Computational Science Laboratory Report CSL-TR-24-1 May 8, 2024

Amit N. Subrahmanya, Andrey A. Popov, Reid J. Gomillion, Adrian Sandu

"Preserving Nonlinear Constraints in Variational Flow Filtering Data Assimilation"

> Computational Science Laboratory "Compute the Future!"

Department of Computer Science Virginia Tech Blacksburg, VA 24060 Phone: (540) 231-2193 Fax: (540) 231-6075 Email: amitns@vt.edu, apopov@vt.edu, rjg18@vt.edu, sandu@cs.vt.edu Web: https://csl.cs.vt.edu





Preserving Nonlinear Constraints in Variational

Flow Filtering Data Assimilation

Amit N. Subrahmanya^{1*}, Andrey A. Popov², Reid J. Gomillion¹, Adrian Sandu¹

^{1*}Computational Science Laboratory, Department of Computer Science, Virginia Tech, 620 Drillfield Dr., Blacksburg, 24061, Virginia, USA.

²Oden Institute for Computational Engineering and Sciences, University of Texas at Austin, 201 E. 24th Street, Austin, 78712, Texas, USA.

*Corresponding author(s). E-mail(s): amitns@vt.edu; Contributing authors: apopov@vt.edu; rjg18@vt.edu; sandu@cs.vt.edu;

Abstract

Data assimilation aims to estimate the states of a dynamical system by optimally combining sparse and noisy observations of the physical system with uncertain forecasts produced by a computational model. The states of many dynamical systems of interest obey nonlinear physical constraints, and the corresponding dynamics is confined to a certain sub-manifold of the state space. Standard data assimilation techniques applied to such systems yield posterior states lying outside the manifold, violating the physical constraints.

This work focuses on particle flow filters which use stochastic differential equations to evolve state samples from a prior distribution to samples from an observation-informed posterior distribution. The variational Fokker-Planck (VFP)—a generic particle flow filtering framework—is extended to incorporate non-linear, equality state constraints in the analysis. To this end, two algorithmic approaches that modify the VFP stochastic differential equation are discussed: (i) VFPSTAB, to inexactly preserve constraints with the addition of a stabilizing drift term, and (ii) VFPDAE, to exactly preserve constraints by treating the VFP dynamics as a stochastic differential-algebraic equation (SDAE). Additionally, an implicit-explicit time integrator is developed to evolve the VFPDAE dynamics. The strength of the proposed approach for constraint preservation in data assimilation is demonstrated on three test problems: the double pendulum, Korteweg-de-Vries, and the incompressible Navier-Stokes equations.

Keywords: Bayesian inference, data assimilation, particle flow filters, nonlinear constraint preservation

MSC Classification: 65C05, 62F15, 62F30, 35R30

1 Introduction

Data assimilation for state estimation attempts to combine uncertain, highdimensional model simulations of a system with noisy (and typically low-dimensional) observations of the said system in a theoretically rigorous manner [1–3]. Data assimilation is widely used to improve forecasts in the areas of weather prediction, hydrology, seismology, and so on. Popular data assimilation methods include the well-known Kalman filter [4], ensemble Kalman filter [5] (and its many variants such as the ensemble transform Kalman filter (ETKF) [6], ensemble adjustment Kalman filter (EAKF) [7], and the local ensemble transform Kalman filter(LETKF) [8]), variational methods such as 3D-Var and 4D-Var [1, 3] like, and particle filters [9, 10].

Particle flow filters solve the assimilation problem by iteratively evolving an ensemble of particles sampling the prior distribution into an ensemble of particles that sample the posterior through stochastic differential equations (SDE). While there are many flavors of particle flow filters [11–16], the main focus is on the authors' Variational Fokker-Planck (VFP) approach [17], a flexible and generalized particle flow framework unifying many well-known particle flow methods. From the perspective of VFP, the particle flow in the state space is driven by a McKean-Vlasov-Itô process whose drift minimizes the Kullback-Leibler divergence between the particle distribution and the posterior, and whose diffusion helps avoid particle collapse. Multiple assumptions can be made on the prior and intermediate probability densities when computing the drift, and at the same time, can be combined with localization techniques to mitigate the curse of dimensionality [18].

Many physical systems of interest obey constraints such as conservation laws for mass, momentum, energy, enstrophy, etc [19, 20]. In this work, the general term *constraint* refers to any preserved fundamental property of a system. This work considers non-linear equality constraints. In general, while the computational models are built to preserve (in a numerical sense) the physical constraints of the system, standard data assimilation algorithms are not. Traditional particle filters, which modify the weights associated with the particles (and not the particle states), preserve nonlinear constraints but are infeasible for high-dimensional problems [10, 18]. On the other hand, ensemble Kalman-like algorithms can conserve linear equality constraints [21], but not non-linear, or inequality constraints. Like most standard filtering methods, the VFP framework [17] makes no guarantees about constraint preservation.

Modifying filtering algorithms to achieve constraint preservation is not a new idea in data assimilation. Early work on preserving positivity constraints was done by Massicotte et. al [22], where they *inexactly* preserve positivity on filtered spectrometric data by heuristically damping the non-negativity of the state by multiplication with a deflation factor $\in [0, 1]$. Simon and Chia [23] discuss Kalman filtering with state equality constraints. They describe a projection-based formulation for linear equality constraints on the states and extend it to non-linear equality constraints by linearization. They also present methods to restrict the Kalman gain such that the solution lies on the constraint manifold. However, they conclude that the method of linearization may result in convergence issues when projecting onto the constraint manifold.

Julier and LaViola Jr. [24] review and analyze two families of non-linear equalityconstrained Kalman Filters, namely the pseudo-observation and projection methods. They identify two types of constraints, one that acts on the entire distribution (Type 1) and the other that acts only on the mean of the distribution (Type 2). They conclude that the computationally simpler pseudo-observation method (which includes constraints as observations with heuristically chosen error covariances) preserves neither

constraint; while the more computationally expensive projection method can preserve both types of constraints. Gupta and Hauser [25] introduce methods for inequality and equality linear constraint by the projection method or by modifying the Kalman gain. Sircoulomb et. al. [26] present a generic framework for dealing with non-linear inequality constraints by iterative projection. Zanetti et. al.[27] modify the Kalman gain operator for quadratic constraints, an approach useful when the norm constraint considers a small subset of the state space. Prakash et. al. [28], propose a method to incorporate box constraints by modifying the ensemble Kalman filter to work on a truncated Gaussian distribution through a constrained optimization. Bavdekar et. al. [29] extend the previous idea to dual state and parameter estimation using the ensemble Kalman filter. Their methodology—to separate the state and (constrained) parameter estimation—resulted in a lower error and a lower variance for the parameter estimate at increased cost (coming from additional model evolutions) when compared with a joint (state and parameter) EnKF approach.

Janjic et. al. [21] deal with positivity preserving constraints in the ensemble Kalman filter by formulating the update as a sequence of quadratic programming problems. Additionally, they show that the ensemble Kalman filters always conserve linear equality constraints assuming a preconsistent forecast. However, they also state that localization prevents this conservation. Zeng and Janjic [20] analyze the effect of localization radii, observed variables, and observation sparsity w.r.t the discretized grid on the conservation of mass, energy, enstrophy, divergence, and noise for the two-dimensional shallow water equations. They observed that a small localization radius or observing only the heights degraded the conservation properties. They also observed that the enstrophy violation, which should be minimal for stable nonlinearities [19, 30], is highly correlated to the localization radius and surprisingly, the ensemble inflation to prevent filter divergence. Zeng, Janjic, et. al. [31] extend the quadratic penalty idea [21] to LETKF.

Li et. al. [32] impose inequality constraints by projecting the unconstrained posterior Gaussian analysis particles onto the constrained region such that the constrained distribution is close in KL divergence to the unconstrained distribution. Albers et. al [33] discuss imposing linear equality and inequality constraints into the ensemble Kalman filter, by constraining the underlying EnKF quadratic optimization problem when there is a constraint violation. They also extend it to the ensemble Kalman Inversion methodology [34] to solve generic inverse problems. Chada et. al. [35] incorporate box constraints into the Ensemble Kalman Inversion (EKI) by using a projected gradient method [36], into the gradient flow structure of the continuous limit EKI. The ensemble Kalman inversion idea relates somewhat to VFP as both methods involve a gradient flow to move samples from a prior distribution to a posterior distribution. Amor et. al [37] present a comprehensive overview of constraint-preserving techniques for ensemble Kalman filters and particle filters.

The new contributions of this work are as follows.

- 1. The incorporation of constraint preservation methodology into the generic VFP framework [17] to obtain physically consistent (w.r.t the constraint) states.
- The first method—VFPSTAB—stabilizes the particle flow drift with an additional term to ensure the states are "close enough" to the constraint manifold.
 VFPSTAB is computationally efficient but inexact for constraint preservation.
- 3. The second method—**VFPDAE**—adds a constraint preservation term along with an algebraic constraint to the flow to ensure the states "exactly" respect the constraint manifold at all times during the filtering process. Both these methods lie in stark contrast with historical methods where the assimilation and constraint preservation steps are treated independently, potentially leading to situations where the posterior and the constraint cannot be reconciled.
- 4. Finally, an implicit-explicit time stepping scheme is developed to evolve VFPDAE flow in pseudo-time efficiently.

The remainder of the paper is organized as follows. Section 2 describes the data assimilation problem, focusing on filtering, and then presents the notation used in this work. Section 3 examines the role of constraints in data assimilation along with a description of the traditional projection and pseudo-observation constraint preservation methods. Section 4 introduces the variational Fokker-Planck method, and develops two constraint preserving extensions, specifically, the VFPSTAB in Section 4.1 and the VFPDAE in Section 4.2. Next, Section 5 describes techniques for regularized covariance estimation in Section 5.1 and then compares VFPDAE and VFPSTAB to modified versions of the ETKF, and LETKF using the double pendulum (Section 5.2), Korteweg-de Vries (Section 5.3) and incompressible Navier-Stokes (Section 5.4) problems. Finally, the conclusions are presented in section 6.

2 Data assimilation

A discrete-time dynamical system model is described by an operator that propagates the states:

$$\mathbf{x}_k = \mathcal{M}_k(\mathbf{x}_{k-1}), \quad k \ge 1, \tag{1}$$

where k indexes the time t_k , $\mathbf{x}_{k-1} \in \mathbb{R}^{n_s}$ represents the states of the system at time t_{k-1} , and the model operator $\mathcal{M}_k : \mathbb{R}^{n_s} \to \mathbb{R}^{n_s}$ evolves the states from time t_{k-1} to time t_k . This paper considers deterministic models \mathcal{M} described by (partial) differential-algebraic equations [38].

The model states must be combined with observations of the unknown reality to track the true states. These observations of an unknown true state are typically sparse (in space and time) and noisy, and given by

$$\mathbf{y}_k = \mathcal{H}_k(\mathbf{x}_{\text{true},k}) + \eta_k^{\text{o}},\tag{2}$$

where $\mathbf{y}_k \in \mathbb{R}^{n_o}$ is the observation at time t_k , $\mathcal{H}_k : \mathbb{R}^{n_s} \to \mathbb{R}^{n_o}$ the observation operator and η_k^o is the observation noise. In a typical data assimilation setting, $n_o \ll n_s$ as it is infeasible to have a larger set of observations.

Filtering alternates between forecasting, where previous analysis states are propagated to a future time using the computational model; and analysis, which combines this forecast with the current observation to obtain new analysis states. This cycle is summarized as

$$\begin{aligned} \mathbf{x}_{\mathrm{f},k} &= \mathcal{M}_k(\mathbf{x}_{\mathrm{a},k-1}), \qquad \text{(Forecast)} \\ \mathbf{x}_{\mathrm{a},k} &= \mathcal{A}(\mathbf{x}_{\mathrm{f},k},\mathbf{y}_k;\boldsymbol{\theta}), \qquad \text{(Analysis)} \end{aligned} \tag{3}$$

where $\boldsymbol{\theta} \in \mathbb{R}^p$ is a p-dimensional set of parameters of the analysis procedure. For example, in the localized ensemble Kalman Filter, the parameters $\boldsymbol{\theta}$ could include ensemble inflation and localization radius. Through the Bayes' rule, we have

$$\mathcal{P}(\mathbf{x}_{\mathrm{a},k}) = \mathcal{P}(\mathbf{x}_{\mathrm{f},k} \mid \mathbf{y}_k) = \frac{\mathcal{P}(\mathbf{y}_k \mid \mathbf{x}_{\mathrm{f},k}) \cdot \mathcal{P}(\mathbf{x}_{\mathrm{f},k})}{\mathcal{P}(\mathbf{y}_k)},\tag{4}$$

where the analysis distribution is the (Bayesian) posterior that combines the observational likelihood $\mathcal{P}(\mathbf{y}_k \mid \mathbf{x}_{\mathrm{f},k})$ and the prior distribution of the forecast $\mathcal{P}(\mathbf{x}_{\mathrm{f},k})$. Equivalently, we refer to the probabilities as $\mathcal{P}(\mathbf{x}_{\mathrm{f},k}) \equiv \mathcal{P}_{\mathrm{f}}(\mathbf{x}_k), \ \mathcal{P}(\mathbf{y}_k \mid \mathbf{x}_{\mathrm{f},k}) \equiv \mathcal{P}_{\mathrm{o}}(\mathbf{y}_k - \mathcal{H}_k(\mathbf{x}_k)), \text{ and } \mathcal{P}(\mathbf{x}_{\mathrm{a},k}) \equiv \mathcal{P}_{\mathrm{a}}(\mathbf{x}_k).$

This paper exclusively considers ensemble approaches to filtering, i.e.

$$\begin{aligned} \mathbf{X}_{\mathrm{f},k} &= \mathcal{M}_k(\mathbf{X}_{\mathrm{a},k-1}), \qquad \text{(Forecast)} \\ \mathbf{X}_{\mathrm{a},k} &= \mathcal{A}(\mathbf{X}_{\mathrm{f},k},\mathbf{y}_k;\boldsymbol{\theta}), \qquad \text{(Analysis)} \end{aligned} \tag{5}$$

where the ensemble

$$\mathbf{X}_{*,k} = \begin{bmatrix} \mathbf{x}_{*,k}^{[1]} & \mathbf{x}_{*,k}^{[2]} & \dots & \mathbf{x}_{*,k}^{[n_e]} \end{bmatrix} \in \mathbb{R}^{n_s \times n_e},$$
(6)

has $n_{\rm e}$ exchangeable random variables, representing samples (not necessarily independent) from the same distribution. Here, it is assumed that $\mathbf{x}_{{\rm f},k}^{[e]} \sim \mathcal{P}_{\rm f}(\mathbf{x}_k)$ and $\mathbf{x}_{{\rm a},k}^{[e]} \sim \mathcal{P}_{\rm a}(\mathbf{x}_k)$ for all e in the index set $\mathcal{I}_{n_{\rm e}} = \{1, 2, \ldots, n_{\rm e}\}$.

3 Constraints

Many dynamical systems of interest, possess natural constraints on their states [21, 23, 25] as in eq. (7). For example, (i) geophysical systems obey conservation laws such as mass, momentum, energy, and so on; (ii) mechanical systems are constrained by lengths or velocities of its components. These constraints, which describe physical or geometric laws, are ignored while filtering as they can be challenging to incorporate into the filtering methodology [23].

For spatially distributed dynamical systems (typically described by PDEs), constraints are classified as global and local [39]. Global constraints preserve information across the entire domain of the problem. For example, the conservation of total energy of a closed system, the conservation of enstrophy for a closed fluid simulation, the length of a pendulum, etc. Local constraints preserve information in a pointwise fashion (i.e. at spatial locations) across the domain. For example, this can be the continuity equation for an incompressible fluid where the divergence is pointwise zero at every spatial location across the domain; or that of non-negativity where the concentration of a chemical is greater than or equal to zero at each point across the domain. Constraints can also be divided into two types based on their origin.

1. Constraints defined by state invariants that do not change with either time or across samples. For models such as eq. (1), the constraints are described by the algebraic relations:

$$\mathbf{g}(\mathbf{x}_{k-1}) = \mathbf{0} \Rightarrow \mathbf{g}(\mathbf{x}_k) = \mathbf{g}\left(\mathcal{M}_k(\mathbf{x}_{k-1})\right) = \mathbf{0},\tag{7}$$

where $\mathbf{g} : \mathbb{R}^{n_s} \to \mathbb{R}^{n_c}$ is the constraint function. If \mathbf{x}_{k-1} were to be modified as $\tilde{\mathbf{x}}_{k-1}$ such that $\mathbf{g}(\tilde{\mathbf{x}}_{k-1}) \neq \mathbf{0}$, the following situations could occur with $\tilde{\mathbf{x}}_k = \mathcal{M}_k(\tilde{\mathbf{x}}_{k-1})$.

- (a) Some model dynamics naturally pull back the states onto the constraint manifold, i.e. $\mathbf{g}(\tilde{\mathbf{x}}_k) = \mathbf{0}$. The states need no special treatment in this scenario and such models are not considered in this work.
- (b) Some models permit a small degree of flexibility, i.e.

$$|\mathbf{g}(\tilde{\mathbf{x}}_{k-1})| \le \varepsilon \quad \Rightarrow \quad |\mathbf{g}(\tilde{\mathbf{x}}_k)| \le \varepsilon, \tag{8}$$

where the constraints are *weakly* preserved.

(c) Sometimes $\tilde{\mathbf{x}}_k$ is not computable if $\mathbf{g}(\tilde{\mathbf{x}}_{k-1}) \neq \mathbf{0}$ or even if $\tilde{\mathbf{x}}_k$ is computable, $\mathbf{g}(\tilde{\mathbf{x}}_k) \neq \mathbf{0}$. For these situations, $\tilde{\mathbf{x}}_{k-1}$ must be modified to lie on the constraint manifold and then evolved.

The following equations are all at time t_k (subscript k has been dropped in the equations for simplicity), the data assimilation problem assumes a preconsistent forecast i.e.

$$\forall e \in \mathcal{I}_{n_e} \quad \mathbf{g}(\mathbf{x}_{\mathrm{f}}^{[e]}) = \mathbf{0}, \quad \text{or} \quad |\mathbf{g}(\mathbf{x}_{\mathrm{f}}^{[e]})| \le \varepsilon.$$
(9)

The goal is to have analysis ensemble members who also live on the constraint manifold as

$$\forall e \in \mathcal{I}_{n_{e}} \quad \mathbf{g}(\mathbf{x}_{a}^{[e]}) = \mathbf{0}, \quad \text{or} \quad |\mathbf{g}(\mathbf{x}_{a}^{[e]})| \le \varepsilon.$$
(10)

Remark 1 The constraints can also vary across time, or differ across the ensemble. We do not delve into these cases, because the underlying ideas regarding these constraints and their preservation are mostly similar to the above.

2. Constraints that preserve information between the forecast and analysis samples. Essentially, the constraints (again, all at time t_k) are defined as

$$\forall e \in \mathcal{I}_{n_{e}} \quad \mathbf{g}^{[e]}(\mathbf{x}_{a}^{[e]}; \mathbf{x}_{f}^{[e]}) = \mathbf{h}(\mathbf{x}_{a}^{[e]}) - \mathbf{h}(\mathbf{x}_{f}^{[e]}), \tag{11}$$

where $\mathbf{h} : \mathbb{R}^{n_s} \to \mathbb{R}^{n_c}$ represents some preservable quantities. The constraint preservation requires

$$\forall e \in \mathcal{I}_{n_{e}} \quad \mathbf{g}^{[e]}(\mathbf{x}_{a}^{[e]}; \mathbf{x}_{f}^{[e]}) = \mathbf{0}, \quad \text{or} \quad |\mathbf{g}^{[e]}(\mathbf{x}_{a}^{[e]}; \mathbf{x}_{f}^{[e]})| \le \varepsilon.$$
(12)

Typically, these constraints are artificial (rather than natural) and enforced based on need.

Problems with both constraint types can be formulated, such as the incompressible Navier Stokes experiment in section 5.4. With a few exceptions, most filtering methods ignore state constraints, and as a consequence, fail to preserve them [20, 23].

Example 1 (Pendulum) Consider an example problem with the simple pendulum to motivate constraint preservation further. Consider the situation where the forecast pendulum ensemble, the truth, the observation, and the constraint manifold are as depicted in Figure 1. The constraint manifold here is defined by the geometric constraints of the pendulum such as its length and the extreme position (dictating the total mechanical energy of the system in the absence of forcing or dissipation). While the ETKF analysis (labeled Standard ETKF in Figure 1) ensemble is closer to the observations, it is completely off the constraint manifold. These inconsistent analysis states are meaningless as the pendulum's length must remain constant after filtering, and hence, a better and physically consistent analysis is preferred (labeled Constrained ETKF in Figure 1).

We now discuss two different approaches to constrain analyses.



Fig. 1: A simple pendulum example where the forecast ensemble members lie on the constraint manifold, but the standard ETKF analysis does not. The goal is to obtain an analysis ensemble that lives on the constraint manifold like the Constrained ETKF.

3.1 Projection onto constraint manifold after filtering

The projection after filtering approach has two main steps. First, a constraint agnostic filtering scheme is applied to the forecast to obtain the unconstrained analysis:

$$\widehat{\mathbf{x}}_{a}^{[e]} = \mathcal{A}(\mathbf{x}_{f}^{[e]}, \mathbf{y}, \boldsymbol{\theta}), \quad \forall e \in \mathcal{I}_{n_{e}}.$$
(13)

Next, each analysis ensemble member is projected onto the manifold using:

$$\mathbf{x}_{\mathbf{a}}^{[e]} = \widehat{\mathbf{x}}_{\mathbf{a}}^{[e]} - \mathbf{G}^{\top}(\widehat{\mathbf{x}}_{\mathbf{a}}^{[e]}) \boldsymbol{\lambda}^{[e]}, \quad \forall e \in \mathcal{I}_{n_e},$$
(14a)

where $\boldsymbol{\lambda}^{[e]} \in \mathbb{R}^{n_{\mathrm{c}}}$ is a solution of the nonlinear equation:

$$\mathbf{0} = \mathbf{g}\left(\mathbf{x}_{\mathrm{a}}^{[e]}\right) = \mathbf{g}\left(\widehat{\mathbf{x}}_{\mathrm{a}}^{[e]} - \mathbf{G}^{\top}(\widehat{\mathbf{x}}_{\mathrm{a}}^{[e]}) \boldsymbol{\lambda}^{[e]}\right), \quad \forall e \in \mathcal{I}_{n_{\mathrm{e}}},$$
(14b)

where the Jacobian of the constraints defines the tangent constraint manifold:

$$\mathbf{G}(\mathbf{x}) \coloneqq \frac{d\mathbf{g}(\mathbf{x})}{d\mathbf{x}}.$$
(15)

The approach eq. (14) is to move the unconstrained analysis states along the adjoint constraint manifold \mathbf{G}^{\top} until it lands onto the constraint manifold. This idea is similar to projected Runge-Kutta methods for differential-algebraic equations [38]. For linear constraints, formulation eq. (14) leads to the Mean Square Method derivation of projection from Simon and Chia [23].

3.2 Pseudo-observations of constraints

The pseudo-observation formulation here follows Julier and LaViola Jr [24]. Formally, the standard observation operator $\mathcal{H}(\mathbf{x})$ is augmented with $\mathbf{g}(\mathbf{x})$ to obtain the extended observation operator $\hat{\mathcal{H}}$:

$$\widehat{\mathcal{H}}(\mathbf{x}) \coloneqq \begin{bmatrix} \mathcal{H}(\mathbf{x}) \\ \mathbf{g}(\mathbf{x}) \end{bmatrix},\tag{16}$$

and the data with a pseudo-observation of the constraint value, to obtain the extended observation $\hat{\mathbf{y}}$:

$$\widehat{\mathbf{y}} \coloneqq \begin{bmatrix} \mathbf{y} \\ \mathbf{0} \end{bmatrix}. \tag{17}$$

For any ensemble method, the observation error covariance \mathbf{R} is augmented with a user-defined covariance $\mathbf{R}^{\mathbf{g}}$ that defines an acceptable amount of constraint violation:

$$\widehat{\mathbf{R}} \coloneqq \begin{bmatrix} \mathbf{R} & \mathbf{0} \\ \mathbf{0} & \mathbf{R}^{\mathbf{g}} \end{bmatrix}.$$
(18)

The extended observation triplet $\widehat{\mathcal{H}}(\mathbf{x})$, $\widehat{\mathbf{y}}$, and $\widehat{\mathbf{R}}$ is used in the standard way in filtering methods. However, combining pseudo-observations with localization is difficult and is not explored here.

4 Ensemble Variational Fokker Planck Filters

The ensemble variational Fokker Planck method [17] for data assimilation evolves a set of samples from the prior distribution to the posterior distribution along a defined stochastic differential equation [17] in pseudo-time. Specifically, the samples move under a McKean-Vlasov-Itô process whose drift minimizes the Kullback-Leibler divergence between the probability density of the samples and the target posterior and whose diffusion is tuned to prevent filter divergence. This framework treats the density as a "parameter" estimated from the samples which can be chosen arbitrarily as long as the log-gradient of the density can be computed exactly or estimated. Formally, the McKean-Vlasov-Itô process defined as

$$d\mathbf{x}_{\tau} = \mathbf{F}(\tau, \mathbf{x}_{\tau}, \mathcal{P}_{\tau}) d\tau + \boldsymbol{\sigma}(\tau, \mathbf{x}_{\tau}, \mathcal{P}_{\tau}) d\mathbf{W},$$
(19)

whose corresponding Fokker-Planck equation is

$$\frac{\partial \mathcal{P}_{\tau}(\mathbf{x})}{\partial \tau} = -\nabla_{\mathbf{x}} \cdot \left(\mathcal{P}_{\tau}(\mathbf{x}) \mathbf{F} - \nabla_{\mathbf{x}} \cdot \left(\frac{\mathcal{P}_{\tau}(\mathbf{x}) \boldsymbol{\sigma} \boldsymbol{\sigma}^{\top}}{2} \right) \right), \tag{20}$$

where τ is the pseudo-time, \mathbf{x}_{τ} is a sample, \mathcal{P}_{τ} is the density of the sample, \mathbf{F} is the KL divergence minimizing optimal drift, $\boldsymbol{\sigma}$ is the diffusion tensor and \mathbf{W} is the Wiener process and $(\nabla_{\mathbf{x}} \cdot) : \mathbb{R}^{n_s \times n_s} \to \mathbb{R}^{n_s}$ is the divergence term. The initial conditions are given as

$$\tau = 0, \ \mathbf{x}_{\tau} = \mathbf{x}_{\mathrm{f}} \ \text{and} \ \mathcal{P}_{\tau} = \mathcal{P}_{\mathrm{f}}(\mathbf{x}).$$
 (21)

When eq. (19) is evolved in pseudo-time, we obtain the analysis samples as

$$\tau \to \infty, \ \mathbf{x}_{\tau} = \mathbf{x}_{\mathrm{a}} \ \text{and} \ \mathcal{P}_{\tau} = \mathcal{P}_{\mathrm{a}}(\mathbf{x}).$$
 (22)

In practice, the integration is terminated at some finite pseudo-time based on a criterion evaluating the convergence of the samples to the posterior. The optimal drift that minimizes the KL-divergence between the ensemble estimate and the posterior (derived in [17]) is given as

$$\mathbf{F}(\tau, \mathbf{x}, \mathcal{P}_{\tau}) = \mathbf{A}(\mathbf{x}, \mathcal{P}_{\tau}) \left(\nabla_{\mathbf{x}} \log \mathcal{P}_{\mathrm{a}}(\mathbf{x}) - \nabla_{\mathbf{x}} \log \mathcal{P}_{\tau}(\mathbf{x}) \right) \\ + \left(\frac{\boldsymbol{\sigma}\boldsymbol{\sigma}^{\top}}{2} \right) \nabla_{\mathbf{x}} \log \mathcal{P}_{\tau}(\mathbf{x}) + \nabla_{\mathbf{x}} \cdot \left(\frac{\boldsymbol{\sigma}\boldsymbol{\sigma}^{\top}}{2} \right),$$
(23)

where $\mathbf{A}(\mathbf{x}, \mathcal{P}_{\tau})$ can be any symmetric positive definite matrix. The term $\nabla_{\mathbf{x}} \log \mathcal{P}_{a}(\mathbf{x})$ pushes the samples toward the posterior, $\nabla_{\mathbf{x}} \log \mathcal{P}_{\tau}(\mathbf{x})$ attempts to keep the samples spread apart (deterministically), and the terms containing $\sigma \sigma^{\top}/2$ account for the stochastic diffusion. For a more rigorous description of this framework, we refer the readers to our previous work [17].

The following simplifying assumptions are made for the remainder of the paper:

- (i) Firstly, $\mathbf{A}(\mathbf{x}, \mathcal{P}_{\tau}) = \mathbf{I}_{n_s}$. This is a trivial choice (theoretically and computationally) and further justification is discussed in [17].
- (ii) Deriving the exact flow based on \mathcal{P}_{τ} , \mathcal{P}_{o} , and \mathcal{P}_{f} would require the estimation of probability densities on manifolds. This idea is not explored in this work as the main focus is on constraint preservation in the VFP framework. Thus, for simplicity, the computational flow assumes the densities \mathcal{P}_{τ} , \mathcal{P}_{o} , and \mathcal{P}_{f} to all be Gaussian. While the Gaussian assumption works in practice, VFP is not restricted by this and can choose any quantifiable gradient log density to derive the flow (discussed in [17]). It is more likely than not, that a Gaussian in \mathbb{R}^{n_s} will not be

Gaussian on the constraint manifold. But practically, the experiments show that this is a good enough approximation for the flow.

(iii) The diffusion is assumed to be some scaled square root (or Cholesky factor) of the forecast covariance estimate $\boldsymbol{\sigma}(\tau, \mathbf{x}_{\tau}, \mathcal{P}_{\tau}) = \alpha \sqrt{\mathbf{P}_{\mathrm{f}}}$ with α being a problem specific tunable parameter. As $\boldsymbol{\sigma}$ is independent of \mathbf{x}_{τ} , we have $\left(\frac{\boldsymbol{\sigma}\boldsymbol{\sigma}^{\top}}{2}\right) = \alpha^{2}\mathbf{P}_{\mathrm{f}}$ and $\nabla_{\mathbf{x}} \cdot \left(\frac{\boldsymbol{\sigma}\boldsymbol{\sigma}^{\top}}{2}\right) = \mathbf{0}$ Note that for $\boldsymbol{\sigma} \in \mathbb{R}^{n_{\mathrm{s}} \times \ell}$, where $\ell < n_{\mathrm{s}}$, the diffusion dynamics are restricted to the forecast ensemble subspace [40].

Example 2 (Gaussian case) Under the Gaussian assumptions,

$$\nabla_{\mathbf{x}} \log \mathcal{P}_{\tau}(\mathbf{x}) = -\mathbf{P}_{\tau}^{-1} \left(\mathbf{x} - \bar{\mathbf{x}}_{\tau}\right), \text{ and}$$

$$\nabla_{\mathbf{x}} \log \mathcal{P}_{a}(\mathbf{x}) = -\mathbf{P}_{f}^{-1} \left(\mathbf{x} - \bar{\mathbf{x}}_{f}\right) - \mathbf{H}^{\top} \mathbf{R}^{-1} \left(\mathcal{H}(\mathbf{x}) - \mathbf{y}\right).$$
(24)

However, there are no guarantees about constraint preservation when considering eq. (19) coupled with eq. (23) and the simplifying assumptions. Stochastic diffusion without any special treatment almost always destroys constraints. Deterministic dynamics with $\sigma = 0$ do not preserve constraints either. Thus, drawing inspiration from the "ODE with invariants" section from Ascher and Petzold [38], we proceed as follows. The first idea is to consider eq. (19) with the constraint $\mathbf{g}(\mathbf{x}) = \mathbf{0}$ as a stochastic differential equation with an invariant. Next, through the equivalence between an ODE with an invariant and a Hessenberg index-2 DAE (see equations 9.38, 9.39 in Ascher [38]), we can reformulate eq. (19) as a Hessenberg index-2 stochastic DAE.

4.1 Ensemble Variational Fokker Planck Filters with Stabilization (VFPSTAB)

In the VFPSTAB approach, the SDE in eq. (19) is stabilized for the constraint $\mathbf{g}(\mathbf{x})$. Following Asher and Petzold [38, equation 9.40], we rewrite eq. (19) as

$$d\mathbf{x}_{\tau} = \left(\mathbf{F}(\tau, \mathbf{x}_{\tau}, \mathcal{P}_{\tau}) - \gamma \mathbf{G}^{\dagger}(\mathbf{x}_{\tau})\mathbf{g}(\mathbf{x}_{\tau})\right) d\tau + \boldsymbol{\sigma}(\tau, \mathbf{x}_{\tau}, \mathcal{P}_{\tau}) d\mathbf{W},$$
(25)

where $\gamma > 0$ is a tunable stabilization parameter and $\mathbf{G}^{\dagger} = \mathbf{G}^{\top} (\mathbf{G}\mathbf{G}^{\top})^{-1}$ is the Moore-Penrose inverse of the Jacobian \mathbf{G} eq. (15). Note that rather than \mathbf{G}^{\dagger} , one may use operators such as $\hat{\mathbf{G}} (\mathbf{G}\hat{\mathbf{G}})^{-1}$ where $\hat{\mathbf{G}}$ is any operator such that $\mathbf{G}\hat{\mathbf{G}}$ is always boundedly invertible [38]. The choice for $\hat{\mathbf{G}}$ defines the projection direction and choosing $\hat{\mathbf{G}} = \mathbf{G}^{\top}$, as in eq. (25) results in orthogonal projection. VFPSTAB is evolved in time using either Euler-Maruyama [41] or Rosenbrock-Euler-Maruyama [17].

4.2 Stochastic Differential Algebraic Equation (VFPDAE)

In the VFPDAE approach, eq. (19) is equipped with constraint $\mathbf{g}(\mathbf{x})$ as a Hessenberg index-2 stochastic differential algebraic equation (SDAE). Formally, by introducing an algebraic variable \mathbf{z} into eq. (19), we can write

$$d\mathbf{x}_{\tau} = \left(\mathbf{F}(\tau, \mathbf{x}_{\tau}, \mathcal{P}_{\tau}) - \mathbf{G}^{\top}(\mathbf{x}_{\tau}) \mathbf{z}\right) d\tau + \boldsymbol{\sigma}(\tau, \mathbf{x}_{\tau}, \mathcal{P}_{\tau}) d\mathbf{W},$$

$$\mathbf{0} = \mathbf{g}(\mathbf{x}_{\tau}),$$

(26)

where $\mathbf{G}(\mathbf{x}_{\tau}) = \nabla_{\mathbf{x}} \mathbf{g}(\mathbf{x}_{\tau})$. The algebraic variable $\mathbf{z} \in \mathbb{R}^{n_c}$ projects the solution along \mathbf{G}^{\top} , compensates not just for the deterministic dynamics but also the stochastic Wiener process, which almost surely pushes the states off the constraint manifold. **Remark 2** The constraint $\mathbf{0} = \mathbf{g}(\mathbf{x}_{\tau} + d\mathbf{x}_{\tau})$, implies the tangent constraint $\mathbf{0} = \mathbf{G}(\mathbf{x}_{\tau})d\mathbf{x}_{\tau}$. Plugging $d\mathbf{x}_{\tau}$ from eq. (26) into $\mathbf{0} = \mathbf{G}(\mathbf{x}_{\tau})d\mathbf{x}_{\tau}$, we get

$$\mathbf{0} = \mathbf{G}(\mathbf{x}_{\tau}) \left(\left(\mathbf{F}(\tau, \mathbf{x}_{\tau}, \mathcal{P}_{\tau}) - \mathbf{G}^{\top}(\mathbf{x}_{\tau}) \mathbf{z} \right) d\tau + \boldsymbol{\sigma}(\tau, \mathbf{x}_{\tau}, \mathcal{P}_{\tau}) d\mathbf{W} \right),$$
(27)

which gives

$$\mathbf{z} \, \mathrm{d}\tau = (\mathbf{G}\mathbf{G}^{\top})^{-1} \mathbf{G}(\mathbf{x}_{\tau}) \left(\mathbf{F}(\tau, \mathbf{x}_{\tau}, \mathcal{P}_{\tau}) \mathrm{d}\tau + \boldsymbol{\sigma}(\tau, \mathbf{x}_{\tau}, \mathcal{P}_{\tau}) \mathrm{d}\mathbf{W} \right).$$
(28)

Plugging $\mathbf{z}d\tau$ back in eq. (26), we get the SDAE system

$$d\mathbf{x}_{\tau} = (\mathbf{I} - \mathbf{G}^{\top} (\mathbf{G} \mathbf{G}^{\top})^{-1} \mathbf{G} (\mathbf{x}_{\tau})) \cdot (\mathbf{F}(\tau, \mathbf{x}_{\tau}, \mathcal{P}_{\tau}) d\tau + \boldsymbol{\sigma}(\tau, \mathbf{x}_{\tau}, \mathcal{P}_{\tau}) d\mathbf{W})$$

$$\mathbf{0} = \mathbf{g}(\mathbf{x}_{\tau}).$$
 (29)

The $(\mathbf{I} - \mathbf{G}^{\top}(\mathbf{G}\mathbf{G}^{\top})^{-1}\mathbf{G}(\mathbf{x}_{\tau}))$ term is projecting the stochastic dynamics orthogonally onto the constraint tangent manifold, while eliminating the algebraic variable \mathbf{z} . Note that $\mathbf{G}\mathbf{G}^{\top}$ must always be non-singular for the system to have index-2 uniformly. **Remark 3** Compared to VFPDAE (as in eq. (26)), VFPSTAB (as in eq. (25)) does not satisfy the constraints exactly, but merely makes the states remain close to the constraint manifold. However, VFPSTAB is computationally cheaper than VFPDAE due to not having to solve a root-finding problem at each evolution step. Both methods improve the preservation of constraints when compared to the standard VFP.

4.2.1 Time integration of stochastic differential-algebraic equations

While theory exists for the numerical solution of SDE [41] and index-2 DAEs [38, 42, 43], very minimal literature exists for index-2 SDAE. Due to this, we combine ideas from SDE and DAE literature to evolve an index-2 SDAE. A direct half-explicit solution to the SDAE in eq. (26) is obtained by taking an Euler-Maruyama step, followed by a projection onto the constraint manifold:

$$\tilde{\mathbf{x}}_{1} = \mathbf{x}_{0} + h\mathbf{F}(\tau_{0}, \mathbf{x}_{0}) + \sqrt{h\boldsymbol{\sigma}(\tau_{0}, \mathbf{x}_{0})}\boldsymbol{\xi},$$
Solve for $\mathbf{z}_{0} : \mathbf{g}(\tilde{\mathbf{x}}_{1} - \mathbf{G}^{\top}(\mathbf{x}_{0}) \mathbf{z}_{0}) = \mathbf{0},$

$$\mathbf{x}_{1} = \tilde{\mathbf{x}}_{1} - \mathbf{G}^{\top}(\mathbf{x}_{0}) \mathbf{z}_{0},$$
(30)

where $h = \tau_1 - \tau_0$. This method projects the solution along the adjoint constraint manifold evaluated at \mathbf{x}_0 (as the constraint does not change over pseudo-time τ).

Alternatively, one may discretize eq. (26)

$$\tilde{\mathbf{x}}_{1} = \mathbf{x}_{0} + h\mathbf{F}(\tau_{0}, \mathbf{x}_{0}) + \sqrt{h\boldsymbol{\sigma}}(\tau_{0}, \mathbf{x}_{0})\boldsymbol{\xi},$$
Solve for $\mathbf{z}_{1} : \mathbf{g}(\tilde{\mathbf{x}}_{1} - \mathbf{G}^{\top}(\tilde{\mathbf{x}}_{1}) \mathbf{z}_{1}) = \mathbf{0},$

$$\mathbf{x}_{1} = \tilde{\mathbf{x}}_{1} - \mathbf{G}^{\top}(\tilde{\mathbf{x}}_{1}) \mathbf{z}_{1}.$$
(31)

The difference between eq. (30) and eq. (31) lies in the evaluation of \mathbf{G} at \mathbf{x}_0 versus $\tilde{\mathbf{x}}_1$. This only changes the projection direction \mathbf{G}^{\top} to be computed at \mathbf{x}_1 . Equation (31) can be seen as an evolve and project approach where the solution is evolved by the posterior flow dynamic using Euler-Maruyama [41] discretization which is then projected onto the constraint manifold. The flow can also be discretized with the Rosenbrock Euler-Maruyama (discussed in our work [17]), followed by the projection:

$$\tilde{\mathbf{x}}_{1} = \mathbf{x}_{0} + h \left(\mathbf{I} - h \mathbf{F}_{\mathbf{x}}(\tau_{0}, \mathbf{x}_{0}) \right)^{-1} \mathbf{F}(\tau_{0}, \mathbf{x}_{0}) + \sqrt{h} \boldsymbol{\sigma}(\tau_{0}, \mathbf{x}) \boldsymbol{\xi},$$
Solve for $\mathbf{z}_{1} : \mathbf{g}(\tilde{\mathbf{x}}_{1} - \mathbf{G}^{\top}(\tilde{\mathbf{x}}_{1}) \mathbf{z}_{1}) = \mathbf{0},$

$$\mathbf{x}_{1} = \tilde{\mathbf{x}}_{1} - \mathbf{G}^{\top}(\tilde{\mathbf{x}}_{1}) \mathbf{z}_{1}.$$
(32)

Note that the rootfinding for \mathbf{z}_1 in eqs. (30) to (32) is typically low-dimensional (of dimension \mathbb{R}^{n_c}) compared to the state dimension, i.e. $n_c \ll n_s$. One may also solve for $\tilde{\mathbf{x}}_1$ using any technique from [16, 44] and project the states onto the constraint manifold. For the listed methods eqs. (30) to (32), the solution $\tilde{\mathbf{x}}_1$ will be half-th order accurate ($\mathcal{O}(\sqrt{h})$). Since the algebraic variable \mathbf{z} is just the projection vector, we ignore whatever order it may be. In our experiments, we found eqs. (31) and (32) to give faster results, and these methods have been used for all experiments. **Remark 4** As in remark 2, we can evolve the projected SDE as

$$\tilde{\mathbf{x}}_1 = \mathbf{x}_0 + \left(\mathbf{I} - \mathbf{G}^\top (\mathbf{G}\mathbf{G}^\top)^{-1} \mathbf{G}(\mathbf{x}_0)\right)^{-1} \left(h\mathbf{F}(\tau_0, \mathbf{x}_0) + \sqrt{h}\boldsymbol{\sigma}(\tau_0, \mathbf{x})\boldsymbol{\xi}\right), \quad (33)$$

and project $\tilde{\mathbf{x}}_1$ onto the constraint manifold. In practice, eq. (33) took longer wall-time to converge to the analysis, without any qualitative difference from the solutions of eqs. (31) and (32). For this reason, the results from eq. (33) are not reported in the paper.

5 Numerical experiments

The VFPSTAB and VFPDAE methodologies are tested on three different test problems, namely: (i) the double pendulum, (ii) the Korteweg-de Vries equation, and (iii) the incompressible Navier Stokes via barotropic vorticity equation. These results are compared to the results from the standard VFP, the standard ETKF, ETKFP (ETKF with projection as described in section 3.1), and ETKFA (ETKF augmented with pseudo-observations as described in section 3.2). For the incompressible Navier-Stokes, we also look at LETKFP and LETKFA which are R-localized [2] versions of ETKFP and ETKFA respectively.

Two different metrics—the spatio-temporal root mean square error (RMSE) of the analysis states and the spatio-temporal root mean square error of the analysis constraints (CRMSE)—are used to evaluate the results. The RMSE at time t_k is defined as

$$\operatorname{RMSE}(k) = \sqrt{\frac{1}{C_1} \sum_{i=\rho}^{k} \sum_{e=1}^{n_e} \left\| \mathbf{x}_{\mathrm{a},i}^{[e]} - \mathbf{x}_{\mathrm{true},i} \right\|_2^2},$$
(34)

where $C_1 = (k - \rho + 1)n_e n_s$ where n_s is the number of states, n_e is the number of ensemble members, and ρ is the spinup steps before which the error statistics are ignored. We chose to not use the traditional spatio-temporal RMSE formula $\sqrt{\frac{1}{(k-\rho+1)n_s}\sum_{i=\rho}^k \|\bar{\mathbf{x}}_{a,i} - \mathbf{x}_{true,i}\|_2^2}$, with $\bar{\mathbf{x}}_{a,i} = \frac{1}{n_e}\sum_{e=1}^{n_e} \mathbf{x}_{a,i}^{[e]}$, since the mean of an ensemble $\bar{\mathbf{x}}_{a,i}$ will not lie on the constraint manifold. In an ideal situation, one would have to compute the mean on the constraint manifold and compute the errors using distances on the constraint manifold. This approach is not pursued due to the complexities of computing distances on a manifold. The CRMSE at time t_k is defined

as

$$\operatorname{CRMSE}(k) = \sqrt{\frac{1}{C_2} \sum_{i=\rho}^{k} \sum_{e=1}^{n_e} \left\| \mathbf{E} \cdot \mathbf{g}(\mathbf{x}_{\mathrm{a},i}^{[e]}) \right\|_2^2},$$
(35)

where $C_2 = (k - \rho + 1)n_e n_c$ and **E** is a chosen scaling matrix.

5.1 Covariance Regularization

Covariance estimation from a small ensemble is an important and challenging issue in data assimilation [45]. Various justifiable heuristics such as inflation [46], localization [47], covariance shrinkage [48], and modified Cholesky [49], have been developed to improve the estimates.

For constraint-preserving systems the covariance in full state space is singular, and the empirically estimated ensemble variances are small along directions normal to the constraint manifold, leading to filter divergence. In the VFP, the computed drift in eq. (24) requires a linear system solution of the covariance matrix. If the covariance is either singular or if its entries are too small, the computed drift will be inaccurate leading to filter divergence over multiple assimilation cycles. For this reason, we propose covariance regularization approaches to prevent VFP filter divergence for the dynamics of eqs. (25) and (26).

The estimated covariance is regularized via static covariance shrinkage for the double pendulum. The estimated covariance \mathbf{P} is shrunk to $\gamma_{sh}\mathbf{P} + (1 - \gamma_{sh})\mathbf{I}$, with $\gamma_{sh} = 0.01$. While there are robust methods to compute γ_{sh} such as the Ledoit-Wolf and the Rao-Blackwell Ledoit-Wolf methods [48], we prefer an experimentally tuned static value for quicker computation.

For high-dimensional PDE systems, it is well-known that the square root of the covariance matrix can approximated by the inverse of a spatial Laplacian-like operator [50, 51] by its structure (specifically, the inverse of the Laplacian-like operator is diagonally dominant, symmetric, and positive definite). The precision matrix estimate for the Kortweg-de Vries and the incompressible Navier-Stokes equations is given

by scaling the discrete spatial Laplacian-like of the system with a scalar which is the inverse of the maximum (across the n_s states) variance of the ensemble. For the incompressible Navier-Stokes problem, we assume that the velocities in the x and ydimensions are uncorrelated. The specific Laplacian-like approximations are described in section 5.3 and section 5.4 respectively.

5.2 Double Pendulum

The double pendulum is a constrained mechanical system [38, 43]. This problem was chosen due to its chaotic dynamical nature and ease of constraint visualization while posing a challenging assimilation problem. We consider the index-2 DAE formulation of the pendulum dynamics

$$\frac{\mathrm{d}^2}{\mathrm{d}t^2} \begin{bmatrix} x_1\\y_1\\x_2\\y_2 \end{bmatrix} = \begin{bmatrix} -\lambda_1 x_1 + \lambda_2 (x_2 - x_1)\\ -\lambda_1 y_1 + \lambda_2 (y_2 - y_1) - g_c\\ -\lambda_2 (x_2 - x_1)\\ -\lambda_2 (y_2 - y_1) - g_c \end{bmatrix},$$
(36)

where $g_c = 9.8$ is the acceleration due to gravity, x_1, y_1, x_2, y_2 are the positions of the two pendulums in Cartesian coordinates with rod tensions λ_1, λ_2 being the solution of the linear system

$$\begin{bmatrix} x_1^2 + y_1^2 & -(x_1\delta_x + y_1\delta_y) \\ -(x_1\delta_x + y_1\delta_y) & 2(\delta_x^2 + \delta_y^2) \end{bmatrix} \begin{bmatrix} \lambda_1 \\ \lambda_2 \end{bmatrix} = \begin{bmatrix} u_1^2 + v_1^2 - g_c y_1 \\ \delta_u^2 + \delta_v^2 \end{bmatrix}, \quad (37)$$

where $u_1 = \frac{dx_1}{dt}$, $v_1 = \frac{dy_1}{dt}$, $u_2 = \frac{dx_2}{dt}$, $v_2 = \frac{dy_2}{dt}$ are the velocities of the two pendulums in the x and y directions, and $\delta_x = (x_2 - x_1)$, $\delta_y = (y_2 - y_1)$, $\delta_u = (u_2 - u_1)$, $\delta_v = (v_2 - v_1)$. The system has $n_c = 5$ constraints on the lengths of the two rods, velocities of the two rods, and total mechanical energy respectively defined as

$$\begin{bmatrix} 0\\0\\0\\0\\0\\0\end{bmatrix} = \begin{bmatrix} 0.5(x_1^2 + y_1^2 - 1)\\0.5(\delta_x^2 + \delta_y^2 - 1)\\x_1u_1 + y_1v_1\\\delta_x\delta_u + \delta_y\delta_v\\0.5(u_1^2 + v_1^2 + u_2^2 + v_2^2) + g_c(y_1 + y_2 + 3) - E_0 \end{bmatrix},$$
(38)

where E_0 is the total mechanical energy at the initial condition.

The masses and lengths are assumed to be unity, i.e. $m_1 = m_2 = 1$ and $l_1 = l_2 = 1$. This index-2 system evolved as a first-order index-2 DAE and using a two-stage, partitioned half explicit Runge-Kutta (PHERK) method [52, 53] for $\Delta t_{ts} = 0.01$. The coefficients of the second-order PHERK method [52, 53] are

$$A = \begin{bmatrix} 0 & 0 \\ 1 & 0 \end{bmatrix}, \tilde{A} = \begin{bmatrix} 0 & 0 \\ \frac{1}{2} & \frac{1}{2} \end{bmatrix}, c = \tilde{c} = \begin{bmatrix} 0 \\ 1 \end{bmatrix}, b = \begin{bmatrix} \frac{1}{2} \\ \frac{1}{2} \end{bmatrix}.$$
 (39)

The different matrices A and \tilde{A} are used to solve for the differential and algebraic variables following the work of Murua [52].

The observation operator $\mathcal{H}(\mathbf{x}) = \mathbf{x}$ observes all states with an unbiased, independent Gaussian error $\eta^{\circ} \sim N(\mathbf{0}, \mathbf{R} = 0.1 \mathbf{I}_8)$ every $\Delta t = 0.1$ time units. The problem is assimilated for 5501 steps (each step at $\Delta t = 0.1$) and ignores the statistics from the first $\sigma = 501$ spinup steps. A reference system state is chosen to be some position of the double pendulum with zero velocities having a total mechanical energy of $E_0 = 56.1741$. This reference is then evolved $n_e = 30$ times at a $\Delta t_{samp} = 0.008$ to create a trajectory of states. The truth and the ensemble are state samples without repetition from this trajectory.

The assimilation methods and their parameters are as follows.

- 1. ETKF, ETKFA, and ETKFP all use a covariance inflation of $\alpha = 1.08$.
- 2. ETKFA uses the augmented error covariance $\mathbf{R}^{\mathbf{g}} = 0.001 \mathbf{I}_5$. This value was handtuned to balance between the RMSE and CRMSE.
- 3. All VFP variants—VFP, VFPSTAB, VFPDAE—have the diffusion parameter

$$\boldsymbol{\sigma} = 0.001 \operatorname{diag} \left(\left[2 \ 2 \ 20 \ 20 \ 2 \ 2 \ 20 \ 20 \right] \right)$$

This was tuned to be a scaling of the climatological (or auto) covariance of the double pendulum system. The VFP system was evolved to the posterior in pseudo-time τ using Euler-Maruyama with a $\Delta \tau = 0.001$. The evolution was terminated when the absolute change in the ensemble mean across pseudo-time steps was under a specified tolerance, here $1e^{-6}$. The covariances \mathbf{P}_{τ} , and $\mathbf{P}_{\rm f}$ were estimated using covariance shrinkage as discussed in Section 5.1.

4. VFPSTAB uses a hand-tuned stabilization constant of $\gamma = 30$ (to balance between RMSE and CRMSE).



Fig. 2: Results for the double pendulum experiments.

Figure 2a compares the spatio-temporal root mean squared errors of various flavors of ETKF and VFP. For this problem, all the ETKF variants behave slightly better than all the VFP variants, with VFPDAE coming very close to plain ETKF. For the ETKF methods, constraint preservation (ETKFA and ETKFP) shows a slightly lower RMSE when compared to plain ETKF.

Figure 2b shows the CRMSE for each method with the scaling $\mathbf{E} = \text{diag}\left(\left[1 \ 1 \ 1 \ 1 \ \frac{1}{E_0}\right]\right)$. The asymptotic CRMSE values of ETKF, ETKFA, VFP, and VFPSTAB are 0.005, 0.003, 0.065, and 0.046 respectively. As expected, ETKFP and VFPDAE preserve the constraints to machine precision. ETKFA and ETKF have the next lowest CRMSE. Due to constraint augmentation, ETKFA has a slightly lower CRMSE than ETKF. Similarly, the CRMSE of VFPSTAB and VFP look identical, but VFPSTAB has a slightly lower CRMSE due to stabilization.

5.3 Korteweg de Vries

The Korteweg-de Vries PDE represents non-linear, non-dissipative waves on shallow water surfaces [54]. While this PDE has an infinite set of integrable constraints [55] on the states, our experiments consider the first three only. The equation is written as:

$$\frac{\partial \mathbf{x}}{\partial t} + 3\frac{\partial (\mathbf{x}^2)}{\partial x} + \frac{\partial^3 \mathbf{x}}{\partial x^3} = 0, \tag{40}$$

where \mathbf{x} is a non-dimensional wave-height displacement, and t is the non-dimensional time. The mass, momentum, and energy integral invariants are respectively given by

$$\phi_{1}(\mathbf{x}) = \int_{\Omega_{x}} \mathbf{x} dx,$$

$$\phi_{2}(\mathbf{x}) = \int_{\Omega_{x}} \frac{1}{2} \mathbf{x}^{2} dx,$$

$$\phi_{3}(\mathbf{x}) = \int_{\Omega_{x}} \left(\frac{\mathbf{x}^{3}}{3} - \left(\frac{\partial \mathbf{x}}{\partial x}\right)^{2}\right) dx.$$
(41)

Thus, the $n_{\rm c} = 3$ constraints, are

$$\mathbf{0} = \mathbf{g}(t, \mathbf{x}) = \phi(\mathbf{x}) - \phi(\mathbf{x}_0). \tag{42}$$

The PDE is spatially discretized using the second-order central finite difference in the periodic domain $x \in [-10, 10]$ with $n_s = 100$ points. This semi-discretized PDE is evolved using the implicit midpoint scheme (with $\Delta t_{ts} = 0.01$) which is known to preserve the invariants [56]. Following [57], the third constraint in eq. (42) is discretized with the first-order forward finite difference.

At time intervals of $\Delta t = 0.01$, every fourth state is observed as $\mathcal{H}(\mathbf{x}) = \mathbf{x}(4:4:100)$ with an unbiased Gaussian observation error with covariance $\mathbf{R} = 0.2 \mathbf{I}_{25}$. The observation covariance corresponds to 10% of the climatological (or auto) covariance. Assimilation is done for 2201 steps, and the reported results ignore the statistics from the first $\rho = 401$ spinup steps. The initial true state $\mathbf{x}_{\text{true},0} = 6 \operatorname{sech}^2(x)$ corresponds to the two soliton problem. The constraints for this are $\phi(\mathbf{x}_0) = \begin{bmatrix} 12 \ 48 \ -211.3815 \end{bmatrix}^{\top}$. The initial ensemble with $n_e = 10$ samples is created by perturbing $\mathbf{x}_{\text{true},0}$ with diagonal random noise and projecting this onto the constraint manifold.

We consider the following hand-tuned data assimilation methods:

- 1. ETKF, ETKFA, and ETKFP all use an inflation of $\alpha = 1.04$ at every assimilation step.
- 2. ETKFA uses an augmented error covariance $\mathbf{R}^{\mathbf{g}} = 10 \mathbf{I}_5$ for the constraints.
- 3. VFPSTAB, and VFPDAE use a hand-tuned diffusion of $\boldsymbol{\sigma} = 1e^{-4}(\Delta + 1e^{-3}\mathbf{I}_{n_s})^{-1}$ where Δ is the discrete Laplacian. As stated in section 5.1, the covariances \mathcal{P}_{τ}^{-1} , and \mathcal{P}_{f}^{-1} are estimated by multiplying $(\Delta + 1e^{-3}\mathbf{I}_{n_s})^2$ with the maximum variance of the intermediate and forecast ensembles respectively. The term $1e^{-3}\mathbf{I}_{n_s}$ is added to the discrete Laplacian as the discrete Laplacian is low-rank due to the periodic domain. Additionally, perturbed observations [58] were used in the VFP

methods to allow for a larger analysis spread with perturbations being drawn from $\mathcal{N}(\mathbf{0}, 0.05^2 \mathbf{R})$. This is because increasing the diffusion results in irregular solutions that do not lie on the constraint manifold.

4. The VFPSTAB and VFPDAE system are solved in pseudo-time τ using Rosenbrock-Euler-Maruyama with a $\Delta \tau = 0.01$. The evolution is terminated when the absolute change in the ensemble mean across pseudo-time steps falls under a specified tolerance, here $1e^{-4}$.

The LETKF [8] (and LETKFP, LETKFA) results are not reported as they are similar to the unlocalized ETKF (and ETKFP, ETKFA) results.



Fig. 3: Results for the Korteweg-deVries experiments where the VFP methods outperform the ETKF methods.

Figure 3a shows the spatio-temporal RMSE in time for various analysis schemes. The VFP methods (VFPDAE and VFPSTAB) perform better than the ETKF variants. Between the ETKF methods, ETKFA seems to perform better than ETKF and ETKFP for the RMSE. This might mean that the ETKFP is attaining a worse solution

26

on the manifold, while ETKFA is obtaining a relatively good solution on the manifold compared to vanilla ETKF. VFPDAE outperforms all these methods. Figure 3b shows the CRMSE values for the different assimilation methods with the scaling matrix $\mathbf{E} = \mathbf{I}_3$. The asymptotic CRMSE values of ETKF, ETKFA, and VFPSTAB are 0.005, 0.003, and 0.0007 respectively. The asymptotic CRMSE values of ETKFP and VFPDAE are at machine epsilon. ETKF is worse at preserving constraints than ETKFA, both of which are worse than VFPSTAB. VFPSTAB is worse than ETKFP and VFPDAE which preserve the constraints exactly.

5.4 Incompressible Navier-Stokes

The incompressible Navier-Stokes equations (in the vorticity-streamfunction form) are used to model large-scale, oceanic flows in shallow basins [59]. Without external forcing, this system must conserve the mass, while also conserving the total energy and enstrophy [19, 30]. This problem was chosen since it has both local and global constraints.

The equations are

$$\frac{\partial\omega}{\partial t} = -J(\psi,\omega) + Re^{-1}\Delta\omega + Ro^{-1}\left(\frac{\partial\psi}{\partial x} + F\right),$$

$$\psi = -\Delta^{-1}\omega,$$
(43)

with the vorticity ω , streamfunction ψ , velocities $(u, v) = \left(\frac{\partial \psi}{\partial y}, -\frac{\partial \psi}{\partial x}\right)$, the Laplacian $\Delta = \frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2}$, the Reynolds number Re = 450, the Rossby number Ro = 0.0036, the Jacobian $J(\psi, \omega) \equiv \frac{\partial \psi}{\partial y} \frac{\partial \omega}{\partial x} - \frac{\partial \psi}{\partial x} \frac{\partial \omega}{\partial y}$, and symmetric double gyre forcing $F = \sin(\pi y)$. The streamfunction and vorticity boundaries on x and y are homogeneous Dirichlet. The settings are chosen from San et. al [59] to produce four gyre circulation over time.

The experiment setup involves solving the vorticity equation in time and then computing the primitive variables (velocities (u, v)) as the forecast. After constraint preserving assimilation on the velocities, the velocities are converted back to the vorticity form as $\omega = \frac{\partial v}{\partial x} - \frac{\partial u}{\partial y}$ with appropriate boundary conditions and marched forward in time. This approach allows the incompressible Navier-Stokes to be solved as an ODE without computational difficulties.

Firstly, the primitive variables must be incompressible due to the conservation of mass, i.e. $\frac{\partial u}{\partial x} + \frac{\partial v}{\partial y} = 0$. Next, additional constraints are constructed to preserve the energy and enstrophy for each ensemble member before and after filtering. This means that each ensemble member will have its own, independent constraint. The incompressibility, energy, and enstrophy constraints are respectively

$$\frac{\partial u_a}{\partial x} + \frac{\partial v_a}{\partial y} = 0,$$

$$\frac{1}{2} \int_{\Omega} \left(u_a^2 + v_a^2 \right) dx \, dy = \frac{1}{2} \int_{\Omega} \left(u_f^2 + v_f^2 \right) dx \, dy,$$

$$\frac{1}{2} \int_{\Omega} \left(\frac{\partial v_a}{\partial x} - \frac{\partial u_a}{\partial y} \right)^2 dx \, dy = \frac{1}{2} \int_{\Omega} \left(\frac{\partial v_f}{\partial x} - \frac{\partial u_f}{\partial y} \right)^2 dx \, dy.$$
(44)

The problem is discretized on the domain $x \in [0, 1]$ and $y \in [-1, 1]$ with central finite differences using 64 and 129 equidistant points in each dimension. This makes $n_{\rm s} = 16512$ from having 2 variables (u and v) on the grid. This discretization results in 8256 local constraints due to incompressibility, and 2 global constraints for the energy and enstrophy, making $n_{\rm c} = 8258$. The Arakawa Jacobian [30] with the RK3 [60] method is used to evolve the semidiscretized PDE with a timestep of $\Delta_{ts} = 0.0109/40$.

The system is observed every $\Delta t = 0.0109$ units equivalent to one real day with this non-dimensionalization [59]. The velocities u and v are observed at $16 \times 16 = 256$ equally spaced grid points in the interior of the domain, with a Gaussian observation error covariance of $\mathbf{R} = 400\mathbf{I}_{512}$. This observation error covariance is approximately 1% of the climatological (or auto) covariance of the velocities. The assimilation was done for 881 steps neglecting the statistics from the first 81 steps as spinup. The true

initial condition is generated by evolving a smooth random vorticity field for 100 days. Random sampling around the truth generates the initial ensemble with $n_{\rm e} = 10$.

The following filtering methods are considered for this problem:

- 1. ETKF and ETKFP with hand-tuned inflations of $\alpha = 1.1$ and $\alpha = 1.05$, respectively.
- 2. LETKF, and LETKFP, with an inflation of $\alpha = 1.05$. Here, LETKF and LETKFP use the Gaspari-Cohn localization [47, 61] with a hand-tuned radius of $r_{loc} = 0.559$ units. The distance used to compute the localization coefficients is the spatial Euclidean distance.
- 3. VFPDAE with diffusion $\boldsymbol{\sigma} = \text{blkdiag}\left(\left[\Delta \Delta\right]\right)^{-1}$, where Δ is the twodimensional discrete Laplacian on the grid. Note that this assumes that the correlation between u and v on the same grid point is 0. From the authors' experience, it was observed that the velocities in the x and y directions were uncorrelated, making this heuristic meaningful. As discussed in section 5.1, \mathcal{P}_{τ}^{-1} , and \mathcal{P}_{f}^{-1} are computed by scaling blkdiag $\left(\left[\Delta^{2} \Delta^{2}\right]\right)$ with the maximum variance of the intermediate and forecast ensembles respectively. Also, to accelerate convergence, we initialize the VFPDAE solution with a projected solution from ETKF. The VFPDAE system initialized at the projected ETKF solution was evolved to the analysis with the flow dynamic for 10 steps using Euler-Maruyama with a $\Delta \tau = 1e^{-6}$. Due to the cost of this method, the pseudo-time evolution was stopped after 10 evolution steps.
- 4. Augmented observations could not be used for the ETKF and LETKF as the constraint is different for each ensemble member.

The results in fig. 4 show that ETKF and ETKFP methods perform worse than their localized variants LETKF and LETKFP. Our method VFPDAE performs better than ETKF and ETKFP, but worse than LETKF and LETKFP. Even with 10 small steps of size $1e^{-5}$, VFPDAE performs better than its initialization method ETKFP.



Fig. 4: Incompressible Navier-Stokes: RMSE for different methods.

Perhaps more evolution pseudo-time could give better results but this was not done due to the high computational cost.



(a) Enstrophy CRMSE for different methods. (b) Divergence CRMSE for different methods.Fig. 5: CRMSE results for the Navier-Stokes experiments.

Figure 5a chooses an $\mathbf{E} = \text{diag}\left(\begin{bmatrix} 0 & \dots & 0 & 1 \end{bmatrix}\right)$ for looking only at the enstrophy. The asymptotic CRMSE values for the ETKF, LETKF, and LETKFP, are 0.03, 0.02,

and 6e-4 respectively. VFPDAE and ETKFP preserve the enstrophy to the numerical limit, while ETKF and LETKF do not. LETKFP preserves enstrophy to the numerical limit at all times except step 840, which makes its CRMSE go up. This happens because the projection method fails to project the LETKF solution onto the constraint manifold.

A more interesting situation is looking at the spatially averaged CRMSE for only the divergence constraint across the states with $\mathbf{E} = \text{diag}\left(\begin{bmatrix}1 & \dots & 1 & 0 & 0\end{bmatrix}\right)$ in Figure 5b. LETKF violates the divergence constraint by a huge margin, with the final time CRMSE being 0.08. It is important to notice that ETKF preserves it naturally (around an asymptotic CRMSE of ten times machine epsilon) due to it being a linear transform of the ensemble as also observed by [21]. ETKFP, LETKFP, and VFPDAE preserve the divergence constraint to the numerical limit.

6 Conclusions

The VFP framework is extended to incorporate general non-linear equality constraints. The two formulated methods—VFPDAE and VFPSTAB—respect the constraint manifold throughout the evolution to the posterior. To preserve constraints *exactly*, one must use the VFPDAE method which requires a somewhat expensive implicit-explicit time-stepping method to evolve the dynamics. VFPSTAB allows one to make computational cost gains at the cost of *inexactly* preserving the constraints. Both these methods require covariance regularization as the constraint manifold can restrict the particles to have small variance, leading to filter divergence over time. While not tested, non-linear inequality and box constraints can be incorporated by drawing on ideas from constrained optimization theory such as using an "active-set" of constraints in the flow, enabling VFP to handle more generic constraints. The same constraintpreserving ideas can be extended to smoothing over a trajectory of observations. Future work will include work on both smoothing and inequality constraints.

Acknowledgements. This work was supported by DOE through award ASCR DE-SC0021313, by NSF through award CDS&E–MSS 1953113, and by the Computational Science Laboratory at Virginia Tech. The authors also thank Dr. Traian Iliescu from the Mathematics Department at Virginia Tech for helpful discussions.

Declarations

Competing interests. All authors certify that they have no affiliations with or involvement in any organization or entity with any financial interest or non-financial interest in the subject matter or materials discussed in this manuscript.

Data availability. Code will be made available on request by email.

References

- [1] Asch, M., Bocquet, M., Nodet, M.: Data Assimilation. Society for Industrial and Applied Mathematics, Philadelphia, PA (2016). https://doi.org/10.1137/1. 9781611974546 . https://epubs.siam.org/doi/abs/10.1137/1.9781611974546
- [2] Reich, S., Cotter, C.: Probabilistic Forecasting and Bayesian Data Assimilation. Cambridge University Press, Cambridge, UK (2015)
- [3] Evensen, G., Vossepoel, F.C., Leeuwen, P.J.: Data Assimilation Fundamentals: A Unified Formulation of the State and Parameter Estimation Problem. Springer, Cham (2022). https://doi.org/10.1007/978-3-030-96709-3 . https://doi.org/10. 1007/978-3-030-96709-3
- [4] Kalman, R.E.: A new approach to linear filtering and prediction problems. Transaction of the ASME- Journal of Basic Engineering 82, 35–45 (1960)
- [5] Evensen, G.: Sequential data assimilation with a nonlinear quasi-geostrophic

model using Monte Carlo methods to forcast error statistics. Journal of Geophysical Research 99(C5), 10143–10162 (1994)

- [6] Bishop, C., Etherton, B.J., Majumdar, S.: Adaptive sampling with the Ensemble Transform Kalman Filter. Part I: Theoretical Aspects. Monthly Weather Review 129, 420–436 (2001)
- [7] Anderson, J.L.: An ensemble adjustment Kalman filter for data assimilation. Monthly Weather Review 129, 2884–2903 (2001)
- [8] Hunt, B.R., Kostelich, E.J., Szunyogh, I.: Efficient data assimilation for spatiotemporal chaos: A local ensemble transform Kalman filter. Physica D: Non-linear Phenomena 230(1), 112–126 (2007) https://doi.org/10.1016/j.physd.2006.
 11.008 . Data Assimilation
- [9] Leeuwen, P.J.: Particle filtering in geophysical systems. Monthly Weather Review 137, 4089–4114 (2009)
- [10] Leeuwen, P.J., Künsch, H.R., Nerger, L., Potthast, R., Reich, S.: Particle filters for high-dimensional geoscience applications: A review. Quarterly Journal of the Royal Meteorological Society 145(723), 2335–2365 (2019) https://doi.org/10. 1002/qj.3551 https://rmets.onlinelibrary.wiley.com/doi/pdf/10.1002/qj.3551
- [11] Pulido, M., van Leeuwen, P.J.: Sequential Monte Carlo with kernel embedded mappings: The mapping particle filter. Journal of Computational Physics 396, 400–415 (2019) https://doi.org/10.1016/j.jcp.2019.06.060
- [12] Hu, C.-C., Leeuwen, P.J.: A particle flow filter for high-dimensional system applications. Quarterly Journal of the Royal Meteorological Society 147(737), 2352–2374 (2021) https://doi.org/10.1002/qj.4028 https://rmets.onlinelibrary.wiley.com/doi/pdf/10.1002/qj.4028

- [13] Pathiraja, S., Reich, S.: Discrete gradients for computational Bayesian inference. Journal of Computational Dynamics 6(2158-2491_2019_2_385), 385 (2019) https: //doi.org/10.3934/jcd.2019019
- [14] Reich, S., Weissmann, S.: Fokker–Planck Particle Systems for Bayesian Inference: Computational Approaches. SIAM/ASA Journal on Uncertainty Quantification 9(2), 446–482 (2021) https://doi.org/10.1137/19M1303162 https://doi.org/10.1137/19M1303162
- [15] Garbuno-Inigo, A., Hoffmann, F., Li, W., Stuart, A.M.: Interacting Langevin Diffusions: Gradient Structure and Ensemble Kalman Sampler. SIAM Journal on Applied Dynamical Systems 19(1), 412–441 (2020) https://doi.org/10.1137/ 19M1251655 https://doi.org/10.1137/19M1251655
- [16] Crouse, D.F., Lewis, C.: Consideration of particle flow filter implementations and biases. Naval Research Laboratory Memo, 1–17 (2019)
- [17] Subrahmanya, A.N., Popov, A.A., Sandu, A.: Ensemble Variational Fokker-Planck Methods for Data Assimilation (2023)
- [18] Hastie, T., Tibshirani, R., Friedman, J.: The Elements of Statistical Learning: Data Mining, Inference, and Prediction. Springer, New York, NY (2009). https:// doi.org/10.1007/978-0-387-84858-7 . https://doi.org/10.1007/978-0-387-84858-7
- [19] Arakawa, A.: Design of the ucla general circulation model. Technical report, University of California, Los Angeles (1972)
- [20] Zeng, Y., Janjić, T.: Study of conservation laws with the local ensemble transform Kalman filter. Quarterly Journal of the Royal Meteorological Society 142(699), 2359–2372 (2016) https://doi.org/10.1002/qj.2829 https://rmets.onlinelibrary.wiley.com/doi/pdf/10.1002/qj.2829

- [21] Janjić, T., McLaughlin, D., Cohn, S.E., Verlaan, M.: Conservation of mass and preservation of positivity with ensemble-type Kalman filter algorithms. Monthly Weather Review 142(2), 755–773 (2014) https://doi.org/10.1175/ MWR-D-13-00056.1
- [22] Massicotte, D., Morawski, R.Z., Barwicz, A.: Incorporation of a positivity constraint into a Kalman-filter-based algorithm for correction of spectrometric data. IEEE Transactions on Instrumentation and Measurement 44(1), 2–7 (1995) https://doi.org/10.1109/19.368111
- [23] Simon, D., Chia, T.L.: Kalman filtering with state equality constraints. IEEE Transactions on Aerospace and Electronic Systems 38(1), 128–136 (2002) https: //doi.org/10.1109/7.993234
- Julier, S.J., LaViola, J.J.: On Kalman filtering with nonlinear equality constraints.
 IEEE Transactions on Signal Processing 55(6), 2774–2784 (2007) https://doi. org/10.1109/TSP.2007.893949
- [25] Gupta, N., Hauser, R.: Kalman Filtering with Equality and Inequality State Constraints (2007)
- [26] Sircoulomb, V., Hoblos, G., Chafouk, H., Ragot, J.: State estimation under nonlinear state inequality constraints. a tracking application. In: 2008 16th Mediterranean Conference on Control and Automation, pp. 1669–1674 (2008). https://doi.org/10.1109/MED.2008.4602024
- [27] Zanetti, R., Majji, M., Bishop, R.H., Mortari, D.: Norm-constrained Kalman filtering. Journal of Guidance, Control, and Dynamics 32(5), 1458–1465 (2009) https://doi.org/10.2514/1.43119 https://doi.org/10.2514/1.43119

- [28] Prakash, J., Patwardhan, S.C., Shah, S.L.: Constrained nonlinear state estimation using ensemble Kalman filters. Industrial & Engineering Chemistry Research 49(5), 2242–2253 (2010) https://doi.org/10.1021/ie900197s https://doi.org/10.1021/ie900197s
- [29] Bavdekar, V.A., Prakash, J., Shah, S.L., Gopaluni, R.B.: Constrained dual ensemble Kalman filter for state and parameter estimation. In: 2013 American Control Conference, pp. 3093–3098 (2013). https://doi.org/10.1109/ACC.2013.6580306
- [30] Kacimi, A., Aliziane, T., Khouider, B.: The Arakawa Jacobian method and a fourth-order essentially nonoscillatory scheme for the beta-plane barotropic equations. International Journal of Numerical Analysis & Modeling 10(3) (2013)
- [31] Zeng, Y., Janjić, T., Ruckstuhl, Y., Verlaan, M.: Ensemble-type Kalman filter algorithm conserving mass, total energy and enstrophy. Quarterly Journal of the Royal Meteorological Society 143(708), 2902–2914 (2017) https://doi.org/10. 1002/qj.3142 https://rmets.onlinelibrary.wiley.com/doi/pdf/10.1002/qj.3142
- [32] Li, R., Jan, N.M., Huang, B., Prasad, V.: Constrained ensemble Kalman filter based on Kullback–Leibler divergence. Journal of Process Control 81, 150–161 (2019)
- [33] Albers, D.J., Blancquart, P.-A., Levine, M.E., Seylabi, E.E., Stuart, A.: Ensemble Kalman methods with constraints. Inverse Problems 35(9), 095007 (2019) https: //doi.org/10.1088/1361-6420/ab1c09
- [34] Schillings, C., Stuart, A.M.: Analysis of the Ensemble Kalman Filter for Inverse Problems. SIAM Journal on Numerical Analysis 55(3), 1264–1290 (2017) https: //doi.org/10.1137/16M105959X https://doi.org/10.1137/16M105959X

- [35] Chada, N.K., Schillings, C., Weissmann, S.: On the incorporation of boxconstraints for ensemble Kalman inversion. Foundations of Data Science 1(4), 433–456 (2019) https://doi.org/10.3934/fods.2019018
- [36] Bertsekas, D.P.: Projected Newton methods for optimization problems with simple constraints. SIAM Journal on control and Optimization 20(2), 221–246 (1982)
- [37] Amor, N., Rasool, G., Bouaynaya, N.C.: Constrained state estimation-a review. arXiv preprint arXiv:1807.03463 (2018)
- [38] Ascher, U.M., Petzold, L.R.: Computer Methods for Ordinary Differential Equations and Differential-Algebraic Equations vol. 61. SIAM, Philadelpha, PA (1998). https://doi.org/10.1137/1.9781611971392 . https://epubs.siam.org/doi/ book/10.1137/1.9781611971392
- [39] Anco, S.C., Cheviakov, A.F.: On the different types of global and local conservation laws for partial differential equations in three spatial dimensions: Review and recent developments. International Journal of Non-Linear Mechanics 126, 103569 (2020) https://doi.org/10.1016/j.ijnonlinmec.2020.103569
- [40] Garbuno-Inigo, A., Nüsken, N., Reich, S.: Affine Invariant Interacting Langevin Dynamics for Bayesian Inference. SIAM Journal on Applied Dynamical Systems 19(3), 1633–1658 (2020) https://doi.org/10.1137/19M1304891 https://doi.org/10.1137/19M1304891
- [41] Kloeden, P.E., Platen, E.: Numerical Solution of Stochastic Differential Equations. Stochastic Modelling and Applied Probability. Springer, Berlin, Heidelberg (1992). https://doi.org/10.1007/978-3-662-12616-5_2 . https://doi.org/ 10.1007/978-3-662-12616-5_2

- [42] Wanner, G., Hairer, E.: Solving Ordinary Differential Equations II: Stiff and Differential-Algebraic Problems vol. 375. Springer, Berlin, Heidelberg (1996). https://doi.org/10.1007/978-3-642-05221-7 . https://doi.org/10.1007/ 978-3-642-05221-7
- [43] Hairer, E., Lubich, C., Roche, M.: The Numerical Solution of Differential-Algebraic Systems by Runge-Kutta Methods vol. 1409. Springer, Berlin, Heidelberg (2006)
- [44] Crouse, D.F.: Particle flow solutions avoiding stiff integration. NRL Technical Report No. NRU5340/FR-2021/1 (2021)
- [45] Matsuo, T.: Covariance modeling in applications of data assimilation to highdimensional earth and geospace systems. Journal of Physics: Conference Series 1036(1), 012019 (2018) https://doi.org/10.1088/1742-6596/1036/1/012019
- [46] Anderson, J.L., Anderson, S.L.: A Monte Carlo implementation of the nonlinear filtering problem to produce ensemble assimilations and forecasts. Mon. Wea. Rev. 127, 2741–2785 (1999)
- [47] Houtekamer, P.L., Mitchell, H.L.: A sequential ensemble Kalman filter for atmospheric data assimilation. Monthly Weather Review 129, 123–137 (2001)
- [48] Chen, Y., Wiesel, A., Eldar, Y.C., Hero, A.O.: Shrinkage Algorithms for MMSE Covariance Estimation. IEEE Transactions on Signal Processing 58(10), 5016– 5029 (2010) https://doi.org/10.1109/TSP.2010.2053029
- [49] Nino-Ruiz, E.D., Sandu, A., Deng, X.: An ensemble Kalman filter implementation based on modified cholesky decomposition for inverse covariance matrix estimation. SIAM Journal on Scientific Computing 40(2), 867–886 (2018) https: //doi.org/10.1137/16M1097031 https://doi.org/10.1137/16M1097031

- [50] Stuart, A.M.: Inverse problems: A Bayesian perspective. Acta Numerica 19, 451– 559 (2010) https://doi.org/10.1017/S0962492910000061
- [51] Bui-Thanh, T., Ghattas, O., Martin, J., Stadler, G.: A computational framework for infinite-dimensional bayesian inverse problems part i: The linearized case, with application to global seismic inversion. SIAM Journal on Scientific Computing 35(6), 2494–2523 (2013) https://doi.org/10.1137/12089586X https://doi.org/10.1137/12089586X
- [52] Murua, A.: Partitioned Runge-Kutta methods for semi-explicit differentialalgebraic systems of index 2. Technical Report EHU-KZAA-IKT-196 (1996)
- [53] Arnold, M.: Half-explicit Runge-Kutta methods with explicit stages for differential-algebraic systems of index 2. BIT Numerical Mathematics 38(3), 415–438 (1998) https://doi.org/10.1007/BF02510252
- [54] De Jager, E.: On the origin of the Korteweg-de Vries equation. arXiv preprint math/0602661 (2006)
- [55] Miura, R.M., Gardner, C.S., Kruskal, M.D.: Korteweg-de Vries Equation and Generalizations. II. Existence of Conservation Laws and Constants of Motion. Journal of Mathematical Physics 9(8), 1204–1209 (2003) https://doi.org/10.1063/1.1664701 https://pubs.aip.org/aip/jmp/articlepdf/9/8/1204/8183110/1204_1_online.pdf
- [56] Ascher, U.M., McLachlan, R.I.: On symplectic and multisymplectic schemes for the KdV equation. Journal of Scientific Computing 25, 83–104 (2005)
- [57] Dutykh, D., Chhay, M., Fedele, F.: Geometric numerical schemes for the KdV equation. Computational Mathematics and Mathematical Physics 53(2), 221–236 (2013) https://doi.org/10.1134/S0965542513020103

- [58] Houtekamer, P.L., Mitchell, H.L.: Data assimilation using an ensemble Kalman filter technique . Monthly Weather Review 126, 796–811 (1998)
- [59] San, O., Staples, A.E., Wang, Z., Iliescu, T.: Approximate deconvolution large eddy simulation of a barotropic ocean circulation model. Ocean Modelling 40(2), 120–132 (2011)
- [60] Gottlieb, S., Shu, C.-W.: Total variation diminishing Runge-Kutta schemes. Mathematics of computation 67(221), 73–85 (1998)
- [61] Gaspari, G., Cohn, S.E.: Construction of correlation functions in two and three dimensions. Quarterly Journal of the Royal Meteorological Society 125, 723–757 (1999)