



THE UNIVERSITY *of* EDINBURGH

This thesis has been submitted in fulfilment of the requirements for a postgraduate degree (e. g. PhD, MPhil, DClinPsychol) at the University of Edinburgh. Please note the following terms and conditions of use:

- This work is protected by copyright and other intellectual property rights, which are retained by the thesis author, unless otherwise stated.
- A copy can be downloaded for personal non-commercial research or study, without prior permission or charge.
- This thesis cannot be reproduced or quoted extensively from without first obtaining permission in writing from the author.
- The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the author.
- When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given.

**Clustering of particles in weakly
divergent flows.**

Aidan Tully

Doctor of Philosophy
University of Edinburgh
August 2025

Abstract

The clustering of particles in turbulent flows is an important mechanism in both natural and industrial processes. It appears across a diverse range of natural phenomena such as rain initiation and growth of rain drops, plankton dispersion in the ocean, and planetesimal formation as a result of aggregation of dust grains. In particular, turbulent particle laden flows in the ocean lead to the clustering of inorganic particulate matter, a topic of active interest with increasing concerns about microplastic and other particulate pollutants.

Divergence in the particle velocity field is the cause of clustering of particles. The divergence is not necessarily a product of the compressibility of the carrier fluid and can occur from indirect causes. This thesis examines the clustering that results from two such indirect causes. The first is the inertia of particles, which makes the particle velocity differ from the fluid velocity and leads to clustering. We examine this process using the Maxey–Riley equation for inertial particle motion in the limit of small Stokes number $St \ll 1$. In this limit, the dynamics in the full position–velocity phase space can be reduced to simple advection in position space, but with a new ‘effective’ velocity that is weakly divergent even though the fluid flow is incompressible. We study the statistics of clustering in this set up using a simple two-dimensional kinematic model in which the fluid velocity is random in time.

The second indirect cause of divergence that is studied concerns particles floating at the surface of the ocean. These particles experience a two-dimensional divergent velocity field, namely the horizontal part of the full three-dimensional non-divergent fluid velocity. We model this horizontal flow using the SQG⁺¹ model. This is an improved-accuracy version of the surface quasigeostrophic (SQG) model which captures the first-order corrections in the Rossby number (Ro) including the weak (horizontal) divergence.

Particles in divergent flows converge to a fractal attractor in the position–velocity

phase space. The projection of this attractor in physical space forms a fractal set. For floating particles, we make use of techniques from chaos theory to quantify the dimension of this set and consider in particular the information dimension D_1 and correlation dimension D_2 . For D_1 , we rely on the Kaplan-Yorke conjecture which expresses D_1 in terms of Lyapunov exponents. D_2 is computed as a correlation sum. The dimension of the attractor is shown to scale quadratically in Ro for both estimates.

For both inertial and floating particles, we examine in detail the rate of clustering as measured by large-deviations statistics of the particle density distribution. A Lagrangian approach is taken to overcome the challenge arising from the finite resolution of Eulerian models. We consider both the rate function describing the long time behaviour of this distribution and its Legendre dual, the free energy, describing the long time behaviour of the moments of the density.

We develop and implement an importance sampling procedure to estimate the free energy accurately. This is necessary to capture the statistics of high moments of the density which are dominated by rare trajectories.

We examine the applicability of the Kraichnan limit of short correlation time of the velocity for both inertial and floating particles. In this limit, we predict a parabolic shape for the free energy. The prediction is shown to be robust even when applied to experiments not strictly in the short correlation time limit.

Lay summary

Particles within flows can be found readily in nature, from raindrops or dust in the atmosphere, to sediment in rivers or pollutants in the ocean. When the fluid that carries these particles is compressible the trajectories of particles moving within the fluid converge. This can happen even if the fluid itself is not compressible but the interaction between the flow and the particle is not simply advection following the fluid velocity field, such as heavy (inertial) particles in the flow.

Inevitably in these compressible, or divergent, flows the particles cluster together. We use numerical models to describe this interaction and the study the statistics of the clustering behaviour that occurs. We do this for two cases. The first is inertial particles, which can be thought of as heavy relative to the fluid and as such react slowly to changes in the velocity field creating a divergent effect. The second is for passive particles in a flow that resembles the surface of the ocean. This has applications in the transport and clustering behaviour of microplastics for example.

We examine the statistics of clustering in these two cases to quantify the rate at which clustering occurs, and how likely a concentration, or density, of particles is to occur. This is extended using an importance sampling algorithm to increase the resolving ability of the statistics to find rare values of density.

We use mathematical tools to describe the space that the clustering happens on, and show that it is fractal in nature, with dimension less than the dimension of the space of the flow. In fact, there is a dependence on the square of a small parameter that controls the strength of the divergence in the flow.

These results are compared to theoretical arguments to both verify them and provide justification for the observed behaviour.

Declaration

I declare that this thesis was composed by myself and that the work contained therein is my own, except where explicitly stated otherwise in the text.

(Aidan Tully)

Acknowledgements

I wish to thank my supervisor, Jacques Vanneste, for his support, guidance and patience over the course of this PhD. I am grateful for everything I have learned from him, and the time he has given to me.

I thank my parents, and sister for their support. The ability to undertake this PhD would not be possible without them.

Finally, to my friends and Alanna, who keep me entertained and sane - thank you, it is not taken for granted.

I love deadlines. I love the whooshing noise they make as they go by.

DOUGLAS ADAMS

Contents

Abstract	v
Contents	xi
1 Introduction	1
2 Background	6
2.1 The motion of particles within a flow	6
2.2 Clustering of particles	7
2.2.1 Measuring the clustering through particle density	7
2.3 Large Deviations theory	13
2.3.1 Zeros of the free energy and its implications for the growth of moments	18
2.4 Importance sampling	24
2.4.1 Importance sampling algorithm	25
2.4.2 Applying importance sampling to a problem with a known dis- tribution	27
2.5 Lyapunov exponents	32
2.5.1 Calculating Lyapunov exponents	32
2.5.2 Dissipative dynamical systems	38
2.6 Summary	42
3 Analytical results	43
3.1 Homogenisation for weakly compressible flows	43
3.1.1 Time independent flow and diffusion	44
3.1.2 Time-periodic flows	47
3.1.3 Removing diffusion from time dependent flow	48

3.2	Modelling weakly compressible flows as Kraichnan flows	51
3.3	Summary	56
4	Inertial particles	58
4.1	Motion of inertial particles	59
4.1.1	Maxey-Riley equations of motion	59
4.1.2	Approximation of dynamics by motion on an invariant manifold	62
4.2	2D Synthetic Random Flow	64
4.2.1	Nested OU process	64
4.2.2	Explicit form of velocity field and effective velocity field	66
4.3	Numerically simulating the advection of inertial particles	68
4.3.1	Euler–Maruyama scheme for the stochastic velocity field	68
4.3.2	Numerical scheme for particle advection	69
4.3.3	Effectiveness of the effective flow approximation	71
4.4	Results	73
4.4.1	Results with Large Deviation theory	74
4.4.2	Importance sampling	78
4.4.3	Comparison of results from importance sampling with brute force	79
4.4.4	Identifying regions of the flow using importance sampling	81
4.4.5	Scaling of the free energy with St and correlation time γ	85
4.5	Conclusion	91
5	Ocean Surface Flow	92
5.1	Model for the ocean surface	93
5.2	Numerical implementation of SQG^{+1}	95
5.3	Clustering of particles in SQG^{+1}	99
5.4	Fractal set for clustering	101
5.4.1	Calculating the fractal dimensions D_1 and D_2	101
5.5	Large deviation statistics for particles on the ocean surface	104
5.5.1	Large Deviation statistics of $\delta(t)$	106
5.5.2	Comparison of clustering regions in the flow with importance sampling	112
5.5.3	Scaling of the free energy with varied Ro	113
5.6	Conclusion	116

6	Summary and discussion	118
6.1	Summary	118
6.2	Discussion of results	119
6.2.1	Potential applications	121
6.3	Further work	122
A	Additional material for Chapter 5	124
A.1	Background material for SQG & SQG ⁺¹	124
A.1.1	Boussinesq equations and scaling	124
A.1.2	Quasi-Geostrophic (QG) equations	128
A.1.3	Surface Quasi-Geostrophic (SQG) equations	133
A.1.4	Numerical implementation of SQG flow	134
A.1.5	Next order in QG and SQG	137
	Bibliography	143

Chapter 1

Introduction

Multiphase flows occur frequently in both natural and industrial settings. A commonly encountered type are suspensions of particulate or droplets, in some carrier flow. The distribution of these particles is key to understanding processes that occur within these flows.

In nature, examples of multiphase flows and associated phenomena can be found in sediment transport in rivers and the coastal ocean, raindrop formation and dust clouds in the atmosphere (Falkovich et al., 2002; Shaw, 2003; Hubbard et al., 2009; Tsai et al., 2005), pollutants or microorganism transport in the ocean (Sundby and Fossum, 1990; Károlyi et al., 2000), and even cosmic dust clouds precursor to planet formation (Wilkinson et al., 2008; Cuzzi et al., 2001). Multiphase flows are common in industrial processes also, such as particulate mixing in chemical processes (Hessel et al., 2005; Motter et al., 2003) or fuel dispersion in combustion engines (Post and Abraham, 2002; Kuo and Acharya, 2012).

In all of these processes, the clustering of particles is an important aspect, sometimes key to achieving the desired outcome for industrial processes or, conversely, an obstacle to overcome. Raindrops rely on collisions to coalesce for example, and the efficiency of engines is reliant on the fuel concentration in combustion chambers. On the other hand, clustering hinders mixing and reaction, and disrupts the manufacture of homogeneous materials. The distribution and transport of pollutants in the ocean such as microplastics is an area of active interest currently also. (van Sebille et al., 2020; Sutherland et al., 2023)

All of these processes share a common structure, comprised of two phases: the particles or droplets, and the carrier flow they are suspended in. If the velocity field as-

sociated with the particles is divergent, an initially random or homogenous distribution will not remain homogeneous: the clustering of the particles will occur inevitably. The source of the divergence can stem from the carrier flow itself or from the interaction of the particles with the flow.

We are interested in characterising the behaviour of clustering in two main ways: first by quantifying the time scale, or rate, at which clustering occurs in these flows, and second by describing the set occupied by the clustered particles. To do this we take a continuous viewpoint and examine the concentration (or, equivalently, probability density) of particles as a function of time and space with the flow.

This problem is non-trivial, as the presence of clustering implies the convergence of particle trajectories over long enough timescales. In fact, there is an attractor in the dynamics in the joint position–velocity phase space of particles (Balkovsky et al., 2001). The attractor is fractal, generally time dependent, and depends on the carrier flow. As we are interested in the physical location of particles, we concern ourselves only with the spatial projection of this attractor from the position–velocity phase space to the position space. This projection is also fractal, which makes modelling the concentration challenging as it becomes singular in the long time limit.

We examine the clustering behaviour in two cases: inertial particles in a two-dimensional random flow, and passive particles in a weakly compressible geophysical flow representing the ocean surface.

Inertial particles are modelled by the Maxey–Riley equation governing the phase-space dynamics of the particles in a prescribed fluid flow. The phase-space dynamics of inertial particles are dissipative and as result clustering is inevitable in this case. (Siggurdsson and Stuart, 2002; Sapsis and Haller, 2010). We describe this clustering using an approximation of the Maxey–Riley equation valid for small particles which reduces the dynamics from phase-space to position space. In this approximation, particles are transported by an effective flow, determined by, but different from the fluid flow. In particular, the effective flow is divergent even for non-divergent fluid flow. We examine the clustering that results from this divergence for a simple synthetic random flow.

The construction of the ocean surface flow is more complicated. We model particles floating on the ocean surface. Because they are confined to the surface by buoyancy, the particles experience a two-dimensional weakly divergent flow which is the horizontal component of the full three-dimensional incompressible flow of the ocean. We model this three-dimensional flow using the SQG⁺ model. This is an improved accuracy

version of the commonly used surface quasigeostrophic (SQG) model (LaCasce and Mahadevan, 2006; Lapeyre and Klein, 2006; Qiu et al., 2016). The SQG model is a special case of the quasigeostrophic model, the standard model of rapidly rotating, strongly stratified geophysical flows, which has the advantage of being governed by two-dimensional equations. The surface flow of the SQG is horizontally non-divergent, but the correction (first-order in the Rossby number) captured by SQG⁺¹ introduces a weak divergence leading to clustering.

In this thesis we make use of large-deviation statistics and techniques taken from the study of dissipative dynamical systems to characterise clustering.

We use large deviations to measure the rate of clustering with time. We do this through the rate function and free energy for the density of fluid parcels or, more precisely, for the log-density δ . The rate function s controls the long-time asymptotics of the PDF for the log-density according to

$$P(\delta, t) \asymp e^{ts(\delta/t)}, \quad (1.1)$$

where \asymp denotes asymptotic equivalence of the logs. The free energy, which is the Legendre transform of s , describes the moments of the density and quantifies the rate of exponential growth or decay of moments.

We use a Lagrangian approach for both our conceptual framework and for numerical computations. Numerically, the Lagrangian approach replaces the solution of the partial differential equations governing the concentration of particles by the solution of ordinary differential equations governing the positions of a large number of particles. To an extent, this approach avoids the issues related to the finite numerical resolution of an Eulerian density field.

The Lagrangian approach enables us to estimate the rate and free-energy functions from the statistics of a large ensemble of particle trajectories. To capture these functions across a broad range of parameters is challenging however because their tails are controlled by rare, extreme trajectories which are sampled poorly. To remedy this, we implement an importance sampling algorithm which makes it possible to accurately estimate rate and free-energy well outside the narrow range captured by straightforward, brute-force sampling.

As mentioned, particles in divergent flows cluster on a set that is fractal in the long-time limit. We characterise this fractal set through its fractal dimensions. We

estimate specific fractal dimensions, namely the information dimension D_1 and correlation dimension D_2 . We compute these numerically through the Lyapunov exponents and the Kaplan–Yorke conjecture for D_1 and through the correlation sum for D_2 .

The focus of the thesis is on flows that are weakly divergent, that is, with velocity field of the form

$$\mathbf{u} = \mathbf{u}^{(0)} + \epsilon \mathbf{u}^{(1)} \quad \text{with} \quad \nabla \cdot \mathbf{u}^{(0)} = 0, \nabla \cdot \mathbf{u}^{(1)} \neq 0. \quad (1.2)$$

The small parameter ϵ that controls the strength of the divergent part of the flow relative to the non-divergent part is different in the two applications we consider. For inertial particles, it is the Stokes number St , whereas for the ocean flow it is the Rossby number Ro . We examine in details the scaling of the large deviation statistics with respect to St and Ro . We also consider these statistics in the limit of short correlation time, or the Kraichnan limit. We make asymptotic predictions which we test against numerical simulations.

Thesis outline

This thesis is organised as follows.

Chapter 2 introduces important concepts that are used throughout the thesis. These include the concept of the Lagrangian particle density, large deviation statistics, the importance sampling algorithm, and Lyapunov exponents.

Chapter 3 contains theoretical results on the scaling of large deviation statistics in a pair of limiting cases. In particular, we obtain asymptotic results in the limit of explicitly weakly divergent flows of the form (1.2), and in the limit of short correlation time flows.

In Chapter 4 we present the results for clustering for inertial particles in a two-dimensional random flow. A simplification of the inertial dynamics to motion on an invariant slow manifold in the phase space is used to model particle motion. The importance sampling algorithm is used to accurately resolve the large deviation statistics which are presented and compared to statistics created without the algorithm. Scaling of the large deviation statistics is investigated in the limit of small Stokes number St , and short correlation time of the random flow.

Chapter 5 focuses on ocean surface flows. The background for the model used to model the surface flow, SQG⁺¹, is presented. The results for the large deviation

statistics with and without importance sampling are shown. The fractal dimension of the attractor for the particle position is estimated and its scaling with the Rossby number Ro is investigated. The scaling of the large deviation statistics with Ro is also examined here.

The concluding Chapter 6 summarises the key findings of the thesis and identifies potential directions for future work.

Chapter 2

Background

2.1 The motion of particles within a flow

Throughout this thesis we examine the motion of particles advected by a flow they are suspended within. While we deal with both the advection of non-inertial and inertial particles, we begin by outlining here the non-inertial case to provide a general framework for subsequent sections. Changes to this framework through the introduction of inertia will be outlined and examined in chapter 4.

Non-inertial particles, also commonly referred to as passive particles, are usually characterised as infinitely small, massless particles. The motion of such passive particles is simply the same as that of fluid parcels. As they are passive, they have no reciprocal effect on the flow they are in, and are simply advected by the motion around them. For a velocity field $\mathbf{u}(\mathbf{x}, t)$ we write their motion as,

$$\dot{\mathbf{x}}(t) = \mathbf{u}(\mathbf{x}, t), \tag{2.1}$$

where $\mathbf{x}(t)$ is the position vector of the particles at a time t , and an over dot $\dot{\mathbf{x}}$ indicates a derivative in time. The evolution of the background flow is left general here, and we are only concerned for now with properties of the flow rather than the the flow explicitly. Equation (2.1) determines the position of a particle in the flow $\mathbf{u}(\mathbf{x}, t)$ uniquely given conditions on the smoothness of the velocity field $\mathbf{u}(\mathbf{x}, t)$. In particular the solution is unique if the velocity field is locally Lipschitz continuous (Boyce, 2017).

For steady flow, particles will follow streamlines of the flow. Streamlines are the integral curves of $\mathbf{u}(\mathbf{x}, t)$ for a fixed t , these curves are tangent to the flow at every point, and never intersect. For steady flows, streamlines can be helpful to visualise the flow

as they remain unchanged in time and represent the path of a fluid element in the flow. As a result particles suspended in the flow will also follow these contours exactly. When the flow is time-dependent (unsteady), particles no longer follow streamlines exactly as and instead follow trajectories given by integrating (2.1) from an initial condition $\mathbf{x}_0 = \mathbf{x}(0)$.

Moving to more than one particle, we can solve equation (2.1) for many particles simultaneously in the flow and investigate the evolution of the particle distribution over time. For an incompressible fluid, an initially homogenous distribution of particles will remain homogenous for later times under only advection as in (2.1). With the introduction of divergence in the flow, no matter the source, clustering of passive particles becomes an inevitability over long enough timescales.

This is the primary focus of this thesis: we examine flows which are weakly divergent and the effect this has on the concentration of particles advected within them. For now we limit ourselves to the case where the source of the divergence stems from the flow itself being compressible, but eventually we will examine the case where divergence is introduced through particle inertia. We aim to characterise the clustering of particles, both through the volume of the domain it occupies, and through the rate at which it occurs. Further, we aim to draw a connection between this and the source of the divergence.

2.2 Clustering of particles

2.2.1 Measuring the clustering through particle density

A natural descriptor to describe the clustering of particles is the concentration of particles. Formally this is the number of particles per unit volume. This can be viewed equivalently (up to a factor in mass) as the density of particles in a volume of fluid in the flow. As the volume is compressed we would say the density of particles increases. To describe a mass of fluid we use the term *fluid parcel*. These are defined as a (fixed) mass of fluid that is identifiable throughout the dynamics. Their volume, or shape, may change over the course of the dynamics, but its mass will remain constant (Batchelor, 2000). These fluid parcels can also be viewed as the immediate volume surrounding a point particle and as such the trajectories of particles and fluid parcels can be used interchangeably. Importantly, we distinguish between the density of particles in a fluid volume, ρ , and the density of the fluid, ρ_f , the particles are suspended in itself.

This concept can be extended across the entire domain to form a distribution for the density of particles in the domain. This can also be viewed as a probability density function (PDF) for particles across the domain, where in higher density areas for example, the chance of finding a particle in the area is larger. We are interested in how this particle density evolves over time.

We can describe the density field directly as a passive scalar advected by the flow. Denoting the density field as $\rho(\mathbf{x}, t)$, this satisfies the continuity equation,

$$\partial_t \rho + \nabla \cdot (\rho \mathbf{u}) = 0, \quad (2.2)$$

which for an incompressible flow reduces to the more familiar form

$$\partial_t \rho = -(\mathbf{u} \cdot \nabla) \rho. \quad (2.3)$$

It is important to note that this equation applies both to the fluid itself for a fluid density ρ_f , and also for the transport of the density of particles within the fluid, $\rho(\mathbf{x}, t)$. The latter only satisfies this equation if molecular motion is ignored within the fluid.

Without molecular motion, under pure advection, turbulent flows like the ones we will be interested in, have strong mixing properties. This mixing causes area elements of the density field to undergo repeated stretching and folding in the flow. The result of this is the creation of a filamented structure within the domain. Through this repeated stretching and folding mechanism the created filaments become increasingly long and thin. Figure 2.1, reproduced from [Welander \(1955\)](#) demonstrates the stretching and distortion of an initially uniform passive scalar patch.

The physics behind natural processes are normally not purely advective, and the absence of molecular diffusion can only be safely ignored for length scales less than the Batchelor length scale, $L < \lambda_B$ ([Batchelor, 1959](#)). We describe molecular diffusion mathematically through a diffusion operator, ∇^2 . Adding diffusion to equation (2.2) results in the advection–diffusion equation

$$\partial_t \rho + \nabla \cdot (\rho \mathbf{u}) = \kappa \nabla^2 \rho, \quad (2.4)$$

where κ is the diffusivity. The fluid density ρ_f no longer satisfies this equation, as this no longer simply the transport of mass by the fluid. As such, we will deal only with the particle density field $\rho(\mathbf{x}, t)$ from here.



Figure 2.1: Illustration of an experiment to demonstrate the deformation of an initially square, coloured fluid element from [Welander \(1955\)](#). The fluid was divided in a vessel using a metal grid and a single square dyed to create the initial square patch before motion was induced through stirring. The entire fluid mass was put under rotation to suppress three dimensional turbulence.

Usually to model the advected scalar field for density, we would solve equation (2.4) numerically on a grid. The introduction of diffusion in the equation helps with this also. In the absence of diffusion, the filaments created from the successive stretching and folding of the advected scalar field would eventually shrink below any chosen grid size, causing numerical instability. Adding a suitable amount of diffusion numerically ensures that the scale of these filaments does not shrink below the scale of the grid. The diffusion acts as a smoothing of the scalar field, reducing sharp gradients in the distribution.

While this approach is simple to implement numerically, even with the adjustment to include diffusion, it has a limited scope. The primary issue is that the resolution for the estimate for the density field is tied directly to how well resolved the filaments are, which at the order of the grid scale has been destroyed by the introduced diffusion. This means we only capture large scale patterns for the density field and lose any information on small scale behaviour. In order to resolve behaviour on finer scales we attempt to examine particles individually rather than as a collective in a scalar field.

A commonly used strategy to improve resolution is through backward integration of the trajectories (2.1). This has the advantage that the solution for the density field relies only on the velocity fields resolution, while still yielding an Eulerian field for the density. Backward integration is efficient for finding a quantity accurately at a single time, but for a time series it requires repeated integration of trajectories. This is the primary drawback for the cases considered in this thesis as we require a time series to investigate the underlying scaling behaviour across the evolution of trajectories in time.

An alternative is to model the density using a Lagrangian approach. By solving (2.1) numerically for a large number of particles, we can also estimate the distribution of $\rho(\mathbf{x}, t)$. We make this estimate by dividing the domain into subdomains, or bins, and counting particles within each subdomain. This can be recalculated at successive times to describe the evolving density field of the particles as their motion is simulated over time. While this is a Lagrangian approach in nature, the density field $\rho(\mathbf{x}, t)$ is still an Eulerian quantity that we are estimating.

The Lagrangian approach is straightforward to calculate but comes with some of the same drawbacks as the scalar field advection. The primary issue is that the accuracy of the estimate for the density field depends directly on both the grid size for the domain and the number of particles simulated. In order to get an accurate estimate, a large

number of particles are needed to properly represent areas of high density relative to low areas and their gradients. At the same time, the resolution of the density field estimate is directly dependent on the number of bins used to cover the domain.

In the limit of infinite particles the estimate from binning particles matches the estimate from the advected scalar field. The advantage in this method though comes from the numerical stability. As the scalar field is no longer directly solved for, filaments shrinking below the grid scale does not cause instability in the simulation.

We will see in section 2.5 that due to the nature of the clustering, the distribution for the particles tends to a fractal set which means sharp gradients and small scales are an inevitability. While this no longer causes instability in the method, it is nonetheless problematic for accurately calculating the density field. Due to this shrinking of scales, the Eulerian density field will eventually fail to be resolved by any chosen fixed grid size and number of particles.

Lagrangian frame for density

In a Lagrangian frame, rather than examining functions of time t at a fixed position \mathbf{x} , the particles label \mathbf{a} is the fixed variable. Generally this label is chosen to be the position at an initial time $t = 0$. The motion of particles at an initial position \mathbf{a} to a position \mathbf{x} at a time t is described by the *flow map*

$$\mathbf{x} = \varphi_t(\mathbf{a}). \quad (2.5)$$

We refer to the coordinate system denoted by \mathbf{a} as the label space for the particles.

We can cast the advection of the density in the Lagrangian frame also. To do this we start with the continuity equation (2.2) in the Eulerian frame

$$\partial_t \rho + \nabla \cdot (\rho \mathbf{u}) = 0. \quad (2.6)$$

Expanding this, we note the velocity field is not necessarily divergence free

$$\partial_t \rho(\mathbf{x}, t) = -(\mathbf{u} \cdot \nabla) \rho(\mathbf{x}, t) - \rho(\mathbf{x}, t) (\nabla \cdot \mathbf{u}), \quad (2.7)$$

and rearranging as

$$D_t \rho(\mathbf{x}, t) = -\rho(\mathbf{x}, t) (\nabla \cdot \mathbf{u}), \quad (2.8)$$

where D_t represents the material derivative $D_t(\cdot) = \partial_t(\cdot) + (\mathbf{u} \cdot \nabla)(\cdot)$. This is still in terms of the Eulerian view of density $\rho(\mathbf{x}, t)$, but we are interested in the Lagrangian description of density

$$R(\mathbf{a}, t) = \rho(\varphi_t(\mathbf{a}), t). \quad (2.9)$$

Writing (2.8) in Lagrangian coordinates with $\mathbf{a} = \varphi_t^{-1}(\mathbf{x})$ is natural as the material derivative is the time derivative at fixed label \mathbf{a} in the Lagrangian frame

$$\partial_t R(\mathbf{a}, t) = -R(\mathbf{a}, t) (\nabla \cdot \mathbf{u})(\varphi_t(\mathbf{a}), t). \quad (2.10)$$

We can solve for the Lagrangian density as

$$R(\mathbf{a}, t) = R(\mathbf{a}, 0) e^{-\int_0^t (\nabla \cdot \mathbf{u})(\varphi_s(\mathbf{a}), s) ds}. \quad (2.11)$$

This shows that the Lagrangian density is the integral of the divergence along particle trajectories, $\varphi_t(\mathbf{a})$. This can be thought of as a fluid parcel moving in the flow, where its density changes based on the compression (or expansion) it feels along the trajectory due to the divergence in the flow $\nabla \cdot \mathbf{u}$

For convenience, we define a new variable, the log-density $\delta(\mathbf{a}, t)$ where

$$\delta(\mathbf{a}, t) = \ln \left[\frac{R(\mathbf{a}, t)}{R(\mathbf{a}, 0)} \right] = - \int_0^t (\nabla \cdot \mathbf{u})(\varphi_s(\mathbf{a}), s) ds. \quad (2.12)$$

The Lagrangian log-density $\delta(\mathbf{a}, t)$ is an important quantity and throughout this thesis we characterise clustering by the statistics of $\delta(\mathbf{a}, t)$. We favour the log-density $\delta(\mathbf{a}, t)$, over the Lagrangian density $R(\mathbf{a}, t)$ due to the exponential growth we would need to consider otherwise when dealing with $R(\mathbf{a}, t)$. This form also has an advantage due to the connection to Large Deviations theory we can draw that we outline in the next section.

With this in mind we expand our notion of density to a system of equations for particle position $\mathbf{x}(\mathbf{a}, t)$ and (log) density $\delta(\mathbf{a}, t)$ from equation (2.12). We write this

as

$$\dot{\mathbf{x}}(\mathbf{a}, t) = \mathbf{u}(\mathbf{x}, t), \quad (2.13a)$$

$$\dot{\delta}(\mathbf{a}, t) = -(\nabla \cdot \mathbf{u})(\varphi_t(\mathbf{a}), t), \quad (2.13b)$$

with the initial condition $\mathbf{x}(\mathbf{a}, t = 0) = \mathbf{a}$. This system is the focus of our attention throughout this thesis. By solving this system of equations we can gather statistics on the log-density and begin to describe the behaviour of clustering within the flow. The exact behaviour of the dynamics depends mainly on the properties of \mathbf{u} , but keeping this general for now we can infer some more general properties of the system.

2.3 Large Deviations theory

Building on the Lagrangian description in the last section, we want to use the system (2.13) to describe clustering in turbulent or other complex flows. In such flows, a large ensemble of particles will all experience different levels of divergence as they move in the flow. With this in mind, it makes sense to describe $\delta(\mathbf{a}, t)$ through the PDF $P(\delta, t)$. Working with the PDF gives us a less detailed description of the distribution of density, as we lose information regarding specific areas of the domain by marginalising over space, but we gain a better understanding of the change in the distribution of typical values of $\delta(t)$ over time that we would usually fail to capture.

To describe $P(\delta, t)$, a simple initial approach could be to use Monte Carlo techniques and directly compute the log-density $\delta(t)$ along trajectories and build an estimate for the PDF of $\delta(t)$. In doing this we observe that for increasing time, the variance of the distribution rapidly shrinks around the mean, similar to the distribution in figure 2.2, presented as part of an illustrative example. We also find that on average the log-density grows in time (i.e. $\mathbb{E}_{\mathbf{a}}(\delta(\mathbf{a}, t)) > 0$). As a consequence, for anything other than very short times we cannot say much about the rate of clustering, other than what the expected value over the domain is, and that it seemingly grows.

This picture is clearly incomplete. We know that not every fluid parcel will compress at the same rate as the divergence field is not necessarily uniform across the domain. As such we expect some deviation from this mean value to be present in the distribution, even in the long time limit. It turns out that the PDF for $\delta(t)$ has a special form that follows from large deviations theory, which we justify shortly. Large deviations

gives an approach to characterising such distributions where there is an increasing concentration of probability around the mean value, or equivalently a shrinking variance in the distribution over time.

In its most basic form, large deviations tells us for a sum of a large number of random variables

$$Z_n = \frac{1}{n} \sum_{i=1}^n X_i, \quad (2.14)$$

the form of the PDF for this sum is ¹

$$P(Z_n = z) \asymp \exp(-ns(z)), \quad (2.15)$$

where $s(z)$ is a non-negative convex function called the *rate function* or Cramér function (Ellis, 2005; Dembo and Zeitouni, 2010). These properties exist for the rate function associated with sample means of random variables such as (2.14) as a consequence of Cramér's Theorem (Touchette, 2009). Importantly, if $P(Z_n = z)$ has a mean at μ , then the minimum of the rate function is also at μ . If the PDF is normalised then this minimum is exactly $s(\mu) = 0$. This can be seen from equation (2.15), as in the $n \rightarrow \infty$ limit the expected value is $\mathbb{E}_{n \rightarrow \infty}[Z_n] = \mu$, i.e. $P(Z_n = \mu) = 1$ for $n \rightarrow \infty$. As a consequence the rate function $s(z)$ must have the property $s(\mu) = 0$. This is also known more familiarly as the Law of Large numbers. This exponentially decaying form of the PDF captures the increasing concentration of probability around the mean value over time.

Large deviation applied to a Gaussian sample mean

It is helpful to briefly consider a simple example of large deviation theory for the sums of random variables like the one in equation (2.14). Through this example we can illustrate more clearly the role of the rate function and how it arises. We consider with the sample mean Z_n for a random variable X

$$Z_n = \frac{1}{n} \sum_{i=1}^n X_i. \quad (2.16)$$

¹The symbol \asymp is used here to represent asymptotic equivalence, i.e. equivalent to first order in their exponents such that $x \asymp y \implies \log(x) \sim \log(y)$ (Cover and Thomas, 1991)

In normal applications the distribution for X_i is generally not known a priori, but for the sake of this example we use a known distribution. We assume the random variables X_i are independent and identically distributed following a Gaussian distribution

$$P(X_i = x_i) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(x_i - \mu)^2}{2\sigma^2}\right), \quad (2.17)$$

with mean μ and variance σ . With this distribution we also know a priori the exact distribution for Z_n is

$$P(Z_n = z) = \frac{\sqrt{n}}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{n(z - \mu)^2}{2\sigma^2}\right), \quad (2.18)$$

as the sum of Gaussian random variables is also itself a Gaussian random variable. This form of the PDF is in terms of a decaying exponential, which is exactly the form of the distribution that large deviations predicts as in equation (2.15). We can identify the rate function for the Gaussian sample mean as

$$s(z) = \frac{(z - \mu)^2}{2\sigma^2}. \quad (2.19)$$

This function is convex and has a single minimum at $s(\mu)$. Due to this minimum in the rate function, as the number of samples n grows, the PDF gets increasingly concentrated around the mean value μ . Figure 2.2 shows the PDF for Z_n from (2.18) for multiple values of n along with the rate function $s(z)$ from (2.19). It is easier to see from this sketch how the negative exponential of the rate function $s(z)$ causes the PDF to concentrate around the mean.

One of the key consequences of this form of the PDF is that values further away from the mean are unlikely to be observed for large n as $P(Z_n)$ decays exponentially with n . Only samples such that $Z_n \approx \mu$ have a non-negligible probability of occurring for $n \gg 1$.

Large deviation applied to the log-density $\delta(t)$

Large deviations for random sums also applies to the statistics of $\delta(t)$. This is because (2.12) shows that $\delta(t)$ is the integral of a random process, the continuous analogue to a sum of random numbers. To apply large deviation theory we require a finite (Lagrangian) correlation time for $\nabla \cdot \mathbf{u}$ which we take for granted. It is reasonable to assume this is the case, given that the type of flows we are interested commonly have a




Figure 2.2: Sketch of the PDF for Z_n from (2.18) for multiple values of $n = \{2, 5, 20, 100\}$ along with the rate function $s(z)$ from (2.19). Gaussian distribution used for the sample mean has parameters $X_i \sim \mathcal{N}(1, 1)$.

finite correlation time due to their turbulent nature. This finite correlation time must also be significantly shorter than the time period of interest to ensure uncorrelated samples. With this, for $t \gg 1$, we say $P(\delta, t)$ is of the form

$$P(\delta, t) \asymp e^{-ts(\delta/t)}. \quad (2.20)$$

Expressing $\delta(t)$ this way is a key step in being able to describe the statistics of clustering. This equation describes in general, the way in which the pdf for $\delta(t)$ is collapsing towards a peak at the mean value. Using this, instead of having to work out $P(\delta, t)$ for each t , we can focus on the determination of the t -independent rate function $s(\xi)$.

We can also use large deviation to create an expression for the moments of the distribution, which can be related back to the rate function. The moments offer an alternative characterisation of the statistics of $\delta(t)$ to the one obtained via the PDF. Working with the moments is more convenient in some cases, such as with importance sampling which we will explore in section 2.4. There is little difference in which perspective is used, as knowledge of the moments is essentially equivalent to knowing the PDF, given that the moment-generating function exists. The existence of the moment-generating function is guaranteed by Cramér's Theorem given that a large deviation principle applies to $P(\delta, t)$, which is the only case we are interested in (Touchette, 2009).

To examine the moments of the density, $\mathbb{E}[R^\alpha(t)]$, we calculate

$$\mathbb{E}[R^\alpha(t)] = \int_0^\infty r^\alpha(t)P(r, t)dr. \quad (2.21)$$

This can be expressed in terms of $\delta(t)$ through the relation $R^\alpha(t) = e^{\alpha\delta(t)}$,

$$\mathbb{E} \left[e^{\alpha\delta(t)} \right] = \int_0^\infty e^{\alpha\delta(t)}P(\delta, t)d\delta. \quad (2.22)$$

Combining with the large deviation approximation for the PDF this can be written as

$$\mathbb{E} \left[e^{\alpha\delta(t)} \right] \asymp \int_0^\infty e^{\alpha\delta(t)}e^{-ts(\delta/t)}d\delta. \quad (2.23)$$

This can be rearranged to give

$$\mathbb{E} \left[e^{\alpha\delta(t)} \right] \asymp \int_0^\infty e^{t[\alpha\xi - s(\xi)]}d\xi, \quad (2.24)$$

where $\xi = \delta/t$. Looking at this integral, Laplace's method tells us the main contribution in the large time limit, $t \gg 1$, comes from the value of ξ for which $\alpha\xi - s(\xi)$ is at its maximum. As the rate function $s(\xi)$ is convex and non-negative, this happens exactly at the critical point of $(\alpha\xi - s(\xi))$. In other words the maximum occurs for the value ξ_* that satisfies

$$s'(\xi_*) = \alpha \quad (2.25)$$

Using this,

$$\int_0^\infty e^{t[\alpha\xi - ts(\xi)]}d\xi \approx e^{t[\alpha\xi_* - s(\xi_*)]}, \quad (2.26)$$

which we write as,

$$\mathbb{E} \left[e^{\alpha\delta(t)} \right] \asymp e^{tf(\alpha)}. \quad (2.27)$$

This function $f(\alpha)$ is commonly referred to as the *free energy*. This expression tells us that the moments of the density grow exponentially with time at a rate that we can quantify through the free energy function $f(\alpha)$. We can also link this back to the rate function.

The free energy $f(\alpha)$ is the Legendre-Fenchel transform of the rate function $s(\xi)$ which is given by the transformation

$$f(\alpha) = \sup_{\xi} (\alpha\xi - s(\xi)). \quad (2.28)$$

The Legendre-Fenchel transform is invertible, and can be easily calculated to obtain the rate function again as

$$s(\xi) = \sup_{\alpha} (\alpha\xi - f(\alpha)). \quad (2.29)$$

The duality of the free energy function with the rate function holds in general for large deviation theory. This means that once we can accurately estimate either the free energy, $f(\alpha)$ or the rate function, $s(\xi)$ we will have an estimate for the other in most cases.

2.3.1 Zeros of the free energy and its implications for the growth of moments

It is clear from (2.28) that $f(\alpha) = 0$ for $\alpha = 0$, given that the underlying probability distribution is normalised and hence $s(\mu) = 0$ (Touchette, 2009). It also turns out that $f(\alpha) = 0$ also for $\alpha = -1$ under some assumptions on the domain. This will be a key property of the free energy functions used in this thesis. We show here why this is the case. Consider the joint system

$$d\mathbf{X} = \mathbf{u}(\mathbf{X}, t) + \sqrt{2\kappa} d\mathbf{W} \quad (2.30)$$

$$d\delta = -(\nabla \cdot \mathbf{u})(\mathbf{X}, t), \quad (2.31)$$

where we include the noise term $\sqrt{2\kappa} d\mathbf{W}$ for generality. (2.30) is simply the SDE for the Lagrangian particle position that is the analogue to the Eulerian advection–diffusion seen in (2.4).

The moments M_α can be considered as averaged over the initial positions \mathbf{x}_0 assumed uniformly distributed:

$$M_\alpha(t) = \int \mathbb{E}_{\delta_0, \mathbf{x}_0} [e^{\alpha\delta(t)}] d\mathbf{x}_0, \quad (2.32)$$

where the expectation is taken over the random noise and conditioned on the initial

conditions $\mathbf{X}(0) = \mathbf{x}_0$ and $\delta(0) = \delta_0$. With this,

$$w_\alpha(\mathbf{x}_0, \delta_0, t) = \mathbb{E}_{\delta_0, \mathbf{x}_0}[e^{\alpha\delta(t)}] \quad (2.33)$$

satisfies the backward Kolmogorov equation

$$\partial_t w_\alpha = \mathbf{u} \cdot \nabla w_\alpha - (\nabla \cdot \mathbf{u}) \partial_\delta w_\alpha + \kappa \nabla^2 w_\alpha, \quad (2.34)$$

where we drop the subscript on \mathbf{x}_0 and δ_0 for convenience. Note that the initial condition

$$w_\alpha(\mathbf{x}, \delta, 0) = e^{\alpha\delta} \quad (2.35)$$

and the δ -independence of the coefficients in (2.34) indicate that

$$w_\alpha(\mathbf{x}, \delta, t) = e^{\alpha\delta} g_\alpha(\mathbf{x}, t), \quad (2.36)$$

where, from (2.34), the function g_α satisfies

$$\begin{aligned} \partial_t g_\alpha &= \mathbf{u} \cdot \nabla g_\alpha - \alpha(\nabla \cdot \mathbf{u})g_\alpha + \kappa \nabla^2 g_\alpha \\ &= \nabla \cdot (g_\alpha \mathbf{u}) - (\alpha + 1)(\nabla \cdot \mathbf{u})g_\alpha + \kappa \nabla^2 g_\alpha. \end{aligned} \quad (2.37)$$

Taking the time derivative of (2.32) and using (2.33), (2.36) and (2.37), we obtain that

$$\partial_t M_\alpha = \int (\nabla \cdot (g_\alpha \mathbf{u}) - (\alpha + 1)g_\alpha + \kappa \nabla^2 g_\alpha) \, d\mathbf{x}. \quad (2.38)$$

In unbounded domains and periodic domains – which we use in chapters 4 and 5 – the integrals of $\nabla \cdot (g_\alpha \mathbf{u})$ and $\nabla^2 g_\alpha$ vanish, which leaves

$$\partial_t M_\alpha = - \int (\alpha + 1)g_\alpha \, d\mathbf{x}. \quad (2.39)$$

The right-hand side vanishes for $\alpha = -1$, hence $M_{-1} = \text{const}$. Recall the moments can be related to the free energy function as

$$M_\alpha(t) \asymp e^{tf(\alpha)}, \quad (2.40)$$

and as a result we can see that for $M_{-1} = \text{const}$ to hold, $f(-1) = 0$ must be the case. The next section highlights the difference in the Eulerian and Lagrangian frame for the

moments, and the implication of these fixed points in the free energy function.

Physical interpretation of the growth of moments

We have seen that the free energy has certain properties, such as $f(-1) = f(0) = 0$. We also know that the free energy function is related to the moments of the density. In this section we will highlight the difference in the Eulerian and Lagrangian frame for the moments, and the physical implication of the zeros of the free energy function.

We start with the Eulerian moments for the density of particles

$$M_E(\alpha, t) = \int \rho^\alpha(\mathbf{x}, t) d\mathbf{x}. \quad (2.41)$$

In order to change perspective to the Lagrangian frame we first outline the relation between Eulerian and Lagrangian frames for volume elements. In the Lagrangian frame, a n -dimensional volume element $d\mathbf{a}$ outlined by particles at positions \mathbf{a} at a time $t = 0$, is mapped to a volume $d\mathbf{x}$ with particles now at positions \mathbf{x} at time t . The two volumes are related through the flow map φ_t

$$d\mathbf{x} = \left| \frac{\partial \varphi_t(\mathbf{a})}{\partial \mathbf{a}} \right| d\mathbf{a} \quad (2.42)$$

where $|A| = \det(A)$, and we identify the Jacobian matrix

$$J_t(\mathbf{a}) = \frac{\partial \varphi_t(\mathbf{a})}{\partial \mathbf{a}}. \quad (2.43)$$

This can be thought of equivalently as a change of variables between position space \mathbf{x} and label space \mathbf{a} , or between Eulerian and Lagrangian coordinates. By conservation of mass, if this initial volume $d\mathbf{a}$ encloses a mass dm , then at time t it also encloses a mass dm . The density of the parcel $R(\mathbf{a}, t)$, with $dm = R(\mathbf{a}, t) d\mathbf{x}$ can be written as

$$R(\mathbf{a}, t) = \frac{1}{|J_t(\mathbf{a})|} R(\mathbf{a}, 0). \quad (2.44)$$

From this, as a brief aside, it is straightforward to show equation (2.10) is valid. We take a derivative in time

$$\partial_t R(\mathbf{a}, t) = -R(\mathbf{a}, 0) \frac{1}{|J_t(\mathbf{a})|^2} \partial_t |J_t(\mathbf{a})|. \quad (2.45)$$

Using Jacobi's formula ($\partial_t |A| = |A| \cdot \text{tr}(A^{-1} \partial_t A)$) this becomes

$$\partial_t R(\mathbf{a}, t) = -R(\mathbf{a}, 0) \frac{1}{|J_t(\mathbf{a})|^2} |J_t(\mathbf{a})| \text{tr}(J_t^{-1}(\mathbf{a}) \partial_t J_t(\mathbf{a})). \quad (2.46)$$

By symmetry of second derivatives, $\partial_t J_t(\mathbf{a}) = \partial_{\mathbf{a}} \dot{\varphi}_t(\mathbf{a})$, and recalling $\mathbf{x} = \varphi_t(\mathbf{a})$ and so $\mathbf{u} = \dot{\varphi}_t(\mathbf{a})$ this reduces to

$$\partial_t R(\mathbf{a}, t) = -R(\mathbf{a}, 0) \frac{1}{|J_t(\mathbf{a})|} \cdot \text{tr} \left(\frac{\partial \mathbf{u}}{\partial \mathbf{x}} \right) \Big|_{\mathbf{x}=\varphi_t(\mathbf{a})} = -R(\mathbf{a}, t) (\nabla \cdot \mathbf{u})(\varphi_t(\mathbf{a}), t), \quad (2.47)$$

matching the expression shown previously in equation (2.10).

Now returning to equation (2.41), changing perspective to the Lagrangian frame of a parcel of fluid moving with the flow this becomes

$$M_E(\alpha, t) = \int R^\alpha(\mathbf{a}, t) \left| \frac{\partial \mathbf{x}}{\partial \mathbf{a}} \right| d\mathbf{a}, \quad (2.48)$$

where $R(\mathbf{a}, t) = \rho(\varphi_t \mathbf{a}, t)$. By changing from Eulerian to Lagrangian coordinates we effectively weight each parcel of fluid we follow by its own volume according to equations (2.44) and (2.42). From here we take $R(\mathbf{a}, 0) = 1$, for ease of notation, as any other value is just a constant that is carried through outside the integral.

Changing from an average over Lagrangian volumes \mathbf{a} to a probabilistic average over densities $r \in [R, R + dR]$, can be done by first identifying the cumulative distribution function (CDF)

$$A(r, t) = \frac{1}{|\Omega|} \int_{\Omega} \Theta(R(\mathbf{a}) - r) d\mathbf{a}, \quad (2.49)$$

where Θ is the Heaviside Step function, and Ω is the domain. From the CDF, we can find the PDF through $P(r, t) = d/dr A(r, t)$, yielding

$$P(r, t) = \frac{1}{|\Omega|} \int_{\Omega} \delta(R(\mathbf{a}) - r) d\mathbf{a}, \quad (2.50)$$

where $\delta(R(\mathbf{a}) - r)$ is the dirac delta function. Finally, using the sifting property of the dirac delta function

$$R^{\alpha-1}(\mathbf{a}) = \int r^{\alpha-1} \delta(R(\mathbf{a}) - r) dr, \quad (2.51)$$

results in (up to the constant Ω)

$$\int R^{\alpha-1} d\mathbf{a} = \int \int r^{\alpha-1} \delta(R(\mathbf{a}) - r) dr d\mathbf{a}, \quad (2.52)$$

$$= \int r^{\alpha-1} p(r, t) dr. \quad (2.53)$$

Combining this with equation (2.48) gives,

$$M_E(\alpha, t) = \int r^{\alpha-1}(\mathbf{a}, t) P(r, t) dr. \quad (2.54)$$

Recalling the log-density $R(t) = e^{\delta(t)}$, the large deviation approximation for $P(\delta, t)$ (2.20), and that $P(\delta, t) d\delta = P(r, t) dr$ by change of variables for probability density functions we have,

$$M_E(\alpha, t) \asymp \int e^{(\frac{\delta}{t})t(\alpha-1)} e^{-ts(\frac{\delta}{t})} d\delta, \quad (2.55)$$

which gives,

$$M_E(\alpha, t) \asymp e^{tf_L(\alpha-1)}. \quad (2.56)$$

It is important to highlight that the free energy here corresponds to the Lagrangian perspective, denoted by f_L and the Eulerian moments are expressed in terms of it. In the remainder of this section, when discussing the free energy (or any other quantity) from the Lagrangian perspective a subscript f_L will be used and when discussing the from the Eulerian perspective a subscript f_E will be used. The Eulerian moments can be expressed in terms of the Eulerian free energy function also through the defining relation (2.27)

$$M_E(\alpha) \asymp e^{tf_E(\alpha)}, \quad (2.57)$$

and similarly for the Lagrangian moments

$$M_L(\alpha) \asymp e^{tf_L(\alpha)}. \quad (2.58)$$

Equations (2.56), (2.57), (2.58) imply the Eulerian moments are related to the La-

grangian moments as

$$M_E(\alpha) = M_L(\alpha - 1). \quad (2.59)$$

We know $M_E(0)$ is the volume of the fluid, which is constant, and also that $M_E(1)$ is the mass of the particles, which is also constant. Therefore $f_L(-1) = f_L(0) = 0$. This matches the result from the previous section.

The rate function is also known to be convex (Touchette, 2009), and the Legendre transform of a convex function, is itself convex. Thus $f_L(\alpha) < 0$ for $-1 < \alpha < 0$, and $f_L(\alpha) > 0$ otherwise. This also tells us that $f'_L(0) > 0$.

As the rate function $s_L(\xi)$ and free energy $f_L(\alpha)$ are Legendre-Fenchel transforms of each other, we have that

$$\xi = f'_L(\alpha), \quad \alpha = s'_L(\xi). \quad (2.60)$$

The minimum of the rate function occurs at

$$s'_L(\xi^*) = 0, \quad (2.61)$$

which corresponds directly to

$$\xi^* = f'_L(0) > 0. \quad (2.62)$$

By definition of the rate function, this value ξ^* is the mean of the PDF $P(\delta, t)$ and is the value that the PDF collapses to as $t \rightarrow \infty$. This implies that $\delta/t \rightarrow \xi^*$ almost surely by the law of large numbers. Recalling that the minimum of the rate function is the mean value of the underlying distribution for the log-density and combining with equation (2.60) means that the slope given by $f'_L(0)$ is the mean compression in the system.

Physically, density must be conserved over the entire domain. According to $f'_L(0) > 0$ though, density is growing on average across trajectories. This is explained by the fact that from a Lagrangian perspective, more particles end up in convergent regions of the flow than end up in divergent regions (Falkovich et al., 2001).

From an Eulerian perspective a similar logic applies. Here $f_E(\alpha) < 0$ for $0 < \alpha < 1$ and thus $f'_E(0) < 0$. This implies through equation (2.62) that the average density

is actually decreasing on across the domain. This is due to the same reason as the Lagrangian perspective. Density is decreasing in regions of the domain that are growing in size, and correspondingly increasing in shrinking areas of the domain.

All of this implies that the system is in fact dissipative in nature. Since the density is shrinking on average across the domain i.e. $\rho(\boldsymbol{x}, t) \rightarrow 0$ as $t \rightarrow \infty$, all the particles must end up clustered on a zero-dimensional set within the flow (Ott, 2002; Holden, 1986). In general, when discussing the free energy we will be referring to the Lagrangian perspective from now and thus will drop the subscript for convenience $f = f_L$,

2.4 Importance sampling

Using large deviation theory we can examine the behaviour of the moments of the density distribution over time and in theory account for the exponentially decaying nature of the tails. To do this we want to estimate either the free energy function $f(\alpha)$ or the rate function $s(\xi)$ for $\delta(t)$. This is done through Monte Carlo sampling to approximate either the moments if we are trying to estimate $f(\alpha)$, or by approximating the PDF to estimate $s(\xi)$. In order to sample the density we will follow Lagrangian trajectories of fluid parcels and calculate the density of each parcel regularly as it is advected by the flow. This is done through equations (2.13).

In general we are interested in tails of the PDF or moments for large $|\alpha|$, but these are difficult to sample by definition as we have seen in the previous section. After a sufficient time, only parcels that experience, on average, moderate compression have a non-negligible chance to be observed, i.e. trajectories such that $\delta(t) \approx \mathbb{E}[\delta(t)] > 0$ as large deviation predicts. This poses an inherent problem with our ability to resolve the PDF numerically.

By the nature of the problem, the fact that large deviation theory applies, means that the tails of this distribution shrink exponentially. Combining this with the error scaling of Monte Carlo methods for n samples being $n^{-1/2}$, means that an extremely large number of samples are needed to estimate the areas of either $f(\alpha)$ or $s(\xi)$ corresponding to the tails of the distribution.

We will choose to characterise the distribution through the moments and thus we are interested in estimating the free energy function $f(\alpha)$. We can then deduce the rate function $s(\xi)$ using $f(\alpha)$ through the Legendre transform. We will see that in order to estimate $f(\alpha)$, even for moderate values of $|\alpha|$, requires an extensive number of samples

to resolve accurately due to the exponential decay in the tails. To work around this we implement an importance sampling method to focus our computation on these under sampled regions.

2.4.1 Importance sampling algorithm

We estimate the free energy function $f(\alpha)$ through the relation

$$f(\alpha) = \lim_{t \rightarrow \infty} \frac{1}{t} \log \mathbb{E} \left[e^{\alpha \delta(t)} \right] \quad (2.63)$$

This comes from simply rearranging equation (2.27) to isolate $f(\alpha)$ and specifying that we are interested in the estimate in the late time limit when the curve has converged to the true value. Numerically we estimate $\mathbb{E} [e^{\alpha \delta(t)}]$ by simulating an ensemble of trajectories $(\mathbf{X}_i(t), \delta_i(t))$, $i = 1, \dots, N$, according to equation (2.13) along their respective paths. We then calculate the expectation as the empirical mean

$$S_N(t) = \frac{1}{N} \sum_{i=1}^N e^{\alpha \delta_i(t)} = \frac{1}{N} \sum_{i=1}^N q_i(t), \quad (2.64)$$

where α is fixed beforehand. This expression is accurate for small values of α , but for $|\alpha| \gtrsim 1$, this method provides poor results for resolving $f(\alpha)$ as we will see later. There is a large error in this brute force estimate due to the main contribution comes from a small number of terms in equation (2.64). These terms are extremely rare realisations of trajectories that give rise to far larger values of $e^{\alpha \delta(t)}$ than a typical observation. These extreme realisations become exponentially more rare with increasing $|\alpha|$ which makes them practically impossible to sample outside of a narrow range around $\alpha = 0$.

In order to combat the error in the brute force estimate, we use an importance sampling algorithm based on an established method of pruning and cloning trajectories of interest (Grassberger, 2002; Tailleur and Kurchan, 2007). The main idea to remedy the problem with a small effective sample size is to replace samples that are dominant in equation (2.64) by several copies of the sample, with a smaller weight attached to them.

Formally to do this, we take the expression for $S_N(t)$ in equation (2.64), and introduce a weighting ω_i to each realisation $q_i(t)$, initially set to unity. This creates a new

expression for the ensemble mean

$$S_N(t) = \frac{1}{N} \sum^N \omega_i e^{\alpha \delta_i(t)} = \frac{1}{N} \sum^N Q_i(t). \quad (2.65)$$

After a set number of time steps in the numerical integration of $\delta(t)$, the combined value $Q_i(t)$ is compared to the ensemble mean $S_N(t)$. If $Q_i(t) > P_c S_N(t)$, where P_c is a parameter used to control the degree of variance reduction desired in the method, the realisation is cloned. This entails an additional particle being created at the location of the original, and being randomly perturbed very slightly from this location. This new particle is integrated forward with the other realisations in further time steps, and follows its own path from this point onwards. To ensure the overall statistics of $S_N(t)$ (and also of $P(\delta, t)$) are unchanged by this, for every cloning event the weighting ω_i is halved in magnitude for both the original and its clone. The perturbation of the particles location is performed as otherwise the two realisations would follow the exact same path going forward. Testing of the algorithm has shown as long as this ‘kick’ is small enough, it has a negligible effect on the final resulting estimate for the free energy curve, but it is necessary to take advantage of the cloning in the importance sampling algorithm.

As well as cloning particles, realisations are pruned with probability $1/2$ if $Q_i(t) < S_N(t)/P_c$. Pruning here is simply removal of the realisation from the ensemble. If the realisation survives, its weight is subsequently doubled to preserve the statistics of S_N . In addition to the systematic cloning and pruning, an additional step of randomly cloning/pruning realisations is used to ensure the population is kept constant.

We can further extend the method to allow multiple clones to be created at a cloning step depending on how much greater than the ensemble average a realisation is. To do this $n = (\lfloor Q_i(t)/P_c S_N(t) \rfloor + 1)$ clones of the original realisation are created, and the weights subsequently divided by $1/n$. The effect of this multiple cloning was found to have no adverse effect on the statistics and worked only to improve efficiency of the sampling.

This method comes with multiple advantages over simply approaching this problem with a brute-force technique. The clearest advantage is the efficiency of the method, allowing the exceptionally rare realisations that dominate (2.64) for each value of α to be targeted specifically for sampling in a simulation aimed at a given α . These realisations were previously practically unobtainable due to the large cost of simulating

enough samples to observe them. The drawback is now that more simulations must be run to create an array of values for $f(\alpha)$, rather than a single simulation to draw sample from. This is not detrimental though as the cost of such simulations is drastically cheaper, and as each is independent, they can be easily run in parallel.

On implementation in each case the algorithm was used for, the effect of the frequency of resampling and the magnitude of the perturbation applied to cloned trajectories are investigated. The results were found to be independent of the magnitude of the kick as long as it was small enough. To avoid any bias from the kick it is set to perturb the location by a randomly drawn magnitude given by $\mathcal{N}(0, L \times 10^{-3})$, where L is the domain size. The resampling frequency requires tuning, if it is undertaken too frequently the dynamics do not get a chance to explore the phase space, and if it is not frequent enough it is ineffective. In general it was found for the cases in this thesis that to perform well, resampling every $\mathcal{O}(10)$ steps was a good range.

2.4.2 Applying importance sampling to a problem with a known distribution

As a demonstration of the effectiveness of importance sampling algorithm, we apply it here to a simple example with a known solution. In particular, this example will demonstrate the ability of the algorithm to expand the resolved range, rather than simply improving accuracy in the part of distribution that can already be estimated.

The equation we chose to examine is Geometric Brownian Motion (GBM). The general form for GBM is

$$dX_t = X_t[\mu dt + \sigma dW_t], \quad (2.66)$$

where W_t denotes the standard Wiener process. Using Ito's formula with the change of variable $Y = \ln X$, this process has the solution $X_t = \exp((\mu - \sigma^2/2)t + \sigma W_t)$. We will for simplicity in this case define $\mu = \sigma^2/2$. Note that with this, the SDE (2.66), is simply the Stratonovich SDE

$$dX_t = \sigma X_t \circ dW_t, \quad (2.67)$$

with solution $X_t = e^{\sigma W_t}$. The moments for this can be readily calculated as

$$\mathbb{E}[X^\alpha(t)] = X_0^\alpha \mathbb{E}[e^{\alpha\sigma W_t}]. \quad (2.68)$$

Using $\mathbb{E}[e^a] = e^{\sigma^2/2}$ for $a \sim N(0, \sigma^2)$ this becomes

$$\mathbb{E}[X^\alpha(t)] = X_0^\alpha e^{(\alpha\sigma)^2 t/2}. \quad (2.69)$$

We know from the large deviations assumption (2.27) that the moments scale like $\mathbb{E}[X^\alpha] \asymp e^{f(\alpha)t}$, and so we identify the free energy function associated with GBM as

$$f(\alpha) = \frac{(\alpha\sigma)^2}{2} \quad (2.70)$$

Calculating a rate function is similarly straightforward, we first need to find the PDF for GBM (given by (2.67)). Using the change of variables $Y_t = \ln X_t$, we find

$$dY_t = \frac{1}{X_t} dX_t = \sigma \circ dW_t. \quad (2.71)$$

Solving gives

$$Y_t = Y_0 + \sigma W_t. \quad (2.72)$$

The PDF for W_t is by definition $\mathcal{N}(0, t)$ and so Y_t is distributed as $Y_t \sim \mathcal{N}(Y_0, \sigma^2 t)$. Thus the PDF for X_t follows a log-normal distribution, which we can write as

$$P_X(x, t) = P_Y(y, t) \frac{dY}{dX} = \frac{1}{\sqrt{2\pi\sigma^2 t}} \exp\left(-\frac{(\ln X_t - \ln X_0)^2}{2\sigma^2 t}\right) \frac{1}{X_t}. \quad (2.73)$$

The exponential contribution will be dominant with time and so the distribution scales as

$$P_X(x, t) \asymp \exp\left(-\frac{\ln^2 X_t}{2\sigma^2 t}\right). \quad (2.74)$$

Recalling again from large deviations theory we expect the PDF to scale as $P(t^{-1} \ln X_t = \xi) \asymp e^{-ts(\xi)}$ for a rate function $s(\xi)$, we can rearrange the distribution to reflect this,

$$P_X(x, t) \asymp e^{-\frac{\ln^2 X_t}{2\sigma^2 t}} = e^{-ts\left(\frac{\ln X_t}{t}\right)}, \quad (2.75)$$

and identify the rate function as

$$s(\xi) = \frac{\xi^2}{2\sigma^2}. \quad (2.76)$$

As a consistency check, we would expect the rate function $s(\xi)$ and free energy $f(\alpha)$ to be Legendre conjugates of each other. From (2.60), the variables are related as $\xi = f'(\alpha) = \sigma^2\alpha$. Applying a Legendre transform to (2.76) gives

$$f(\alpha) = \sup_{\xi} (\alpha\xi - s(\xi)) = \sigma^2\alpha^2 - \frac{(\sigma^2\alpha)^2}{2\sigma^2} = \frac{(\alpha\sigma)^2}{2}, \quad (2.77)$$

which matches what we found previously in equation (2.70).

With these two analytical results for the rate function and free energy, we can evaluate the improvement of the importance sampling algorithm over regular brute-force sampling. To do this we numerically solve sample paths of GBM, and form a distribution of the values at a final time t . From this we can recover a numerically calculated rate function with the relation

$$s(\xi) = - \lim_{t \rightarrow \infty} \frac{1}{t} \log(P(x, t)), \quad (2.78)$$

where $\xi = t^{-1} \ln x$ and $P(x, t)$ is estimated by binning the trajectories. We also recover a free energy function from this via the Legendre transform (2.28). Note that this will be a finite time approximation only to the true $t \rightarrow \infty$ rate function.

The paths are computed using the exact solution $X_t = e^{\sigma W_t}$, where the $W_t = \sqrt{\Delta t} \eta$ where $\eta \sim \mathcal{N}(0, 1)$.

A second estimate for the free energy is then computed using the importance sampling algorithm outlined in the previous section. Rather than a single ensemble of trajectories, a number of smaller batches of trajectories will need to be computed, with each batch focusing on a single value of α . 15 values of α evenly spaced in the interval $\alpha \in [-2, 2]$ were used in this example. Combining these values we construct an estimate for the entire free energy function $f(\alpha)$. Each of these values of $f(\alpha_i)$ are computed as outlined in section 2.4.1.

The comparison of the importance sampling algorithm to brute-force sampling is shown in figure 2.3. The first plot, labelled a), shows the computed estimates for both the brute-force and importance sampled rate functions relative to the underlying PDF (the brute-force rate function shown is directly computed from this PDF). This is pre-

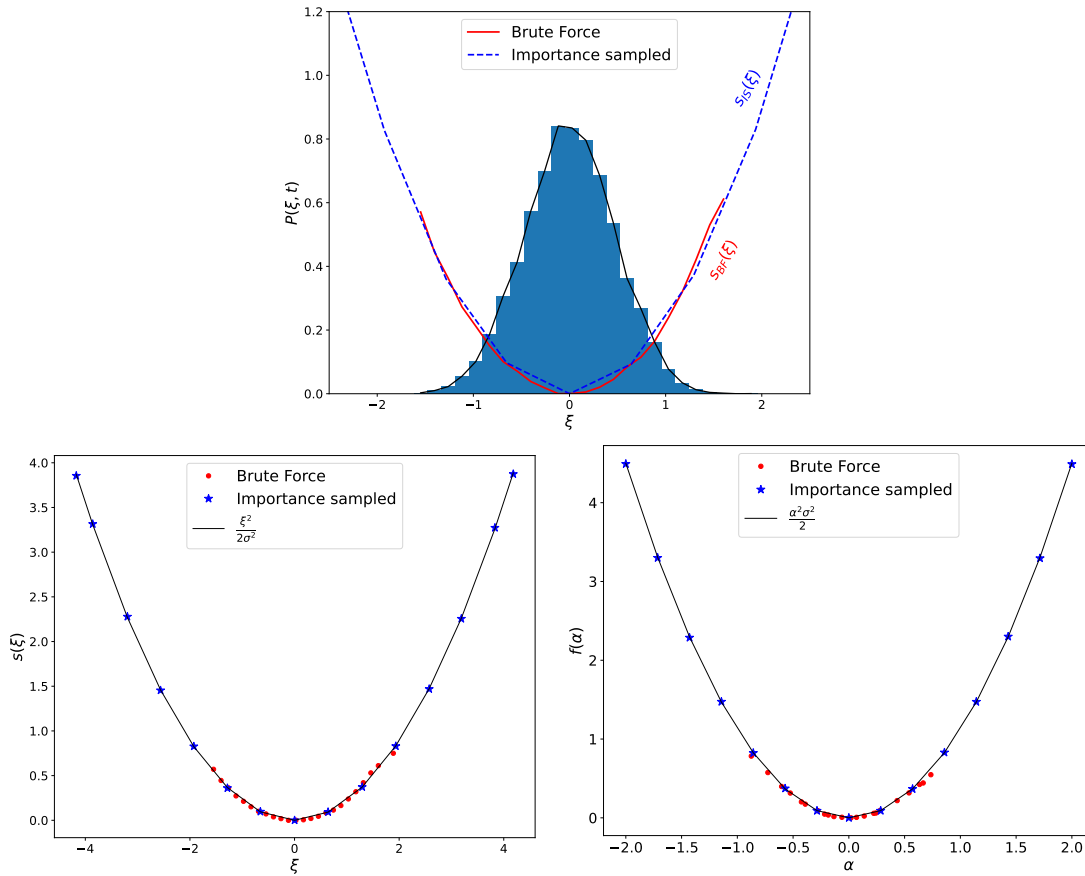


Figure 2.3: Figure a) shows the PDF for $X_t = \exp(\sigma W_t)$ after $N = 10,000$ steps, with $\sigma = 1.5$, $\Delta t = 10^{-3}$ scaled to $\xi = \ln X_t/t$. Rate functions computed with both brute-force sampling ($s_{BF}(\xi)$) and the importance sampling algorithm ($s_{IS}(\xi)$) are superimposed. Figures b) and c) show the computed values of the rate function and free energy respectively for X_t using both brute-force and importance sampling, compared to the analytical expressions in (2.70) and (2.76)

sented to demonstrate the exponential decay of the tails and the difficulty in sampling the tail regions normally, while showing the improvement in resolving information in these regions with the importance sampling algorithm.

The importance sampled rate function in a) is computed through the numerical Legendre transform of the importance sampled free energy. To use the Legendre transform we take advantage of the relation (2.60), $\xi = f'(\alpha)$ to find a set of discrete values of ξ that correspond to the set of points $\{\alpha_1, \dots, \alpha_n\}$ that we have a value for $f(\alpha_i)$. To achieve this numerically we take the simple gradient approximated by the second order central difference

$$\xi_i = \frac{f(\alpha_{i+1}) - f(\alpha_{i-1})}{\alpha_{i+1} - \alpha_{i-1}}, \quad (2.79)$$

where we use the first order approximation for boundary values, $i = 1$ and $i = n$. With these values for ξ we use the Legendre transform to find the rate function at these points,

$$s(\xi_i) = \xi_i \alpha_i - f(\alpha_i). \quad (2.80)$$

The plots labelled b) and c), show the comparison of computed data points for both $f(\alpha)$ and $s(\xi)$ explicitly, with the exact solution found in equations (2.70) and (2.76). These plots again show the ability of the importance sampling algorithm to capture a much larger domain than the traditional brute-force sampling. The brute-force sampled points are all concentrated at the minimum of the rate function. This is a consequence of the law of large numbers, as the distribution they are sampling decays to its mean, which is the expected value $Y_0 = \ln(X_0 = 1) = 0$.

For these results, both curves were constructed with the same number of total trajectories, with 15,000 trajectories for the brute-force simulation, and 15 batches of 1000 trajectories for the importance sampling algorithm (one batch for each point shown in plots b) and c) in figure 2.3). The importance sampling algorithm can be used in this way to achieve a greater accuracy in the already resolvable domain of the free energy curve, for a lower computational cost. The cut-off parameter P_c was set as 2 for this experiment, as is the case in the particle results shown later in the thesis. The algorithm can also be used to extend the resolved domain of the free energy curve rather than simply improve the accuracy. By increasing the number of trajectories per

batch and spreading over a larger domain, an accurate estimate of the free energy curve over a much larger domain can be resolved. The total cost will increase but a much more efficient return on accuracy is achieved compared to brute-force sampling. This is how we will use the algorithm in this thesis, with the intent to resolve areas of the free energy curve that would otherwise be prohibitively expensive, rather than simply improving accuracy in already resolvable regions.

2.5 Lyapunov exponents

2.5.1 Calculating Lyapunov exponents

We have introduced the concept of Lagrangian density in section 2.2.1 and described how it changes with time by equation (2.12). In particular we can quantify the mean compression relatively easily using this equation. With large deviation statistics we can extend this picture to capture the compression rate in general $\xi = \delta/t$ and not just the mean compression.

The compression rate ξ describes the rate at which the volume occupied by nearby particles gets compressed. In general, one can be interested not only in this change in volume but more generally in the change of length, area or n-volume. This is measured by the finite-time Lyapunov exponents (FTLE) which we will briefly outline in this section. We have stated in section 2.3.1 how the change in volume elements is described by the Jacobian matrix

$$d\mathbf{x} = \left| \frac{\partial \varphi_t(\mathbf{a})}{\partial \mathbf{a}} \right| d\mathbf{a}. \quad (2.81)$$

More generally, we are interested in the evolution of the Jacobian matrix itself. Looking at the separation between two initially close points \mathbf{a} and $\mathbf{a} + \delta\mathbf{a}$ as they evolve under the flow map to $\mathbf{x}(t) = \varphi_t(\mathbf{a})$ and $\mathbf{x}(t) + \delta\mathbf{x}(t) = \varphi_t(\mathbf{a} + \delta\mathbf{a})$ respectively. We can write the first terms of the Taylor expansion for $\varphi_t(\mathbf{a} + \delta\mathbf{a})$ at \mathbf{x} as

$$\varphi_t(\mathbf{a} + \delta\mathbf{a}) = \varphi_t(\mathbf{a}) + \frac{\partial \varphi_t}{\partial \mathbf{a}} \delta\mathbf{a}, \quad (2.82)$$

which simplifies to

$$\delta \mathbf{x}(t) = \frac{\partial \varphi_t}{\partial \mathbf{a}} \delta \mathbf{a}, \quad (2.83)$$

$$= J_t(\mathbf{a}) \delta \mathbf{a}. \quad (2.84)$$

Equation (2.83) highlights the role of the Jacobian matrix $J_t(\mathbf{a})$ in describing the separation of nearby particles over time.

We want to examine the rate of change of the separation vectors $\delta \mathbf{x}(t)$. To do this we use the equations for a single pair of particles

$$\dot{\mathbf{x}} = \mathbf{u}(\mathbf{x}, t), \quad (2.85)$$

$$\dot{\mathbf{x}} + \dot{\delta \mathbf{x}} = \mathbf{u}(\mathbf{x} + \delta \mathbf{x}, t), \quad (2.86)$$

with the initial conditions as before, $\mathbf{x}(0) = \mathbf{a}$, $\delta \mathbf{x}(0) = \delta \mathbf{a}$. Taylor expanding and rearranging in a similar way to (2.83) yields

$$\dot{\delta \mathbf{x}} = (\nabla \mathbf{u}) \delta \mathbf{x}, \quad (2.87)$$

where $(\nabla \mathbf{u})_{i,j} = \partial_j u_i$. This ODE for the separation vector is referred to as the *tangent dynamics*. Equation (2.87) is a linear system of ODEs and so we can express the general solution using a fundamental matrix $\Phi(\mathbf{a}, t)$ as $\delta \mathbf{x}(t) = \Phi(\mathbf{a}, t) \mathbf{c}$ for some arbitrary constant vector \mathbf{c} . It is straightforward to identify $\mathbf{c} = \delta \mathbf{a}$ from the initial condition, and that by definition $\Phi(\mathbf{a}, 0) = \mathbb{I}$. Combining this we have $\delta \mathbf{x}(t) = \Phi(\mathbf{a}, t) \delta \mathbf{a}$ which when compared to equation (2.83) implies

$$\Phi(\mathbf{a}, t) = J_t(\mathbf{a}). \quad (2.88)$$

We can generalise all of this to more than one tangent vector by defining the matrix

$$Z(t) = [\delta \mathbf{x}_1(t), \dots, \delta \mathbf{x}_n(t)]^T \quad (2.89)$$

where the rows are the individual separation vectors. These vectors also outline an initial volume at $Z_0 = [\delta \mathbf{a}_1, \dots, \delta \mathbf{a}_n]$ such as the one in equation (2.81). We denote this volume at time t by its determinant $|Z(t)|$. The tangent dynamics in n dimensions also

reads

$$\dot{Z} = (\nabla \mathbf{u})Z, \quad (2.90)$$

where the solutions still satisfy the fundamental matrix as

$$Z(t) = J_t(\mathbf{a})Z_0, \quad J_0(\mathbf{a}) = \mathbb{I}. \quad (2.91)$$

For convenience we will assume that each initial separation vector $\delta \mathbf{a}_i$ is of unit length and thus $\|Z_0\| = 1$. There is no loss of generality through doing this as any other length would simply amount to a constant factor being present.

Geometrically, the fundamental matrix (which is the Jacobian matrix in this case) describes the evolution of an n -volume in the dynamics phase space. This change is a result of the deformation caused by the surrounding flow (as described by the Jacobian matrix). Starting with the length of one vector of this n -volume, or equivalently two nearby particles, the distance d typically grows exponentially in time as

$$d_1(t) = e^{\lambda_1 t}. \quad (2.92)$$

Similarly, the area that is outlined by two vectors grows like

$$d_2(t) = e^{(\lambda_1 + \lambda_2)t}. \quad (2.93)$$

The volume outlined by three vectors grows like $e^{(\lambda_1 + \lambda_2 + \lambda_3)t}$, and so on. This can be generalised as the k -volume outlined by the k vectors growing as

$$d_k(t) = e^{\sum_{i=1}^k (\lambda_i)t}. \quad (2.94)$$

This can be used to estimate the sum of the first k Lyapunov exponents by the growth rate of a k -volume element in the dynamics. The term k -volume here meaning simply the volume outlined by k vectors in an n -dimensional space, similar to how a pair of vectors would outline an area in three dimensional space. We define the finite time Lyapunov exponent (FTLE) as

$$\lambda_i = \frac{1}{t} \ln \frac{d_i(t)}{d_i(0)}. \quad (2.95)$$

The FTLE depends implicitly on the local aspects of the flow that it is initialised at. The long time limit of this equation defines the Lyapunov exponents

$$\hat{\lambda}_i = \lim_{t \rightarrow \infty} \frac{1}{t} \ln \frac{d_i(t)}{d_i(0)}. \quad (2.96)$$

The exponents are ordered in order of magnitude, with the largest first. We have included the true definition for Lyapunov exponents for completeness here, but in the flows we encounter in this thesis we will assume that after a sufficiently long time, the FTLEs will provide a reasonable estimate of the actual Lyapunov exponents we seek and so we only concern ourselves with the FTLEs.

The Lyapunov exponents are intrinsically related to the expansion or contraction of directions in the phase space. As such we can say a volume $|Z(t)|$ grows (or contracts) exponentially with a rate equal to the sum of the Lyapunov exponents, or

$$\frac{\partial}{\partial t} |Z(t)| = \frac{\partial}{\partial t} |Z(0)| e^{\sum_{i=1}^n (\lambda_i) t} = |Z(t)| \sum_{i=1}^n (\lambda_i) \quad (2.97)$$

Equivalently, a change in a volume element is given by

$$\frac{\partial}{\partial t} |Z(t)| = \frac{\partial}{\partial t} \left(|\delta \mathbf{x}_1(t)| \cdot |\delta \mathbf{x}_2(t)| \cdots |\delta \mathbf{x}_n(t)| \right), \quad (2.98)$$

which with the chain rule and Jacobi's formula ($(\partial_t |A| = |A| \cdot \text{tr}(A^{-1} \partial_t A))$) gives on the LHS

$$\begin{aligned} & \left(|\delta \mathbf{x}_1(t)| \cdot |\delta \mathbf{x}_2(t)| \cdots |\delta \mathbf{x}_n(t)| \right) \left(\frac{1}{\delta \mathbf{x}_1(t)} \frac{\delta \mathbf{x}_1(t)}{\partial t} + \frac{1}{\delta \mathbf{x}_2(t)} \frac{\delta \mathbf{x}_2(t)}{\partial t} + \cdots + \frac{1}{\delta \mathbf{x}_n(t)} \frac{\delta \mathbf{x}_n(t)}{\partial t} \right) \\ &= |Z(t)| \left(\frac{\delta \mathbf{u}_1(t)}{\delta \mathbf{x}_1(t)} + \frac{\delta \mathbf{u}_2(t)}{\delta \mathbf{x}_2(t)} + \cdots + \frac{\delta \mathbf{u}_n(t)}{\delta \mathbf{x}_n(t)} \right) \\ &= |Z(t)| (\nabla \cdot \mathbf{u}). \end{aligned} \quad (2.99)$$

The Lyapunov exponents are the result of the average growth of volume elements over time, and thus to compare equation (2.97) and (2.99) requires averaging in time along the trajectory the divergence experienced in equation (2.99) before equating them to find

$$\sum_{i=1}^n \lambda_i = \overline{\nabla \cdot \mathbf{u}}. \quad (2.100)$$

This average along trajectories is equivalent to the log-density defined in equation (2.12) as the integral along a trajectory for the divergence. Thus, there is an explicit link between the sum of the Lyapunov exponents and the log-density (2.12). We know from section 2.3.1 that the system is dissipative ($\dot{\delta}(t) \approx -\overline{\nabla \cdot \mathbf{u}} > 0$), and so, on average, volumes contract over time. This implies through equation (2.100) that the sum of the Lyapunov exponents in the system will be negative ($\sum_{i=1}^n \lambda_i < 0$) also. We write this link explicitly as

$$\frac{\delta}{t} = - \sum_{i=1}^n (\lambda_i) \quad (2.101)$$

It is worth noting here that there is also a large deviations theory for the Lyapunov exponents, in a similar way that we have outlined for $\delta(t)$ in section 2.3. Here there is a rate function that dictates the scaling of the joint PDF of all of the finite-time Lyapunov exponents (Tailleur and Kurchan, 2007; Laffargue et al., 2013).

To estimate the Lyapunov exponents individually, we take advantage of the growth rate of a k -volume being the sum of the first k Lyapunov exponents. By using the $(n \times k)$ matrix of separation vectors defined in (2.89) we know the change of volume for this is given by

$$|Z(t)| = |J_t(\mathbf{a})| |Z_0|, \quad (2.102)$$

where we define the initial condition for both $J_0, Z_0 = \mathbb{I}$. Through equation (2.94) we can conclude that

$$|J_t(\mathbf{a})| |Z_0| = \exp \left(\sum_{i=1}^k (\lambda_i) t \right) \quad (2.103)$$

We can use the reduced QR decomposition for $Z(t)$ to compute this

$$|J_t(\mathbf{a})| |Z_0| = |QR| = |R|, \quad (2.104)$$

as by definition $|Q| = 1$. By the QR decomposition, R is an upper triangular matrix and so its determinant is given by the product of its diagonal, $|R| = \prod_{i=1}^k r_{ii}$. Combining this with (2.104) and (2.103), gives the first k Lyapunov exponents in terms of the

components of R

$$\sum_{i=1}^k (\lambda_i) = \lim_{t \rightarrow \infty} \frac{1}{t} \sum_{I=1}^k \ln(r_{ii}). \quad (2.105)$$

There are two main issues with computing the Lyapunov exponents numerically. The first is that for long times, which we need to estimate the FTLE, the columns become linearly dependent as they tend to the stretching direction of the largest exponent λ_1 . The other problem is again a consequence of the stretching, which is exponential and so very quickly becomes singular. Conveniently we can overcome both of these issues with one solution - the Gram-Schmidt (GS) process. This is a standard way to compute FTLE (Schuster, 2005; Parker and Chua, 1989)

Practically, to compute the Lyapunov exponents numerically for the system in (2.91) we repeatedly perform QR decompositions through Gram-Schmidt orthonormalisation at discrete times along the trajectory. We will see why this helps with the singularity problem shortly, but the repeated orthogonalisation immediately solves the problem with the linearly dependent columns. We outline the basic steps for GS.

The Jacobian has the telescoping property

$$J_t(\mathbf{a}) = J_\tau(\varphi_{t_{N-1}}(\mathbf{a}))J_\tau(\varphi_{t_{N-2}}(\mathbf{a})) \cdots J_\tau(\varphi_{t_1}(\mathbf{a}))J_\tau(\mathbf{a}), \quad (2.106)$$

which allows us to naturally break the evolution into N discrete steps.

Beginning with an arbitrary orthonormal frame Z_0 (any non-normal frame will simply include a constant that will be carried through to the end), the very first decomposition looks like

$$J_\tau(\mathbf{a})Z_0 = Q_1R_1. \quad (2.107)$$

From here, the evolution is applied to the new orthonormal frame created, Q_1 , which is then decomposed again as

$$J_\tau(\varphi_{t_1}(\mathbf{a}))Q_1 = Q_2R_2, \quad (2.108)$$

and so in general we have

$$J_\tau(\varphi_{t_i}(\mathbf{a}))Q_i = Q_{i+1}R_{i+1}. \quad (2.109)$$

Combining (2.106) and (2.109) yields

$$J_t(\mathbf{a})Z_0 = [J_\tau(\varphi_{t_{N-1}}(\mathbf{a})) \cdots J_\tau(\varphi_{t_1}(\mathbf{a}))J_\tau(\mathbf{a})] Z_0, \quad (2.110)$$

$$= [J_\tau(\varphi_{t_{N-1}}(\mathbf{a})) \cdots J_\tau(\varphi_{t_1}(\mathbf{a}))] Q_1 R_1, \quad (2.111)$$

$$= [J_\tau(\varphi_{t_{N-1}}(\mathbf{a})) \cdots J_\tau(\varphi_{t_2}(\mathbf{a}))] Q_2 R_2 R_1, \quad (2.112)$$

$$= Q_N R_N \cdots R_2 R_1. \quad (2.113)$$

This telescoping behaviour is key, as it allows us to store inside R_i , the information of successive normalisations of the vector frame that we perform as by the property of determinants $\det(AB) = \det(A)\det(B)$ we have

$$|J_t(\mathbf{a})Z_0| = |Q_N R_N \cdots R_2 R_1| \quad (2.114)$$

$$|J_t(\mathbf{a})| = |R_N| \cdots |R_2| \cdot |R_1|, \quad (2.115)$$

where $|Z_0| = |Q_N| = 1$. This relates to the Lyapunov exponents as

$$|J_t(\mathbf{a})| = e^{\sum_{k=1}^n (\lambda_k)t} = \prod_{i=1}^N |R_i|, \quad (2.116)$$

$$\sum_{k=1}^n (\lambda_k)t = \sum_{i=1}^N \ln |R_i|. \quad (2.117)$$

Note that as R_i are upper triangular, their determinant is simply the product of the diagonal entries. In storing the normalisations at multiple times in the evolution like this the distances are no longer singular, and we can recover the total distance by recombining the normalisations at the end.

With these tools, we can estimate the change in the volume of a set of k vectors in n dimensions as they evolve under φ_t by simply keeping track of the diagonal elements of R , which are the magnitudes of the orthogonalised frame immediately before it is normalised. This in turn gives us a direct way to compute the Lyapunov exponents through the relation (2.116).

2.5.2 Dissipative dynamical systems

The clustering of particles that we see in this thesis is a consequence of divergence present in the flow. This divergence means the system described by 2.13 is dissipative and so volumes in the phase space contract over time. This dissipative nature of the

system implies an eventual limiting behaviour of trajectories to an attractor, or limit set which is of zero volume. Particle trajectories converge in the limit $t \rightarrow \infty$ to this attractor in the phase space for the dynamics which itself is generally some form of fractal set that evolves over time (Sommerer and Ott, 1993; Toschi and Bodenschatz, 2009). In turn this creates areas within the (2D) physical space that particles are more likely to reside on. This space is simply a projection of the underlying fractal attractor, that itself is fractal with a dimension less than 2.

A natural measure to consider when talking about fractal sets is their dimension. The fractal dimension characterises the complexity of the set as the scale it is examined at changes. In general there is no single definition for a fractal dimension and multiple ‘types’ of fractal dimension exist. Generally, these definitions of fractal dimension use techniques to approximate the complexity (or detail) of the set as a function of the scale (Mandelbrot, 1977). We examine in detail two of these definitions for this work.

We will examine the information dimension D_1 , and the correlation dimension D_2 , which are the second and third of a family of measures dubbed D_q commonly used to quantify the dimension of fractal attractors (Grassberger and Procaccia, 1983). The general form for D_q is given by

$$D_q = \frac{1}{1-q} \lim_{\epsilon \rightarrow 0} \frac{\ln I(q, \epsilon)}{\ln(1/\epsilon)}, \quad (2.118)$$

$$I(q, \epsilon) = \sum_{i=0}^{N_\epsilon} \mu_i^q \quad (2.119)$$

$$\mu_i = \lim_{T \rightarrow \infty} \frac{\eta(C_i, \mathbf{x}_0, T)}{T} \quad (2.120)$$

where $I(q, \epsilon)$ is the sum of the hypercubes of length ϵ needed to cover the attractor in phase space, and μ_i is a typical measure for a hypercube C_i , where η is the amount of time a trajectory initialised from \mathbf{x}_0 spends in a hypercube C_i over a time interval $t \in [0, T]$. This general form is abstract and not straightforward to implement, so instead we focus on a subset of the measures D_1 and D_2 that are commonly used. The first measure in the sequence, D_0 , is the box counting dimension which describes the number of hypercubes required to cover the attractor for a given cube size. We do not consider D_0 here due to the fact that for strange attractors like the ones examined here, it is often more informative to consider how important a given cube is, or how often it is visited rather than that it is simply visited at all as the D_0 describes. D_1 describes the rate that information about the state of a dynamical system changes

with scale, or equivalently the growth rate of the Shannon entropy with progressively finer discretisations. D_2 measures the number of particles in a hypersphere of radius ϵ . This can be viewed alternatively as the probability of finding two particles within a given distance of each other in phase-space, by centring the hypersphere on a given particle and counting the number of particles within it. In this manner D_2 can be seen as a natural measure of clustering to consider also. A key property of the family of measures D_q , is the bounding they induce on one another where in general (Ott, 2002)

$$D_{q_1} \leq D_{q_2} \text{ if } q_1 > q_2, \quad (2.121)$$

and so in our case D_2 provides a lower bound for D_1 . Another commonly used measure of fractal dimension is the Hausdorff dimension. This is similar to the box counting dimension but rather than counting hypercubes, it counts the number of open balls at a given radius needed to completely cover the attractor or fractal set. The Hausdorff dimension D_H bounds the box counting dimension D_0 from below ($D_H \leq D_0$). In general the Hausdorff dimension and D_0 are identical and only in fringe cases concerning countable sets does the box counting become inconsistent and the inequality becomes strict (Ott, 2002). It is due to this that the Hausdorff measure is more popular than D_0 as a descriptor for a ‘covering style’ metric. An example of its use to estimate the fractal dimension of an attractor for inertial particles can be found in (Bec, 2003), where the set up is similar to the one explored later in section 4.1.2. In general we have the following relations which gives some context to the metrics used in this section (D_1 , D_2) relative to the popular metric D_H ,

$$D_0 \geq D_1 \geq D_2, \quad (2.122a)$$

$$D_0 \geq D_H \quad (2.122b)$$

Information dimension D_1

The information dimension D_1 can be re-expressed (Ott, 2002) by defining the limit $D_1 = \lim_{q \rightarrow 1} D_q$ and using L’Hopitals rule. This alternate form is commonly used to express D_1 and is given by

$$D_1 = \lim_{\epsilon \rightarrow 0} \frac{\sum_{i=1}^{N_\epsilon} \mu_i \ln \mu_i}{\ln \epsilon}, \quad (2.123)$$

where N_ϵ is the number of hypercubes of size ϵ needed to cover the space, and μ_i is a measure associated with a hypercube C_i , which can be thought of as the probability of a trajectory being in a particular hypercube. This form for D_1 is difficult to calculate and so to calculate D_1 numerically, we take advantage of the Kaplan-Yorke conjecture. [Kaplan and Yorke \(1979\)](#) introduces a quantity known as the Lyapunov dimension, which given Lyapunov exponents $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n$ reads

$$D_L = K + \frac{1}{|\lambda_{K+1}|} \sum_{j=1}^K \lambda_j, \quad (2.124)$$

where K is the largest index such that $\sum_{j=1}^K \lambda_j \geq 0$. The conjecture states that the information dimension and the Lyapunov dimension are equal for most attractors.

In the case we examine in chapter 5, the phase space is simply the velocity space and so $\lambda_i = \lambda_{1,2}$. In this, the fractal dimension D_1 is given through the Kaplan-Yorke conjecture in 2 dimensions as

$$D_1 = 1 + \frac{\lambda_1}{|\lambda_2|}. \quad (2.125)$$

For compressible flows in two dimensions, like the ones we examine in this thesis, we have that

$$D_1 < 2, \quad (2.126)$$

and the fractal dimension measured by D_1 , and in turn D_2 by equation (2.122a), is always less than the dimension of the full physical space $D = 2$.

Correlation dimension D_2

The correlation dimension D_2 is given by ([Grassberger, 1983](#))

$$D_2 = \lim_{r \rightarrow 0} \frac{\ln C(r)}{\ln r}, \quad (2.127)$$

where $C(r)$ is the correlation integral for a distance r given by

$$C(r) = \lim_{N_p \rightarrow \infty} \frac{1}{N_p^2} \sum_{i=1, j>i}^{N_p} \Theta(r - |\mathbf{x}_i - \mathbf{x}_j|), \quad (2.128)$$

where Θ is the Heaviside step function, N_p is the total number of particles, and $|\mathbf{x}_i - \mathbf{x}_j|$ represents the pairwise distance between particles i, j which takes into account the periodic boundary of the domain. In the case of a finite number of particles, this is approximated as the correlation sum

$$C(r) = \frac{1}{N_p(N_p - 1)} \sum_{i=1, j>i}^{N_p} \Theta(r - |\mathbf{x}_i - \mathbf{x}_j|). \quad (2.129)$$

From (2.127), we can see that the correlation sum scales as $C(r) \sim r^{D_2}$.

These measures of the fractal dimension will be used in chapter 5 to estimate the dimension of the set that particles cluster on. Details on their implementation is found in section 5.4.

2.6 Summary

In this section we have outlined the background theory for the numerical advection of particles in weakly divergent flows, and a selection of measures such as the log-density $\delta(t)$, or the fractal dimension, $D_{\{1,2\}}$ to measure the clustering of these particles in the flow.

In summary, the primary measure of clustering we use is the log-density of particles in the Lagrangian frame, $\delta(t)$, which is given by equation (2.12). For the flows we are interested in for this thesis, the moments of the PDF of $\delta(t)$ have an exponential form given by large deviation theory in equation (2.27) as $\mathbb{E}[e^{\alpha\delta(t)}] \asymp \exp(tf(\alpha))$, where $f(\alpha)$ is called the free energy function. The free energy function has a key property we use to check our estimates for $f(\alpha)$, where zeros are expected at both $f(-1) = f(0) = 0$. The free energy also has the property that the slope at $\alpha = 0$ is the mean compression experienced along trajectories, i.e. $f'(0) = \langle \delta(\mathbf{a}, t) \rangle_{\mathbf{a}}$. In order to estimate the scaling of the moments in time more accurately, we calculate the free energy function using an importance sampling algorithm outlined in section 2.4.

We have described the concepts from dynamical systems theory also to further quantify the clustering of particles. We use this to estimate the dimension of the fractal attractor for the particle trajectories in phase space through the quantities D_1 and D_2 . D_1 is estimated in this work using the Kaplan-Yorke conjecture as shown in equation (2.125), and D_2 is estimated using through the its scaling relation to the correlation sum: $C(r) \sim r^{D_2}$, which is given in equation (2.129).

Chapter 3

Analytical results

3.1 Homogenisation for weakly compressible flows

Our main interest is in characterising clustering in complex, turbulent flows. Although these flows are best modelled as random to reflect the chaotic nature of the dynamics, we start by considering simpler flows. We begin with flows that are time independent and spatially periodic and then move to periodic flows in both space and time.

To represent the weakly compressible nature of the flows that we are investigating, we assume the velocity field has the form

$$\mathbf{u} = \mathbf{u}_0 + \epsilon \mathbf{u}_1 \quad \text{with} \quad \nabla \cdot \mathbf{u}_0 = 0, \quad \nabla \cdot \mathbf{u}_1 \neq 0, \quad (3.1)$$

where $\epsilon \ll 1$ is some small parameter. Recall from (2.8) that the density ρ satisfies $D_t \rho(\mathbf{x}, t) = -\rho(\mathbf{x}, t) (\nabla \cdot \mathbf{u})$ in the Lagrangian frame, so we can combine (3.1) with this to form the system of ODEs

$$\dot{\mathbf{x}}(t) = \mathbf{u}_0(\mathbf{x}, t) + \epsilon \mathbf{u}_1(\mathbf{x}, t), \quad (3.2a)$$

$$\dot{\delta}(t) = -\epsilon (\nabla \cdot \mathbf{u}_1)(\mathbf{x}, t), \quad (3.2b)$$

where $\delta = \ln \rho$ is the log-density. Note that this is the explicit small parameter form of equation (2.13). This form implies that the particle position changes for $t = \mathcal{O}(1)$ but that δ evolves on a slower time scale of at least $t = \mathcal{O}(\epsilon^{-1})$.

We are interested in the statistics of the density of particles in such a flow and the natural object to examine is the pdf $p(\delta, t)$ of δ . We have seen already the dynamics for the particles are not simple and so $p(\delta, t)$ is also a complicated object, even with

the best assumptions we can make on \mathbf{u} (random in t , mixing, ergodic). For $t \rightarrow \infty$, the pdf $p(\delta, t)$ is supported on a time-dependent fractal attractor, and we struggle to say anything about it. For large, but still finite time we expect some large deviation behaviour as we have already shown. With this we expect the pdf to take the form

$$p(\delta, t) \asymp e^{-ts(\delta/t)}, \quad (3.3)$$

for a rate function $s(\delta/t)$. With multiple timescales for the dynamics like the ones in (3.2a), we want to say something about the timescale on which the pdf varies. We would assume that it is unlikely that the pdf for density varies slowly (i.e. a time scale $t = \mathcal{O}(1)$). The most basic statement we want to make is that the time scale for the pdf varies on a scale $p(\delta, t) = P(\delta, T)$ with $T = \epsilon^2 t$, rather than in $t = \mathcal{O}(1/\epsilon)$. Physically, we can see this through the non-uniformity in sampling caused by equations (3.2a). The divergence in the flow is of order $\mathcal{O}(\epsilon)$, and $\delta(t) = \epsilon(\nabla \cdot \mathbf{u}_1)$ is sampled along trajectories. These trajectories themselves experience the effect of the divergence of strength $\mathcal{O}(\epsilon)$, and so we are not uniformly sampling the domain. Instead the sampling is biased by a factor of ϵ , and hence we can expect the pdf to scale at order $T = \epsilon^2 t$, rather than in $T = \epsilon t$. To show this we begin with a basic model that is a particular case of (3.2a) where the velocity is time independent and periodic in space. We then progressively build towards more complex models using (3.2a).

3.1.1 Time independent flow and diffusion

The first model that we consider is for particles in a time-independent diffusive flow. We assume that the domain, Ω , is periodic in space. We start with a diffusive flow as it ensures that the pdf for particle position is smooth in space and is needed as a regularisation. The joint system for position \mathbf{x} and density δ is

$$\dot{\mathbf{x}} = \mathbf{u}_0(\mathbf{x}) + \epsilon \mathbf{u}_1(\mathbf{x}) + \sqrt{2\kappa} \dot{\mathbf{W}}, \quad (3.4a)$$

$$\dot{\delta} = -\epsilon(\nabla \cdot \mathbf{u}_1)(\mathbf{x}), \quad (3.4b)$$

where κ is the diffusivity and \mathbf{W} is standard Brownian motion. For clearer notation, the variable τ will be used to denote the unscaled time from here. It is important to note that the system is a coupled system, and although we are interested in the pdf $p(\delta, \tau)$, it does not make sense to examine the pdf for δ solely and not consider the

influence of \mathbf{x} . As such, we examine the joint pdf $p(\mathbf{x}, \delta, \tau)$ for the position and log-density. The form of (3.4) makes it suitable to apply the method of homogenisation (Evans, 2013; Pavliotis, 2014). The Fokker-Planck equation for the joint pdf $p(\mathbf{x}, \delta, \tau)$ is

$$p_\tau = -\mathbf{u}_0 \cdot \nabla p + \kappa \nabla^2 p - \epsilon \nabla \cdot (\mathbf{u}_1 p) + \epsilon (\nabla \cdot \mathbf{u}_1) p_\delta. \quad (3.5)$$

We take advantage of the smallness of ϵ to obtain an approximation to (3.5) using the method of multiple scales. We define the fast and slow time scales $t = \tau$ and $T = \epsilon^2 \tau$ respectively, which gives derivatives in time as

$$\frac{d}{d\tau} = \frac{d}{dt} + \epsilon^2 \frac{d}{dT}. \quad (3.6)$$

We also expand

$$p(\mathbf{x}, \delta, \tau) = p(\mathbf{x}, \delta, t, T) = p_0(\mathbf{x}, \delta, t, T) + \epsilon p_1(\mathbf{x}, \delta, t, T) + \epsilon^2 p_2(\mathbf{x}, \delta, t, T) + \mathcal{O}(\epsilon^3). \quad (3.7)$$

Inserting this expansion and the multiscale derivative into the Fokker-Planck equation (3.5), and then isolating the $\mathcal{O}(1)$ terms gives

$$\partial_t p_0 = -\mathbf{u}_0 \cdot \nabla p_0 + \kappa \nabla^2 p_0. \quad (3.8)$$

The solution we seek is to the steady state of the problem at the fast time scale, as a result of the statistical stationarity of (3.4) at this order. As such, from equation (3.8) and $\partial_t p_0 = 0$, the solution p_0 must also be independent of \mathbf{x} :

$$p_0 = p_0(\delta, T). \quad (3.9)$$

The $\mathcal{O}(\epsilon)$ terms give

$$\partial_t p_1 = -\mathbf{u}_0 \cdot \nabla p_1 + \kappa \nabla^2 p_1 - \nabla \cdot (\mathbf{u}_1 p_0) + (\nabla \cdot \mathbf{u}_1) p_{0\delta}. \quad (3.10)$$

We are looking for a steady state solution that is independent of the fast time t , and

so

$$-\mathbf{u}_0 \cdot \nabla p_1 + \kappa \nabla^2 p_1 = \nabla \cdot (\mathbf{u}_1 p_0) - (\nabla \cdot \mathbf{u}_1) p_{0\delta}, \quad (3.11)$$

$$= (\nabla \cdot \mathbf{u}_1)(p_0 - p_{0\delta}), \quad (3.12)$$

where we use $p_0 = p_0(\delta, T)$. Due to the periodic domain, all terms have zero spatial average across the domain and as a result the problem has a solution, given by

$$p_1(\mathbf{x}, \delta, T) = \phi(\mathbf{x})(p_0(\delta, T) - p_{0\delta}(\delta, T)), \quad (3.13)$$

where $\phi(\mathbf{x})$ is the solution of the elliptic cell problem

$$-\mathbf{u}_0 \cdot \nabla \phi + \kappa \nabla^2 \phi = \nabla \cdot \mathbf{u}_1. \quad (3.14)$$

Finally at $\mathcal{O}(\epsilon^2)$, we have

$$\partial_t p_2 + \partial_T p_0 = -\mathbf{u}_0 \cdot \nabla p_2 + \kappa \nabla^2 p_2 - \nabla \cdot (\mathbf{u}_1 p_1) + (\nabla \cdot \mathbf{u}_1) p_{1\delta}. \quad (3.15)$$

Assuming $\partial_t p_2 = 0$ for a steady state, and averaging (3.15) over the domain leaves only

$$\partial_T p_0 = \langle (\nabla \cdot \mathbf{u}_1) p_{1\delta} \rangle = \langle (\nabla \cdot \mathbf{u}_1) \phi \rangle (p_{0\delta} - p_{0\delta\delta}), \quad (3.16)$$

where $\langle \cdot \rangle$ denotes the average over the domain Ω . To solve this, we first multiply equation (3.14) by ϕ and integrate over the domain. Note that the term $\phi(\mathbf{u}_0 \cdot \nabla)\phi = \nabla \cdot ((\mathbf{u}_0 \phi^2)/2)$ which integrates to zero due to net zero flux over the (periodic) boundary. After this, all that remains is

$$\int_{\Omega} (\nabla \cdot \mathbf{u}_1) \phi = \int_{\Omega} \kappa \phi \nabla^2 \phi, \quad (3.17)$$

$$= \kappa \int_{\Omega} \nabla \cdot (\phi \nabla \phi) - \kappa \int_{\Omega} (\nabla \phi)^2, \quad (3.18)$$

$$= -\kappa \langle |\nabla \phi|^2 \rangle =: -D, \quad (3.19)$$

where we identify $D > 0$ as the effective diffusivity. Using this we rewrite (3.16) as

$$\partial_T p_0 = -D p_{0\delta} + D p_{0\delta\delta}. \quad (3.20)$$

This is the result of homogenisation. It reduces the Fokker-Planck equation for the joint pdf to a Fokker-Planck equation for the marginal pdf $p(\delta, t)$. This equation shows that δ approximately evolves according to the constant coefficient SDE

$$d\delta = Ddt + \sqrt{2D}dW. \quad (3.21)$$

The drift term implies that δ grows on average, consistent with section 2.3.1. Remarkably, the drift coefficient is exactly the same as the diffusivity.

We can work out the free energy function $f(\alpha)$ from (3.20). Recall that it satisfies $e^{tf(\alpha)} \asymp \mathbb{E}[\rho^\alpha] = \mathbb{E}[e^{\alpha\delta}]$. Equation (3.20) yields

$$\frac{d}{dT}\mathbb{E}[e^{\alpha\delta}] = \int e^{\alpha\delta} \partial_T p_0(\delta, T) d\delta \quad (3.22)$$

$$= \int e^{\alpha\delta} (-Dp_{0\delta} + Dp_{0\delta\delta}) d\delta, \quad (3.23)$$

which integrating by parts gives

$$\frac{d}{dT}\mathbb{E}[e^{\alpha\delta}] = \int e^{\alpha\delta} (D\alpha + D\alpha^2)p_0 d\delta = D\alpha(\alpha + 1)\mathbb{E}e^{\alpha\delta}, \quad (3.24)$$

and so we can identify the free energy as

$$f(\alpha) = \epsilon^2 D\alpha(\alpha + 1) \quad (3.25)$$

returning to the original time variable τ .

The free energy curve in this case takes a clearly parabolic shape in α , uniquely determined by the slope at the origin given by $f'(0) = \epsilon^2 D$. This form for the free energy also vanishes for $\alpha = \{0, 1\}$ as expected from the argument in section 2.3.1.

3.1.2 Time-periodic flows

The introduction of periodic time dependence does not change much from the time independent case. The governing equations are almost identical, but instead $\mathbf{u} = \mathbf{u}_0(\mathbf{x}, t) + \epsilon\mathbf{u}_1(\mathbf{x}, t)$. The pdf is expanded in a similar way again, and again $p_0 = p(\delta, T)$ is independent of the fast time t . The differences start with $p_1 = p_1(\mathbf{x}, \delta, t, T)$ which now has fast time dependence and in turn the cell problem now needs to account for

the $\partial_t p_1$ term explicitly. With this, p_1 is now given by

$$p_1(\mathbf{x}, \delta, t, T) = \phi(\mathbf{x}, t)(p_0(\delta, T) - p_{0\delta}(\delta, T)), \quad (3.26)$$

where $\phi(\mathbf{x}, t)$ now also has time dependence and is the solution of the cell problem

$$-\phi_t - \mathbf{u}_0 \cdot \nabla \phi + \kappa \nabla^2 \phi = \nabla \cdot \mathbf{u}_1. \quad (3.27)$$

From here, the only difference is the averaging for (3.15) is now in both time and space rather than solely in space. The time dependent case has the same result, with the pdf varying on a timescale $T = \epsilon^2 \tau$.

3.1.3 Removing diffusion from time dependent flow

Removing diffusion from the system complicates matters, as the cell problem in equation (3.14) is ill-defined for a diffusivity $\kappa = 0$, and so we need to hope that the flow is sufficiently chaotic to act as a diffusion. Any fully justified results would require that the flow have properties such as exponential mixing and ergodicity, which themselves are difficult to show, in order to say anything conclusive. In light of this, we take a heuristic approach to attempt to demonstrate a time dependence of the pdf $p(\delta, T)$ of order $T = \epsilon^2 \tau$. Removing the diffusion from (3.4) and adding time dependent flow $\mathbf{u} = \mathbf{u}(\mathbf{x}, t)$,

$$\dot{\mathbf{x}} = \mathbf{u}_0(\mathbf{x}, \tau) + \epsilon \mathbf{u}_1(\mathbf{x}, \tau), \quad (3.28a)$$

$$\dot{\delta} = -\epsilon(\nabla \cdot \mathbf{u}_1)(\mathbf{x}, \tau). \quad (3.28b)$$

which now resembles the system outlined in (3.2a). The Fokker-Planck equation for this is

$$p_\tau = -\mathbf{u}_0 \cdot \nabla p - \epsilon \nabla \cdot (\mathbf{u}_1 p) + \epsilon(\nabla \cdot \mathbf{u}_1)p_\delta. \quad (3.29)$$

Expanding $p = p_0 + \epsilon p_1 + \dots$ with $t = \tau$, $T = \epsilon^2 \tau$ as before, and then organising at each order gives for $\mathcal{O}(1)$

$$\partial_t p_0 + \mathbf{u}_0 \cdot \nabla p_0 = 0. \quad (3.30)$$

We again look for a steady state solution at the fast time scale. We focus on the steady state here for p_0 as physically we would expect p_0 to represent the average behaviour over the domain, and any deviation from this to occur at a higher order. As such, the solution is independent of t again, and in turn, must be independent of \mathbf{x} to give $p_0 = p(\delta, T)$. At $\mathcal{O}(\epsilon)$

$$\partial_t p_1 = -\mathbf{u}_0 \cdot \nabla p_1 - \nabla \cdot (\mathbf{u}_1 p_0) + (\nabla \cdot \mathbf{u}_1) p_{0\delta}. \quad (3.31)$$

This no longer has a ‘nice’ solution like equation (3.13) as the cell problem is ill-defined without a diffusivity from the noise term in (3.28a). To write a solution for this, it is helpful to think again in a Lagrangian frame. We introduce the volume-preserving flow map φ_t of the leading order flow \mathbf{u}_0 only. We can view the system through the variables of label space (\mathbf{a}, t) rather than the position space (\mathbf{x}, t)

$$\mathbf{X}(\mathbf{a}, t) = \varphi_t(\mathbf{a}) = \mathbf{x}, \quad \mathbf{X}(\mathbf{a}, 0) = \mathbf{a}. \quad (3.32)$$

This also has a corresponding pdf such that

$$p(\mathbf{x}, \delta, t, T) = \tilde{p}(\mathbf{a}, \delta, t, T) \quad (3.33)$$

We define the push-forward (φ_{t*}) and pull-back (φ_t^*) notation of a function f by φ_t that allow us to denote moving in and out of label space more clearly:

$$(\varphi_{t*} f)(\mathbf{x}) = f(\varphi_t^{-1}(\mathbf{x})), \quad (\varphi_t^* f)(\mathbf{x}) = f(\varphi_t(\mathbf{x})). \quad (3.34)$$

Note that, when applied to multivariable functions other than \mathbf{x} , the push-forward (or pull-back) apply only to \mathbf{x} and leave the other variables unchanged. Equation (3.31) can be written in label space as

$$\partial_t p_1 + \mathbf{u}_0 \cdot \nabla p_1 = -(\nabla \cdot \mathbf{u}_1)(\mathbf{x}, t) (p_0 - p_{0\delta}), \quad (3.35)$$

$$\partial_t \tilde{p}_1 = -(\nabla \cdot \mathbf{u}_1)(\mathbf{X}, t) (p_0 - p_{0\delta}), \quad (3.36)$$

which using the pull-back notation is equivalent to

$$\partial_t \tilde{p}_1 = -\varphi_t^* [(\nabla \cdot \mathbf{u}_1)(\mathbf{a}, t) (p_0 - p_{0\delta})]. \quad (3.37)$$

The solution to p_1 in position space is therefore

$$p_1 = -\varphi_{t*} \int_{-\infty}^t \varphi_s^* [(\nabla \cdot \mathbf{u}_1)(\mathbf{x}, s) \, ds (p_0 - p_{0\delta})]. \quad (3.38)$$

At $\mathcal{O}(\epsilon^2)$ equation (3.29) is

$$\partial_t p_2 + \mathbf{u}_0 \cdot \nabla p_2 = -\partial_T p_0 + (\nabla \cdot \mathbf{u}_1) p_{1\delta} - \nabla \cdot (\mathbf{u}_1 p_1). \quad (3.39)$$

Averaging this over the domain eliminates the LHS and divergence on the RHS. Using (3.38), this reduces to

$$\partial_T p_0 = -D p_{0\delta} + D p_{0\delta\delta}, \quad (3.40)$$

with

$$D = \langle (\nabla \cdot \mathbf{u}_1)(\mathbf{x}, t) \varphi_{t*} \int_{-\infty}^t \varphi_s^* (\nabla \cdot \mathbf{u}_1)(\mathbf{x}, s) \, ds \rangle, \quad (3.41)$$

where $\langle \cdot \rangle$ is an average over space and time, differing from equation (3.16) due to the time dependence that has been introduced. A similar set of results to (3.25) holds for this case also, highlighting the tendency for $p(\delta, T)$ to vary on a scale $T = \epsilon^2 \tau$. As a final note, the form for D here is not easy to work with, but it can be expressed in an alternate way that will prove to be more useful. As φ_t is volume preserving, we can apply the pull-back to D to give

$$D = \langle \varphi_t^* [(\nabla \cdot \mathbf{u}_1)(\mathbf{x}, t)] \int_{-\infty}^t \varphi_s^* [(\nabla \cdot \mathbf{u}_1)(\mathbf{x}, s) \, ds] \rangle. \quad (3.42)$$

$$= \langle [(\nabla \cdot \mathbf{u}_1)(\varphi_t(\mathbf{a}), t)] \int_{-\infty}^t [(\nabla \cdot \mathbf{u}_1)(\varphi_s(\mathbf{a}), s) \, ds] \rangle. \quad (3.43)$$

Averaging this along trajectories in the Lagrangian frame then gives

$$D = \lim_{T \rightarrow \infty} \frac{1}{T} \int_{-\infty}^T dt \int_{-\infty}^t \langle \varphi_t^* [(\nabla \cdot \mathbf{u}_1)(\mathbf{x}, t)] \varphi_s^* [(\nabla \cdot \mathbf{u}_1)(\mathbf{x}, s)] \rangle \, ds, \quad (3.44)$$

which is the average of the time integral of the autocorrelation of the divergence of the velocity field evaluated along a trajectory. We can calculate this form of D explicitly, and provides us with a way to estimate D from simulated trajectories in some velocity field.

The key behaviour here is due to the difference in the time scale associated with the divergent and non-divergent parts of the velocity field. The leading order flow is responsible for mixing and occurs on a relatively fast time scale ($t = \tau$), with a short correlation time. The weakly divergent part of the flow is shown to act over a slower time scale which is proportional to the square of the small parameter $T = \epsilon^2\tau$ rather than simply being linear in ϵ . With this fact, even though the exact dynamics of the system may be too complex to determine exactly, we can make basic statements about the statistics of the log-density, such as the parabolic shape of the free-energy curve shown in equation (3.25) which holds in the final case just shown also. The parameter D , which is the autocorrelation along trajectories, also corresponds to the mean compression in the system as the equation (3.40) is the Fokker-Planck equation associated with the SDE

$$d\delta = -Ddt + \sqrt{2D}dW \quad (3.45)$$

The drift of this SDE is exactly the mean compression rate we expect from the system $\mathbb{E}\delta(t) = Dt$. It is clear from both this and (3.44) that D must be finite to be realistic for any potential application. This is the same form again as the one found in the last section, and all of the points made there, hold in this case also.

3.2 Modelling weakly compressible flows as Kraichnan flows

As an alternative perspective, instead of modelling the dynamics with an explicitly multiscale flow for the velocity field, we can instead use a general stochastic flow which is white in time, but smooth in space. These are called Kraichnan flows (Kraichnan, 1968). Formulating the dynamics with a Kraichnan flow like this, allows us to draw results on the expected behaviour of flows with short correlation time. This will be relevant in particular for the clustering of inertial particles which we examine in chapter 4.

Formally we express the defining properties of the velocity field for a Kraichnan flow as

$$\mathbb{E}[dU_i(\mathbf{x}, t)] = 0 \quad (3.46a)$$

$$\mathbb{E}[dU_i(\mathbf{x}, t) dU_j(\mathbf{y}, s)] = \delta(t - s)C_{ij}(\mathbf{x}, \mathbf{y}) dt ds, \quad (3.46b)$$

where $\mathcal{C}(\mathbf{x}, \mathbf{y})$ is the flow's (spatial) correlation function.

To model the system of particles advected by a Kraichnan flow physically, we examine the system in the limit of short correlation time rather than immediately as white noise. Approximating white noise with a smoother process with short correlation time, and gradually removing the regularity through the zero-limit correlation time is a common method. In these cases, a limiting SDE is found that is a Stratonovich SDE rather than an Itô SDE. This process is usually called the Wong-Zakai theorem (Wong and Zakai (1965a), Wong and Zakai (1965b), Pavliotis (2014)) We consider the limit $\epsilon \rightarrow 0$ of the equation

$$dX_i = \epsilon^{-1} U_i(\mathbf{X}) dt, \quad \text{with} \quad U_i = A_k(t) \phi_{ik}(\mathbf{X}), \quad (3.47)$$

where

$$dA_k = -\epsilon^{-2} \gamma A_k dt + \epsilon^{-1} \sqrt{2\gamma} dW_k. \quad (3.48)$$

γ has dimension $[\gamma] = [\text{time}]^{-1}$, so that both ϵ and A_k are non-dimensional. ϕ_{ik} are basis functions for the velocity field, which for example could be trigonometric which will be the case in chapter 4. The A_k are the scalar coefficients for the basis and are defined in this way in particular to resemble the coefficients chosen for the flow in chapter 4, which is an adapted Ornstein-Uhlenbeck process (Gardiner, 2009). The A_k are defined also such that in the late time limit $t \gg 1$, that we are interested in studying the clustering behaviour at, they are independent zero mean Gaussian with

$$\mathbb{E}[A_k(t) A_l(t)] = \delta_{kl}. \quad (3.49)$$

We also keep the correlation function for ϕ_{ik} general, following the notation in equation (3.46b) as

$$\mathbb{E}[\phi_{ik}(\mathbf{x}) \phi_{jk}(\mathbf{y})] = \mathcal{C}_{ij}(\mathbf{x}, \mathbf{y}). \quad (3.50)$$

The Lagrangian log-density satisfies the equation

$$d\Delta = -\epsilon^{-1} \partial_i U_i(\mathbf{x}, t) dt. \quad (3.51)$$

We derive a reduced model for the coupled system made up of equations (3.47),(3.48), (3.51) in the limit of $\epsilon \rightarrow 0$. To begin, the generator associated with the stochastic process (3.48) is

$$\epsilon^{-2}\mathcal{L} = \gamma(-a_k\partial_{a_k} + \partial_{a_k}a_k). \quad (3.52)$$

With this, the backward Kolmogorov equation for the joint process consisting of (3.47),(3.48), (3.51) is

$$\partial_t u = (\epsilon^{-1}\phi_{ik}(\mathbf{x})a_k\partial_{x_i} + \epsilon^{-2}\mathcal{L} - \epsilon^{-1}\partial_i\phi_{ik}(\mathbf{x})a_k\partial_\delta) u. \quad (3.53)$$

We expand the probability distribution as

$$u = u_0(\mathbf{x}, \delta, t) + \epsilon u_1(\mathbf{x}, \delta, a, t) + \mathcal{O}(\epsilon^2). \quad (3.54)$$

We don't examine the slow time scale explicitly here, as we expect in the limit $\epsilon \rightarrow 0$ this is exactly white noise. The leading order term is independent of $a = \{a_k\}$ to ensure that the $\mathcal{O}(\epsilon^{-2})$ equation $\mathcal{L}u_0 = 0$ is satisfied. At $\mathcal{O}(\epsilon^{-1})$ the backward Kolmogorov equation reads

$$\mathcal{L}u_1 = -\phi_{ik}(\mathbf{x})a_k\partial_{x_i}u_0 + \partial_i\phi_{ik}(\mathbf{x})a_k\partial_\delta u_0. \quad (3.55)$$

As the RHS is linear in a_k , this has straightforward solution

$$u_1 = \gamma^{-1}a_k(\phi_{ik}(\mathbf{x})\partial_{x_i}u_0 - \partial_i\phi_{ik}(\mathbf{x})\partial_\delta u_0). \quad (3.56)$$

At $\mathcal{O}(1)$, we have

$$\partial_t u_0 = \mathcal{L}u_2 + \phi_{ik}(\mathbf{x})a_k\partial_{x_i}u_1 - \partial_i\phi_{ik}(\mathbf{x})a_k\partial_\delta u_1. \quad (3.57)$$

This can be solved by requiring that $\mathcal{L}u_2$ is orthogonal to the null space of the adjoint of \mathcal{L} or equivalently $\int(\mathcal{L}u_2)p(a) da = 0$, where $p(a)$ is the solution of $\mathcal{L}^*p = 0$, where \mathcal{L}^* denotes the adjoint. First substituting (3.56), we multiply (3.57) by $p(a)$ and then integrating over $a = \{a_k\}$. The LHS is straightforward and is simply

$$\int \partial_t u_0 p(a) da = \partial_t u_0. \quad (3.58)$$

For the RHS we use the fact that the solution $p(a_k)$ to $\mathcal{L}^*p = 0$ is the stationary pdf for the A_k . As a consequence,

$$\int (a_k a_l) p(a) da = \mathbb{E}[A_k A_l] = \delta_{k,l}, \quad (3.59)$$

and thus altogether we obtain

$$\begin{aligned} \partial_t u_0 = \gamma^{-1} & \left[\phi_{ik}(\mathbf{x}) \partial_{x_i} (\phi_{jk}(\mathbf{x}) \partial_{x_j} u_0) - \phi_{ik}(\mathbf{x}) \partial_{x_i} (\partial_j \phi_{jk}(\mathbf{x}) \partial_\delta u_0) \right. \\ & \left. - \partial_i \phi_{ik}(\mathbf{x}) \partial_\delta (\phi_{jk}(\mathbf{x}) \partial_{x_j} u_0) + \partial_i \phi_{ik}(\mathbf{x}) \partial_j \phi_{jk}(\mathbf{x}) \partial_{\delta\delta} u_0 \right]. \end{aligned} \quad (3.60)$$

This is the backward Kolmogorov equation associated with the SDEs

$$\dot{\mathbf{X}}_i = (2\gamma)^{-1/2} \phi_{ik}(\mathbf{X}) \circ dW_k, \quad (3.61a)$$

$$\dot{\Delta} = -(2\gamma)^{-1/2} \partial_i \phi_{ik}(\mathbf{X}) \circ dW_k, \quad (3.61b)$$

where $\circ dW_k$ is Stratonovich noise, which is the natural choice when looking at the limit of a short correlation time process. This is the Wong-Zakai theorem mentioned at the start of this section. These are the equations for the joint position log-density system as the limit of a short correlation flow.

The Fokker-Planck equation for the full system (3.61) corresponding to the BKE (3.60) is

$$\begin{aligned} \gamma \partial_t p &= \partial_{x_i} (\phi_{jk}(\mathbf{x}) \partial_{x_i} (\phi_{ik}(\mathbf{x}) \partial_{x_j} p) - \partial_i \phi_{ik}(\mathbf{x}) \partial_{x_j} (\phi_{jk}(\mathbf{x}) \partial_\delta p) \\ &\quad - \partial_{x_i} (\partial_j \phi_{jk}(\mathbf{x}) \phi_{ik}(\mathbf{x}) \partial_\delta p) + \partial_i \phi_{ik}(\mathbf{x}) \partial_j \phi_{jk}(\mathbf{x}) \partial_{\delta\delta} p \\ &= \partial_{x_i} (\phi_{jk}(\mathbf{x}) \partial_{x_i} (\phi_{ik}(\mathbf{x}) \partial_{x_j} p)) - \partial_{x_j} (\partial_i \phi_{ik}(\mathbf{x}) (\phi_{jk}(\mathbf{x}) \partial_\delta p)) + \partial_{x_j} (\partial_i \phi_{ik}(\mathbf{x})) \phi_{jk}(\mathbf{x}) \partial_\delta p \\ &\quad - \partial_{x_i} (\partial_j \phi_{jk}(\mathbf{x}) \phi_{ik}(\mathbf{x}) \partial_\delta p) + \partial_i \phi_{ik}(\mathbf{x}) \partial_j \phi_{jk}(\mathbf{x}) \partial_{\delta\delta} p. \end{aligned} \quad (3.62)$$

We are primarily interested in the pdf for the log-density Δ , the marginal pdf

$$p_\Delta(\delta, t) = \int p(\mathbf{x}, \delta, t) d\mathbf{x}. \quad (3.63)$$

Integrating (3.62) with respect to \mathbf{x} , and noting that p decays quickly as $|\mathbf{x}| \rightarrow \infty$ so

terms containing a derivative in \mathbf{x} are eliminated, we are left with

$$\gamma \partial_t p_\Delta = \partial_{x_j} (\partial_i \phi_{ik}(\mathbf{x})) \phi_{jk}(\mathbf{x}) \partial_\delta p_\Delta + \partial_i \phi_{ik}(\mathbf{x}) \partial_j \phi_{jk}(\mathbf{x}) \partial_{\delta\delta} p_\Delta. \quad (3.64)$$

For homogeneous flows, the correlation function can be written as

$$\mathcal{C}_{ij}(\mathbf{x}, \mathbf{x}') = \mathcal{C}_{ij}(\mathbf{x} - \mathbf{x}') = \phi_{ik}(\mathbf{x}) \phi_{jk}(\mathbf{x}'), \quad (3.65)$$

and noting the anti-symmetry $\partial_x \mathcal{C}_{ij}(\mathbf{x} - \mathbf{x}') = -\partial_{x'} \mathcal{C}_{ij}(\mathbf{x} - \mathbf{x}')$, we identify from (3.64)

$$D := -\partial_{ij} \mathcal{C}_{ij}(0) > 0, \quad (3.66)$$

hence

$$\partial_t p_\Delta = -D\gamma^{-1} \partial_\delta p_\Delta + D\gamma^{-1} \partial_{\delta\delta} p_\Delta. \quad (3.67)$$

This is the Fokker-Planck equation for a simple SDE,

$$d\Delta = D\gamma^{-1} dt + \sqrt{2D\gamma^{-1}} dW. \quad (3.68)$$

The drift of this SDE corresponds with the mean compression rate we expect from the system, $\mathbb{E}\Delta = Dt/\gamma$.

The system must be examined as the joint position-density system in this way, as examining the density on its own leads to the incorrect characterisation. For example, the density of particles satisfies the equation

$$\dot{\Delta} = -\partial_i U_i(\mathbf{x}, t). \quad (3.69)$$

Naively, examining the pdf for the particle density $p(\delta, t)$ using (3.47) we would arrive at the Fokker-Planck equation

$$\partial_t p = -\partial_i \phi_{ik} \partial_j \phi_{jk} \partial_{\delta\delta} p \quad (3.70)$$

$$= D \partial_{\delta\delta} p, \quad (3.71)$$

where $D = -\partial_{ij} \mathcal{C}_{ij}(0)$. This is clearly not the complete picture. If this were true, we would observe a simple diffusion of probability for the density of particles in the

domain. In this picture, there is no evidence of preferential concentration, or a mean rate of compression in the system, which we know should be present. This result ignores the fact that ϕ_{ik} is a function of position \mathbf{X} also and so it is crucial the Fokker-Planck equation be applied to the joint position-density system rather than only equation (3.69).

Similarly to the previous section, we can identify the free energy function from (3.68) also

$$\frac{d}{dt}\mathbb{E}e^{\alpha\Delta} = \frac{d}{dt} \int e^{\alpha\delta} p_{\Delta}(\delta, t) d\delta = \int e^{\alpha\delta} (-D\gamma^{-1}\partial_{\delta}p_{\Delta} + D\gamma^{-1}\partial_{\delta\delta}p_{\Delta}) d\delta \quad (3.72)$$

$$(3.73)$$

Integrating by parts,

$$\frac{d}{dt}\mathbb{E}e^{\alpha\Delta} = \int e^{\alpha\delta}\gamma^{-1}D(\alpha + \alpha^2)p_{\Delta} d\delta = \gamma^{-1}D(\alpha + \alpha^2)\mathbb{E}e^{\alpha\Delta}, \quad (3.74)$$

hence $\mathbb{E}e^{\alpha\Delta} \propto e^{tf(\alpha)}$ with the free energy

$$f(\alpha) = \gamma^{-1}D\alpha(\alpha + 1). \quad (3.75)$$

As expected from section 2.3.1, $f(-1) = f(0) = 0$. We also again show the free energy curve in this case takes a clearly parabolic shape in α , uniquely determined by the slope at the origin given by $\gamma^{-1}D$.

This coefficient D is the same D that we saw in the previous section. This implies that, up to scaling in γ or ϵ^2 , the underlying behaviour of the statistics of $\delta(t)$ is dependent on this single parameter D given by the autocorrelation of the divergence along trajectories.

3.3 Summary

The results outlined in this chapter will be useful in the subsequent sections. In both cases the divergence is controlled by a small parameter (either Stokes (St) or Rossby (Ro) number for inertial or rotational effects respectively) and so they directly correspond to multiscale flow with a weak divergence term as seen in the first half of this chapter. The results for Kraichnan flows as a limit of short correlation time flow in the second portion of this chapter relate to the work on inertial particles found in chapter

4, where the flow in this chapter used to advect particles is also driven by a short correlation time stochastic process.

We have shown in both the white noise limit in this section and the weakly divergent flow in section 3.1.3 that the free energy will take a parabolic form, which scales in either the inverse of the correlation time of the flow in the white noise case, or the square of the parameter that controls the weak divergence in section 3.1.3. In both cases the free energy is uniquely determined by the slope at the origin. given by the time integrated autocorrelation of the divergence D along trajectories. This behaviour will be investigated for both flows examined in this thesis.

It is worth noting here also that these results are all in terms of the free energy, but as a consequence of the Legendre transform, the parabolic behaviour and scaling also applies to the rate function. In general, the equivalence of information contained in the entire collection of moments to the full PDF, implies that a result for one large deviation statistic, also applies to the other.

In the next two chapters, we first examine the case of inertial particles in a random flow, where inertia is the source of weak divergence in the system. We then examine the dynamics of passive particles in a weakly compressible flow, based on the commonly used SQG model which is widely used to model ocean surface flow.

Chapter 4

Inertial particles

Often particles in flows are modelled using passive particles. These passive particles are infinitely small, massless particles whose motion is simply given by the motion of the flow in which they are advected. As they are passive, they have no reciprocal effect on the flow itself. When dealing with particles with a finite mass and size they cannot simply be modelled as passive particles. Instead we must take into account the effect of inertia on these particles.

Introducing the effects of inertia increases the complexity of the dynamics. With this change, the particles no longer follow the motion of the flow they are within like passive tracers, but instead have their own velocity that differs from that of the surrounding flow. This added complexity means that, in order to model such particles, we must also keep track of the velocity \boldsymbol{v} as well as the position \boldsymbol{x} of each individual particle. With this addition, the phase space for the dynamics is no longer simply the position space, but a 4-dimensional phase space we call the position-velocity phase space.

In this chapter we aim to describe the clustering of inertial particles in physical space using the statistics of the log-density $\delta(t)$, and draw comparisons to the theoretical work outlined in chapter 3. In particular we aim to investigate the relationship between $\delta(t)$ and both the small parameter that controls inertial effects, the Stokes number St , and the correlation time of the flow, γ . To do this we make use of an approximation of the full phase space dynamics, taking advantage of an invariant slow manifold to describe the motion. This reduces the dynamics to a 2D physical phase space, by treating the inertial particles in an incompressible flow as passive particles in a compressible flow. Details of the approximation are outlined in section 4.1.2. With this approximation, and the importance sampling techniques outlined in chapter 2 we simulate inertial

particles in a random flow and compare the behaviour of $P(\delta, t)$ with the results in chapter 3.

There is existing work to characterise clustering of inertial particles in the full position-velocity phase space. A recent review (Bec et al., 2024) discusses work on the statistics of the full phase space models. In particular Bec (2003) examines the existence of a fractal attractor in the full phase space, and its projection to the position space, finding the fractal dimension of this projected attractor to scale as St^2 , and in the limit of $\text{St} \rightarrow 0$, to equal the dimension of the position space. This is closely related to the discussion in section 2.5.2. Investigation of the fractal dimensions $D_{\{1,2\}}$ outlined in section 2.5.2 for the approximated inertial particles is not undertaken in this thesis, but the scaling found in Bec (2003) states that we would expect them also to scale quadratically in St , as long as the approximation is reasonable (i.e. $\text{St} \ll 1$).

Previous work also exists for passive particles and turbulent transport in random compressible flows. For example, the review Falkovich et al. (2001) discusses the statistics of clustering from a Lagrangian perspective, and Balkovsky et al. (2001) discuss the concentration of inertial particles modelled as passive particles in compressible flow, and their independence of the statistics of the velocity field.

The work in this chapter aims to complement the previous work in the area, with a focus on exploring the large deviation statistics of the log-density, $\delta(t)$, relative to the Stokes number, St , and correlation time of the carrier flow γ .

4.1 Motion of inertial particles

4.1.1 Maxey-Riley equations of motion

The velocity of the particles is commonly modelled using the Maxey-Riley equations Maxey and Riley (1983); Renaud and Vanneste (2020). For a flow $\mathbf{u}(\mathbf{x}, t)$, the position $\mathbf{x}(t)$ and velocity $\mathbf{v}(t)$ of a spherical particle satisfy

$$\dot{\mathbf{x}} = \mathbf{v} \tag{4.1a}$$

$$\begin{aligned} \rho_p \dot{\mathbf{v}} = & \rho_f \frac{D\mathbf{u}}{Dt} + (\rho_p - \rho_f)\mathbf{g} - \frac{9\nu\rho_f}{2a^2} \left(\mathbf{v} - \mathbf{u} - \frac{a^2}{6}\Delta\mathbf{u} \right) - \frac{\rho_f}{2} \left(\dot{\mathbf{v}} - \frac{D}{Dt} \left(\mathbf{u} + \frac{a^2}{10}\Delta\mathbf{u} \right) \right) \\ & - \frac{9\rho_f}{2a} \sqrt{\frac{\nu}{\pi}} \int_0^t \frac{1}{\sqrt{t-s}} \left(\dot{\mathbf{v}}(s) - \frac{d}{ds} \left(\mathbf{u} + \frac{a^2}{6}\Delta\mathbf{u} \right)_{\mathbf{x}=\mathbf{x}(s)} \right) ds. \end{aligned} \tag{4.1b}$$

In the above, $\frac{D\phi}{Dt} = \frac{\partial\phi}{\partial t} + \mathbf{u} \cdot \nabla\phi$ represents the material derivative, ρ_f and ρ_p are the density of the fluid and particle respectively, \mathbf{g} is the gravitational force, ν is the kinematic viscosity, and a is the radius of the particle.

The terms in order on the right hand side represent: the force on the particle from the flow, the buoyancy of the particle, the Stokes drag on the particle, the added mass effect, which is the result of part of the carrier fluid moving with the particle, and the Basset–Boussinesq memory term, which describes the development of a boundary layer around the moving particle, and relies on the history of the particles' trajectory. The terms with a factor $a^2\Delta\mathbf{u}$ are commonly referred to as Faxén corrections, which are corrections to Stokes' Law and the added mass effect that come into effect for larger particles. This accounts for the additional mass entrained by larger particles as they move through the flow.

Re-organising and non-dimensionalising equation (4.1b) with a characteristic length scale L (the domain size or flow scale), velocity U and time scale based on these $T = L/U$, yields

$$\dot{\mathbf{x}} = \mathbf{v} \tag{4.2a}$$

$$\begin{aligned} \dot{\mathbf{v}} = & \frac{3\rho_f}{2\rho_p + \rho_f} \frac{D\mathbf{u}}{Dt} + \frac{2(\rho_p - \rho_f)}{2\rho_p + \rho_f} \mathbf{g} \\ & - \frac{9\nu\rho_f}{a^2(2\rho_p + \rho_f)} \left(\mathbf{v} - \mathbf{u} - \frac{a^2}{6} \Delta\mathbf{u} \right) + \frac{\rho_f}{2\rho_p + \rho_f} \left(\frac{D}{Dt} \left(\frac{a^2}{10} \Delta\mathbf{u} \right) \right) \\ & - \frac{9\rho_f}{2a(2\rho_p + \rho_f)} \sqrt{\frac{\nu}{\pi}} \int_0^t \frac{1}{\sqrt{t-s}} \left(\dot{\mathbf{v}}(s) - \frac{d}{ds} \left(\mathbf{u} + \frac{a^2}{6} \Delta\mathbf{u} \right)_{\mathbf{x}=\mathbf{x}(s)} \right) ds. \end{aligned} \tag{4.2b}$$

From this we identify the Stokes number, St , and a modified density ratio, β , that occurs in multiple terms;

$$St = \left(\frac{a}{L} \right)^2 \frac{UL}{3\beta\nu}, \tag{4.3}$$

$$\beta = \frac{3\rho_f}{2\rho_p + \rho_f}. \tag{4.4}$$

The Stokes number can also be expressed as $St = \left(\frac{a^2}{3\beta L^2} \right) Re$, where $Re = UL/\nu$ is the Reynolds number associated with the flow. The parameter β measures the influence of the added mass effect in the Maxey–Riley equation (4.1b). Note that $\beta \in (0, 2)$, where $\beta = 0$ corresponds to the limit of infinitely heavy particles and $\beta = 2$ to infinitely light particles. The Stokes number St provides a measure of the effect of inertia in

the dynamics. It describes the ratio of the time scale of a particle to react to the flow around it, relative to the time scale of the flow itself. It is commonly written as

$$\text{St} = \frac{\tau_p}{\tau_f} = \tau_p \frac{U}{L}. \quad (4.5)$$

With these quantities, equation (4.2) becomes

$$\dot{\mathbf{x}} = \mathbf{v} \quad (4.6a)$$

$$\begin{aligned} \dot{\mathbf{v}} = & \beta \frac{D\mathbf{u}}{Dt} + \frac{2\beta}{3\rho_f}(\rho_p - \rho_f)\mathbf{g} - \frac{1}{\text{St}} \left(\mathbf{v} - \mathbf{u} - \frac{a^2}{6}\Delta\mathbf{u} \right) + \frac{\beta}{3} \left(\frac{D}{Dt} \left(\frac{a^2}{10}\Delta\mathbf{u} \right) \right) \\ & - \frac{1}{\sqrt{\text{St}}} \sqrt{\frac{3\beta}{\pi}} \int_0^t \frac{1}{\sqrt{t-s}} \left(\dot{\mathbf{v}}(s) - \frac{d}{ds} \left(\mathbf{u} + \frac{a^2}{6}\Delta\mathbf{u} \right)_{\mathbf{x}=\mathbf{x}(s)} \right) ds. \end{aligned} \quad (4.6b)$$

For the idealised case considered here, we make some assumptions that allow much of the equation to be reduced for simplicity. We assume the particles are small relative to the scale of the flow L (i.e. $a \ll L$). With this we can assume the Faxén corrections, which are of order a^2/L^2 , are negligible. The Basset–Boussinesq history term is non-trivial to deal with, and so it is often neglected (Haller and Sapsis, 2008; Michaelides, 1997) by noting the ratio between this term and the Stokes drag is $\sqrt{\text{St}}$. We are interested only in $\text{St} \ll 1$ and so this term is neglected. Similarly, the ratio between the gravitational term and the Stokes drag is proportional to $\beta\text{St} \ll 1$ and so we neglect gravitational forces ($\mathbf{g} = 0$). This leaves us with the simplified equations of motion

$$\dot{\mathbf{v}} = -\frac{1}{\text{St}} (\mathbf{v} - \mathbf{u}) + \beta \frac{D\mathbf{u}}{Dt}, \quad (4.7)$$

To utilise the Maxey-Riley equation, or the approximation to it in equation (4.7), we formalise the joint system for inertial particle motion in a flow $\mathbf{u}(\mathbf{x}, t)$ as

$$\dot{\mathbf{x}} = \mathbf{v}, \quad (4.8a)$$

$$\dot{\mathbf{v}} = -\frac{1}{\text{St}} (\mathbf{v} - \mathbf{u}) + \beta \frac{D\mathbf{u}}{Dt}. \quad (4.8b)$$

It is important to emphasise that as a key difference to passive tracers, the dynamics for inertial particles takes place on a $2N$ -dimensional phase space, with N the number of spatial dimensions.

We note that the dynamics are dissipative even if the carrier flow \mathbf{u} is divergent-free.

Indeed, the phase-space divergence is

$$\partial_{\mathbf{x}}\dot{\mathbf{x}} + \partial_{\mathbf{v}}\dot{\mathbf{v}} = -1/\text{St} < 0. \quad (4.9)$$

As the dynamics are uniformly contracting with time, we expect phase-space trajectories to converge to an attracting manifold. In the next section we examine the form of this manifold and how it relates the dynamics of passive tracers to that of inertial particles in the same carrier flow.

4.1.2 Approximation of dynamics by motion on an invariant manifold

We aim to further simplify the Maxey-Riley equations in the case of small Stokes number $\text{St} \ll 1$. With this in mind we can examine the simplified equation of motion using perturbation theory and treat $\text{St} = \epsilon$ as a small parameter. This is a common approximation and other examples can be found in (Ferry and Balachandar, 2001; Renaud and Vanneste, 2020)

The following has been shown in Haller and Sapsis (2008); Rubin et al. (1995), and is included here for completeness. We write (4.8) as the singularly perturbed system

$$\dot{\mathbf{x}} = \mathbf{v}, \quad (4.10a)$$

$$\epsilon\dot{\mathbf{v}} = \mathbf{u}(\mathbf{x}, t) - \mathbf{v} + \epsilon\beta \frac{D\mathbf{u}}{Dt}. \quad (4.10b)$$

We introduce the standard fast time scale $\tau = t/\epsilon$. Writing as an autonomous system in $(\mathbf{x}, \mathbf{v}, t)$

$$\mathbf{x}' = \epsilon\mathbf{v}, \quad (4.11a)$$

$$\mathbf{v}' = \mathbf{u}(\mathbf{x}, t) - \mathbf{v} + \epsilon\beta \frac{D\mathbf{u}(\mathbf{x}, t)}{Dt}, \quad (4.11b)$$

$$t' = \epsilon \quad (4.11c)$$

where $\frac{d\mathbf{x}}{d\tau} = \mathbf{x}'$. Note that by simply setting $\epsilon = 0$ in equations (4.10) gives $\mathbf{v} = \mathbf{u}$, i.e. non-inertial particles. A better approximation is presented in Haller and Sapsis (2008) where it is shown that this system admits an invariant slow manifold given by:

$$M_\epsilon = \left\{ (\mathbf{x}, \mathbf{v}, t) : \mathbf{v} = \mathbf{v}_0(\mathbf{x}, t) + \epsilon\mathbf{v}_1(\mathbf{x}, t) + \dots + \epsilon^n\mathbf{v}_n(\mathbf{x}, t) + \mathcal{O}(\epsilon^{n+1}) \right\}. \quad (4.12)$$

As M_ϵ is invariant we differentiate with respect to the fast time τ to get

$$\mathbf{v}'|_{M_\epsilon} = \epsilon \left[\frac{\partial \mathbf{v}_0}{\partial t} + \frac{\partial \mathbf{v}_0}{\partial \mathbf{x}} \cdot \mathbf{v}_0 \right] + \epsilon^2 \left[\frac{\partial \mathbf{v}_1}{\partial t} + \frac{\partial \mathbf{v}_1}{\partial \mathbf{x}} \cdot \mathbf{v}_0 + \frac{\partial \mathbf{v}_0}{\partial \mathbf{x}} \cdot \mathbf{v}_1 \right] + \mathcal{O}(\epsilon^3). \quad (4.13)$$

Constraining (4.11b) to M_ϵ yields

$$\mathbf{v}'|_{M_\epsilon} = \mathbf{u}(\mathbf{x}, t) - (\mathbf{v}_0 + \epsilon \mathbf{v}_1 + \dots) + \epsilon \beta \frac{D\mathbf{u}(\mathbf{x}, t)}{Dt}. \quad (4.14)$$

Equating both, matching orders of ϵ and moving back to the original time scale, we have to $\mathcal{O}(\epsilon)$

$$\mathbf{v} = \mathbf{u} - \epsilon(1 - \beta) \frac{D\mathbf{u}}{Dt} + \mathcal{O}(\epsilon^2). \quad (4.15)$$

Changing to $\epsilon = St$ gives us the *effective flow*

$$\mathbf{u}_e = \mathbf{u} - St(1 - \beta) \frac{D\mathbf{u}}{Dt} + \mathcal{O}(St^2). \quad (4.16)$$

This tells us that inertial particles behave (to first order) as if they are non-inertial particles advected by the flow \mathbf{u}_e . In other words, solving the system

$$\dot{\mathbf{x}} = \mathbf{u}_e(\mathbf{x}, t), \quad (4.17)$$

for passive particles, results in dynamics that mimic the dynamics of inertial particles in the flow \mathbf{u} .

Note that while the original flow may be incompressible ($\nabla \cdot \mathbf{u} = 0$), this is no guarantee that the new effective flow will be. It is in fact rare for this to be the case outside of the trivial case of neutrally buoyant particles ($\beta = 1$), where the dynamics reduces again to passive tracers. This compressibility of the flow allows the clustering effect we know is characteristic of the inertial dynamics to be captured in this simplified model. We will only be interested in the heavy particle limit corresponding to $\beta \rightarrow 0$, and to reflect this we set $\beta = 0$ from here. This is common practice when examining heavy particles (Bec, 2003; Monchaux et al., 2012; Rubin et al., 1995).

Using this effective flow allows for a reduction in the cost of simulating inertial particle dynamics by approximating the full dynamics with the motion of non-inertial particle advected by this effective flow, for only a minor cost to accuracy. Section 4.3.3

demonstrates the scope of the approximation and its ability to resolve the full inertial dynamics well.

The biggest advantage to this approximation though is not in the slightly reduced cost of simulation, but in the closed form of the divergence of the effective velocity field $\nabla \cdot \mathbf{u}_e$. Without this, calculating the metric chosen to quantify clustering in this work, $\delta(t)$ in (2.12), would be substantially harder from equation (4.8). Combining this with the importance sampling algorithm we intend to use also would be difficult and so the approximation provided by (4.16) makes the problem tractable.

4.2 2D Synthetic Random Flow

4.2.1 Nested OU process

We now make use of the reduced, effective-flow description to study the clustering that results from the effect of inertia. We consider a model flow that is random in time and has a simple single-scale spatial structure. Using this flow we examine the large-deviation statistics of δ .

The flow is a two-dimensional, synthetic random flow. The flow is divergence-free and we define it via the streamfunction

$$\Psi(x, y, t) = \frac{L^2}{T} \left[A_1(t) \cos\left(\frac{2\pi x}{L}\right) + A_2(t) \sin\left(\frac{2\pi x}{L}\right) \right. \quad (4.18)$$

$$\left. + A_3(t) \cos\left(\frac{2\pi y}{L}\right) + A_4(t) \sin\left(\frac{2\pi y}{L}\right) \right], \quad (4.19)$$

with

$$u = -\partial_y \Psi, \quad v = \partial_x \Psi. \quad (4.20)$$

Here L is the domain length, and A_i are independent stochastic processes. Each process is driven by noise and solves a ‘nested’ Ornstein-Uhlenbeck process (Pavliotis, 2014; Gardiner, 2009) defined by the stochastic differential equations

$$dA_i = -\gamma A_i dt + \sigma B_i dt \quad (4.21a)$$

$$dB_i = -\Gamma B_i dt + \Gamma dW_{it}, \quad (4.21b)$$

where γ , σ and Γ are positive parameters, and W_{it} are independent Wiener processes.

We do not specify the A_i themselves as Ornstein-Uhlenbeck processes but use the ‘nesting’ construction to ensure that the flow is smooth in time (up to C^1). This is needed to ensure that the derivative du/dt is properly defined, as it appears in \mathbf{u}_e . Note that in the limit $\Gamma \rightarrow \infty$ white noise is recovered with $B dt = dW$, and A becomes a simple OU process in this limit.

We set $\mathbb{E}A_i^2 = 1$ for $t \gg 1$, and as a result the number of parameters in system can be reduced. To show this we first write an expression for dA^2 . Using the fact that $d(uv) = (u + du)(v + dv) - uv$

$$d(A^2) = 2A dA - (dA)^2 = -2\gamma A^2 dt + 2\sigma AB dt, \quad (4.22)$$

where we drop the subscript on both processes for now. We also need expressions for dAB and subsequently for dB^2

$$d(B^2) = 2B(-\Gamma B dt + \Gamma dW_t) + \Gamma^2 dt, \quad (4.23)$$

$$d(AB) = dA B + A dB + dA dB \quad (4.24)$$

$$= -\gamma AB dt + \sigma B^2 dt \quad (4.25)$$

$$- \Gamma AB dt + \Gamma A dW_t. \quad (4.26)$$

Taking expectations for each and differentiating in time

$$\frac{d}{dt}\mathbb{E}(A^2) = -2\gamma\mathbb{E}(A^2) + 2\sigma\mathbb{E}(AB), \quad (4.27)$$

$$\frac{d}{dt}\mathbb{E}(B^2) = -2\Gamma\mathbb{E}(B^2) + \Gamma^2, \quad (4.28)$$

$$\frac{d}{dt}\mathbb{E}(AB) = -(\gamma + \Gamma)\mathbb{E}(AB) + \sigma\mathbb{E}(B^2). \quad (4.29)$$

We are interested in the stationary solution corresponding to $t \gg 1$ which yields

$$\mathbb{E}(B^2) = \frac{\Gamma}{2}, \quad (4.30)$$

$$\mathbb{E}(AB) = \frac{\sigma}{(\gamma + \Gamma)}\mathbb{E}(B^2) = \frac{\sigma\Gamma}{2(\gamma + \Gamma)}, \quad (4.31)$$

$$\mathbb{E}(A^2) = \frac{\sigma}{\gamma}\mathbb{E}(AB) = \frac{\sigma^2\Gamma}{2\gamma(\gamma + \Gamma)}. \quad (4.32)$$

Hence $\mathbb{E}A^2 = 1$ requires that

$$\sigma = \sqrt{\frac{2\gamma(\gamma + \Gamma)}{\Gamma}}, \quad (4.33)$$

which in the white noise limit for the nested process $\Gamma \rightarrow \infty$, we find $\sigma \rightarrow \sqrt{2\gamma}$. This is the expected result if we had begun with the OU process $dA = -\gamma A dt + \sigma dW_t$.

In what follows, we choose $\Gamma = \gamma$ for simplicity. This leaves $\gamma = \Gamma = \sigma^2/2$ as the sole remaining parameter. This single parameter will be used as a proxy for the inverse correlation time in the flow.

We can non-dimensionalise the stream function (and in turn the velocity field) as we did for the simplified equations of motion (4.8) with a length L (the domain size), and time $T = L/U$, based on the time for a particle to move a distance L . With this we have the scaled variables $\mathbf{x}^* = \mathbf{x}/L$ and $t^* = Ut/L$. Recasting the stream function in terms of these new variables we have

$$\Psi^*(x^*, y^*, t^*) = A_1(t^*) \cos x^* + A_2(t^*) \sin x^* + A_3(t^*) \cos y^* + A_4(t^*) \sin y^*. \quad (4.34)$$

From here, the asterisk on scaled variables will be dropped for ease of notation. This non-dimensionalisation extends to the correlation time scale controlled by γ which has units $\gamma^* = \gamma L/U$.

4.2.2 Explicit form of velocity field and effective velocity field

Using equations (4.34) and (4.21) we can write an explicit form of the velocity field and the derivatives needed for the effective velocity field (4.16). The non-dimensional velocity field is

$$\mathbf{u}(\mathbf{x}, t) = \begin{pmatrix} -\partial_y \Psi \\ \partial_x \Psi \end{pmatrix} = \begin{pmatrix} A_3(t) \sin y - A_4(t) \cos y \\ -A_1(t) \sin x + A_2(t) \cos x \end{pmatrix}. \quad (4.35)$$

Qualitatively, this is a vortical flow that consists of a pair of cells within the domain with time varying aspect ratio. A pair of example streamlines for the flow at fixed times are shown in figure 4.1. As t changes, this behaves like a change in the aspect ratio, and the location of the cell centre according to the values of the A_i .

To form the effective flow for inertial particles \mathbf{u}_e (4.16), we also need the material derivative. The convective derivative is straightforward to compute and simply results




Figure 4.1: Streamlines corresponding to two different realisations of the steady velocity field (4.35), demonstrating the effect of changing A_i and showing the cellular structure of the flow in a periodic domain.

in non-linear terms. The time derivative is non-trivial as it relies on the time derivative of the stochastic processes A_i . This is the reason for the nesting of the stochastic processes, which allows us to express the derivatives of the A_i as

$$\frac{dA_i}{dt} = -\gamma A_i + \sigma B_i. \quad (4.36)$$

The full acceleration term is then given by

$$\frac{\partial \mathbf{u}}{\partial t} = \begin{pmatrix} (-\gamma A_3(t) + \sigma B_3(t)) \sin y + (\gamma A_4(t) - \sigma B_4(t)) \cos y \\ (\gamma A_1(t) - \sigma B_1(t)) \sin x + (-\gamma A_2(t) + \sigma B_2(t)) \cos x \end{pmatrix} \quad (4.37)$$

and the convective derivative by

$$\mathbf{u} \cdot \nabla \mathbf{u} = \begin{pmatrix} (-A_1(t) \sin x + A_2(t) \cos x) (A_3(t) \cos y + A_4(t) \sin y) \\ (-A_1(t) \cos x - A_2(t) \sin x) (A_3(t) \sin y - A_4(t) \cos y) \end{pmatrix}. \quad (4.38)$$

A demonstration of the effectiveness of the approximation of passive particles in \mathbf{u}_e , compared to the full inertial particle dynamics can be found in section 4.3.3.

4.3 Numerically simulating the advection of inertial particles

In this section we outline the methods used to numerically integrate both the stochastic processes A_i and B_i and the dynamics for advection of particles in a flow. We also briefly outline a verification of the accuracy of the approximation of the full four-dimensional dynamics of inertial particles by the two-dimensional dynamics on the invariant manifold.

In general we will be computing the simple advection of passive tracer particles which is given by the ODE

$$\dot{\mathbf{x}}(t) = \mathbf{u}(\mathbf{x}, t). \quad (4.39)$$

We will deal primarily with the case where $\mathbf{u}(\mathbf{x}, t)$ is given by the effective flow approximation for inertial particle motion \mathbf{u}_e as in equation (4.16), but this equation is left general as it applies to the case for advection by the base flow (4.35) only.

4.3.1 Euler–Maruyama scheme for the stochastic velocity field

To begin to simulate the advection of particles in a flow, whether they are inertial or not, we first need to implement a method to compute the flow that is responsible for the advection. In both cases, we can note the one way coupling of the dynamics. The evolution of the flow is independent of the particle position and is solely time dependent. These coefficients can thus be computed independently of information on the particle position for any time step. The coefficients $A_i(t_n)$ are computed numerically using a standard Euler-Maruyama scheme

$$A_i(t_{n+1}) = A_i(t_n) - \gamma A_i(t_n) \Delta t + \sigma B_i(t_n) \Delta t, \quad (4.40)$$

$$B_i(t_{n+1}) = B_i(t_n) - \Gamma B_i(t_n) \Delta t + \Gamma \sqrt{\Delta t} \xi_i, \quad (4.41)$$

$$t_{n+1} = t_n + \Delta t, \quad (4.42)$$

where Δt is the time step and $\boldsymbol{\xi}$ is a vector of i normal random variables $\xi_i \sim \mathcal{N}(0, 1)$ that is sampled at every time step. In all of the experiments performed in this chapter we set $\gamma = \Gamma = 1$, with $\sigma = \sqrt{2\gamma} = \sqrt{2}$ as specified in section 4.2.1. The process is initialised with a random vector of initial conditions $A_i(t = 0) \sim \mathcal{N}(0, 1)$. The time

step for all simulations of the stochastic processes is $\Delta t = 5 \cdot 10^{-4}$.

With the values for $A_i(t_n)$ computed in this way, both the standard velocity field \mathbf{u} in (4.35) and the effective velocity field for inertial particles $\mathbf{u}_e = \mathbf{u} - \text{St}(1 - \beta) D_t \mathbf{u}$ can be treated as completely discrete-in-time functions. This discrete-in-time velocity field does not then need to be updated ‘live’ during the simulation and can be then treated as a function of space only $\mathbf{u}_n(\mathbf{x}) = \mathbf{u}(\mathbf{x}, t_n)$ at each time step, where the $A_i(t_n)$ are constants during each time step. This collection of coefficients can be computed a priori, outside of the integration along particle trajectories. By pre-computing the coefficients this way, we also ensure the same realisation of the velocity field can be used over multiple different experiments, while still being able to run them in parallel easily. Simply fixing the random seed is not enough, as when using the importance sampling algorithm, it also draws random samples to clone (or prune) with $p = 0.5$. This will be important when advection is combined with the importance sampling algorithm later where multiple experiments with an identical flow is critical.

4.3.2 Numerical scheme for particle advection

To integrate equation (4.39) for the particle position, we use a standard fourth-order Runge–Kutta (RK4). For a specified initial condition $\mathbf{x}_0 = \mathbf{x}(t = 0)$, the RK4 scheme is implemented as

$$\mathbf{x}_{n+1} = \mathbf{x}_n + \frac{\Delta\tau}{6} [K_1 + 2K_2 + 2K_3 + K_4], \quad (4.43)$$

$$\tau_{n+1} = \tau_n + \Delta\tau, \quad (4.44)$$

where $\mathbf{x}_n = \mathbf{x}(\tau_n)$ denotes the position vector at a time $\tau_n = n\Delta\tau$. The $K_{1,\dots,4}$ are given by

$$K_1 = \mathbf{u}(\mathbf{x}_n, \tau_n), \quad (4.45a)$$

$$K_2 = \mathbf{u}\left(\mathbf{x}_n + \frac{\Delta\tau}{2} K_1, \tau_n + \frac{\Delta\tau}{2}\right), \quad (4.45b)$$

$$K_3 = \mathbf{u}\left(\mathbf{x}_n + \frac{\Delta\tau}{2} K_2, \tau_n + \frac{\Delta\tau}{2}\right), \quad (4.45c)$$

$$K_4 = \mathbf{u}(\mathbf{x}_n + \Delta\tau K_3, \tau_n + \Delta\tau), \quad (4.45d)$$

where $\mathbf{u}(\mathbf{x}, t)$ is given by either of the flows outlined in 4.2.2. Importantly the time step used here $\Delta\tau = 2\Delta t = 10^{-3}$ is exactly twice the time step of the pre-computed stochas-

tic coefficients in section 4.3.1, ensuring the half steps have a value for A_i associated to them also.

Advection of inertial particles

Equation (4.43) will be used to advect passive particles in the flow (this is, both in the base flow \mathbf{u} and the effective flow \mathbf{u}_e used to approximate inertial particles). We also implement the phase-space dynamics for the advection of fully inertial particles in order to assess accuracy of the effective-flow approximation in 4.3.3. To do this we need to include the equation for the particles velocity also from equation (4.8). The full four-dimensional dynamics for inertial particles' is given by

$$\dot{\mathbf{x}} = \mathbf{v}, \quad (4.46a)$$

$$\dot{\mathbf{v}} = \frac{1}{\text{St}}(\mathbf{u} - \mathbf{v}). \quad (4.46b)$$

This is slightly more involved to solve numerically, as this is now a second order ODE represented as a first order system of ODEs, which are therefore coupled. We apply RK4 to this system but as \mathbf{x} and \mathbf{v} are coupled variables we must solve for them both at each increment simultaneously. We do this by vectorising the method in section 4.3.2 for a phase space vector of the position-velocity: $\mathbf{X} = (\mathbf{x}, \mathbf{v})$ as

$$\mathbf{X}_{n+1} = \mathbf{X}_n + \frac{\Delta\tau}{6} [\mathbf{K}_1 + 2\mathbf{K}_2 + 2\mathbf{K}_3 + \mathbf{K}_4], \quad (4.47)$$

$$\tau_{n+1} = \tau_n + \Delta\tau, \quad (4.48)$$

where the \mathbf{K}_i are given by

$$\mathbf{K}_1 = \mathbf{F}(\mathbf{X}_n, \tau_n), \quad (4.49a)$$

$$\mathbf{K}_2 = \mathbf{F}\left(\mathbf{X}_n + \frac{\Delta\tau}{2}\mathbf{K}_1, \tau_n + \frac{\Delta\tau}{2}\right), \quad (4.49b)$$

$$\mathbf{K}_3 = \mathbf{F}\left(\mathbf{X}_n + \frac{\Delta\tau}{2}\mathbf{K}_2, \tau_n + \frac{\Delta\tau}{2}\right), \quad (4.49c)$$

$$\mathbf{K}_4 = \mathbf{F}(\mathbf{X}_n + \Delta\tau\mathbf{K}_3, \tau_n + \Delta\tau), \quad (4.49d)$$

and \mathbf{F} is the vector-valued function

$$\mathbf{F}(\mathbf{x}, \mathbf{v}, \tau) = \begin{pmatrix} \mathbf{v}(\tau) \\ \frac{1}{\text{St}}(\mathbf{u}(\mathbf{x}, t) - \mathbf{v}(\tau)) \end{pmatrix}. \quad (4.50)$$

4.3.3 Effectiveness of the effective flow approximation

To demonstrate the effectiveness of the approximation we compare the results of the full dynamics for inertial particles in a flow according to equation (4.46), with the resulting dynamics from advecting particles as tracers with the effective flow \mathbf{u}_e given by equation (4.16) using the random flow (4.35) as a base flow. As a control, we also show the result of tracer particle advection with the unaltered random flow \mathbf{u} according to (4.35). All three experiments start with the same initial condition of a uniformly random distribution of particles in the domain, i.e. $(x, y) \sim U(0, 2\pi)$. The results are then obtained by integrating equations (4.16), (4.16) and (4.8) according to the methods outlined in sections 4.3.2 and 4.3.2.

Figure 4.2 shows these results, with plots of the particle location at various time steps. There is close agreement between the location of particles in both the effective flow and the fully inertial particles. Some discrepancy is expected as \mathbf{u}_e is only an approximation up to $\mathcal{O}(\text{St}^2)$ and so cannot capture the motion completely. This becomes more apparent at later times, when successive error has a chance to accumulate and begins to have a visible effect. In contrast, the motion of particles advected purely by the base random flow \mathbf{u} is markedly different for later times relative to the other two panes. While earlier in the evolution, there is resemblance between the two non-inertial flows, they visibly diverge with time. This is what we expect to see, given that the effective flow \mathbf{u}_e and base flow \mathbf{u} only differ by what is essentially a small correction in St .

A key feature that these plots demonstrate also is the effect of inertia (or the divergent term $\text{St}(1-\beta) D_t \mathbf{u}$ for the effective flow) introduces to the dynamics. At late times clustering is very apparent in the plots in the second and third column, whereas there is little evidence of any clustering in the first column corresponding to tracers advected in the non-divergent random flow (4.35). The first column demonstrates a flow that is well-mixing and the particles tend to a homogenous distribution. In contrast, the distribution of particles in the second and third columns tends to a strongly inhomogeneous distribution with large concentrations of particles residing on filamentary structure and large voids present in the distribution. This inhomogeneous, filamentary structure is the key behaviour that inertia introduces into the system and it is encouraging to see it captured by the effective flow approximation \mathbf{u}_e .

Overall, the approximation provided by \mathbf{u}_e is a good match to the full inertial dy-




Figure 4.2: Distribution of $N = 10^4$ particles in the domain at various times using different flows or inertial/non-inertial particles. The first column shows non-inertial particles in \mathbf{u} , the second column shows inertial particles in \mathbf{u} and the final column shows non-inertial particles in the effective flow \mathbf{u}_e , both for a Stokes number of $St = 0.1$ and correlation time $\gamma = 1$. The times are the same across each row and are $t = 30$, $t = 70$, $t = 100$ for the first, second and third row respectively.




Figure 4.3: Distribution of $N = 10^5$ particles in the domain at various times advected by the effective flow \mathbf{u}_e , including the initial distribution in the top-left, for $St = 0.1$, $\gamma = 1$. The plots show (left to right starting at the top-left) the distribution at times $t = 0, 4, 10, 50, 150, 185$.

namics given by (4.8), even if it is not a perfect substitute. Any pointwise error that the figure would indicate exists, is not as large as the error on the measurement of the statistics of δ that occur in practice. The simplification it affords to the computation throughout this work is also crucial, with the closed form for $\delta(t)$ it provides. Regardless, we acknowledge the limitations of using the approximation, but assume it is sufficient to capture the effects we are interested in examining, while making the use of importance sampling tractable in this case.

4.4 Results

In this section we present the results of experiments to measure the large-deviation statistics of the density of particles in the flow outlined in section 4.2.1. Until the final section where we examine the result for both varied St and γ , we fix the values $St = 0.1$ and $\gamma = 1$.

The results are presented in two ways: with plots of the rate function (and free energy) corresponding to the PDF (and moments) calculated numerically for the Lagrangian density distribution, and occasionally through plots of the domain showing

the physical locations of particles in an Eulerian interpretation. We will use the former as evidence for claims regarding clustering behaviour and the latter for illustrative purposes only generally. All particles in this section are advected passively using the effective flow approximation \mathbf{u}_e for the full inertial dynamics.

To begin we show in figure 4.3 the position of $N = 10^5$ particles at successive times along their trajectories, advected according to the method outlined in section 4.3.2. The figure demonstrates how the evolution of an initially uniform distribution of particles becomes rapidly inhomogeneous.

We have mentioned that the dynamics are dissipative, and in section 4.1.2 stated that there is a time dependent attractor for the dynamics in phase space that the particles tend to. There is also an attractor for the position of inertial particles, that is a projection of the phase space attractor onto physical space. This attractor is fractal in nature as a consequence of the dissipative nature of the of phase-space dynamics. The existence of such an attractor has been investigated for inertial particles in (Bec, 2003; Bec et al., 2024).

Figure 4.3 demonstrates the short relaxation time for the dynamics to converge toward the attractor present in the physical space (relative to the correlation time scale of the flow $\gamma = 1$). This can be seen along the top row of plots in figure 4.3 where even for very early times $t = 4$, the evidence of the clustering can be seen relative to the initial condition $t = 0$. At late times, in the plots along the bottom row, the flow does not qualitatively change as dramatically. We see that all of these plots contain similar structure – there are voids and also particles on a filamentary structure in the domain. There is still a non-zero level of compression is occurring at this time, but is just not visible at this scale.

4.4.1 Results with Large Deviation theory

To say anything quantitative about the clustering of particles and the density of the resulting distribution of particles, we turn to the large deviation theory. Following the groundwork laid out in section 2.3, we aim to describe the shrinking density distribution through large deviation statistics like the rate function or free energy.

To start, we introduce the finite-time rate function $s_t(\xi)$ that we use to approximate the true rate function defined in the limit $t \rightarrow \infty$. This will be how we calculate rate functions numerically from trajectories. The finite-time rate function is calculated




Figure 4.4: Estimate for the rate function $s(\xi)$ in the absence of any importance sampling methods. Left panel: estimate of $s(\xi)$ obtained at a finite value of t indicated by the colour of each curve. Right panel: same as left panel but for the smaller range $300 < t < 500$. The estimate of $s(\xi)$ in this range are averaged to obtain an average estimate indicated by the black dots and curve in the right panel.

through the relation (2.20)

$$P(\delta, t) \asymp e^{-ts(\delta/t)}. \quad (4.51)$$

We achieve this numerically by binning the values of $\delta(t)$ at a given late time ($t \gg \gamma$) to create a PDF, and then calculating the rate function:

$$s_t(\xi = \frac{\delta}{t}) = -\frac{1}{t} \ln P(\delta, t). \quad (4.52)$$

We will for ease of notation immediately drop the subscript t from this notation, setting $s_t(\xi) = s(\xi)$, as this will be the only version of the rate function we use in practice.

Figure 4.4 shows the estimate for the rate function calculated by averaging over multiple realisations of the rate function calculated at regular intervals. The rate function here is calculated for the same experiment that is shown in figure 4.3. The left plot shows a large time interval over which rate functions are repeatedly calculated for the distribution.

The right plot shows a shorter sub-interval at late time, with an the empirical average of the curves superimposed. This averaged version will be our best estimate for the rate function and is how we will compare results in this section. We use this averaged version to reduce the effect of any noise in the results that could occur for an individual realisation of the density PDF. There is an initial transient period for the

dynamics which is not shown in these figures, but this can be seen in the early times for the particle distribution shown in figure 4.3.

We can note some important aspects from these versions of the rate function at this stage. We would expect a mean compression overall in the system due to the inertial effects causing clustering, which we can see in figure 4.3, and that we have seen evidence for a mean compression through the flow being weakly compressible in equation (4.16).

We recall also from section 2.3 that the minimum of the rate function corresponds to the mean value of the PDF $P(\delta, t)$. This mean value is the average rate of exponential growth of density in the domain from an Eulerian perspective (or decrease in density if it is negative). Equivalently, it is the average compression (expansion if positive) experienced by parcels along trajectories. Figure 4.4 shows that the minimum is a (small) positive value, which matches the conditions for clustering to occur.

We observe also the shrinking of the range of ξ over which the rate function is resolved as t increases. This shrinking is centred around the minimum value of the rate function. This is no surprise as by definition when large deviation applies, the PDF is shrinking to its mean value over time. As a consequence, sampling the tails of the PDF - which correspond to the outer areas of the rate function - is becoming progressively more and more difficult and this is the cause of the shrinking effect.

This means that there is a trade-off between choosing a t that is large enough that the finite-time rate function is a good approximation to $s(\xi)$ but small enough that $s(\xi)$ can be approximated over a substantial range of ξ .

We can also estimate the free energy function $f(\alpha)$ from the distribution. In a similar way to equation (4.52), we define the finite time free energy $f_t(\alpha)$ that we use to approximate the true free energy that exists in the limit $t \rightarrow \infty$,

$$\mathbb{E}[e^{\alpha\delta}] \asymp e^{tf(\alpha)}. \quad (4.53)$$

In practice we calculate the finite time approximation to this by choosing an array of values for α to calculate f , and for each value we use the relation

$$f_t(\alpha) = \frac{1}{N} \sum_{i=1}^N e^{\alpha\delta_i(t)} \quad (4.54)$$

We once again for ease of notation drop the subscript t from the finite time free energy function and set $f_t(\alpha) = f(\alpha)$ as this will be the only version we use to calculate free




Figure 4.5: Estimate for the free energy function in the absence of any importance sampling methods, created by averaging over multiple estimates for the free energy calculated over a time interval similar to figure 4.4. Left panel: evolution of the individual free energy functions over a long time interval. Right panel: same as left panel but for a shorter interval at late times which is used to calculate the average curve that is superimposed on the left plot in black.

energy functions.

Figure 4.5 shows the estimate for this curve in a similar way to the rate function above. The left plot is again a long time interval showing the evolution of the free energy estimate, and the right plot is a shorter sub-interval with the empirical average of the curves superimposed.

We note a similar effect to the plots of the rate function. For large $|\alpha|$, especially negative values, we see that over time sampling becomes more difficult for these regions of the function. These values of large $|\alpha|$ correspond to higher order moments of the PDF $P(\delta, t)$. We saw in section 2.3.1 that there is a unique correspondence between the rate function and the free energy given by the Legendre transform. The parts that have shrunk for the rate function due to poor sampling are related to the areas in the free energy that are poorly resolved. The portions of the free energy that have constant slope indicate they are dominated by a single sample when we attempt to calculate them through the distribution for $P(\delta, t)$. This effect can be seen especially for negative moments, $\alpha < 0$. As a consequence, this poor sampling breaks the assumption of convexity we have for the free energy function and are unlikely to accurately reflect the true behaviour of the free energy function.

In addition to the slope of the curve, we observe the marked points for $\alpha = -1, 0$. We expect at the points $f(-1) = f(0) = 0$ based on section 2.3.1. The free energy curve intercepts the marked point at $f(0) = 0$ but not the point at $\alpha = -1$ where we can

clearly see that $f(-1) \neq 0$ as we would have expected based on the results in section 2.3.1. This highlights the poor sampling that occurs for negative moments $f(\alpha < 0)$.

The trajectories that contribute to most ($\alpha < -1$) negative moments are the ones that experience predominantly positive divergence over their lifetime, which is unlikely to occur by definition. This is the same portion of the rate function that has negative ξ . This is because any particle on a trajectory experiencing positive divergence would be ‘pushed’ away from this area of the domain, and in turn is attracted to areas of negative divergence.

This is the predominant issue with sampling this area of the free energy curve (and in turn the corresponding area of the rate function as a consequence of the Legendre duality). This effect is a symptom of the primary cause of the spatial sampling bias: regions where particles are unlikely to go are sparsely sampled by definition. To expand on the results and attempt to overcome this we incorporate the importance sampling algorithm introduced in section 2.4.1 to improve the accuracy of the large deviation curves, especially in these tail regions.

4.4.2 Importance sampling

We now examine the large deviation statistics with the importance sampling algorithm from section 2.4.1. This is set up much the same as the Geometric Brownian motion example in section 2.4.2.

We run multiple simulations in parallel for an array of 20 values of α evenly spaced in the interval $\alpha \in [-3, 3]$, each to resolve the free energy $f(\alpha)$ more accurately at that single value. Combining these values allows an improved estimate for the entire curve to be made. Each simulation uses the same set of coefficients as seen in the brute force simulations already, allowing a direct comparison to be made between the results. The initial condition is also identical for all importance sampled simulations, although we now use only $N = 10^4$ particles in each simulation, compared to $N = 10^5$ particles in the brute force simulations in the previous section. As discussed in section 2.4.1, the importance sampling parameters were varied to ensure they do not introduce unwanted bias to the importance sampling. Both the cloning parameter, and the perturbative kick applied to clones was varied to ensure no bias is introduced. A value of $P_c = 2$ was enough to ensure adequate importance sampling, and the kick was set to $\mathcal{N}(0, L \times 10^{-3})$, where $L = 2\pi$ is the domain size. Resampling was undertaken every $n = 20$ time steps.




Figure 4.6: Comparison of estimates for the free energy function with and without importance sampling methods. Left panel: same as right panel in figure 4.5. Right panel: estimate for the free energy function using importance sampling methods. Created by averaging over a time interval up to late times. A shorter time interval at late times is used to calculate the average curve that is superimposed on the left plot in black.

This was also tested through varying the number of steps for a test simulation to ensure the results were independent of the frequency.

4.4.3 Comparison of results from importance sampling with brute force

To compare the resulting estimate for the free energy curve created using the importance sampling algorithm to the one created with the brute force technique, we first compare the change in the estimate over time. The right panel of figure 4.6 shows the series of curves calculated for the importance sampled free energy over a the same interval as the brute force results shown on the left pane. The most striking difference between the pair is in the portion of the plot corresponding to negative α . Here the brute force curve completely fails as time progresses, with the failure to cross through $f(-1) = 0$ at later times in the experiment, and the constant slope with increasingly negative α that indicates poor sampling of the true free energy curve.

In contrast the importance sampled version in figure 4.6, shows a curve more indicative of what we would expect from the free energy. The curve is no longer a constant slope for negative α , and instead shows a more parabolic shape which aligns with what we predict in chapter 3. The importance sampled curve also correctly passes through the point $f(-1) = 0$ for all times shown in the plot. The estimate is generally stable for times greater than roughly $t = 100$, with all estimates after this time being in a small




Figure 4.7: Comparison of average estimates for the free energy function on a single plot that are the same as figure 4.6. Inset is a magnified region near the origin to highlight close agreement of the estimates in this region.

band around the average value. There are slight exceptions towards the tails but this area in particular is difficult to sample accurately, even with importance sampling. The range of α for which a reasonable estimate of the curve can be made is also larger than the range that can be achieved with brute force calculation. It should be noted also though that in both cases the range for the curve is specified before the experiment, but reliable results are not guaranteed for the entire range, but only obtained for a subset of this range. It is only through trial and error in creating the results shown here that the range in which the curves produce stable results is known. Even still, the greater range afforded by the importance sampling method is a positive result.

Figure 4.7 shows the estimate for the free energy computed as the empirical average over the time interval $t = 450 - 500$ for both the brute force method (red) and importance sampling (blue). This plot highlights the difference between the two methods to estimate the free energy, and indicates for which range of α the brute force calculation provides a reasonable estimate.

Figure 4.7 shows a very close agreement between the curves for a small range of alpha, roughly between $-0.5 < \alpha < 1$. Outside this range the estimate for the free energy from the brute force method is unreliable. Inside this range we would expect




Figure 4.8: Distribution of $N = 10^4$ particles in the domain at a single time, $t = 450$, for different values of α corresponding to importance sampled trajectories. Left panel: $\alpha = -2$. Centre panel: $\alpha = 0$. Right panel: $\alpha = 2.5$

the method to perform well though as the realisations that are sampled to make up this portion of the curve (the area near $\alpha = 0$) correspond to the typical behaviour in the system. As a consequence there are plenty of trajectories to sample for this range.

4.4.4 Identifying regions of the flow using importance sampling

In this section we briefly outline an interesting by-product of the importance sampling algorithm. In order to implement the importance sampling algorithm for the particles in the effective flow \mathbf{u}_e we need to simulate the advection multiple times in the same flow, one simulation for each α value we wish to target. As we state at the start of section 4.4.2, all of the simulations start with an identical initial condition and are evolved for the same set of coefficients A_i that make up the synthetic flow. Each simulation differs only by the value of α that is targetted, and it is this which influences the resulting particle trajectories.

The importance sampling algorithm selects particles that predominantly contribute to a particular moment or, equivalently, to a particular value of $\xi = \delta/t$. The particle positions obtained with importance sampling show the location of particles conditioned to this value of δ over their history.

We already have seen that the value of α dictates the moment of the PDF that is examined. This extends further in that each value of α corresponds to a unique value of $\delta(t)$ in the PDF $P(\delta, t)$. This can be seen by the Legendre duality of the free energy and rate function. As such, each simulation that is run for the importance sampling is targetting a specific section of the PDF that matches a density δ along trajectories. A similar phenomenon is observed in the case of large deviation statistics




Figure 4.9: Same as figure 4.8, but for a different time, $t = 190$.

for Lyapunov exponents to in [Tailleur and Kurchan \(2007\)](#), where through resampling, different regions in space corresponding to rare realisations could be identified.

The plots in figures 4.8 and 4.9 show the location of particles at two different times, for different values of α . The central panels correspond to $\alpha = 0$ and therefore shows the true position of the particles, obtained without any cloning or pruning.

Recall from section 2.4.1 a particle is cloned (or pruned) if its contribution to the expectation is higher (lower) than most other particles. Each particle contributes to the expectation taken through the ensemble mean

$$S_N = \frac{1}{N} \sum_{i=0}^N Q_i(t) \quad (4.55)$$

where each Q_i is given by the weighted history of $\delta(t)$ over a single trajectory,

$$Q_i(t) = w_i e^{\alpha \delta_i(t)}. \quad (4.56)$$

For a value of $\alpha = 0$, all trajectories have an equal value of $Q_i = w_i = 1$, and so no trajectory will ever reach the threshold to be cloned (or pruned) i.e. $Q_i > P_c S_N = 2$ will never occur as $Q_i = 1$.

Interestingly this is equivalent to the importance sampling algorithm attempting to sample the mass in the domain. We know that two fixed points exist in the free energy curve for the Eulerian moments at $\alpha = 0, 1$. These represent the volume ($\alpha = 0$) and mass ($\alpha = 1$) in the system as we have shown in section 2.3.1.

For $\alpha = 0$, the importance sampling, is biased toward sampling the mass in the domain, i.e. particles as they are, since each particle has the same initial mass. This aligns with sampling for $\alpha = 0$ being equivalent to no sampling at all. The same idea




Figure 4.10: Distribution of $N = 10^4$ particles in the domain at a single time (same as figure 4.8), for different values of α corresponding to importance sampled trajectories. Both panels correspond to zeros of the free energy and uniformly sample the volume (left) or mass (right) in the domain. Left panel: $\alpha = -1$. Right panel: $\alpha = 0$

holds for the $\alpha = -1$ moment, where the volume is uniformly sampled. This can be seen in figure 4.10, where we show the position of sampled trajectories for $\alpha = -1$ and $\alpha = 0$.

The left plots in both figures show a negative value of α , which correspond to values less than the mean in $P(\delta, t)$ and for negative enough values, tend to be regions of low particle concentration in this case. Comparing the left plots to the centre plots. The particles reside on two almost entirely distinct regions, with the negative α particles being heavily biased toward the voids seen in the $\alpha = 0$ plot.

It is important to remember though that sampling low α *does not* correlate with regions of low concentration in general. The importance sampling algorithm biases the trajectories according to their history, not the current state of the flow. It is just by the nature of this flow that these two things align, at least visually in this case. As the flow has a relatively long correlation time, $\mathcal{O}(\gamma) = 1$, the regions that trajectories seeking positive divergence sample, does not move very quickly relative to the trajectory. As a result a well defined border for these regions, that generally eject particles, can be created as seen in the comparison between the left and centre plots.

A similar observation can be made for the rightmost plots in both figures 4.8 and 4.9. Here the importance sampling algorithm is targetting trajectories that experience higher than usual compression over their history. Here the region occupied by particles is not distinct from the one in the centre plot. Instead the region occupied by the

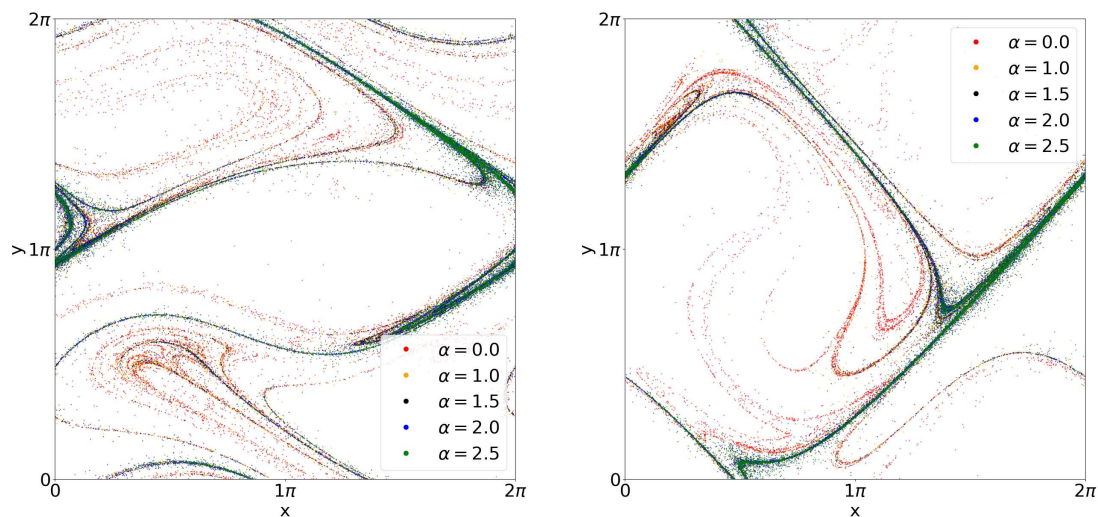


Figure 4.11: Multiple distributions of $N = 10^4$ particles in the domain at a single shared time, each colour represents a different value of α corresponding to importance sampled trajectories. Left panel: $t = 250$. Right panel: $t = 450$ (same as figure 4.8)

particles is a subset of the one shown in the centre plot. It is not surprising to see this, as the sampling bias inherent in the brute force calculation, is biased towards trajectories which historically experience net convergence, as we have seen from the rate functions earlier in this chapter. Biasing the calculation further again to trajectories experiencing much higher compression naturally leads to sampling a subset of the original set. In particular the areas that are shown to be sampled in the rightmost plot for positive α , correspond to the main filaments in the centre plot, implying compression expressed in these areas is higher than average. These areas also directly border the main voids seen in the centre plots. This implies that the filaments created through the dynamics are the areas with the highest compression in the flow.

Figure 4.11 shows the effect of increasing $\alpha \geq 0$ by superimposing snapshots from multiple importance sampled simulations for a single flow, at a single time. Heuristically, particles appear to congregate on smaller subsets of the domain with increasing α , all of which appear to be approximately a subset of the locations of particles for the non-importance sample trajectories corresponding to $\alpha = 0$. This shrinking of the space that particles cluster on makes sense as for increasing α , as the number of trajectories that would sample such an area of $P(\delta, t)$ decreases by definition. As a result, these rare regions that correspond to high historic divergence are densely sampled through the importance sampling algorithm. The likelihood for particles not to be some subset of $\alpha = 0$ (outside small fluctuations in the flow) is expected to be low, and to decrease with increasing α . This is because $\alpha = 0$ corresponds to the mean value of divergence

along trajectories, and so it is unlikely to find an importance sampled trajectory for some above average historic compression ($\alpha > 0$) sampling a region corresponding to trajectories with a less than average historic compression ($\alpha < 0$), or even a region of expansion ($\alpha < -1$).

The observations in this chapter are not meant as rigorous results, but more of an interesting observation of the correlation between historic positive and negative divergence along trajectories and regions within the flow. It also serves as a demonstration of the physical effect of the importance sampling algorithm, by showing the change in where in the flow trajectories actually sample.

4.4.5 Scaling of the free energy with St and correlation time γ

Equipped with the importance sampling algorithm, and the increased accuracy of the free energy curve it yields, we can investigate the theoretical results of chapter 3. In this section we run multiple simulations while varying the correlation time of the flow γ , and also the Stokes number of the particles St . These two parameters have direct parallels to the parameters in chapter 3, that are discussed here.

Varying the correlation time γ

We first examine the correlation time of the flow, γ . This was presented in chapter 3 as the parameter for which in its limit $\gamma^{-1} \rightarrow 0$ describes Kraichnan flows.

The nested OU flow as defined in (4.21) is smooth in space with correlation time, γ . This velocity field has the same form as a Kraichnan flow in the limit $\gamma \rightarrow \infty$. In section 3.2, we find that in this limit, the resulting free energy for $\delta(t)$ in such a flow has the form

$$f(\alpha) = \frac{D}{\gamma} \alpha(\alpha + 1), \quad (4.57)$$

where D is the time integrated autocorrelation of the divergence of the flow. This prediction tells us that up to a scaling in γ we should have a parabolic shape for the free energy, with the slope at the origin dictated by the value of D . This value for D can be evaluated as the Eulerian autocorrelation in the Kraichnan limit.

This is due to time scale associated with the flow being much faster than the time scale for particle advection as a consequence of being in the Kraichnan limit for the flow. In this case, the Eulerian and Lagrangian form are interchangeable, as the divergence

sampled along a trajectory over a short time difference is approximately the same as sampling at a fixed location i.e.

$$(\nabla \cdot \mathbf{u}_e)(\varphi_t(\mathbf{x}), t) (\nabla \cdot \mathbf{u}_e)(\varphi_{t-\tau}(\mathbf{x}), t - \tau) \approx (\nabla \cdot \mathbf{u}_e)(\mathbf{x}, t) (\nabla \cdot \mathbf{u}_e)(\mathbf{x}, t - \tau). \quad (4.58)$$

The ability to evaluate the autocorrelation in an Eulerian frame here rather than in the Lagrangian form as stated in section 3.1 is not always possible. In general these two forms are not interchangeable and the Lagrangian form should be used, but in this case due to the short correlation time of flow we use the Eulerian expression.

With the importance sampling algorithm, and the ability to accurately predict a free energy curve $f(\alpha)$ we attempt to verify this numerically. Using multiple values of $\gamma = \{1, 2, 2, 4, 5\}$, and fixing $St = 0.1$ we simulate the particle advection by \mathbf{u}_e using the importance sampling algorithm. This entails creating one simulation per value of α to form $f(\alpha)$, similar to the results in the previous sections. This is then repeated for each value of γ . In total this consists of 21 values of α evenly spaced between $\alpha = -3$ and $\alpha = 2$ for 5 values of γ for a total of 105 simulations. For each value of γ , the resulting free energy curve is averaged in a similar way to the ones presented in section 4.4.2 to yield a single best estimate.

The results are shown in figure 4.12, after being scaled by a factor of γ as in (4.57). A reference curve of the same predicted shape $f(\alpha) = D \alpha(\alpha+1)$ for a value $D = 0.0494$ is shown on the plot also. We will explain the origin of this value later in this section.

There are two main points for the results in this figure: the applicability of the scaling $\gamma^{-1} f(\alpha)$, and the prediction for the value of D shown on the plot. Beginning with the scaling in γ , we see that the agreement for all of the curves is generally good, especially for values of $\gamma \geq 3$.

As a related point, the swift collapse of the curves to the reference curve in black also shows how forgiving the Kraichnan limit is. Already for values of $\gamma = 3 \ll \infty$ the approximation made under the assumption of the limit $\gamma \rightarrow \infty$ is very reasonable relative the reference curve in figure 4.12. This implies that the results of section 3.2 provide a good approximation even in scenarios outside, but close to the white noise limit.

This brings us to the second of the two main points, the calculation of the value for D in the plot. D is the time integrated autocorrelation of the divergence in the flow, which we can estimate numerically for this flow.




Figure 4.12: Estimates of the free energy function for fixed Stokes number St and varying correlation time γ . Reference curve showing parabolic shape in the limit uniquely determined with a value of $D = 0.0494$ for the slope at the origin. y -axis is scaled in γ^{-1} .

As the autocorrelation $R(\tau)$ is a function of the lag time $\tau = t - s$ only, we simply choose a single point in the domain (e.g. $\mathbf{x} = (0, 0)$) and compute a very long time series for $\nabla \cdot \mathbf{u}_e(\mathbf{0}, t)$ in order to calculate $R(\tau)$. Numerically, this is the discrete calculation

$$R(\tau_i) = \frac{1}{N-i} \sum_{j=1}^{N-i} (\nabla \cdot \mathbf{u}_e)(\mathbf{0}, t_j) (\nabla \cdot \mathbf{u}_e)(\mathbf{0}, t_j - \tau_i). \quad (4.59)$$

where $T = N\Delta t$ is the total length of the time series with step size Δt , $t_j = j\Delta t$ and $\tau_i = i\Delta t$. Once we have a result for $R(\tau)$, we simply sum over the lag times to find an estimate for D

$$D = 2\Delta t \sum_{i=1}^{N_\tau} R(\tau_i), \quad (4.60)$$

where the factor of two arises by the symmetry of the autocorrelation around $\tau = 0$.

Using this, we find the value used to create the reference curve $D\alpha(\alpha+1)$ shown in figure 4.12, $D = 0.0494$. The close agreement between this curve, and the free energy estimates computed, especially for higher values of γ is remarkable. This shows that,

accounting for the scaling of the free energy with correlation time γ , the statistics of the log-density $\delta(t)$ are determined entirely by the autocorrelation of the divergence.

Varying the Stokes number St

We now turn our attention to the Stokes number St . This is the parameter used to describe the effect of inertia. So far, in order to mimic the effect of inertia without simulating the entire position-velocity phase space, we have been using the effective flow approximation

$$\mathbf{u}_e = \mathbf{u} - St \frac{D\mathbf{u}}{Dt}, \quad (4.61)$$

where we have set $\beta = 0$ explicitly from (4.16). Here, $\nabla \cdot \mathbf{u} = 0$ and the value of St controls the strength of the divergence present in $\nabla \cdot (D_t \mathbf{u})$. We can see this flow is of the same form as the one discussed in section 3.1 described by (3.1):

$$\mathbf{u} = \mathbf{u}_0 + \epsilon \mathbf{u}_1 \quad (4.62)$$

with $\epsilon = St$. As a result of section 3.1 we found that the free energy curve for a flow of this form should scale in St as

$$f(\alpha) = St^2 D \alpha (\alpha + 1) \quad (4.63)$$

where D is once again the time integrated autocorrelation of the divergence in the flow and is independent of St . This form as presented in section 3.1 was purely heuristic, and was by no means a rigorous result. In this section we investigate the result numerically in order to confirm its applicability.

We again utilise the importance sampling algorithm to create estimates for the free energy curve like the ones in the last section, but this time for multiple values of $St = \{0.02, 0.05, 0.1, 0.2, 0.4\}$, and setting $\gamma = 3$. This comprises of running multiple simulations for an array of values of α and St . Accounting for the cross over with the previously run values of $\gamma = 3, St = 0.1$, this is a total of 84 new simulations. The resulting curves are averaged again to create a family of curves much like the ones produced for varying γ previously.

These curves are plotted in figure 4.13, after they have been scaled by St^2 as predicted in (4.63).





Figure 4.13: Estimates of the free energy function for varied Stokes number St and fixed correlation time $\gamma = 3$. Reference curves showing parabolic shape in the limit uniquely determined with a value of $D = 4.94$ and $D = 3$ for the slope at the origin. y-axis is scaled in St^{-2} .

Figure 4.14: Same as figure 4.13 but with $\gamma = 1$

Again the primary result is about the parabolic shape for the free energy curves. The difference this time though, is now there is a clear relationship between decreasing St and approaching a parabolic shape. Here reference curves are plotted for a pair of values for the diffusion coefficient D to demonstrate that the curves for a higher St are not just scaling poorly but are fundamentally not a true parabola in their tails. This aligns with what we have stated in chapter 3. The underlying assumption for the result (4.63), is that St is a small parameter in the context of the multiscale flow as in equation (4.62). It makes sense then, as we move toward larger values of St , that the expected parabolic shape starts to fail.

Even with this lack of pure parabolic shape for larger St , it is still significant that the dominant behaviour of the free energy is captured by the form $f(\alpha) = St^2 D \alpha(\alpha + 1)$. The free energy curves all being of the same magnitude once scaled by St^2 is evidence that the divergence acts over the square of the small parameter in (4.62), rather than simply linearly.

A further point to highlight, is that the free energy curves for smaller St , are approximated well by a parabola with coefficient $D = 4.94$. This is the same coefficient calculated in the previous section for the time integrated autocorrelation ($D = 0.0494$) that is subsequently scaled by a factor of $St^2 = (0.1)^2$, which is the value of St that

was used for its computation. This is a surprise, since the Eulerian autocorrelation is not necessarily the same as the Lagrangian form in general.

The fact that the estimates for $f(\alpha)$ with varied St are well approximated for this value of D is explained by the value of $\gamma = 3$ used here. This value of γ is large enough that the dynamics are well approximated by a Kraichnan flow, rather than any consequence of the value of St , and so the Lagrangian and Eulerian form of the autocorrelation here can be treated as equivalent. This can be shown by examining figure 4.14 where in this case for a smaller γ and the same values of St the approximation is not as strong for the correctly scaled value of $D = 4.94$. This once again shows the wide scope of what can be considered to be in this limit, where $\gamma = 3$ is sufficient for the two versions of the autocorrelation to be good approximations of each other as seen here.

This fact once again highlights that accounting for the scaling, this time in St , the behaviour of the statistics of the log-density is influenced primarily by a single parameter D , the time integrated autocorrelation of the divergence.

4.5 Conclusion

In this work we have computed large deviation statistics for the log-density of inertial particles $\delta = \ln \rho$, which we have improved the accuracy of by means of an importance sampling algorithm described in section 2.4.1. This allows us to sample the exponentially rare trajectories that influence the large deviation statistics.

We model inertial particles through an approximation that describes the motion on an invariant manifold. The effectiveness of this approximation was examined and found to agree well. By using this approximation we have a closed form for the weak divergence in the velocity field experienced by inertial particles. This weak divergence is proportional to the Stokes number St which is assumed to be small. The approximation is used with a two dimensional random flow controlled by a stochastic process with correlation time γ .

Two limits corresponding to results from chapter 3 are examined. These are for $St \ll 1$ and $\gamma \rightarrow \infty$. In both cases the predicted quadratic shape of the free energy was found to be true, alongside the scaling of the free energy in St^2 , and linearly in γ . In particular the Kraichnan limit of $\gamma \rightarrow \infty$, applied to values not strictly in this limit.

Chapter 5

Ocean Surface Flow

We have looked at the clustering of inertial particles in a 2D synthetic flow, and while inertial particles do occur in nature, the flow that we used to advect them was highly idealised. We want to utilise the tools and results developed in chapters 2 and 3 for a more physical application. In this chapter we aim to examine passive particles in a two-dimensional weakly divergent flow, that can be used to model the clustering of floating particles on the surface of the ocean. We then relate this clustering on the ocean surface to the theoretical work outlined in chapter 3, matching the scaling in the small parameter, ϵ from chapter 3, to the relevant small parameter for the context of the ocean, the Rossby number Ro . An estimate for the fractal dimension of the physical space which the particles cluster on is made, based on the methods outlined in section 2.5.2 and its scaling with Ro is also examined.

Lagrangian particles on the ocean surface is a topic of active interest, with existing work on Lagrangian drifters found, for example, in Maalouly et al. (2023), Zhang et al. (2019) or Berti et al. (2011). In particular Maalouly et al. (2023) examines particles advected by SQG⁺¹ and measurements of clustering for those particles, drawing similar conclusions to this work, but with more focus on the correlation between regions of vorticity and clustering.

We choose not to examine inertial effects in the context of the ocean, instead focusing on the advection of passive particles in a compressible flow in this chapter. This choice is made because for inertial effects to be non-negligible in the ocean, the particles would need to be of a substantial size. We can estimate this size with some basic dimensional analysis. As we saw in section 4.1.1, inertial effects are usually measured

by the Stokes number St . The Stokes number is given by

$$St = \frac{2}{9} \frac{a^2 U \rho_p}{\nu L \rho_f}, \quad (5.1)$$

where a is the particle radius, L is the typical domain length scale, U the typical velocity of the flow, ρ_p and ρ_f are the density of the particles and carrier fluid respectively, and ν is the kinematic viscosity. Ignoring the prefactor, and taking neutrally buoyant particles ($\rho_p = \rho_f$), which we will be examining in this chapter, we can use typical values for ocean flows to estimate the particle size needed for inertial effects to be important ($St = \mathcal{O}(1)$). The kinematic viscosity of seawater is approximately $10^{-6} \text{ m}^2 \text{ s}^{-1}$, we would expect a typical ocean velocity to be roughly 1 ms^{-1} , and the domain we will examine to be of the order of 10 km. Using these values, in order for the Stokes number to be of order unity, we would need particles of radius

$$a = \sqrt{\frac{[St][\nu][L]}{[U]}} = \sqrt{\frac{[10^0][10^{-6}][10^4]}{[10^{-1}]}} \approx 0.3 \text{ m}. \quad (5.2)$$

The size of particles we wish to investigate are of a much smaller size, for example particles of size $a = 1 \text{ cm}$, yield a Stokes number of $St = 0.01$, indicating that inertial effects are no longer relevant. While we do not examine the dynamics of inertial particles in the ocean in this thesis, there are existing studies for the effect of inertia in the ocean and 2D turbulence in general, for example, [Beron-Vera et al. \(2019\)](#) and [Olascoaga et al. \(2020\)](#) for inertial particles in the ocean, and [Boffetta et al. \(2003\)](#) or [Goto and Vassilicos \(2006\)](#) for inertial particles in 2D turbulence. To this end, we seek the simplest model that can describe the clustering observed at the surface of the ocean.

5.1 Model for the ocean surface

A common choice for this objective is the Surface Quasi-Geostrophic (SQG) model ([Blumen, 1978](#)). This is a special case of Quasi-Geostrophic (QG) equations ([Charney, 1971](#)), that apply only to the (two-dimensional) surface flow rather than the full three-dimensional flow in the ocean. Quasi-Geostrophy is the approximation of the equations of motion for a rotating stratified fluid with a small Rossby (Ro) number. The Rossby number is a dimensionless measure of the ratio of inertial forces to the Coriolis force

(rotation) given by

$$\text{Ro} = \frac{U}{fL}, \quad (5.3)$$

where U and L are respectively the characteristic velocity and length scales of the flow, and f is the Coriolis frequency. This has been shown to be a good approximation for ocean dynamics (Lapeyre, 2017), but it is non-divergent and thus clustering cannot occur in this model .

Hakim et al. (2002) extends the SQG model to the SQG⁺¹ model by introducing a higher order term to SQG that can be divergent. This is based on an extension to the original QG equations (Muraki et al., 1999) that was expanded to SQG also. The next order correction introduces weak divergence into the system, much like the effective flow approximation for inertial particles in section 4.1.2 introduces a correction term that is divergent.

The SQG⁺¹ model relies on substantial background material around the QG and SQG equations and so in order to frame SQG⁺¹ appropriately and interpret the results in the context of geophysical fluids, we briefly outline the main points for the transition from QG to SQG, and then the extensions of QG⁺¹ and SQG⁺¹ also can be found in appendix A. The material in appendix A is all standard material, and largely follows the work in Vallis (2017)

In summary of appendix A, the SQG⁺¹ equations read

$$[\partial_t + u\partial_x + v\partial_y]\theta^s = 0 \quad (5.4a)$$

$$\hat{\Phi}_0^s(k, l) = \frac{\hat{\theta}^s(k, l)}{\kappa}, \quad (5.4b)$$

where θ is the surface temperature, $\hat{\Phi}^s$ is the Fourier transform of the leading order term of the expanded streamfunction, and $\kappa = \sqrt{k^2 + l^2}$ is the modulus of the wave numbers. A superscript θ^s , denotes a variable evaluated on the surface $z = 0$. The velocities in (5.4a) are given by the relations

$$v^s = \partial_x \Phi_0^s + \epsilon(\partial_x \Phi_1^s - \partial_z G_1^s), \quad (5.5a)$$

$$-u^s = \partial_y \Phi_0^s + \epsilon(\partial_y \Phi_1^s - \partial_z F_1^s). \quad (5.5b)$$

The exact forms for Φ , F , and G are listed in equation (A.78), but the details of how

they arise is non-trivial and is left for appendix A. The system has the significant property that while the inversion at both orders is three dimensional, the calculation of the surface velocity at both orders can be computed solely from the surface temperature field θ^s .

With SQG⁺¹, we have a model for ocean surface flow that is two dimensional and takes into account higher order behaviour beyond the leading order in Rossby number. We identify the leading order and correction velocity fields $(u^s, v^s) = (u_0^s + \epsilon u_1^s, v_0^s + \epsilon v_1^s)$ in (5.5) for ease of notation in subsequent sections. The higher order correction allows the flow to be divergent through the F and G terms in (5.5), meaning clustering behaviour can be observed even though the leading order velocity field is by definition incompressible. The divergent part of the flow on the surface is given explicitly by

$$\nabla \cdot \mathbf{u}^s = \epsilon(\partial_{xz}F_1^s - \partial_{yz}G_1^s). \quad (5.6)$$

5.2 Numerical implementation of SQG⁺¹

SQG⁺¹ is implemented in a similar way to the implementation of SQG, with modifications made to the governing equations to reflect the changes made in section A.1.5 and to introduce forcing to the system. The core two step process of inversion and advection is maintained as described in section A.1.4, as are the periodic boundary conditions.

Inversion now consists of recovering Φ_0 from θ^s using (5.4b) and using Φ_0 to build the velocity fields using (5.5). These velocity fields are used to advect θ^s in (5.4a).

We do not examine only freely decaying turbulence and so introduce forcing to the system, with the aim to reach a statistically steady state. A random, white in time, forcing term is added which acts over a band of width $2\delta_f$, centred at κ_f with an intensity A_F . Dissipation is required in the system to balance the energy input, which acts at both large and small scales. A hypodiffusion term, $\eta(\nabla_H^2)^{-2}$, is used to remove energy from the system at large scales, and a hyperdiffusion term, $\nu(\nabla_H^2)^4$, to dissipate energy at small scales. The governing equation for the system now reads, in wavenumber space,

$$\partial_t \hat{\theta}^s = -\mathcal{F}[u^s \partial_x \theta^s + v^s \partial_y \theta^s] - \nu \kappa^4 \hat{\theta}^s - \eta \kappa^{-2} \hat{\theta}^s + A_F 1_F, \quad (5.7)$$




Figure 5.1: Evolution of forced 2D turbulence with SQG resulting in a stationary flow. Panels are snapshots of the theta perturbation $\theta_p = \bar{\theta}^s - \theta^s$ at times $t = 500$ (top left), $t = 600$ (top right), $t = 800$ (bottom left), $t = 1000$ (bottom right).

where 1_F is indicator function for the set $F = \{\kappa : \kappa_f - \delta_f \leq \kappa \leq \kappa_f + \delta_f\}$.

Equation (5.7) is split into multiple stages to be solved numerically. Similar to the SQG case, the advection is solved pseudo-spectrally with a fourth order Runge-Kutta scheme. The forcing term is evaluated next in spectral space as a simple addition to $\hat{\theta}^s$. Finally the dissipative terms are resolved exactly in wavenumber space.

The initial condition is handled by an initially specified Φ_0 field. This is in turn used to create an initial θ^s with the use of the inversion relation (A.75). In the cases examined in this thesis we use an initial condition $\Phi_0 = \theta^s = 0$ which is allowed to ‘spin up’ through the forcing acting repeatedly on the domain before it reaches a steady state against the dissipative terms.

Forced 2D turbulence with SQG

In order to compare the implementation of SQG with SQG⁺¹, we set up a simple simulation of forced 2D turbulence, by introducing the forcing terms in section 5.2 to the model outlined in section A.1.4. This is a straightforward addition. Periodic boundary conditions are used with a spatial resolution of 1024×1024 wavenumbers to form a $L = 2\pi$ square domain with a temporal discretisation of $\Delta t = 5 \times 10^{-3}$. The dissipation constants are set at $\nu = 10^{-14}$ and $\eta = 5 \times 10^{-1}$. Forcing is carried out over a $2\delta_f = 2$ wide band of wavenumbers centred at $\kappa_f = 6$ with an intensity $A_F = (0.01 n_x^4 \Delta t)^{1/2} \approx 10^2$. The initial condition $\Phi_0 = \theta^s = 0$ is allowed to ‘spin up’ through the forcing acting repeatedly on the domain before it reaches a steady state against the dissipative terms.

Figure 5.1, shows the temperature field at multiple times, which we use to compare to SQG⁺¹ in the following section.

Forced 2D turbulence with SQG⁺¹

Using the setup outlined at the start of section 5.2, we use both the forcing as well as the large and small scale dissipation terms so that the flow reaches a statistically stationary state for which we use as the base for subsequent results. We once again use a periodic $L = 2\pi$ domain, discretised with a spatial resolution of $(n_x)^2 = 1024 \times 1024$ wavenumbers and a timestep $\Delta t = 5 \times 10^{-3}$. The dissipation constants are also set at $\nu = 10^{-14}$ and $\eta = 5 \times 10^{-1}$. An array of values for the Rossby number is used in subsequent sections. For now, we set $Ro = 0.05$ to demonstrate properties of the flow. This is the value used for all results where Rossby is not varied. An important point to note is that the Rossby number defined here is based on the lengthscale of the domain and not the scale of the vortices that emerge and so is an underestimate for the Rossby number associated with any individual vortex.

Forcing is carried out identically to the SQG case, with a $2\delta_f = 2$ wide band of wavenumbers centred at $\kappa_f = 6$ with an intensity $A_F = (0.01 n_x^4 \Delta t)^{1/2} \approx 10^2$. The initial condition $\Phi_0 = \theta^s = 0$ is allowed to ‘spin up’ through the forcing acting repeatedly on the domain before it reaches a steady state against the dissipative terms.

Figure 5.2 shows a series of snapshots of the resulting θ_p field to demonstrate what a typical snapshot looks like. The field is characterised by vortices of multiple scales interspersed by a fine filamentary structure, which plays a role in the creation of the new




Figure 5.2: Evolution of forced 2D turbulence with SQG^+ , with $\text{Ro} = 0.04$, resulting in a stationary flow. Panels are snapshots of the theta perturbation $\theta_p = \bar{\theta}^s - \theta^s$ at times $t = 500$ (top left), $t = 600$ (top right), $t = 800$ (bottom left), $t = 1000$ (bottom right).

vortices at small scales (Lapeyre, 2017). The difference between the temperature field for SQG in figure 5.1 and SQG⁺¹ for in figure 5.2 is not vastly different, as we would expect given SQG⁺¹ is simply a correction to SQG of order $\mathcal{O}(\text{Ro})$ between the two simulations. Differences are difficult to see. The primary difference is the compressibility of SQG⁺¹, which is the most important aspect for the clustering of particles. More detail on the subtle differences between SQG and SQG⁺¹ can be found in Hakim et al. (2002), but are not directly relevant to this work.

In section 5.3 we examine the evolution of passive tracers seeded in the flow and investigate the clustering behaviour of these particles. Section 5.4 looks at estimating the dimension of the attractor in the flow as a function of Rossby and section 5.5.1 examines the rate of particle clustering using large deviations theory and the effect of varying Rossby on the underlying probability distribution for particle concentration.

5.3 Clustering of particles in SQG⁺¹

In this section we consider the dynamics of Lagrangian tracers that are advected by the surface flow introduced in section 5.2. The motion of these passive tracers is entirely two dimensional and takes place on the surface field modelled by SQG⁺¹. The dynamics are given simply by

$$\dot{\boldsymbol{x}}_i = \boldsymbol{v}(\boldsymbol{x}_i, t), \quad (5.8)$$

where \boldsymbol{x}_i is the position of the i^{th} particle and $\boldsymbol{v}(\boldsymbol{x}, t)$ is the full surface velocity $\boldsymbol{v} = \boldsymbol{v}_0^s + \epsilon \boldsymbol{v}_1^s$. We drop the superscript s as we are only concerned with the surface fields from now on.

Particles are seeded in the flow once it has reached a steady state at a time t_0 . The $N = 128^2$ particles are initially positioned in a uniform grid spanning the entire domain. The particles are then allowed to evolve according to (5.8). At a given timestep each particle samples a number of quantities at its location, including the divergence field, which is used in this section to quantify clustering. As the flow is defined on a grid, and is not a continuous field like the inertial particle effective flow \boldsymbol{u}_e , interpolation is needed to advect the particles, and also to sample quantities in the domain. This is achieved through simple linear interpolation in space. Otherwise the implementation is identical to the one outlined in section 4.3.2. Care has been taken to ensure the results




Figure 5.3: Initial seeding of $N = 128^2$ particles in stationary SQG⁺ flow at $Ro = 0.05$ at a time $t_0 = 500$ (left) and the positions at a time $t = 975$ (right), overlaid on the RMS-normalised vorticity field at each time.

are robust, even with the interpolation for the velocity field. This proved not to be an issue, given the high resolution of the velocity field on the $N = 1024^2$ grid.

Figure 5.3 shows the initial seeded particles at time $t_0 = 500$ and a snapshot at a later time where the clustering of particles can be seen taking place by the non-uniform distribution of particles in the domain. There seems to be a subtle link between the vorticity field and regions of clustering, but the relationship is not clear.

The results we examine in this chapter focus on a varied Rossby number, Ro , which we treat as a small parameter. Unlike chapter 4, we do not have fine control over the correlation time for this model and so cannot comment on the effect a varied correlation time has on clustering behaviour. Furthermore, SQG⁺ (and SQG) both are inherently long correlation time flows (this is clearly demonstrated by the persistent vortices present in both flows) and so the results relating to Kraichnan flows, presented in section 3.2, do not apply here directly. Instead we look to use theory on the dimension of the underlying fractal attractor from section 2.5.2, and on explicitly multiscale flows from section 3.1 to characterise the clustering of particles.

5.4 Fractal set for clustering

We can see from figure 5.3 that particles are distributed in the domain in some non-uniform way. In order to describe this distribution of particles we turn to techniques from dynamical systems outlined in section 2.5 to quantify the space that particles reside on.

5.4.1 Calculating the fractal dimensions D_1 and D_2

We introduced the two fractal dimensions D_1 and D_2 we will use in section 2.5. Now we will outline their computation in the context of the SQG⁺ flow used here.

Calculating D_1

As we have mentioned in chapter 3 the phase space for the flow is simply the velocity space and so $\lambda_i = \lambda_{1,2}$. We take advantage of the fact that a volume element in physical space will expand (or contract as is the case in this system) at a rate proportional to $e^{(\lambda_1+\lambda_2)t}$. We already possess a good estimate for the contraction of space in the system from the mean compression which can be calculated from the Lagrangian sampled divergence of the flow, $\delta(t)$.

The maximal Lyapunov exponent can be estimated from the separation of two trajectories that are initially close together. In the long time limit, the exponential rate of separation of two initially close trajectories tends to the maximal Lyapunov exponent. With an estimate for both $\lambda_1 + \lambda_2$ and λ_1 we have enough information to calculate the D_1 using the Kaplan-Yorke conjecture as in (2.124).

To estimate the maximal Lyapunov exponent λ_1 , we will use the technique of Gram-Schmidt (GS) outlined in section 2.5. This is straightforward in our case as we are only interested in the stretching of a single vector rather than k -volumes and so we can skip the orthogonalising step. Recall we are interested in the tangent dynamics

$$Z(t) = J_t(\mathbf{a})Z_0 \quad (5.9)$$

where $Z(t)$ is the separation vector. As we are already simulating a number of Lagrangian trajectories in the domain, we now calculate $J_t(\mathbf{a})$ along them in addition to $\delta(t)$. This can be interpreted as each particle being the anchor of a separation vector as it is advected around the domain.




Figure 5.4: Sketch of the repeated normalisation at regular intervals for the separation vector $\delta \mathbf{x}$ used in the calculation of the maximal finite time Lyapunov exponent λ_1 .

Specifically, this consists of integrating (5.9) for a time τ from an initially unit length normalised separation $Z(0)/\|Z(0)\|$ to reach a state $Z^*(\tau)$, where we normalise as $\zeta_1 = \|Z^*(\tau)\|$. Equation (5.9) is then integrated again for a time τ but starting from the normalised initial condition $Z^*(\tau)/\|Z^*(\tau)\|$ to reach a state $Z^*(2\tau)$. This is subsequently normalised again with $\zeta_2 = \|Z^*(2\tau)\|$ and so on (Wolf et al., 1985). The stars indicate that these are not the ‘true’ solution to equation (5.9) but the solution to the piecewise dynamics including the normalisation step. This integration and normalisation is repeated a number of iterations k , where the separation can be expressed as

$$Z(k\tau) = Z^*(k\tau) \cdot \zeta_{k-1} \cdots \zeta_2 \cdot \zeta_1. \quad (5.10)$$

Using this we can estimate λ_1 as

$$\lambda_1 \approx \frac{1}{k\tau} \ln \frac{\|Z(k\tau)\|}{\|Z_0\|}, \quad (5.11)$$

$$= \frac{1}{k\tau} \ln \prod_{j=1}^k \zeta_j, \quad (5.12)$$

$$= \frac{1}{k\tau} \sum_{j=1}^k \ln \zeta_j. \quad (5.13)$$

A sketch of the process can be found in figure 5.4, illustrating the main idea of the algorithm. The choice of normalisation interval τ can impact the calculation, if it is too short then the underlying dynamics are not truly sampled, but if it is too long, overflow is more likely to occur in the calculation. In general the calculation is quite robust given a long enough time series which is the case in the work described here.

Calculating $\lambda_1 + \lambda_2$ is straightforward. Using the already computed values for $\delta(t)$, we simply take an average over particles to estimate the mean compression in the flow over time $\langle \delta(t) \rangle$. $\lambda_1 + \lambda_2$ is then computed as

$$\lambda_1 + \lambda_2 = \frac{1}{t} \langle \delta(t) \rangle. \quad (5.14)$$

Combining the estimates for λ_1 and $\lambda_1 + \lambda_2$, the information dimension is calculated as

$$D_1 = 1 + \frac{\lambda_1}{|\lambda_2|} = 2 + \frac{\lambda_1 + \lambda_2}{\lambda_1}. \quad (5.15)$$

Calculating D_2

The calculation of D_2 numerically is straightforward but intensive, as it requires calculating the pairwise distance for every particle at a given time. Recall the correlation sum scales with separation r as

$$C(r) \sim r^{D_2}, \quad (5.16)$$

with the finite correlation sum given as

$$C(r) = \frac{1}{N_p(N_p - 1)} \sum_{i=1, j>i}^{N_p} \Theta(r_p - |\mathbf{x}_i - \mathbf{x}_j|). \quad (5.17)$$

This can be readily calculated for the particle distribution for the positions that are already computed along trajectories. Using an array of logarithmically spaced values for r , we can estimate the exponent D_2 for the particle distribution.

Scaling of D_1 and D_2 with Ro

The clustering exhibited by tracer particles in SQG^{+1} flow scales with the divergence which is directly related to the Rossby number Ro . Here we estimate the fractal dimension of the attractor with varying Ro . Multiple simulations of SQG^{+1} flow are used, as outlined in section 5.3 with the only difference between simulations being Ro .

Figure 5.5 shows the scaling of the fractal dimensions D_1 and D_2 with Ro . For increasing Ro , both dimensions decrease monotonically implying the reduction in size of the structures within the flow in which particles congregate. This equivalently implies

an increase in the intensity of clustering behaviour with increasing Ro number. The dimension is still close to 2 and so we would expect particles to fill the space well, even if the set they reside on has zero area.

Using a fitted curve, we note the behaviour is remarkably close to quadratic in Ro . This demonstrates an important relation between the concepts discussed in chapter 3. The relation between D_1 and the Lyapunov exponents in equation (5.15), shows the fractal dimension depends directly on the mean compression in the flow

$$\lambda_1 + \lambda_2 = \frac{1}{t} \mathbb{E} \delta(t). \quad (5.18)$$

We also know from section 3.1.3 that the mean compression in the system $\mathbb{E} \delta(t)$ is given by the (scaled) autocorrelation of $\delta(t)$ along trajectories

$$\mathbb{E} \delta(t) = Ro^2 Dt. \quad (5.19)$$

Through both of these relations we see that the fractal dimension of the attractor is controlled directly by the autocorrelation of $\delta(t)$ along trajectories. Combining this with equation (5.15), we have the relation

$$D_1 = 2 - \frac{DRo^2}{\lambda_1}. \quad (5.20)$$

We expect λ_1 to be independent of Ro , and so the fractal dimension D_1 has a quadratic dependence on the small parameter Ro . Implicitly D_2 must also be quadratic in Ro , due to the relation of the fractal dimensions $D_1 \geq D_2$. This agrees with, and expands on existing work (Maalouly et al., 2023) on estimating D_2 which was found to scale with Ro^2 , but with no explanation for the cause.

5.5 Large deviation statistics for particles on the ocean surface

Section 5.4 explores the space the set that particles cluster on but provides no information on the rate at which this occurs, or the behaviour of the particle density distribution within the flow through time. Using the theory and methods outlined in chapter 2 we examine in this section the long term behaviour of the particle density distribution, in a similar way to the inertial particle case.




Figure 5.5: Values of the fractal dimension estimated using the information dimension D_1 (red) and correlation dimension D_2 (blue). Lines of best fit (black dash) constrained to be of the form $D = 2 - aRo^2$ calculated for each dimension superimposed.

As outlined in section 2.3, applying large deviations theory to the log-density $\delta(t)$ defines a rate function $s(\delta(t)/t)$ and free energy function $f(\alpha)$ that describe the late time behaviour of the probability distribution $P(\delta, t)$ and moments $\mathbb{E}[\delta(t)^\alpha]$. The rate function and free energy can be estimated for the ocean surface approximated through SQG⁺¹ by sampling the divergence along a large number of Lagrangian trajectories.

Recall the finite time rate function is given by

$$s(\xi = \frac{\delta}{t}) = -\frac{1}{t} \ln P(\delta, t). \quad (5.21)$$

We also estimate the finite time free energy function by choosing an array of values for α to calculate f , and for each value we use the relation

$$f(\alpha) = \frac{1}{N} \sum_{i=1}^N e^{\alpha \delta_i(t)} \quad (5.22)$$

Using the expression for the log-density $\delta(t)$ in (2.12), we need the divergence of the surface velocity field, including corrections in Rossby, along particle paths. The components of the divergence can be calculated directly from (5.5) and are given by

$$\nabla \cdot \mathbf{u} = \epsilon \nabla \cdot \mathbf{u}_1 = \epsilon (\partial_{xz} F_1^s - \partial_{yz} G_1^s).$$

Using this, we can calculate the compression along Lagrangian trajectories and examine the rate of clustering as described by large deviation statistics.

5.5.1 Large Deviation statistics of $\delta(t)$

Figure 5.6a shows the rate function calculated by estimating the distribution for the log-density $P(\delta, t)$ from the particle trajectories over a long time interval after a transient period from the initial seeding. The flow is already in a statistically steady state when this initial seeding takes place. Figure 5.6b shows a shorter time interval of the series shown in figure 5.6a alongside the average of the rate functions over the interval superimposed in black. This average is used to give a better idea of the underlying rate function we wish to estimate over the interval as each individual estimate is susceptible to variance in the flow itself.

From figure 5.6a the rapid shrinking of the estimated rate function for early times can be seen, consistent with the convergence to the mean in the PDF predicted by




Figure 5.6: Estimate for the rate function $s(\xi)$ in the absence of any importance sampling methods for SQG^{+1} . Left panel: estimate of $s(\xi)$ obtained at the finite value of t indicated by the colour of each curve. Right panel: same as left panel but for the smaller range $950 < t < 1000$.

large deviations. At later times, the estimate can be seen shrinking more slowly around the minimum over long times. This minimum also has positive drift as the mean compression experienced along trajectories increases with time, consistent with our expectation for clustering. As we have seen, this minimum value represents the peak of the log-density distribution.

The shrinking of the rate function toward the minimum is indicative of poor sampling. These shrinking areas in the range of ξ that is resolved correspond to tails of the underlying $P(\delta, t)$ which are exponentially decaying, and as such are increasingly harder to sample. Over short time intervals relative to this decay, such as the interval shown in 5.6b, the rate function has sufficiently converged over the interval to estimate the function. In this case we observe a rather flat rate function with a minimum that is quite close to 0, so the average behaviour is the result of a wide range of behaviours, with a substantial fraction of particles experiencing expansion.

We can also create an estimate for the free energy function from the sampled divergence along trajectories. Figure 5.7 shows the free energy function computed for the same time intervals as the rate functions shown previously. The convergence to the limiting behaviour at late times can once again be seen, but as $|\alpha|$ is increased, the unreliability of the sampling becomes clearer as these higher order moments are difficult to capture relative to the moments at the core area of the free energy corresponding to $|\alpha| \leq 1.5$.

The zeros at $f(-1) = 0$ and $f(0) = 0$ are marked on the plots, and we would expect




Figure 5.7: Estimate for the free energy function $f(\alpha)$ in the absence of any importance sampling methods for SQG^{+1} . Left panel: estimate of $f(\alpha)$ obtained at the finite value of t indicated by the colour of each curve. Right panel: same as left panel but for the smaller range $950 < t < 1000$.

the estimate for the free energy curve to pass through these points. The $f(0)$ point is trivial to capture from the form of $f(\alpha)$ in equation (5.22), however $f(-1)$ is particularly challenging to capture as discussed in section 2.3. The challenge in capturing $f(-1)$, and indeed capturing negative α at all, can be seen in the figure 5.7b, with the estimate falling below $f(0) = 0$ for most sampled times and the erratic evolution of the curve over time relative to that of positive α .

For further insight, we can compute the free energy by taking the Legendre transform of the calculated rate function, which allows us to compare the estimate of each statistic with the other. The plots in figure 5.8 show the empirically averaged rate function computed in figure 5.6, with the Legendre transform of the averaged curves in figure 5.7.

While similar, these curves disagree, in particular in the estimate for the mean compression $\langle \xi \rangle$ in the system. We have shown in section 2.3.1, the slope of the free energy at the origin, $f'(0)$, corresponds to the value of ξ at the minimum of the rate function, i.e. $f'(0) = \arg \min_{\xi} s(\xi)$. This minimum is the mode of the log-density distribution which in the long term limit is identically the mean. The disagreement seen in figure 5.8 stems from the fact that the estimate for the free energy curve measures the true mean, and the estimate for the rate function estimates the mode of the log-density distribution.




Figure 5.8: Comparison of average estimates for the large deviation statistics. Left panel: empirically averaged rate function computed in figure 5.6b with the rate function computed through the Legendre transformed free energy in figure 5.7b. Minimum values of the rate function marked on each estimate. Right panel: empirically averaged free energy computed in figure 5.7b with the free energy computed through the Legendre transformed rate function in figure 5.6b. Inset showing the magnified view of the slope at the origin for each estimate.

Importance sampling

We now examine the large deviation statistics with the importance sampling algorithm from section 2.4.1. We run multiple simulations in parallel for an array of 20 values of α evenly spaced in the interval $\alpha \in [-2, 1]$ each to resolve the free energy $f(\alpha)$ more accurately at that single value. Combining these values allows an improved estimate for the entire curve to be made. As before, each simulation uses an identical realisation of the SQG^{+1} flow, which is also the same flow used for the brute force estimates just shown. Analysis to ensure the simulations are independent of importance sampling parameters have been undertaken. Similar to the inertial case, they are chosen to avoid unwanted bias, and the cloning parameter is again set as $P_c = 2$ and the kick set to $\mathcal{N}(0, L \times 10^{-3})$. The initial condition for particles is also identical to the non-importance sampled case, comprising of a uniform grid of 128×128 particles across the domain.

Figure 5.9 shows the evolution of the importance sampled estimate for the free energy over time in a similar manner to the normally sampled case. The most striking difference is for negative α which no presents a consistent relaxing to a stable estimate for the free energy function for late times, in contrast to the erratic estimate when importance sampling was not used. A shorter interval for late times within the simulation is shown in 5.9b, highlighting for a shorter timescale, $f(\alpha)$ is approximately constant and overall the estimate has a minimal drift after an initial period of convergence.




Figure 5.9: Estimate for the free energy function $f(\alpha)$ using importance sampling methods for SQG^{+1} . Left panel: estimate of $f(\alpha)$ obtained at the finite value of t indicated by the colour of each curve. Right panel: same as left panel but for the smaller range $950 < t < 1000$.

Taking the average over a short interval, such as the one in 5.9b, we compare the regularly sampled and importance sampled estimates in figure 5.10. Both estimated curves are in good agreement for values of α close to 0, as we would expect them to be due to the relative ease of sampling this region. The most obvious improvement comes by comparing the curves outside the region of $\alpha \leq |0.5|$. Here, the brute force sampled curve begins to exhibit linear behaviour, indicative of undersampling. This linear behaviour implies a single realisation, or a single trajectory in this case, is responsible for the entire linear section of the estimate, echoing the exponential decay of rare events associated to the higher order moments that large deviations predicts. In contrast the importance sampled curve has an increasingly steeper slope with increasing $|\alpha|$ as we would expect, implying increasingly extreme events are also increasingly rare occurrences.

This lack of information can be seen also in figure 5.8b, where the Legendre transform of the rate function is markedly limited for negative α relative to the directly calculated brute force free energy curve. This is a sign of undersampling also as the information for these more negative α points comes from the presence of steeper negative slopes we would expect to see for smaller values of ξ . These regions in turn correspond to the area on the left hand side of the rate function, in figure 5.6a. Here we see the shrinking of the rate function around its minimum, especially for these values of ξ less than the value at the minimum, as these events relate to the tails of $P(\delta, t)$ are sampled less and less frequently across trajectories with time. The importance sampling algorithm works to combat this directly and provide a more accurate $P(\delta, t)$.

One obvious drawback is the incorrect crossing of the $f(\alpha) = 0$ axis near $\alpha = -1$. Recalling that this point on the free energy curve represents the preservation of area over the domain, which we expect to be conserved, the crossing at values less than $\alpha = -1$ implies some loss of area conservation over the domain with time. The fact that this error only becomes visibly apparent when coupled with the importance sampling is surprising, but potentially due to the manner in which the importance sampling algorithm specifically targets trajectories depending on the value of α it is tuned to. When targetting points on $f(\alpha)$ close to $\alpha = -1$, the ensemble of these trajectories will need to be representative of a conservation of global area rather than simply local area, a non-trivial task.

The importance sampling algorithm performs well, but not as well as it did for inertial particles. This is understandable given the increased complexity of the dynamics in this section compared to the idealised random flow in chapter 4.

The primary cause is the sensitivity to deviations in the background flow. While the flow used is statistically stationary as a whole, events such as merging vortices or collisions of vortices can cause sudden high levels of divergence in the flow, which in turn the importance sampling algorithm will heavily weight trajectories near this event, creating additional realisations around it and with high probability removing realisations elsewhere in the domain. This effectively reduces the statistics to a very localised area of the domain until sufficient time passes for a uniform redistribution through the chaotic dynamics to occur.

While these events are rare, sampling over long times increases the chance of sampling occurring over such an event. It is worth noting here also that for larger values of Rossby number Ro , the flow becomes more violent and variance within the stationary regime is larger. As a consequence, the occurrences of these spurious events also become more frequent and sampling becomes challenging.

While the importance sampling algorithm underperforms in the region close to $\alpha = -1$, it still provides a closer approximation as a whole to the true free energy curve we wish to capture. We can use this approximation to infer further properties of the system, such as how the curve scales with changing Rossby in the underlying SQG⁺¹ flow.

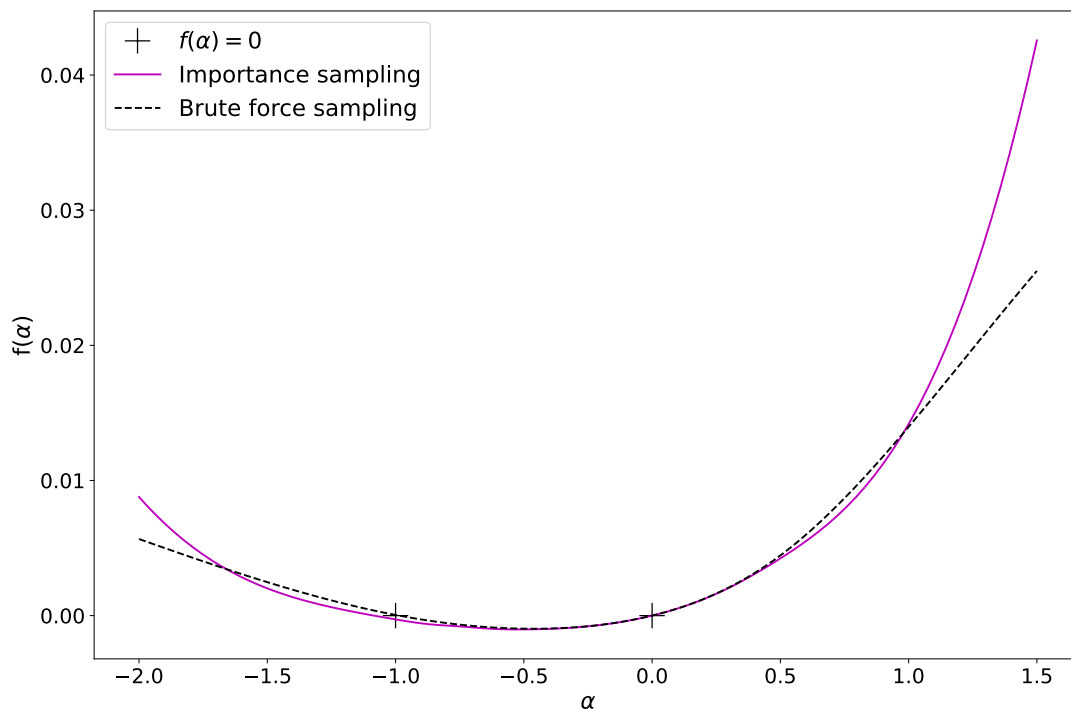


Figure 5.10: Comparison of average estimates for the free energy function for SQG⁺¹ on a single plot. Empirical average for each is computed from figures 5.7b and 5.9b. Zeros of the free energy $f(\alpha) = 0$ marked also.

5.5.2 Comparison of clustering regions in the flow with importance sampling

Figure 5.11 shows multiple snapshots of the particle distribution at the same instant in time for different values of α , along with a snapshot of the non-importance sampled simulation for comparison. These snapshots demonstrate the typical behaviour of the importance sampling algorithm.

From the snapshots it is possible to attempt to identify regions in the flow for different α 's similar to the inertial particles in section 4.4.4. These are regions that contain particles that historically are compressed at a greater level than average, or regions with particles that experienced historically lower levels than average, depending on the value of α .

A particularly noticeable trend is the clustering of particles on regions of positive vorticity in the $\alpha = 1$ simulation. This demonstrable preference for clustering in cyclonic regions is deceptive and is in fact an aversion to anticyclonic regions instead. This would align with the findings in Hakim et al. (2002) regarding cyclonic/anti-cyclonic asymmetry. They state there is a correlation with surface convergence (divergence)

and filaments between anticyclonic (cyclonic) vortices. This causes particles to cluster on the filaments between anticyclonic vortices rather than inside them, and hence appearing as if they preferentially reside on cyclonic regions. The effect is harder to see for negative α but particles here also sample areas between vortices as these are areas of positive divergence, when the cyclonic/anticyclonic argument is reversed. The structure for positive and negative α is similar, as they both are attempting to sample filaments between high vorticity regions, but the filaments they sample are distinct.

We note also the same behaviour for $\alpha = \{-1, 0\}$ that was seen for inertial particles in figure 4.8. Here the distribution of particles corresponding to $\alpha = -1$ appears generally devoid of any clustering. This is due to $\alpha = -1$ sampling uniformly for the volume of the domain. The distribution for $\alpha = 0$ also closely resembles the brute force sampling as we would expect.

We have seen in figure 5.10 that the uniform sampling for $\alpha = -1$ is not exact, which aligns with the small regions of clustering that can be seen around particularly high vorticity areas.

5.5.3 Scaling of the free energy with varied Ro

With the tools to estimate the large deviation statistics statistics, we can apply them to investigate the theoretical results of chapter 3. In this section we run multiple simulations while varying the Rossby number for the SQG^{+1} model.

In section 3.1, we discussed weakly compressible flows of the form,

$$\mathbf{u} = \mathbf{u}_0 + \epsilon \mathbf{u}_1 \quad \text{with} \quad \nabla \cdot \mathbf{u}_0 = 0, \quad \nabla \cdot \mathbf{u}_1 \neq 0. \quad (5.23)$$

As a result of section 3.1 we found that the free energy curve for a flow of this form should scale in Ro as

$$f(\alpha) = \text{Ro}^2 D \alpha (\alpha + 1) \quad (5.24)$$

for some constant effective diffusion D . This applies to SQG^{+1} where the correction to SQG is small and scales in Ro . In the SQG^{+1} case though, the constant D is implicitly a function of the small parameter Ro also due to the manner in which \mathbf{u} is defined. This comes from the inversion step defined in (A.75) where even though \mathbf{u}_0 itself does not explicitly depend on Ro , through the inversion of θ^s which is advected by the full




Figure 5.11: Distribution of $N = 128^2$ particles in stationary SQG⁺ flow at $Ro = 0.05$ at a time $t_0 = 975$ overlaid on the RMS-normalised vorticity field for different values of α .




Figure 5.12: Importance sampled free energy function for multiple Rossby numbers between 0.001 and 0.08. A reference curve for the predicted form of $f(\alpha)$ is superimposed as a black dashed line.

$\mathbf{u} = \mathbf{u}_0 + \epsilon \mathbf{u}_1$, the leading order term \mathbf{u}_0 is implicitly dependent on Ro .

As such, the free energy for changing Ro will not scale exactly as Ro^2 but with some contribution from this implicit dependence also. This is in contrast to the synthetic flow used in section 4 in which precise control over the diffusion coefficient D and the small parameter ϵ was possible. In that case the \mathbf{u}_0 is prescribed and so the scaling in St is definitive in the limit. This is not the case with SQG^{+1} where \mathbf{u}_0 comes from solving nonlinear Ro -dependent PDEs.

Figure 5.12 shows the free energy curves calculated as outlined in 5.5.1 for multiple Rossby numbers between 0.001 and 0.08. We see a strong collapse of the curves in the scaling of Ro when compared with what an unscaled version would show. This collapse is not perfect though. For small Ro the behaviour is close to parabolic, but with increasing Ro , we observe a departure from the parabolic shape indicated by the reference (black dash) lines superimposed on the figure. This departure from the parabolic form means moments grow much faster for these values of Ro and so there is large intermittency in these cases.

This departure for large Ro is likely explained by the argument for \mathbf{u}_0 being Ro dependent, but the validity of SQG^{+1} as a model for these larger values of Ro may also




Figure 5.13: Magnitudes of the leading order flow \mathbf{u}_0 , correction term $\text{Ro}\mathbf{u}_1$ and total surface velocity $\mathbf{u} = \mathbf{u}_0 + \text{Ro}\mathbf{u}_1$ of SQG⁺¹ over time, for a small ($\text{Ro} = 0.005$) and large ($\text{Ro} = 0.08$) Rossby number

play a part. The model assumes \mathbf{u}_1 to be a correction to the leading order flow \mathbf{u}_0 , but for larger values of Ro this is not always the case. Figure 5.13 shows the relative magnitudes of \mathbf{u}_0 and \mathbf{u}_1 and the sum of the pair. For small Ro , \mathbf{u}_1 is clearly a small correction, but for larger Ro , $\mathcal{O}(\mathbf{u}_0) \approx \mathcal{O}(\mathbf{u}_1)$. In this case the assumption in (5.24) also likely breaks down further as the divergent part of the flow $\epsilon \mathbf{u}_1$ is no longer weakly divergent. With this in mind, the curves representing the behaviour of large Ro (i.e. $\text{Ro} \gtrsim 0.05$) are not an accurate representation scaling we would expect to see from equation (5.24). Furthermore, for large Ro and large $|\alpha|$ (i.e. $\alpha < -2$ and $\alpha > 1$) the results are exceedingly difficult to sample due to the combination of the breakdown of the approximation as seen in equation (5.13) and the already difficult sampling for large $|\alpha|$, even with the importance sampling algorithm. The stark departure from the scaling and parabolic behaviour we would expect to see from equation (5.24) in these ranges, that can be seen in the results shown in figure 5.12, is therefore not unexpected and likely not an accurate representation of the true dynamics and rather a result of poorly resolved trajectories and undersampling.

5.6 Conclusion

We have estimated the the fractal dimension of the attractor for ocean surface flow modelled with SQG⁺¹. Two different fractal dimensions were calculated the information dimension D_1 and the correlation dimension D_2 . We show a general relation for the dependence of the fractal dimension on the Rossby number: $D_1(\text{Ro}) = 2 - D\text{Ro}^2/\lambda_1$.

This highlights both the general scaling in terms of the small parameter Ro of the physical space in which particles cluster, and again the important role of the autocorrelation of $\delta(t)$ along trajectories in determining clustering properties. At a basic level, this also implies an increase in the intensity of clustering with increasing Ro number, although with a dimension still close to 2. This emphasises that the spatial structure of the particle distribution is complex, as the particles are still relatively space filling. This can be misleading, as it still shows that floating particles will be hard to find, as the set they reside on has zero area, but when they are encountered they will likely be plentiful.

We find that large deviation statistics for varied Ro do not scale perfectly in Ro^2 as predicted in chapter 3. For small Ro the behaviour is close to parabolic, but with increasing Ro , we observe a departure from the parabolic shape. This is likely due to the breakdown of the approximation $\mathbf{u} = \mathbf{u}_0 + \epsilon\mathbf{u}_1$. As the leading order term \mathbf{u}_0 is also implicitly a function of the small parameter Ro . This departure from the parabolic form for large Ro means moments grow much faster for these values there is large intermittency in these cases.

Chapter 6

Summary and discussion

6.1 Summary

We have examined the clustering of particles as a result of weak divergence in the velocity field they are advected by. Two different applications are studied, with a different source of divergence in each case.

The first application studied was for inertial particles. The effect of inertia on the particles causes the particle velocity to differ from the fluid velocity. This effect causes the particles to cluster as the dynamics in position–velocity phase space are dissipative. A reduced model for the dynamics on an invariant manifold is used to simplify the equations. The reduced order model consists of advection by an effective velocity which includes a weakly divergent correction term proportional to the Stokes number St , which is assumed to be small. A two-dimensional random flow, controlled by an altered OU-process that is C^1 differentiable in time, with a correlation time γ , is used as the base flow.

The second application concerns particles floating on the surface of the ocean. These particles experience a two-dimensional weakly divergent velocity field. This velocity field is the horizontal part of a full three-dimensional non-divergent fluid velocity. This corresponds to the surface of the ocean being weakly divergent in a two-dimensional sense, while the dynamics as a whole for the ocean are incompressible. We model this using the SQG⁺ model. This is an improved-accuracy version of the commonly used surface quasigeostrophic (SQG) model. SQG⁺ includes a first-order correction in the Rossby number Ro . Unlike the leading-order, SQG velocity, the first-order correction is divergent. The resulting weak divergence leads to clustering.

A pair of theoretical results describing the long-term statistics of particles in weakly divergent flows are outlined in section 3. These results examine two distinct limits for divergent flows. The first is the limit of a weakly divergent flow, of the form $\mathbf{u} = \mathbf{u}_0 + \epsilon \mathbf{u}_1$, with $\nabla \cdot \mathbf{u} = \epsilon \nabla \cdot \mathbf{u}_1 \neq 0$. The second limit is that of short correlation time for the flow.

These results conclude that the large deviation statistics for the density of particles has specific asymptotic form for large time. This form is given by the rate function $s(\xi)$ and free-energy function $f(\alpha)$ which are found to be parabolic in the limit for both cases. The slope at the origin is determined by the autocorrelation of the flow. Since $f(-1) = f(0) = 0$ also, this determines the entire parabola.

The free energy is also found to scale linearly in the correlation time γ , and in the square of the small parameter ϵ , which corresponds to the Stokes number St for inertial particles, and the Rossby number Ro for the ocean surface flow.

6.2 Discussion of results

The theoretical work in chapter 3 has some limitations as we have seen in the results of this thesis. The long time statistics of particles in divergent flows is complicated by the convergence to a fractal attract. As mentioned in section 3.1, in the absence of diffusion, for a fixed (random) flow the particle-position PDF, is not a smooth function. For time-dependent flows, the support of this function is also time-dependent. We rely on the flow being sufficiently mixing that it can replace this diffusion, in which case averaging over time or realisations can lead to smooth PDFs. It is difficult to prove anything rigorous without this, but to make the mixing condition precise, let alone verify it for specific flows, is challenging.

In both applications, we numerically compute large deviation statistics for the log-density of particles $\delta = \ln \rho$. The accuracy of these estimates are improved by utilising the importance sampling algorithm outlined in section 2.4.1 to sample exponentially rare trajectories predicted by large deviations theory.

The importance sampling algorithm enable us to test the results of chapter 3 numerically. This is undertaken for a series of values of St and γ in the inertial particle case, and for a series of Rossby numbers Ro in the ocean surface case.

For inertial particles there is good agreement with the shape predicted in these results for both St and γ , and with the correct scaling in γ and St^2 . The Kraichnan

limit in particular was found to be robust for values away from this limit.

While the results in this thesis are well behaved, care must be taken when dealing with the short correlation time limit $\gamma \rightarrow \infty$ for inertial particles. The effective flow \mathbf{u}_e that is used to approximate the motion of particles is not valid for short enough correlation time, as the acceleration term in the approximation becomes too large. Further, the Maxey–Riley equation assumes a small Stokes number St , defined as the ratio of the Stokes drag time scale (τ_d), to the time scale of the flow. In the Kraichnan limit, a suitable time scale for the flow is its correlation time $\tau_c = \gamma t$, which itself should be much shorter than the advective time scale for the flow L/U . Combining these, we find that a suitable Kraichnan regime for inertial particles should therefore assume $\tau_d \ll \tau_c \ll U/L$. This can be achieved for multiple choices for each parameter, and there is no obvious choice of the relative size of each time scale that leads to natural distinguished limits.

For the ocean surface flow, the large deviation statistics for varied Ro do not scale perfectly in Ro^2 as predicted in chapter 3, but they do approach a parabolic form in the limit of $Ro \rightarrow 0$. This is due to the SQG⁺¹ not being strictly of the form $\mathbf{u} = \mathbf{u}_0 + \epsilon \mathbf{u}_1$. Due to the implementation of the model, specifically the inversion step for the stream function from the surface buoyancy field, the leading order velocity \mathbf{u}_0 is implicitly a function of Ro also, that is, the leading-order flow in SQG⁺¹ differs from the SQG flow. This means the results of chapter 3 are limited in their ability to predict the shape outside of small Ro where this effect is minimal. For larger Ro the parabolic shape, or the scaling proportional to Ro^2 cannot be expected to hold for the free energy function outside of very small values of $|\alpha|$.

We estimate the fractal dimension of the attractor for the ocean surface flow. Two metrics for the fractal dimension are used, the information dimension D_1 and the correlation dimension D_2 . Both metrics are found to scale proportionally to the square of the Rossby number, with a relation found to fit $D_{\{1,2\}}(Ro) = 2 - a_{\{1,2\}} Ro^2$. This highlights a relation between the mean compression in the flow and the fractal dimension of the space in which particles cluster. This can be seen through equation (5.20), which relates the mean compression, which scales as Ro^2 (as shown by the scaling argument $f'(0) = Ro^2 D$, in chapter 3), to the quantity $\lambda_1 + \lambda_2$ in the Kaplan-Yorke calculation for D_1 . This holds even for larger Ro unlike the parabolic shape of the free energy as the mean compression is still proportional to Ro^2 . This would rely on λ_1 being independent of Ro , which is not known for definite.

In the context of floating particles in the ocean, such as plastics or other pollutants, we have shown that particles cluster on zero-measure fractal sets. In particular though, their dimension is still close to 2 ($D \sim 2 - \text{Ro}^2$). This emphasises that the spatial structure of the particle distribution is complex, as the particles are still relatively space filling, even if the set they reside on has zero area. This suggests that floating particles are hard to find but that, when one is found, it is likely part of a cluster.

6.2.1 Potential applications

The work in chapter 4 is for idealised random flows, and as such is difficult to link directly to potential applications. With more realistic flows, the work could be readily extended to applications such as pollutants in the atmosphere or mixing in industrial flows (Tsai et al., 2005; Hessel et al., 2005). The dissipative dynamics of inertial particles means that clustering is inevitable, and as such the measurement of clustering through the distribution of the density of particles, as outlined in this thesis, is an important factor in working to accurately describing such phenomena.

The application of the work in chapter 5 is more evident. The dynamics of particles in the ocean is a constant topic of interest. With new advances in measurements for ocean flows (Morrow et al., 2019), the ability to link theoretical work, like the work from this thesis, to experimental field work is becoming more accessible, but is still limited to large scales. In order to link the work in chapter 5 to related field work, both the quality of the measurements for the ocean surface flow, and the particle drifter location would need to be of high resolution to allow a direct link to the grid based models used. Examining an area of 100km^2 with a $N = 1024^2$ grid for example would need measurements with resolution of approximately 100m, which is still several orders of magnitude finer the available data sets, which is of a resolution of $\approx 15\text{km}$. Measuring clustering on a global scale is feasible with this resolution, but in this case the SQG⁺¹ would breakdown as a global model for the ocean.

Lagrangian drifter data is readily available and a topic of active interest (Saraceno et al., 2024; Miron et al., 2021). The problem in these cases is that the devices which are used as drifters to obtain measurements of position are usually large resulting in an increased Stokes number St . With this increase in size, the inertial effects on these drifters become important and the assumptions of passive particles made in chapter 5 begins to break down and a direct comparison can no longer be drawn. Measurements

of the presence of small particles in the ocean, in particular for plastic have been found to be heterogenous (van Sebille et al., 2015, 2020). This implies the presence of a fractal set that clustering occurs on, as predicted in chapter 5, could be realistic in the applications to microplastic in the ocean.

6.3 Further work

There are a number of points of further work that are natural extensions of the results presented in this thesis.

One is the systematic derivation of the Kraichnan limit of the Maxey–Riley equation. The issue in attempting to identify such a limit lies with reconciling the effects of taking the limits $\gamma \gg 1$, and $St \ll 1$. The decrease in St causes a shrinking in the divergence as $\nabla \cdot \mathbf{u}_e \sim St$. As the divergence is weak, long time scales are needed to measure its effect, and with long time scales, the effects of the finite correlation time of the flow appear. In the short correlation time limit, the approximation for inertial particles begins to break down also. This is a result of the gradients in the velocity field causing the acceleration term to grow rapidly. In order to deal with these effects, a distinguished limit is needed that follows the regime outlined above, $\tau_d \ll \tau_c \ll U/L$. There are multiple potential choices for the three time scales that are applicable, but there is no obvious choice that can be made for the relative size of these time scales which would lead to natural distinguished limits. Work was carried out on identifying such a limit but has ultimately been unsuccessful so far, and as such it is not included in this thesis, but has been deferred as a topic of future work.

The main results on the spatial characterisation of clustering has been through the dimension of the fractal set that particles cluster on. Extending this to two point statistics, similar to (Haynes and Vanneste, 2005), is a natural extension and could provide further insight on the structure of the attractor.

In the work for this thesis, all initial conditions for particles have been drawn from a uniform distribution across the domain. A natural question may be around the case for non-homogenous initial distributions, or initially localised distributions. In the context of the results generated from flows discussed in chapter 3, each of these flows have strong mixing properties, and the combination of this with the divergence in the flow means that initially local distributions should not behave any differently over the timescales of interest.

Instead of the initial distribution of particles, a potential avenue for future work could be for bidisperse or polydisperse particles. The work in this thesis readily generalises to this, as we already assume non-interaction of particles in the flow. To examine polydisperse flows, multiple simulations of particles with varied Stokes number St to represent the varied particle size could be used. The phenomenon could then be represented by taking advantage of the non-interaction of particles and superimposing each of these simulations into a single domain. The case of interacting particles is more complex, and would be significantly different to the work in this thesis, and as such is not considered.

Appendix A

Additional material for Chapter 5

A.1 Background material for SQG & SQG⁺¹

The material in the following sections is all standard material, and largely follows the work in Vallis (2017). A common starting place for such an outline is the Boussinesq equations. These are the equations of motion for a fluid with small density variations throughout, which is appropriate in the ocean. They are sometimes referred to as buoyancy driven flows as it ignores density differences except when they are associated with gravity. The Boussinesq momentum equation, together with mass conservation and the equation for buoyancy are often used as the default equations of motion for ocean general circulation models (Vallis, 2017).

A.1.1 Boussinesq equations and scaling

The 3D Boussinesq equations are given by

$$\frac{D\mathbf{u}}{Dt} + \mathbf{f} \times \mathbf{u} = -\nabla\phi + b\hat{\mathbf{k}}, \quad (\text{A.1a})$$

$$\nabla \cdot \mathbf{u} = 0, \quad (\text{A.1b})$$

$$\frac{Db}{Dt} = 0. \quad (\text{A.1c})$$

where $\mathbf{u} = (u, v, w)$, $\phi = p/\rho_0$ is the pressure potential scaled by the density, b is the buoyancy and $f = f_0\hat{\mathbf{k}}$ is the f-plane approximation for the Coriolis parameter.

Before examining their dimensions, it is helpful to split the momentum equation

(A.1a) into its horizontal (x, y) and vertical (z) constituents,

$$\frac{D\mathbf{v}}{Dt} + \mathbf{f} \times \mathbf{v} = -\nabla_H \phi, \quad (\text{A.2a})$$

$$\frac{Dw}{Dt} + \frac{\partial \phi}{\partial z} = b, \quad (\text{A.2b})$$

where $\mathbf{v} = (u, v)$ and $\nabla_H = (\partial_x, \partial_y)$. As we wish to examine motion over a particular scale, we first need to non-dimensionalise the Boussinesq equations (A.1). From these dimensionless equations we can then employ asymptotics to extract the important terms for a chosen scale. The most straightforward scales to identify are a typical horizontal length L and velocity scale U for the motion. We also identify a scale for the mean depth H . Choosing a scale like this is equivalent to saying $\mathbf{v} = \mathcal{O}(U)$, for example. Using these we can create non-dimensional variables

$$(x, y) = L(\hat{x}, \hat{y}), \quad z = H\hat{z}, \quad \mathbf{v} = U\hat{\mathbf{v}}, \quad (\text{A.3})$$

where the non-dimensional variables are denoted with a hat. By definition non-dimensional variables are $\mathcal{O}(1)$. Using the length and velocity scales, we can create a characteristic time scale for advection as $T = L/U$ and we choose a scale f to ensure the non-dimensional Coriolis parameter is $\mathcal{O}(1)$,

$$t = \frac{L}{U}\hat{t}, \quad \mathbf{f} = f\hat{\mathbf{f}}. \quad (\text{A.4})$$

As the Coriolis parameter is assumed constant here, scaling it in this way means $\hat{\mathbf{f}} = (0, 0, \hat{f}) = \hat{\mathbf{k}}$ is simply the unit vector in the vertical direction. We will continue to use the variable \hat{f} in this section until the final scaling to indicate the presence of the Coriolis parameter in the general scalings of other variables.

Using mass conservation (A.1b) and the scalings in (A.3), the vertical velocity scales as

$$\begin{aligned} \frac{\partial w}{\partial z} &= -\left(\frac{\partial u}{\partial x} + \frac{\partial v}{\partial y}\right), \\ w &= U\left(\frac{H}{L}\right)\hat{w}, \end{aligned} \quad (\text{A.5})$$

where we identify H/L as the aspect ratio of the model.

Next, to examine stratification, we assume that the variation in buoyancy in the

vertical direction caused by (constant) stratification would differ from the variation in the horizontal directions. As such we write the buoyancy as

$$b = \bar{b}(z) + b'(x, y, z, t), \quad (\text{A.6})$$

and then the thermodynamic equation (A.1c) becomes

$$\frac{Db'}{Dt} + N^2 w = 0, \quad (\text{A.7})$$

where $N^2 = \partial_z \bar{b}$. The same idea is applied to the pressure potential $\phi = \bar{\phi}(z) + \phi'(x, y, z, t)$, where $\bar{\phi}$ is in hydrostatic balance with \bar{b} by construction. With this, the momentum equations become

$$\frac{D\mathbf{v}}{Dt} + \mathbf{f} \times \mathbf{v} = -\nabla_H \phi' \quad (\text{A.8a})$$

$$\frac{Dw}{Dt} + \frac{\partial \phi'}{\partial z} = b'. \quad (\text{A.8b})$$

Using the new scaled variables the horizontal momentum equation scales as

$$\frac{\partial \mathbf{v}}{\partial t} + \mathbf{v} \cdot \nabla \mathbf{v} + \mathbf{f} \times \mathbf{v} = -\nabla_H \phi. \quad (\text{A.9})$$

$$\frac{U^2}{L} \quad \frac{U^2}{L} \quad fU \quad \frac{\Phi}{L}$$

Here we identify an important scaling parameter, the Rossby number Ro , which is the ratio of the advective terms ($D\mathbf{v}/Dt$) to the rotational term ($\mathbf{f} \times \mathbf{v}$)

$$\frac{U^2}{fL} = \frac{U}{fL} = \text{Ro}. \quad (\text{A.10})$$

The Rossby number is a measure of the importance of rotation in the flow, with smaller values meaning the effect of rotation is important to consider. This value is typically small for large scale flows in the ocean, with typical values $\mathcal{O}(10^{-2})$ (Olbers et al., 2012) and as such we will treat Ro as small through this chapter.

As the advective terms are smaller than the Coriolis term by a factor of Ro , in order to avoid the trivial state of $f\mathbf{v} = 0$, the Coriolis and pressure terms must balance

$$|\mathbf{f} \times \mathbf{v}| \sim |\nabla_H \phi|, \quad (\text{A.11})$$

and so the pressure must scale as

$$\phi' = fUL\hat{\phi}. \quad (\text{A.12})$$

With this the vertical momentum has dimensions

$$\frac{Dw}{Dt} + \frac{\partial\phi'}{\partial z} = -b'. \quad (\text{A.13})$$

$$\frac{U^2 H}{L^2} \quad \frac{fUL}{H} \quad B$$

In a similar way to the horizontal momentum equation, the advective term is a factor of $\text{Ro}(H/L)$ smaller than the pressure gradient. This is the condition for hydrostatic balance in the flow. As $\text{Ro} = \mathcal{O}(10^{-2})$ for the typical large-scale flows we examine, hydrostatic balance holds and the buoyancy has dimension

$$b' = \frac{fUL}{H}\hat{b}. \quad (\text{A.14})$$

Using this scaling for the buoyancy and the altered thermodynamic equation (A.7) we identify the second scaling parameter

$$\frac{\frac{\partial b'}{\partial z}}{N^2} \sim \frac{fUL}{H^2 N^2} = \text{Ro} \frac{L^2}{L_D^2}, \quad (\text{A.15})$$

where $L_D = NH/f$ is the radius of deformation. This is related to the Rossby number and is a measure for the length scale that the effect of rotation becomes as important as the effects due to buoyancy in the flow. We will assume that $L/L_D = \mathcal{O}(1)$, meaning the scale of motion is not significantly different to the deformation radius. This also gives information on the stratification in the flow, which for $L/L_D = \mathcal{O}(1)$ implies the variation in stratification is small relative to the background stratification.

From here we can move to deriving the QG and subsequently the SQG⁺¹ equations we use for this chapter, with the following summarised non-dimensional variables and scalings

$$(\hat{x}, \hat{y}) = L^{-1}(x, y), \quad \hat{z} = H^{-1}z, \quad \hat{\mathbf{v}} = U^{-1}\mathbf{v}, \quad \hat{w} = \frac{L}{UH}w, \quad (\text{A.16})$$

$$\hat{t} = \frac{U}{L}t, \quad \hat{\mathbf{f}} = f^{-1}\mathbf{f} = \hat{\mathbf{k}}, \quad \hat{\phi} = \frac{1}{fUL}\phi', \quad \hat{b} = \frac{H}{fUL}b'. \quad (\text{A.17})$$

where we specifically ensure $\text{Ro} = U/fL \ll 1$ and $L/L_D = \mathcal{O}(1)$.

A.1.2 Quasi-Geostrophic (QG) equations

Now that the choice of scaling outlined above means we are examining motion for flows in geostrophic balance, we move to reduce the complexity of the full Boussinesq equations which has 5 variables to solve for, to a model that will be in a single variable. What follows in the next sections is standard material that largely follows from material in Vallis (2017).

With the above scalings, the non-dimensional Boussinesq equations are

$$\text{Ro} \frac{D\hat{\mathbf{v}}}{Dt} + \hat{\mathbf{k}} \times \hat{\mathbf{v}} = -\nabla_H \hat{\phi}, \quad (\text{A.18a})$$

$$\text{Ro} \frac{D\hat{w}}{Dt} + \frac{\partial \hat{\phi}}{\partial z} = \hat{b}, \quad (\text{A.18b})$$

$$\nabla \cdot \hat{\mathbf{u}} = 0, \quad (\text{A.18c})$$

$$\text{Ro} \frac{D\hat{b}}{Dt} + \left(\frac{L_D}{L}\right)^2 \hat{w} = 0, \quad (\text{A.18d})$$

where we have separated the momentum equation into its horizontal (A.18a) and vertical (A.18b) components from the start. We drop the hats on non-dimensional variables from here unless it is explicitly said otherwise. We expand \mathbf{u} in the usual way with the small parameter Ro

$$\mathbf{u} = \mathbf{u}_0 + \text{Ro} \mathbf{u}_1 + \text{Ro}^2 \mathbf{u}_2 + \mathcal{O}(\text{Ro}^3) \quad (\text{A.19})$$

and in a similar way for ϕ and b . The expansions are substituted into (A.18). At leading order the horizontal part of the momentum equation reads

$$\hat{\mathbf{k}} \times \mathbf{u}_0 = -\nabla_H \phi_0, \quad (\text{A.20})$$

stating that the system is in geostrophic balance at this order. Note that as $\hat{\mathbf{k}} \times \mathbf{u}_0 = (-v_0, u_0)$ and $\nabla_H \times (-v_0, u_0) = \nabla_H \cdot \mathbf{v}_0$ we can use (A.20) to see that

$$\nabla_H \cdot \mathbf{v}_0 = -\nabla_H \times (\nabla_H \phi_0) = 0, \quad (\text{A.21})$$

and conclude that the leading order horizontal flow is non-divergent. The vertical com-

ponent of the momentum equation at this order simply recovers hydrostatic equilibrium

$$\frac{\partial \phi_0}{\partial z} = b_0. \quad (\text{A.22})$$

Combining (A.21) with mass continuity (A.18c) we also obtain

$$\frac{\partial w_0}{\partial z} = 0, \quad (\text{A.23})$$

meaning that the leading order vertical velocity is constant. Furthermore as the vertical velocity is zero on the upper (or lower) boundary, it is zero everywhere. The first order $\mathcal{O}(\text{Ro})$ momentum equation reads

$$\frac{\partial \mathbf{u}_0}{\partial t} + (\mathbf{u}_0 \cdot \nabla) \mathbf{u}_0 + \hat{\mathbf{k}} \times \mathbf{u}_1 = -\nabla \phi_1 + b_1 \hat{\mathbf{k}}. \quad (\text{A.24})$$

Taking the curl of this gives the vorticity equation

$$\frac{\partial \boldsymbol{\omega}_0}{\partial t} + (\mathbf{u}_0 \cdot \nabla) \boldsymbol{\omega}_0 + \nabla \times (\hat{\mathbf{k}} \times \mathbf{u}_1) = -\nabla^\perp b_1, \quad (\text{A.25})$$

where $\boldsymbol{\omega}_0 = \nabla \times \mathbf{u}_0$ and $\nabla^\perp = \hat{\mathbf{k}} \times \nabla$. We are only interested in the vertical component of the vorticity equation which reads

$$\frac{\partial \zeta_0}{\partial t} + (\mathbf{u}_0 \cdot \nabla) \zeta_0 = -\nabla_H \cdot \mathbf{v}_1, \quad (\text{A.26})$$

where $\zeta_0 = \boldsymbol{\omega}_0 \cdot \hat{\mathbf{k}} = \partial_x v_0 - \partial_y u_0$. Combining this with mass conservation again gives

$$\frac{D_0}{Dt} \zeta_0 = \frac{\partial}{\partial z} w_1, \quad (\text{A.27})$$

where we use $D_0/Dt = \partial_t + \mathbf{u}_0 \cdot \nabla$. At leading order (A.18d) is identically zero because $w_0 = 0$. At the next order it reads

$$\frac{D_0 b_0}{Dt} + \left(\frac{L_D}{L} \right)^2 w_1 = 0. \quad (\text{A.28})$$

Eliminating w_1 between (A.27) and (A.28) leads to

$$\frac{D_0}{Dt} \zeta_0 = -F \frac{\partial}{\partial z} \left[\frac{D_0 b_0}{Dt} \right], \quad (\text{A.29})$$

where we define $F = (L/L_D)^2$. Rearranging yields

$$\frac{D_0}{Dt} \zeta_0 = -F \left[\frac{\partial}{\partial t} \left(\frac{\partial b_0}{\partial z} \right) + \left(\frac{\partial \mathbf{u}_0}{\partial z} \cdot \nabla \right) b_0 \right]. \quad (\text{A.30})$$

Combining the vertical derivative of geostrophic balance (A.20), and hydrostatic balance (A.22), we can eliminate the final term using

$$\hat{\mathbf{k}} \times \frac{\partial \mathbf{u}_0}{\partial z} = -\nabla_H b_0, \quad (\text{A.31})$$

which is referred to as the thermal wind equation. With this equation (A.30) reads

$$\frac{D_0}{Dt} \left[\zeta_0 + F \frac{\partial}{\partial z} (b) \right] = 0. \quad (\text{A.32})$$

Finally we can define a streamfunction for the horizontal velocity \mathbf{v}_0 due it being divergence free as found in (A.21). Geostrophic balance provides a natural way to define a streamfunction since

$$\mathbf{v}_0 = \hat{\mathbf{k}} \times \nabla \phi_0,$$

and so the streamfunction is

$$\psi = \phi_0. \quad (\text{A.33})$$

We use this to define the horizontal velocity \mathbf{v}_0 , vorticity ζ_0 and buoyancy b in terms of ψ

$$\mathbf{v}_0 = \nabla^\perp \psi, \quad \zeta_0 = \nabla_H^2 \psi, \quad b_0 = \frac{\partial \psi}{\partial z}. \quad (\text{A.34})$$

Using these, we write

$$\frac{D_0}{Dt} \left[\nabla_H^2 \psi + \frac{\partial}{\partial z} \left(F \frac{\partial \psi}{\partial z} \right) \right] = 0. \quad (\text{A.35})$$

Bringing back the dimensions and choosing to scale the vertical axis such that

$$\hat{z} = \left(\frac{N}{f} \right)^2 z, \quad (\text{A.36})$$

yields the equations

$$\partial_t q + J(\psi, q) = 0, \quad (\text{A.37a})$$

$$q = \nabla^2 \psi, \quad (\text{A.37b})$$

where $J(A, B) = \partial_x A \partial_y B - \partial_y A \partial_x B$. The quantity q is the *quasi-geostrophic potential vorticity* (QGPV). These equations are the governing equations for 3D quasi-geostrophic flow and, up to boundary conditions explored below, control the dynamics of 3D QG. In these equations, q is the dynamical variable, and it is through this that we deduce ψ through the Poisson equation.

The QGPV q , just introduced is related to the Ertel, or exact, potential vorticity (PV), Q which, for a hydrostatic fluid, is given by

$$Q = (\mathbf{f} + \nabla \times \mathbf{v}) \cdot \nabla b \quad (\text{A.38})$$

$$= -\partial_z v \partial_x b + \partial_z u \partial_y b + (f + \zeta) \partial_z b \quad (\text{A.39})$$

where only the horizontal velocity \mathbf{v} appears due to hydrostatic balance. The potential vorticity Q is also a materially conserved quantity. The above derivation of the QG equations from the Boussinesq equations (A.18) can also be done using the exact PV (A.38) as a starting point. We briefly highlight this to show that the quasi-geostrophic PV is an approximation to (A.38) in the QG limit.

By expanding the buoyancy as a reference state in the vertical and a perturbation $b = \tilde{b}(z) + b'(x, y, z, t)$, this becomes

$$Q = -\partial_z v \partial_x b' + \partial_z u \partial_y b' + (f + \zeta) \partial_z \tilde{b} + (f + \zeta) \partial_z b' \quad (\text{A.40})$$

$$= fN^2 + \zeta N^2 + f \partial_z b' + \zeta \partial_z b' - \partial_z v \partial_x b' + \partial_z u \partial_y b'. \quad (\text{A.41})$$

Dividing by N^2 and grouping by relative size yields

$$Q = f + \left(\zeta + \frac{f}{N^2} \partial_z b'\right) + \frac{1}{N^2} (\zeta \partial_z b' - \partial_z v \partial_x b' + \partial_z u \partial_y b') \quad (\text{A.42})$$

$$= \tilde{Q} + q_0 + q_1, \quad (\text{A.43})$$

where the terms, decrease in relative size. We identify $q' = q_0 + q_1$ as the disturbance PV around a constant reference state \tilde{Q} . The final term q_1 is small so we ignore it for

now also, and it is only included here for completeness. The term q_0 is the same, up to scaling, as the quasi-geostrophic PV identified in equation (A.35) and thus we treat it as the leading order approximation to the exact disturbance PV q' . This equivalence is dependent on the prior assumption of constant $N = \partial_z \tilde{b}(z)$ made in section A.1.1. Without this, the presence of non-uniform stratification complicates the relationship (Wagner and Young, 2015). From here onward we refer to q' simply as the potential vorticity rather than the quasi-geostrophic potential vorticity.

Returning to equation (A.37), to solve this PDE, we need boundary conditions for ψ . Due to the aspect ratio of the ocean, we consider only the boundaries at the surface ($z = 0, -H$) and treat the domain as horizontally periodic. Using the thermodynamic equation (A.18d) and evaluating at the ocean surface ($z = 0$) and ocean floor ($z = -H$), where the vertical velocity is zero, we have that

$$\left[\text{Ro} \frac{D\hat{b}}{Dt} + F\hat{w} \right]_{z=\{0, -H\}} = \frac{D\hat{b}}{Dt} \Big|_{z=\{0, -H\}} = 0. \quad (\text{A.44})$$

The same is done for the hydrostatic equation (A.18b) also

$$\left[\text{Ro} \frac{D\hat{w}}{Dt} + \frac{\partial \hat{\phi}}{\partial z} \right]_{z=\{0, -H\}} = \frac{\partial \hat{\phi}}{\partial z} \Big|_{z=\{0, -H\}} = \hat{b}. \quad (\text{A.45})$$

Combining this with (A.37) gives the full set of quasi-geostrophic equations

$$\partial_t q + J(\psi, q) = 0, \quad (\text{A.46a})$$

$$q = \nabla^2 \psi, \quad (\text{A.46b})$$

$$\partial_t \theta + J(\psi, \theta) = 0 \quad \text{at } z = 0, -H, \quad (\text{A.46c})$$

$$\partial_z \psi = \theta \quad \text{at } z = 0, -H, \quad (\text{A.46d})$$

where $\theta = b/f$ is a scaled version of the buoyancy, that can also be interpreted as temperature. We use the scaling $f = 1$ in this work and so these coincide exactly, but in general this interpretation holds for $f \neq 1$. The QG equations describe the evolution of the conserved scalar q , the quasi-geostrophic potential vorticity, as it is advected by the 2D horizontal velocity $\mathbf{v} = \nabla^\perp \psi$ in the interior, and the buoyancy θ on the top and bottom surfaces. From the potential vorticity q and surface buoyancies, we can infer ψ by solving a Poisson equation. Other dynamical variables of interest, such as the velocity, vorticity or pressure are recovered through the relations in equation (A.34).

A.1.3 Surface Quasi-Geostrophic (SQG) equations

We now examine a special case of the QG equations, by setting the potential vorticity to be identically zero. The interior equation is trivially satisfied when $q \equiv 0$ and the flow is entirely governed by the coupled system of ψ and θ on the boundaries. By also assuming the system is deep this can be further reduced to only depend on the surface boundary θ . The resulting system is known as surface quasi-geostrophic (SQG) flow (Blumen, 1978; Held et al., 1995):

$$\partial_t \theta^s + J(\psi^s, \theta^s) = 0, \quad (\text{A.47a})$$

$$\Delta \psi = 0, \quad (\text{A.47b})$$

$$\partial_z \psi|_{z=0} = \theta^s, \quad (\text{A.47c})$$

$$\psi \rightarrow 0 \text{ as } z \rightarrow -\infty, \quad (\text{A.47d})$$

where a superscript θ^s is used to denote the quantity on the surface ($z = 0$), and $\Delta \psi = \nabla^2 \psi$. This system can be readily solved as a joint system for ψ and θ at the surface. Omitting any time dependence, we take the two-dimensional Fourier transform of the streamfunction. The Laplacian then becomes an ODE in the vertical coordinate,

$$\Delta \hat{\psi}(x, y, z) = (-k^2 - l^2 + \partial_{zz}) \hat{\psi}(k, l, z) = 0. \quad (\text{A.48})$$

Grouping the wave numbers gives

$$\partial_{zz} \hat{\psi}(k, l, z) = \kappa^2 \hat{\psi}(k, l, z), \quad (\text{A.49})$$

where $\kappa^2 = k^2 + l^2$. The general solution for $\hat{\psi}$ is

$$\hat{\psi}(k, l, z) = Ae^{\kappa z} + Be^{-\kappa z}, \quad (\text{A.50})$$

which is straightforward to solve given the boundary conditions. The Fourier transform of the surface boundary condition relating ψ and θ gives us the first constant

$$\begin{aligned} \partial_z \hat{\psi}|_{z=0} &= \hat{\theta}^s \\ A &= \frac{\hat{\theta}^s(k, l)}{\kappa}, \end{aligned}$$

where $\kappa > 0$ is a positive variable. Noting that the streamfunction decays with depth gives $B = 0$. With this we have the relation between the 3D streamfunction and buoyancy on the surface in Fourier space

$$\hat{\psi}(k, l, z) = \frac{\hat{\theta}^s(k, l)}{\kappa} e^{\kappa z}. \quad (\text{A.51})$$

We are only interested in the surface dynamics ($z = 0$) where this reduces to

$$\hat{\psi}(k, l, z = 0) = \frac{\hat{\theta}^s(k, l)}{\kappa}. \quad (\text{A.52})$$

Moving back to real space the κ in the denominator defines a fractional Laplacian of the streamfunction:

$$\Delta^{\frac{1}{2}} \psi^s(x, y) = \theta^s(x, y). \quad (\text{A.53})$$

With this we have reduced the full 3D dynamics to a flow that is entirely driven by the 2D distribution of buoyancy on the surface:

$$\partial_t \theta^s + J(\psi^s, \theta^s) = 0, \quad (\text{A.54a})$$

$$\Delta^{\frac{1}{2}} \psi^s = \theta^s. \quad (\text{A.54b})$$

These two equations define the SQG system. It is noteworthy that this describes a three dimensional system by computation carried out in only two dimensions. The problem in general is still three dimensional even though the advection at the surface is two dimensional, as ψ which defines the velocity is three dimensional. The reduction in complexity allows efficient computation of a flow that models the ocean surface well, without needing to compute fields in three dimensions only to examine a single two dimensional plane of the whole field. A review of SQG with details of the applicability to modelling ocean dynamics, alongside some properties of SQG as a model can be found in [Lapeyre \(2017\)](#).

A.1.4 Numerical implementation of SQG flow

The numerical implementation of SQG has two stages, namely, *inversion* and *advection*. Inversion consists of using the current value of θ^s to recover ψ^s with (A.52), and in turn the horizontal velocity field \mathbf{v}^s . θ^s is then advected using this horizontal velocity

field. A dissipation is added in order to control a build up of energy at sub-grid scales which would cause numerical instability. This is implemented as a hyperdiffusion term of the form

$$\nu \nabla_H^8 \theta^s = \nu (\nabla_H^2)^4 \theta^s = \nu (\partial_{xx} + \partial_{yy})^4 \theta^s = \nu \kappa^4 \theta^s, \quad (\text{A.55})$$

where ν is used to control the intensity of the dissipation. Taking the 2D Fourier transform of the theta advection in (A.54a) and combining it with the hyperdiffusion gives the governing equation in wavenumber space that is solved numerically

$$\partial_t \hat{\theta}^s = -\mathcal{F} [u^s \partial_x \theta^s + v^s \partial_y \theta^s] - \nu \kappa^4 \hat{\theta}^s, \quad (\text{A.56})$$

where \mathcal{F} and \mathcal{F}^{-1} denotes a Fourier transform and its corresponding inverse transform respectively.

An example of an elliptical vortex evolved in the SQG flow as seen in [Held et al. \(1995\)](#) is used to demonstrate the method. The initial condition for θ^s is

$$\theta^s = e^{-(x^2+4y^2)}. \quad (\text{A.57})$$

The initial ellipse is evolved in a $L = 2\pi$ periodic box. Dissipation is kept low with $\nu = 10^{-18}$ and acts only to avoid instability at small scales. The time step is set at $\Delta t = 10^{-2}$. Snapshots of the initial ellipse and subsequent times are presented in figure [A.1](#). A main feature of SQG in comparison to regular 2D turbulent flow is the greater presence of filamentary instabilities. These are seen forming in the outer filaments in the final panel of figure [A.1](#). This is due to the weaker influence of vortices on the velocity surrounding them.

The simulation in figure [A.1](#) is undertaken with a spatial grid resolution of $(512)^2$ with periodic boundary conditions. The nonlinear advection is solved pseudo-spectrally using a fourth order Runge-Kutta scheme, and the hyperdiffusion is solved exactly in wavenumber space. Initial conditions are either specified for the streamfunction in wavenumber space, $\hat{\psi}$, or for an initial physical space θ^s distribution. With either one of these, the other can be initialised through the inversion relation (A.52). The code used to create the QG flow seen here and the subsequent models in this chapter were adapted from an open-source python package [PyQG \(PyQG, 2016\)](#)

An example of forced turbulence is also presented later in section [5.2](#) where it is




Figure A.1: Evolution of an elliptical vortex: the surface buoyancy θ^s is shown at times $t = 0$ (top left), $t = 5$ (top right), $t = 20$ (bottom left), $t = 30$ (bottom right).

used as a comparison with a variation of SQG introduced in the following section.

A.1.5 Next order in QG and SQG

While using the special case of SQG is advantageous in reducing the need to calculate the full velocity field in order to examine the surface dynamics only, it still lacks the critical component of divergence that is needed for clustering to occur. Both QG and SQG flow are incompressible by definition and will not exhibit clustering behaviour. QG, and hence SQG, are the leading order result of the asymptotic expansions in Rossby number in section A.1.2. To include divergence in the flow, we need to look at an order above this. The next-order correction to QG theory was first introduced in Muraki et al. (1999). We will use it as our simple model for compressible surface currents in the ocean. We briefly outline the derivation of this model.

The original (leading order in $\epsilon \equiv \text{Ro}$) QG equations describe the relation between the velocities and temperature, (u, v, w, θ) starting with the Boussinesq equations (A.18). The next order corrections are more easily expressed through a recast set of variables (Φ, F, G) . These new variables are motivated by rearranging the horizontal components of the non-dimensional Boussinesq equations for momentum (A.18a) and hydrostatic balance (A.18b) as

$$v = \partial_x \phi + \epsilon \frac{Du}{Dt}, \quad (\text{A.58a})$$

$$-u = \partial_y \phi + \epsilon \frac{Dv}{Dt}, \quad (\text{A.58b})$$

$$\theta = \partial_z \phi. \quad (\text{A.58c})$$

Manipulating these into this form emphasises the gradient nature of the equations at leading order. The next order terms in (A.58) are not guaranteed to be the gradient of a potential. Any part of the flow that is not a gradient introduces divergence. It is natural to use a 3D Helmholtz decomposition to represent the original variables in a way that isolates the gradient and divergent parts. This is achieved by the representation

$$\begin{pmatrix} v \\ -u \\ \theta \end{pmatrix} = \nabla \Phi + \nabla \times \begin{pmatrix} F \\ G \\ 0 \end{pmatrix} = \begin{pmatrix} \partial_x \Phi - \partial_z G \\ \partial_y \Phi + \partial_z F \\ \partial_z \Phi + \partial_x G - \partial_y F \end{pmatrix}. \quad (\text{A.59})$$

To ensure consistency with QG in the limit $\epsilon \rightarrow 0$, we relate the potential Φ to the

pressure at leading order in QG, i.e. $\Phi = \phi_0$ at leading order. In a similar way, both F and G must be of order ϵ to ensure consistency. Using the continuity equation we can also find an expression for the vertical velocity

$$w = \partial_x F + \partial_y G. \quad (\text{A.60})$$

The vertical velocity is the same order as F and G which aligns with the assumption that movement in the vertical is small ($\mathcal{O}(\epsilon)$) relative to the horizontal motion. From here on we will define $w \mapsto \epsilon w$ to reflect the small nature of w , unless explicitly stated otherwise. Bringing this together with (A.59) gives a set of relations between the original set of variables in the Boussinesq equations (u, v, w, θ) and the new potential variables (Φ, F, G) ,

$$\begin{aligned} u &= -\partial_y \Phi - \partial_z F, & \theta &= \partial_z \Phi + \partial_x G - \partial_y F, \\ v &= \partial_x \Phi - \partial_z G, & \epsilon w &= \partial_x F + \partial_y G. \end{aligned} \quad (\text{A.61})$$

The Boussinesq equations can be restated in terms of these new variables in a structure that resembles the QG equations but without any approximation. Starting with the two new terms with no counterpart in the original QG version, F and G . Using (A.61) and the non-dimensional thermodynamic equation,

$$\frac{D\theta}{Dt} + w = 0, \quad (\text{A.62})$$

we compare the curl of both equations (A.58) and (A.59) to obtain a Poisson equation for both

$$\nabla^2 F = \epsilon \left[-\partial_x \left(\frac{D\theta}{Dt} \right) + \partial_z \left(\frac{Dv}{Dt} \right) \right], \quad (\text{A.63})$$

$$\nabla^2 G = \epsilon \left[-\partial_y \left(\frac{D\theta}{Dt} \right) - \partial_z \left(\frac{Du}{Dt} \right) \right], \quad (\text{A.64})$$

with a boundary condition of $(\partial_x F + \partial_y G)^s = 0$. The equation for the potential Φ is similar in nature to that of ψ in equation (A.46b). It relates the leading order gradient potential behaviour to the PV, but now takes the exact form of the PV in (A.38) which is conserved exactly:

$$\nabla^2 \Phi = q' - \epsilon \nabla \cdot [\theta(\nabla \times \mathbf{v})] = \partial_x v - \partial_y u + \partial_z \theta, \quad (\text{A.65})$$

with a boundary condition of $(\partial_z \Phi + \partial_x G - \partial_y F)^s = \theta^s$ and

$$\epsilon \nabla \cdot [\theta(\nabla \times \mathbf{v})] = \epsilon(\nabla \times \mathbf{v}) \cdot \theta. \quad (\text{A.66})$$

As a consistency check, we can see that setting $\epsilon = 0$ results in the same form as the original QG equations.

From here we expand the potentials in Rossby number as follows

$$\Phi(\mathbf{x}, t) = \Phi_0 + \epsilon \Phi_1 + \epsilon^2 \Phi_2 + \mathcal{O}(\epsilon^3), \quad (\text{A.67a})$$

$$F(\mathbf{x}, t) = \epsilon F_1 + \epsilon^2 F_2 + \mathcal{O}(\epsilon^3), \quad (\text{A.67b})$$

$$G(\mathbf{x}, t) = \epsilon G_1 + \epsilon^2 G_2 + \mathcal{O}(\epsilon^3), \quad (\text{A.67c})$$

where both F and G are zero at leading order to match the leading order approximation in QG. The variables q and θ^s are not explicitly expanded as power series. The inversion of potential vorticity q in (A.65) is non-trivial when expanded to a power-series. In short, an expansion is considered for $q = q_0 + \epsilon q_1 + \dots$ but it is computed separately at each order from Φ , at each discrete time. The unexpanded version of q is advected at each time step, and not separately by order like the other variables. We will not concern ourselves in this thesis with the exact details though, which can be found in (Muraki et al., 1999), as they are complex and not relevant as we will be interested only in the SQG case of $q \equiv 0$, that we will see in the next section.

Using these expanded variables in the reformulated Boussinesq equations, we are interested only in the first order corrections so we discard any $\mathcal{O}(\epsilon^2)$ terms and are left with the equations listed below.

We begin with the potentials

$$\nabla^2 \Phi_0 = q_0, \quad (\text{A.68a})$$

$$\nabla^2 \Phi_1 = q_1 - [(\nabla^2 \Phi_0) \partial_{zz} \Phi_0 - |\nabla(\partial_z \Phi_0)|^2], \quad (\text{A.68b})$$

$$\nabla^2 F_1 = 2J(\partial_z \Phi_0, \partial_x \Phi_0), \quad (\text{A.68c})$$

$$\nabla^2 G_1 = 2J(\partial_z \Phi_0, \partial_y \Phi_0), \quad (\text{A.68d})$$

with boundary conditions

$$w^s = 0, \quad (\text{A.69a})$$

$$\partial_z \Phi_0^s = \theta_0^s, \quad (\text{A.69b})$$

$$(\partial_z \Phi_1 + \partial_x G_1 - \partial_y F_1)^s = \theta_1^s, \quad (\text{A.69c})$$

$$(\partial_y G_1 - \partial_x F_1)^s = 0, \quad (\text{A.69d})$$

where the final boundary condition applies to both (A.68c) and (A.68d). From (A.61) the velocity components and θ are given by

$$v = \partial_x \Phi_0 + \epsilon(\partial_x \Phi_1 - \partial_z G_1), \quad (\text{A.70a})$$

$$-u = \partial_y \Phi_0 + \epsilon(\partial_y \Phi_1 - \partial_z F_1), \quad (\text{A.70b})$$

$$\theta = \partial_z \Phi_0 + \epsilon(\partial_z \Phi_1 + \partial_x G_1 - \partial_y F_1), \quad (\text{A.70c})$$

$$\epsilon w = \epsilon(\partial_x F_1 - \partial_y G_1). \quad (\text{A.70d})$$

The full velocity fields are used to advect the full $q = q_0 + \epsilon q_1$ and θ^s :

$$\frac{D}{Dt} q = 0 \quad \text{and} \quad \frac{D}{Dt} \theta^s = 0, \quad (\text{A.71})$$

with

$$\frac{D}{Dt} = \partial_t + u\partial_x + v\partial_y + \epsilon w\partial_z. \quad (\text{A.72})$$

These equations form the full set of QG⁺¹ equations that extend the QG equations to an additional order of accuracy in ϵ . The structure of the equations closely resembles the structure of the QG equations in (A.46). They exactly reduce to the QG equations for $\epsilon = 0$, but now include the possibility for divergence introduced by the higher order corrections in the velocity field which are necessary for clustering to occur.

Next order SQG

As we are interested only in examining a 2D flow representing the ocean surface, we now extend the QG⁺¹ equations to SQG⁺¹ in a similar way as done with QG and SQG in section A.1.3. By setting $q = 0$ identically we can satisfy the interior equation in (A.71). This once again reduces the dynamics to advection of surface temperature θ^s

with $w = 0$ in (A.72). The adjusted equations for QG^{+1} to account for $q = 0$ are presented below, and further details can be found in Hakim et al. (2002) which we follow. Apart from the removal of the interior equation for q and the altered material derivative, the only change to the equations comes to the potentials and their boundary conditions. They now read:

$$\nabla^2 \Phi_0 = 0, \quad \partial_z \Phi_0^s = \theta^s - \bar{\theta}^s, \quad (\text{A.73a})$$

$$\nabla^2 \Phi_1 = |\nabla(\partial_z \Phi_0)|^2, \quad \partial_z \Phi_1^s \sim \bar{\theta}^s, \quad (\text{A.73b})$$

$$\nabla^2 F_1 = 2J(\partial_z \Phi_0, \partial_x \Phi_0), \quad F_1^s = 0, \quad (\text{A.73c})$$

$$\nabla^2 G_1 = 2J(\partial_z \Phi_0, \partial_y \Phi_0), \quad G_1^s = 0, \quad (\text{A.73d})$$

where $\bar{\theta}^s$ is the mean temperature across the surface. The solution to (A.73a) follows a similar process to the original SQG equations with ψ , resulting in a relation that resembles (A.51):

$$\hat{\Phi}_0(k, l, z) = \frac{\hat{\theta}^s(k, l)}{\kappa} e^{-\kappa z}. \quad (\text{A.74})$$

We are again interested only in the surface ($z = 0$) so this becomes

$$\hat{\Phi}_0^s(k, l) = \frac{\hat{\theta}^s(k, l)}{\kappa}. \quad (\text{A.75})$$

The remaining correction potentials can be reduced to Laplace inversions at the surface. As the inhomogeneous right hand side of each relation contains only derivatives of Φ_0 , which we know satisfies a Laplace problem from (A.73a), we can specify particular solutions of the form

$$\Phi_1 = \frac{1}{2} \partial_z \Phi_0 \partial_z \Phi_0 + \tilde{\Phi}_1, \quad (\text{A.76a})$$

$$F_1 = \partial_y \Phi_0 \partial_z \Phi_0 + \tilde{F}_1, \quad (\text{A.76b})$$

$$G_1 = -\partial_x \Phi_0 \partial_z \Phi_0 + \tilde{G}_1. \quad (\text{A.76c})$$

where a tilde $\tilde{\Phi}_1, \tilde{F}_1, \tilde{G}_1$ denotes a homogenous solution to the Laplace problem with boundary conditions that ensure the particular solutions (Φ_1, F_1, G_1) satisfy (A.73):

$$\nabla^2 \tilde{\Phi}_1 = 0, \quad \partial_z \tilde{\Phi}_1^s \sim \bar{\theta}^s - \partial_z \Phi_0 \partial_{zz} \Phi_0, \quad (\text{A.77a})$$

$$\nabla^2 \tilde{F}_1 = 0, \quad \tilde{F}_1^s = -\partial_y \Phi_0 \partial_z \Phi_0, \quad (\text{A.77b})$$

$$\nabla^2 \tilde{G}_1 = 0, \quad \tilde{G}_1^s = \partial_x \Phi_0 \partial_z \Phi_0. \quad (\text{A.77c})$$

This remarkable reduction to Laplace inversions greatly simplifies the calculations required for the surface velocities, especially when compared to the calculations needed for the full velocities in the QG^{+1} case.

Practically, these potentials are computed on the surface as

$$\partial_z F_1^s = \partial_z (\partial_y \Phi_0 \partial_z \Phi_0) - \mathcal{F}^{-1} \left[\kappa \mathcal{F} [\partial_y \Phi_0 \partial_z \Phi_0] \right], \quad (\text{A.78a})$$

$$\partial_z G_1^s = -\partial_z (\partial_x \Phi_0 \partial_z \Phi_0) - \mathcal{F}^{-1} \left[\kappa \mathcal{F} [\partial_x \Phi_0 \partial_z \Phi_0] \right], \quad (\text{A.78b})$$

$$\partial_x \Phi_1^s = \partial_{xz} \Phi_0 \partial_z \Phi_0 - \mathcal{F}^{-1} \left[\frac{ik}{\kappa} \mathcal{F} [\bar{\theta}^s - \partial_z \partial_x \Phi_0 \partial_{zz} \Phi_0] \right], \quad (\text{A.78c})$$

$$\partial_y \Phi_1^s = \partial_{yz} \Phi_0 \partial_z \Phi_0 - \mathcal{F}^{-1} \left[\frac{il}{\kappa} \mathcal{F} [\bar{\theta}^s - \partial_z \partial_y \Phi_0 \partial_{zz} \Phi_0] \right], \quad (\text{A.78d})$$

where \mathcal{F} and \mathcal{F}^{-1} are the Fourier and inverse Fourier transform respectively. With this, the entire system can be expressed in terms of the leading order potential Φ_0 and the surface temperature θ^s . Using these potentials the surface velocity components can be calculated via

$$v^s = \partial_x \Phi_0^s + \epsilon (\partial_x \Phi_1^s - \partial_z G_1^s), \quad (\text{A.79a})$$

$$-u^s = \partial_y \Phi_0^s + \epsilon (\partial_y \Phi_1^s - \partial_z F_1^s). \quad (\text{A.79b})$$

Bibliography

- Balkovsky, E., Falkovich, G., and Fouxon, A. (2001). Intermittent distribution of inertial particles in turbulent flows. *Phys. Rev. Lett.*, 86:2790–2793.
- Batchelor, G. K. (1959). Small-scale variation of convected quantities like temperature in turbulent fluid part 1. general discussion and the case of small conductivity. *Journal of Fluid Mechanics*, 5(1):113–133.
- Batchelor, G. K. (2000). *An Introduction to Fluid Dynamics*. Cambridge Mathematical Library. Cambridge University Press.
- Bec, J. (2003). Fractal clustering of inertial particles in random flows. *Physics of Fluids*, 15(11):L81–L84.
- Bec, J., Gustavsson, K., and Mehlig, B. (2024). Statistical models for the dynamics of heavy particles in turbulence. *Annual Review of Fluid Mechanics*, 56(Volume 56, 2024):189–213.
- Beron-Vera, F. J., Olascoaga, M. J., and Miron, P. (2019). Building a maxey–riley framework for surface ocean inertial particle dynamics. *Physics of Fluids*, 31(9):096602.
- Berti, S., Santos, F. A. D., Lacorata, G., and Vulpiani, A. (2011). Lagrangian drifter dispersion in the southwestern atlantic ocean. *Journal of Physical Oceanography*, 41(9):1659 – 1672.
- Blumen, W. (1978). Uniform potential vorticity flow: Part i. theory of wave interactions and two-dimensional turbulence. *Journal of Atmospheric Sciences*, 35(5):774–783.
- Boffetta, G., Lillo, F., and Gamba, A. (2003). Large scale inhomogeneity of inertial particles in turbulent flow. *Physics of Fluids*, 16.

- Boyce, W. E. (2017). *Boyce's elementary differential equations and boundary value problems*. John Wiley & Sons, Inc., Hoboken, New Jersey, eleventh edition / william e. boyce, richard c. diprima, douglas b. meade. edition.
- Charney, J. G. (1971). Geostrophic turbulence. *Journal of Atmospheric Sciences*, 28(6):1087 – 1095.
- Cover, T. and Thomas, J. (1991). *Elements of Information Theory*. Wiley.
- Cuzzi, J. N., Hogan, R. C., Paque, J. M., and Dobrovolskis, A. R. (2001). Size-selective concentration of chondrules and other small particles in protoplanetary nebula turbulence. *The Astrophysical journal*, 546(1):496–508.
- Dembo, A. and Zeitouni, O. (2010). *Large deviations techniques and applications*. Stochastic Modelling and Applied Probability ; 38. Springer-Verlag, Berlin ;, second edition. edition.
- Ellis, R. S. (2005). *Entropy, large deviations, and statistical mechanics*. Classics in Mathematics. Springer, Berlin, Heidelberg, reprint of the 1985 edition edition.
- Evans, L. C. (2013 - 2013). *An introduction to stochastic differential equations*. American Mathematical Society, Providence, Rhode Island.
- Falkovich, G., Fouxon, A., and Stepanov, M. G. (2002). Acceleration of rain initiation by cloud turbulence. *Nature*, 419:151–154.
- Falkovich, G., Gawedzki, K., and Vergassola, M. (2001). Particles and fields in fluid turbulence. *Reviews of Modern Physics*, 73(4):913–975.
- Ferry, J. and Balachandar, S. (2001). A fast eulerian method for disperse two-phase flow. *International Journal of Multiphase Flow*, 27(7):1199–1226.
- Gardiner, C. W. (2009). *Stochastic methods : a handbook for the natural and social sciences*. Springer series in synergetics. Springer, Berlin, 4th ed. edition.
- Goto, S. and Vassilicos, J. (2006). Self-similar clustering of inertial particles and zero-acceleration points in fully developed two-dimensional turbulence. *Physics of Fluids*, 18.
- Grassberger, P. (1983). Generalized dimensions of strange attractors. *Physics Letters A*, 97(6):227–230.

- Grassberger, P. (2002). Go with the winners: a general monte carlo strategy. *Computer Physics Communications*, 147(1):64–70. Proceedings of the Europhysics Conference on Computational Physics Computational Modeling and Simulation of Complex Systems.
- Grassberger, P. and Procaccia, I. (1983). Measuring the strangeness of strange attractors. *Physica D: Nonlinear Phenomena*, 9(1):189–208.
- Hakim, G. J., Snyder, C., and Muraki, D. J. (2002). A new surface model for cyclone–anticyclone asymmetry. *Journal of the Atmospheric Sciences*, 59(16):2405 – 2420.
- Haller, G. and Sapsis, T. (2008). Where do inertial particles go in fluid flows? *Physica D: Nonlinear Phenomena*, 237(5):573–583.
- Haynes, P. H. and Vanneste, J. (2005). What controls the decay of passive scalars in smooth flows? *Physics of Fluids*, 17(9):097103.
- Held, I. M., Pierrehumbert, R. T., Garner, S. T., and Swanson, K. L. (1995). Surface quasi-geostrophic dynamics. *Journal of Fluid Mechanics*, 282:1–20.
- Hessel, V., Löwe, H., and Schönfeld, F. D. (2005). Micromixers—a review on passive and active mixing principles. *Chemical Engineering Science*, 60:2479–2501.
- Holden, A. V., editor (1986). *Chaos*. Princeton University Press, Princeton.
- Hubbard, J., Haglund, J., and Ezekoye, O. (2009). Simulation of the evolution of particle size distributions containing coarse particulate in the atmospheric surface layer with a simple convection-diffusion-sedimentation model. *Atmospheric Environment*, 43(29):4435–4443.
- Kaplan, J. L. and Yorke, J. A. (1979). Chaotic behavior of multidimensional difference equations. In Peitgen, H.-O. and Walther, H.-O., editors, *Functional Differential Equations and Approximation of Fixed Points*, pages 204–227, Berlin, Heidelberg. Springer Berlin Heidelberg.
- Kraichnan, R. H. (1968). Small-scale structure of a scalar field convected by turbulence. *The Physics of Fluids*, 11(5):945–953.
- Kuo, K. K. and Acharya, R. (2012). *Applications of turbulent and multi-phase combustion*. Wiley, Hoboken, N.J.

- Károlyi, G., Péntek, A., Scheuring, I., Tél, T., and Toroczkai, Z. (2000). Chaotic flow: The physics of species coexistence. *Proceedings of the National Academy of Sciences - PNAS*, 97(25):13661–13665.
- LaCasce, J. H. and Mahadevan, A. (2006). Estimating horizontal and vertical velocities from sea surface temperature.
- Laffargue, T., Lam, K.-D. N. T., Kurchan, J., and Tailleur, J. (2013). Large deviations of lyapunov exponents. *Journal of Physics A: Mathematical and Theoretical*, 46(25):254002.
- Lapeyre, G. (2017). Surface quasi-geostrophy. *Fluids*, 2:7.
- Lapeyre, G. and Klein, P. (2006). Dynamics of the upper oceanic layers in terms of surface quasigeostrophy theory. *Journal of physical oceanography*, 36(2):165–176.
- Maalouly, M., Lapeyre, G., Cozian, B., Mompean, G., and Berti, S. (2023). Particle dispersion and clustering in surface ocean turbulence with ageostrophic dynamics. *Physics of Fluids*, 35:126601.
- Mandelbrot, B. B. (1977). *Fractals : form, chance & dimension*. W. H. Freeman, San Francisco.
- Maxey, M. R. and Riley, J. J. (1983). Equation of motion for a small rigid sphere in a nonuniform flow. *The Physics of Fluids*, 26(4):883–889.
- Michaelides, E. E. (1997). Review—The Transient Equation of Motion for Particles, Bubbles, and Droplets. *Journal of Fluids Engineering*, 119(2):233–247.
- Miron, P., Beron-Vera, F., Helfmann, L., and Koltai, P. (2021). Transition paths of marine debris and the stability of the garbage patches. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 31:033101.
- Monchaux, R., Bourgoïn, M., and Cartellier, A. (2012). Analyzing preferential concentration and clustering of inertial particles in turbulence. *International Journal of Multiphase Flow*, 40:1–18.
- Morrow, R., Fu, L.-L., Arduin, F., Benkiran, M., Chapron, B., Cosme, E., d’Ovidio, F., Farrar, J. T., Gille, S. T., Lapeyre, G., Le Traon, P.-Y., Pascual, A., Ponte, A., Qiu, B., Rasclé, N., Ubelmann, C., Wang, J., and Zaron, E. D. (2019). Global

- observations of fine-scale ocean surface topography with the surface water and ocean topography (swot) mission. *Frontiers in Marine Science*, Volume 6 - 2019.
- Motter, A. E., Lai, Y.-C., and Grebogi, C. (2003). Reactive dynamics of inertial particles in nonhyperbolic chaotic flows. *Physical review. E, Statistical physics, plasmas, fluids, and related interdisciplinary topics*, 68(5):563071–563075.
- Muraki, D. J., Snyder, C., and Rotunno, R. (1999). The next-order corrections to quasigeostrophic theory. *Journal of the Atmospheric Sciences*, 56(11):1547 – 1560.
- Olascoaga, M. J., Beron-Vera, F. J., Miron, P., Triñanes, J., Putman, N. F., Lumpkin, R., and Goni, G. J. (2020). Observation and quantification of inertial effects on the drift of floating objects at the ocean surface. *Physics of Fluids*, 32(2):026601.
- Olbers, D., Eden, C., and Willebrand, J. (2012). *Ocean dynamics*. Springer, Berlin ;.
- Ott, E. (2002). *Chaos in Dynamical Systems*. Cambridge University Press, 2 edition.
- Parker, T. S. and Chua, L. O. (1989). *Practical Numerical Algorithms for Chaotic Systems*. Springer.
- Pavliotis, G. A. (2014 - 2014). *Stochastic processes and applications : diffusion processes, the Fokker-Planck and Langevin equations*. Texts in applied mathematics, volume 60. Springer, New York.
- Post, S. L. and Abraham, J. (2002). Modeling the outcome of drop–drop collisions in diesel sprays. *International Journal of Multiphase Flow*, 28(6):997–1019.
- PyQG (2016). Pyqg github. <https://github.com/pyqg/pyqg>.
- Qiu, B., Chen, S., Klein, P., Ubelmann, C., Fu, L.-L., and Sasaki, H. (2016). Reconstructability of three-dimensional upper-ocean circulation from swot sea surface height measurements. *Journal of physical oceanography*, 46(3):947–963.
- Renaud, A. and Vanneste, J. (2020). Dispersion of inertial particles in cellular flows in the small-stokes, large-péclet regime. *Journal of Fluid Mechanics*, 903:A2.
- Rubin, J., Jones, C. K. R. T., and Maxey, M. (1995). Settling and asymptotic motion of aerosol particles in a cellular flow field. *Journal of Nonlinear Science*, 5(4):337–358.

- Sapsis, T. and Haller, G. (2010). Clustering criterion for inertial particles in two-dimensional time-periodic and three-dimensional steady flows. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 20(1):017515.
- Saraceno, M., Bodnariuk, N., Ruiz-Etcheverry, L. A., Berta, M., Simionato, C. G., Beron-Vera, F., and Olascoaga, M. (2024). Lagrangian characterization of the southwestern atlantic from a dense surface drifter deployment. *Deep Sea Research Part I: Oceanographic Research Papers*, 208:104319.
- Schuster, H. G. (2005). *Deterministic chaos: An introduction*. John Wiley & Sons, Ltd.
- Shaw, R. A. (2003). Particle-turbulence interactions in atmospheric clouds. *Annual Review of Fluid Mechanics*, 35(1):183–227.
- Sigurgeirsson, H. and Stuart, A. M. (2002). A model for preferential concentration. *Physics of Fluids*, 14(12):4352–4361.
- Sommerer, J. C. and Ott, E. (1993). Particles floating on a moving fluid: A dynamically comprehensible physical fractal. *Science*, 259(5093):335–339.
- Sundby, S. and Fossum, P. (1990). Feeding conditions of arcto-norwegian cod larvae compared with the rothschild–osborn theory on small-scale turbulence and plankton contact rates. *Journal of plankton research*, 12(6):1153–1162.
- Sutherland, B. R., DiBenedetto, M., Kaminski, A., and van den Bremer, T. (2023). Fluid dynamics challenges in predicting plastic pollution transport in the ocean: A perspective. *Phys. Rev. Fluids*, 8:070701.
- Tailleur, J. and Kurchan, J. (2007). Probing rare physical trajectories with lyapunov weighted dynamics. *Nature Physics*, 3(3):203–207.
- Toschi, F. and Bodenschatz, E. (2009). Lagrangian properties of particles in turbulence. *Annual Review of Fluid Mechanics*, 41(Volume 41, 2009):375–404.
- Touchette, H. (2009). The large deviation approach to statistical mechanics. *Physics Reports*, 478(1):1–69.
- Tsai, M.-Y., Elgethun, K., Ramaprasad, J., Yost, M. G., Felsot, A. S., Hebert, V. R., and Fenske, R. A. (2005). ‘the washington aerial spray drift study: Modeling pes-

- ticide spray drift deposition from an aerial application'. *Atmospheric Environment*, 39(33):6194–6203.
- Vallis, G. K. (2017). *Atmospheric and Oceanic Fluid Dynamics: Fundamentals and Large-Scale Circulation*. Cambridge University Press, 2 edition.
- van Sebille, E., Aliani, S., Law, K. L., Maximenko, N., Alsina, J. M., Bagaev, A., Bergmann, M., Chapron, B., Chubarenko, I., Cózar, A., Delandmeter, P., Egger, M., Fox-Kemper, B., Garaba, S. P., Goddijn-Murphy, L., Hardesty, B. D., Hoffman, M. J., Isobe, A., Jongedijk, C. E., Kaandorp, M. L. A., Khatmullina, L., Koelmans, A. A., Kukulka, T., Laufkötter, C., Lebreton, L., Lobelle, D., Maes, C., Martinez-Vicente, V., Morales Maqueda, M. A., Poulain-Zarcos, M., Rodríguez, E., Ryan, P. G., Shanks, A. L., Shim, W. J., Suaria, G., Thiel, M., van den Bremer, T. S., and Wichmann, D. (2020). The physical oceanography of the transport of floating marine debris. *Environmental Research Letters*, 15(2):023003.
- van Sebille, E., Wilcox, C., Lebreton, L., Maximenko, N., Hardesty, B. D., van Franeker, J. A., Eriksen, M., Siegel, D., Galgani, F., and Law, K. L. (2015). A global inventory of small floating plastic debris. *Environmental Research Letters*, 10(12):124006.
- Wagner, G. and Young, W. (2015). Available potential vorticity and wave-averaged quasi-geostrophic flow. *Journal of Fluid Mechanics*, 785:401–424.
- Welander, P. (1955). Studies on the General Development of Motion in a Two-Dimensional, Ideal Fluid. *Tellus*, 7(2):141–156.
- Wilkinson, M., Mehlig, B., and Uski, V. (2008). Stokes trapping and planet formation. *The Astrophysical journal. Supplement series*, 176(2):484–496.
- Wolf, A., Swift, J. B., Swinney, H. L., and Vastano, J. A. (1985). Determining lyapunov exponents from a time series. *Physica D: Nonlinear Phenomena*, 16(3):285–317.
- Wong, E. and Zakai, M. (1965a). On the convergence of ordinary integrals to stochastic integrals. *The Annals of Mathematical Statistics*, 36(5):1560–1564.
- Wong, E. and Zakai, M. (1965b). On the relation between ordinary and stochastic differential equations. *International Journal of Engineering Science*, 3(2):213–229.

Zhang, Z., Qiu, B., Klein, P., and Travis, S. (2019). The influence of geostrophic strain on oceanic ageostrophic motion and surface chlorophyll. *Nature Communications*, 10(1):2838.