



Article submitted to journal

**Keywords:**

sampling, statistical estimation,  
thermostat, statistical fluid dynamics

**Author for correspondence:**

Keith W. Myerscough

e-mail: [keith@myerscough.nl](mailto:keith@myerscough.nl)

# Observation-based correction of dynamical models using thermostats

Keith W. Myerscough<sup>1</sup> Jason Frank<sup>2</sup>

Benedict Leimkuhler<sup>3</sup>

<sup>1</sup> Department of Computer Science, KU Leuven,  
3001 Leuven, Belgium

<sup>2</sup> Mathematical Institute, Utrecht University, P.O. Box  
80010, 3508 TA Utrecht, The Netherlands

<sup>3</sup> School of Mathematics, University of Edinburgh,  
James Clerk Maxwell Building, Kings Buildings  
Edinburgh EH9 3JZ

Models used in simulation may give accurate short-term trajectories but distort long term (statistical) properties. In this work, we augment a given approximate model with a control law (a “thermostat”) that gently perturbs the dynamical system to target a thermodynamic state consistent with a set of prescribed (possibly evolving) observations. As proof of concept, we provide an example involving a point vortex fluid model on the sphere, for which we show convergence of equilibrium quantities (in the stationary case) and the ability of the thermostat to dynamically track a transient state.

## 1. Introduction

In applications of modern computational science the underpinning physical laws (and equations of motion) are often well established yet the detailed behavior is unpredictable on long time scales due to the presence of deterministic chaos. Examples of this arise in molecular dynamics [Frenkel and Smit, 2002, Schlick, 2010], in polymer simulation [Ilg et al., 2002] and in the modelling of turbulent fluids in the atmosphere and ocean [Davidson, 2004, Holmes et al., 1998], where long simulations are used to extract statistical information (e.g. the statistics of rare transitions between basins in molecular dynamics, or slow relaxation processes in fluids).

© The Authors. Published by the Royal Society under the terms of the Creative Commons Attribution License <http://creativecommons.org/licenses/by/4.0/>, which permits unrestricted use, provided the original author and source are credited.

In this type of *dynamical sampling* the common requirement is that simulated paths are sufficiently accurate to calculate measures of dynamical mixing such as two-point temporal correlation functions [Hayes and Jackson, 2007].

In Hamiltonian systems such as molecular dynamics, it is common to run ensembles of microcanonical (i.e. constant energy) simulations. For such simulations, backward error analysis [Hairer et al., 2006, Leimkuhler and Reich, 2005] suggests that a long simulation should not be viewed as the approximation of a particular trajectory of the system but rather as an exact solution of a perturbed continuum process described by modified equations. In the case of dynamical sampling of complex systems, the statistics of simulation data are biased because they sample an invariant measure of the modified equations (rather than those of the target system), i.e., the time-discretization error induces an effective statistical bias.

Statistical bias may also arise due to spatial discretization. For example, in the setting of geophysical fluid dynamics, a comparison of discretizations of the quasi-geostrophic equations reveals that the long time mean potential vorticity field and pointwise fluctuation statistics are heavily dependent on discrete conservation laws such as energy, enstrophy, and material conservation of vorticity [Abramov and Majda, 2003, Dubinkina and Frank, 2007, 2010]. The discretization bias may be controlled by reducing timestep or by incorporating a Metropolis condition [Metropolis et al., 1953, Roberts and Tweedie, 1996], possibly increasing the computational cost. The use of a Metropolis condition can also increase the difficulty of computing accurate dynamical properties. The need for computations to be accurate with respect to the dynamical process as well as to the long term statistics thus poses difficult challenges to the simulator.

In this paper we consider a method of perturbing the dynamics of the system to correct for statistical bias. Our method is based on the concept of a thermostat, by which we mean a scheme based on a control law designed to facilitate sampling of a target probability density function (pdf). Although originally proposed in molecular dynamics (see, e.g. the Nosé-Hoover method [Nosé, 1984a,b]), thermostats can be extended to handle a wide variety of systems and have been used in a wide variety of applications. For example, in [Dubinkina et al., 2010], a thermostat was used as a model reduction technique for a vortex model of a fluid (suppressing the detailed interactions of a few strong vortices with a weak vortex field), an approach we draw on in the numerical experiments of this article. In another recent article [Ding et al., 2014], thermostats have been suggested as a means of sampling incompletely specified systems (with noisy gradients), with applications in machine learning for Bayesian inference in the “big data” context.

In the general modelling scenario, thermostats are significantly hampered by the need to specify a functional form of the pdf a priori. In this article, we assume that, instead of the pdf, what is actually available is a partial set of *expectations* of observables with respect to the unknown probability measure, which may arise from experiment or other types of modelling. In this setting, information theory (in particular entropy maximization [Jaynes, 1957a,b]) offers tools for constructing densities consistent with observations that are close to a given prior distribution (see [Giffin and Urmiezius, 2014, Majda and Gershgorin, 2011, Panagiotis, 2005] for examples in geophysical modelling, [Dama et al., 2013, Shell, 2008] for discussion in molecular modelling). Our algorithm for computing the parameters finds Lagrange multipliers (one for each observable) that modify the probability density, using iterative techniques like those in [Agmon et al., 1979, Davis and Gutiérrez, 2012, Haken, 2006], where each step of iteration involves ensemble simulation that reweights samples from the prior distribution appropriately. The density resulting from such an entropy maximization is then used as the basis for designing a thermostat that ensures the extended dynamical system samples the distribution and consequently exhibits the correct (long-term) statistical averages, while mildly perturbing the short-term dynamical behavior of the system. The complexity and generality of our iterative procedure (sampling→parameterization→sampling. . .) means that convergence has to be verified through computer experiment.

Entropy maximization also forms the basis for closure approximations in the works of Ilg et al. [2002], Samaey et al. [2011], where the result from entropy maximization is used in determining evolution equations for resolved macroscopic quantities. This approach is distinct in philosophy from the current work in which the result of entropy maximization is used to define a statistical ensemble and evolution equations for *microscopic* dynamics that are consistent with macroscopic observations.

While we motivate the method in the setting of equilibrium sampling of a fixed target distribution, we can easily apply our scheme to the case of a more general model with unknown or even undefined steady state. For this purpose, an ensemble of simulations are run in finite duration bursts, approximating the Lagrange multipliers via an iterative procedure based on the simulation history. This procedure allows entropy maximization to be performed with ensemble averages that evolve in time, providing a flexible method of nonlinear data assimilation.

The remainder of this article is organized as follows. In the next section (Section 2), we review the idea of a thermostat which is the key tool we employ to control the invariant distribution. Section 3 describes the combination of thermostats with the maximum entropy framework for correcting the density to reflect thermodynamic constraints. Section 4 describes the adaptive procedure. In Section 5, we apply and evaluate the method in the setting of a system of point vortices on the surface of a sphere, which represents a simple geophysical model with multiple statistically relevant first integrals.

## 2. Thermostats

We introduce the thermostat technique for generic Hamiltonian dynamical systems of the form

$$\frac{dy}{dt} = f(y) = B(y)\nabla H(y), \quad y(t) \in \mathcal{D}, \quad B(y) = -B(y)^T, \quad H(y) : \mathcal{D} \rightarrow \mathbb{R}, \quad (2.1)$$

possessing a divergence-free vector field<sup>1</sup>  $\nabla \cdot f \equiv 0$ . Invariance of the Hamiltonian  $H$  along solutions of (2.1) follows from  $\frac{d}{dt}H(y(t)) = \nabla H \cdot \frac{dy}{dt} = \nabla H \cdot B\nabla H = 0$ , due to skew-symmetry of  $B(y)$ . Additional first integrals may exist, denoted by  $I_\ell(y) : \nabla I_\ell \cdot f \equiv 0$ ,  $\ell = 1, \dots, L$ .

Thermostats were introduced in molecular dynamics to model the trajectories of molecules in a fluid at constant temperature. In the common statistical mechanical framework, it is assumed that the trajectories of a system of particles in thermal equilibrium with a reservoir at constant temperature sample the canonical distribution, with smooth invariant density proportional to  $\exp(-\beta H)$  where  $\beta$  is a parameter reciprocal to temperature scaled by Boltzmann's constant. The trajectories of the Hamiltonian system (2.1) are restricted to a level set of energy. Hence, to model a system at constant temperature, it is necessary to perturb the vector field to make trajectories ergodic with respect to the canonical distribution. The most common way of achieving this is by adding stochastic and dissipative terms satisfying a fluctuation-dissipation relation (Langevin dynamics), which is rigorously ergodic [Mattingly et al., 2002].

However, DelSole [DelSole, 2000] points out that additive stochastic forcing of trajectories leads to inaccurate dynamical quantities since autocorrelation functions are strongly perturbed. For smooth deterministic Hamiltonian dynamics, normalized velocity autocorrelation functions are of the form  $1 - c\tau^2$ ,  $c > 0$  in the zero-lag limit  $\tau \rightarrow 0$ , whereas the autocorrelation of a variable that is directly forced by white noise must take the form  $\exp(-\kappa\tau)$ ,  $\kappa > 0$  in the same limit. This implies that direct stochastic perturbation leads to auto-correlation functions that have nonzero slope and opposite curvature at zero lag [Leimkuhler et al., 2011].

An alternative approach, due to Nosé [Nosé, 1984a,b] and Hoover [Hoover, 1985], augments the phase space with an additional thermostat variable  $\xi(t)$  (driven by a differential equation) in order to ensure that the extended dynamics on  $\mathbf{R}^d \times \mathbf{R}$  preserves an equilibrium density whose marginal on  $\mathbf{R}^d$  is the target (e.g. Gibbs) density. The deterministic thermostats of Nosé and Hoover are not ergodic, but they can be combined with stochastic forcing of the thermostat

<sup>1</sup>The latter condition is automatic for systems (2.1) with constant  $B$ . Strictly speaking, the approach described here is applicable to any system with divergence-free vector field  $\nabla \cdot f \equiv 0$  possessing one or more first integrals.

variable  $\xi$ , leading to the so-called Nosé-Hoover-Langevin method [Samoletov et al., 2007]. For generic divergence-free dynamical systems, the Generalized Bulgac-Kusnezov (GBK) [Bulgac and Kusnezov, 1990, Leimkuhler, 2010] thermostat provides the augmented system:

$$dy = f(y)dt + \xi \varepsilon g(y)dt \quad (2.2a)$$

$$d\xi = \varepsilon h(y)dt - \gamma \xi dt + \sqrt{2\gamma} dw, \quad (2.2b)$$

where  $\varepsilon > 0$  and  $\gamma > 0$  are parameters,  $w(t)$  is a scalar Wiener process, and  $g$  and  $h$  are vector fields which we discuss in more detail below. Given a target density  $\rho(y) \propto \exp(-A(y))$ ,  $A: \mathcal{D} \rightarrow \mathbf{R}$ , denote the augmented product density by  $\tilde{\rho}(y, \xi) = \rho(y) \cdot \mu(\xi)$ , with  $\mu$  a univariate normal distribution with mean zero and standard deviation one. It is easily checked that  $\tilde{\rho}$  is stationary under the Fokker-Planck operator associated with (2.2) provided

$$h(y) = \nabla \cdot g - g \cdot \nabla A \quad (2.3)$$

and provided the ‘potential’  $A$  is only a function of  $y$  through the first integrals of  $f$ . The (extended) target measure  $\tilde{\rho}$  is ergodic if the vector fields  $f$  and  $g$  satisfy a Hörmander condition [Bajars et al., 2013]. Up to verification of the Hörmander condition, the choice of  $g$  is free. In the example of Section 5 we consider a specific form for  $g$ .

The parameter  $\varepsilon$  controls the strength of the thermostat relative to that of the vector field  $f$ . This affects the *rate* at which the invariant measure is sampled. It has been demonstrated in [Bajars et al., 2011, 2013, Leimkuhler et al., 2011] that GBK/NHL thermostating gives a weak perturbation of microcanonical dynamics: typical autocorrelation functions match the exact ones to leading order, i.e. have the form  $1 - c\tau^2 + O(\tau^3)$ , as  $\tau \rightarrow 0$ .

Designing a thermostat relies on knowing the functional form of the target distribution, which may not always be available. In this article we relax this requirement by constructing the target distribution iteratively and adaptively, based on Jaynes’ principle of least-biased ensemble prediction.

### 3. Bias correction method

Suppose that we are given a simplified dynamical model for the evolution of some projection, i.e. a ‘‘coarse graining’’ of the phase variables (coarse grained variables  $y(t) \in \mathbf{R}^d$ ) which is in the form (2.1). Although the original system is complex and its details unknown, we assume that we can obtain in some way (e.g. through measurement) a collection of observations of mean values of functions of the reduced variables. That is, there are functions  $C_k: \mathbf{R}^d \rightarrow \mathbf{R}$ ,  $k = 1, 2, \dots, K$  and given values  $c_k$ ,  $k = 1, 2, \dots, K$ , such that

$$c_k = \langle C_k(y) \rangle, \quad k = 1, \dots, K, \quad (3.1)$$

where  $\langle C_k(y) \rangle$  represents averaging with respect to the true, empirical invariant measure of the dynamical system. The use of a GBK thermostat (2.2) limits our implementation to observables that are functions of the conserved quantities  $\{H, I_\ell, \ell = 1, \dots, L\}$ , that is,  $C_k(y) = C_k(H(y), I_1(y), \dots, I_L(y))$ ,  $k = 1, \dots, K$ . This restriction is only due to the choice of thermostat, the entropy maximization has no such restriction. Our goal is to find a perturbed dynamical model for the reduced variables which (a) is compatible with the thermodynamic constraints (3.1), and (b) mildly perturbs the dynamics compared to those of the native model.

#### (a) Entropy maximization

Empirical information theory generalizes the principle of insufficient reason, by proposing the least-biased probability density consistent with a set of observations. This idea was first proposed by Jaynes [1957a,b] and is treated in the monographs of Dewar et al. [2014], Haken [2006] and in the context of geophysical fluids by Majda and Wang [Majda and Wang, 2006]. The least-biased density is defined as the probability density  $\rho(y)$  that maximizes the information entropy

functional  $\mathcal{S}[\rho] = - \int_{\mathcal{D}} \rho(y) \log \rho(y) dy$ , subject to a set of constraints given by observations. When  $\mathcal{D}$  is a compact set and there are no observations, the minimizer is the uniform density  $\rho \equiv |\mathcal{D}|^{-1}$ . The entropy  $\mathcal{S}$  is the unique measure of uncertainty that is positive valued, monotonically increasing as a function of uncertainty, and additive for independent random variables. With observable functions  $\{C_k(y) | k = 1, 2, \dots, K\}$  let

$$\mathbb{E}_{\rho} C_k = \int_{\mathcal{D}} C_k(y) \rho(y) dy \quad (3.2)$$

denote expectation in the (as yet undetermined) density  $\rho$ . Defining Lagrange multipliers  $\lambda_k$ ,  $k = 1, \dots, K$ , associated with the observables  $C_k$ , we obtain the solution  $\rho$  from a stationarity condition for the Euler-Lagrange functional

$$W(\rho, \lambda_1, \dots, \lambda_K) := \left[ \mathcal{S}[\hat{\rho}] - \sum_{k=1}^K \lambda_k (\mathbb{E}_{\hat{\rho}} C_k(y) - c_k) \right].$$

Conditions for a unique solution to exist are discussed in [Balian, 2006]. When it exists, the maximum entropy solution satisfies

$$\rho(y) = \lambda_0 \exp(-\lambda_1 C_1(y) - \dots - \lambda_K C_K(y)),$$

where  $\lambda_0$  is chosen to satisfy  $\int_{\mathcal{D}} \rho dy = 1$ , and  $\lambda_k$  is chosen such that  $\mathbb{E}_{\rho} C_k(y) = c_k$ .

In some cases, besides the observations, we may be given prior statistical information on the process  $y(t)$ . The Kullback-Leibler divergence, or relative entropy,

$$\mathcal{S}_{\pi}[\rho(y)] = \int \rho(y) \ln \frac{\rho(y)}{\pi(y)} dy$$

represents a (non-symmetric) distance between measures. Note the change of sign, which makes the relative entropy more like a distance measure; we stick to this convention. It quantifies the information lost in approximating  $\rho(y)$  by  $\pi(y)$ .

Suppose  $y$  is a random variable with distribution (law)  $y \sim \rho$ , where  $\rho$  is unknown. Suppose further, that we are given a prior distribution  $\pi$ , presumed to be close to  $\rho$ , and a set of  $K$  observations (3.1). Following Jaynes [Jaynes, 1957a,b], the least-biased distribution  $\rho$  consistent with the observations  $c_k$  and prior  $\pi$  solves the constrained minimization problem

$$\rho = \arg \min_{\hat{\rho}} \left[ \mathcal{S}_{\pi} - \lambda_0 \left( 1 - \int \hat{\rho}(y) dy \right) - \sum_{k=0}^K \lambda_k \left( c_k - \int C_k(y) \hat{\rho}(y) dy \right) \right],$$

where the  $\lambda_k$  are Lagrange multipliers to enforce the condition that the expectations (3.2) agree with the observations (3.1). The solution to the variational problem is

$$\rho(y) = \lambda_0 \exp(-\lambda_1 C_1(y) - \dots - \lambda_K C_K(y)) \pi(y), \quad (3.3)$$

where the Lagrange multipliers  $\lambda_k$  are chosen consistently with the observations (3.1) and  $\lambda_0$  is a normalization constant so that  $\rho$  is a probability density function.

Calculation of the Lagrange multipliers is discussed in [Agmon et al., 1979, Davis and Gutiérrez, 2012, Haken, 2006]. We use the following algorithm based on re-weighting. Assume we are given a sequence of samples  $y^n$ ,  $n = 1, \dots, N$ , distributed according to a known prior distribution  $\pi(y)$ , i.e.  $y^n \sim \pi$ . The expectation under  $\pi(y)$  of a function  $\Phi(y)$  has the consistent and unbiased estimator

$$\hat{\Phi}^{\pi} = \frac{1}{N} \sum_{n=1}^N \Phi(y^n).$$

Given the posterior distribution  $\rho(y)$  of the form (3.3), compute the expectation  $\mathbb{E}_{\rho} \Phi$  by re-weighting of the integral

$$\mathbb{E}_{\rho} \Phi = \int \Phi(y) \rho(y) dy = \lambda_0 \int \Phi(y) e^{-\sum_{i=1}^K \lambda_i C_i(y)} \pi(y) dy = \lambda_0 \mathbb{E}_{\pi} \{ \Phi(y) \lambda_0 e^{-\sum_{i=1}^K \lambda_i C_i(y)} \},$$

yielding an unbiased estimator for  $\mathbb{E}_\rho \Phi$  given by

$$\widehat{\Phi}^\rho = \frac{\lambda_0}{N} \sum_{n=1}^N \Phi(y^n) e^{-\sum_{i=1}^K \lambda_i C_i(y^n)}.$$

We wish to ensure that the estimators  $\widehat{C}_k^\rho$  for observables  $C_k$  match their empirical values  $c_k$ , i.e.

$$c_k = \widehat{C}_k^\rho = \frac{\lambda_0}{N} \sum_{n=1}^N C_k(y^n) e^{-\sum_{i=1}^K \lambda_i C_i(y^n)}, \quad k = 1, \dots, K.$$

We can use this fact to define a Newton-Raphson iteration to determine the Lagrange multipliers  $\lambda_k$ . Define the residual  $r$  with components

$$r_k(\lambda) = c_k - \frac{\lambda_0}{N} \sum_{n=1}^N C_k(y^n) e^{-\sum_{i=1}^K \lambda_i C_i(y^n)}, \quad k = 1, \dots, K,$$

with  $\lambda = (\lambda_1, \dots, \lambda_K)$  and  $r = (r_1(\lambda), \dots, r_K(\lambda))$ . Note that  $\lambda_0$  can be viewed as a function of  $\lambda_1, \lambda_2, \dots, \lambda_K$  chosen from the normalization condition, i.e.,

$$\lambda_0 = \left[ \sum_{n=1}^N e^{-\sum_{j=1}^K \lambda_j C_j(y^n)} \right]^{-1}.$$

The Jacobian matrix  $J = (J_{kj})$  of the vector function  $r$  is determined as

$$J_{kj}(\lambda) := \frac{\partial r_k}{\partial \lambda_j} = \frac{\lambda_0}{N} \sum_{n=1}^N C_k(y^n) C_j(y^n) e^{-\sum_{i=1}^K \lambda_i C_i(y^n)} \quad j, k = 1, \dots, K.$$

The iteration then proceeds as  $\lambda^{\alpha+1} \leftarrow \lambda^\alpha - J^{-1}(\lambda^\alpha) r(\lambda^\alpha)$ .

## 4. Adaptive determination of Lagrange multipliers

When the observations deviate far from the prior, the Lagrange multipliers take larger values and the weights of the samples diverge, i.e. some weights become very large and others become very small. This results in a larger variance on the estimates and consequently in poorer results for the Lagrange multipliers. Moreover, the approach of the previous section excludes applications where (i) the statistical knowledge is expected to improve as the simulation progresses, (ii) the average observables are known to vary slowly with time, or (iii) it is unfeasible to construct a large enough ensemble distributed in the prior. For these cases we consider using the simulation data of a small ensemble (propagated in short bursts of  $M$  timesteps) for updating the Lagrange multipliers for mean observation data. This results in an adaptive algorithm for obtaining the Lagrange multipliers on-the-fly during simulation. In this situation, the sample set used is drawn from a distribution closer to the target, resulting in better estimations for the residuals and Jacobian.

Let us assume we have a (moderately sized) ensemble of initial conditions according to some distribution

$$\rho_0(y) = \pi(y) \exp \left( - \sum_{k=1}^K \lambda_k^0 C_k(y) \right) \quad (4.1)$$

In our numerical experiments, we assume  $\lambda_k^0 = 0, \forall k$  and hence  $\rho_0 = \pi$ , but we present the general form here nevertheless. Rather than first computing the Lagrange multipliers corresponding to observations and then running the simulation over the time interval of interest, we perform a short burst of the thermostatted simulation for  $P$  ensemble members, using the Lagrange multipliers  $\lambda_k^0$  in the thermostat. The system states at the end of this burst, denoted by  $y^P(M \Delta t)$

for ensemble member  $p$  at time  $M\Delta t$ , are used for an estimator of the observables in the distribution  $\rho_0$ :

$$\hat{C}_k^{(1)} = \frac{1}{P} \sum_{p=1}^P C_k(y^p(M\Delta t)) \lambda_0^1 \exp\left(\sum_{i=1}^K (\lambda_i^0 - \lambda_i^1) C_i(y^p(\Delta t))\right), \quad k = 1, \dots, K, \quad (4.2)$$

where the superscript (1) indicates that it is an estimator for a distribution with Lagrange multipliers  $\lambda_k^1$ . Note that the samples being distributed in  $\rho_0$  and not in  $\pi$  leads to the appearance of the Lagrange multipliers  $\lambda_k^0$ , this is a simple importance sampling procedure. A Newton-Raphson procedure similar to that in Section 3 is then used to find  $\lambda_k^1$  using the estimator (4.2). The resulting multipliers are used in the second burst of thermostatted simulations, sampling  $\rho_1(y) = \pi(y) \exp\left(-\sum_{k=1}^K \lambda_k^1 C_k(y)\right)$ . Now, crucially, this estimator from the ensemble after two simulation bursts is combined with that at the end of the first burst. As both estimators are unbiased, we may use any weighted sum of the two (with the sum of the weights equal to one) to find another estimator. For simplicity we weight both estimators equally. We write the combination of estimators for the general case after  $L$  bursts of the simulation, using the system states at the end of each previous burst, that is  $y^p(\ell M \Delta t)$  for  $\ell = 1, 2, \dots, L$ , as an estimator for the densities  $\rho_{\ell-1}$  with the Lagrange multipliers  $\lambda_k^{\ell-1}$ . This estimator takes the form

$$\hat{C}_k^{(L)} = \frac{1}{LP} \sum_{p=1}^P \sum_{\ell=1}^L C_k(y^p(\ell M \Delta t)) \lambda_0^{\ell-1} \exp\left(\sum_{i=1}^K (\lambda_i^{\ell-1} - \lambda_i^L) C_j(y^p(\ell M \Delta t))\right), \quad k = 1, \dots, K. \quad (4.3)$$

As with the non-adaptive scheme, a Newton-Raphson procedure is used to find Lagrange multipliers such that  $\widehat{C}_k^{(L)} = c_k$  with the residuals

$$r_k^L = \widehat{C}_k^{(L)} - c_k \quad k = 1, \dots, K, \quad (4.4)$$

and the gradients

$$\frac{\partial r_k^m}{\partial \lambda_j^m} = \frac{1}{mP} \sum_{p=1}^P \sum_{\ell=0}^{m-1} C_k(y_{\ell M}^p) C_j(y_{\ell M}^p) \lambda_0^\ell \exp\left(\sum_{i=1}^K (\lambda_i^\ell - \lambda_i^m) C_i(y_{\ell M}^p)\right), \quad j, k = 1, \dots, K. \quad (4.5)$$

By performing this iteration after each burst, and using the result for the subsequent burst, the Lagrange multipliers are computed “on-the-fly” without the need for constructing a large ensemble in the prior.

Let us consider two limiting cases to justify the approach. (1) By increasing the number of ensemble members  $P$ , one recovers the non-adaptive scheme, where the Lagrange multipliers are found based on the initial ensemble after a single burst. On the other hand, (2) if we have a modest ensemble  $P$ , it is necessary that the burst is long enough for the thermostatted dynamics with the adaptive parameters to equilibrate on this time scale. The error due to the transient of the thermostat is removed in the limit of many time steps per burst, i.e. as  $M \rightarrow \infty$ . We present numerical verification of the practical use of the scheme in a specific application setting in Section (c).

### (a) Adaptive algorithm

There are two important modifications to the algorithm described above that are included in the numerical implementation of this method:

- If the Lagrange multipliers change rapidly, the thermostat may require a long time to equilibrate. This requires a larger value for  $M$ , increasing the simulation time required before including new samples. The effect is noticeable at the beginning of a simulation, due to two factors: (i) the small sample size leads to inaccurate expectations for the

observables, and (ii) the initial values for Lagrange multipliers may be far from their equilibrium. By limiting the rate of change of the Lagrange multipliers, these problems are circumvented.

- In equations (4.4) and (4.5) all previous values  $\lambda_k^\ell$  are included. In a long simulation, this leads to a growing computational demand. By taking only a fixed number ( $q$ ) of recent steps the cost is reduced. In the case that the initial samples cannot accurately be drawn from the prior, this has the further advantage that these inaccuracies are eventually forgotten.

The resulting method is summarized in Algorithm 1.

```

Given initial conditions according to prior  $\pi(y)$ 
Set initial Lagrange multipliers to zero
for  $m \leftarrow 1$  to  $n$  do
  for  $j \leftarrow 1$  to  $M$  do
    | advance simulation one time step using current value for the Lagrange multipliers
  end
  store relevant simulation observables for time step  $im$ .
  while  $|\langle C_k(y) \rangle_\lambda - c_k| > \text{tolerance}$  do
    | compute residual using reweighted samples at times  $M \times \max(m - q, 0), \dots, mM$ 
    | compute residual gradient using the same data
    | update Lagrange multiplier estimation
  end
  limit the change in Lagrange multiplier (if necessary)
end

```

**Algorithm 1:** Adaptive determination of Lagrange multipliers “on-the-fly”

## 5. Application to reduced modelling of point vortices

Here we apply the least-biased correction method to a system of point vortices on the sphere, which has a Poisson structure and multiple conserved quantities, including total energy and angular momentum and a set of Casimirs.

### (a) Point vortex system

Point vortices represent singular solutions to two-dimensional, incompressible fluid flow in which each vortex, positive or negative, is centred at a specified position. Point vortices have been studied extensively in the books by Cottet and Koumoutsakos [2000], Majda and Bertozzi [2002], Newton [2001]. We work on a spherical domain Jamalodeen and Newton [2006], Laurent-Polz [2005], Myerscough and Frank [2016], Newton and Sakajo [2007], Newton and Shokraneh [2006], viewing the vortex position as a vector in  $\mathbf{R}^3$  with unit norm. The equations of motion have a Lie-Poisson structure

$$\Gamma_i \dot{\mathbf{x}}_i = \mathbf{x}_i \times \nabla_{\mathbf{x}_i} H \quad i = 1, 2, \dots, M, \quad (5.1)$$

where the Hamiltonian is given by

$$H = - \sum_{i=1}^M \sum_{j=1}^{i-1} \frac{\Gamma_i \Gamma_j}{4\pi} \ln(2 - 2\mathbf{x}_i \cdot \mathbf{x}_j)$$

and each  $\Gamma_i$  represents the circulation of a single point vortex.

By introducing  $\mathbf{y} = (\mathbf{x}_1^T, \mathbf{x}_2^T, \dots, \mathbf{x}_M^T)^T$ , equation (5.1) has the form (2.1) with the block-diagonal structure matrix

$$B(\mathbf{y}) = \begin{bmatrix} \Gamma_1^{-1} \hat{\mathbf{x}}_1 & & & \\ & \Gamma_2^{-1} \hat{\mathbf{x}}_2 & & \\ & & \ddots & \\ & & & \Gamma_M^{-1} \hat{\mathbf{x}}_M \end{bmatrix},$$

where  $\hat{\mathbf{x}}_i$  is the  $3 \times 3$  skew-matrix satisfying  $\hat{\mathbf{x}}_i \mathbf{a} := \mathbf{x}_i \times \mathbf{a}$ , for all  $\mathbf{a} \in \mathbf{R}^3$ . The Poisson bracket for the system is given equivalently by

$$\{F, G\} = \sum_{i=1}^M \frac{1}{\Gamma_i} \nabla_{\mathbf{x}_i} F \cdot (\mathbf{x}_i \times \nabla_{\mathbf{x}_i} G) \quad \text{or} \quad \{F, G\} = \nabla F(\mathbf{y})^T B(\mathbf{y}) \nabla G(\mathbf{y}).$$

This Poisson structure generalizes the rigid body Poisson structure and also occurs in ferromagnetic spin lattices [Faddeev and Takhtajan, 1987, Frank, 2004, Frank et al., 1997] and elastic rods (e.g. [Kehrbaum and Maddocks, 1997]).

The vortex positions are defined in Cartesian coordinates, but initial positions  $\mathbf{x}_i(0)$  are chosen on the sphere. Because each  $|\mathbf{x}_i|$  is a Casimir of the Poisson bracket, the vortices remain on the sphere. Furthermore, the rotational symmetry of the sphere gives rise to three Noether momenta, which are expressed by the angular momentum vector

$$\mathbf{J} = \sum_{i=1}^M \Gamma_i \mathbf{x}_i. \quad (5.2)$$

The GBK thermostat (2.2) is only applicable to divergence free systems  $\nabla \cdot f \equiv 0$ . It is straightforward to check that this condition holds for the spherical point vortex model.

A numerical integrator for the point vortex system is constructed as in [Myerscough and Frank, 2016] based on splitting the differential equations into integrable subproblems (see related ideas in [Patrick, 2000, Zhang and Qin, 1993]). The backward error analysis of symplectic integrators has an analogous development for Poisson systems, implying approximate conservation of the Hamiltonian [Hairer et al., 2006]. The Casimirs ( $|\mathbf{x}_i| \equiv 1$ ) and the momentum  $\mathbf{J} = \sum \Gamma_i \mathbf{x}_i$  are exactly preserved by the integrator.

## (b) Thermostat perturbation vector

The choice of the perturbation vector field  $g$  in (2.2) is flexible, but the vector fields  $f$  and  $g$  should satisfy a Hörmander condition [Bajars et al., 2013]. It is vital that the Casimirs should be respected by the perturbation vector field. Double-bracket dissipation [Bloch et al., 1996] preserves Casimirs of the original system and is thus an ideal candidate for  $g$ :

$$\tilde{g}_i(\mathbf{x}_i) = \sum_{j \neq i} \mathbf{x}_i \times \mathbf{x}_i \times \frac{\Gamma_j}{4\pi} \frac{\mathbf{x}_j}{1 - \mathbf{x}_i \cdot \mathbf{x}_j}. \quad (5.3)$$

This “double-bracket thermostat” also has the advantage that it can be split into pairwise interactions along with the original dynamics  $f$ , simplifying computation. The denominator in (5.3) leads to a stiff differential equation when like-signed vortices approach one another, restricting the step size of an explicit splitting method. For this reason we use a slightly modified scheme (but still Casimir-preserving) defined by

$$g_i(\mathbf{x}_i) = \sum_{j \neq i} \mathbf{x}_i \times \mathbf{x}_i \times \frac{\Gamma_j}{4\pi} \mathbf{x}_j. \quad (5.4)$$

Details of the numerical integration of these dynamics may be found in the appendix.

The thermostat (2.2) is designed to sample a target density  $\rho(y) \propto e^{-A(y)}$  on the phase space of  $y$ . The thermostat variable  $\xi$  is normally distributed, yielding the extended distribution

$\rho \propto e^{-A(y) - \frac{1}{2}\xi^2}$ . The perturbation vector field  $g$  must additionally ensure that the thermostatted system is ergodic in the target density. Because the target measure is positive for all open sets on the phase space, hypoellipticity of the Fokker-Planck equation associated with (2.2) is sufficient to prove uniqueness of the invariant measure [Kliemann, 1987]. Hypoellipticity follows from Hörmander's controllability condition [Rey-Bellet, 2006]. The condition has been tailored to GBK thermostats in [Bajars et al., 2013], but it is difficult to verify in practice. Here we verify using simulation that single trajectories have statistics that agree with the target distribution.

### (c) Maximum entropy model

To apply the methods proposed in Section 3 in the setting of a reduced model for point vortices, we use point vortices distributed evenly over the surface of the sphere as the prior  $\pi$ . The use of a thermostat for the bias correction limits us to using observables that are function of the first integrals. Furthermore we assume the angular momentum has no directional preference, a choice we motivate in the numerical comparisons of Section 6. In that case, only the magnitude  $|\mathbf{J}|$  of the momentum vector needs to be considered.

To accurately reproduce statistics from a full model with a moderate number of point vortices, it is necessary to modify the canonical density with a term quadratic in the Hamiltonian, that is, a density of the form  $\rho(y) \propto \exp(-\beta H(y) - \gamma H(y)^2)$  [Dubinkina et al., 2010]. Motivated by [Dubinkina et al., 2010], we choose observations that include linear and quadratic functions of  $H$  and  $|\mathbf{J}|^2$ .

We consider the following set of observables:

$$C_1 = H, \quad C_2 = |\mathbf{J}|^2, \quad C_3 = H^2, \quad C_4 = |\mathbf{J}|^4, \quad C_5 = H|\mathbf{J}|^2, \quad (5.5)$$

and denote the corresponding Lagrange multipliers by  $\beta_H, \beta_J, \gamma_H, \gamma_J, \gamma_{HJ}$ .

The least-biased density consistent with observations of the  $\mathbb{E}C_k$  is given by

$$\tilde{\rho}(H) = e^{-\beta_H H - \beta_J |\mathbf{J}|^2 - \gamma_H H^2 - \gamma_J |\mathbf{J}|^4 - \gamma_{HJ} H |\mathbf{J}|^2}, \quad (5.6)$$

where the Lagrange multipliers are to be found via the procedure detailed in Section (a) or that of Section 4.

## 6. Numerical comparison

Here we apply Algorithm 1 to a reduced model of point vortices similar to the configuration used in [Bühler, 2002, Dubinkina et al., 2010]. We distinguish between three models. The *full model* consists of a system (5.1) of  $M_{\text{full}} = 288$  point vortices, of which 8 strong vortices of circulation  $\Gamma_j = \pm 1$  and 280 weak vortices of circulation  $\Gamma_j = \pm \frac{1}{5}$ . Both strong and weak classes are comprised of equal numbers of positively and negatively oriented point vortices. The *reduced model* consists of (5.1) with just  $M = 8$  strong vortices. Finally, the *corrected model* consists of a thermostatted system (2.2) with unperturbed vector field  $f$  given by (5.1) for  $M = 8$  strong vortices, perturbation vector field  $g$  given by (5.4), and equilibrium measure defined by the least-biased density (5.6). Additionally, we compare with Metropolis-Hastings samples from the least-biased density (5.6) to help distinguish between errors incurred due to the maximum-entropy model and those due to sampling bias of the (discretized) thermostatted dynamics.

We run five long simulations of the full model with total energies chosen from the set  $H_{\text{full}} \in \{-2, -1, 0, 1, 2\}$ . The total angular momentum vector is fixed at  $\mathbf{J}_{\text{full}} = \mathbf{0}$  such that there is no directional preference for the momentum of the reduced model embedded in the full model. For each run we determine the time averages of the observables (5.5) for the subset of strong vortices. When computing the Hamiltonian  $H$  of this subsystem, we include only the internal coupling between strong vortices. The time averages are tabulated in Table 1).

Using these averages, we compute the Lagrange multipliers using the algorithm described in Section 4 with prior distribution  $\pi$  taken to be uniform on the sphere. The Lagrange multipliers are

Table 1: Full model observations and (in parentheses) corrected values of first integrals

	$\langle H \rangle$	$\langle  \mathbf{J} ^2 \rangle$	$\langle H^2 \rangle$	$\langle  \mathbf{J} ^4 \rangle$	$\langle H \mathbf{J} ^2 \rangle$
$H_{\text{full}} = -2$	-0.33 (-0.38)	4.59 (4.45)	0.22 (0.23)	-0.63 (-0.98)	34.58 (31.45)
$H_{\text{full}} = -1$	-0.11 (-0.18)	4.78 (4.68)	0.10 (0.12)	0.38 (-0.01)	37.55 (35.88)
$H_{\text{full}} = 0$	0.02 (-0.04)	4.63 (4.56)	0.08 (0.08)	0.90 (0.60)	35.26 (34.30)
$H_{\text{full}} = 1$	0.17 (0.15)	4.74 (4.75)	0.13 (0.12)	1.73 (1.61)	37.76 (37.44)
$H_{\text{full}} = 2$	0.31 (0.28)	4.87 (5.00)	0.22 (0.21)	2.49 (2.46)	39.26 (41.74)

also recorded in Table 2. The magnitude of  $\gamma_{\{H,J,HJ\}}$  indicates that including these observables impacts the resulting density. In other words, all observables add information to the least-biased density, with the possible exception of the lowest energy case with  $H_{\text{full}} = -2$ .

We then run simulations of the corrected model using the computed parameters. Table 1 also records expectations from the thermostat-corrected model.

Table 2: Lagrange multipliers for each energy level.

	$\beta_H$	$\beta_J$	$\gamma_H$	$\gamma_J$	$\gamma_{HJ}$
$H_{\text{full}} = -2$	5.98	-0.20	0.69	$0.41 \times 10^{-3}$	-0.04
$H_{\text{full}} = -1$	2.89	-0.03	2.67	$9.77 \times 10^{-3}$	-0.33
$H_{\text{full}} = 0$	-0.76	0.20	3.38	$9.97 \times 10^{-3}$	-0.37
$H_{\text{full}} = 1$	-3.54	0.37	4.29	$15.31 \times 10^{-3}$	-0.54
$H_{\text{full}} = 2$	-6.42	0.53	4.45	$14.05 \times 10^{-3}$	-0.51

By analogy with canonical statistical mechanics, we may think of the weak vortices that are ignored in the reduced model as forming a reservoir with which our reduced model exchanges energy and angular momentum. Experience with canonical statistical mechanics of point vortices in the plane [Bühler, 2002, Dubinkina et al., 2010] suggests that for small reservoir sizes the canonical distribution must be modified with higher order terms to agree with the full system statistics. Table 3 gives the multipliers for different numbers  $M_{\text{full}}$  of vortices in the full system, confirming that those corresponding to  $\gamma_H$ ,  $\gamma_J$  and  $\gamma_{HJ}$  are more significant for smaller  $M_{\text{full}}$ ; as the number of weak vortices grows, the strong vortex system approaches the canonical density.

Table 3: Lagrange multipliers as a function of  $M_{\text{full}}$ , all for  $H_{\text{full}} = 0$ .

	$\beta_H$	$\beta_J$	$\gamma_H$	$\gamma_J$	$\gamma_{HJ}$
$M_{\text{full}} = 36$	-1.51	1.35	27.75	$117.15 \times 10^{-3}$	-3.07
$M_{\text{full}} = 72$	-4.27	0.82	8.71	$37.79 \times 10^{-3}$	-1.12
$M_{\text{full}} = 144$	-0.97	0.32	6.70	$25.80 \times 10^{-3}$	-0.82
$M_{\text{full}} = 288$	-0.76	0.20	3.38	$9.97 \times 10^{-3}$	-0.37
$M_{\text{full}} = 576$	-1.09	0.13	0.87	$3.08 \times 10^{-3}$	-0.10

**Remark 6.1.** The form of the distribution (5.6) implies there is a nonzero probability for any configuration of the strong vortex system with finite energy and momentum. When considering the strong vortex system in contact with a reservoir of weak vortices, its energy is bounded by the energy in the total system plus or minus how much energy can be stored in the reservoir and in the interactions between strong and weak vortices. For the momentum, a similar condition holds, but here there are no interaction terms. If it is not possible for the reservoir to add or subtract sufficient energy from the reservoir, it will not be possible for the subsystem of strong vortices to attain all possible configuration with nonzero probability.

For the energy it is readily verified that as long as the reservoir contains three vortices not all of equal sign, the reservoir may supply or remove an arbitrary amount of energy, since close approaches of two like-signed (resp. opposite-signed) vortices produces arbitrarily large positive (resp. negative) energies.

For the momentum, the result is more interesting. The system of  $M$  strong vortices, all with strength  $\pm\Gamma_{\text{strong}}$ , may exist in configurations with an angular momentum  $\sup |\mathbf{J}_{\text{red.}}| = M\Gamma_{\text{strong}}$ . Note the supremum, as the configuration with exactly  $\sup |\mathbf{J}_{\text{red.}}| = M\Gamma_{\text{strong}}$  will have infinite energy. For the reservoir it holds that  $\sup |\mathbf{J}_{\text{weak}}| \leq (N - M)\Gamma_{\text{weak}}$ . For the reservoir to supply sufficient angular momentum it is necessary that

$$M\Gamma_{\text{strong}} \leq (M_{\text{full}} - M)\Gamma_{\text{weak}} \Leftrightarrow \frac{M_{\text{full}} - M}{M} \geq \frac{\Gamma_{\text{strong}}}{\Gamma_{\text{weak}}}.$$

In the thermal bath simulations discussed in this section  $M = 8$  and  $\Gamma_{\text{strong}}/\Gamma_{\text{weak}} = 5$ , this means  $M_{\text{full}}$  should satisfy  $M_{\text{full}} \geq 48$ . The smallest system considered ( $M_{\text{full}} = 36$ ) does not, explaining its eccentric parameter values in Table 3.

### (a) Equilibrium results

In this section we compare statistical properties of the corrected model with those of the full and reduced models. In Figures 1–3 we show histograms of a number of solution features for the 8 vortex model: the distributions of  $H$  and  $|\mathbf{J}|$ , as well as typical distances between like- and opposite-signed vortices, a metric also used by Bühler [Bühler, 2002]. In each histogram, the statistics corresponding to the strong vortices in the full model, the reduced model, thermostat-corrected reduced model, and Metropolis-Hastings samples are displayed. Figures 1, 2 and 3 correspond to approximate total energies  $H_{\text{full}} \approx -2, 0$  and  $2$ , respectively.

The full and reduced model simulations are performed with a time step of  $5 \times 10^{-3}$  and run up to  $T = 5 \times 10^4$ , taking  $10^5$  samples spaced evenly in time. For the Metropolis-Hastings method we use  $10^6$  samples. The same figures also show results from the thermostatted system (dash-dot lines), run with a time step of  $10^{-3}$  up to  $T = 10^6$ , taking  $10^6$  samples. The parameters in (2.2) were set to be  $\varepsilon = 10$  and  $\gamma = 0.1$ . These results confirm that the thermostatted system samples the least-biased density closely.

The reduced model is Hamiltonian and the Poisson integrator ensures that the energy is conserved with a standard deviation of order  $10^{-3}$  and the angular momentum constant to machine precision. Both cases correspond to approximate delta-distributions in the upper histograms in Figures 1–3. Note that due to the high skewness of the distribution for  $|\mathbf{J}|$ , the observed mean differs significantly from the median and mode, implying some ambiguity in choosing the angular momentum for an appropriate initial condition for the reduced model.

A simple Hamiltonian reduced model is incapable of sampling the energy and angular momentum spectra, as these are first integrals, hence the reduced model shows significant bias in statistics such as vortex separation. On the other hand, the thermostat-corrected model faithfully samples the least-biased probability density, as indicated by the good agreement in the histograms of the corrected model and Metropolis-Hastings samples. The least-biased density approximates the strong-vortex statistics well, particularly in the negative to moderate total energy regime. At large positive total energies, the strong vortex energy and angular momentum distributions are still well-represented by the least-biased PDF, but some bias in the vortex separations can be observed. The closeness of the thermostat results to those from the Metropolis-Hastings sampling indicate the error lies in the choice of least-biased density, not in the thermostat sampling.

## (b) Dynamic consistency

The results in the previous section confirm that the thermostatted simulations lead to equilibrium distributions of observables  $H$  and  $|J|$  similar to those of the full system. In this section we address the issue of the degree to which our equilibrium correction mechanism disturbs dynamics, as encoded in autocorrelation functions and diffusivity. Diffusivity was considered by [Chavanis, 2001] for a system of identical point vortices and by [Cotter and Pavliotis, 2009] for a wide array of problems with scale separation. We emphasize that the values of the thermostat parameters  $\varepsilon$  and  $\gamma$  have no impact on the equilibrium statistics presented in the previous section, and only affect the rate at which the least-biased PDF is sampled. Parameters with a larger deviation from the unperturbed dynamics lead to faster equilibration of the distribution, giving rise to a modeling choice.

### (i) Autocorrelation functions

Given a sequence of  $L$  equally spaced observation times  $t_i \in [0, T]$  for  $i \in [0, L]$ , and the values of the relevant observable (in our case vortex position)  $u_i = u(t_i)$  at those times, the discrete autocorrelation function is defined by

$$\nu_i^u = \frac{1}{L-i} \sum_{j=i}^L u(t_j)u(t_{j-i}).$$

A normalized autocorrelation function  $\hat{\nu}^u$  is given by dividing each  $\nu_i^u$  by  $\nu_0^u$ , i.e.  $\hat{\nu}_i^u = \nu_i^u / \nu_0^u$ .

We average the autocorrelation function over all  $3M$  (strong) vortex coordinates  $\{x_m, y_m, z_m\}_{m=1..M}$ . Three symmetries in the problem justify this averaging: the vortex numbering is arbitrary; the choice of reference frame is arbitrary and the sign of the vortices appears in the dynamics as a reversal of time, to which the autocorrelation is insensitive. Additionally, the observables  $H$  and  $|J|$  are isotropic.

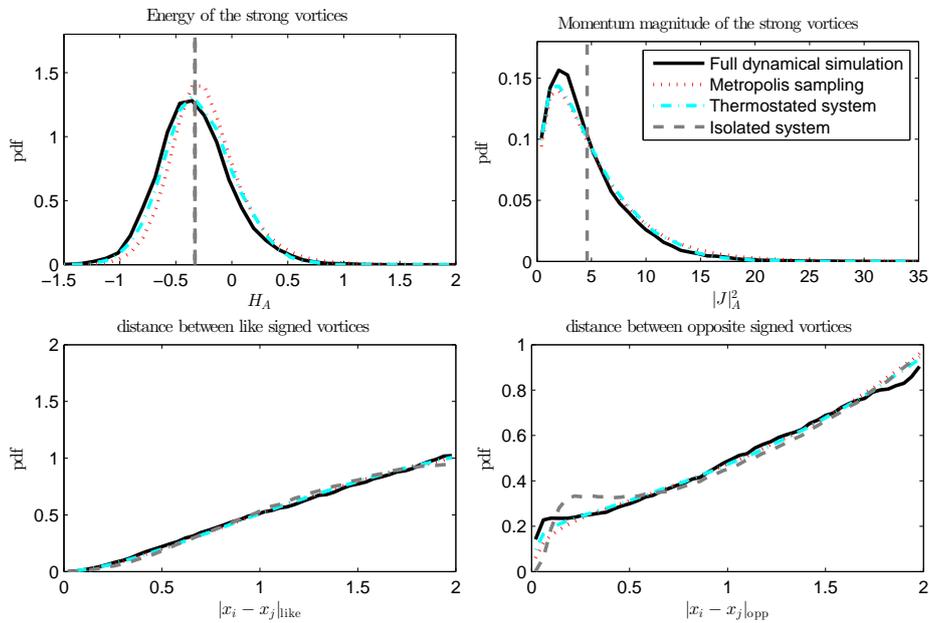


Figure 1: Histograms for  $H_{\text{full}} \approx -2$ . The upper left and right panels compare strong vortex energy and angular momentum. The lower left (resp. right) panel compares the distance between like (resp. opposite) signed strong vortices. The parameters are specified in the text.

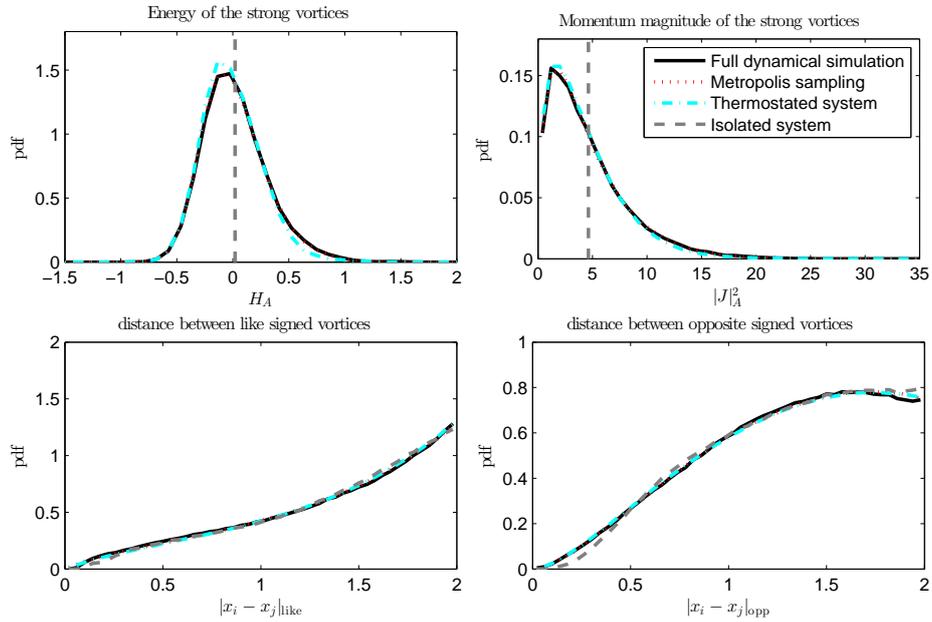


Figure 2: Histograms for  $H_{\text{full}} \approx 0$ . Same panel layout as Figure 1.

Furthermore we ensure that the phase space is well sampled by averaging the autocorrelation functions over an ensemble of  $P$  solutions. We then find the average autocorrelation function

$$\nu_i = \frac{1}{3MP(L-i)} \sum_{p=1}^P \sum_{m=1}^M \sum_{j=i}^L x_m^p(t_j) x_m^p(t_{j-i}) + y_m^p(t_j) y_m^p(t_{j-i}) + z_m^p(t_j) z_m^p(t_{j-i}) \quad (6.1)$$

where a superscript  $p$  represents the solution from ensemble member  $p$ . The normalized average autocorrelation function is given by

$$\hat{\nu}_i = \frac{1}{MP(L-i)} \sum_{p=1}^P \sum_{m=1}^M \sum_{j=i}^L \mathbf{x}_m^p(t_j) \cdot \mathbf{x}_m^p(t_{j-i}), \quad (6.2)$$

Where we have used that the Casimirs  $C_i = \mathbf{x}_i(t) \cdot \mathbf{x}_i(t) = 1 \forall i, t$  are conserved exactly, also in the distretized equations.

In Figure 4 we compare auto-correlation functions for the strong vortices in the full and reduced models as well as for the thermostat-corrected model over a range of parameters  $\varepsilon$  and  $\gamma$ . The thick solid black line represents the result for an (unthermostatted) system in contact with 280 weak ( $\Gamma_B = \pm \frac{1}{5}$ ) vortices, with a total energy  $H_{\text{full}} = 0$ . The results present the average over an ensemble of 1000 runs. For each simulation the initial placement of each strong vortex was taken uniformly over the sphere and the weak vortices were placed such that the full system satisfied  $H_{\text{full}} = 0$  and  $\mathbf{J}_{\text{full}} = 0$ . The thick dashed black line represents the results for an ensemble of simulations of the isolated system, with everything else unchanged. The other lines represent results for thermostatted simulations using the Lagrange multipliers as given in Table 2 for the case of  $H = 0$ . As expected, when the thermostat is weak the autocorrelation functions are similar to those of the isolated system. The stronger thermostat parameters presented here match the autocorrelation well for a short time period, but all result in excessive decorrelation for longer lag times. In the following Section on the diffusivity constant, we discuss the performance of different parameter values in more detail.

## (ii) Diffusivity

For general multiscale dynamical systems with a separation of slow and fast dynamics, it is often desirable to model fast forces by a diffusion process, resulting in stochastic differential equation of the form [Hornung, 1997, Pavliotis and Stuart, 2008]

$$dX = f(X)dt + K(X)dW,$$

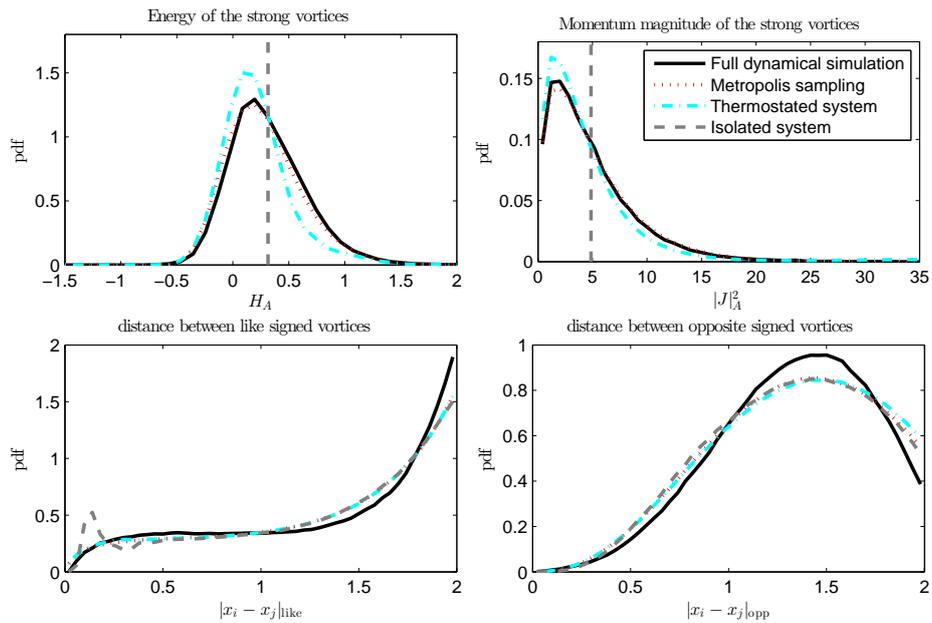


Figure 3: Histograms for  $H_{\text{full}} \approx 2$ . Same panel layout as Figure 1.

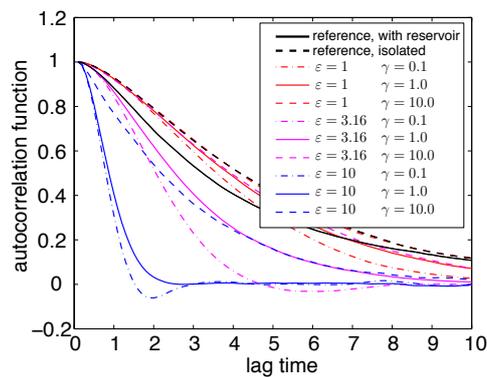


Figure 4: Comparison of autocorrelation of the vortex coordinates. The bold lines are two reference cases: the full model (solid) and the reduced model (dashed). The thin lines indicate autocorrelation functions of the thermostat-corrected model for indicated values of parameters  $\varepsilon$  and  $\gamma$ .

where  $f$  represents the slow dynamics,  $W$  is a Wiener process and  $K$  is the diffusivity. The value of the diffusivity  $K$  can be estimated by sampling solutions to the original, multiscale, problem and applying Kubo's formula

$$K_{\Delta\tau} = \frac{\langle \Delta X \Delta X \rangle}{2\Delta\tau},$$

where  $\Delta X$  represents displacement during the sampling interval  $\Delta\tau$ . Choosing the correct sampling interval is a notorious problem; for a comparison see [Pavliotis and Stuart, 2008].

If we take the average diffusivity for each vortex coordinate we find

$$K_{\Delta\tau} = \frac{1}{6M\Delta\tau} \sum_{m=1}^M \langle \Delta \mathbf{x}_m \cdot \Delta \mathbf{x}_m \rangle.$$

We assume the observations are given at the same times  $t_i$  as before and that the sampling time is an integer multiple of the observation interval, i.e.  $\Delta\tau = k \frac{T}{L}$ . With an ensemble of  $P$  simulations the diffusivity estimator would then be

$$\begin{aligned} K_{\Delta\tau} &= \frac{1}{6MP\Delta\tau} \sum_{p=1}^P \sum_{m=1}^M (\mathbf{x}_m^p(\Delta\tau) - \mathbf{x}_m^p(0)) \cdot (\mathbf{x}_m^p(\Delta\tau) - \mathbf{x}_m^p(0)) \\ &= \frac{1}{6MP\Delta\tau} \sum_{p=1}^P \sum_{m=1}^M 2 - 2\mathbf{x}_m^p(0) \cdot \mathbf{x}_m^p(\Delta\tau) \\ &= \frac{1}{3\Delta\tau} - \frac{1}{3MP\Delta\tau} \sum_{p=1}^P \sum_{m=1}^M \mathbf{x}_m^p(0) \cdot \mathbf{x}_m^p(\Delta\tau), \end{aligned}$$

where again a superscript  $p$  denotes the solution from ensemble member  $p$ . Averaging over all time series data yields the estimator:

$$K_{\Delta\tau} = \frac{1}{3\Delta\tau} - \frac{1}{3MP(L-k)\Delta\tau} \sum_{p=1}^P \sum_{j=1}^{L-k} \sum_{m=1}^M \mathbf{x}_m^p(t_j) \cdot \mathbf{x}_m^p(t_j + \Delta\tau) = \frac{1 - \hat{\nu}_k}{3\Delta\tau},$$

with  $\Delta\tau = k \frac{T}{L} = t_k$ . This *shift-averaged estimator* incorporates all possible time intervals of length  $\Delta\tau$  available from the data; it is shown by Cotter and Pavliotis [2009] to improve the quality of the estimator.

Figure 5 presents the diffusivity constant as a function of the interval time  $\Delta\tau$  for different parameter values for  $\varepsilon$  and  $\gamma$ . The results are taken from the same simulations used for the autocorrelation functions, as described in Section .

For  $\varepsilon$  small, the thermostat perturbation is weak, and both autocorrelation functions and diffusivity approach those of the reduced model with constant  $H$ ,  $\mathbf{J}$ . Also, the autocorrelations are insensitive to the parameter  $\gamma$  in this regime. For values of  $\varepsilon < 1$  the autocorrelations and diffusivities become indistinguishable from those of the isolated model. Hence even though the dynamics samples the least-biased density on long time scales, its short time dynamics is similar to an unperturbed model. For moderate  $\varepsilon$ , dependence on  $\gamma$  becomes more pronounced, and a diffusivity closer to that of the full model can be achieved. For even larger values of  $\varepsilon$ , the diffusivity becomes much more sensitive to the value of  $\gamma$ , as indicated in Figure 5c.

Figure 5d has been included to illustrate two important properties by plotting the diffusion parameter against the sampling interval on a log-log scale. First, for large sampling interval the estimator shows an inverse linear tendency. This corresponds simply to the decorrelation of the vortex dynamics. Second, as the sampling interval goes to zero, the estimator of the diffusivity constant shows polynomial behavior. This is in agreement with known results for the GBK thermostat [Frank and Gottwald, 2011] and is an improvement on Langevin thermostats, which tend to a constant value for short sampling intervals, in disagreement with a deterministic reference [DelSole, 2000]. See also the discussion in Section 2.

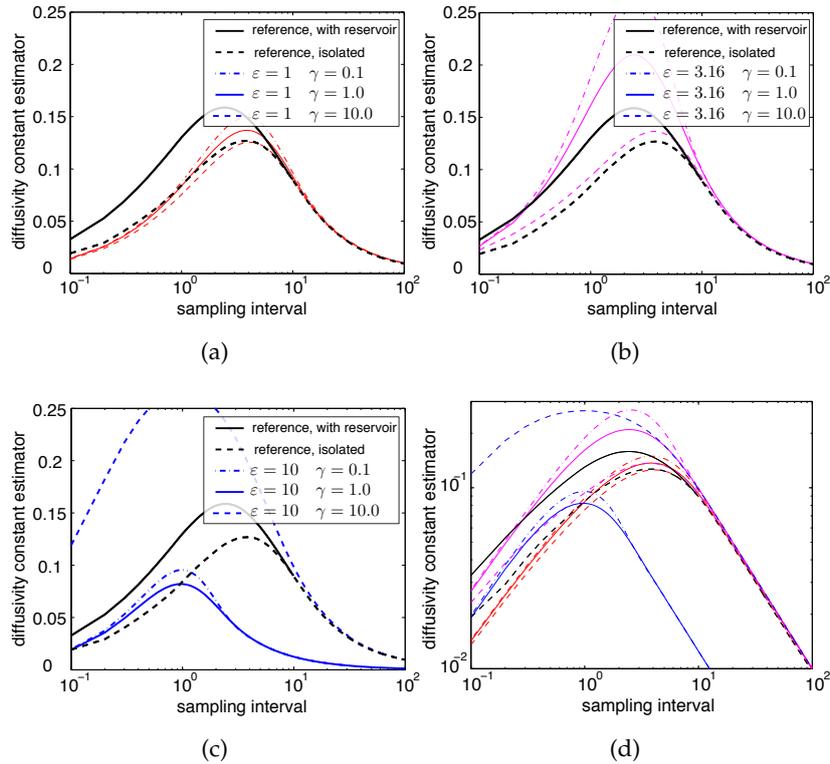


Figure 5: Comparison of average diffusivity constant as a function of sampling intervals. In all figures, the bold lines indicate two reference cases: the full model (solid) and the reduced model (dashed). Subfigures (a), (b) and (c) show thermostatted simulation results for  $\epsilon$  equal to  $10^0$ ,  $10^{0.5}$  and  $10^1$  respectively. The value of  $\gamma$  is represented by dash-dot ( $10^{-1}$ ), solid ( $10^0$ ) or dashed ( $10^1$ ) lines. A combined log-log plot of all parameter values is given in subfigure (d).

### (c) Adaptive determination of multipliers

In this section, we apply the scheme of Section 4 to the point vortex system. Consider the same reduced point vortex model of 8 vortices with  $\Gamma = \pm 1$ . We perform a thermostatted simulation of this system, which we feed with observation data taken from a simulation of the 288-vortex system including the weak vortices that form the thermal bath. The observation from simulation with the thermal bath may be replaced by observation from another source—if available—such as empirical observations or a prediction of the trend for the observables. In particular, we use the observed data in three different ways for computing the Lagrange multipliers that determine the equilibrium distribution used by the thermostat:

- (i) The first approach mirrors the non-adaptive method of Section 3. We use the *long time averages* from equilibrated thermal bath simulation. The equilibrium distribution varies during the thermostatted simulation, but during the whole simulation it converges to the same equilibrium density as before (like those denoted in Table 2).
- (ii) The second approach does not assume long time averages are available a priori. Instead, the *running mean* from the thermal bath simulation is used. This means the iterative procedure (4.4) drives the system toward a varying target during simulation. Eventually, the target approaches the long time average used above.
- (iii) In the final implementation, we relax the notion of statistical equilibrium for the strong system and consider the statistical state to be slowly varying. This is implemented by

using a *time-local averaged* observations from the thermal bath simulations. As a result, there is no stationary equilibrium density.

For testing the adaptive scheme, we only consider the thermal bath simulations where  $H_{\text{full}} = 0$ . Also, we consider just two observables  $C_1 = H$  and  $C_2 = |\mathbf{J}|^2$ , associated with the Lagrange multipliers  $\beta_H$  and  $\beta_J$ . Including more observables is possible, but the Newton-Raphson iteration is sensitive to noise in the estimators, requiring a larger ensemble. We initialize with an ensemble of  $P = 100$  initial conditions drawn from the uniform prior, applying Algorithm (a). The time step in the thermostatted simulation is chosen as  $1 \times 10^{-2}$  and the method described in Section 4 for updating the Lagrange multipliers is applied every time unit, i.e.  $M = 100$ . Between subsequent updates of the multipliers, the maximum difference is limited by  $|\Delta\lambda_k| \leq 0.1$ . When using equilibrium statistics, this limit only affects the beginning of the simulation, when the small sample size used leads to a large variance in the estimators.

We present the results for the three different uses of the observed data in Figures 6–8.

- (i) In Figure 6 the long time mean is taken and used throughout.
- (ii) In Figure 7 the running mean is used. This reflects the situation where we have no a priori knowledge of the observations, and are continuously feeding new real-time data into the simulation.
- (iii) In Figure 8 a time-localized average of the observable is used. The averaging has a time-scale of a 100 time units.

In each of the Figures, the left-hand panel shows the target mean for the energy solid cyan (light), the instantaneous ensemble mean in black dots and the running mean of the ensemble in dashed red (dark). Each centre panel shows the same means for the momentum magnitude and the right hand panel shows the Lagrange multipliers for the energy (cyan, light) and momentum (black). Superimposed on the Lagrange multiplier for the energy is an indicator (red, darker dots) showing when the rate of change in the estimate for the Lagrange multiplier has been limited.

When using either a long time mean observation or a running mean observation, the simulation results tend towards the correct long-time averages. When using time-local averages the simulation averages appear to tend towards a similar value. In all three cases the instantaneous ensemble mean remains close to the (moving) target for both energy and momentum. This is especially notable for the third case, where the target varies over time, but the simulation ensemble mean follows closely, with only a little lag.

The inaccuracies during the first approximately 100 time units indicate that the prior does not match the observed state well. This results in the (negative) growth of  $\beta_H$  being limited briefly at the beginning of each simulation. Subsequently, both Lagrange multipliers appear to oscillate irregularly about some mean value for the first two cases. In the case of a shifting target, the Lagrange multipliers vary in time more erratically, resulting in the limiter being active for a few brief periods of the simulation.

## 7. Conclusion

We have proposed a thermostat-based method for perturbing trajectories of numerical simulations to correct for stationary or time-dependent thermodynamic observations. We have also described an adaptive procedure that uses the thermostatted simulation data in finding the least-biased density, removing the need to compute the density a priori.

Our numerical experiments with the vortex model confirm that the distributions of the observed energy  $H$  and angular momentum magnitude  $|\mathbf{J}|$  can be well approximated using the thermostat technique. Other equilibrium metrics such as the distribution of distances between like- and opposite-signed vortices are also in agreement across a range of total energy values of the full system, although some discrepancies occur at large positive energies.

We also investigated the degree to which correction of trajectories for expectations may affect dynamical information in the form of autocorrelation functions and diffusivity. By decreasing the perturbation parameter  $\varepsilon$  of the thermostat, the autocorrelation functions of the unperturbed, reduced system may be precisely recovered. As  $\varepsilon$  is increased, one may increase the diffusivity to values that agree with the full system. This is consistent with results reported in [Frank and Gottwald, 2011] in the context of molecular dynamics where it was shown that the GBK thermostat used here approaches Langevin dynamics in the limit of large stochastic forcing.

**Data Accessibility.** The code for thermostated simulation of the point vortex system can be found at [https://www.dropbox.com/sh/ekhj4ohtetj35ri/AABNmPHki63cn16hpnaQ\\_\\_OYa?dl=0](https://www.dropbox.com/sh/ekhj4ohtetj35ri/AABNmPHki63cn16hpnaQ__OYa?dl=0).

**Authors' Contributions.** The authors have contributed equally to this work.

**Competing Interests.** The authors have no competing interests.

**Funding.** The work of K.W.M. was supported by a grant financed by the Netherlands Organisation for Scientific Research (NWO), as part of research programme 613.001.009. B.L. was supported by grant EP/K039512/1, Grant SI2-CHE: ExTASY (Extensible Tools for Advanced Sampling and analysis) from the Engineering and Physical Sciences Research Council (UK), and a faculty fellowship from the Alan Turing Institute (ATI).

## A. Time integration of the thermostated system

In this Appendix we detail the method of integration for the perturbation dynamics and thermostat variable.

The unperturbed dynamics are formed by the combination of pairwise interactions:

$$f(\mathbf{x}) = B(\mathbf{y}) \nabla_{\mathbf{y}} H(\mathbf{y}) = -B(\mathbf{y}) \sum_{i=1}^M \sum_{j=1}^{i-1} \nabla_{\mathbf{y}} \frac{\Gamma_i \Gamma_j}{4\pi} \ln(2 - 2\mathbf{x}_i \cdot \mathbf{x}_j) = \sum_{i < j} f_{ij}. \quad (\text{A } 1)$$

This fact is exploited in our earlier work [Myerscough and Frank, 2016] to construct a splitting method for integration of the point vortex system with favourable (geometric) properties. The modified double bracket thermostat is also a combination of pairwise vortex interactions. Hence it is natural to apply the same splitting to  $g$  as was applied to  $f$ . In fact, the pairwise thermostat interaction can be executed along with the unperturbed dynamics. Let us denote by  $\phi_{\Delta t}^{i,j}$  the time  $\Delta t$  flow map associated with  $f_{ij}$ . Similarly we define by  $\gamma_{\Delta t}^{i,j}$  the time  $\Delta t$  flow map associated with the perturbation dynamics of a vortex pair  $(i, j)$ . The flow map of the dynamics of the thermostat variable  $\xi$  is represented by  $\chi_{\Delta t}$ . A symmetric composition of these flows is given by

$$\Phi_{\Delta t} = \prod_{(i,j) \in C} \left( \phi_{\Delta t/2}^{i,j} \circ \gamma_{\Delta t/2}^{i,j} \right) \circ \chi_{\Delta t} \circ \prod_{(i,j) \in C^*} \left( \gamma_{\Delta t/2}^{i,j} \circ \phi_{\Delta t/2}^{i,j} \right),$$

where  $C$  is an ordering of pairs and  $C^*$  is the reverse ordering. While advancing  $\xi$  by  $\chi_{\Delta t}$ , it is assumed that  $\mathbf{y}$  is fixed and therefore  $h(\mathbf{y})$  is a constant. This means the dynamics of  $\xi(t)$  is just an Ornstein-Uhlenbeck process with mean  $\frac{h}{\gamma}$  and unit variance. The exact solution is given by

$$\chi_{\Delta t} \xi_0 = \xi_0 e^{-\gamma \Delta t} + \frac{h}{\gamma} (1 - e^{-\gamma \Delta t}) + e^{-\gamma \Delta t} W(e^{2\gamma \Delta t} - 1).$$

The integration of the perturbed dynamics for a vortex pair  $(i, j)$  means integrating the system of ordinary differential equations given by

$$\begin{aligned} \dot{\mathbf{x}}_i &= \frac{\Gamma_j}{4\pi} \mathbf{x}_i \times \mathbf{x}_i \times \mathbf{x}_j = \frac{\Gamma_i + \Gamma_j}{4\pi} \mathbf{x}_i \times \mathbf{x}_i \times \mathbf{x}_j + \frac{\Gamma_i - \Gamma_j}{4\pi} \mathbf{x}_i \times \mathbf{x}_i \times \mathbf{x}_j, \\ \dot{\mathbf{x}}_j &= \frac{\Gamma_i}{4\pi} \mathbf{x}_j \times \mathbf{x}_j \times \mathbf{x}_i = \frac{\Gamma_i + \Gamma_j}{4\pi} \mathbf{x}_j \times \mathbf{x}_j \times \mathbf{x}_i - \frac{\Gamma_i - \Gamma_j}{4\pi} \mathbf{x}_j \times \mathbf{x}_j \times \mathbf{x}_i. \end{aligned}$$

The equations have been split into a symmetric and an anti-symmetric part. A symmetric composition may once more be applied to integrate the two parts. In our case however, the

thermostatted system consists only of vortices with strength  $\pm\Gamma_{\text{strong}}$  and thus each vortex pair interaction is either fully symmetric or fully antisymmetric.

By introducing  $\alpha_S = \frac{\Gamma_i + \Gamma_j}{4\pi}$  the symmetric part may be written as

$$\dot{\mathbf{x}}_i = \alpha_S \mathbf{x}_i \times \mathbf{x}_i \times \mathbf{x}_j = \alpha_S (\mathbf{x}_j \times \mathbf{x}_i) \times \mathbf{x}_i, \quad (\text{A } 2)$$

$$\dot{\mathbf{x}}_j = \alpha_S \mathbf{x}_j \times \mathbf{x}_j \times \mathbf{x}_i = -\alpha_S (\mathbf{x}_j \times \mathbf{x}_i) \times \mathbf{x}_j. \quad (\text{A } 3)$$

These dynamics are symmetric with respect to the plane equidistant between the two vortices. This is because the dynamics are rotations in opposite direction about the same vector  $\mathbf{x}_j \times \mathbf{x}_i$  and this vector must lie in said plane. Furthermore the dynamics of (A 2)–(A 3) does not allow the vortex pair to pass through the position where they are exactly opposed, as in this case their cross product is zero. All in all this means the final position of the two vortices is uniquely determined by their chord distance.

The change in the distance between the two vortices can be represented by the change of their inner product

$$\begin{aligned} \frac{\partial}{\partial t} (\mathbf{x}_i \cdot \mathbf{x}_j) &= \dot{\mathbf{x}}_i \cdot \mathbf{x}_j + \mathbf{x}_i \cdot \dot{\mathbf{x}}_j \\ &= 2\alpha_S \left( (\mathbf{x}_i \cdot \mathbf{x}_j)^2 - 1 \right) \end{aligned}$$

The solution to this differential equation is given by

$$\mathbf{x}_i \cdot \mathbf{x}_j |_{t=\Delta t} = \tanh(2\alpha\Delta t - \text{artanh}(\mathbf{x}_i \cdot \mathbf{x}_j) |_{t=0}).$$

In the anti-symmetric part we introduce  $\alpha_A = \frac{\Gamma_i - \Gamma_j}{4\pi}$  to write

$$\dot{\mathbf{x}}_i = \alpha_A \mathbf{x}_i \times \mathbf{x}_i \times \mathbf{x}_j,$$

$$\dot{\mathbf{x}}_j = -\alpha_A \mathbf{x}_j \times \mathbf{x}_j \times \mathbf{x}_i.$$

By rearranging the order of the cross product we achieve

$$\dot{\mathbf{x}}_i = \alpha_A (\mathbf{x}_j \times \mathbf{x}_i) \times \mathbf{x}_i = \hat{a}_A \mathbf{x}_i, \quad (\text{A } 4)$$

$$\dot{\mathbf{x}}_j = \alpha_A (\mathbf{x}_j \times \mathbf{x}_i) \times \mathbf{x}_j = \hat{a}_A \mathbf{x}_j. \quad (\text{A } 5)$$

The vector  $\mathbf{a}_A$  implicitly defined by (A 4)–(A 5) can simply be shown to be constant under the antisymmetric flow. This means the dynamics of the anti-symmetric part may be integrated by using Rodrigues's formula.

## References

- R. V. Abramov and A.J. Majda. Statistically relevant conserved quantities for truncated quasigeostrophic flow. *Proceedings of the National Academy of Sciences*, 100(7):3841–3846, 2003.
- N. Agmon, Y Alhassid, and R.D. Levine. An algorithm for finding the distribution of maximal entropy. *Journal of Computational Physics*, 30(2):250–258, 1979.
- J. Bajars, J. Frank, and B. Leimkuhler. Stochastic-dynamical thermostats for constraints and stiff restraints. *Eur. Phys. J. Special Topics*, 200:131–152, 2011.
- J. Bajars, J.E. Frank, and B. Leimkuhler. Weakly coupled heat bath models for Gibbs-like invariant states in nonlinear wave equations. *Nonlinearity*, 26(7):1945, 2013.
- Roger Balian. *From microphysics to macrophysics: methods and applications of statistical physics*, volume 2. Springer Science & Business Media, 2006.
- A. Bloch, P.S. Krishnaprasad, J.E. Marsden, and T.S. Ratiu. The Euler-Poincaré equations and double bracket dissipation. *Communications in Mathematical Physics*, 175(1):1–42, 1996.
- O. Bühler. Statistical mechanics of strong and weak point vortices in a cylinder. *Phys. of Fluids*, 14(7):2139–2149, 2002.
- A. Bulgac and D. Kusnezov. Canonical ensemble averages from pseudomicrocanonical dynamics. *Phys. Rev. A*, 42:5045–5048, Oct 1990. doi: 10.1103/PhysRevA.42.5045.

- P.-H. Chavanis. Kinetic theory of point vortices: Diffusion coefficient and systematic drift. *Physical Review E*, 64(2):026309, 2001.
- C.J. Cotter and G.A. Pavliotis. Estimating eddy diffusivities from noisy Lagrangian observations. *Communications in Mathematical Sciences*, 7(4):805–838, 2009.
- G.-H. Cottet and P.D. Koumoutsakos. *Vortex Methods*. Cambridge University Press, first edition, 2000.
- J.F. Dama, A.V. Sinitskiy, M. McCullagh, J. Weare, B. Roux, A.R. Dinner, and G.A. Voth. The theory of ultra-coarse-graining. 1. general principles. *Journal of Chemical Theory and Computation*, 9(5): 2466–2480, 2013. doi: 10.1021/ct4000444. URL <http://dx.doi.org/10.1021/ct4000444>. PMID: 26583735.
- P.A. Davidson. *Turbulence: An Introduction for Scientists and Engineers*. Oxford University Press, 2004.
- S. Davis and G. Gutiérrez. Conjugate variables in continuous maximum-entropy inference. *Physical Review E*, 86(5):051136, 2012.
- T. DelSole. A fundamental limitation of Markov models. *Journal of the Atmospheric Sciences*, 57(13): 2158–2168, 2000.
- R.C. Dewar, C.H. Lineweaver, R.K. Niven, and K. Regenauer-Lieb, editors. *Beyond the Second Law: Entropy Production and Non-equilibrium Systems*. Understanding Complex Systems. Springer, 2014.
- N. Ding, Y. Fang, R. Babbush, C. Chen, R.D. Skeel, and H. Neven. Bayesian sampling using stochastic gradient thermostats. In *Advances in Neural Information Processing Systems (NIPS)*, 2014.
- S. Dubinkina and J. Frank. Statistical mechanics of Arakawa’s discretizations. *Journal of Computational Physics*, 227(2):1286–1305, 2007.
- S. Dubinkina and J. Frank. Statistical relevance of vorticity conservation in the hamiltonian particle-mesh method. *Journal of Computational Physics*, 229(7):2634–2648, 2010.
- S. Dubinkina, J. Frank, and B. Leimkuhler. Simplified modelling of a thermal bath system, with application to a fluid vortex system. *Multiscale Model Simul.*, 8(5):1882–1901, 2010.
- L.D. Faddeev and L.A. Takhtajan. *Hamiltonian methods in the theory of solitons*. Springer Berlin, 1987.
- J. Frank. Geometric space–time integration of ferromagnetic materials. *Applied numerical mathematics*, 48(3):307–322, 2004.
- J. Frank and G.A. Gottwald. The Langevin limit of the Nosé-Hoover-Langevin thermostat. *Journal of Statistical Physics*, 143(4):715–724, 2011.
- J. Frank, W. Huang, and B. Leimkuhler. Geometric integrators for classical spin systems. *J. Comput. Phys.*, 133(1):160–172, 1997. ISSN 0021-9991. doi: 10.1006/jcph.1997.5672.
- D. Frenkel and B. Smit. *Understanding Molecular Simulation*. Academic Press, San Diego, second edition edition, 2002. ISBN 978-0-12-267351-1. doi: <http://dx.doi.org/10.1016/B978-012267351-1/50000-6>. URL <http://www.sciencedirect.com/science/article/pii/B9780122673511500006>.
- A. Giffin and R. Uffink. Simultaneous state and parameter estimation using maximum relative entropy with nonhomogenous differential equation constraints. *Entropy*, 16(9):4974–4991, 2014.
- E. Hairer, C. Lubich, and G. Wanner. *Geometric Numerical Integration: Structure-Preserving Algorithms for Ordinary Differential Equations*. Springer-Verlag, 2006.
- H. Haken. *Information and Self-Organization: A Macroscopic Approach to Complex Systems*. Springer, 2006.
- W. Hayes and K.R. Jackson. A fast shadowing algorithm for high dimensional ODE systems. *SIAM Journal on Scientific Computing*, 29(4):1738–1758, 2007.
- P. Holmes, J.L. Lumley, and G. Berkooz. *Turbulence, coherent structures, dynamical systems and symmetry*. Cambridge University Press, 1998.
- W.G. Hoover. Canonical dynamics: equilibrium phase-space distributions. *Physical Review A*, 31(3):1695, 1985.
- U. Hornung. *Homogenization and porous media*, volume 6. Springer Verlag, 1997.

- Patrick Ilg, Iliya V Karlin, and Hans Christian Öttinger. Canonical distribution functions in polymer dynamics.(i). dilute solutions of flexible polymers. *Physica A: Statistical Mechanics and its Applications*, 315(3):367–385, 2002.
- M.I. Jamaloodeen and P.K. Newton. The N-vortex problem on a rotating sphere: II. Heterogeneous Platonic solid equilibria. *Proc. R. Soc. A*, 462(2075):3277–3299, 2006.
- E.T. Jaynes. Information theory and statistical mechanics. *Physical review*, 106(4):620, 1957a.
- E.T. Jaynes. Information theory and statistical mechanics. II. *Phys. Rev.*, 108:171–190, Oct 1957b. doi: 10.1103/PhysRev.108.171.
- S. Kehrbaum and J.H. Maddocks. Elastic rods, rigid bodies, quaternions and the last quadrature. *Phil. Trans. R. Soc. Lond. A*, 355(1732):2117–2136, 1997.
- Wolfgang Kliemann. Recurrence and invariant measures for degenerate diffusions. *The annals of probability*, pages 690–707, 1987.
- F. Laurent-Polz. Point vortices on a rotating sphere. *Regular and Chaotic Dynamics*, 1(10):39–58, 2005.
- B. Leimkuhler. Generalized Bulgac-Kusnezov methods for sampling of the Gibbs-Boltzmann measure. *Physical Review E*, 81(2):026703, 2010.
- B. Leimkuhler and S. Reich. *Simulating Hamiltonian Dynamics*. Cambridge Univ. Press, Cambridge, first edition, 2005.
- B. Leimkuhler, E. Noorzadeh, and O. Penrose. Comparing the efficiencies of stochastic isothermal molecular dynamics models. *J. Stat. Phys.*, 143(5):921–942, 2011.
- A.J. Majda and A.L. Bertozzi. *Vorticity and Incompressible Flow*. Cambridge University Press, first edition, 2002.
- A.J. Majda and B. Gershgorin. Improving model fidelity and sensitivity for complex systems through empirical information theory. *Proceedings of the National Academy of Sciences*, 108(25): 10044–10049, 2011.
- A.J. Majda and X. Wang. *Nonlinear Dynamics and Statistical Theories for Basic Geophysical Flows*. Cambridge University Press, 2006.
- J.C. Mattingly, A.M. Stuart, and D.J. Higham. Ergodicity for SDEs and approximations: locally Lipschitz vector fields and degenerate noise. *Stochastic processes and their applications*, 101(2): 185–232, 2002.
- N. Metropolis, A.W. Rosenbluth, M.N. Rosenbluth, A.H. Teller, and E. Teller. Equation of state calculations by fast computing machines. *The Journal of Chemical Physics*, 21(6):1087–1092, 1953.
- K.W. Myerscough and J. Frank. Explicit, parallel Poisson integration of point vortices on the sphere. *Journal of Computational and Applied Mathematics*, 304:100–119, 2016.
- P.K. Newton. *The N-Vortex Problem: Analytical Techniques*. Springer-Verlag, New York, first edition, 2001.
- P.K. Newton and T. Sakajo. The N-vortex problem on a rotating sphere: III. Ring configurations coupled to a background field. *Proc. R. Soc. A*, 463(2080):961–977, 2007.
- P.K. Newton and H. Shokraneh. The N-vortex problem on a rotating sphere: I. Multi-frequency configurations. *Proc. R. Soc. A*, 462(2065):149–169, 2006.
- S. Nosé. A molecular dynamics method for simulations in the canonical ensemble. *Molecular physics*, 52(2):255–268, 1984a.
- S. Nosé. A unified formulation of the constant temperature molecular dynamics methods. *The Journal of Chemical Physics*, 81(1):511–519, 1984b.
- S. Panagiotis. A maximum likelihood algorithm for the estimation and renormalization of exponential densities. *Journal of Computational Physics*, 208(2):691–703, 2005.
- G.W. Patrick. Dynamics of perturbed relative equilibria of point vortices on the sphere or plane. *Journal of Nonlinear Science*, 10(3):401–415, 2000.
- G.A Pavliotis and A.M. Stuart. *Multiscale Methods: Averaging and Homogenization*. Springer, 2008.
- L. Rey-Bellet. Ergodic properties of Markov processes. *Open Quantum Systems II*, pages 1–39, 2006.
- Gareth O Roberts and Richard L Tweedie. Exponential convergence of langevin distributions and their discrete approximations. *Bernoulli*, pages 341–363, 1996.

- Giovanni Samaey, Tony Lelievre, and Vincent Legat. A numerical closure approach for kinetic models of polymeric fluids: exploring closure relations for fene dumbbells. *Computers & Fluids*, 43(1):119–133, 2011.
- A. Samoletov, C. Dettmann, and M. Chaplain. Thermostats for “slow” configurational modes. *Journal of Statistical Physics*, 128(6):1321–1336, 2007.
- T. Schlick. *Molecular Modeling and Simulation: An Interdisciplinary Guide (Interdisciplinary Applied Mathematics)*. Springer, 2nd edition, 2010. ISBN:978-1441963505.
- M.S. Shell. The relative entropy is fundamental to multiscale and inverse thermodynamic problems. *The Journal of Chemical Physics*, 129(14):144108, 2008. doi: <http://dx.doi.org/10.1063/1.2992060>. URL <http://scitation.aip.org/content/aip/journal/jcp/129/14/10.1063/1.2992060>.
- M.-Q. Zhang and M.-Z. Qin. Explicit symplectic schemes to solve vortex systems. *Computers Math. Applic.*, 26(5):51–56, 1993.

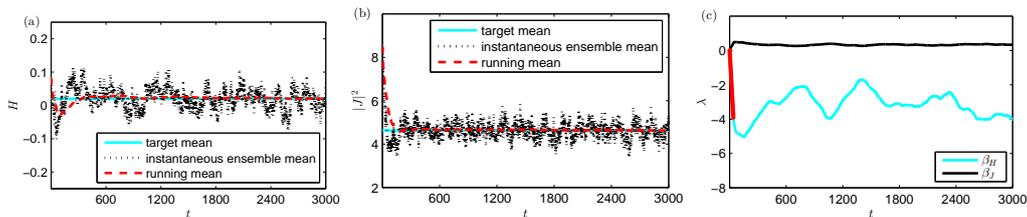


Figure 6: Results when using *long-time mean* observations as a target while adaptively determining the Lagrange multipliers. Target observations for Hamiltonian (a) and momentum magnitude (b) are overlaid with the instantaneous ensemble mean (black dotted) and the running ensemble mean (red solid) from simulation. (c): Lagrange multipliers, the red dots indicate time steps at which their rate of change was limited.

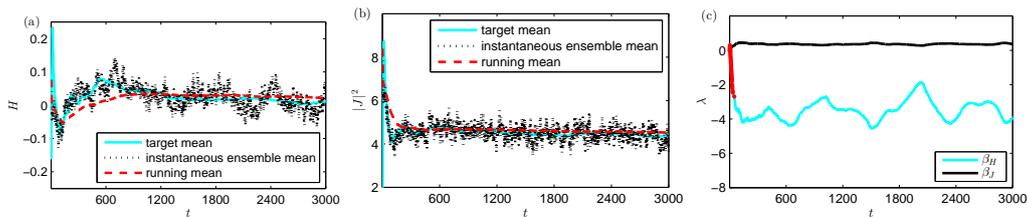


Figure 7: Results when using *running mean* observations as a target while adaptively determining the Lagrange multipliers.

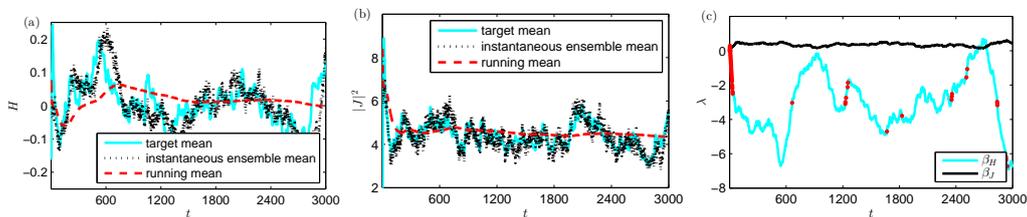


Figure 8: Results when using *time-local averaged* observations as a target while adaptively determining the Lagrange multipliers.