17)$$

for $(V^{n-1}, \tilde{V}_{n-1}^n)$ assumed to be typical elements of the sample space of either of the two probability measures above (e.g., all such paths are assumed to possess relevant continuity properties at t_{n-1}).

In the context of particle filtering and IS one aims to design g in a way that the proposed trajectories are in locations where Π_n is higher. This in turn implies that the importance weights in (17) will exhibit much less variance than the ones from the prior signal dynamics, hence the design of g is critical for generating effective Monte Carlo approximations.

4 Particle Filtering

We are interested in approximating the distribution Π_n using a particle filter approach. We present in Algorithm 1 a naive particle filter algorithm that provides the particle approximations:

$$\Pi_n^N = \sum_{j=1}^N \mathscr{W}_n^j \delta_{V^j} \quad \text{or} \quad \bar{\Pi}_n^N = \frac{1}{N} \sum_{j=1}^N \delta_{\tilde{V}^j}.$$

Such a particle filter will be typically overwhelmed by the dimensionality of the problem and will not be able to provide accurate solutions with a moderate computational cost. When $g = 0$ in (16), the algorithm corresponds to a standard bootstrap particle filter. For the latter, it is well known in the literature ([5, 35]) that it exhibits weight degeneracy in the presence of large dissimilarity between $\Pi_{n-1} \otimes \mathbb{V}_{n-1}^n$ and Π_n , which can be caused in our context by the high dimensionality of the state space and the complexity of the SPDE dynamics. When g is well designed then the particles can be guided in areas of larger importance weights and the algorithmic performance can be considerably improved, but this modification may still not be sufficient for obtaining a robust and efficient algorithm.

Algorithm 1 A naive particle filter

- Initialise $V_0^i \sim \pi_0$, $1 \leq i \leq N$.
- For $n \geq 1$
 1. For $i = 1, \dots, N$: sample independently

$$\tilde{V}_{n-1}^{n,i} \sim \mathbb{Q}_{n-1}^n(\cdot | V^{n-1}(t_{n-1})^i).$$

2. For $i = 1, \dots, N$: compute importance weights

$$\mathscr{W}_n^i \propto p(Y_{t_n} | \tilde{V}(t_n)^i) \cdot \frac{d\mathbb{V}_{n-1}^n}{d\mathbb{Q}_{n-1}^n}(\tilde{V}_{n-1}^{n,i} | V^{n-1}(t_{n-1})^i), \quad \text{s.t.} \quad \sum_{i=1}^N \mathscr{W}_n^i = 1.$$

3. For $i = 1, \dots, N$: resample

$$V^{n,i} \sim \sum_{j=1}^N \mathscr{W}_n^j \delta_{(V^{n-1,j}, \tilde{V}_{n-1}^{n,j})}(\cdot).$$

In the remainder of this section, we will discuss how to improve upon this first attempt to tackle the high-dimensional filtering problem at hand using the following ingredients: (i) specifying a particular form of g in (16) that results in gains of efficiency, (ii) using adaptive tempering, and (iii) MCMC moves. Guided proposals and tempering are employed to bridge the dissimilarity between $\Pi_{n-1} \otimes \mathbb{V}_{n-1}^n$ and Π_n . The MCMC steps are required for injecting additional diversity in the particle population, which would otherwise diminish gradually due to successive resampling and tempering steps. The method is summarised in Algorithm 2. In the following subsections we explain in detail our implementation of (i)-(iii) mentioned above.

4.1 Likelihood-Informed Proposals

In the importance weight of (17) we are using a Girsanov Theorem and assume absolute continuity between SPDEs (16) and (9) when started at the same position. Under the assumption

$$\mathbb{P}\left[\int_0^T \|g(t, V(t))\|^2 dt < \infty\right] = 1, \quad (18)$$

absolute continuity indeed holds and we have Radon-Nikodym derivative

$$\begin{aligned} \log \frac{d\mathbb{V}_{n-1}^n}{d\mathbb{Q}_{n-1}^n}(V_{n-1}^n | V^{n-1}(t_{n-1})) \\ = - \int_{t_{n-1}}^{t_n} \langle g(t, V(t)), Q^{\frac{1}{2}} dW(t) \rangle_0 - \frac{1}{2} \int_{t_{n-1}}^{t_n} \|g(t, V(t))\|_0^2 dt, \end{aligned}$$

where

$$\langle u, v \rangle_0 := \langle Q^{-\frac{1}{2}} u, Q^{-\frac{1}{2}} v \rangle \equiv \sum_{k \in \mathbb{Z}^2 \setminus \{0\}} \frac{1}{\sigma_k^2} \langle u, \psi_k \rangle \langle v, \psi_k \rangle;$$

see [9, Theorem 10.14] and [9, Lemma 10.15] for details. It remains to provide an effective design for g . One can use proposals developed for problems whereby a finite-dimensional SDE generates linear Gaussian observations and one is interested to perform a similar IS method, see e.g. [18, 42, 31, 32, 38]. In this paper we use the proposal employed in [18] and set

$$g(t, V(t)) = Q^{\frac{1}{2}} F^* (\Sigma + (t_n - t) F Q F^*)^{-1} (Y_{t_n} - FV(t)), \quad t \in (t_{n-1}, t_n], \quad (19)$$

where F^* denotes the adjoint of F . The guiding function g could be interpreted as an one-step Euler approximation of the h -transform needed to evolve $V(t)$ conditional on the acquired observation Y_{t_n} within the interval $(t_{n-1}, t_n]$. It is not hard to verify (18) for this choice of g . Since Σ, Q are invertible then $(\Sigma + (t_n - t) F Q F^*)^{-1}$ exists via the Sherman-Morrison-Woodbury identity and $Q^{\frac{1}{2}} F^* (\Sigma + (t_n - t) F Q F^*)^{-1}$ is a bounded linear operator. Then (18) holds from [9, Proposition 10.18] and [26, Proposition 2.4.9] that imply that there exists a $\delta > 0$ such that

$$\sup_{t \in [0, T]} \mathbb{E} \left[\exp \left(\delta \|g(t, V(t))\|^2 \right) \right] < \infty,$$

which implies (18).

For the finite-dimensional SDE case more elaborate guiding functions can be found in [42, 38] and some of these could be potentially extended so that they can be used in the SPDE setting instead of (19). The advantage of using g in (19) is that it provides a simple functional and can perform well for problems where $t_n - t_{n-1}$ is of moderate length, as also confirmed in the numerical examples of Section 5.

4.2 Bridging Π_{n-1} and Π_n With Tempering

In a high-dimensional setting, using well-designed likelihood-informed proposals might not suffice for the purposes of providing importance weights that do not degenerate. Hence, we also make use of a tempering scheme introduced in [21] to bridge the potentially large dissimilarity between the informed proposal $\Pi_{n-1} \otimes \mathbb{Q}_{n-1}^n$ and the target Π_n . Algorithm 2 targets a sequence of distributions $(\Pi_p)_{p=1, \dots, n}$ and provides corresponding particle approximations $(\Pi_p^N)_{p=1, \dots, n}$. When moving from Π_{n-1} to Π_n , it introduces a user-specified intermediate sequence of target distributions $(\Pi_{n,l})_{l=0, \dots, \tau_n}$. Each $\Pi_{n,l}$ can be defined using a tempering approach so that

$$\frac{d\Pi_{n,l}}{d(\Pi_{n-1} \otimes \mathbb{Q}_{n-1}^n)}(V^n) \propto \left(\frac{d\mathbb{V}_{n-1}^n}{d\mathbb{Q}_{n-1}^n}(V_{n-1}^n) \cdot p(Y_{t_n} | V^n(t_n)) \right)^{\phi_{n,l}},$$

for inverse temperatures $0 = \phi_{n,0} < \phi_{n,1} < \dots < \phi_{n,\tau_n} = 1$. Note that by construction we have $\Pi_{n,\tau_n} \equiv \Pi_n$. We can determine the temperatures $\phi_{n,l}$ and their number τ_n on the fly via an adaptive procedure. Assume we are at the n -th step of the algorithm, have completed $l - 1$ tempering steps, and we have equally weighted particles. The next temperature is determined by expressing the weights as a function of ϕ

$$\begin{aligned} \mathscr{W}_{n,l}^j(\phi) &\propto \left(\frac{dV_{n-1}^n}{dQ_{n-1}^n}(V_{n-1}^{n,j}) p(Y_{t_n} | V^n(t_n)^j) \right)^{\phi - \phi_{n,l-1}}, \quad \phi_{n,l-1} < \phi \leq 1, \\ \sum_{i=1}^N \mathscr{W}_{n,l}^i(\phi) &= 1, \end{aligned} \quad (20)$$

and determining $\phi_{n,l}$ via a requirement based on a quality criterion for the particle population. We use here the Effective Sample Size (ESS), and set

$$\phi_{n,l} = \inf \left\{ \phi \in (\phi_{n,l-1}, 1] : ESS_{n,l}(\phi) := \frac{1}{\sum_{j=1}^N \{\mathscr{W}_{n,l}^j(\phi)\}^2} \leq \alpha N \right\}, \quad (21)$$

(under the convention that $\inf \emptyset = 1$) with a user-specified fraction $\alpha \in (0, 1)$. Equation (21) can be easily solved numerically using for instance a standard bisection method. This approach leads to a particle approximation for $\Pi_{n,l}$, say

$$\Pi_{n,l}^N = \sum_{i=1}^N \mathscr{W}_{n,l}^i(\phi_{n,l}) \delta_{V^{n,i}};$$

we then propose to resample from $\Pi_{n,l}^N$ so that one ends up with equally weighted particles.

4.3 Adding Particle Diversity With MCMC kernels

Successive resampling due to the tempering steps leads to sample impoverishment unless the method re-injects sampling diversity. To achieve this, we propose using a small number of iterations from a MCMC procedure that leaves $\Pi_{n,l}$ invariant. Let $\mathcal{K}_{n,l}$ denote the Markov probability kernel obtained from the MCMC procedure, given as

$$\mathcal{K}_{n,l}[V' \in \cdot | V] = \alpha(V, V') \mathcal{Q}[V' \in \cdot | V] + \delta_V(\cdot) \left(1 - \int \alpha(V, V') \mathcal{Q}[dV' | V] \right),$$

with \mathcal{Q} denoting the proposal kernel and α the acceptance probability in a standard Metropolis-Hastings framework. We use a particular MCMC design similar to [19] that is well defined on function spaces (based on theory for MCMC on general state spaces [37]). The design is often referred to as preconditioned Crank-Nicolson, abbreviated here to pCN; see [36, 8] for a detailed review. For a basic description: let Λ be a probability measure that is absolutely continuous with respect to $\Pi_{n,l}$ with Radon-Nikodym derivative $(d\Pi_{n,l}/d\Lambda)(V) =: \vartheta(V)$. Similar to [36, 8, 19] we specify the proposal kernel \mathcal{Q} to satisfy detailed balance with respect to Λ , i.e. $\mathcal{Q}(dV' | V) \Lambda(dV) = \mathcal{Q}(dV | V') \Lambda(dV')$. Then, using $\alpha_{n,l}(V, V') = 1 \wedge (\vartheta(V')/\vartheta(V))$ provides a kernel $\mathcal{K}_{n,l}$ which is $\Pi_{n,l}$ -invariant (by [37, Theorem 2]).

Next we discuss the design of \mathcal{Q} for our problem. First, we stress that for the purpose of particle filtering we are mainly interested on the invariance property of $\mathcal{K}_{n,l}$ and not necessarily its ergodic properties on the full space. With this in mind \mathcal{Q} can be a Markov kernel that generates proposals V' with $V'_s = V_s$ for $s \leq t_{n-1}$. This allows for on-line computation at each n, l ; at the same time reversibility holds as Proposition 1 and Theorem 2 in [37] still hold for such proposals. From a practical perspective, we are adding noise to the path of the hidden signal only for within $[t_{n-1}, t_n]$. Then, we need to specify Λ and \mathcal{Q} . One possibility is to take under consideration that $V_{n-1}^n = V_{n-1}^n(W)$ with W being the driving noise that generated V_{n-1}^n (note that we can assume that $W(t_{n-1}) = 0$ without

loss of generality, since the V -path uses the increments of W). Suppose also that both V_{n-1}^n and W are stored in the computer's memory and let $\Lambda = \Pi_{n-1} \otimes \mathbb{Q}_{n-1}^n$ so that

$$\vartheta(V) = \frac{d\Pi_{n,l}}{d\Lambda}(V) = \left(\frac{dV_{n-1}^n}{d\mathbb{Q}_{n-1}^n}(V) p(Y_{t_n}|V(t_n)) \right)^{\phi_{n,l}}.$$

To simulate from a Λ -preserving proposal one first generates a new noise sample W'

$$W(s)' = \rho W(s) + \sqrt{1 - \rho^2} \xi(s), \quad t_{n-1} < s \leq t_n, \quad \xi \sim \mathbb{W}, \quad (22)$$

where $W(s)$ is the noise driving V . To return to the original space, we use the new noise W' to solve for V' in (16). A standard calculation can show that $W' \sim \mathbb{W}$, which in turn implies that for the part of the proposal V' in $(t_{n-1}, t_n]$, $(V_{n-1}^n)' \sim \mathbb{Q}_{n-1}^n$ holds. Reversibility with respect to Λ is ensured using a simple conditioning and marginalization argument presented in the Appendix. The MCMC implementation using the m -th iterate of $\mathcal{K}_{n,l}$ for some $m \geq 1$, denoted $\mathcal{K}_{n,l}^m$, is shown in Algorithm 3.

Remark 1. We note that there are some extensions not mentioned here: a) similar in spirit with [12] the lower bound on the time we start adding noise (here t_{n-1}) could be made smaller and this can be beneficial in terms of adding diversity, but for the sake of simplicity we do not pursue this further; b) similar to [25], one could improve the mixing of $\mathcal{K}_{n,l}$ by estimating the mean of $\Pi_{n,l}$ from the particles and for low k -s use this information in the proposal in (22).

Remark 2. The convergence of Algorithm 2 has been studied in [4, 17].

Algorithm 2 Adaptive Particle Filtering Algorithm

- At $n = 0$. For $i = 1, \dots, N$, sample i.i.d. $V_0^i \sim \pi_0$, and set $\mathscr{W}_0^i = 1/N$.
- At time $n \geq 1$.
 1. For $i = 1, \dots, N$: sample independently

$$X_n^i \sim \mathbb{Q}_{n-1}^n(\cdot | V^{n-1,i}(t_{n-1}))$$

2. Set $l = 0$, $X_{n,0}^i = X_n^i$, $\Pi_{n,0} = \Pi_{n-1} \otimes \mathbb{Q}_{n-1}^n$, $\phi_{n,0} = 0$.
3. While $\phi_{n,l} < 1$
 - (a) Set $l \leftarrow l + 1$
 - (b) Specify $\Pi_{n,l}$, $\phi_{n,l}$ based on the ESS computation in (21)
 - (c) For $i = 1, \dots, N$
 - i. Compute weights $\mathscr{W}_{n,l}^i$ as in (20)
 - ii. Resample and move particles:

$$X_{n,l}^i \stackrel{i.i.d.}{\sim} \sum_{j=1}^N \frac{\mathscr{W}_{n,l}^j}{\sum_{k=1}^N \mathscr{W}_{n,l}^k} \mathcal{K}_{n,l}^m(\cdot | X_{n,l-1}^j)$$

4. If $\phi_{n,l} = 1$ return $V^{n,i} = (V^{n-1,i}, X_{n,l}^i)$, $\tau_n = l$; otherwise go back to Step 3.
-

5 Numerical examples

We solve SPDE (9) for $\nu = 0.1$ and $f = 0$ numerically using the exponential Euler scheme [22] for the finite-dimensional projection (11). For (11), we use a Fourier truncation with $L = 64$ i.e. $-64 \leq$

$k_1, k_2 \leq 64$. For π_0 we use $\beta = 0.5$, $\alpha = 3$ and $\mu = v_0^\dagger$, with v_0^\dagger being a random sample from $\mathcal{N}(0, A^{-\alpha})$ that is also used as the true signal to generate the observations. To determine Q we use $\sigma_k = \sqrt{2\delta\nu}|k|^{-3}$ with $\delta = 1$. For the observation equation in (13) we use $\Sigma = 0.8I$ and for the observer in (14) we place the observers' locations x_l on a uniform square grid with equal spacing and set r to be small (smaller than $2\pi/L$). Thus, we can make the likelihood more informative by decreasing the observation noise or by increasing the grid size. As the information in the likelihood increases, one expects a larger number of tempering steps (and slower total execution times). When no tempering is used this will lead to a much lower value for the ESS.

We present results from two types of experiments with simulated observations. In the first case we will look at a batch of $n = 5$ observations from a dense grid (16×16). We use this short run to illustrate the efficiency and performance of the methodology. The length of the data-set allows using multiple independent runs for the same observations. In the second experiment we use a large number of observations ($n = 100$) obtained from a 8×8 grid using both Gaussian and Student-t distributed additive noise. We show that the method performs well for the longer time and performance is similar for both Gaussian and non-Gaussian observations.

We begin with the case of $n = 5$ and dense observation grid (16×16). In Table 1 we present results for $N = 100$ and $\delta t_n = 0.4$ comparing a naive bootstrap PF, a PF that uses the informed proposal (16) for IS but without tempering (both based on Algorithm 1), a PF that uses tempering when sampling from the stochastic NSE dynamics in (9), and a PF that uses both tempering and (16) for IS. We show the number of tempering steps per batch of observations, the ESS at each observation time t_p , and L^2 -errors between the true signal vorticity w^\dagger and the estimated posterior mean \hat{w} at each epoch, i.e.

$$\int_{\Omega} \|\hat{w}(x, t_i) - w^\dagger(x, t_i)\|^2 dx.$$

When tempering is used we present in Figure 1 selected typical PDFs and scatter plots for a few chosen frequencies k -s. The advantage of using (16) is clear when noting the dispersion of the particles relative to sampling from (9). In Table 1 it is evident that when using tempering, IS resulted to about half of the tempering steps than when sampling from (9). In both cases, the tuning of the MCMC steps lead to the same acceptance ratio (around 0.2 at the final tempering step). We use $m = 20$ MCMC iterations per tempering for $n = 1$. For $n > 1$, plain tempering uses $m = 20$ and IS with (16) uses $m = 10$. In addition, the IS-tempering case uses a larger step size ρ for the MCMC (0.5 rather than 0.9). This results in lower total computational cost and runtimes when IS is used despite the added computations imposed by computing g in (16). We also note that a lower number of tempering steps is beneficial in addressing potential path degeneracy issues.

We proceed with the second numerical experiment, where we use only a single run of a PF with both tempering and IS for $N = 100$ and $n = 1, \dots, 100$. The dynamics for the state are as before, but for the observations we use a 8×8 equally spaced observation grid and look at two cases for the distribution of the noise Ξ_n : a zero mean Gaussian and zero mean Student-t distribution with 4 degrees of freedom. In both cases $\Sigma = 0.8I$. In Figure 2 we plot the estimated vorticity posterior mean for $n = 10, 50, 100$, in each case together with the vorticity of the true signal. We also provide in Figure 3 a plot of the ratio of the posterior variance of the vorticity of V_{t_n} over the unconditional variance when obeying the probability law determined by the stochastic NSE dynamics in (9). In Figure 4 we plot ESS, L^2 -errors as before and number of tempering steps per iteration. In both cases the performance is fairly stable with time and the algorithm provides good posterior mean estimates.

6 Discussion

We have presented a particle filtering methodology that uses likelihood-informed IS proposals, tempering and MCMC moves for signals obeying the stochastic NSE observed with additive noise. The approach is computationally intensive and requires a significant number of particles N , but we believe the cost is quite moderate relatively to the dimensionality of the problem. The use of tempering and

$n =$	number of tempering steps					ESS					MCMC step sizes	
	1	2	3	4	5	1	2	3	4	5	$n = 1$	$n \geq 1$
IS-PF-tempering	5.6	4.7	4.4	4	4.3	64.87	73.88	63.02	57.01	53.03	$\rho_0 = 0.9$	$\rho = 0.5$
PF-tempering	10.1	7.7	7.4	7.6	8.1	77.73	70.62	68.50	75.88	82.02	$\rho_0 = 0.98$	$\rho = 0.9$
IS-PF	n/a					1.16	1.92	1.56	2.11	1.90	n/a	
bootstrap PF	n/a					1.00	1.00	1.01	1.13	1.06	n/a	

$n =$	L^2 -errors				
	1	2	3	4	5
IS-PF-tempering	0.1901 (0.0012)	0.2612 (0.0002)	0.2080 (0.0003)	0.1607 (0.0001)	0.2664 (0.0005)
PF-tempering	0.4305 (0.0110)	0.3205 (0.0029)	0.2498 (0.0054)	0.2339 (0.0023)	0.3806 (0.0137)
IS-PF	0.3078 (0.0033)	0.4536 (0.0166)	0.4193 (0.0062)	0.3345 (0.0062)	0.4553 (0.0023)
bootstrap PF	0.8458 (0.0185)	1.1253 (0.1493)	0.8604 (0.0810)	0.9643 (0.0260)	1.1542 (0.0467)

Table 1: Average results for number of tempering steps, ESS and L^2 -errors (with standard deviations in parenthesis) obtained from 10 independent runs of each algorithm. In all cases we use $N = 100$, $\delta t_n = 0.4$. For the tempering based PF without guided proposal we use $m = 20$ MCMC steps and for IS with tempering we use $m = 10$ for $n > 1$ and $m = 20$ for $n = 1$. We also show MCMC tuning parameters for Algorithm 3.

MCMC steps is crucial for this high-dimensional application. The inclusion of likelihood-informed proposals results in higher efficiency and ESS , less tempering steps and higher step sizes for the MCMC steps - thus, overall, in lower computational cost. The IS proposals were designed using a Gaussian noise assumption for the observations, but we demonstrated numerically that they are still useful and efficient for observation noise obeying a Student-t distribution with heavier tails. In addition, as δt_n increases using proposals as in (16) will be more beneficial.

In the experiments presented in Section 5 the effective dimensionality of the problem is determined by ν, γ and σ_k . More challenging parameterizations than the ones presented here could be dealt with by increasing N or via a more advanced numerical method for the solution of the SPDE. In addition the MCMC steps can be made even more effective by including information for the empirical mean and covariance of the particles in a spirit similar to [25]. Other potentially useful extensions for challenging cases could be regenerating parts of the path of V before t_{n-1} as in [12] or using different number of particles for different ranges of k as in [24]. Furthermore we note that we did not make use of parallelization, but this is certainly possible for many parts of Algorithm 2 and can bring significant execution speed-ups in applications. Finally, future work could aim to extend this methodology by designing suitable IS proposals for non-linear observation schemes or observations obtained from Lagrangian drifters or floaters.

Acknowledgements

FPLI was supported by EPSRC and the CDT in the Mathematics of Planet Earth under grant EP/L016613/1. AJ was supported by an AcRF tier 2 grant: R-155-000-161-112. AJ is affiliated with the Risk Management Institute, the Center for Quantitative Finance and the OR & Analytics cluster at NUS. AB was supported by the Leverhulme Trust Prize.

A More simulation results

We first present some numerical results when tempering is not used. We will consider a perfect initialization for each particle with v_0^\dagger . Whilst this is an extremely favorable scenario that is unrealistic

Algorithm 3 An MCMC Procedure for $\bar{X}_{n,l}^i \sim \mathcal{K}_{n,l}^m(\cdot | X_{n,l}^i)$.

- Initialise: set $V^{(0)} = X_{n,l}^i$ and let $W^{(0)} = W_{n,l}^i$ be the Wiener process generating $X_{n,l}^i$.
- For $k = 1, \dots, m$: let $V = V^{(k-1)}$, $W = W^{(k-1)}$.

– Sample a new noise

$$W(s)' = \rho W(s) + \sqrt{1 - \rho^2} \xi(s), \quad s \in (t_{n-1}, t_n], \quad \xi \sim \mathbb{W}.$$

– Obtain solution of SPDE (16) with W' the driving noise, i.e.

$$dV'(s) = (-\nu AV'(s) - B(V'(s), V'(s))) dt + Q^{\frac{1}{2}} g(s, V'(s)) ds + Q^{\frac{1}{2}} dW'(s), \\ t \in (t_{n-1}, t_n].$$

– Compute acceptance ratio

$$\alpha = 1 \wedge \frac{\left(\frac{d\mathbb{V}^n}{d\mathbb{Q}^n}(V') p(Y_{t_n} | V'(t_n)) \right)^{\phi_{n,t}}}{\left(\frac{d\mathbb{V}^n}{d\mathbb{Q}^n}(V) p(Y_{t_n} | V(t_n)) \right)^{\phi_{n,t}}}.$$

– With probability α set $V^{(k)} = V'$, $W^{(k)} = W'$; otherwise reject proposal, set $V^{(k)} = V$, $W^{(k)} = W$.

- Return $\bar{X}_{n,l}^i = V^{(k)}$ and $\bar{W}_{n,l}^i = W^{(k)}$.
-

in practice, it shows a clear benefit in using IS and tempering. For $N = 200$, $\nu = 0.01$ and $\delta t_n = 0.2$, we present some scatter plots in Figure 5 for the experiment with $n = 5$ seen earlier with a 16×16 block of observations. Notably, the estimated posterior means for the vorticity seem to exhibit good performance; see Figure 6. Indicatively, the ESS here is 34 for IS and 3 for the bootstrap case. Even in this extremely favorable scenario, the ESS is low and this strongly motivates the use of tempering to improve the efficiency of the particle methodology. In results not shown here, we also experimented the size of time increment $\delta t_n = t_n - t_{n-1}$ a naive particle filter (Algorithm 1) can handle with perfect initialization. When the likelihood-informed proposals in (16) are used, the method produces accurate results for δt_n up to 0.2 – 0.25. This is in contrast to when sampling from the dynamics, where the bootstrap version of Algorithm 1 can handle only up to 0.15.

Finally, in Figure 7 we present some estimated PDFs from the long experiment of Section 5 with $n = 100$.

References

- [1] A. BAIN AND D. CRISAN, *Fundamentals of stochastic filtering*, Springer, 2008.
- [2] A. F. BENNETT, *Inverse modeling of the ocean and atmosphere*, Cambridge University Press, 2005.
- [3] A. BESKOS, D. CRISAN, A. JASRA, K. KAMATANI, AND Y. ZHOU, *A stable particle filter for a class of high-dimensional state-space models*, *Advances in Applied Probability*, 49 (2017), pp. 24–48.
- [4] A. BESKOS, A. JASRA, N. KANTAS, AND A. THIERY, *On the convergence of adaptive sequential Monte Carlo methods*, *The Annals of Applied Probability*, 26 (2016), pp. 1111–1146.

- [5] M. BOCQUET, C. A. PIRES, AND L. WU, *Beyond Gaussian statistical modeling in geophysical data assimilation*, Monthly Weather Review, 138 (2010), pp. 2997–3023.
- [6] A. J. CHORIN AND P. KRAUSE, *Dimensional reduction for a Bayesian filter*, Proceedings of the National Academy of Sciences of the United States of America, 101 (2004), pp. 15013–15017.
- [7] A. J. CHORIN, M. MORZFELD, AND X. TU, *Implicit particle filters for data assimilation*, Communications in Applied Mathematics and Computational Science, 5 (2010), pp. 221–240.
- [8] S. L. COTTER, G. O. ROBERTS, A. M. STUART, AND D. WHITE, *MCMC methods for functions: modifying old algorithms to make them faster*, Statistical Science, 28 (2013), pp. 424–446.
- [9] G. DA PRATO AND J. ZABCZYK, *Stochastic equations in infinite dimensions*, Cambridge University Press, 2008.
- [10] P. DEL MORAL, *Feynman-Kac Formulae*, Springer, 2004.
- [11] P. DEL MORAL, *Mean field simulation for Monte Carlo integration*, CRC Press, 2013.
- [12] A. DOUCET, M. BRIERS, AND S. SÉNÉCAL, *Efficient block sampling strategies for sequential Monte Carlo methods*, Journal of Computational and Graphical Statistics, 15 (2006), pp. 693–711.
- [13] A. DOUCET, N. DE FREITAS, AND N. GORDON, *Sequential Monte Carlo methods in practice*, Springer Science & Business Media, 2001.
- [14] B. FERRARIO, *Stochastic Navier-Stokes equations: Analysis of the noise to have a unique invariant measure*, Annali di Matematica Pura ed Applicata, 177 (1999), pp. 331–347.
- [15] F. FLANDOLI, *Dissipativity and invariant measures for stochastic Navier-Stokes equations*, Nonlinear Differential Equations and Applications NoDEA, 1 (1994), pp. 403–423.
- [16] C. FOIAS, O. MANLEY, R. ROSA, AND R. TEMAM, *Navier-Stokes equations and turbulence*, vol. 83, Cambridge University Press, 2001.
- [17] F. GIRAUD AND P. DEL MORAL, *Nonasymptotic analysis of adaptive and annealed Feynman-Kac particle models*, Bernoulli, 23 (2017), pp. 670–709.
- [18] A. GOLIGHTLY AND D. J. WILKINSON, *Bayesian inference for nonlinear multivariate diffusion models observed with error*, Computational Statistics & Data Analysis, 52 (2008), pp. 1674–1693.
- [19] V. H. HOANG, K. J. LAW, AND A. M. STUART, *Determining white noise forcing from Eulerian observations in the Navier-Stokes equation*, Stochastic Partial Differential Equations: Analysis and Computations, 2 (2014), pp. 233–261.
- [20] M. JARDAK, I. NAVON, AND M. ZUPANSKI, *Comparison of sequential data assimilation methods for the Kuramoto-Sivashinsky equation*, International journal for numerical methods in fluids, 62 (2010), pp. 374–402.
- [21] A. JASRA, D. A. STEPHENS, A. DOUCET, AND T. TSAGARIS, *Inference for Lévy-driven stochastic volatility models via adaptive sequential Monte Carlo*, Scandinavian Journal of Statistics, 38 (2011), pp. 1–22.
- [22] A. JENTZEN AND P. E. KLOEDEN, *Overcoming the order barrier in the numerical approximation of stochastic partial differential equations with additive space-time noise*, in Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences, vol. 465, 2009, pp. 649–667.

- [23] A. M. JOHANSEN, *On block, tempering and particle mcmc for systems identification*, in Proceedings of 17th IFAC Symposium on System Identification, IFAC, 1998.
- [24] A. M. JOHANSEN, N. WHITELEY, AND A. DOUCET, *Exact approximation of Rao–Blackwellised particle filters*, IFAC Proceedings Volumes, 45 (2012), pp. 488–493.
- [25] N. KANTAS, A. BESKOS, AND A. JASRA, *Sequential Monte Carlo methods for high-dimensional inverse problems: A case study for the Navier–Stokes equations*, SIAM/ASA Journal on Uncertainty Quantification, 2 (2014), pp. 464–489.
- [26] S. KUKSIN AND A. SHIRIKYAN, *Mathematics of two-dimensional turbulence*, vol. 194, Cambridge University Press, 2012.
- [27] K. LAW, A. STUART, AND K. ZYGALAKIS, *Data assimilation: a mathematical introduction*, vol. 62, Springer, 2015.
- [28] A. LEE, C. YAU, M. B. GILES, A. DOUCET, AND C. C. HOLMES, *On the utility of graphics cards to perform massively parallel simulation of advanced Monte Carlo methods*, Journal of Computational and Graphical Statistics, 19 (2010), pp. 769–789.
- [29] A. J. MAJDA AND J. HARLIM, *Filtering complex turbulent systems*, Cambridge University Press, 2012.
- [30] N. PAPADAKIS, É. MÉMIN, A. CUZOL, AND N. GENGEMBRE, *Data assimilation with the weighted ensemble Kalman filter*, Tellus A, 62 (2010), pp. 673–697.
- [31] O. PAPASPILIOPOULOS AND G. ROBERTS, *Importance sampling techniques for estimation of diffusion models*, Statistical methods for stochastic differential equations, 124 (2012), pp. 311–340.
- [32] O. PAPASPILIOPOULOS, G. O. ROBERTS, AND O. STRAMER, *Data augmentation for diffusions*, Journal of Computational and Graphical Statistics, 22 (2013), pp. 665–688.
- [33] S. SÄRKKÄ AND E. MOULINES, *On the L_p -convergence of a Girsanov theorem based particle filter*, in Acoustics, Speech and Signal Processing (ICASSP), 2016 IEEE International Conference on, IEEE, 2016, pp. 3989–3993.
- [34] S. SÄRKKÄ, T. SOTTINEN, ET AL., *Application of Girsanov theorem to particle filtering of discretely observed continuous-time non-linear systems*, Bayesian Analysis, 3 (2008), pp. 555–584.
- [35] C. SNYDER, T. BENGTTSSON, P. BICKEL, AND J. ANDERSON, *Obstacles to high-dimensional particle filtering*, Monthly Weather Review, 136 (2008), pp. 4629–4640.
- [36] A. M. STUART, *Inverse problems: a Bayesian perspective*, Acta Numerica, 19 (2010), pp. 451–559.
- [37] L. TIERNEY, *A note on Metropolis-Hastings kernels for general state spaces*, Annals of Applied Probability, (1998), pp. 1–9.
- [38] F. VAN DER MEULEN AND M. SCHAUER, *Bayesian estimation of incompletely observed diffusions*, arXiv preprint arXiv:1606.04082, (2016).
- [39] P. J. VAN LEEUWEN, *Nonlinear data assimilation in geosciences: an extremely efficient particle filter*, Quarterly Journal of the Royal Meteorological Society, 136 (2010), pp. 1991–1999.
- [40] P. J. VAN LEEUWEN, Y. CHENG, AND S. REICH, *Nonlinear Data Assimilation*, Springer, 2015.
- [41] J. WEARE, *Particle filtering with path sampling and an application to a bimodal ocean current model*, Journal of Computational Physics, 228 (2009), pp. 4312–4331.
- [42] G. A. WHITAKER, A. GOLIGHTLY, R. J. BOYS, AND C. SHERLOCK, *Improved bridge constructs for stochastic differential equations*, Statistics and Computing, (2016), pp. 1–16.

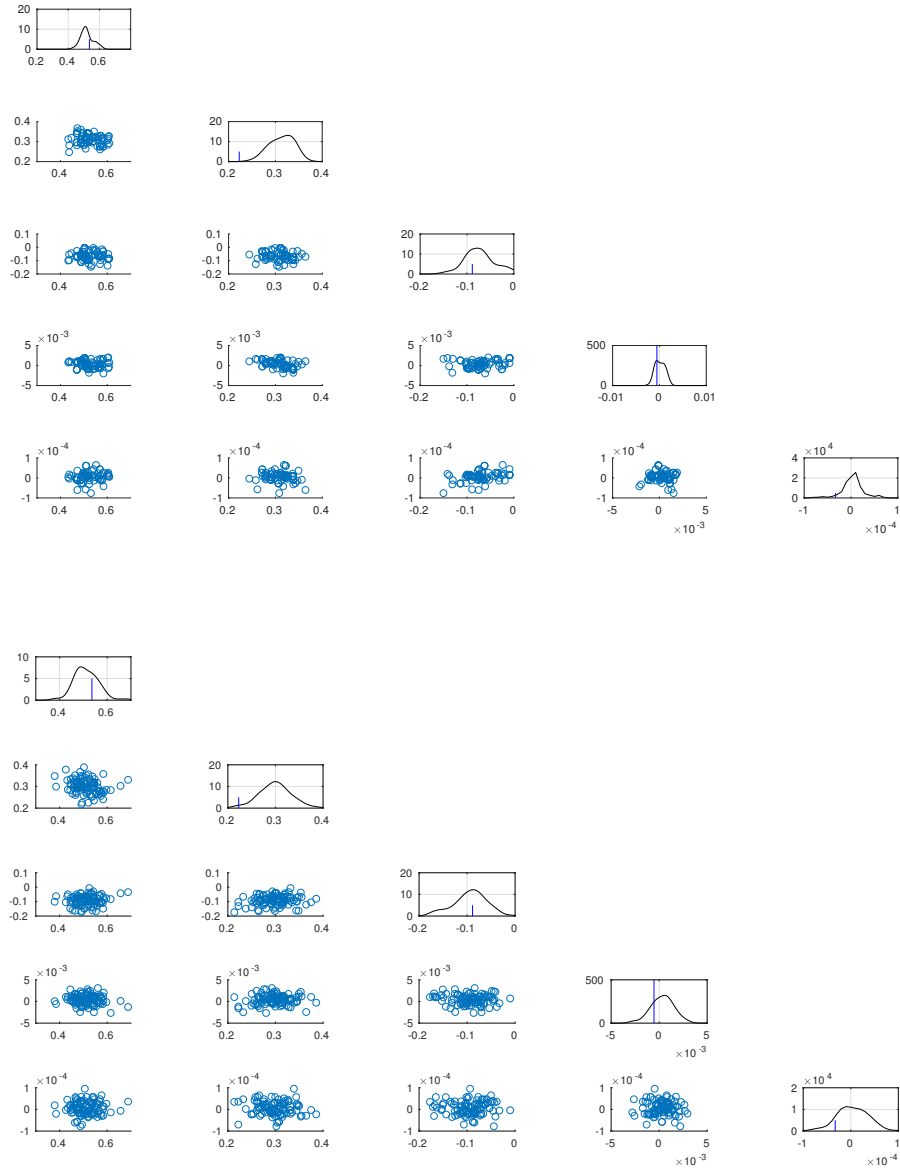


Figure 1: PDF and scatter plots for $\text{Real}(u_k)|\mathcal{Y}_n$ at $n = 5$ for $k = (1, 0), (1, 1), (1, -1), (2, 5), (9, 9)$. Top is bootstrap (sampling with (9)) and bottom is IS (with (16)) and both use tempering. Vertical lines in PDF plots are true signal values used to generate the observations.




Figure 2: Vorticity plots showing posterior mean of $p(\nabla \times V_t | \mathcal{Y}_n)$ and true signal: top row $n = 10$, middle $n = 50$, bottom $n = 100$. The left column contains posterior means from Gaussian observation noise, the right one from Student- t noise and in the middle is w_t^\dagger (true signal vorticity).




Figure 3: Variance plots: top row $n = 10$, middle $n = 50$, bottom $n = 100$. We present heat maps of the ratio of the posterior variance of π_{t_n} over the variance for the law of the signal dynamics against k ; left part is for Gaussian noise and right for Student-t.




Figure 4: Results for single run of PF with tempering and IS for 8×8 grid. Top panels are L^2 -errors and number of tempering steps against n . Dotted lines are for Gaussian observation noise and solid for Student-t. In the bottom panels we present ESS against SMC iteration for Gaussian (left) and Student-t (right) errors. Execution times were 4.8018×10^5 and 4.17196×10^5 seconds respectively.

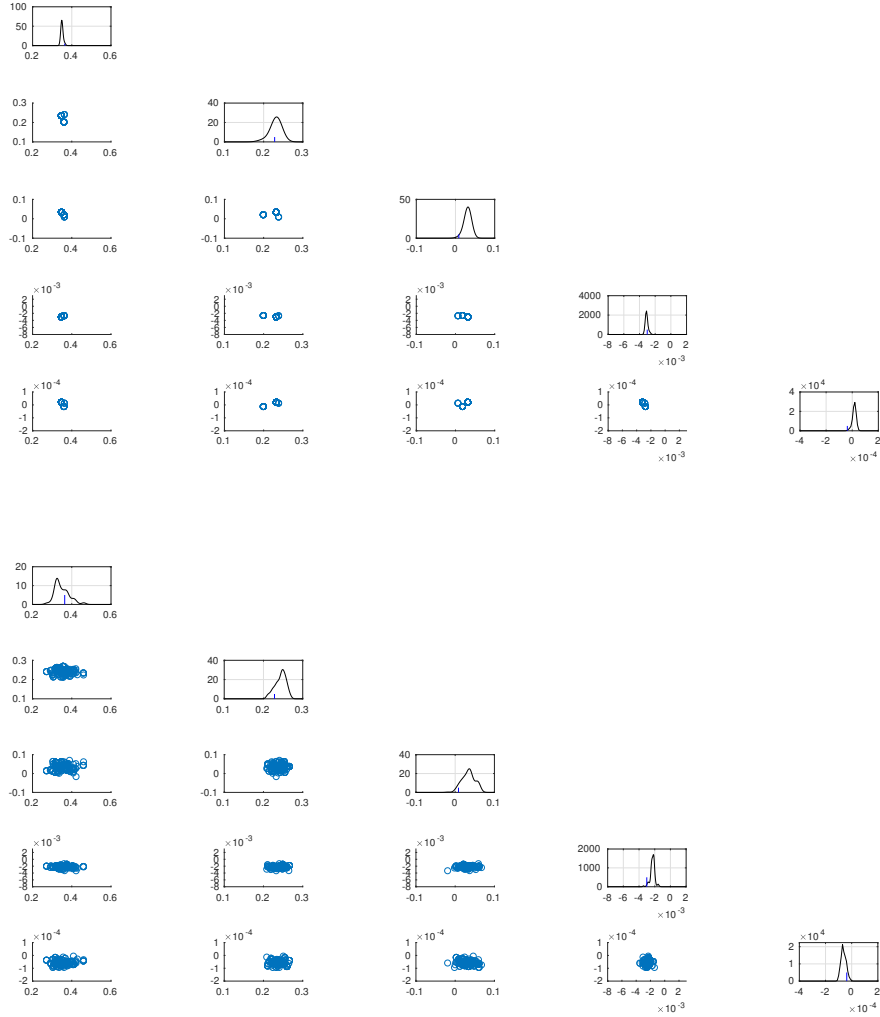


Figure 5: Scatter plots at $n = 5$ for perfect initialization $k = (1, 0), (1, 1), (1, -1), (2, 5), (9, 9)$. Top is bootstrap and bottom is IS with (16).




Figure 6: Vorticity plots for $n = 5$ and perfect initialization: left is posterior mean from bootstrap PF, middle is real signal v_t^\dagger , and right is posterior mean of PF of Algorithm 1 and IS with (16).

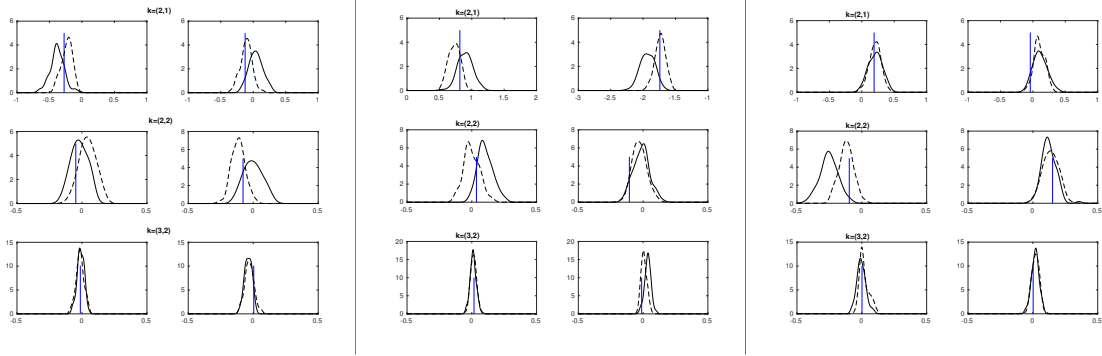


Figure 7: PDFs of vorticity for $n = 10, 50, 100$ from left to right. In each panel for n , left side displays real part and right is imaginary; top row is $k = (2, 1)$, middle $k = (2, 2)$ and bottom is $k = (3, 2)$. Dotted line is for Gaussian observation noise and solid for Student-t, vertical lines are true signal values used to generate the observations.